

Computer Science and Artificial Intelligence Laboratory

The [Computer Science and Artificial Intelligence Laboratory \(CSAIL\)](#) pioneers approach to computing that improve how people work, play, and learn. MIT CSAIL serves the MIT community, the country, and society at large by creating a positive future enhanced by computer science through contributions of ideas, artifacts, and people:

- Integration of computing into the fabric of everyday life
- Acceleration of discovery in engineering, sciences, humanities, arts, and society
- Development of novel paradigms, models, techniques, software, and machines
- Creation of mathematical and scientific foundations for computing
- Training students and postdoctoral researchers pursuing computer science research

Lab members conduct research in almost all aspects of computer science, including artificial intelligence (AI), the theory of computation, systems, machine learning (ML), and computer graphics, as well as exploring revolutionary new computational methods for advancing health care, manufacturing, energy, and human productivity.

CSAIL has a long history of technological innovations that have affected how people interact and do business. CSAIL is known as the incubator for some of the greatest technological advances of the past 30 years that were true life-changers, including the internet, personal computing, mobile computing, open source software, microprocessors, robotic surgery, and social networking.

CSAIL's current research addresses some of the grand challenges of the 21st century, including developing personalized learning, securing cyberspace, advancing health informatics, reverse engineering the brain, enhancing virtual reality, developing tools for scientific discovery, improving urban infrastructure, and ensuring the health of our environment. Computing is central to solving these challenges and CSAIL contributes to making computing more capable by addressing fundamental algorithmic and systems questions at the core of computing, and broadening the scope of computing to address the important social challenges that confront us.

Key CSAIL initiatives currently underway include tackling the challenges of big data, developing new models for wireless and mobile systems, securing computers and the cloud against cyberattacks, rethinking the field of artificial intelligence, and developing the next generation of robots. Advanced, software-based medical instrumentation and medical informatics systems to aid clinical decision making is being investigated. Advancements in biological research are also under way, including developments in the field of computational biology and the application of machine learning to the interpretation of complete genomes and understanding gene regulation.

CSAIL leadership spearheaded several other efforts that benefited the broader MIT community, they are as follows:

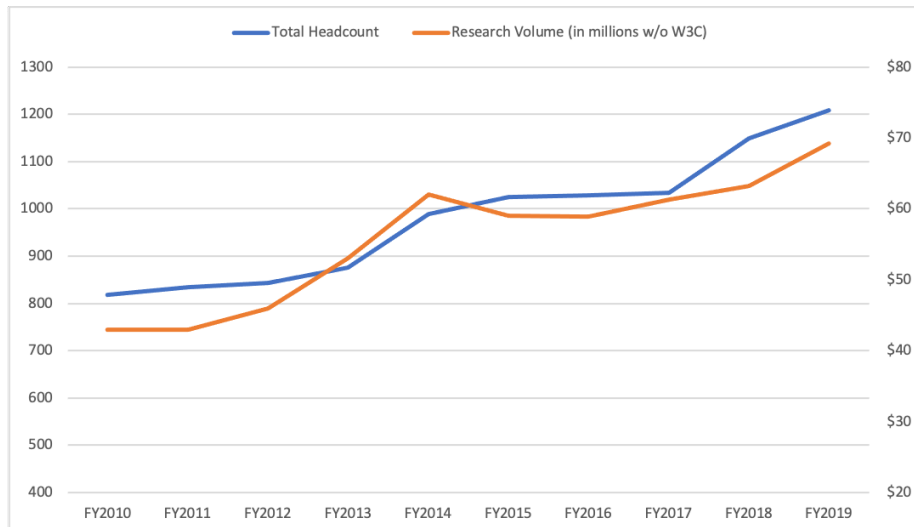
Achievements

CSAIL Initiatives and Corresponding Impact, FY2019

Initiatives	Impact
Communication: created the resources to support the web presence of principal investigators (PIs)	Modern and compelling web stories that bring the work of PIs to the world
Space: created processes to enable community growth and improved the common spaces through new furniture; also created a new CSAIL lobby installation	The new spaces are more inviting and help build community.
Advisory Board Meeting: hosted the yearly advisory board meeting	Feedback from the advisors is for CSAIL to remain a rock of stability to provide research excellence which in turn will help the MIT Stephen A. Schwarzman College of Computing succeed; continue to pursue exciting industry collaborations
Lab-wide research initiatives: iFlyTek (new), Microsoft (new), Air Force AI Accelerator (new), Quanta (renewal), Wistron (renewal)	Launched funded initiatives to address basic research challenges in AI: resilience and unsupervised learning
International symposia: TEDxMIT; The Second AI; and the Future of Work Symposium (with MIT Sloan School of Management's Institute on the Digital Economy)	These events were campus collaborations to raise visibility through presented research; TEDxMIT featured 12 great MIT women researchers
Hot Topics in Computing	This MIT-wide forum brought students and faculty together to discuss important current issues in computing: self-driving cars, fake news, privacy, security, AI, and diversity
Sustainable dining: eliminated paper and plastic products from CSAIL events	The use of china and silverware reduces the garbage generated by CSAIL
Stephen A. Schwarzman College of Computing planning	Held CSAIL focus groups to plan how to evolve CSAIL's internal organization to support governance, identity, and agility in research at a level in between CSAIL (which is too large) and individual PIs (which is too narrow)

CSAIL Growth

Over the past 10 years CSAIL has had a 48% growth in headcount, and a 61% growth in research volume.



With a total combined research volume (primary, secondary, and the World Wide Web Consortium) of \$72,817,580 for FY2019, CSAIL is the Institute's largest interdisciplinary research laboratory. CSAIL continues to have the highest research volume among MIT Interdisciplinary Departmental Laboratories. CSAIL's growth over the last year has been relatively steady, culminating with a 9% growth between FY2018 and FY2019. CSAIL manages over 500 active research awards and over 110 PIs with appointments across 11 MIT departments. Through AY2019 we had 554 graduate students with residential advisor (RA) appointments in CSAIL, and 236 Undergraduate Research Opportunities Program (UROP) students.

CSAIL research is sponsored by a large number of diverse sources, from US government contracts to the private sector. United States government sponsors include the following:

- Air Force Research Laboratory and the Air Force Office of Scientific Research
- Army Research Office
- Defense Advanced Research Project Agency
- Department of Defense Research and Engineering
- Department of Education
- Department of Energy
- Food and Drug Administration
- Intelligence Advanced Research Projects Activity, National Institutes of Health
- National Institute of Justice
- National Science Foundation
- Space and Naval Warfare Systems Center
- US Navy (including the Office of Naval Research and Naval Air Systems Command)

US and international nonfederal sponsors include the following:

- Accenture
- Advanced Technology Laboratories
- Aptima
- BAE Systems
- BBN Technologies
- Boeing
- BMW of North America
- Delta Electronics Foundation
- Ford Motor Company
- Foxconn Technology Group
- Honda R&D
- IBM
- iFlyTek
- Intel Corporation
- Jaguar Land Rover Limited
- JD.com
- Lockheed Martin
- Microelectronics Advanced Research Corporation
- Mitsubishi Electric Corporation
- National ICT Australia Limited
- Nippon Electric Company
- Nippon Telegraph and Telephone Corporation
- Nissan Motor Company

- Northrop Grumman Corporation
- Omron
- Pfizer
- Ping An Technology
- Qatar Computing Research Institute
- Quanta Computer
- Rakuten
- Raytheon
- Steelcase
- Systems & Technology Research
- Takeda
- Samsung Electronics
- Siemens
- Suzhou Industrial Park
- Toyota Research Institute
- Wistron Corporation

Other organizations sponsoring research include the following:

- Aarhus University
- Battelle Memorial Institute
- DSO National Laboratories
- Epoch Foundation
- Hong Kong University of Science and Technology
- Industrial Technology Research Institute
- Nanyang Technical University
- Singapore-MIT Alliance

Research Projects

Within CSAIL we have many single- and multi-investigator projects, as well as a number of virtual centers and large-scale projects. The large-scale projects and collaborations include the following:

Air Force AI Innovators

MIT and the United States Air Force established an Artificial Intelligence Innovation Accelerator (AIIA) program to achieve breakthroughs in computational intelligence. AIIA will support basic research on advanced AI algorithms, systems, and capabilities, and a wide variety of applications with positive impact on society and national defense. CSAIL signed an agreement for \$75 million on May 19, 2019, and commenced a request for proposals.

Toyota-CSAIL Joint Research Center

Today, a car crash occurs every 5 seconds in the United States. Globally, road traffic injuries are the eighth leading cause of death, with about 1.24 million lives lost every year. In addition to this terrible human cost, these crashes take an enormous economic toll. The National Highway Traffic Safety Administration has calculated the economic cost in the United States at about \$277 billion per year. Putting a dent in these numbers is an enormous challenge—and it's one that is motivating the research of the Toyota-CSAIL Joint Research Center, which was kicked off in September 2015. The center is in collaboration with the Toyota Research Institute (TRI) led by Gill Pratt.

Imagine if your car could sense if you were having a bad day and turned on your favorite album to improve your mood. Imagine if your car could communicate with your refrigerator, figure out that you are out of milk, and suggest where to stop on your way home to buy more. Imagine if your car knew that you forgot to call your parents yesterday and issued a gentle reminder on the way home. And making that call was easy because you could turn the driving over to the car on a boring stretch of highway. These are just a few of the possibilities when we bring together cars and computer science; topics motivating the research at the Toyota-CSAIL Joint Research Center.

The objective of the Toyota-CSAIL program is to advance AI and robotics research, develop a safe and intelligent car, and improve mobility and transportation by advancing the science of autonomy and machine intelligence. The CSAIL researchers are working on new tools for collecting and analyzing navigation data with the objective to learn from humans; perception and decision-making systems for safe navigation; systems that can handle difficult driving situations: congestion, high-speed driving, and inclement weather; predictive models that can anticipate the behavior of humans and vehicles; more intelligent user interfaces; and human-support robots, whereby robots can work safely and seamlessly with and around humans in performing tasks, establish effective communication with humans, and recognize human intent.

More specifically, the projects and principal investigators active in the Toyota-CSAIL Joint Research Center during AY2019 are as follows:

- Geordi: A Driver's Assistant for Risk-Bounded Maneuvering (PI: Brian Williams)
- Driver-Friendly Bilateral Control for Suppressing Traffic Instabilities (PI: Berthold Horn)
- Machines that Can Introspect (PIs: Nick Roy and Boris Katz)
- Uhura: A Driver's Personal Coach for Managing Risk (PI: Brian Williams)
- Understanding Human Gaze (PIs: Antonio Torralba and Wojciech Matusik)
- Exploring the World of High Definition Touch (PIs: Edward Adelson and John Leonard)
- Formal Verification Meets Big Data Intelligence to Address the Trillion Miles Challenge (PI: Armando Solar-Lezama)
- The Car Can Explain! (PIs: Gerald Sussman, Lalana Kagal, Daniel Weitzner, and Hal Abelson)
- Crossing the Vision-Language Boundary (PIs: James Glass and Antonio Torralba)
- Analysis by Synthesis Revisited: Visual Scene Understanding by Integrating Probabilistic Programs and Deep Learning (PI: Josh Tenenbaum)
- Wi-Fi-Based Obstacle Detection for Robot Navigation (PIs: Dina Katabi and Daniela Rus)

- Using Deep Learning to Speed Up Deep Learning (PIs: Saman Amarasinghe and Fredo Durand)
- Decision Making for Parallel Autonomy in Clutter: Addressing Intent, Interactions, Rules of the Road, and Safety (PIs: Daniela Rus and Sertac Karaman)
- A Parallel Autonomy System: Data-Driven and Model-Based Parallel Autonomy with Robustness and Safety Guarantees (PIs: Sertac Karaman and Daniela Rus)
- Driver Perception and the Car-to-Driver Handoff (PI: Ruth Rosenholtz)
- Simulation and Verification for Vision-in-the-Loop Control (PI: Fredo Durand)
- Tools and Data to Revolutionize Driving (PIs: Sertac Karaman and Daniela Rus)
- Robotic Manipulation Data Engine (PI: Alberto Rodriguez)
- Dense, Freeform Tactile Feedback for Manipulation and Control (PI: Wojciech Matusik)
- Sensible Deep Learning for 3D Data (PI: Justin Solomon)
- All Terrain Mobility and Navigation (PI: Sangbae Kim)
- Inner Vision: Camera Based Proprioception for Soft Robots (PIs: Edward Adelson and Daniela Rus)
- A Safety Interlock for Self-Driving Cars (PIs: Daniel Jackson and Armando Solar-Lezama)
- Automation for Everyone (PI: Brian Williams)

For more information on the center's projects and publications, visit the [Toyota-CSAIL Joint Research Center's website](#).

Wistron-CSAIL Research Collaboration

Good health—both mental and physical—is one of the most pressing social and economic issues of the day. A healthier population makes for a happier society and a more productive economy. Today, people are surrounded by an explosion of sophisticated and increasingly affordable information devices, from laptop computers, e-book readers, and smart glasses, to mobile phones, smart watches, and fitness trackers. We monitor stock prices, weather forecasts, and traffic patterns through websites and apps, share our thoughts and experiences through emails, on Facebook and Twitter, and increasingly learn within online communities. These technologies open so many new opportunities for improving how we live, work, and play. But how do they empower us, and at what costs? Recent studies show that we consume 11 to 14 hours of technology each day. And this often involves multitasking, which in turn retrains our brains, reduces concentration, and increases stress (e.g., studies show that the brains of heavy technology users show similar patterns to those who suffer from substance abuse disorders.) Finding ways to reduce stress and technology's negative impact on a workforce is critical to our well-being in the future.

The multiyear research program between Wistron and CSAIL focuses on rethinking how we compute and communicate in the digital age to ensure that health and well-being are at the core of our lives, and our use of technology accelerates this objective. Some of the questions we pose and the answers we seek include: How to design the next generation of computers and communication systems to minimize our body's exposure to electromagnetic radiation? How should we rethink computer and communication architectures for sustainability? How to develop systems that deliver appropriate lighting? How to develop systems that reengineer email? How to develop algorithms that can help with information overload? How to use computing and communication in support of individual and community well-being? And, how to build computer and communication systems that are friendlier to our environment?

Our vision is to develop new computing and communication hardware and software platforms and supporting algorithms for modeling, controlling, and making decisions that will bring wellness to our use of technology. One thrust of this program focuses broadly on the computer and communication platforms. The second thrust focuses on using these novel platforms to promote healthier living. More specifically, the three projects that are currently active in the Wistron-CSAIL Research Collaboration are:

- Individual Prediction and Interpretation of Risk: Predicting Trajectories of Chronic Disease and Recovery (PIs: Polina Golland and Peter Szolovits)
- Smart Homes that Monitor Breathing, Heart Rate, and Life Quality (PI: Dina Katabi)
- Using Machine Learning to Build Better Clinical Support Tools (PI: John Guttag)

Microsoft-CSAIL Trustworthy AI Research Collaboration

The mission of the Microsoft-CSAIL Trustworthy AI Collaboration (TRAC) is to extend understanding and to develop methods and tools for fielding AI technologies that are robust, secure, reliable, understandable, and safe to deploy. This call solicits proposals from teams that comprise both CSAIL and MSR researchers and aim to pursue this mission. Topics of particular interest are:

- Robust ML: Development of training techniques that are able to deliver ML models that are more resistant to input perturbations and distribution shifts
- ML (for) Security: Research that lies at the intersection of systems security and machine learning; this includes both addressing the security aspects of ML system deployment and applying ML methods to traditional computer security problems
- Understanding and Resisting ML Attacks: Research on identifying and addressing potential vulnerabilities of ML systems to malicious manipulations; work includes addressing the safety, security, and privacy challenges that arise in the context of ML systems deployed on edge devices
- Human-ML Collaboration: Approaches to making ML models more understandable to humans; tools for enabling human-ML teaming

- **Reliable Reinforcement Learning:** Design of (deep) reinforcement learning primitives that are dependable and deliver policies that come with reliability guarantees
- **ML for Health Care:** Development of ML tools and requirements that are suitable for use in high-stakes and complex decision making, and exploring these issues in the context of health care
- **Safety and Robustness in the Open World:** Approaches to learning in an open environment in a way that is able to account for potential biases and blind spots
- **Policy Perspective on ML:** Work on devising best practices and laws around the responsible development and fielding of ML technologies, including potential regulation of the uses

The current active projects are as follows:

- **Robustness Meets Algorithms** (PI: Ankur Moitra)
- **Explainability and Interpretability at Scale** (PI: Stefanie Jegelka)
- **Uncertainty and Robustness at Scale** (PI: Aleksander Madry)
- **Machine Learning with Theoretical Guarantees** (PI: Tamara Broderick)
- **Compression for Interpretability at Scale** (PI: Daniela Rus)

Qatar Computing Research Institute

In 2012, CSAIL began a \$35 million research collaboration with Qatar Computing Research Institute (QCRI) to collaborate on a wide-range of research topics in computer science. Originally a seven-year program, the program was extended for an eighth year in 2018. In the seventh year of the program, AY2019, the program included the following five projects:

- **Audio-Visual Facial Reanimation** (Professor Wojciech Matusik): This project focuses on the generation of video facial expression given an audio signal.
- **Arabic Speech and Language Processing** (Senior Research Scientist James Glass): This project aims to develop speech and language processing technologies that will support natural interaction via spoken language.
- **Database Management** (Adjunct Professor Michael Stonebraker and Professor Samuel Madden): This project investigates three data management tasks, build an end-to-end system (Data Civilizer) to support the data discovery and preparation needs of data scientists; study resource elasticity in online transaction processing database management systems; and use program synthesis for entity resolution and copy detection.
- **Understanding Health Habits from Social Media Pictures** (Professor Antonio Torralba): The major goal of the project is to understand food habits from social media images, including, training machine learning models for image auto-tagging and content extraction from noisy hashtags; predicting population-level

health statistics in the United States and Qatar; monitoring temporal and regional trends in food consumption and its implications; and learning models that can achieve in-depth analysis of food images through the use of large-scale cooking and recipe data collected from the web.

- **Accurate Map Making Using Mobile Sensor Data** (Professor Hari Balakrishnan, Professor Samuel Madden, Assistant Professor Mohammad Alizadeh, and Adjunct Professor David DeWitt): This project aims to develop accurate map-making techniques using crowd-sourced methods to overcome challenges related to creating and maintaining street maps, especially in a rapidly developing environment such as Doha, Qatar, leveraging data primarily from mobile phones and investigating current limitations due sensor noise, outages, and data sparsity.

Quanta-CSAIL Research Collaboration

Quanta and CSAIL started a new, five-year, \$12.5 million collaboration. In this phase of the collaboration the parties will be focusing on long-term research to create future platforms focusing on the delivery of improved health care using computer science in general, and AI in particular.

The group is currently pursuing the following three projects:

John Guttag—Using Machine Learning to Curb Infectious Disease

Curbing infectious disease requires a multiprong attack, including:

- The development of new vaccinations and treatments
- Improved public health measures
- Better health care at all levels from primary to acute care

Machine learning can be used to develop new knowledge and algorithms that can assist with all of these aspects. The work proposed here is aimed at the second and third aims. We propose to use machine learning to build models that will further our understanding of what makes people more or less susceptible to different kinds of infections, and how different kinds of infections are spread. We will then design and test interventions based on these models.

Regina Barzilay—Learning to Assess Breast Cancer Risk to Enable Early Detection and Prevention

Our goal is the development and clinical deployment of machine learning-based models for risk assessment and early detection of breast cancer. This includes both the first occurrence of the disease, and its recurrence in both local and metastatic form. The detection will be informed by rich information about a patient including their imaging, genomics sequencing, and other test results. Beyond just predicting the occurrence of the disease, we are interested to characterize its expected severity and biological characteristics. This would enable physicians to pursue aggressive treatment strategies, such as chemoprevention and detect cancer in its earliest stage.

Collin Stultz—Revolutionizing the Care of Patients with Cardiovascular Disease

The intersection of modern cardiovascular medicine and information technology presents an opportunity to establish a new paradigm for the treatment of patients with cardiovascular disease. During the course of this work, we will therefore develop and implement new machine learning methods that will make personalized cardiovascular care a reality. We begin by developing algorithms that use all of the information in the medical record to identify patients at high risk of adverse cardiovascular events, including episodes of congestive heart failure, myocardial infarction (i.e., a heart attack), and cardiovascular death. Furthermore, we recognize that it is not enough to simply identify high-risk patients or to warn caregivers of impending adverse events. Risk stratification and early warning systems are most powerful when coupled with intelligent, patient-specific therapeutic recommendations. We will therefore use sophisticated machine learning methods to build models that suggest personalized therapeutic interventions that minimize adverse outcomes in patients with cardiovascular disease. Moreover, we strive to build models that not only have predictive power, but that also provide clinically meaningful insights that help explain how each model arrives at a particular result. This proposal represents a collaborative effort between MIT and Massachusetts General Hospital and seeks to transform the way in which clinical care is provided to patients with cardiovascular disease.

Brought on by the confluence of knowledge, computing, and data, health care worldwide is undergoing a revolution. Our collective knowledge of human health, diseases, and treatments are improving, and yet the number of sick people is still increasing by leaps and bounds.

Computing in general, and AI in particular, can help. The vast amount of data that we can collect and interrogate are rapidly increasing. The confluence of big data, cloud computing, and machine learning techniques offer us the hope that we may be able to drastically improve the odds of disease eradication and human suffering.

The primary focus of this collaboration between Quanta Computer and MIT CSAIL addresses many aspects of this crisis. First, we will explore how computation can enable researchers to significantly improve the diagnoses and treatments of some of the most pervasive diseases, such as cancer. Second, we will apply AI techniques to improve the institution-specific delivery of patient care. Third, we will investigate how the interface between a patient and a primary care physician can be streamlined, so that the data collection and resulting analytics can be more reliable.

While the focus of this collaboration will largely be defined by the troika of patient, hospital, and doctor, other aspects, such as privacy and security, nutrition, and self-care may also be investigated as time and resources permit.

iFlyTek—CSAIL Research Collaboration

In 2018 CSAIL launched a five-year collaboration with iFlyTek, in the field of artificial intelligence and natural language processing.

We are currently pursuing the following projects:

- Toward Unsupervised Speech Processing (PI: James Glass). Much of the recent progress in artificial intelligence has been powered by the use of large annotated datasets. In these situations, training data comes in the form of input-output pairs, such as a speech waveform being paired with an orthographic transcript. Machine learning models are able to use these data to learn their parameters to optimize an overall likelihood or objective function. Deep learning models are but the most recent approach that uses this supervised training framework in areas ranging from automatic speech recognition, machine vision, natural language processing, machine translation, and so forth.
- Inadvertent Interfaces and Their Use in Inferring Cognitive Status (PI: Randall Davis)
- A Framework for Building More Humanlike AI: Integrating Insights from Cognitive Development, Cognitive Neuroscience and Probabilistic Programs (PI: Josh Tenenbaum)

Industrial Outreach

CSAIL Alliances

The CSAIL Alliance Program

The CSAIL Alliance Program (CAP) is a gateway into the lab for industry, organizations, and governmental institutions seeking a closer connection to the work, researchers, and students of CSAIL. The program provides a proactive and comprehensive approach to developing strong connections with all CSAIL has to offer. Leading organizations come to CSAIL to learn about our research, to recruit talented graduate students, and to explore collaborations with our researchers. Through this program, we are able to better provide our members with access to our latest thinking and our deep pool of exceptional human and informational resources. Overall, CAP supports the mission of CSAIL by connecting our researchers, students, and technological advances to industry and organizations across the globe.

Levels of membership: Student engagement, affiliate, and partner

The CAP program provides a proactive and comprehensive approach to connect members to the whole lab—all 60 research groups, spanning robotics, natural language processing, networks, databases, cryptography, web science, and more. CAP has three levels: Student Engagement—focused on connecting with students and postdocs for career opportunities; Affiliate—which provides lab visits, access to the annual meeting, recruiting assistance, research briefings, and professional education discounts; and, Partner—which includes all of the benefits of the Affiliate level as well as more expanded options with added access to research initiative meetings, custom faculty-led seminars, and expanded recruiting options.

Member companies

Currently there are over 80 member companies, including global brands such as Apple, BASF, Google, Samsung, JP Morgan, NASDAQ, and Microsoft. Members are headquartered in North America, South America, Europe, and Asia, and represent a wide variety of industry verticals.

Online Courses and Professional Development

CSAIL Alliances also produces and manages online professional development courses in partnership with MIT's Professional Education, MIT's Office of Digital Learning, edX, Get Smarter, and Harvard Extension School. The following is a list of the programs to date, a brief description, number of offerings to date, and total enrollments to date. Total enrollment is now approximately 35,000 online learners:

Enrollments in CSAIL-Produced Online Courses

Course title	Description	Offered	Enrolled
Artificial Intelligence: Implications for Business Strategy	Offered in partnership with the Sloan School of Management, this course focuses on the organizational and managerial implications of AI technologies.	19	10,974
Machine Learning: Implementation in Business	This course is offered in partnership with the Sloan School of Management and aims to demystify machine learning for the business professional—offering a firm, foundational understanding of the advantages, limitations, and scope of machine learning from a management perspective.	3	341
Tackling the Challenges of Big Data*	Survey state-of-the-art topics in big data, looking at data collection (smartphones, sensors, the web), data storage, and processing (scalable relational databases, Hadoop, Spark, etc.), extracting structured data from unstructured data, systems issues (exploiting multicore, security), analytics (machine learning, data compression, efficient algorithms), visualization, and a range of applications.	10	11,431
Tackling the Challenges of Big Data—Taiwan*	The original course, translated into traditional Chinese.	1	1,296
Tackling the Challenges of Big Data—Illumno	The original course, translated into Spanish and Portuguese, and offered through Illumno in collaboration with universities in South America.	3	1,032
Cybersecurity: Technology, Application, and Policy*	This six-week online course provides a holistic look at cybersecurity technologies, techniques, and systems.	8	3,233
HCI: Human Computer Interaction for User Experience Design	The six-week course was produced in partnership with Get Smarter and includes eight CSAIL researchers who review a host of cutting-edge HCI concepts, including voice activated and spatially-aware computers, as well as speech and vision tools.	8	644
Introduction to the Challenges and Opportunities of Big Data, the Internet of Things, and Cybersecurity	A semester-long course combining our big data, internet of things (IoT), and cyber courses offered for credit through Harvard Extension School.	4	374
Startup Success: How to Start a Technology Company in Six (not so easy) Steps*	This course discusses the lessons learned by Michael Stonebraker and Andy Palmer during their startup endeavors over a 30-year period. The lessons are distilled into six steps that any entrepreneur can follow to get a company going. Topics include the generation and assessment of ideas, the challenges of building a prototype, the recruitment of a talented team, the closing of the first financing round, and pursuing growth with the right business leadership.	2	359
Internet of Things: Roadmap to a Connected World	The course introduces both the broad range of IoT technologies and the most recent developments in the space; offered through Illumno in FY2019.	9	3,955
Internet of Things: Roadmap to a Connected World— Illumno	The original course, translated into Spanish and Portuguese, and offered through Illumno in collaboration with universities in South America	3	608
Total		70	34,247

*courses not offered FY2019

SystemsThatLearn@CSAIL

The next decade will usher in a new frontier of sophisticated systems that perform complex, humanlike tasks, with complex inferences and predictions. Using data gathered from diverse sensors and mobile devices, computing power spread across embedded devices and data centers, as well as ubiquitous network connectivity, we will need new tools to realize the potential of learning systems. We are already seeing practical applications of these systems in areas such as autonomous vehicles and personalized health care that have the potential to transform industries and societies.

The goal of SystemsThatLearn@CSAIL is to accelerate the development of systems and applications that learn. We intend to accomplish this goal through combining our expertise in systems and machine learning to create new applications for understanding complex relationships unearthed by analyzing the avalanche of available data.

Presently, however, software systems that incorporate machine learning are difficult to build, deploy, and maintain. They require a large and highly skilled workforce. Unlike traditional enterprise systems, once built, they often require thousands of hours of ongoing (sometimes daily) maintenance to ensure that their predictions and behavior continue to be accurate and useful. Integrating ML systems into traditional enterprise architecture, testing and deployment processes are too complex, partly due to organizational silos that exist between systems engineers and data scientists. In application, many problems in large-scale software systems involve optimizations that benefit from predictions, such as scheduling, compilation, query planning, routing, data cleaning and congestion control. Today, it is hard to apply machine learning tools to design this type of system software.

Our approach to designing, training, and deploying humanlike tools will focus on the following four areas of investigation:

Heterogeneous architectures

The data and features that drive learning in these systems and applications increasingly come from diverse, distributed infrastructure, including phones, sensors, or other bandwidth and power impoverished endpoints. Thus, even acquiring data for learning may require adaptive allocation of computation over heterogeneous infrastructure. Furthermore, the rise in heterogeneous hardware, such as GPUs and multicore processors, which excel at certain aspects of the learning pipeline, suggests a diversity of computational resources will be brought to bear.

Predictable composition

Successfully designing and training machine learning methods for the desired task once data is available (i.e., programming at the level of learning components and reasoning about the behavior of the composition of such components), calls for skill and expertise that is not yet well-supported or automated.

Distributed execution

In terms of the underlying infrastructure, complex machine learning methods also demand considerable parallel resources to train effectively. Once trained, models may be deployed either on massive parallel infrastructures (e.g., data centers) or may have to be reduced and distributed back to the heterogeneous components to be utilized where needed (e.g., mobile devices), requiring new distributed algorithms and execution frameworks.

Seamless integration of training and deployment

Many machine learning solutions today are trained and deployed in well-separated phases of training and testing (deployment), but this will change. Learning will increasingly become an ongoing, integrated process. The tighter integration of learning and computer systems offers exciting possibilities in terms of new capabilities, but requires us to overcome challenging hurdles pertaining to programming abstractions, maintenance and monitoring, analysis, and performance guarantees. This includes, among other things, safeguards and ways of containing learning functions in the event that something does not operate as expected as well as approaches to learning on untrusted infrastructures.

In addition to building better systems for machine learning, we believe our focus on the deployability of models will help us advance machine learning itself by developing new models designed to further the above democratization goals, while still providing excellent prediction accuracy. We expect many new tools and practices to be developed.

SystemsThatLearn@CSAIL is a large, multi-PI research program to accelerate the development of this next generation of systems. The primary focus is on developing a common infrastructure, specifically in the form of software that includes the following new theoretical advances:

- Tools to help data scientists and engineers understand their models, train them, monitor their results, and retrain models efficiently
- Useful models focusing on efficiently deploying models in distributed and datacenter settings, reusing and redeploying models, as well as creating development environments good for training and deployment
- Developing heterogeneously deployable models (i.e., models that can be decomposed across heterogeneous devices), or lower fidelity models that can run on sensors or smartphones and also on more powerful servers as well as developing models that are more interpretable
- Tools for statistical monitoring and performance prediction where machine learning is used to understand the performance of complex systems
- Tools and methods to implement and run systems that learn over an untrusted infrastructure

SystemsThatLearn@CSAIL is led by Professors Sam Madden and Tommi Jaakkola and includes 37 CSAIL researchers. It is structured as an industry consortium. On March 29, 2017 we launched this initiative with five founding members: BT, Microsoft, NOKIA

Bell Labs, Salesforce, and Schlumberger. Additional members added in FY2018 include: BASF, ElementAI, EY, and JP Morgan Chase. In FY2019 Facebook was added as a member. In FY2019, there were grants for 10 projects.

FinTech@CSAIL

FinTech (financial technology) is disrupting many aspects of financial services, banking, and insurance, among other industries. Not only will infrastructure and operations be disrupted, but new technologies, business models, services, and even industries will be launched. FinTech holds promise not only for verified transaction systems such as blockchain, but also for technologies involving AI, security, data analytics and value extraction, machine learning, trust verification, risk management, and privacy advances as well. The financial sector has many unique attributes and at the core of a company's success in this sector is trust, security, value, and efficiency. Current technology roadmaps aren't perceived as providing sufficient guidance for FinTech, and new players and technologies are constantly emerging. The shifting demands of customers are evident and pose both risk and opportunity. To stay competitive and ahead of the curve, companies need access to innovation, thought leadership, new technologies, and high-caliber talent.

FinTech@CSAIL will bring together industry, thought leaders, innovators, academics, disruptive technology development, and startup companies that are reinventing global financial services. Through FINTECH@CSAIL, we will work closely with industry partners in leveraging innovation from cutting-edge research to develop the next generation of impactful technologies that will open up new business models, broaden access, gain new data insights, and improve security.

The breadth of research at CSAIL uniquely positions the lab to address a wide variety of challenges in the space. FinTech@CSAIL will include 15 CSAIL researchers, who have pioneered the fields of secure computation, machine learning, artificial intelligence, data analytics and risk management. The goal is to advance the state of the art in collaboration with select industry partners to address the hardest problems facing the finance industry today.

Through the rigorous research of our faculty coupled with our tradition of collaborating with industry, FinTech@CSAIL will address relevant business problems with long-term vision. Additionally, FinTech@CSAIL will draw from across the lab as well as other focused research initiatives in Cybersecurity@CSAIL and SystemsThatLearn@CSAIL—all of which come together for a very powerful collaboration to address the challenging issues that are emerging. We anticipate many opportunities for direct, active collaboration and knowledge sharing through events, projects, and directed research.

By leveraging the research ecosystem at CSAIL, we will work to address the following:

- New approaches to efficient shared public ledger systems and digital currency
- Technologies to provide secure multiparty computation as well as secure and private data extraction

- Advanced data analytics to apply to risk management, and prediction
- Scalable, trusted systems
- Security of ML and AI systems
- ML and AI in automatic advising, compliance, and augmented assistance
- Applying natural language processing in the context of financial contracts
- Speech recognition applications
- Security of legacy systems
- Efficient processing and automation of tasks
- Anonymization of data and privacy
- Movement of datasets and sharing of datasets securely
- New lending and payment technologies

Addressing the changing role of banks as new technologies form the landscape of the future of the financial industry

FinTech@CSAIL was launched on July 18, 2018. The initiative is led by Professors Andrew Lo, Silvio Micali, and Shafi Goldwasser. In FY2019 Gary Gensler joined as an associate director. Founding members include: NASDAQ, Ant Financial, Citigroup, Ryan Software, the London Stock Exchange Group, and Ripple Labs; State Street Bank, the Canadian Imperial Bank of Commerce, the Bank of International Settlements, Fidelity, ConsenSys, and Itaú have also joined. In FY2019 there were grants for 10 projects.

Cybersecurity@CSAIL

Cyber systems cover communications, banking, data processing, purchasing, energy infrastructure, transportation, and defense—nearly every aspect of our lives. Consequently, cyber-based attacks have become more frequent, and more devastating. The present weaknesses in both hardware and software continue to threaten not only the confidentiality of private data and the integrity of data at large, but also the availability of the critical operating systems organizations use to support internal operations, manage assets, secure logistics, sales, and even personnel. Today these cybersecurity challenges rise across virtually all industry sectors and organizations are dealing with an ever-increasing number of attacks.

Through Cybersecurity@CSAIL, we are not just designing technology for specific tasks, but working toward solutions for the whole security spectrum. We approach security from all sides: programming languages, software verification, computer architecture, crypto, systems, and policy. Our goal is to create security by default and remove program error as a source of vulnerability. We are designing new theoretical and practical foundations of secure computing that integrate security in the design process.

Our objective is to design protocols to make attacks more difficult, retain function despite attacks, and allow a system to recover quickly from an attack. CyberSecurity@CSAIL intends to maintain an interdisciplinary focus that brings together thought leaders from industry and government, as well as MIT faculty, researchers, and students conducting research across the security spectrum in hardware, software, encryption, and theory specifically addressing the challenges of ensuring operating system security, secure code, hardware designs for optimal security, defense tools, securing the cloud, multiparty protocols, and usability of encrypted data.

Cybersecurity@CSAIL is an industry consortia model launched in March 2015. Our industry partners provide valuable perspective on the challenges faced across several industry verticals. Our partners include Boeing, Raytheon, BBVA, Akamai, and State Farm. Initiatives are only meant to last three to five years. Cybersecurity@CSAIL concluded its fourth and final year on June 30, 2019. In the final round, grants were made for 11 projects.

Internet Policy Research Initiative

Communication and information networks have become fundamental to our increasingly digital economy and society. Despite this importance, the technologists and policymakers who play key roles in supporting the transition to these networks approach issues from different perspectives and often do not speak the same language. This disconnect can lead to uninformed policymaking and misdirected research efforts. As such, there is a pressing need to bridge the gap between technical and policy communities.

The mission of the Internet Policy Research Initiative (IPRI), an Institute-wide initiative, is to work with policymakers and technologists to remedy this issue and increase the trustworthiness and effectiveness of interconnected digital systems, like the internet. We accomplish this via a three-pronged approach: targeted engineering and public policy research, educational programs for students and policymakers, and outreach programs to build policy communities that facilitate communication.

IPRI's Research Efforts Cover Six Categories

Cybersecurity

IPRI's cybersecurity research focuses on the technical and policy aspects of cybersecurity issues as they relate to the communication networks and software systems affecting the global society and economy. This multidisciplinary research area encompasses encryption policy, accountability, cryptography, data sharing, securing core economic and social infrastructure, measuring cyber risk, and more. Major projects include the following:

- Keys Under Doormats: which has been cited favorably in leading policy documents, such as the House Encryption Working Group report.
- An interdisciplinary CSAIL/Sloan Cyber Risk project, which uses a new cryptographic platform to pool sensitive cyber risk and loss data from firms with the goal of providing more accurate market data on cyber risks.

AI policy

Artificial intelligence and machine learning technologies are becoming increasingly prevalent in not only advertising and research, but also in traditionally regulated spaces, such as health care, finance, transportation, and employment. The AI Policy team at IPRI is focused on increasing the trustworthiness of AI and ML systems by enhancing accountability and expanding public awareness. Current research areas include studying the role of AI in financial decision making, increasing access to new training datasets with policy, working with stakeholders on AI principles, and shaping global internet policymaking via policymaker engagement and informing the public debate. As part of this work, IPRI co-hosted the first AI Policy Congress with the MIT Quest for Intelligence on January 15, 2019. At this event, OECD members, leading ML experts, global policymakers, and industry experts discussed how we should govern AI systems and how to enable AI systems to meet society's needs, both domestically and internationally.

Privacy

IPRI's work also focuses on privacy policy and its critical role in trustworthiness. The Privacy group has published work on such topics as privacy and security in home assistants, exposure minimization, and the data-sharing practices of smartphone apps. Current projects include the development of privacy-aware databases. IPRI and CSAIL are also launching a data governance and privacy consortium in 2019 covering challenging issues, such as database systems, applied cryptography, AI and ML, data portability and new information architectures, and human-computer interaction.

Networks

The Advanced Network Architecture (ANA) group works to understand and shape the future of the internet. They achieve this goal with the understanding that the future of the internet is defined by the economic, social, regulatory, legal, and political concerns involving the internet. As such, the ANA group is organized around five themes: internet architecture, internet security, internet economics, internet policy, and network management. In 2018, the ANA group had its first major release of data about interconnection congestion on the internet.

Decentralized web

The Decentralized Information Group (DIG) focuses on data and systems governance (primarily on the web) and explores both policy and technical issues. Current projects include a decentralized privacy-preserving platform for clinical research, evaluating the trustworthiness of autonomous systems, studying the relationship between privacy and machine learning, developing explanations for complex machines and models, securely aggregating distributed data, and developing smart contracts for data sharing.

App Inventor

The core goal of the MIT App Inventor team is to empower young people to develop useful apps that serve as novel digital solutions to the problems they face in their lives, communities, and world.

World Wide Web Consortium

The World Wide Web Consortium (W3C) was founded at MIT in 1994 by the inventor of the web, Tim Berners-Lee. W3C is responsible for developing and maintaining the

standards that make the web work and for ensuring the long-term growth of the web. Nearly 450 member organizations, including most of the world’s leading technology companies, are working to enhance the capabilities of web documents and create an open web platform for application development—available across a wide range of devices—enabling billions of people to collaborate and share data and information. In recent years, a great many factors (people, devices, bandwidth, policy decisions, etc.) have extended the reach of the web in society. Video, social networking tools, user-generated content, location-based services, and web access from mobile devices are transforming many industries—including mobile, television, publishing, automotive, entertainment, gaming, and advertising. This transformation has led to greater demands on W3C and other organizations to build robust technology that meets society’s needs in areas such as privacy, security, accessibility, and multilingual content.

Core Technology Focus

W3C standards define an open web platform for application development that has the unprecedented potential to enable developers to build rich interactive experiences, powered by vast data stores that are available on any device. Although the boundaries of the platform continue to evolve, industry leaders speak in unison about how HTML5 is the cornerstone for this platform. The full strength of the platform relies on many more technologies that the W3C and its partners are creating, including cascading style sheets, scalable vector graphics, and web open font format. The W3C’s real-time communications spec is powering the increased tendency for people to videoconference using native support in their browser.

In recent years, publicly noted security and privacy breaches have resulted in unprecedented attention to fixing web security and privacy. W3C addresses that both with specific solutions (such as the recently completed web authentication spec that authenticates users without passwords), and by conducting reviews of every W3C standard for security and privacy. The growth of e-commerce has focused new attention on standardizing payment and e-commerce approaches. And with the internet of things arriving, our web of things project aims to address semantic interoperability to prevent IoT from driving silos at the application level. With immersive technologies, there is a strong focus on web solutions for virtual reality and augmented reality.

The growing impact of the web is also driving W3C to expand its agenda and the size of its community. In 2011, W3C launched its Community and Business Groups, and after eight years over 11,000 people participate. By making it easier for people to participate, W3C has increased the relevance and quality of its work and brought more innovators to the table for prestandards and standards track work.

Industry Impact and Broadening the Set of Participants

In recent years, web technology is not only used by consumers and companies for information sharing, but increasingly the web is the delivery mechanism for companies to deliver their services. Examples of that include telecommunications (where web access is a key service), entertainment (which is increasingly delivered over the web), publishing (whose standards organization—the International Digital Publishing Forum—recently merged into W3C), and retail and financial services (both impacted with an increase of

payments on the web). This has caused a diversification in the membership of W3C, and also has enriched the technical agenda to address new technical issues that arise. For example, web browser companies, credit card companies, and other fintech stakeholders are using W3C to streamline web payments through a browser.

Recognition

In early 2019, W3C was awarded with its second Technical Emmy in four years—this one for video streaming on the web.

Organization

In early 2019, W3C's Steering Committee approved a resolution that began a process that explores W3C becoming an independent legal entity.

Research Highlights

In addition to the large-scale collaborative projects and center research, numerous individual and multi-investigator projects are under way. A sampling of the work is highlighted below:

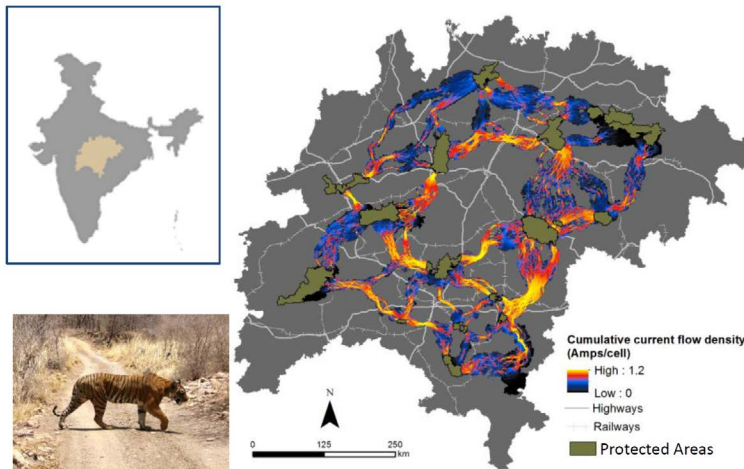
Julia

Co-created by Alan Edelman, the Julia programming language reached a major milestone.


- The Julia programming language reached 1.0 at JuliaCon 2018 multithreading support for parallel computing
- Julia 1.0 also allows users to target heterogeneous architecture support including Nvidia GPUs
- In July 2019, Julia was featured in *Nature*

Circuitscape—Conservation Science at Scale

Circuitscape is an open source connectivity modelling program (now written in Julia) that is used by conservation scientists worldwide to inform policy.



Professor Trishna Dutta and her colleagues combined Circuitscape with least-cost corridor methods to map pinch points within corridors connecting protected areas for tigers in central India. Areas with high current flow are most important for tiger movements and keeping the network connected.




Projected climate-driven range shifts of 2,903 species in response to climate change using Circuitscape. Arrows represent the direction of modelled movements from unsuitable climates to suitable climates via routes that avoid human land uses. From: Lawler et al., 2013. [Explore the full animation of these results.](#)

Climate Modelling Alliance

Leveraging recent advances in the computational and data sciences, a consortium of engineers from MIT, Caltech, the Naval Postgraduate School and the Jet Propulsion Lab aim to build a new climate model from the ground up. The consortium, dubbed the Climate Modeling Alliance, plans to fuse Earth observations and high-resolution simulations into a model that represents important, small-scale features, such as clouds and turbulence, more reliably than existing climate models. The goal is a climate model that projects future changes in critical variables such as cloud cover, rainfall, and sea ice extent more accurately – with uncertainties at least half the size of those in existing models.

PuMaS—Pharmacometric and Pharmacodynamic Simulations



PuMaS is a new tool for pharmaceutical modeling and simulation written in the Julia programming language, which allows for integrating the workflows across the spectrum of preclinical and clinical analysis.

Neural Ordinary Differential Equations

Neural Ordinary Differentiation Equations are tools used to simulate the dynamics of ordinary differential equations (ODEs) using artificial neural networks. ODEs have many applications, for instance, they can be used as surrogate models for large, complex engineering or natural systems; they also can be used to reduce the complexity of systems via nonlinear model order reduction. DiffEqFlux.jl is a Julia package used to train neural ODEs, and supports delayed differential equations, differential algebraic equations, and stochastic jump differential equations.

Crossing the Vision-Language Boundary

As humans and cognitive machines interact, it will be important for both parties to be able to speak with each other about things that they perceive in their local environment. Our ongoing research investigates deep learning methods for learning semantic concepts across both audio and visual modalities. Contextually correlated streams of sensor data from multiple modalities (for example, an image accompanied by a spoken caption describing that image) are used to train deep neural network models capable of discovering patterns using otherwise unlabeled training data.

The current structure of the model, called DAVeNet (Deep Audio Visual Embedding Network), consists of one deep neural network that scans over image pixels, and a second network that scans over speech waveform samples. Both networks learn to represent the information contained in their respective inputs by a high-dimensional embedding vector. The vectors are coerced into learning a joint audio-visual semantic embedding space by showing the model both positive and negative matching pairs of images and speech descriptions.

One of the outcomes of the model is that it is able to jointly learn properties of objects as represented in both the visual and audio modalities. For example, these networks are able to pick out instances of the spoken word “water” from within continuous speech signals and associate them with images containing bodies of water. The figure shows three colored regions in an image and its spoken caption (the transcript is for illustration only) that illustrate how the model has learned the audio-visual correspondences for an airplane (yellow), trees (red), and mountain (green). The networks learn these associations directly from the raw data, without the use of conventional speech recognition, text transcriptions, or any expert linguistic knowledge whatsoever. The models also learn these correspondences from random initialization.

The results of this research should have wide-scale applications to future cognitive machines that perceive and interact with people. They also have relevance to language acquisition, since conventional speech recognition technology requires thousands of hours of annotated data to train, which is a limiting factor to cover the nearly 7,000 languages of the world, and which stand in stark contrast to our human ability to acquire language. This research involves MIT staff and students: David Harwath, Adrià Recasens, Angie Boggust, Emmanuel Azuh, Antonio Torralba, and James Glass.

D. Fox Harrell, professor of digital media and artificial intelligence, was recently appointed as the director of the [MIT Center for Advanced Virtuality](#), a new center for research into virtual reality, augmented reality, and related technologies.

The Center for Advanced Virtuality (MIT Virtuality for short) pioneers innovative experiences using technologies of virtuality—computing systems that construct imaginative experiences atop our physical world. MIT Virtuality’s approach to engineering and creative practices pushes the expressive potential of technologies of virtuality and simulates social and cognitive phenomena, while intrinsically considering their social and cultural impacts.

Several endeavors are underway at MIT Virtuality. One such endeavor is a project called A Framework for Self-Reflection in Massive Online Workplace Learning Systems, supported by J-WEL in MIT Open Learning. This project has focused on learning about sexism in the workplace in a framework that also supports learner reflections during and after educational experiences, resulting in measurable learning gains. This has been implemented in a massive open online course and interactive narrative prototype called [Grayscale](#). Another MIT Virtuality research project has focused on modeling racial and ethnic socialization for interactive storytelling. This research examined the effectiveness of the use of an embedded computational model of racial and ethnic socialization (RES) for virtual identity representation in an interactive virtual reality (VR) narrative games. A particular outcome is a novel VR system called [Passage Home VR](#) and a user study conducted with the system. The findings of the study suggest a significant relationship between players’ RES, their in-game behaviors, and their narrative interpretations. The results have practical implications for the design of videogames aiming to support players from diverse racial and ethnic backgrounds. MIT Virtuality also co-organized an international workshop on the impacts of AI on indigenous cultures (for aims such as language preservation), called “AI and Society Workshop: Indigenous Protocol and AI,” supported by the Canadian Institute for Advanced Research.

Harrell also leads his research group, the Imagination, Computation, and Expression Laboratory (known as the ICE Lab and established at MIT in 2010), which applies AI and cognitive science approaches to the research and development of interactive narratives, videogames, virtual reality, social media, and related forms of digital media. Outcomes of recent ICE Lab projects include a three-year, MIT CSAIL-Qatar Computing Research Institute collaboration researching culturally specific, everyday uses of virtual identities in social media and videogames (with the Middle East and North Africa [MENA] region as a case study). This research analyzed a dataset of over 42,000 publicly available social media profiles using computational approaches (archetypal analysis), augmented with social science-based semi structured interviews. This resulted in a set of five needs/values exhibited by Qatari users supporting their creativity in effectively using virtual identities: Arab cultural features, self-expression, social connections, social monitoring, and physical and virtual identity contrasts; it also resulted in guidelines to support developers of virtual identity systems in better serving these users while preserving their cultural values and creative agency.



A version of the MazeStar Computer Science Learning Platform localized for MENA region cultures based on computationally elicited values



Deployment of MazeStar at QCRI in Doha, Qatar

This was deployed to create a MENA region localized version of another ICE Lab outcome called *MazeStar*, a National Science Foundation (NSF)-supported computer science learning curriculum and educational computer game creation platform that engages students in learning computer science concepts while seeing themselves as powerful STEM learners and doers.

Active Model Learning

Our overarching goal is to construct robots that can act capably and robustly in complex, uncertain, and highly variable domains over long periods of time. Especially in the context of robots that work with humans, it is critical for humans to be able to quickly and effectively train robots to perform novel types of tasks, and for the robots to be able to immediately generalize and broadly apply the abilities they have learned.

A great deal of existing work in robotics has focused on how robots can learn new sensorimotor skills, such as pouring or stirring. However, there are no effective approaches for combining these motor skills to achieve long-horizon, high-level task objectives, such as cooking a meal or making a cup of coffee.

On the other hand, our group has developed several methods for constructing plans to achieve such high-level task objectives given well-characterized skills, such as picking up and placing objects. In recent work, we have developed an approach to integrate newly learned skills into our planning framework. The key bridging step is to automatically learn abstract models of these new skills, based on active experimentation by the robot. These models characterize conditions under which the skill is likely to succeed, and the effects of the skill on the state of the world.

After learning a new skill (such as pouring), it can be combined with existing skills (such as picking and placing), so that the robot can pick up a cup, pour the contents out and place the cup back down. It is only through this ability to combine primitive skills that robots can truly become useful in highly complex and variable human environments.

Laboratory Sponsored Activities

CSAIL Outreach

CSAIL regularly encourages the online community to submit questions about computer science and academia to its researchers in a series of Reddit Ask Me Anything (AMA) sessions. CSAIL's AMAs have spurred approximately 8,000 questions and comments, as well as more than 300,000 page views.

Media Outreach


CSAIL has a combined following of 170,000 users across Twitter, Instagram, Facebook, and YouTube.

Growth in media coverage and viewership:

- There has been a 900% increase in media hits since 2013 (567 versus 53)
- The amount of CSAIL video that YouTube users have watched is now more than 2 million minutes, which adds up to more than four years

- CSAIL's versus MIT News' YouTube viewership comparison: CSAIL has posted a total of 10 videos with an average of 37,000 views; while MIT News has posted a total of 70 videos with an average of 13,000 views

CSAIL Twitter followers by year



Distinguished Lecture Series

The Dertouzos Lecture Series has been a tradition since 1976, featuring some of the most influential thinkers in computer science. Three speakers presented lectures during the AY2019 Dertouzos Distinguished Lecture Series:

- Vladimir Vapnik, Columbia University, presented “Learning Using Statistical Invariants (Revision of Machine Learning Problem),” on September 26, 2018
- John Hennessey, Stanford University, presented “The End of Road for General Purpose Processors and the Future of Computing,” on March 6, 2019
- Kai-Fu Lee, Sinovation Ventures and author of *AI Superpowers*, presented “The Era of Artificial Intelligence,” on May 3, 2019

CSAIL also continued a series of Hot Topics in Computing, a community speaker series convening experts in computing to discuss the potential, perception, and problems associated with the proliferation of computation and machines:

- Kate Starbird, University of Washington, presented “Muddied Waters: Online Disinformation during Crisis Events,” on September 12, 2018
- John Leonard, MIT, presented “Challenges and Opportunities for Self-Driving Cars,” on October 24, 2018
- Katie Bouman, Caltech, presented “Imaging a Black Hole with the Event Horizon Telescope,” on April 22, 2019
- Elizabeth Cobbs and James Shelley (producers), presented *CyberWork and the American Dream* (film screening and discussion), on April 24, 2019

CSAIL also hosted the inaugural TEDxMIT event on May 28, 2019, featuring a series of talks from 14 women in science and technology across the Institute.

Organizational Changes

Professor Daniela Rus, has continued in her role as director of CSAIL. The director's duties include developing and implementing strategies designed to keep CSAIL growing and evolving, fundraising, determining laboratory policies, and examining promotion cases.

CSAIL's leadership team includes an associate director and a chief operating officer (COO), and the executive cabinet. These leaders are appointed by the laboratory's director and assist her with her duties. Professors Daniel Jackson and Charles Leiserson are the current associate directors. Professor Leiserson also serves as chief operating officer, providing leadership and strategy for how we conduct our operations and initiatives, enabling the director to allocate more time to strategic planning. Victor Zue holds the role of director of greater China relations, managing the engagements and oversight of various important CSAIL international contracts and international contract negotiations.

The CSAIL executive cabinet met weekly to review and advise the director on policy, processes, activities within the laboratory, and preparation for transitions related to the Schwarzman College of Computing. Members of the FY2019 executive cabinet include: Edward Adelson, Saman Amarasinghe, Randall Davis, James Glass, Daniel Jackson, Charles Leiserson, Sam Madden, Wojciech Matusik, Ronitt Rubinfeld, Daniela Rus, Bruce Tidor, and Victor Zue.

The CSAIL enterprise services team manages lab operations. There are seven units (Administrative Assistants; CSAIL Alliance Program; Communications; Finance; HR; Special Projects; and TIG) reporting to the CSAIL COO on all operational matters. Carmen Finn is the assistant director for administration; Agnes Chow also served as acting administrative officer during AY2019 while Finn was on medical leave. Lori Glover is managing director of the CSAIL Alliance Program. John Costanza is the assistant director for infrastructure, overseeing information technology infrastructure and user support, building operations, and communications, and serves on the space committee.

Bruce Tidor oversees the space committee and manages the allocation of space within CSAIL. The space committee also implements improvements to the facilities that will increase the quality of the environment for the laboratory's faculty, staff, and students. The space committee also includes Assistant Director John Costanza.

New faculty who started during AY2019 include the following:

- Manya Ghobadi, assistant professor, October 1, 2018
- Phillip Isola, assistant professor, July 1, 2018
- Arvind Satyanarayan, assistant professor, July 1, 2018)

Faculty who took leave during AY2019 included:

- Regina Barzilay, sabbatical, fall 2018
- Tim Berners-Lee, sabbatical, fall 2018 to fall 2019
- Tamara Broderick, sabbatical, fall 2018
- Michael Carbin, sabbatical, fall 2018
- Adam Chlipala, sabbatical, fall 2018; fall 2019
- David Gifford, professional leave, fall 2018
- Shafi Goldwasser, professional leave, fall 2018 to fall 2019
- Tommi Jaakkola, sabbatical, fall 2018; fall 2019
- John Leonard, TRI 80%
- Sam Madden, sabbatical, fall 2018 to spring 2019
- Silvio Micali, sabbatical, fall 2018; professional leave, fall 2019
- Robert Miller, sabbatical, fall 2018 to spring 2019
- Nir Shavit, sabbatical, fall 2018 to spring 2019
- Armando Solar-Lezama, sabbatical, fall 2018
- Gerald Sussman, sabbatical, fall 2018
- Russ Tedrake, TRI 80%
- Virginia Williams, sabbatical, fall 2018
- Ryan Williams, professional leave, fall 2018

Faculty changes and promotions during AY2019 included:

- Mohammad Alizadeh, associate professor, 2019
- Tamara Broderick, associate professor, 2019
- Stefanie Jegelka, associate professor, 2019
- Aleksander Madry, achieved tenure, 2019
- Wojciech Matusik, full professor, 2019
- Daniel Sanchez, tenure, 2019
- Virginia Williams, tenure, 2019

Awards and Honors

Our faculty and staff have achieved many awards, including the following:

- Edward Adelson: IEEE Fellow, 2018
- Mohammad Alizadeh: Microsoft Research Faculty Award, 2018

- Anant Agarwal: Yidan Prize Education Development Laureate, 2018
- Bonnie Berger: International Society for Computational Biology, Senior Scientist Award, 2019
- Regina Barzilay: Susan G. Komen Scholar, 2018; and Top 100 AI Leaders in Drug Discovery and Advanced Health Care, 2019
- Tim Berners-Lee: *Financial Times* Boldness in Business Person of the Year, 2019
- Judy Brewer: SIGACCESS Award for Outstanding Contributions to Computing and Accessibility, 2018
- Tamara Broderick: International Conference on AI and Statistics Notable Paper Award, 2019
- Michael Carbin: Google Faculty Award, 2019; and International Conference on Representation Learning Best Paper Award, 2019
- Adam Chlipala: Most Influential ICFP Paper Award, 2018; and Ruth and Joel Spira Award for Excellence in Teaching, 2019
- David Clark: Association of American Publishers, Prose Award; History, Science, and Medicine, 2019
- Constantinos Daskalakis: Rolf Nevanlinna Prize, 2018; Simons Foundation's Investigator Award for Mathematics, 2018; and ACM Grace Murray Hopper Award, 2019
- Erik Demaine: Margaret MacVicar Faculty Fellow award, 2019
- Srinivas Devadas: Distinguished Alumnus Award, Indian Institute of Technology, 2019
- Carmen Finn: MIT Infinite Mile Award, 2019
- Shafi Goldwasser: Honorary Degree, University of Waterloo, 2019; and Honorary Degree, Oxford University, 2019
- Tommi Jaakkola: Top 100 AI Leaders in Drug Discovery and Advanced Health Care, 2019
- David Karger: American Academy of Arts and Sciences, 2019
- Dina Katabi: TED Annual Conference featured talk, 2018; and Carnegie Corporation Great Immigrants recognition, 2019
- Manolis Kellis: Top 100 AI Leaders in Drug Discovery and Advanced Health Care, 2019
- Nancy Lynch: Best Paper Award 22nd International Conference, 2018
- Aleksander Madry: Presburger Award for Young Scientists, 2018
- Rob Miller: Visual Languages and Human-Centric Computing Most Influential Paper Award, 2018

- Robert Morris: Member, National Academy of Engineering, 2019
- Stefanie Mueller: NSF Career Award, 2019
- Ronitt Rubinfeld: EECS Seth J. Teller Award, 2019
- Daniela Rus: Constellation’s Business Transformation 150, 2018; President’s Medal, Stevens Institute of Technology, 2018; RoboSoft Best Paper Award, 2019; RoboSoft Best Poster Award, 2019; Now Awards Silver Medal, 2019; and ICRA Best Paper Award Finalist, 2019
- Peter Shor: Micius Quantum Prize, 2019
- Julian Shun: MIT Research Support Committee Award, 2018; ACM Best Paper Award, 2018; DOE Early Career Award, 2018; and NSF Career Award, 2019
- David Sontag: The Burgess (1952) and Elizabeth Jamieson Prize for Excellence in Teaching, 2019
- Vivienne Sze: MIT Harold E. Edgerton Faculty Achievement Award, 2019; Top 100 AI Leaders in Drug Discovery and Advanced Health Care, 2019; and The Burgess (1952) and Elizabeth Jamieson Prize for Excellence in Teaching, 2019
- Russ Tedrake: Conference on Robot Learning Best Paper Award, 2018; Amazon Robotics “First Ever” Best Technical Paper Award, 2018; and “First Ever” Research Paper of the Year, 2018
- Ryan Williams: Google Faculty Research Award, 2019; and SAT Conference Best Paper Award, 2019
- Alan Willsky:, IEEE Jack S. Kilby Signal Processing Medal, 2019

Key Statistics for AY2019

Category	Total	Women percentage
Faculty	103	17%
Postdoctoral associate/fellow	87	8%
Principal research scientist	10	60%
Research staff	46	22%
Senior research scientist	4	25%
Administration, technical, and support staff	87	60%
PhD students	457	22%
Master of Engineering students	97	40%
UROP	236	34%
Visitors	82	20%
Total personnel	1,209	27%

Daniela Rus

Director, Computer Science and Artificial Intelligence Laboratory