MIT
Schwarzman
College of
Computing

MIT Schwarzman
College of Computing

As computing permeates every aspect of our lives, it is increasingly becoming part of the intellectual fabric of every field. In order to keep pace, conventional education and research pathways must evolve. Through a threefold mission that engages students and faculty across campus, the MIT Schwarzman College of Computing aims to lead the advancement of computing education and research, infuse computing into other disciplines, and facilitate consideration of social and ethical issues in computing.

As we mark the third year of the college, I am extremely proud of where we stand today. In that time, the college has undertaken several initiatives related to our mission, including creating the 6-4 Artificial Intelligence and Decision-Making major and establishing MicroMasters to address the need for artificial intelligence education and to serve as a model for a new national curriculum. The cross-cutting Common Ground for Computing Education now offers our "bilingual" students a selection of 11 individual subjects and curricula that combines computing and other disciplines and fuses different perspectives of multiple departments. Original pedagogical materials developed through the Social and Ethical Responsibilities of Computing by members from across computing, data sciences, humanities, arts, and social sciences are having substantial impact at MIT and beyond. During the 2021–2022 academic year, the materials were integrated into 14 subjects in seven departments and programs, enrolling over 2,100 students.

Furthermore, we continued our ongoing commitment to expand MIT's academic capacity by adding new faculty. To date, we have recruited 27 outstanding scholars through a mix of the college's 50 new positions and existing ones in the departments and schools. We also engaged principal investigators across all five schools and the college with a number of interdisciplinary computing research opportunities. In addition, we collaborated with the MIT Sloan School of Management to launch novel online professional education classes in computing and business with an industry focus. We have much more to look forward to in the years ahead, including the opening next summer of our new building, which will provide state-of-the-art space for computing education and research and serve as an interdisciplinary hub for computing at MIT.

I am thankful to the growing community of faculty, students, researchers, staff, and friends for its support and tireless effort to help move the college forward. People at MIT are known for rolling up their sleeves and flocking to get the work done. I've been fortunate to witness this firsthand. We have made a transformative impact thus far, and we could not have achieved without the invaluable contributions of this community.

Sincerely,

Daniel P. Huttenlocher SM '84, PhD '88
Dean, MIT Schwarzman College of Computing
Henry Ellis Warren (1894) Professor of Electrical Engineering and Computer Science

College Vision

The mission of the MIT Stephen A. Schwarzman College of Computing is to address the opportunities and challenges of the computing age by supporting the rapid evolution and growth of computer science (CS) and artificial intelligence (AI); facilitating collaborations between computing and other disciplines; and advancing multidisciplinary efforts in the study and practice of socially responsible computing, melding insights from technologists with those from across the humanities, arts, and social sciences.

Computing—hardware, software, algorithms, AI—is becoming part of the *conceptual fabric* of every domain, changing how problems are framed and the methods for solving them. The work of the college is stretching the frontiers of core computing and other disciplines that will transform and possibly yield new fields.

Advance core computing research and education, notably computer science and artificial intelligence	Infuse the forefront of computing with other fields to advance both	Incorporate social and ethical considerations given broad impact

Computing-Focused Academics and Research

Coalescing while also connecting more to the rest of MIT

Academic

MIT EECS — Electrical Engineering | Computer Science | Artificial Intelligence + Decision-making

IDSS — MIT Institute for Data, Systems, and Society

MIT Statistics + Data Science Center

MIT Technology & Policy Program

MIT | the center for computational science & engineering

Operations Research Center

Research

MIT CSAIL

MIT Quest for Intelligence

LIDS — Laboratory for Information & Decision Systems

MIT-IBM Watson AI Lab

MIT Jameel Clinic

Note: The Department of Electrical Engineering and Computer Science is joint with the School of Engineering. The Operations Research Center is joint with the MIT Sloan School of Management. The MIT-IBM Watson AI Lab is joint with the School of Engineering.

The fall 2021 SuperUROP Showcase, spotlighting advanced undergraduate research, was held in a hybrid Zoom and in-person session.

ts progi
ghlights p
ts highli
hlights p
ts progi
ghlights p
ts highli
hlights p
ts progi
ghlights p
ts highli
hlights p
ts progi
ghlights p
ts highli

Program Highlights

The MIT Stephen A. Schwarzman College of Computing continues to offer and develop exciting and new computing activities. These programs address the demand for CS and AI expertise and also build on the platforms for the study and practice of social and ethical responsibilities of computing and for multi-departmental computing education.

The AI+D major will prepare students both to go on to graduate school and for careers in machine learning, data science, and other AI-related areas in tech and other industries.

# New Major: Artificial Intelligence + Decision-Making

Preparations have been underway for a new undergraduate major in Artificial Intelligence + Decision-Making (AI+D) launching in fall 2022. Offered by the Department of Electrical Engineering and Computer Science (EECS), the major prepares students to use computational tools to infer, predict, and act on models of real-world problems. Furthermore, the new major allows MIT to develop and scale the curriculum in this area, which holds many possibilities for academic innovation.

The discipline of AI+D develops techniques for the analysis and synthesis of systems that interact with an external world via perception, communication, and action, and that learn, make decisions, and adapt in a changing environment. It includes study of the impact of such systems on the people and societies they interact with. Driven by recent technical innovations, massive availability of data, and the deployment of "intelligent" systems, ranging from Siri to autonomous vehicles to medical AI, this area has undergone huge growth in societal importance. It has opened up an enormous range of important intellectual questions, from foundational math to algorithms, applications, and social impact. With this major, MIT is able to educate students broadly and deeply in these intellectual areas, give them the practical skills to develop and deploy intelligent systems carefully and effectively, and evaluate critically the systems and claims produced by others.

This degree program will help meet the incredible demand for coursework in AI+D. For example, the introductory subject in machine learning has recently had an average enrollment of 500 students per semester or 12.5% of the undergraduates at MIT. EECS is working closely with the Social and Ethical Responsibilities of Computing to integrate social, ethical, and policy considerations into the intellectual fabric of the major. The department is also engaging with the Common Ground for Computing Education to establish additional introductory classes with different emphases. The plan is to develop additional courses under the major which will be of interest to majors in other MIT departments as well.

Through this new program, the college and EECS department will prepare students both for graduate school and for careers in machine learning, data science, and other AI-related areas in tech and other industries, e.g., automotive, financial services, logistics, health care, government agencies, and nonprofits.
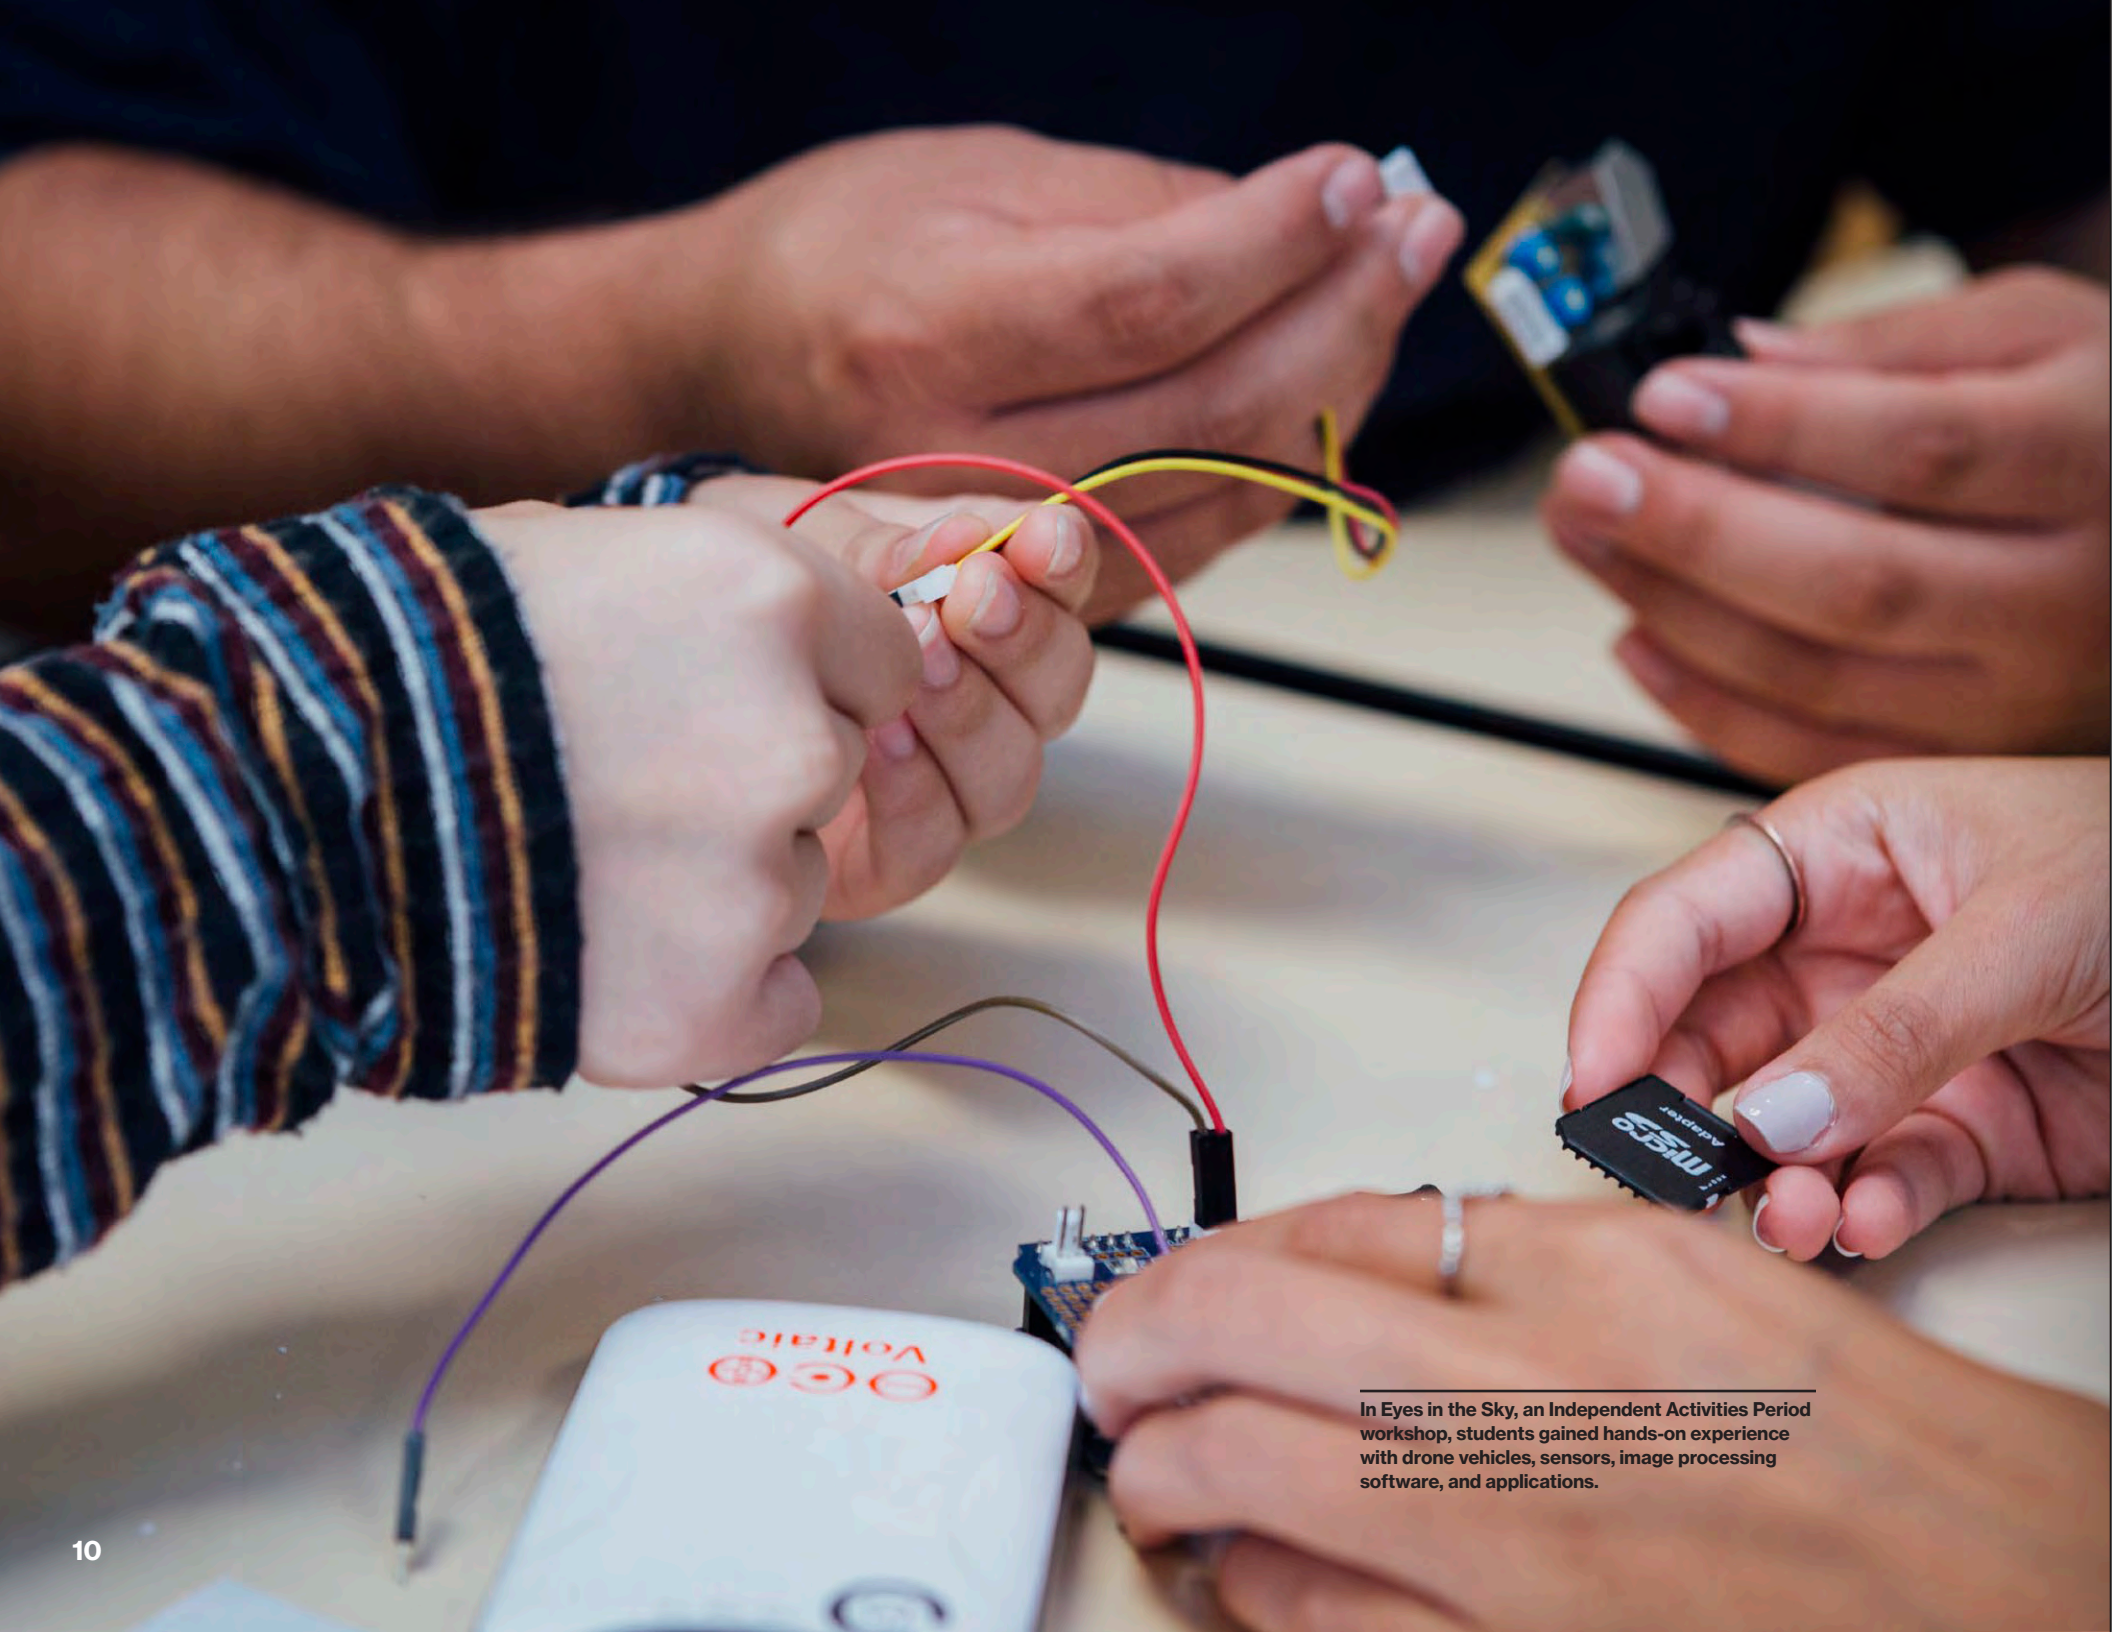
In Eyes in the Sky, an Independent Activities Period workshop, students gained hands-on experience with drone vehicles, sensors, image processing software, and applications.

Social and Ethical Responsibilities of Computing

Social and Ethical Responsibilities of Computing (SERC) was established through the MIT Schwarzman College of Computing to bring faculty from a range of disciplines together to develop resources to guide responsible research practices, create original pedagogical materials, catalyze new research collaborations, and forge partnerships with external stakeholders. SERC has been led by associate deans David Kaiser, the Germeshausen Professor of the History of Science and professor of physics, and Julie Shah '04, SM '06, PhD '11, the H. N. Slater Professor in Aeronautics and Astronautics and head of the Interactive Robotics Group at the Computer Science and Artificial Intelligence Laboratory (CSAIL).

Each semester, SERC invites about 20 faculty members to join action groups dedicated to advancing SERC projects, such as developing new curricular materials. Participants typically team up into smaller groups of two or three to hone their shared interest over the course of the term and develop new materials. Faculty are drawn from MIT's five schools and the college.

Some 80 faculty members have participated since the program started, and in the last year alone, more than 2,100 students took courses that included new SERC content. Those students aren't all engineers; about 500 encountered original SERC content through courses offered in the School of Humanities, Arts, and Social Sciences, the MIT Sloan School of Management, and the School of Architecture and Planning.

Central to SERC is the principle that ethics and social responsibility in computing should be integrated into all areas of teaching and research at MIT, so they become just as central as the technical parts of the curriculum. Technology touches every industry, so students and researchers in all disciplines should have training that helps them understand tools like AI, and to think deeply about how to mitigate potential harms as well as foster broad benefits.

> **"By sharing how MIT faculty and students use the materials, we're creating pathways for educators around the world to adapt the materials for maximum relevance to their students."**
>
> **Curt Newton,** Director, MIT OpenCourseWare

SERC Scholars

The inaugural 2021–2022 cohort of SERC Scholars, representing undergraduate and graduate students from across all five MIT schools and the college, have worked to advance SERC's broad mission of incorporating humanist, social science, social responsibility, and policy/civic perspectives into MIT's teaching, research, and implementation of computing. There were 54 SERC scholars throughout the year working alongside four postdocs.

SERC Scholars collaborate in multidisciplinary teams with team members from across computing, data sciences, humanities, arts, and social sciences; develop and pilot new SERC course materials in collaboration with postdocs and faculty; engage with external partners to advance AI in the public interest; and conduct research in an exciting, new cross-disciplinary area. Additional opportunities for SERC Scholars include regular informal meetings and talks with pioneering leaders from academia, industry, and the public sectors; internships advancing computing in the public interest; SERC undergraduate research projects; and initiatives in computing and public policy.

A new resource for teaching responsible technology development

Curricular materials created through SERC are now freely available via MIT OpenCourseWare (OCW). The collection includes original active learning projects, homework assignments, in-class demonstrations, and other resources and tools.

SERC has been bringing together cross-disciplinary teams of faculty, researchers, and students to generate the original content. Most of the materials featured on OCW were produced by participants in SERC's semester-long Dean's Action Groups on Active Learning Projects, in which faculty from the humanities, arts, and social sciences are paired with faculty in computing and data sciences to collaborate on new projects for each of their existing courses.

Organized by topic areas, including privacy and surveillance; inequality, justice, and human rights; AI and algorithms; social and environmental impacts; autonomous systems and robotics; ethical computing and practice; and law and policy, the site also spotlights materials from the MIT Case Studies in Social and Ethical Responsibilities of Computing, an ongoing, peer-reviewed series that examines social, ethical, and policy challenges of present-day efforts in computing. The Case Studies series is available free via open-access publishing.

Learning to think critically about machine learning

Students in the MIT course 6.036 Introduction to Machine Learning study the principles behind powerful models that help physicians diagnose disease or aid recruiters in screening job candidates.

Now, thanks to the SERC framework, these students are also critically engaging with the implications of these AI tools, which sometimes come with their share of unintended consequences.

A team of SERC Scholars worked with Leslie Kaelbling, the Panasonic Professor of Computer Science and Engineering, and the 6.036 teaching assistants to infuse weekly labs with material covering ethical computing, data and model bias, and fairness in machine learning.

More than 500 students who were enrolled in 6.036 during the 2021 spring term grappled with these ethical dimensions alongside their efforts to learn new computing techniques. For some, it may have been their first experience thinking critically in an academic setting about the potential negative impacts of machine learning.

"I was inspired and energized by this process, and I learned so much, not just the technical material, but also what you can achieve when you collaborate across disciplines. Just the scale of this effort felt exciting. If we have this cohort of 500 students who go out into the world with a better understanding of how to think about these sorts of problems, I feel like we could really make a difference."

SERC Scholar Marion Boulicault PhD '21, the 2021 recipient of the MAC3 Society and Ethics in Computing Research Award, who was a graduate student in the Department of Linguistics and Philosophy and is now a postdoc in the Schwarzman College of Computing.

Ankur Moitra SM '09, PhD '11, right, the Norbert Wiener Professor of Mathematics at MIT, talks with students during a class in the Common Ground course Linear Algebra and Optimization.

Common Ground for Computing Education

Computing permeates every aspect of our lives and impacts education and research across all academic disciplines. Common Ground capitalizes on this unprecedented opportunity by facilitating coordinated computing education across the Institute. The initiative brings multiple departments together to develop and teach new courses and launch new programs that blend computing with other disciplines.

"MIT has a vibrant community in interdisciplinary research. These classes are designed to make those same cross-cutting connections, giving students unique and exciting opportunities in their education."

Jennifer Donath, Program Manager, Common Ground

> **"It is the combination of math with science and computer science. It sings when you put it all together. For the students, it is really a skills-based class. We want to provide students with skills that can be used almost everywhere in their studies later on, and then in so many other domains as well."**

Laurent Demanet, professor of applied mathematics, who designed Introduction to Computational Science and Engineering with David Darmofal SM '91, PhD '94, the Jerome C. Hunsaker Professor of Aeronautics and Astronautics

Leadership

Common Ground is led by Deputy Dean for Academics Asu Ozdaglar SM '98, PhD '03, the MathWorks Professor of Electrical Engineering and Computer Science and head of the Department of Electrical Engineering and Computer Science (EECS).

Oversight for this initiative is provided by the Common Ground Standing Committee, co-chaired by Ozdaglar and Jeffrey C. Grossman, the Morton and Claire Goulder and Family Professor of Environmental Systems and head of the Department of Materials Science and Engineering. The committee is composed of faculty from MIT's five schools, the college, and MIT Open Learning.

In March, Robert C. Miller '95, MNG '95, Distinguished Professor in EECS, was named the inaugural education officer for Common Ground. In this role, Miller, who also co-chairs the Common Ground subcommittee on Programming/Computational Thinking, coordinates Common Ground subject offerings and staffing with departments, as well as teaching assignments for shared faculty in the Schwarzman College of Computing.

Piloting Success

After successful pilots in the 2020–2021 academic year, the first three Common Ground subjects—Introduction to Computational Science and Engineering, Linear Algebra and Optimization, and Modeling with Machine Learning: from Algorithms to Applications—were added permanently to programs of study in brain and cognitive sciences, civil and environmental engineering, EECS, energy, mathematics, supply chain management, transportation, and urban planning.

Highlights

- **Introduction to Computational Science and Engineering** teaches students how to combine computational techniques with physical modeling to solve real-world problems, from landing a spacecraft to placing cell phone towers.

  The first problem set involves using a computer to simulate the Mars landing. Students learn to use computer programs for simulation, optimization, and uncertainty quantification.

These foundational principles are framed with tangible examples designed to be relatable to students who aren't necessarily computer science majors. Most students in the course in the fall 2021–2022 semester were either studying aeronautics and astronautics or math.

- **Physical Systems Modeling and Design Using Machine Learning** is one of five discipline-specific modules students can take as part of the course Modeling with Machine Learning: from Algorithms to Applications. Other modules focus on molecular engineering, nuclear science and engineering, sustainable systems, and supply chains.

In Physical Systems, George Barbastathis, professor of optics and mechanical engineering, highlights how complementary physics-based engineering and data science are. Physical laws present a number of ambiguities and unknowns, ranging from temperature and humidity to electromagnetic forces. Data science can be used to predict these physical phenomena. Meanwhile, having an understanding of physical systems helps ensure the resulting output of an algorithm is accurate and explainable.

"What's needed is a deeper combined understanding of the associated physical phenomena and the principles of data science, machine learning in particular, to close the gap," adds Barbastathis. "By combining data with physical principles, the new revolution in physics-based engineering is relatively immune to the 'black box' problem facing other types of machine learning."

In addition to choosing a module, students in the Modeling with Machine Learning: from Algorithms to Applications course also take a core class taught by Regina Barzilay, the School of Engineering Distinguished Professor for AI and Health, and Tommi Jaakkola PhD '97, the Thomas Siebel Professor of Electrical Engineering and Computer Science and the Institute for Data, Systems, and Society.

New Pilots

A call for proposals in fall 2021 resulted in approval of three new pilots. In coordination with the Common Ground Standing Committee, faculty have continued to develop the subjects for the 2022–2023 academic year:

- **Computational Imaging: Physics and Algorithms, Fall 2022**

    Instructors: George Barbastathis, professor of optics and mechanical engineering; James M. LeBeau, associate professor of materials science and engineering; Rajeev Ram, professor of electrical engineering; and Sixian You, Alfred Henry and Jean Morrison Hayes Career Development Professor and assistant professor of electrical engineering.

- **Julia: Solving Real-World Problems with Computation, Fall 2022**

    Instructors: Alan Edelman PhD '89, professor of applied mathematics and Julia Lab research group leader; Raffaele Ferrari, the Cecil and Ida Green Professor of Oceanography; Youssef Marzouk '97, SM '99, PhD '04, professor of aeronautics and astronautics; and John Williams, professor of civil and environmental engineering.

- **Interactive Data Visualization and Society, Spring 2023**

    Instructors: Catherine D'Ignazio SM '14, the Sherman Fairchild Career Development Professor and associate professor of urban studies and planning; Arvind Satyanarayan, assistant professor of computer science; and Sarah Williams MCP '05, the Norman B. and Muriel Leventhal Professor of Architecture and Planning.

The Standing Committee also laid groundwork for additional pilots in optimization, causal inference, creative computing, and digital humanities and began coordinating with departments to form teaching teams for each.
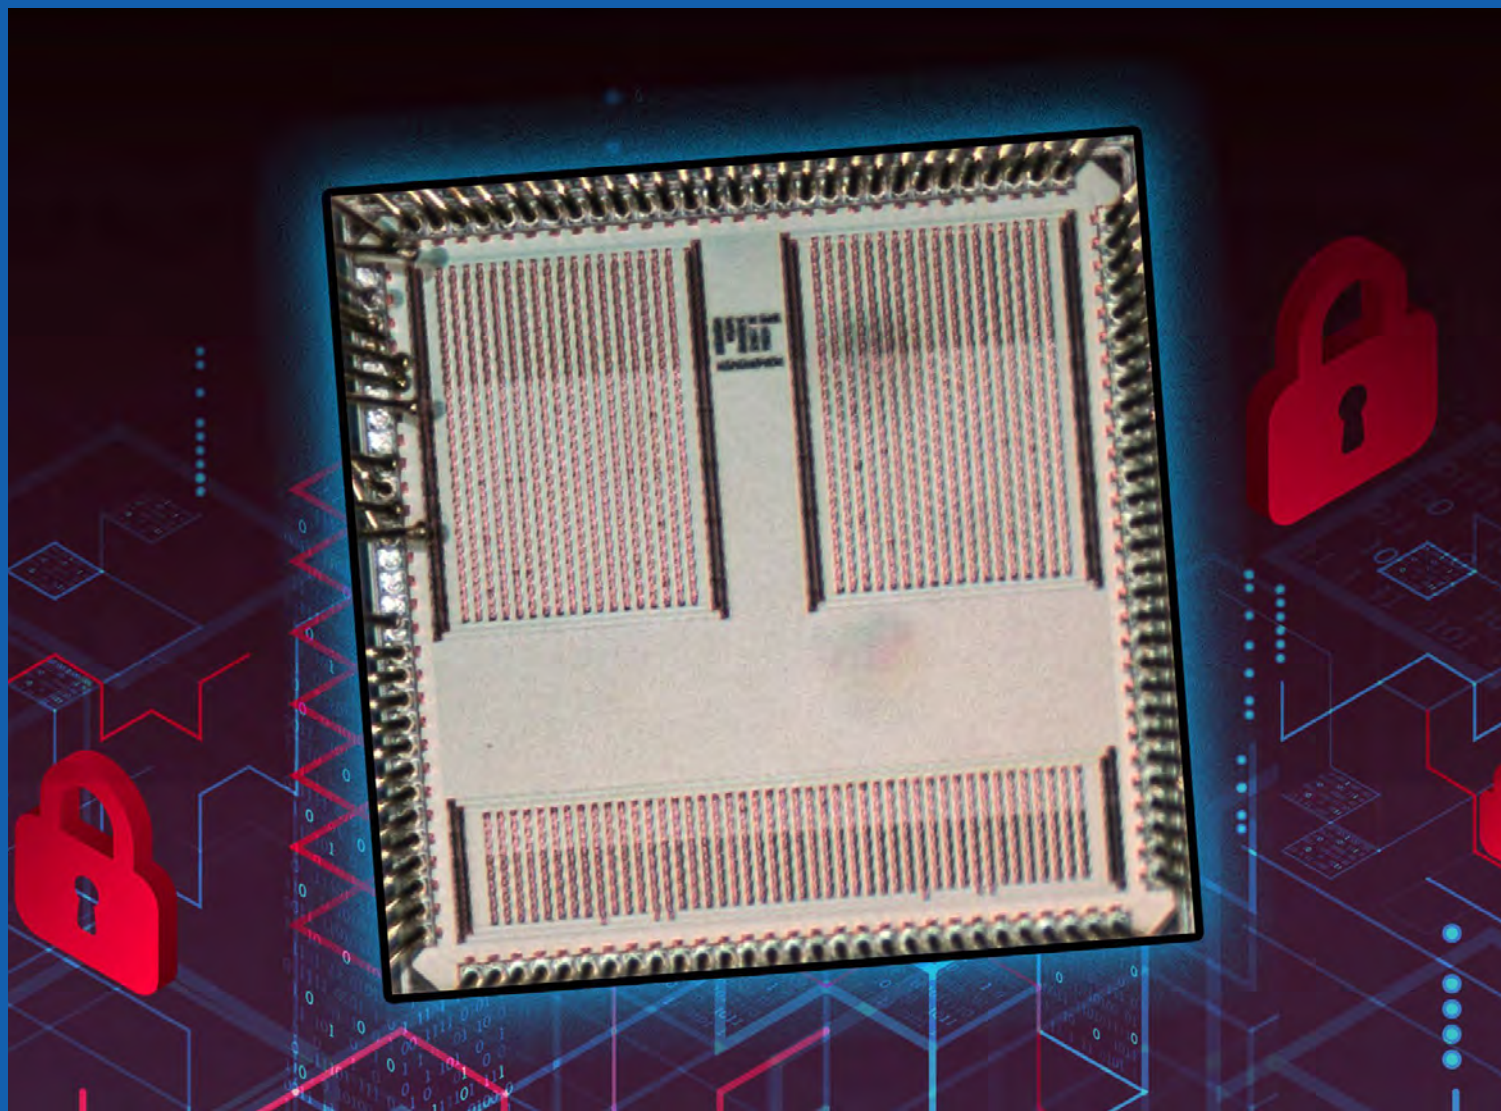
MIT researchers built a lower-energy chip that can prevent hackers from extracting hidden information from a smart device.

Research Impact

Research projects at the MIT Schwarzman College of Computing bring together MIT faculty and students to advance computing in a quest for a better world. As humanity works to address the challenges presented by climate change, disease, poverty, inequality, and other great problems, computing offers powerful new solutions, with MIT scholars at the forefront applying big ideas across disciplines.
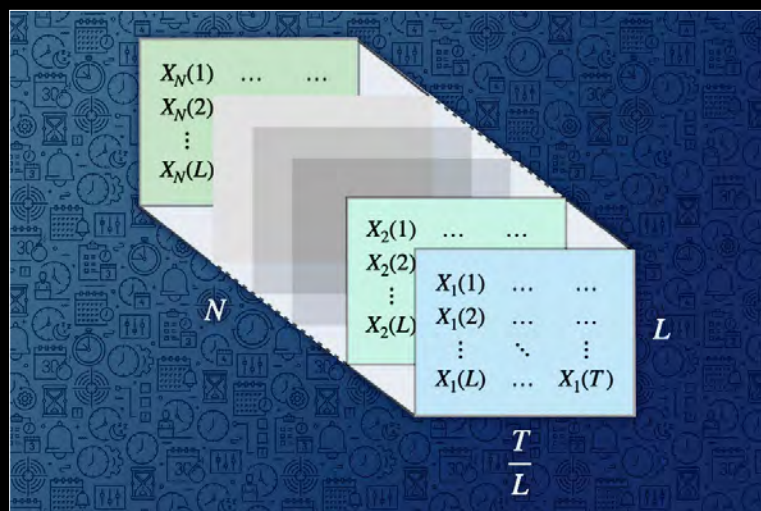
Researchers designed a user-friendly interface that helps nonexperts make forecasts using data collected over time.

A tool for predicting the future

Whether trying to predict tomorrow's weather, forecast future stock prices, identify missed opportunities for sales in retail, or estimate a patient's risk of developing a disease, one likely will need to interpret time-series data, a collection of observations recorded over time. Making predictions using time-series data typically requires several data-processing steps and the use of complex machine-learning algorithms that have such a steep learning curve they aren't readily accessible to nonexperts.

To make these powerful tools more user-friendly, MIT researchers in EECS, with affiliations in the MIT Institute for Data, Systems, and Society (IDSS), and the Laboratory for Information & Decision Systems (LIDS) developed a system that directly integrates prediction functionality on top of an existing time-series database. Their simplified interface does the complex modeling behind the scenes so a nonexpert can easily generate a prediction in only a few seconds. The new system is more accurate and more efficient than state-of-the-art deep learning methods when predicting future values and filling in missing data points.

A security technique to fool would-be cyber attackers

Multiple programs running on the same computer that share the same memory hardware are vulnerable to having their secrets stolen by a malicious program through a "memory timing side-channel attack." This malicious program notices delays when it tries to access a computer's memory because the hardware is shared among all programs using the machine. It can then interpret those delays to obtain another program's secrets, like a password or cryptographic key.

A team of EECS and CSAIL researchers has devised a new approach that allows memory sharing to continue while providing strong security against this type of side-channel attack. In addition to providing better security while enabling faster computation, the technique could be applied to a range of different side-channel attacks that target shared computing resources, the researchers say.

Seeing into the future: Personalized cancer screening with AI

MIT scientists from EECS, CSAIL, and the MIT Jameel Clinic worked alongside those at institutions across the globe to demonstrate that AI-risk models, paired with AI-designed screening policies, can offer significant and equitable improvements to cancer screening.

Current guidelines state women aged 45 to 54 should get mammograms every year. While personalized screening has long been thought of as the answer, tools that can leverage the troves of data to do this lag behind.

Out of this came Tempo, a technology for creating risk-based screening guidelines. Using an AI-based risk model that looks at who was screened and when they got diagnosed, Tempo will recommend a patient return for a mammogram at a specific time in the future, such as six months or three years.

The model was trained on a large screening mammography dataset from Massachusetts General Hospital (MGH) and was tested on held-out patients from MGH as well as external datasets from Emory University, Karolinska Institute (Sweden), and Chang Gung Memorial (Taiwan) hospitals. Using the team's previously developed risk-assessment algorithm Mirai, Tempo obtained better early detection than annual screening while requiring 25% fewer mammograms overall at Karolinska. At MGH, it recommended roughly a mammogram a year, and obtained a simulated early detection benefit of roughly four-and-a-half months better.
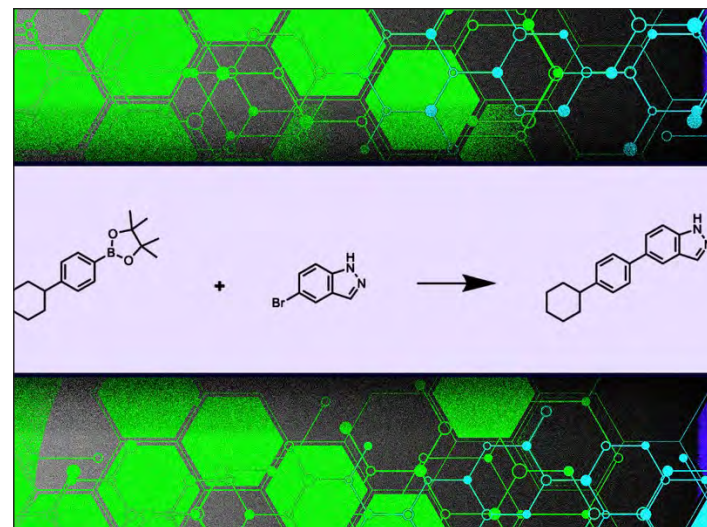
MIT researchers devised a new AI technique that only proposes candidate molecules that can actually be produced in a lab.

A smarter way to develop new drugs

An MIT researcher in EECS and the MIT Department of Chemical Engineering has developed with his team a machine-learning model that proposes new molecules for the drug discovery process, while ensuring these molecules can be synthesized in a laboratory.

The method guarantees that molecules are composed of materials that can be purchased and that the chemical reactions that occur between those materials follow the laws of chemistry.

Their system also takes less than one second to propose a synthetic pathway.

Toward a smarter electronic health record

An AI-enhanced system enables doctors to spend less time searching for clinical information and more time treating patients.

Electronic health records have been widely adopted with the hope they would save time and improve the quality of patient care. But due to fragmented interfaces and tedious data entry procedures, physicians often spend more time navigating these systems than they do interacting with patients.

Researchers at MIT EECS, CSAIL, Institute for Medical Engineering & Science (IMES), and Jameel Clinic along with partners at Beth Israel Deaconess Medical Center are combining machine learning and human-computer interaction to create a better electronic health record (EHR). They developed MedKnowts, a system that unifies the processes of looking up medical records and documenting patient information into a single, interactive interface.

Driven by AI, this smart EHR automatically displays customized, patient-specific medical records when a clinician needs them. MedKnowts also provides autocomplete for clinical terms and auto-populates fields with patient information to help doctors work more efficiently.

Systems scientists find clues to why false news snowballs on social media

The spread of misinformation on social media is a pressing societal problem that tech companies and policy makers continue to grapple with, yet those who study this issue still don't have a deep understanding of why and how false news spreads.

A researcher in the Department of Civil and Environmental Engineering with affiliations in IDSS and LIDS, working alongside a graduate student in the Social and Engineering Systems program in IDSS and a research scientist in LIDS, developed a theoretical model of a Twitter-like social network to study how news is shared and explore situations where a non-credible news item will spread more widely than the truth. Agents in the model are driven by a desire to persuade others to take on their point of view. The key assumption in the model is that people bother to share something with their followers if they think it is persuasive and likely to move others closer to their mindset. Otherwise they won't share.

The researchers found that in such a setting, when a network is highly connected or the views of its members are sharply polarized, news that is likely to be false will spread more widely and travel deeper into the network than news with higher credibility. This theoretical work could help social media companies adapt networks to limit the spread of false information.

MedKnowts is a smart electronic health record
system that can help doctors work more efficiently.

Assessing the impact of climate change and Covid-19 on human well-being

MIT researchers in the Department of Urban Studies and Planning worked with a team including a graduate student from the MIT Technology and Policy Program (TPP) to conduct a policy-relevant study that paired natural language processing techniques with social media data to generate novel insights on how pressing threats like climate change and Covid-19 impact the well-being of individuals across the world. The team analyzed language terms used in hundreds of millions of social media posts in about 100 countries. This analysis found a pronounced drop in positive public sentiment after the pandemic set in during early 2020, with a subsequent, incremental, halting return to prepandemic status.

Sentiments that people around the world express online have much potential to serve as an effective alternative to conventional survey-based methods when policy makers seek to measure how people are feeling about current events.

Embodied intelligence

MIT researchers from across EECS, the Department of Brain and Cognitive Sciences (BCS), CSAIL, and the McGovern Institute for Brain Research are working on broadly addressing how we perceive the world around us and integrate this information to plan and complete tasks. Scientific goals include research into how perception, planning, and action interface, and how we learn efficiently from small datasets and the creation of behavioral benchmark tasks. Engineering goals include task-based benchmarks derived from natural intelligence and methods for composing artificial systems to achieve those tasks.

This is one of a number of research missions being fostered and funded by the MIT Quest for Intelligence. Each mission brings together a team of scientists and engineers to pose and answer foundational questions of natural intelligence where current AI falls short, to build engineered systems as scientific hypotheses to advance these studies, and to execute tests of those systems to ensure that scientific progress is iteratively guided by natural intelligence results and real-world AI engineering challenges.

A research mission launched by the MIT Quest for Intelligence broadly addresses how we perceive the world around us and integrates this information to plan and complete tasks.

Machines that see the world more as humans do

Computer vision systems sometimes make inferences about a scene that fly in the face of common sense. For example, if a robot were processing a scene of a dinner table, it might completely ignore a bowl that is visible to any human observer, estimate that a plate is floating above the table, or misperceive a fork to be penetrating a bowl rather than leaning against it. Move that computer vision system to a self-driving car and the stakes become much higher.

MIT researchers from EECS, BCS, and the MIT-IBM Watson AI Lab have developed a framework that helps machines see the world more as humans do. Their new AI system for analyzing scenes learns to perceive real-world objects from just a few images, and perceives scenes in terms of these learned objects.

The researchers built the framework using probabilistic programming, an AI approach that enables the system to cross-check detected objects against input data to see if the images recorded from a camera are a likely match to any candidate scene. This safeguard enables the system to detect and correct many errors that plague the deep-learning approaches also used for computer vision. Probabilistic programming also makes it possible to infer probable contact relationships between objects in the scene, and use common-sense reasoning about these contacts to infer more accurate positions for objects.

A dataset that was used to create crash-risk maps covered 7,500 square kilometers in Los Angeles, New York City, Chicago, and Boston. Among the four cities, Los Angeles was the most unsafe, since it had the highest crash density, followed by New York City, Chicago, and Boston.

Deep learning helps predict traffic crashes before they happen

A deep model was trained on historical crash data, road maps, satellite imagery, and GPS to enable high-resolution crash maps that could lead to safer roads. The model, developed by scientists from EECS, CSAIL, and the Qatar Center for Artificial Intelligence, produced risk maps that describe the expected number of crashes over a period of time in the future, to identify high-risk areas and predict future crashes.

By capturing the underlying risk distribution that determines the probability of future crashes, the maps can help find safer routes, enable auto insurance companies to provide customized insurance plans based on driving trajectories of customers, help city planners design safer roads, and even predict future crashes, researchers say.

Getting dressed with help from robots

Robots don't have the same hardwired behavioral awareness and control as people, so secure collaboration with humans requires methodical planning and coordination. You can likely assume your friend can fill up your morning coffee cup without spilling on you, but for a robot, this seemingly simple task requires careful observation and comprehension of human behavior.

EECS, CSAIL, and the Department of Aeronautics and Astronautics scientists recently created a new algorithm to help a robot find efficient motion plans to ensure physical safety of its human counterpart. In this case, the bot helped put a jacket on a human, which could potentially prove to be a powerful tool in expanding assistance for those with disabilities or limited mobility.

MIT CSAIL PhD student Shen Li gets safe assistance with dressing using his team's human-aware motion planning algorithm.

AI predicts patients' race from their medical images

The miseducation of algorithms is a critical problem. When AI mirrors unconscious thoughts, racism, and biases of the humans who generated these algorithms, it can lead to serious harm. Removing sensitive features from the data seems like a viable tweak. But what happens when it's not enough?

Examples of bias in natural language processing are boundless—but MIT scientists representing EECS, CSAIL, IMES, and Jameel Clinic have investigated another important, largely underexplored modality: medical images. Using both private and public datasets, the team found that AI can accurately predict self-reported race of patients from medical images alone. Using imaging data of chest X-rays, limb X-rays, chest CT scans, and mammograms, the team trained a deep learning model to identify race as white, Black, or Asian, even though the images themselves contained no explicit mention of the patient's race. This is a feat even the most seasoned physicians cannot do, and it's not clear how the model was able to do this.

The fact that algorithms "see" race can be dangerous. But an important and related fact is that, when used carefully, algorithms can also work to counter bias. Just like any tool, algorithms can be a force for evil or a force for good: which one depends on those using them and the choices made when building algorithms.

Exact symbolic AI for faster, better assessment of fairness

The justice system, banks, and private companies use algorithms to make decisions that have profound impacts on people's lives. Unfortunately, those algorithms are sometimes biased, disproportionately impacting people of color as well as individuals experiencing economic hardship when they apply for loans or jobs, or even when courts decide what bail should be set while a person awaits trial.

MIT EECS and CSAIL researchers say their new probabilistic programming language can quickly and accurately assess the fairness of AI algorithms.

Probabilistic programming is an emerging field at the intersection of programming languages and AI that aims to make AI systems much easier to develop, with early successes in computer vision, common-sense data cleaning, and automated data modeling. Probabilistic programming languages make it much easier for programmers to define probabilistic models and carry out probabilistic inference, that is, work backward to infer probable explanations for observed data.

The system developed by the MIT researchers gives fast, exact solutions to probabilistic inference questions such as "How likely is the model to recommend a loan to someone over age 40?" or "Generate 1,000 synthetic loan applicants, all under age 30, whose loans will be approved."

Dexterous robotic hands manipulate thousands of objects with ease

Scientists from EECS and CSAIL, in the ever-present quest to get machines to replicate human abilities, created a system that can reorient more than 2,000 different objects, with a robotic hand facing both upwards and downwards. This ability to manipulate anything from a cup to a tuna can to a Cheez-It box could help the hand quickly pick and place objects in specific ways and locations, and even generalize to unseen objects.

This deft "handiwork," which is usually limited to single tasks and upright positions, could be an asset in speeding up logistics and manufacturing, helping with common demands such as packing objects into slots for kitting, or dexterously manipulating a wider range of tools. The team used a simulated, anthropomorphic hand with 24 degrees of freedom, and showed evidence that the system could be transferred to a real-world robotic system in the future.
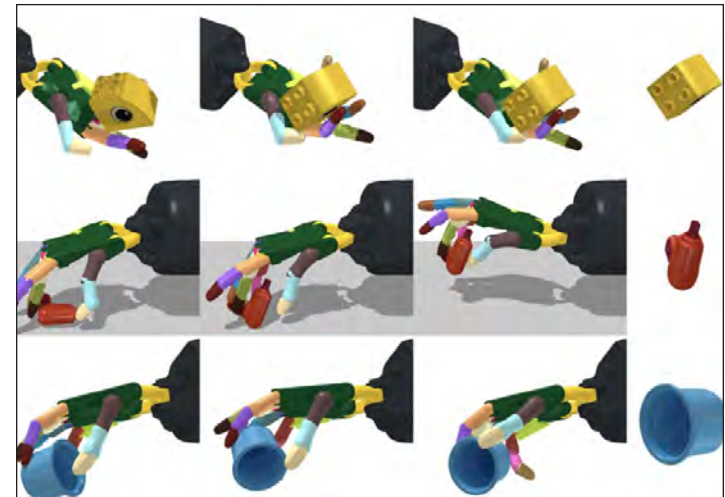
A new system can reorient more than 2,000 different objects, with the robotic hand facing both upwards and downwards.

A comprehensive study of technological change

The societal impacts of technological change can be seen in many areas, from messenger RNA vaccines and automation to drones and climate change. The pace of that technological change can affect its impact, and how quickly a technology improves in performance can be an indicator of its future importance. For decision makers like investors, entrepreneurs, and policy makers, predicting which technologies are fast improving (and which are overhyped) can mean the difference between success and failure.

New research from MIT aims to assist in the prediction of technology performance improvement using US patents as a dataset. The team built on previous work from researchers in IDSS and the MIT Sociotechnical Systems Research Center which found that by looking for patent overlap between US and international patent-classification systems, they could quickly identify patents that best represent a technology. These researchers ultimately created a correspondence of all patents within the US patent system to a set of 1,757 technology domains, and quantitatively assessed each domain for its improvement potential.

A one-up on motion capture

From *Star Wars* on, many popular films contain scenes made possible by motion capture technology, which records movement of objects or people through video, and involves complicated interactions between physics, geometry, and perception. Applications extend beyond Hollywood to the military, sports training, medical fields, and computer vision and robotics, allowing engineers to understand and simulate action happening within real-world environments.

As this can be a complex and costly process, often requiring markers placed on objects or people and recording the action sequence, researchers are working to shift the burden to neural networks, which could acquire these data from a simple video and reproduce them in a model.

A team of researchers from EECS, CSAIL, and BCS working alongside those at the MIT-IBM Watson AI Lab has developed a trained neural network pipeline with the ability to infer the state of the environment and the actions happening, the physical characteristics of the object or person of interest, and its control parameters. The technique can outperform other methods in simulations illustrating different types of dynamics and interactions under various environmental conditions. Further, the methodology allows for imitation learning, predicting and reproducing the trajectory of a flying quadrotor helicopter from a video.

The team tested their method in four simulated systems: the quadrotor (a flying rigid body that doesn't have any physical contact), a cube (a rigid body that interacts with its environment, like a die), an articulated hand, and a rod (deformable body that can move like a snake).
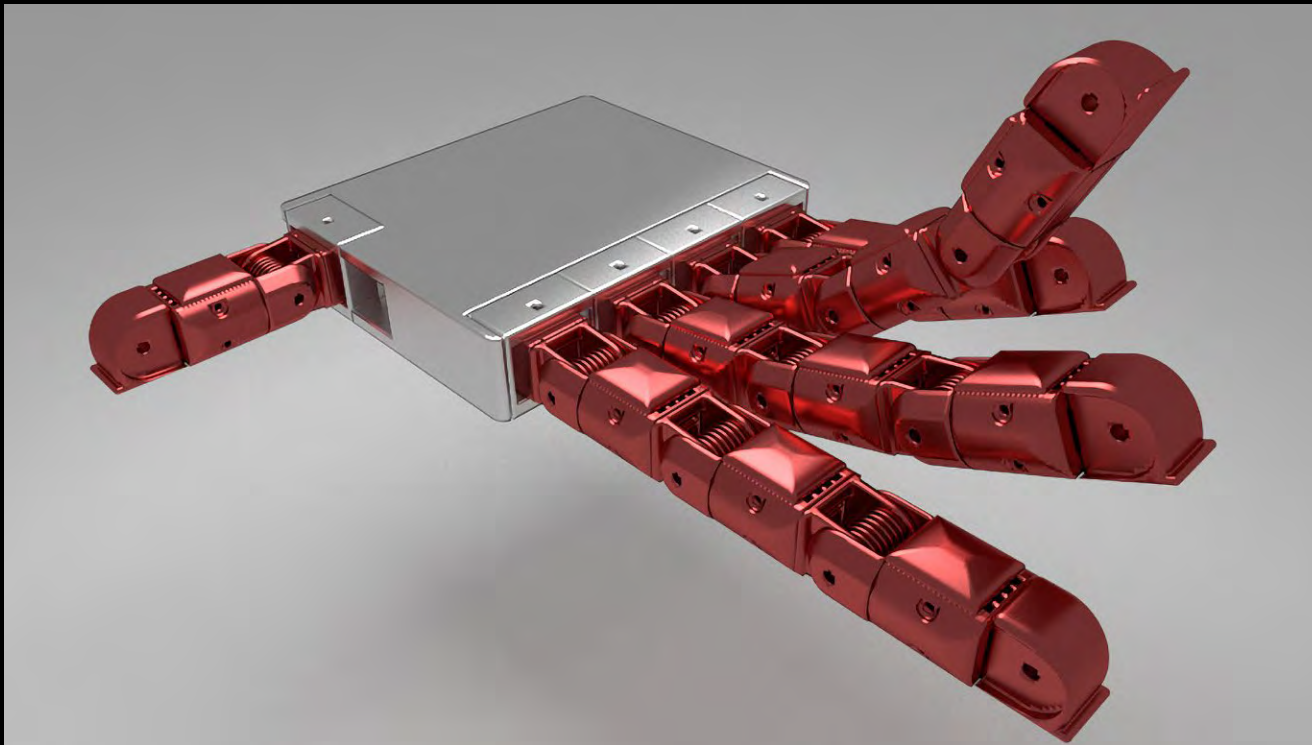
MIT researchers used a method based on a trained neural network pipeline to predict the action sequence, joint stiffness, or movement of an articulated hand like this one from a target image or video.

Accelerating the discovery of new materials for 3-D printing

The growing popularity of 3-D printing for manufacturing all sorts of items, from customized medical devices to affordable homes, has created more demand for new 3-D printing materials designed for very specific uses.

To cut down on the time it takes to discover these new materials, researchers at CSAIL have developed a data-driven process that uses machine learning to optimize new 3-D printing materials with multiple characteristics, like toughness and compression strength.

By streamlining materials development, the system lowers costs and lessens its environmental impact by reducing the amount of chemical waste. The machine-learning algorithm could also spur innovation by suggesting unique chemical formulations that human intuition might miss.

## Promoting climate justice through modeling and analysis

Action on climate could deliver major health and economic benefits to people in the United States, but will such benefits accrue to low-income and minority communities that are harmed the most by environmental pollution? A new study led by researchers at the MIT Joint Program on the Science and Policy of Global Change explores how, under a federally coordinated carbon dioxide emissions cap-and-trade program aligned with the US Paris Agreement pledge and implemented through Section 115 of the Clean Air Act, the EPA might allocate emissions cuts among states.

As a part of this effort, a TPP student is using integrated modeling and analysis to assess the influence of strategies to mitigate climate change on air pollution and environmental justice, adding to the ongoing debate about the equity-related impacts of carbon pricing policies in the United States.

Researchers at MIT and the chemical company BASF have developed a data-driven system that accelerates the process of discovering new 3-D printing materials that have multiple mechanical properties.

To represent a 3-D scene from a 2-D image, a light field network encodes the 360-degree light field of the 3-D scene into a neural network that directly maps each camera ray to the color observed by that ray.
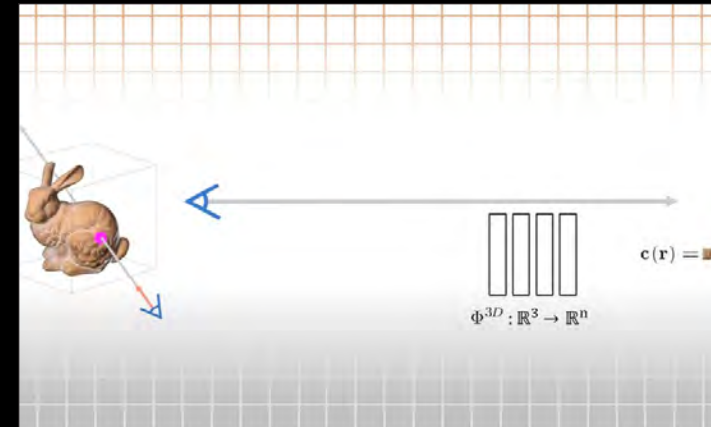
A new language for quantum computing

Quantum computers have the potential for computational breakthroughs in classically unsolvable tasks, like cryptographic and communication protocols, search, and computational physics and chemistry.

Unlike traditional computers that use bits, quantum computers use qubits to encode information as zeros or ones or both at the same time. Harnessing forces from quantum physics, these refrigerator-sized machines can process a great deal of information but they're far from flawless. As with regular computers, the right programming languages are needed to properly compute on quantum computers.

Scientists representing EECS, CSAIL, and the Research Laboratory of Electronics created their own programming language for quantum computing called Twist that can describe and verify which pieces of data are entangled to prevent bugs in a quantum program. Twist is seen as paving the way to languages that make the unique challenges of quantum computing more accessible to programmers.

Technique enables real-time rendering of scenes in 3-D

A machine that needs to interact with objects in the world, like a robot designed to harvest crops or assist with surgery, must be able to infer properties about a 3-D scene from observations of the 2-D images it's trained on. While scientists have had success using neural networks to infer representations of 3-D scenes from images, these machine-learning methods aren't fast enough to make them feasible for many real-world applications.

A new technique demonstrated by researchers at EECS, CSAIL, BCS, and elsewhere is able to represent 3-D scenes from images about 15,000 times faster than some existing models.

The method represents a scene as a 360-degree light field, which is a function that describes all the light rays in a 3-D space, flowing through every point and in every direction. The light-field networks the researchers developed can reconstruct a light field after only a single observation of an image, and they are able to render 3-D scenes at real-time frame rates.

Students share a laugh over a computer on campus, spring 2021.

nity   hi
ts commu
ty   highli
mmunity
nity   hi
ts commu
ty   highli
mmunity
nity   hi
ts commu
ty   highli
mmunity
nity   hi
ts commu
ty   highli

Community Highlights

The MIT Schwarzman College of Computing derives its energy from world-class faculty and researchers at the nexus of computing and students channeling their passion for computer science and AI into solving the world's biggest challenges. Through exciting research and innovative activities, they are addressing today's most consequential questions facing our world.

33

Faculty

The Schwarzman College of Computing continues to increase MIT's academic capacity in computing and AI, welcoming a new cohort of faculty. Those in core computing positions are appointed in the Department of Electrical Engineering and Computer Science (EECS) and shared positions are appointed in both an academic department and in the college.

New faculty in core computing positions

Sara Beery will join EECS's Faculty of AI+D as an assistant professor in September 2023. Beery received her PhD in Computing + Mathematical Sciences at Caltech, advised by Pietro Perona. Her research focuses on building computer vision methods that enable global-scale environmental and biodiversity monitoring across data modalities, tackling real-world challenges including strong spatiotemporal correlations, imperfect data quality, fine-grained categories, and long-tailed distributions.

Christina Delimitrou will join MIT as an assistant professor in the Faculty of Computer Science in EECS in September 2022. She arrives from Cornell University, where she was an assistant professor. Her main interests are in computer architecture and computer systems. She works on improving the design, management, and programming interfaces of modern cloud systems by leveraging practical machine-learning techniques to solve large-scale systems problems. Delimitrou earned a PhD in electrical engineering at Stanford University, where she worked with Christos Kozyrakis.

Priya Donti will join MIT as an assistant professor in the Faculties of Electrical Engineering and AI+D in EECS in academic year 2023–2024. Her work focuses on machine learning for forecasting, optimization, and control in high-renewable power grids. Specifically, her research explores methods to incorporate the physics and hard constraints associated with electric power systems into deep learning models. Donti was a PhD candidate in the Computer Science Department and the Department of Engineering & Public Policy at Carnegie Mellon University, co-advised by Zico Kolter and Inês Azevedo.

Gabriele Farina joins MIT's EECS Faculty of AI+D in academic year 2023–2024 as an assistant professor. He was a PhD candidate in the Computer Science Department at Carnegie Mellon University, where he worked under the supervision of Tuomas Sandholm. His work lies at the interface between AI, computer science, operations research, and economics. His primary research interests lie in learning and optimization methods for multi-agent decision making, with a focus on the computation of optimal strategies and equilibria in large-scale strategic interactions.

Kuikui Liu will join MIT as an assistant professor in the Faculty of Computer Science in EECS in September 2022. He was a PhD student in the Theory Group at the Paul G. Allen School for Computer Science and Engineering at the University of Washington, where he was advised by Shayan Oveis Gharan. Liu's research interests are in high-dimensional geometry and analysis of Markov chains. He develops and uses mathematical tools from fields such as high-dimensional expanders, geometry of polynomials, and statistical physics.

The fall 2021 SuperUROP Showcase, spotlighting advanced undergraduate research, was held in a hybrid Zoom and in-person session.

New faculty in shared positions

Ericmoore Jossou will join MIT as an assistant professor in a shared position between the Department of Nuclear Science and Engineering and the Faculty of Electrical Engineering in EECS in July 2023. He is currently a postdoctoral research associate in the Nuclear Science and Technology Department of Brookhaven National Laboratory (BNL). Prior to joining BNL in 2019, he received his doctorate in mechanical engineering with specialization in materials science from the University of Saskatchewan. His research interests focus on applying computational and experimental techniques to understand the fundamental thermodynamic and kinetic driving forces influencing microstructural evolution at multi-length scales, particularly in extreme environments.

Sherrie Wang will join MIT as an assistant professor in a shared position between the Department of Mechanical Engineering and IDSS in academic year 2023–2024. Professor. She currently is a Ciriacy-Wantrup Postdoctoral Fellow at UC Berkeley, hosted by Solomon Hsiang and the Global Policy Lab. She develops machine learning for Earth observation data. Her primary application areas are improving agricultural management and forecasting climate phenomena. She obtained her PhD in Computational and Mathematical Engineering from Stanford University in 2021, where she was advised by David Lobell.

## Faculty named to endowed chairs

Costis Daskalakis has been named the inaugural Avanessians Professor in EECS. He is a theoretical computer scientist who works at the interface of game theory, economics, probability theory, statistics, and machine learning. He has resolved long-standing open problems about the computational complexity of the Nash equilibrium, the mathematical structure and computational complexity of multi-item auctions, and the behavior of machine-learning methods such as the expectation-maximization algorithm. He has obtained computationally and statistically efficient methods for statistical hypothesis testing and learning in high-dimensional settings, as well as results characterizing the structure and concentration properties of high-dimensional distributions. His current work focuses on multi-agent learning, learning from biased and dependent data, causal inference, and econometrics. He holds a diploma in electrical and computer engineering from the National Technical University of Athens and a PhD in electrical engineering and computer science from UC Berkeley.

Manish Raghavan has been named the inaugural Drew Houston Career Development Professor. He is in a shared position between the college, the Faculty of Computer Science in EECS, and the MIT Sloan School of Management. His primary interests lie in the application of computational techniques to domains of social concern, including online platforms, algorithmic fairness, and behavioral economics, with a particular focus on the use of algorithmic tools in the hiring pipeline. Most recently he was a postdoctoral fellow at the Harvard Center for Research on Computation and Society. He completed a PhD in computer science at Cornell after receiving a BS in electrical engineering and computer science from UC Berkeley.

Faculty Spotlights

Fadel Adib SM '13, PhD '17
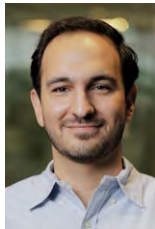
A busy commuter is ready to walk out the door, only to realize they've misplaced their keys and must search through piles of stuff to find them. Rapidly sifting through clutter, they wish they could figure out which pile was hiding the keys.

Fadel Adib SM '13, PhD '17, Doherty Chair of Ocean Utilization, associate professor and director of the Signal Kinetics Group at the MIT Media Lab and in EECS, is among a team of researchers at MIT who have created a robotic system that can do just that. The system, RFusion, is a robotic arm with a camera and radio frequency antenna attached to its gripper. It fuses signals from the antenna with visual input from the camera to locate and retrieve an item, even if the item is buried under a pile and completely out of view.

"This idea of being able to find items in a chaotic world is an open problem that we've been working on for a few years," says Adib. "Having robots that are able to search for things under a pile is a growing need in industry today. Right now, you can think of this as a Roomba on steroids, but in the near term, this could have a lot of applications in manufacturing and warehouse environments."

Fotini Christia

Community policing is meant to combat citizen mistrust of the police force. The concept was developed in the mid-20th century to help officers become part of the communities for which they are responsible. The hope was that such presence would create a partnership between citizens and the police force, leading to reduced crime and increased trust.

But a recently published study of six different sites in the Global South showed no significant positive effect associated with community policing across a range of countries.

"We found no reduction in crime or insecurity in these communities, and no increase in trust in the police," says Fotini Christia, an author of the paper, which was published in *Science*. Christia is the Ford International Professor in the Social Sciences in the Department of Political Science and the director of the Sociotechnical Systems Research Center in IDSS.

David Gamarnik PhD '98

The notion that some computational problems in math and computer science can be hard should come as no surprise. There is, in fact, an entire class of problems deemed impossible to solve algorithmically. Just below this class lie slightly "easier" problems that are less well-understood.

David Gamarnik PhD '98, Nanyang Technological University Professor of Operations Research at the MIT Sloan School of Management and IDSS and a member of the Operations Research Center, is focusing his attention on the latter, less-studied category of problems, which are more relevant to the everyday world because they involve randomness, an integral feature of natural systems. He and his colleagues have developed a potent tool for analyzing these problems called the overlap gap property. Gamarnik described the new methodology in a recent paper in the *Proceedings of the National Academy of Sciences*.

Dina Katabi SM '99, PhD '03

Dina Katabi SM '99, PhD '03 is the Thuan and Nicole Pham Professor in EECS and director of the MIT Center for Wireless Networks and Mobile Computing.

She is designing the next generation of smart wireless devices that will sit in the background of a given room, gathering and interpreting data, rather than being wrapped around one's wrist or worn elsewhere on the body.

The noninvasive device that sits in a person's home can help track breathing, heart rate, movement, gait, time in bed, and the length and quality of sleep. The device could be used in the homes of seniors and others to help detect early signs of serious medical conditions and as an alternative to wearables.

Russ Tedrake

As robots evolve, society's collective imagination forever ponders what else robots can do, with recent fascinations coming to life as self-driving cars or robots that can walk and interact with objects as humans do.

In fall 2021, EECS launched a new course, 6.800 Robotic Manipulation, to help engineering students broadly survey the latest advancements in robotics while troubleshooting real industry problems. It's a unique course that can provide an inroad into robotics for students with no robotics experience at all, designed by Russ Tedrake, the Toyota Professor of Electrical Engineering and Computer Science, Aeronautics and Astronautics, and Mechanical Engineering at MIT.

Tedrake developed the course after robotic manipulation became the new focus of his research at the Toyota Research Institute and MIT's Robot Locomotion Group. His students learn fundamental algorithmic approaches to build robot systems capable of autonomously manipulating objects in unstructured environments. The course culminates in a final project in which students can explore any problem in robotic manipulation that fascinates them.

Step 1: phonome clustering
+
semantics from online status

Step 2: pairing within clusters
* (stem, past)

phc + highest 0.8
weight

0.8

$\boxed{0.25}$

$\boxed{P\left(W_{past} \mid W_{stem}, Phon_{stem}, Sem_{stem}\right)}$

w/ estimation

$W_{past} \cdot (C_1 C_2 \cdots C_n)$ where $C \in \mathcal{E}$
$n \to \infty$

where

$I = \{(x_i, y_i)\}$ $d_{11}, d_{12}, \ldots, d_i$

Writing on a white board is Candace Ross SM '18, PhD '22, whose MIT research focused on AI problems and their connection to cognition and perception.

Students in Common Ground for Computing Education course Linear Algebra and Optimization, fall 2021.

Students

Graduate Student Spotlights

- **Monica Agrawal** is a PhD student in EECS. Her research focuses on the development of machine-learning algorithms that could unlock the potential of electronic health records to power personalized, real-world studies of comparative effectiveness. She is tackling the issue from three interconnected angles: understanding the basic building blocks of clinical text, enabling the structuring of clinical timelines with only minimal labeled data, and redesigning clinical documentation to incentivize high-quality structured data at the time of creation.

- **Sarah Cen** is a PhD student in EECS and affiliated with LIDS. She explores the interplay between humans and AI systems to help build accountability and trust. She studies algorithm design in areas including social media algorithms, the fairness of matching markets, and the impact of policy interventions in complex scenarios. Cen has worked on a wide range of projects during her time at LIDS, many of which tie directly to her interest in the interactions between humans and computational systems.

- **Martin Nisser,** a PhD student in EECS with an affiliation in CSAIL, wants to democratize hardware by making it easier to build and customize. The cost of entering space is currently too high for the average citizen, but Nisser's work may help change that. His work on self-assembling robots could be key to reducing the costs that help determine the price of a ticket. As a member of the HCI Engineering Group at CSAIL, Nisser has partnered with the MIT Space Exploration Initiative to continue studying self-assembly in space. He is particularly interested in how his research can improve sustainability and make advanced technology more affordable.

- **Erin Walk** is a PhD student in the Social and Engineering Systems program, which is part of IDSS. She studies the intersection between data science, policy, and technology. Her research has demonstrated the value and bias inherent in social media data, with a focus on how to mine those data to better understand the conflict in Syria. After she graduates, Walk plans to explore issues that have a positive, tangible impact on policy outcomes and people, perhaps in an academic lab or in a nonprofit organization. Two such issues that particularly intrigue her are internet access and privacy for underserved populations.

- **Emma Gibson PhD '22,** as a doctoral student in MIT's Operations Research Center, sought to improve patient care by helping health care facilities in sub-Saharan Africa use their limited resources more effectively. She plans to join a technology and analytics health care practice in South Africa, where her initial focus will be on public health care institutions, including a hospital in Johannesburg that is the third-largest in the world. In this role, Gibson will work to fill in gaps in African patient data for medical operational research and develop ways to use these data more effectively to improve health care in resource-limited areas.

- **Jonathan Zong SM '20** found a lack of systems that earn and maintain public trust in large-scale online research, so he made one himself. Zong, a PhD student in EECS with an affiliation in CSAIL, created Bartleby, a system for debriefing research participants and eliciting their views about social media research that involved them. Using Bartleby, he says, researchers can automatically direct each of their study participants to a website where they can learn about their involvement in research, view what data researchers collected about them, and give feedback. Most importantly, participants can use the website to opt out and request to delete their data.

Undergraduate Spotlights

- **Dolapo Adedokun '23**, after completing his degree in EECS, will travel to Ireland to undertake an MS in intelligent systems at Trinity College Dublin as MIT's fourth student to receive the prestigious George J. Mitchell Scholarship. A talented jazz guitarist, he has translated his talent for improvisation into his creative work with audio synthesis technology. He says, "Jazz, and in particular improvisation, has taught me so much about what it means to be creative: to be willing to experiment, take risks, build upon the work of others, and accept failure—all skills that I wholeheartedly believe have made me a better technologist and leader."

- **Emily Caragay '22**, who majored in EECS, was selected by TPP this year as MIT's representative to the Presidential Fellows Program run by the Center for the Study for the Presidency and Congress, which aims to inspire students to become public servants. Next year, she will be pursuing her master's in computer science at MIT. Her research project is going to be at the intersection of computer science and policy, looking at how to make software more accountable to policy and law.

- **Julia Caravias '22** majored in computer science, economics, and data science and minored in sustainability and statistics. Awarded a Fulbright Fellowship, she will conduct research at the Technical University of Munich in Germany on optimization models for the efficient allocation of electric vehicle charging stations. At MIT, Caravias pursued her interest in using computational models to inform sustainable development through her research at the Environmental Solutions Initiative and D-Lab. She was also a member of the varsity women's volleyball team, Terrascope, and Alpha Phi sorority.

- **Zoë Marschner '23** is a third-year student majoring in computer science and math and was awarded a Goldwater Scholarship for the 2022–2023 academic year. She has worked in the Geometric Data Processing Group with Professor Justin Solomon. Of Marschner's research, Solomon says, "her work is easily on par with that of my most advanced graduate students." Marschner describes geometry processing as "the coolest field in the world, since it lets you take interesting results from math and apply them to interesting problems in computer science. People see graphics all the time, such as in digital drawing programs or Pixar's 3-D animated movies, so it's really a very far-reaching field."

- **David Li '22** has been awarded a 2022 Marshall Scholarship to pursue graduate studies at Cambridge and Oxford universities in England. At MIT, he majored in EECS with minors in mechanical engineering and economics. Working in the Zhang Lab at the Broad Institute, which develops molecular and cellular tools for manipulating biological systems, Li helped engineer CRISPR technologies for genome editing as well as new approaches for directed evolution and Covid-19 testing. As a Marshall Scholar, he will complete an MPhil in biological science at the MRC Laboratory of Molecular Biology through Cambridge, and then an MS in neuroscience at Oxford. On his return, he intends to pursue a PhD in bioengineering.

A ceremony on Briggs Field on May 26, 2022, honored graduates of master's and doctoral programs in the MIT School of Engineering and Schwarzman College of Computing.

■ **Syamantak Payra '22**, recipient of a 2022 Paul and Daisy Soros Fellowship for New Americans, graduated this spring with majors in electrical engineering and computer science, and minors in public policy and entrepreneurship and innovation. At MIT, he worked in the research lab of Professor Yoel Fink PhD '00, creating digital sensor fibers to be woven into intelligent garments that can assist in diagnosing illnesses, and in the lab of Professor Joseph Paradiso PhD '81, contributing to next-generation space-suit prototypes that better protect astronauts on spacewalks. He will pursue a PhD in engineering at Stanford University, with the goal of creating new biomedical devices that can help improve daily life for patients worldwide.

Commencement

On May 27, the MIT graduating class of 2022 gathered for the first in-person Commencement exercises in three years. President L. Rafael Reif told graduates, "After you depart for your new destinations, I want to ask you to hack the world — until you make the world a little more like MIT: More daring and more passionate. More rigorous, inventive, and ambitious. More humble, more respectful, more generous, more kind. And because the people of MIT also like to fix things that are broken, as you strive to hack the world, please try to heal the world, too."

On May 26, graduates of master's and doctoral programs in the MIT School of Engineering and Schwarzman College of Computing were honored in a ceremony on Briggs Field. Dario Gil SM '00, PhD '03, senior vice president and director of research at IBM, joined as the inaugural speaker.

Engaging Global Leadership

AI Policy Forum

The AI Policy Forum (AIPF) is an undertaking by the MIT Schwarzman College of Computing to bring together leaders in government, industry, academia, and civil society from around the world to develop approaches to deal with the societal challenges posed by the rapid advances and increasing applicability of AI. The overarching goal of the AIPF is to bridge deep, technical understanding of AI and its challenges with the practice and tradeoffs of governing.

Early this year, the AIPF "Deep Dive" convened governmental officials from around the world as well as MIT and external faculty to participate in a series of virtual meetings under Chatham House Rules. The goal was to discuss the bidirectional learning on AI policy with a focus on concrete AI policy challenges and specific policy tradeoffs while gathering input that would inform the direction of future AIPF work and events. This was an opportunity to exchange knowledge and perspectives between scientists and a variety of policy makers from around the world.

In May, the AIPF also put together the second AI Policy Forum Symposium that featured an array of panelists covering cross-sector topics such as the design of AI laws and AI auditing, as well as sector-specific issues like AI in health care and mobility.

The AIPF's ongoing work as well as its task forces are developing thoughtful policy responses to complex AI policy questions. Specifically, through its task force on AI auditing, the AIPF is exploring how AI monitoring can be accomplished at scale, determining what government infrastructure, human capital, and other systems might be necessary to achieve effective execution of this burgeoning government function. In its work on social media, the AIPF is analyzing possible reforms in this high-profile sector, including potential taxation regimes to curb negative externalities and ways to govern social media boosting (the content that platforms promote on user feeds). As part of its finance task force, the AIPF is planning a workshop on AI in consumer lending to better understand potential alternatives to explainable machine learning in consumer lending regulations. Overall, these examples reflect the AIPF's approach to its work, which is to ask focused questions in specific subject areas in order to make meaningful headway towards substantive policy solutions.

The AI Policy Forum is a global collaboration to move AI principles to practice.

Climate Implications of Computing and Communications workshop

The voracious appetite for energy from the world's computers and communications technology presents a clear threat for the globe's warming climate. That was the blunt assessment from presenters in an intensive two-day Climate Implications of Computing and Communications workshop in March hosted by MIT's Climate and Sustainability Consortium, MIT-IBM Watson AI Lab, and the Schwarzman College of Computing.

Information and communications technology already account for more than two percent of global energy demand, which is on a par with the aviation industries' emissions from fuel.

"If we continue with the existing trajectory of compute energy, by 2040, we are supposed to hit the world's energy production capacity," says Bilge Yildiz PhD '03, the Breene M. Kerr Professor in the MIT departments of nuclear science and engineering and materials science and engineering, a presenter at the workshop.

Top: Clockwise from upper left: Dean Dan Huttenlocher moderates an AIPF symposium on AI laws with panelists Jonathan Zittrain, Eva Kaili, and Bitange Ndemo. Bottom: Panelists at an AIPF symposium on auditing and monitoring AI systems were, clockwise from upper left, Aleksander Mądry, Diane Lye, Renaud Vedel, Elham Tabassi, and moderator Luis Videgaray.

Drew Houston Fireside Chat

When the cloud storage firm Dropbox decided to shut down its offices at the onset of the Covid-19 pandemic, cofounder and CEO Drew Houston '05 had to send the company's nearly 3,000 employees home and tell them they were not coming back to work anytime soon.

In the early days of the pandemic, Houston says, Dropbox reacted as many others did to ensure that employees were safe and customers were taken care of. "It's surreal," he says. "There's no playbook for running a global company in a pandemic over Zoom. For a lot of it we were just taking it as we go."

Houston talked about his experience leading Dropbox through a public health crisis and how Covid-19 has accelerated a shift to distributed work in a fireside chat in October 2021 with Dan Huttenlocher, dean of the MIT Stephen A. Schwarzman College of Computing. Houston also spoke about his gift to MIT to endow the first shared professorship between the MIT Schwarzman College of Computing and the MIT Sloan School of Management, as well as provide a catalyst startup fund for the college.

AI and Our Human Future

From health care to transportation to social media, the rapid adoption of AI across many fields has transformed the economic and social structures of people's daily lives. In doing so, AI has not only impacted our society in unforeseen ways, but has also altered how we experience reality.

In January, the college hosted a live-streamed conversation between Dean Huttenlocher and Eric Schmidt, former CEO and chairman of Google. The discussion, moderated by Asu Ozdaglar, deputy dean of academics for the college and MathWorks Professor of EECS, focused on how AI is changing our relationship with knowledge and society and what this technology means for us all.

A conversation between Dropbox CEO Drew Houston (left) and MIT Schwarzman College of Computing Dean Dan Huttenlocher spanned a range of topics, from the impact of Covid-19 on the workplace to Houston's gift in support of a shared professorship between the college and the MIT Sloan School of Management, along with a catalyst startup fund.

Industry Collaborations

Efficient AI and Computing Technologies Conference

The 2022 MIT Efficient AI and Computing Technologies Conference in May showcased the most recent developments and tangible impacts of computing technologies in AI hardware and software for improved efficiency and explainability, as well as their applications in diverse technological areas. The conference, presented by the MIT Schwarzman College of Computing and the MIT Industrial Liaison Program, featured insights from MIT faculty, MIT Startup Exchange entrepreneurs, and industry executives.

Amazon-MIT Science Hub

With the goal of increasing connections and fostering the growth of the research community in computing, the Science Hub launched in October 2021 as a new collaboration between MIT and Amazon. The Science Hub supports research, education, and outreach efforts in areas of mutual interest, beginning with AI and robotics in the first year.

Administered at MIT by the Schwarzman College of Computing, the Science Hub ensures that the benefits of artificial intelligence and robotics innovations are shared broadly — both through education and by advancing research — and that participation in the research is broadened to encompass diverse, interdisciplinary scholars and other innovators.

The collaboration with Amazon is enabling a broad set of programs across MIT, including annual fellowships for graduate students and postdoctoral researchers and seed funds for MIT faculty. Also supported are events and activities that accelerate AI and robotics research in ways to make it more accessible, such as research symposia that are open to other academic institutions and the public. Additional sponsored research funds are provided to stand up research projects led by MIT faculty members.

MIT AI Hardware Program

The MIT AI Hardware Program is a new academia and industry collaboration aimed at defining and developing translational technologies in hardware and software for the AI and quantum age. A collaboration between the MIT School of Engineering and MIT Schwarzman College of Computing, the cross-disciplinary effort innovates technologies that will deliver enhanced energy efficiency systems for cloud and edge computing.

Based on use-inspired research involving materials, devices, circuits, algorithms, and software, the MIT AI Hardware Program convenes researchers from MIT and industry to facilitate the transition of fundamental knowledge to real-world technological solutions. The program spans materials and devices, as well as architecture and algorithms enabling energy-efficient and sustainable high-performance computing.

The five inaugural members of the MIT AI Hardware Program are Amazon, a global technology company whose hardware inventions include the Kindle, Amazon Echo, Fire TV, and Astro; Analog Devices, a global leader in the design and manufacturing of analog, mixed signal, and DSP integrated circuits; ASML, an innovation leader in the semiconductor industry, providing chipmakers with hardware, software, and services to mass produce patterns on silicon through lithography; NTT Research, a subsidiary of NTT that conducts fundamental research to upgrade reality in game-changing ways that improve lives and brighten our global future; and TSMC, the world's leading dedicated semiconductor foundry.

Students gather at MIT for a Break Through Tech AI event in September 2022. The program helps participants gain the skills they need to get jobs in the fastest-growing areas of tech: data science, machine learning, and artificial intelligence.

Diversity, Equity, and Inclusion

The MIT Schwarzman College maintains a strong commitment to improving and creating programs and activities to help broaden participation in computing classes and degree programs, increase the diversity of top faculty candidates in computing fields, and ensure that faculty search and graduate admissions processes have diverse slates of candidates and interviews.

Perspectives and action in the college's departments and centers

Break Through Tech AI

Aimed at driving diversity and inclusion in AI, the Schwarzman College of Computing launched Break Through Tech AI, a new program to bridge the talent gap for women and underrepresented genders in AI positions in industry.

Break Through Tech AI provides skills-based training, industry-relevant portfolios, and mentoring to qualified undergraduate students in the Greater Boston area in order to position them more competitively for careers in data science, machine learning, and AI. The free, 18-month program also provides each student with a stipend for participation to lower the barrier for those typically unable to engage in an unpaid, extracurricular educational opportunity.

The program kicked off in summer 2022 with an eight-week, skills-based online course and in-person lab experience that teaches industry-relevant tools to build real-world AI solutions. Students learned how to analyze datasets and use several common machine-learning libraries to build, train, and implement their own machine-learning models in a business context.

Students are matched with machine-learning challenge projects for which they convene monthly at MIT and work in teams to build solutions and collaborate with an industry advisor or mentor throughout the academic year, resulting in a portfolio of résumé-quality work. The participants are also paired with young professionals in the field to help build their network, prepare their portfolio, practice for interviews, and cultivate workplace skills.

33

Students taking part in first year of Break Through Tech AI

The third annual GW6 Research Summit in November 2021 featured the latest research of women and other underrepresented genders in EECS at MIT. Panelists included MIT President Emerita Susan Hockfield (upper right).

Thriving Stars

In October 2021, EECS announced the Thriving Stars initiative to improve gender representation in MIT's largest doctoral graduate program.

"All types of representation are vital to EECS at MIT, and Thriving Stars will unify multiple disparate efforts focusing on women and other underrepresented genders," says EECS head Asu Ozdaglar, MIT Schwarzman College of Computing deputy dean of academics, and the MathWorks Professor of EECS.

Thriving Stars takes a holistic and concerted approach to achieving this goal by providing support and information to students throughout every step of their PhD journey: from recruitment to admission all the way to graduation. The initiative focuses on efforts to navigate the application process and showcase research opportunities, interdisciplinary collaborations, and internships, as well as the diverse array of career opportunities accessible to doctoral EECS graduates.

GW6 Research Summit

The EECS group Graduate Women in Course 6 (GW6) held its third annual research summit in November 2021. The summit featured the latest research of women and other underrepresented genders in MIT EECS, along with an opportunity to network, share experiences, and learn.

Much of the summit was devoted to a whirlwind of three-minute "lightning" talks, with graduates presenting mini-crash courses in their latest research pursuits. Talks spanned a wide range of research areas, from photonics to wireless communication to machine learning, reflecting the department's diverse spectrum of electrical engineering, computer science, and AI and decision-making disciplines. In addition to technical talks, the summit also hosted a panel on navigating research careers.

Research Initiative on Combating Systemic Racism

An MIT-wide effort launched by IDSS uses social science and computation to address systemic racism.

The IDSS Research Initiative on Combating Systemic Racism works to bridge the social sciences, data science, and computation by bringing researchers from these disciplines together to address systemic racism across housing, health care, policing, education, employment, and other sectors of society.

The initiative aims to create a visible presence at MIT for cutting-edge computational research with a racial equity lens and across societal domains that will attract and train students and scholars. It is led by Fotini Christia, the Ford International Professor in the Social Sciences in the Department of Political Science, associate director of IDSS, and director of the Sociotechnical Systems Research Center, and Munther Dahleh, IDSS director and William A. Coolidge Professor of Electrical Engineering and Computer Science.

The steering committee for this research initiative is composed of underrepresented minority faculty members from across MIT's five schools and the MIT Schwarzman College of Computing. Members will serve as close advisors to the initiative as well as share the findings of their work beyond MIT's campus. MIT Chancellor Melissa Nobles heads the committee.

MEnTorEd Opportunities in Research (METEOR)

METEOR is a postdoctoral fellowship program sponsored by MIT CSAIL to support exceptional researchers in computer science and AI and to broaden participation in the field. Fellows conduct cutting-edge research as part of a vibrant community and are hosted by research groups with synergistic scientific interests.

The goal of METEOR is to enable outstanding individuals to do research in the lab and prepare for careers in academia or industry. The program also features in-depth opportunities for mentorship from researchers in both CSAIL and industry. Postdoctoral researchers selected are awarded a one-year funded appointment (renewable for a second year). It is hoped this multiyear program will grow and broaden beyond CSAIL.

The program welcomed two new participants this past year who are staying on for a second year. Four additional postdocs will be joining the program during the 2022–2023 academic year. The cohort includes five women, among them one Latina as well as one African American. These postdocs are paired with CSAIL principal investigator mentors and have the opportunity to attend seminars and receive training on the application process for academic and industrial careers. They also participate in group activities with their cohort and the larger lab.

A New Hub for Computing Excellence

Designed to be a physical manifestation of the college's mission, the MIT Schwarzman College of Computing building will engage faculty, students, and collaborators across a broad array of disciplines in computing and AI education, research, and innovation, and will serve as a foundation for computing breakthroughs at MIT.

Foundation work was completed over the winter and a tower crane erected on the project site. From there, the steel went up and the last beam was placed in the building structure this spring. This milestone in the evolution of this building project was marked with a topping-off ceremony that took place on May 23. The project remains on track, with completion anticipated in summer of 2023.

For a tour of the new space, visit: betterworld.mit.edu/scctour

Left: Lobby construction in progress.
Right: A rendering by Skidmore, Owings & Merrill of the space when completed.

Thank You

Informed and inspired by wisdom across disciplines, the MIT Schwarzman College of Computing is fostering breakthroughs in AI, delivering the power of AI tools to researchers in every field, and advancing pioneering work on AI's ethical use and societal impact. The college is equipping students to be as fluent in computing and AI as they are in their own fields of study, and ready to use these digital tools wisely and humanely to help make a better world. I am deeply proud of all the college is doing to keep MIT — and the nation — on the leading edge of computing and AI, and profoundly grateful to the alumni and friends whose generosity supports this vital work.

L. Rafael Reif
MIT President

MIT Schwarzman
College of Computing