Computer Science and Artificial Intelligence Laboratory

The Computer Science and Artificial Intelligence Laboratory (CSAIL) is focused on developing the architectures and infrastructures of tomorrow's information technology, and on creating innovations that will yield long-term improvements in how people live and work. Lab members conduct research in almost all aspects of computer science, including artificial intelligence, the theory of computation, systems, machine learning, and computer graphics, and explore revolutionary new computational methods for advancing healthcare, manufacturing, energy, and human productivity.

CSAIL researchers focus on finding innovative ways to make systems and machines operate faster, better, safer, easier, and more efficiently for the benefit of humanity. Our projects fall into three areas of inquiry:

- Artificial intelligence (AI): We seek to understand and develop both living and artificial systems capable of intelligent reasoning, perception, and behavior.
- Systems: We seek to discover new principles, models, metrics, and tools of both hardware and software-based computer systems.
- Theory: We seek to understand the mathematics of computation and its wide-ranging, real-world consequences.

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 research is sponsored by a large number of diverse sources, from US government contracts to the private sector. United States government sponsors include: the Air Force Research Laboratory and the Air Force Office of Scientific Research; the Army Research Office; the Defense Advanced Research Project Agency; the Department of Defense Research and Engineering; the Food and Drug Administration; the US Department of Education; the Department of Energy; the Intelligence Advanced Research Projects Activity; the National Institutes of Health; the National Institute of Justice; the National Science Foundation; the Navy (including the Office of Naval Research, and Naval Air Systems Command); and the Space and Naval Warfare Systems Center. US and international nonfederal sponsors include: Accenture, LLP; Boeing; BMW of North America, LLC; Ford Motor Company; Foxconn Technology Group; Intel Corporation; Jaguar Land Rover Limited; Lockheed Martin Advanced Technology Laboratories; Microelectronics Advanced Research Corporation; Mitsubishi Electric Corporation, Nissan Motor Company, Ltd.; Nippon Electric Company; Nippon Telegraph and Telephone Corporation; Northrop Grumman Corporation; Ping An Technology; Qatar Computing Research Institute; Quanta Computer, Inc.; Samsung Electronics; Siemens; Toyota Research Institute; and Wistron Corporation. Other organizations sponsoring research include Aarhus University; Battelle Memorial Institute; Delta Electronics Foundation; DSO National Laboratories; Epoch Foundation; The Hong Kong University of Science and Technology; IBM; Industrial Technology Research Institute; Nanyang Technical University; Pfizer Inc.; and the 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:

Toyota-CSAIL Joint Research Center

Toyota established a collaborative research center with CSAIL and Stanford in 2015 to further the development of autonomous vehicle technologies, with the goal of reducing traffic casualties and potentially developing a vehicle incapable of causing a vehicular accident.

Today, a car crash occurs on average every five seconds in the United States. Globally, road traffic injuries are the 10th leading cause of death, with approximately 1.25 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. This new center works in collaboration with the newly formed Toyota Research Institute (TRI) led by Gill Pratt.

Imagine if your car could tell you were having a bad day, and turned on your favorite album to improve your mood. What if your car could talk to your refrigerator, figure out that you're out of milk, and suggest where to stop on your way home? Or 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, and they are 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 objectives to learn from humans; study perception and decision-making systems for safe navigation; design systems that can handle difficult driving situations such as congestion, high-speed driving, and inclement weather; and create predictive models that can anticipate the behavior of humans and vehicles; and more intelligent user interfaces.

More specifically the projects and principal investigators (PIs) that are currently active in the Toyota-CSAIL Joint Research Center 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)
- Using Vision and Language to Read Minds (PI: Nickolas Roy; Co-PI: Boris Katz)
- Uhura: A Driver's Personal Coach for Managing Risk (PI: Brian Williams)
- Predicting a Driver's State-of-Mind (PI: Antonio Torralba; Co-PI: Wojciech Matusik)
- Exploring the World of High Definition Touch (PI: Ted Adelson; Co-PI: John Leonard)
- Formal Verification Meets Big Data Intelligence to Address the Trillion Miles Challenge (PI: Armando Solar-Lezama)
- The Car Can Explain! (PI: Gerald Sussman; Co-PIs: Daniel Weitzner, Hal Abelson, and Lalana Kagal)
- Crossing the Vision-Language Boundary for Contextual Human-Vehicle Interaction (PI: James Glass; Co-PI: Antonio Torralba)
- Analysis by Synthesis Revisited: Visual Scene Understanding by Integrating Probabilistic Programs and Deep Learning (PI: Joshua Tenenbaum)
- Wi-Fi-Based Obstacle Detection for Robot Navigation (PI: Dina Katabi; Co-PI: Daniela Rus)
- Drinking from the Visual Firehose: High-Frame-Rate, High-Resolution Computer Vision for Autonomous and Assisted Driving (PI: Saman Amarasinghe; Co-PIs: John Leonard, and Fredo Durand)
- Decision Making for Parallel Autonomy in Clutter (PI: Daniela Rus; Co-PI: Sertac Karaman)

- A Parallel Autonomous Driving System (PI: John Leonard; Co-PIs: Sertac Karaman and Daniela Rus)
- Uncovering the Pain Points in Driving (PI: Ruth Rosenholtz; Co-PIs: Fredo Durand, William Freeman, Aude Oliva, and Antonio Torralba)
- Simulation and Verification for Vision-in-the-Loop Control (PI: Fredo Durand)
- Tools and Data to Revolutionize Driving (PI: John Leonard; Co-PI: Daniela Rus)

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 health trackers. We monitor stock prices, weather forecasts, and traffic patterns through websites and apps, share our thoughts and experiences through emails, 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). Finding ways to reduce stress and technology's negative impact on a workforce is critical to our future well-being.

The multiyear research program between Wistron and CSAIL focuses on rethinking how we compute and communicate in the digital age to ensure that: (1) health and well-being are at the core of our lives; and (2) 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)
- Personally Authored Wellness Applications (PIs: David Karger and Daniel Jackson)
- Smart Homes that Monitor Breathing, Heart Rate, and Life Quality (PI: Dina Katabi)

Ping An Research Projects

Many organizations today struggle to make intelligent use of all the data they have collected and how to structure their systems most efficiently. In addition, companies dealing with their customer base are continually looking for ways to not only automate but improve customer experience.

To help achieve their goal at excelling in the financial services space, Ping An sought to partner with CSAIL to address certain technical challenges faced by its business. The Ping An research collaborative began in 2014. It was structured as a three-year research engagement, with three specific research projects.

Semantic Summarization for Financial Data: A problem is proposed whereby there is a large set of database tables that contain data regarding customers, but tools do not yet exist for the required number of fields and records. To address this problem, we will build on our experience with summarizing data from mobile sensors, such as smart phones, which are playing an increasingly important role in our lives, and can be the source of very large and useful data about the users carrying them. Our goal is to develop systems and algorithms that take large data streams and convert them into semantic summaries. (PI: Daniela Rus)

Speech Recognition: This project focuses on Chinese speech recognition methods for customer support. A speech recognition capability would be useful to process the millions of existing Ping An human to human, customer to agent, telephone-based call-center communications. There are also potential scenarios for future customer to computer spoken interactions via mobile devices as well. There is tremendous potential for speech technology benefits via spoken interfaces but in the scope of this three-year project, we will focus on current spoken language systems and automatic speech recognition research investigating the use of deep, neural network-based methods for multilingual speech recognition. (PI: James Glass)

Speaker Verification: Speaker verification is the task of processing a speech recording and deciding whether or not it belongs to a putative speaker. For customer telephony applications, speaker verification is useful as a complementary verification method to existing approaches such as caller ID, personal identification numbers, and so on. Speaker verification is closely related to the problem of speaker identification (or recognition), which seeks to determine the identity of a recording from a large candidate pool. There are several potential uses for speaker verification technology for call-center customer support. The majority of these capabilities are directly relevant to the millions of existing human to human, customer to agent telephone-based call-

center communications, though there are potential uses for future human to computer spoken interactions as well. Speaker verification technology complements existing methods for verifying customer identity. This project will focus on speaker divarication to automatically separate customer and agent speech turns during a single recorded dialogue. (PI: James Glass)

The collaboration has been very successful. The initial three-year term is expiring and discussions are underway to possibly extend or expand the engagement.

Qatar Computing Research Institute

In 2012, CSAIL started a seven-year, $35 million research collaboration with Qatar Computing Research Institute (QCRI) to collaborate on a wide-range of research topics in computer science. Five years into the collaboration, we are currently pursuing the following eight projects:

- Arabic Speech and Language Processing for Cross-Language Information Search and Fact Verification (Senior Research Scientist James Glass, and Professors Regina Barzilay and Tommi Jaakkola): This project aims to develop key technologies to enable cross-lingual search for verified facts and claims in multimedia content using questions posed in written and spoken language.
- Content-Adaptive Video Retargeting (Associate Professor Wojciech Matusik): This project aims to develop a complete system for delivering high-quality stereoscopic broadcast video, focusing on the real-time video of sporting events—soccer in particular.
- Database Management (Adjunct Professor Michael Stonebraker, Professor Samuel Madden, and Associate Professor Armando Solar-Lezama): This project investigates three data management tasks, (1) build an end-to-end system (Data Civilizer) to support the data discovery and preparation needs of data scientists; (2) study resource elasticity in online transaction processing database management systems; and (3) 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, (1) using deep image auto-tagging models for the analysis of food perception gap; (2) learning a joint embedding space for cooking recipes and food images; and (3) vision-based estimation of population-level health from social media images.
- Understanding and Developing for Cultural Identities Across Platforms (Professor Fox Harrell): This project aims to (1) develop and deploy computational tools and new techniques to understand the Qatari Gulf Cooperation Council's (GCC) use of virtual identity systems; (2) elicit and articulate best practices empowering Qataris to enact regional values and norms; and (3) develop a GCC-specific novel application.

- A Vertically-Integrated Approach to Resource-Efficient Shared Computing (Associate Professor Daniel Sanchez): This project aims to investigate an integrated node- and cluster-level architecture that provides both near-peak utilization and guaranteed performance in shared clusters, focusing on QCRI/QF workloads and infrastructure needs.
- Urban Data Analytics to Improve Mobility for Growing Cities in the Context of Mega Events (Visiting Associate Professor Marta Gonzalez): This project aims to evaluate the impact of large-scale events and levels of accessibility based on travel times from different origins within a city, to sequentially propose travel demand management strategies to mitigate the traffic congestion during the 2022 FIFA World Cup in Qatar.
- Accurate Map Making using Mobile Sensor Data (Professors Hari Balakrishnan and Samuel Madden, Assistant Professor Mohammad Alizadeh, and Adjunct Professor David DeWitt): This project aims to develop accurate, cartographic 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.

Centers and Initiatives

CSAIL Alliance Program

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

CAP provides access to all 50 research groups, spanning robotics, natural language processing, networks, databases, cryptography, and web science, among others. CAP membership currently has two levels: affiliate and partner. Both levels include lab visits, access to the annual meeting, recruiting assistance, research briefings and professional education discounts. Partner members, however, have expanded benefits over affiliate members, including greater access to research initiative meetings; custom, faculty-led seminars; and expanded recruiting options.

Currently there are over 60 member companies, such as Apple, Google, Samsung, and Microsoft. Members are headquartered in North America, South America, Europe, and Asia, and represent a wide variety of industry verticals.

CAP also produces and manages online professional development courses in partnership with MIT's Professional Education, Office of Digital Learning, and edX. The following is a list of the programs to date. Total enrollment is now approaching 20,000 online learners:

Course title	Description	Offerings to date	Total enrolled to date
Tackling the Challenges of Big Data	Survey of 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
Cybersecurity: Technology, Application, and Policy	This six-week online course provides a holistic look at cybersecurity technologies, techniques, and systems.	5	2,966
Start-Up 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 start-up 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	This course introduces both the broad range and most recent developments of Internet of Things (IoT) technologies.	4	3,598
Total enrollments in CSAIL-produced online courses		22	19,650

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 datacenters, 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, which 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 by combining our expertise in systems and machine learning to create new applications for understanding complex relationships unearthed by analyzing the avalanche of data available today.

Presently, however, software systems that incorporate machine learning are difficult to build, deploy, and maintain, and require a large and highly skilled workforce. Unlike traditional enterprise systems, once built, they often require thousands of hours of on-going—sometimes daily—maintenance to ensure that their predictions and behavior continue to be accurate and useful. Integrating machine-learning 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 difficult to apply machine-learning tools to design this type of system software.

Our approach to designing, training, and deploying, 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 many-core 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, that is, 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 offer exciting possibilities in terms of new capabilities, but requires us to overcome challenging hurdles pertaining to programming abstractions, maintenance, 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-investigator research program designed 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 new theoretical advances and tools to help data scientists and engineers understand their models, train them, monitor their results, and retrain models efficiently. The program designs useful models, focusing on efficiently deploying models in distributed and datacenter settings, reusing and redeploying models, as well as creating development environments suited for training and deployment. There is a focus on developing heterogeneously deployable models, that is, 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. Researchers create tools for statistical monitoring and performance prediction where machine learning is used to understand the performance of complex systems, as well as tools and methods to implement and run systems that learn over an untrusted infrastructure.

SystemsThatLearn@CSAIL—structured as an industry consortium—is led by Professors Samuel Madden and Tommi Jaakkola and includes 37 CSAIL researchers. On March 29, 2017 we launched this initiative with six founding members: BT, Microsoft, NOKIA Bell Labs, Salesforce, Schlumberger, and ScotiaBank.

bigdata@CSAIL

The MIT Big Data Initiative at CSAIL, bigdata@CSAIL, was originally launched in 2012 in parallel with Intel's selection of MIT and CSAIL to host its Intel Science and Technology Center for Big Data, and as a research consortia with industry members including EMC, Facebook, Microsoft, and Shell. The initial term was three years but, with industry support, was extended for an additional two years. It will be concluding June 2017 with the research direction moving to SystemsThatLearn@CSAIL. The faculty director is Professor Samuel Madden and there are 38 affiliated researchers from MIT who are world leaders in parallel architecture, massive-scale data processing, databases, algorithms, machine learning, visualization, and user interfaces addressing the following four broad research themes:

- Computational Platforms: The goal of these platforms was to make it easy for developers of big data applications to write programs much as they would on a single-node computational environment, and be able to rapidly deploy those applications on tens or hundreds of nodes. Additionally, as the computation and storage requirements of applications change, these platforms should be able to dynamically and elastically adapt to those changes.
- Scalable Algorithms: We developed a range of algorithms designed to deal with very large volumes of data, and to process that data in parallel. A particular focus is on algorithms for summarizing, comparing, searching, and querying massive data sets.

- Machine Learning and Understanding: We design novel, machine-learning applications focused on machine understanding in specific domains. For example, in work on scene understanding in images we built tools that automatically label parts of an image, or classify an image as belonging to a certain category or categories based on the types of objects that appear in them. As a second example, we used natural language processing to convert massive quantities of text tweets and text reviews on the web into structured information about products, restaurants, and services, which indicate the type of content in some text (e.g., a restaurant review), and an assessment of the sentiment of the text (e.g., positive), and so forth.
- Privacy and Security: Much of the mining and analysis involved in a big data context involves sensitive, private information. Therefore, we worked on technologies and policies for protecting and anonymizing users, and allowing them to retain control over their data.

Bigdata@CSAIL was a great success and many new tools and technologies were not only created but also adopted by our industry partners to address key challenges in their business. The Living Lab was successful in helping the technologies move from the lab to industry applications and we anticipate transitioning the Living Lab to a technology accelerator that can be used for all CSAIL research initiatives.

cybersecurity@CSAIL

Cybersystems cover communications, banking, data processing, purchasing, power and energy infrastructure, transportation, and defense—nearly every aspect of our lives. Consequently, cyberattacks 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, and secure logistics, sales, and personnel. Today these cybersecurity challenges spread across virtually all industry sectors and organizations are dealing with an ever increasing amount 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, cryptography, 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 cyberattacks more difficult, retain function despite such attacks, and allow a system to recover quickly after an attack. Cybersecurity@CSAIL intends to maintain an interdisciplinary focus that brings together thought leaders from industry and government with 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, cloud security, multiparty protocols, and usability of encrypted data.

Cybersecurity@CSAIL is an industry consortia model launched in March of 2015. Our industry partners provide valuable perspectives on the challenges faced across several industry verticals. Our partners include Akamai, BAE Systems, BBVA, Boeing, BP, Raytheon, and State Farm. Present projects undertaken though this initiative include the following:

- Certified Secure File System Applications (Professors Nickolai Zeldovich and Frans Kaashoek)
- Security Monitors for Industrial Control Systems (Howard Shrobe)
- Adversarial Analyses in Cybersecurity (Erik Hemberg and Una-May O'Reilly)
- Cyber Security in Multirobot Networks (Professor Daniela Rus)
- Security Monitors for Industrial Control Systems (Howard Shrobe and Professor Brian Williams)
- Co-Adversarial Dynamics of the Reconnaissance Phase in Advanced Persistent Threats in Software Defined Networking (Una-May O'Reilly and Erik Hemberg)

Internet Policy Research Initiative

The mission of the Internet Policy Research Initiative (IPRI) is to work with policy makers and technologists to increase the trustworthiness and effectiveness of interconnected digital systems. We accomplish this with targeted engineering and public policy research, various educational programs geared students and policy makers, and outreach programs to build policy communities that facilitate communication, education, and information exchange.

Communication and information networks are a fundamental infrastructure for our increasingly digital economy and society. Technologists and policy makers both play key roles in supporting this transition, yet they approach issues from different perspectives and often do not speak the same language. This can lead to not-fully-informed policy making or misdirected research efforts. There is a pressing need to bridge the gap between technical and policy communities because of society's reliance on this critical infrastructure. IPRI's core research efforts cover the following six distinct categories:

- Cybersecurity: IPRI's cybersecurity research focuses on security issues related to communication networks and software systems as they affect the economy and society as a whole. The work covers encryption policy, core infrastructure, and securing the Internet of Things.
- Privacy: IPRI maintains a strong focus on privacy policy, including its critical role in trustworthiness. Research projects include developing and strengthening privacy infrastructure, evaluating the relationships between security and privacy, and assessing the complex terrain of citizens' rights and state authority. Privacy is a key international focus area for the Initiative, with various projects coming out of the "Privacy Bridges" report and its 10 recommendations.

- Networks: IPRI research studies the communication networks that support the economy and society and includes topics related to assessing the reliability of services, reducing the cost of content delivery, measuring network performance, facilitating disclosure of mechanisms for an open Internet and spectrum licensing, and analyzing the future of Internet architecture.
- Critical Infrastructure: The digital systems that control critical infrastructure in the United States and most other countries are easily penetrated and architecturally weak. These vulnerabilities have been evident for a long time, but policy makers and system operators have tended to focus on short-term fixes and tactical improvements. The IPRI research stream focuses on developing and supporting short- and long-term recommendations that are applicable to critical infrastructure in the United States and abroad.
- Machine Understanding: With the increasing computational capability and amounts of devices such as sensors, neural networks, algorithms, and data, autonomous machines are doing incredible things. But with so many complicated parts and opaque algorithms, how are these machines arriving at their decisions? Are we able to audit their behavior, challenge their decisions in an adversary proceeding, and understand their methods and outcomes? As humans relinquish control to systems and machines, we need to be assured of the reliability and rationality of the systems and machines. As a society, when something goes wrong we assign responsibility and determine liability to activate legal frameworks—however, with machines operating and functioning on their own, how will they be held responsible? If their behavior is inadequate or inappropriate, the autonomous machine should be able to be corrected or disabled. IPRI's machine understanding team is developing the methodology and supporting technology for autonomous machines to explain themselves in a clear and concise way that humans can easily understand
- Internet Experience: IPRI research on the Internet experience includes the study of Internet governance, reflection on the role and evolution of information and communications technology (ICT) in society, and analysis of how key sectors use ICTs now and how they may in the future. Additional research tracks focus on socially-linked data, accountable systems, and the security of autonomous vehicles. As part of the innovative technical research, IPRI's App Inventor group provides a platform for developing policy-aware applications.

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. Over four hundred member organizations—including most of the world's leading technology companies—are working to enhance the capabilities of web documents and create the Open Web Platform for application development, available across a wide range of devices, enabling more people than ever before 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 into society. Video, social networking tools, user-generated content, location-based services, and Internet 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 the 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 many devices. Although the boundaries of the platform continue to evolve, industry leaders speak in unison about how HTML5 (published in October 2014) is the cornerstone for this platform. But 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, web open font format, real-time communications, the Semantic Web stack, and a variety of application programming interfaces. The platform continues to grow, and the W3C community, in turn, is growing to meet the demand.

With the completion of HTML5, there are many new areas of focus. Publicly noted security breaches have resulted in unprecedented attention to fixing cybersecurity. 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.

The demand is also driving W3C to expand its agenda and the size of its community. W3C launched Community and Business Groups in 2011. After six years over 8,000 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 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.

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:

Communication with Strong Anonymity

Srini Devadas

In an era of mass surveillance, maintaining anonymity on the Internet is an important yet very difficult challenge. Tor, the only widely deployed anonymity system, unfortunately fails to provide anonymity against an adversary who can globally monitor the Internet. Our two recent systems, Riffle and Atom, instead aim to provide cryptographic guarantees of anonymity, even against such a powerful adversary. Our first work, Riffle, uses a small number of servers, only one of which needs to be honest, to provide anonymity against a powerful adversary who monitors all of the Internet and controls all but one server and a large number of users. Riffle employs two major cryptographic primitives, verifiable shuffle and private information retrieval (PIR) to do so. More specifically, verifiable shuffle is used to protect the senders of the messages, meaning that the adversary cannot learn the origin of any message. PIR is used to protect the receivers of the messages, meaning the adversary cannot learn the recipient of any message. In our work, we propose improvements to the two primitives to scale our system to hundreds of thousands of users using a handful of commodity machines.

Unfortunately, Riffle and other prior systems that provide strong anonymity only scales vertically, meaning they can only scale to more users by making each server more powerful. This can become very expensive, or even impossible, when there are more than millions of users in the system. Atom, on the other hand, is a system with strong anonymity properties that also scales horizontally. That is, each server in Atom only handles a small fraction of the total messages, and as a result, adding more servers to the network increases the performance. At the same time, Atom provides similar levels of anonymity as Riffle, and ensures that a user is anonymous among all users against a powerful adversary.

Fetal MRI

Polina Golland

In collaboration with Professor Adalsteinsson's group and clinical colleagues at Children's Hospital led by Dr. Grant, Professor Golland's group aims to develop MRI-based biomarkers of placental function. The researchers are using MRI to characterize how well oxygen and other nutrients are transferred from the maternal blood stream to the fetus. The collaborative team demonstrated that MRI-based signals can be used to visualize normal function of the placenta and its dysfunctions. The results of this research have been recently published in Scientific Reports.

References

Liao, Ruizhi, Esra Abaci Turk, M. Zhang, Jie Luo, Ellen Grant, Elfar Adalsteinsson, and Polina Golland. "Temporal Registration in In-Utero Volumetric MRI Time Series." *Medical Image Computing and Computer Assisted Intervention* 9902, (2016): 54–62. doi: 10.1007/978-3-319-46726-9_7.

Luo, Jie, Esra Abaci Turk, Carolina Bibbo, Borjan Gagoski, Drucilla J. Roberts, Mark Vangel, Clare M. Tempany-Afdhal et al. "*In Vivo* Quantification of Placental Insufficiency by BOLD MRI: A Human Study." *Scientific Reports* 7, no. 3713 (2017). doi: 10.1038/s41598-017-03450-0.

Turk, Esra Abaci, Jie Luo, Borjan Gagoski, Javier Pascau, Carolina Bibbo, Julian N. Robinson, P Ellen Grant, Elfar Adalsteinsson, Polina Golland, and Norberto Malpica. "Spatiotemporal Alignment of In Utero BOLD-MRI Series." *Journal of Magnetic Resonance Imaging* 46, no. 2 (2017): 403–412. doi: 10.1002/jmri.25585.

Theory of Distributed Systems

Nancy Lynch

This year two students, Mohsen Ghaffari and Tsvetomira (Mira) Radeva, completed their PhD theses in our group. Ghaffari's thesis is entitled "Improved Distributed Algorithms for Fundamental Graph Problems." It is a tour de force, containing fast algorithms for computing maximal independent sets, for solving a large number of graph connectivity problems, and for running several algorithms concurrently in a single graph-based network. The preliminary conference versions of the results of this thesis were awarded six Best Paper or Best Student Paper prizes. Substantial new work was required this year to complete the results, for example, for developing a general algorithm transformation method for running distributed algorithms concurrently. Many other new graph network algorithms were developed this year by Mohsen and also by postdocs Stephan Holzer, Merav Parter, and Hsin-Hao Su; these addressed problems of graph edge coloring, orientation, random contractions, minimum spanning tree, and building fault-tolerant network structures.

Mira Radeva's thesis is entitled "A Symbiotic Perspective on Distributed Algorithms and Social Insects." It contains new algorithms and analysis for a variety of insect colony problems including foraging, task allocation, and house-hunting (finding and agreeing on a new nest). The newest work in the thesis involves the careful introduction of uncertainty (in estimation of local colony size) into our model of house-hunting, and shows that one of our algorithms is able to tolerate this uncertainty without much loss of efficiency. Furthermore, we showed that a density estimation algorithm we developed last year satisfies our limits on uncertainty, and therefore can be plugged into the noise-tolerant, house-hunting algorithm with no further loss.

In the area of wireless communication, a highlight was our development of a new aggregation tree abstraction layer to aid in writing algorithms for fading channel wireless network models, also known as signal to interference and noise ratio (SINR) models. This abstraction supports spatial reuse of wireless channels and provides a simple basis for writing algorithms for information dissemination and aggregation. We also have new algorithms for SINR models that depend on power adjustments; one of these (for leader election, by mechanical engineering student Evangelia Anna (Lilika) Markatou and her collaborators) won the Best Paper award at SIROCCO 2017.

In addition, this year, PhD student Cameron Musco and postdoc Merav Parter completed and published two papers on algorithms for stochastic spiking neural network models. These algorithms solve simple problems of recognition, comparison, and focus that might be solved in actual brain networks. We analyzed the algorithms in terms of network size and time for stabilization. We also proved near-matching lower bound results. We believe these serve as good initial examples of a new type of algorithmic work that one can do to help in understanding brain network behavior. Finally, postdoc Kishori Konwar and collaborators have produced a new layered algorithm for efficiently maintaining atomic read-and-write memory in a fault-prone distributed client-server system. The algorithm uses modern data coding techniques. The two layers allow mixing and matching of coding and concurrency control methods.

Self-Supervising Networks

Antonio Torralba

Computer vision is undergoing a revolution. One of the key reasons for the recent successes in computer vision is the access to massive annotated data sets that have become available in the last few years. Several of those data sets have been created at MIT. Unfortunately, creating these data sets is expensive and labor intensive. To overcome the limits of data set building, researchers are now taking inspiration from how humans learn. Humans do not require massive annotated data sets in order to learn to perceive the world, in fact, most of the time, kids learn by themselves. In our work, we have shown how using inputs from different sensory modalities (vision and audition) one can build systems that learn by themselves without requiring human annotated data.

Sound conveys important information about the objects in our surroundings. In our 2016 paper, "Ambient Sound Provides Supervision for Visual Learning," we show that ambient sounds can be used as a supervisory signal for learning to see. Although human annotations are indisputably useful for learning, they are expensive to collect. The correspondence between ambient sounds and video is, by contrast, ubiquitous and free. To demonstrate this, we trained a deep neural network to predict the sound associated with a video frame. We show that, through this process, the system learns to recognize objects that are often associated with characteristic sounds.

Laboratory Sponsored Activities

CSAIL Outreach

CSAIL's Hour of Code: CSAIL regularly hosts a presentation and demo fair in conjunction with the global Hour of Code movement, inviting local high-school students to learn more about a wide array of computer science research. Hundreds of local STEM students have attended, with CSAIL receiving support from high-profile public figures such as author John Green and musician will.i.am. of the Black Eyed Peas.

Reddit "Ask Me Anything": CSAIL regularly encourages the online community to submit questions about computer science and academia in a series of Reddit "Ask Me Anything" (AMA) sessions involving the lab's researchers. CSAIL's AMAs have spurred approximately 6,000 comments and questions, as well as more than 200,000 page views.

Middle East Education through Technology: CSAIL has been a long-time supporter of the Middle East Education through Technology (MEET) program, an innovative educational initiative aimed at creating a common professional language between young Israelis and Palestinians. MEET enables its participants to acquire advanced technological and leadership tools while empowering them to create positive social change within their own communities. Many MIT students volunteer to teach MEET summer courses at the Hebrew University in Jerusalem. CSAIL continues to provide financial support for the program.

Dertouzos Distinguished Lecture Series

The Dertouzos Lecture Series has been a tradition since 1976, featuring some of the most influential thinkers in computer science. Two speakers presented lectures during the AY2017 Dertouzos Distinguished Lecture Series. On November 2, 2016, Associate Professor Sarah Parcak of University of Alabama-Birmingham presented "Hacking Archaeology: Beyond Shovels or iSandbox?"; and on February 8, 2017, Bill Thies, senior researcher at Microsoft Research India presented "Frugal Innovations for a Developing World."

CSAIL Research Highlights

On November 3, 2016, CSAIL hosted a research highlights party for CSAIL graduate students and postdoctoral associates, featuring 16 students presenting their research to members of the CSAIL community. It was an engaging event to build community, support the professional development of our students, and learn entirely new things about research at CSAIL.

Organizational Changes and Personel

Professor Daniela Rus has continued her role as director of CSAIL. The director's duties include developing and implementing strategies designed to grow and evolve CSAIL, fund raise, determine laboratory policies, and examine promotion cases.

CSAIL's leadership team includes two associate directors, a chief operating officer (COO), and an executive cabinet. These leaders assist the director with her duties. These positions are appointed by the laboratory's director. Professors Daniel Jackson and Polina Golland became the associate directors in October 2014. Professor Golland's term ended June 2017. Professor Charles Leiserson is associate director and chief operating officer, providing leadership and strategy for how we conduct our operations and events, enabling the director to allocate more time to strategic planning. Professor Victor Zue holds the role of director of international relations, managing the engagements and oversight of various important CSAIL international contracts and international contract negotiations.

Additionally, the CSAIL executive cabinet meets twice per month to review and advise the director on policy, processes, and activities within the laboratory. Members of the executive cabinet include, Hal Abelson, Edward (Ted) Adelson, Saman Amarasinghe, Regina Barzilay, Randall Davis, David Gifford, Polina Golland, Daniel Jackson, Charles Leiseron, Samuel Madden, Ronitt Rubinfeld, Daniela Rus, Nir Shavit, and Victor Zue.

The CSAIL enterprise services team manages lab operations. There are seven units—Administrative Assistants, CSAIL Alliance Program, Communications, Finance, Human Resources, Special Projects, and the Infrastructure Group—reporting to the CSAIL COO on all operational matters. Carmen Popovici is the acting assistant director for administration. John Costanza is the assistant director for infrastructure, overseeing information technology infrastructure and user support, building operations, and communications. Lori Glover is managing director of the CSAIL Alliance Program. Victoria Palay is the senior manager of special projects.

Saman Amarasinghe is the current space czar, overseeing the space committee and managing 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.

Awards and Honors

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

- Anant Agarwal and Padma Shri, distinguished service award, Government of India
- Hari Balakrishnan, member, American Academy of Arts and Sciences; Test-of-Time Paper Award, Association for Computing Machinery SIGMOBILE
- Tim Berners-Lee, A. M. Turing Award, Association for Computing Machinery
- Erik Demaine, fellow, Association for Computing Machinery
- Srini Devadas, W. Wallace McDowell Award, Institute of Electrical and Electronics Engineers
- Fredo Durand, fellow, Association for Computing Machinery; Computer Graphics Achievement Award, Association for Computing Machinery SIGGRAPH
- William Freeman, fellow, Association for Computing Machinery
- Shafi Goldwasser, honorary doctorate, Bar Ilan University; member, Russian National Academy of Science; metal of distinction, Barnard College; honorary doctorate, University of Haifa
- Tommi Jaakkola, fellow, Association for the Advancement of Artificial Intelligence
- Daniel Jackson, Impact Paper Award, Association for Computing Machinery Special Interest Group on Software Engineering; fellow, Association for Computing Machinery
- Dina Katabi, member, National Academy of Engineering; Best Paper, Association for Computing Machinery SIGCOMM
- Manolis Kellis, Faculty Research Innovation Fellowship, MIT Electrical Engineering and Computer Science
- Butler Lampson, member, National Cyber Security Hall of Fame

- Tom Leighton, member, National Inventors Hall of Fame
- Charles Leiserson, Best Paper, Association for Computing Machinery SIGPLAN
- Aleksander Madry, Research Award, Google
- Ankur Moitra, fellow, David and Lucille Packard Foundation
- Stefanie Mueller, 30 Under 30: Science, *Forbes*
- Aude Oliva, Vannevar Bush Faculty Fellows, Department of Defense
- Daniela Rus, Joseph F. Engelberger Robotics Award for Education, Robotic Industries Association; member, American Academy of Arts and Sciences
- Julie Shah, fellow, Radcliffe Institute for Advanced Study at Harvard University
- Howard Shrobe, fellow, American Association for the Advancement of Science
- Michael Sipser, fellow, MacVicar
- Justin Solomon, 30 Under 30: Science, *Forbes*

Key Statistics for Academic Year 2017

Faculty: 94 (16% women)
Research staff: 29 (20% women)
Administration, technical, and support staff: 95 (51% women)
Postdocs: 83 (16% women)
Visitors: 73 (16% women)
Paid Undergraduate Research Opportunities Program participants: 155 (26% women)
Master of engineering students: 92 (24% women)
Graduate students: 392 (18% women)

More information about the Computer Science and Artificial Intelligence Laboratory can be found at our website.

Daniela Rus
Director, Computer Science and Artificial Intelligence Laboratory