Computer Science and Artificial Intelligence Laboratory

The Computer Science and Artificial Intelligence Laboratory (CSAIL) studies computing and its application to finding solutions to many of the most challenging problems of our lives, our work, and our world. We employ a long-term lens to engineer innovative solutions in an effort to unlock the secrets of human intelligence, extend the functional capabilities of machines, explore human–machine interactions, and enhance the human experience. Our mission is to be an innovative and passionate community driving the future of computing and making the world a better place.

With approximately 50 research groups working on hundreds of diverse projects, CSAIL researchers focus on finding innovative ways to make systems and machines operate faster, better, more safely, more easily, and more efficiently for the benefit of humanity. Our projects fall into three areas of inquiry:

> Artificial intelligence—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, 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 the 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 important social challenges. Current research explores cloud computing, mobile computing, the next generation of computing devices, autonomy, intelligence, and the application of sensor technology to traffic congestion, medical monitoring, and climate and environmental observations. Robotic locomotion and the human–robot interface are being investigated, as are advanced software-based medical instrumentation and medical informatics systems to aid clinical decision making. 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 the understanding of gene regulation.

CSAIL research is sponsored by a large number of diverse entities, from US government contracts to the private sector. US 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 (DARPA), the Department of Defense Research and Engineering, the US Department of Education, the US 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 the Naval Air Systems Command), and the Space and Naval Warfare Systems Center. US and international industrial sponsors include Boeing, Ford, Foxconn, General Electric, Intel Corporation, Lockheed Martin Advanced Technology Laboratories, Microelectronics Advanced Research Corporation, Nippon Telegraph and Telephone Corporation, Nokia, Northrop Grumman Corporation, Quanta Computer, Inc., Shell, Siemens, and State Farm Insurance. Other organizations sponsoring research include Aarhus University, Battelle Memorial Institute, Delta Electronics Foundation, DSO National Laboratories, Epoch Foundation, Industrial Technology Research Institute, Nanyang Technological University, and the Singapore–MIT Alliance.

Research Projects

CSAIL comprises many single- and multi-investigator projects as well as a number of virtual centers and large-scale projects. These projects and collaborations include:

T-Party

The T-Party project is a 10-year, $45.5 million research collaboration with Quanta Computer, Inc. The project has just completed its seventh year and is currently funding four major research areas:

- Cloud computing technologies: new, scalable operating systems for clusters of multi-core computers (professor Anant Agarwal); investigation into scalability limitations of standard operating systems on multicore platforms (professor Frans Kaashoek); scalable data stores supporting high-traffic websites (professor Robert Morris); security frameworks that limit data dissemination and allow recovery from attacks (professor Nickolai Zeldovich); new database architectures that increase concurrency of cloud-based implementations (professor Samuel Madden); trusted storage built using untrusted commodity cloud storage (professor Srinivas Devadas).
- Human-computer interfaces for interactions between mobile devices and the cloud: continued evolution of a framework for natural language understanding and dialogs (Dr. Stephanie Seneff); an embeddable web-based interface for multimodal interactions—speech and gesture (Dr. James Glass); creating easy-to-use interfaces using web automation (professor Robert Miller).
- Telepresence: Novel data structures that allow efficient client manipulation of large server-side data streams (professor Lucas Ward); applying machine learning techniques to medical telemetry to aid in the appropriate diagnosis of chronic diseases, and medical telepresence incorporating multiple video, audio, and telemetry channels between ambulance and hospital (professor John Guttag).

- Computational photography: many unique and effective algorithms to manipulate digital images (professor Fredo Durand); algorithms to extract information content or important features from images (professor William Freeman).

Basic Research in Interactive Computing

In January 2008, CSAIL started a three-year collaboration with Foxconn for basic research encompassing several areas of computer science, ranging from networking, human–computer interactions, computer graphics and vision, to theory. This research is predicated on the belief that computers and information devices are fast becoming interactive; they interact with other computers, with their environments, and, above all, with humans. Each form of interaction adds a new dimension to the challenge of modeling and understanding the behavior of computer systems, as well as to the task of building and using these systems.

The first phase of this collaboration concluded in December 2010. Phase two started in January 2011, this time supporting four projects and six principal investigators (PIs) as shown below:

- Advanced Hand-Tracking and Gesture-Based Interaction (professor Randal Davis);
- Factor-Analysis-Based Speech and Language Analysis (Dr. James Glass);
- Activity Modeling and Recognition: Using Phones as Sensors (professor Daniela Rus); and
- 3-D Capture and Display for Mobile Devices (professors Fredo Durand, William T. Freeman, and Antonio Torralba).

Nine Foxconn engineers in two shifts have visited CSAIL during this academic year to facilitate technology transfer and to receive training for advanced research and development in computer science. In addition, video presentations for the four projects were made, via video conferencing from CSAIL to Foxconn in Taiwan, on December 7–8, 2011, to summarize the research results.

Research in Cyber Security

In September 2009, CSAIL established a five-year research collaboration with Northrop Grumman Information Technology Inc. in the research area of cyber security. In so doing, MIT joined Carnegie Mellon University and Purdue University as part of the Northrop Grumman Cyber Security Consortium.

This project has just completed its third year. During that year, four projects and five PIs were supported by Northrop Grumman:

- An Argument Engine for System Security and Dependability (professor Daniel Jackson);
- Secure Audit Trails (professor Barbara Liskov);

- Secure Cloud Applications Using Practical Data Flow Assertions (professors Srinivas Devadas and Nickolai Zeldovich); and
- An Argument Engine for System Security and Dependability (Dr. Jeffrey Jaffe, W3C).

This project is progressing well, and specific proposals for funding for the fourth year are now being considered.

The MIT Angstrom Project: Universal Technologies for Exascale Computing

The CSAIL-led Angstrom team was one of four teams selected by DARPA for funding under the ubiquitous high-performance computing program. The team is charged with the task of rethinking computing and creating a fundamentally new computing architecture to meet the challenges of computing in the 2020 timeframe. The CSAIL team is the only university-led team of the four, consisting predominantly of university researchers, most of whom are MIT faculty. Angstrom is a strongly interdisciplinary program involving faculty from MIT's CSAIL, Microsystems Technology Laboratories, Research Laboratory of Electronics, and Material Processing Center; industry partners Freescale Semiconductor, Mercury Federal Systems, and Lockheed ATL; and faculty from the University of Maryland's Department of Electrical and Computer Engineering. Project Angstrom's goal is to create the fundamental technologies needed for future extreme-scale computers. Extreme-scale computers face a number of major challenges, the most difficult of which are energy efficiency, scalability, programmability, and dependability. To address these challenges, the Angstrom team is reexamining every layer of system design and interfaces, including circuits, hardware architecture, operating systems, runtime software systems, compilers, programming languages, and applications.

Project Angstrom's vision for addressing the major challenges of extreme-scale computing is based on two key foundations: creating both a revolutionary self-aware computational model, called SEEC, and a fully distributed factored architecture for both hardware and software. Self-aware systems will not only make computers smarter, they could also prove essential for managing ever-more complex computers in the future. Over the past decade, computer engineers have added more and more basic computing units, called cores, to computers. Today's computers have four to eight cores, but future machines will use anywhere from dozens to thousands of cores. That would make the task of splitting up computational tasks among the cores, which programmers now do explicitly, nearly impossible. A self-aware system takes that burden off the programmer, adjusting the program's core use automatically. Other systems will do the same for memory usage and network usage.

In the past year, Project Angstrom completed an overall system design. The design includes the architectural design of a 1,000-core processor that has self-awareness at all levels of abstraction to provide near-optimal results with minimal programmer effort. Project Angstrom's self-aware computing was selected by the editors of *Scientific American* as one of "10 world changing ideas" in the December 2011 issue. In 2012–2013, the Angstrom design will be refined and the implementation effort will begin.

Self-Aware Computational Model

Project Angstrom's Self-Aware Computational Model (SEEC) is a goal-oriented computational model that radically increases developer productivity by abstracting traditional procedural programming into goals (e.g., "achieve the best possible chess move within 10 seconds while burning less than 20 watts") that are targeted in the self-aware, factored system. SEEC attempts to enable systems that are orders of magnitude more energy efficient and dependable than today's by explicitly incorporating energy and resiliency goals into the hardware, operating system, compiler, and language. A major goal of this research is to create and to evaluate algorithms and interfaces for SEEC using methods based on machine learning and control theory.

Distributed Factored Architecture

Our factored approach targets energy-efficient multicores that are scalable to thousands of cores. For example, distributed power converters will scale, because they eliminate centralized control bottlenecks and allow fine-grain voltage control; this facilitates SEEC, which demands individual control of the voltage, clock, and body bias of each core. Similarly, the factored software targets levels of resiliency and scalability that meet the demands of billions of threads. For example, the self-aware factored operating system will factor operating system functions into services (e.g., scheduling service or fault tolerance service) that are implemented by a dynamic fleet of cooperating servers. Accordingly, our second major goal is to invent fully distributed architectural mechanisms and factored software approaches.

SEEC and factoring are the two overarching themes of Project Angstrom. These two concepts, instantiated in several novel mechanisms, will provide solutions for the four major extreme-scale challenges.

Rethinking Artificial Intelligence

Rethinking artificial intelligence (AI) is exploring a number of unconventional approaches to the fundamental goal of developing a computational model of human-like intelligence. The four projects in this research initiative come at the problem from a number of directions.

Work on learning to abstract starts from the premise that abstraction is a key obstacle in learning. To learn abstraction from observations, we need rich probabilistic representations that support the creation of templates and that enable efficient assessment of the predictive value of templates, i.e., the ability to recognize quickly when and where templates may apply. Two key insights afford stepping stones toward these goals: (i) templates must be designed in a way that naturally enables step-wise generalization, and (ii) structural hashing methods that make possible fast hierarchical recognition and tracking of templates must be developed. We believe our framework is particularly well suited for learning abstractions across different modalities and hope to demonstrate cross-modal transfer in the context of learning to perform visually guided actions.

Work on algorithmic encoding seeks to understand how to guarantee computational results with unreliable, weakly communicating components—that is, how to adaptively coax a set of noisy, uncertain components, into collectively achieving the task at hand. This is a long-standing problem with major implications across a wide range of fields, including AI, fault-tolerant computing, error-correcting codes, bioengineering, and neuroscience. Unlike previous work in fault tolerance, however, we do not conceive of the task as guaranteeing that every operation will succeed (or that recovery is possible when an operation fails). Instead, we assume that errors are commonplace and heterogeneous, and seek to redefine how algorithms should be written in this context.

Work in probabilistic programming combines Bayesian probability with universal computation. Instead of modeling joint distributions over a set of random variables, probabilistic programs model distributions over the execution histories of programs, including programs that analyze, transform, and write other programs. Probabilistic programs, as the first computationally universal systems for probabilistic modeling, have capacities for abstraction and recursion that can support reflective reasoning, natural language semantics, and the learning of model structure from probabilistic inference. They can naturally express core common-sense knowledge about physical objects, intentional agents, and their causal interactions, and are starting to be used in cognitive science to model people's intuitive physical and psychological reasoning. They thus have the potential to provide a new approach to AI by bringing the full power of probabilistic inference to bear on classic AI problems that have been largely left out of the statistical AI revolution of the past two decades.

Work in propagators starts from the presumption that we need to develop a computational account of the sort of processing that takes place in our brains and we need to understand why that kind of processing is so powerful, robust, agile, and adaptable. Engineers tend to think in terms of information flowing mostly in one direction, through functionally circumscribed modules, with only an occasional feedback loop sending information in the other direction. As the biological sciences indicate, the brain is very different: for any neural pathway carrying information up, there is another pathway carrying information down or sideways. The first step toward our goal is to develop new, more brain-like ways of computing.

The propagator model suggests a fundamental shift in viewpoint: the things manipulated by basic computing elements are not fixed values, but are instead information about values—information that is continually refined as new information becomes available. Propagators are autonomous independent machines, interconnected by shared cells through which they communicate. Each propagator continuously examines the cells it is connected to, and adds information to some cells based on computations it can make from information it can get from others. Cells accumulate information; propagators produce information. This makes a natural computational structure for building powerful systems that fill in details.

Explorations in Cyber International Relations

The Explorations in Cyber International Relations project is a collaboration among CSAIL, the Political Science Department and the Sloan School at MIT, and the Kennedy

School of Government at Harvard University. It is one of the projects funded under the US Department of Defense Minerva Research Initiative. The project is motivated by the hypothesis that the emergence of cyberspace as a phenomenon calls for new theories of international relations. The current field of international relations emerged in the era of nuclear deterrence and focuses on the role of key state actors. It emphasizes concepts of balance of power and a 20th-century world order dominated by a few major players. In contrast, cyberspace (or the internet, more specifically) is constructed by private-sector players, with little state involvement, and empowers a wide range of actors, from large and small states to individuals. At the same time, with growing fears of cyberwar and tensions between states (such as the US and China) over how the internet should be regulated, it is no longer reasonable for governments to ignore the influence and importance of the internet.

The collaboration between CSAIL and the Political Science Department has the objective of developing a framework or model that incorporates both traditional elements of international relations and a technically sound model of cyberspace. This model will provide a means by which to categorize and distinguish different phenomena in cyberspace, so as to position them within the correct scope of international relations analysis and permit the development of relevant theory. Phenomena include moments of conflict (such as cyberattacks in Georgia and Estonia), attempts at global cooperation (such as the Internet Governance Forum), and ongoing international tensions over theft of intellectual property and repression of free speech.

We continue the work, reported last year, on novel extensions to game theory that remove some of the key limiting assumptions of that approach, and on technology for the automated extraction of meaning from narrative, such as news reports and government statements. We have developed a new way of representing the structure of cyberspace (in contrast to the traditional layered model used to describe technology such as the internet). This new approach, which we call control point analysis, captures the dynamic interaction among parts, as opposed to the static structure of layers. Control point analysis seems to be an effective basis for cross-disciplinary collaboration in the study of cyberspace. We have made an extensive study of the different ways in which personal identity is tracked and managed on the internet, with the goal of understanding how anonymity and accountability can both be accommodated in suitable contexts. We are undertaking an extensive study of different governance mechanisms for the internet, contrasting the older, bottom-up organizations and the emerging top-down, state-centered organizations, such as the International Telecommunications Union.

The Ford–MIT Alliance

The Ford–MIT Alliance, an Institute-wide initiative, was established in 1997. In 2012, the alliance is completing the final year of its third phase of funding, with renewal of funding, for a fourth phase to begin in 2013, in progress. The alliance is the MIT’s longest-running large-scale commitment from industry and represents a significant acknowledgment by the Ford Motor Company of MIT’s relevance and impact on its research. Since 1997, the alliance has funded more than 100 projects of varying duration and budget across MIT, with a total investment by Ford to date of more than $40 million.

The Ford–MIT Alliance research portfolio has a strong connection to CSAIL's research areas. CSAIL faculty are currently conducting research projects with Ford–MIT Alliance funding in four areas:

- Multiple vehicle networking;
- Algorithms for driving history analysis;
- Novel speech interfaces; and
- Crowd computing.

CSAIL hosts a testbed of two Ford vehicles, a 2011 Ford Explorer and a 2012 Ford Focus, for use in these projects.

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. Three hundred seventy-five member organizations, including most of the world's leading technology companies, are working to enhance the capabilities of web documents and to create an open web platform for application development, available across a wide range of devices, eventually enabling planet-wide collaboration.

Recent Focus

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 communications, television, publishing, automotive, entertainment, games, 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.

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 nearly in unison about how HTML5 will be the cornerstone for this platform. But the full strength of the platform relies on many more technologies that W3C and its partners have been creating, including cascading style sheets (CSS), scalable vector graphics (SVG), web open font format (WOFF), real-time communications (WebRTC), the Semantic Web stack, extensible markup language (XML), and a variety of application programming interfaces (APIs). The platform continues to grow, and the W3C community, in turn, is growing to meet the demand.

The demand is also driving W3C to expand its agenda and the size of its community. W3C launched Community and Business Groups in 2011. Just 11 months after launch, more than 1,500 people are participating in over 90 groups. 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 pre-standards and standards track work.

The Open Web Platform

The open web platform, a cornerstone of the W3C standardization effort, is a suite of technologies that is transforming business practices, creating new business models, and allowing for greater innovation on the web while reducing product lifecycle costs compared with other systems. W3C is designing royalty-free technologies that:

- Provide a rich interface feature set, including styles, interaction, and media;
- Enrich apps through APIs for device capabilities and user data;
- Integrate data and services (mashups, integration of existing databases and services);
- Run on any device (computer, telephone, television, consumer electronics, automobile, etc.) and support interaction through a variety of input and output modes;
- Meet network and communications demands (cross-origin resource sharing, real-time communications);
- Satisfy performance and distribution requirements, enable rapid development and deployment, and facilitate maintenance; and
- Address diverse social requirements for privacy, security, multilingual content, and accessibility.

New Centers and Initiatives

Intel Science and Technology Center in Big Data

MIT PIs Samuel Madden and Michael Stonebraker are heading a new Intel Science and Technology Center in Big Data (ISTC) based at CSAIL. Professors Madden and Stonebraker are leading a team of 20 researchers (from Brown University, MIT, Portland State University, Stanford University, University of Tennessee, University of California, Santa Barbara, and University of Washington) in a project to build new software tools to process and manage massive amounts of data. Specifically, the center is focused on new data management systems and new computer architectures that together can help users process data that exceeds the scale, rate, or sophistication of data processing that existing systems provide. The center will develop new technologies in a number of areas, including data-intensive scalable computing, machine learning, computer architecture and domain sciences (genomics, medicine, oceanography, imaging, and remote sensing).

The center is focused on five major research themes:

- Big data databases and analytics: developing new software platforms for processing massive amounts of data and applying analytics beyond what conventional relational systems can do. Examples include database systems focused on arrays (massive scale linear algebra) and graphs (for link and social network analysis).

- Big data math and algorithms: designing and implementing algorithms for linear algebra, signal processing, search, and machine learning that scale to tens or hundreds of machines and petabytes of data.
- Big data visualization: designing visualizations and interfaces that allow users to interact with massive data sets, on displays ranging from phones to video walls.
- Big data architecture: understanding how next-generation hardware innovations, such as many-core chips, non-volatile random-access memories, and reconfigurable hardware, affect the design of data processing systems.
- Streaming big data: Building data processing systems that facilitate rapid processing and ingestion of data streams.

bigdata@csail

A collaborative initiative between CSAIL and a number of industry partners, the goal of bigdata@csail is to identify and develop the technologies needed to solve the next generation of data challenges, which will require the ability to scale well beyond what today's computing platforms, algorithms, and methods can provide. Specifically, bigdata@csail is intended to enable people to truly leverage big data by developing platforms that are reusable, scalable, and easy to deploy across multiple application domains.

Our approach includes two key aspects: first, we will collaborate closely with industry to provide real-world applications and drive impact; second, we view the big data problem as fundamentally multi-disciplinary. Our team includes faculty and researchers across many related technology areas, including algorithms, architecture, data management, machine learning, privacy and security, user interfaces, and visualization, as well as domain experts in finance, medical, smart infrastructure, education, and science.

Founding members include American International Group, EMC, Intel, Microsoft, SAP AG, and Thomson Reuters.

EdX: Planet-Scale Online Learning

EdX is an online learning initiative of MIT and Harvard. EdX offers free online courses to students worldwide, aiming to transform education in quality, efficiency, and scale through technology and research, and to promote unique learning opportunities for students wherever there is access to the internet. EdX, and its predecessor *MITx*, were incubated at CSAIL in fall 2011 and spring 2012. Anant Agarwal, former director of CSAIL, serves as the first president of edX.

A CSAIL team built the prototype cloud-based platform for edX, which included interactive videos, computer graded exercises and tests, an online simulation-based laboratory, a discussion forum, a wiki, and a free online textbook. CSAIL members also taught the inaugural course for edX, called 6.002x Circuits and Electronics.

More than 154,000 students from around the world enrolled in this course. Students could register free and interact with each other and with the course staff on a discussion forum. A group of 20 on-campus students served as beta testers in a blended model on-campus class. EdX has now spun off as a not-for-profit startup from MIT and CSAIL, and has moved to space in Kendall Square at 11 Cambridge Center, which is a block from CSAIL.

Qatar Computing Research Institute

MIT and the Qatar Foundation have signed a multi-year agreement for funding collaborative research in computer science. The agreement covers collaborations between CSAIL and the Qatar Computing Research Institute (QCRI). It provides a total of $35 million of funding for CSAIL projects over seven years. The objective of the agreement is to create connections between researchers in Qatar and those at MIT, and advance the state of the art in various areas in computer science.

Under this agreement, CSAIL and QCRI have processed an initial fast-track funding cycle. A call for proposals was sent to the PIs. A steering committee that includes members from both CSAIL and QCRI reviewed the submissions and picked six proposals for funding. The funded proposals cover a wide range of topics: news and social media analytics, opinion mining in Arabic multimodal dialogue systems, multilingual speech informatics, data integration in database systems, generating event summaries from social media streams, and high-resolution genetic mapping. A second round of funding is currently in progress, with proposals due by the end of August 2012.

The projects developed under this agreement further CSAIL's research and educational mission, leading to innovative technologies and supporting many students and postdoctoral researchers.

MIT Center for Wireless Networks and Mobile Computing

The goal of the MIT Center for Wireless Networks and Mobile Computing (Wireless@ MIT) is to develop next-generation wireless network technologies and mobile computing systems. The features of the center are:

- An interdisciplinary focus that brings together more than15 MIT professors and their groups, conducting research in networking, communication and information theory, systems, security, hardware, algorithms, and societal applications (transportation, health care, and autonomous systems), and
- A strong industrial partnership and an emphasis on influencing and affecting standards and products.

The center's activities include the following topics:

- Spectrum utilization and mobile network connectivity: The center focuses on developing new solutions that enable communication-intensive applications across large numbers of wireless and mobile devices. These solutions include advanced interference management techniques, better coding and transmission schemes that increase spectral efficiency, scalable protocols for heterogeneous radio networks, and novel protocol architectures for mobility. The research covers licensed cellular networks, unlicensed bands, emerging frequencies, and dynamic spectrum access.

- Mobile applications: The center explores and develops distributed application designs that work across mobile devices and computing and storage clouds. This research includes new programming models, abstractions, and communication protocols for mobile environments. Studied applications include transportation systems and vehicle-to-vehicle networks, scalable mobile video, personalized content delivery, sensor-intensive tasking and coordination applications.
- Security: The center aims to make security configuration of wireless devices simple and effortless, and hence accessible to ordinary users. The center also works on improving the security of embedded wireless systems and enhancing the privacy of users of mobile applications. This research incorporates physical layer techniques, modern cryptographic protocols, and system software.
- Energy: The center develops hardware and software designs for energy-efficient mobile systems. These include low-power handsets, energy-scavenging sensors, and wireless medical devices.

The center's directors are professors Hari Balakrishnan and Dina Katabi. Between January and July 2012, seven companies have joined the center: Amazon, Cisco, Intel, Mediatek, Microsoft, STMicroelectronics, and Telefonica. The center is planning a public launch and an initial member retreat in October 2012.

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 described below.

Quantum and Classical Complexity Theory

With four PhD students—Andy Drucker, Michael Forbes, Alex Arkhipov, and Adam Bouland—and a postdoctoral researcher, Thomas Vidick, professor Scott Aaronson and his research group have been continuing their work on understanding the fundamental capabilities and limits of quantum computers and their relation to classical computers. In collaboration with Alex Arkhipov on the computational complexity of linear optics, they proposed an experiment involving a network of beam-splitters that probably wouldn't yield a universal quantum computer, but would probably be easier to build, and for which there was evidence that the experiment would perform a sampling task that is intractable using a classical computer. This work was featured in MIT News, and two experimental groups, that of Jeremy O'Brien in Bristol, UK, and that of Andrew White in Queensland, Australia, are currently working to implement it.

Scott Aaronson, with MIT undergraduate Paul Christiano, worked on a new scheme for public-key quantum money using hidden subspaces, which solved many of the problems with previous schemes. It proved possible to base the security for this scheme on a relatively conventional cryptographic assumption.

Vidick's work on the power of multi-prover quantum proof systems, and especially his recent proof with Tsuyoshi Ito that MIP* contains NEXP, showed that entangled provers can convince a polynomial-time verifier of all languages in nondeterministic exponential time. This solved a long-standing open problem and was arguably the most important advance in quantum complexity theory of the past year. Drucker's work on ruling out quantum and classical "instance compression" schemes, recently highlighted on Lance Fortnow's Computational Complexity blog, solved a long-standing open problem in the field. This research is under the direction of associate professor Scott Aaronson.

Bedrock: An Environment for Developing Software with Mathematical Correctness Guarantees

Today it is taken for granted that deployed software will contain many security and reliability problems that remain hidden until triggered by real usage, often with serious consequences. In contrast, most engineering disciplines have developed mathematical analysis methods whereby systems can be certified to exhibit low probabilities of catastrophe. The field of software engineering needs similar foundations to overcome the crisis in software quality. Many theoretical frameworks for program certification have been developed, but relatively few practical tools exist for applying them. As a realistic software project contains detail beyond the capacity of a human to remember, it is essential that software packages be developed to drive the certification process via analysis of the source code of other software packages.

Professor Chlipala's group has been developing such a tool, called Bedrock, that is intended for use with some of the most fundamental and widely used pieces of software, such as operating systems and programming language implementations. Through a mixture of automated analysis and human-written mathematical proof, Bedrock can produce a formal certificate explaining why a program adheres to a mathematical specification of behavior. A certificate can be checked by a relatively small, trustworthy checking program that is much simpler than the programs being analyzed and that can be reused for the analysis of arbitrary programs. Thus, just as mathematical analysis of a new bridge rests on such foundations as differential equations, the Bedrock approach reduces a software correctness argument to the austere foundations of logic and discrete mathematics.

Bedrock is distinguished from related tools by its focus on reducing programmer effort. A fully rigorous program correctness certificate must consider many details of program code and behavior, but most details are uninteresting. Bedrock applies novel algorithms to fill in most of the details without human intervention. In certification of key operating system pieces, such as a process scheduler and a memory manager, we have demonstrated that our automation procedures can reduce by 100 times the complexity of the argument that the programmer must fill in manually, compared with past work. The Bedrock project also considers how programming languages ought to be designed differently to make certificate production easier, studying novel ways to let programmers add new features to a programming language when the features are accompanied by suitable certificates.

Ongoing work focuses on scaling Bedrock to certification of larger software systems. Current plans are to certify realistic software platforms for internet servers and autonomous vehicle controllers. This research is under the direction of assistant professor Adam Chlipala.

Integrating Planning, Perception, and Control for Robots in the Real World

If a robot is to operate in a complex environment over a period of hours or days, it must be able to plan actions involving large numbers of objects and long time horizons. It must also be able to plan and carry out actions in the presence of uncertainty, both in the outcome of its actions and in the actual state of the world. We are developing an approach to robot planning that addresses these challenges by integrating several key ideas:

- Planning in belief space, that is, the space of probability distributions over the underlying state space. This enables a principled approach to planning in the presence of state uncertainty.
- Planning with simplified models and re-planning as necessary. This enables planning efficiently when outcomes are uncertain while still enabling action choices based on looking ahead into likely outcomes.
- Combining logical and geometric reasoning. This enables detailed planning in large state spaces involving many objects.
- Hierarchical planning with interleaved execution. This enables plans with very long time horizons by breaking up the planning problem into a sequence of smaller problems.

Based on these ideas, we are constructing a system for planning, perception, and execution on a Willow Garage PR2 robot. The software developed in this project will be freely available for easy porting to a wide variety of robots. The techniques developed in this project are more widely applicable: we are applying them not only in robots but also to problems of persistent surveillance and of understanding the intentions of humans, in order to build software and hardware agents to assist humans in a variety of contexts. This work is supported by the National Science Foundation, the Air Force Office of Scientific Research, the Office of Naval Research, the Singapore-MIT Game Lab, the Ford–MIT Alliance, and the MIT–SUTD International Design Center. This research is under the direction of professors Leslie Kaelbling and Tomas Lozano-Perez.

Current Projects of the Clinical Decision-Making Group

Medical Natural Language Processing

As part of several national-scale collaborations, we are working on extracting the clinically significant data from narrative text records found in patients' electronic medical records. These texts include referring notes, emergency department notes, nurses' and doctors' notes during in-patient care, visit notes from outpatient encounters, specialty reports such as pathology and radiology, etc. Our goal is to extract (1) mentions of entities such as diseases, symptoms, medications, procedures, lab test, test results, etc.; (2) characterizations of each entity, such as its duration, severity, dosage details, etc.; (3) attributions about whether such a mention is certain, negated, or about someone other than the patient; and (4) relations among the mentioned entities, e.g., "drug x is given to control symptom y" vs. "drug x caused symptom y as a side effect." We use mainly statistical machine learning models, augmented by pattern matchers and medical dictionaries and thesauri.

Our techniques are, for certain specific tasks, more effective than asking human annotators to read the narratives and extract the same information. Across all the tasks listed above, recall and precision vary considerably, depending on the nature of the data to be extracted, the specific kind of data being processed, and the amount of available training data. We have also applied the same language processing methods to the problem of de-identifying clinical text so that it can be used for research purposes without compromising patient privacy. Those programs achieve F-measures of 98 percent or more for the most critical data elements, such as patient names, though they are less effective for items such as hospital names, where we have less training data.

Current outstanding problems include improved extraction of the timeline on which these entities exist; better methods to identify co-references among different mentions of the same entity; and detection of "cut and paste" duplication of data and of internal contradictions and inconsistencies. Overall, we also seek to improve the accuracy of the extraction techniques and to minimize the amount of required training data by using active learning, co-training, and bottom-up analysis of unlabeled data.

Predictive Modeling

As part of a collaborative project with professors Mark and Verghese in Electrical Engineering & Computer Science (EECS) and a number of doctors at Beth Israel Deaconess Medical Center, we have collected comprehensive data about 42,000 patient admissions to intensive care units (ICUs) at that hospital. The resulting MIMIC data set contains all laboratory test data, medications, nursing notes, specialty reports, and discharge summaries, as well as minute-by-minute summaries of data from ICU monitoring equipment. For about 4,500 patients, we have also collected detailed high-frequency waveform data from the monitors. Our efforts with these data have been to estimate, from the numerical data, how patient acuity (measured by probability of death within 30 days) varies as the patient's pathophysiologic state changes endogenously and through responses to therapies. In addition, we predict other useful shorter-term clinical endpoints, such as whether a patient is becoming hypotensive or septic, and whether it is possible to wean the patient successfully from a ventilator, intra-aortic balloon pump, or vasopressors. We also study the problems of glycemic control in intensive care. Our predictive algorithms achieve an area under the ROC curve of about 0.9 for mortality and 0.8 for other endpoints. We continue to try to improve these predictive models by developing more informative summary features from the patient data, by adding data derived from narrative text, and by doing bottom-up cluster analysis of patient states. We are developing models that track the temporal evolution of a patient's state, which we believe will lead to more accurate and useful predictive models. We are also very interested in testing our models prospectively; if they continue to do well, we would make their output available to treating clinicians to see if that information is considered helpful and (eventually) improves outcomes.

Translational Biomedical Informatics

It is a truism that a clinical phenotype is determined by some function of a patient's genetic endowment and the environment to which he or she has been exposed over a lifetime. As electronic records create rich data sets collected to support clinical care, we hypothesize that such data, with essentially zero marginal cost, can sometimes substitute for expensively collected clinical trial data. With our colleagues at the Informatics for Integrating Biology and the Bedside project at Partners Healthcare System, we are working on making this insight real for specific diseases and on creating improved methods to make future studies easier. For example, we are investigating the relationship between genes and autoantibody laboratory data to predict a patient's having rheumatoid arthritis (RA). We have helped to develop models based on numerical data and on data derived from natural language processing (NLP) to identify RA patients with very high (> 97 percent) specificity for participation in a gene-wide association study (GWAS), and are now working to discover associations among the genetic variables of those patients using Bayesian network methods. The selection techniques led to a highly pure population of RA patients, yielding a positive predictive value (PPV) of 94 percent for our selection criteria. Our subsequent GWAS, using about 1,500 controls and 1,800 cases, closely reproduced the genetic associations determined in a meta-analysis of numerous controlled trials. We also showed that genetic risk scores based on the number of minor alleles found in each patient's GWAS clearly formed distinct distributions in patients with positive anti-citrullinated protein antibody (ACPA) titers, compared with controls, in each of four ethnic subpopulations.

As next-generation sequencing techniques come online, we anticipate having ever-richer sources of genetic data that will permit increased refinement in such studies, and, we hope, help elucidate the genetic bases of disease.

Doctor-Patient Communication

We have been conducting a study to see whether the state of the art of speech understanding and NLP could allow a computer to listen to the conversation between a doctor and a patient, interpret the clinically significant components of that conversation, abstract them into a format typically used to record notes from such encounters, and create a draft of the final encounter note that memorializes the visit. There is extensive evidence that doctors often misremember or forget facts that they learn in a patient encounter, so this project aims to improve the quality of primary data. Our approach is to use a commercial speech-understanding tool (Dragon Naturally Speaking, or DNS) to transcribe both sides of an encounter to text; then to use NLP techniques to extract the relevant clinical content from those transcripts; then to organize those facts into a format typically used by our clinician colleagues to record their data; and then to create a draft encounter note from those intermediate forms. We work with the Pediatric Environmental Health Clinic at Children's Hospital Boston, where we have recorded more than 100 encounters.

We have developed effective techniques for extracting facts from the transcripts, but are severely limited by the low quality of the transcripts generated by the speech system. Contributing to these problems is the heterogeneity of patients' language skills and habits, accents, poor microphone use, and background noise in a busy clinic. Other weaknesses have been our inability to train the speech system specifically to each patient's voice, which seems to be a critical need for it to perform well. In addition, we hypothesize that the language models for dictation and conversation are quite different, and are studying the idea. DNS is optimized to the task of dictation, but we use it for interpreting a conversation. Based on preliminary analysis of language perplexity and word tuple frequencies, it does appear that the tasks are quite different, which may explain our poor results with DNS. It will be important to train new DNS models based on human-transcribed conversations (we will not have the resources to do this adequately in the current project).

Other Topics

In addition to the above active research projects, our group has also worked on issues of personally controlled health records (an idea we may have pioneered in 1994), technical support for maintaining confidentiality in patient records, and policy and technology considerations in the establishment of health information exchanges. Our earlier work focused on development of diagnostic and therapeutic expert systems, probabilistic inference, methods for computer programs to be able to explain their own reasoning, qualitative physiologic and probabilistic reasoning, and knowledge representation issues. This research is under the direction of professor Peter Szolovits.

Laboratory Sponsored Activities

CSAIL Outreach

Imara

The overall goal of Imara is to find and implement long-term, sustainable solutions to make educational technology and resources available to domestic and international communities. In concert with partners both inside and outside of the Institute, we hope our work will enable us to help bridge the digital divide—the perceived gap between those who have access to the latest information technologies and those who do not.

CommuniTech

CommuniTech is a domestic outreach initiative that attempts to heal the digital divide in the local community. The program provides economically disadvantaged adults with the tools they need to gain access to valuable information they can use to better their lives and the lives of their families.

CommuniTech focuses on both teaching basic computer technology skills and providing accessibility to computer hardware. Families Accessing Computer Technology (FACT) is a six-week course, taught by MIT student volunteers, that is designed to teach basic computer skills to underprivileged adults. Clients gain basic computer proficiency, skills to offer prospective employers, and a stronger connection to the world at large.

Clients who complete the class receive computers from FACT's partner program, the CSAIL Used Computer Factory (UCF). UCF refurbishes donated computers by installing new operating systems and productivity software. UCF also plays a small part in greening MIT by recycling unwanted equipment that might otherwise end up in landfills.

Lacotec Laare, Kenya

The Laare Community Technology Centre (LCTC) in Kenya was founded by Eric Mibuari ('06) in 2005 with the assistance of MIT's Public Service Center. The aim of this community center is to increase general computer awareness and literacy in the Laare community by providing cheap, local, and accessible training on the use of computers.

The LCTC particularly targets high-school-educated youth, seeking to equip them with basic computing skills that they can apply for personal use, in industry, and in education. So far, the LCTC has offered various levels of training to more than 500 students. The LCTC takes particular note of the economic difficulties of its potential trainees, and strives to charge the lowest feasible fees. Recently, it has expanded to include a primary school information technology educational initiative that targets elementary school students.

Middle East Education through Technology

CSAIL's support of the Middle East Education through Technology (MEET) program has continued over the past year. MEET is 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 host http://meet.mit.edu and to provide technical support to the MEET program.

TEK

Several CSAIL members have been supporting the Time Equals Knowledge (TEK) project, which empowers low-connectivity communities by providing a full internet experience using email as the transport mechanism.

Although the internet has revolutionized information delivery for most of us, for many communities in the developing world it remains an economic and technological challenge to access online resources. High charges for telephone service and internet access can quickly grow unaffordable, and low-bandwidth connections limit the amount of material that can be viewed per session. Because phone lines are limited, it is often difficult to time-share between internet and voice. Furthermore, unreliable network and power infrastructures can sometimes block internet access altogether.

Compared with direct web access, email can be much cheaper, more reliable, and more convenient in developing areas. The TEK client operates as a proxy on the user's machine, enabling users to browse downloaded pages using a standard web browser.

New searches are automatically encoded as emails and sent to the TEK server, which queries the web and returns the contents of resulting pages via email. TEK is free software distributed under the GNU lesser general public license.

OpenCourseWare Outreach Initiative

Adnan Esmail, a mechanical engineering graduate student, has combined MIT's OpenCourseWare (OCW) initiative with support from CSAIL to bring the Institute's educational resources to the Indian subcontinent. The OCW outreach initiative consists of mirror sites that make the course information available to those in nations with poor bandwidth infrastructure.

The program has been realized in partnership with Aligarh University in Aligarh, India, and Lahore University of Management Sciences, in the capital of Pakistan's Punjab province. It runs on hard drives that have been generously donated by Maxtor for OCW proliferation. This expansion of access gives talented students without the technological resources they need a chance to reach their full educational potential.

Seminar Series

Three speakers gave presentations during the 2011–2012 Dertouzos Distinguished Lecture Series. They were:

- Michael Kearns, University of Pennsylvania, "Experiments in Social Computation";
- Andrea Wong, Sony Pictures Television & Entertainment, "Electrical Engineering and the World of Entertainment";
- Leslie Valiant, Harvard University, "A Computational Theory of Cortex and Hippocampus."

Organizational Changes

In July 2011, Anant Agarwal became the 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. Chris Terman became the codirector and assisted the director with his duties.

CSAIL's new administration included an active executive cabinet which met twice per month to review and advise the director on policy, processes, and activities within the laboratory. Members of the executive cabinet included Regina Barzilay, Jack Costanza, Randy Davis, Srini Devadas, John Fisher, Piotr Indyk, John Guttag, Nancy Lynch, Sam Madden, Martin Rinard, Karen Shirer, and Bruce Tidor.

John Guttag continued in the role of "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.

Jack Costanza continued as the assistant director for infrastructure, overseeing information technology infrastructure and user support, building operations, and communications. Karen Shirer continued in her role as assistant director for administration, overseeing finance and human resources.

Elizabeth Bruce continued as director of industry partnership. She oversees the CSAIL Industry Affiliates Program, a corporate membership program that offers companies the opportunity to have access to CSAIL's faculty and students through annual conferences, recruiting events, and onsite visits.

Awards and Honors

Our faculty and staff won many awards this year, including the following:

- Scott Aaronson: National Science Foundation Alan T. Waterman Award
- Hal Abelson: Association for Computing Machinery (ACM) Karl V. Karlstrom Outstanding Educator Award; ACM Special Interest Group on Computer Science Education Outstanding Contributions to Computer Science Education Award
- Anant Agarwal: Included in *Scientific American,* "10 World Changing Ideas" (2011)
- Arvind Mithal: American Association for the Advancement of Science (AAAS) Research Fellow; Institute of Electrical and Electronics Engineers (IEEE) Computer Society Harry Goode Memorial Award
- Bonnie Berger: AAAS, Research Fellow; National Institutes of Health Director's Margaret Pittman Lecture for Outstanding Scientific Achievement and Lectureship
- Tim Berners-Lee: IEEE, Intelligent Systems Magazine's Hall of Fame Inductee
- David Clark: Oxford Internet Institute Lifetime Achievement Award
- Fernando Corbato: Computer History Museum, Research Fellow
- Constantinos Daskalakis: Microsoft Research, Faculty Fellow
- Jack Dennis: ACM SIGOPS Hall of Fame Award
- Alan Edelman: SIAM, Research Fellow
- David Gifford: ACM, Research Fellow
- Polina Golland: MICCAI Young Investigator Publication Impact Award
- Piotr Indyk: Included in *Technology Review*, "10 Emerging Technologies That Will Change the World" (2012)
- Frans Kaashoek: AAAS, Research Fellow
- Leslie Kaelbling: MIT MacVicar Faculty Fellow
- Dina Katabi: ACM SIGCOMM Best Paper Award; MIT EECS Faculty Research Innovation Fellowship Award

- Manolis Kellis: Athens Information Technology Center of Excellence for Research and Education Niki Award
- Charles Leiserson: ACM *Symposium on Parallelism in Algorithms and Architectures* Best Paper Award
- Barbara Liskov: Carnegie Mellon University and Tokyo University of Technology Katayanagi Award for Research Excellence; Invent Now, National Inventors Hall of Fame Inductee; National Academy of Sciences, Member
- Andrew Lo: Journal of Investment Management, Harry M. Markowitz Award; *Time* Magazine, 100 Most Influential People in the World (2012); MIT–Sloan Teacher of the Year
- Tomas Lozano-Perez: IEEE, Research Fellow
- Nancy Lynch: ACM Committee on Women, Athena Lecturer; Radcliffe Institute, Research Fellow
- Wojciech Matusik: Sloan Foundation, Research Fellow
- Una-May O'Reilly: EVOApplications Conference Best Paper Award
- Li-Shiuan Peh: ACM, Distinguished Scientist
- Nir Shavit: ACM and EATCS, Edsger W. Dijkstra Prize in Distributed Computing
- Peter Shor: AAAS, Research Fellow

Key Statistics for Academic Year 2012

The following are key statistics for the year:

- Faculty: 91 (16 percent women)
- Research staff: 46 (22 percent women)
- Administration, technical, and support staff: 68 (63 percent women)
- Postdoctoral researchers: 43 (10 percent women)
- Visitors: 126 (13 percent women)
- Paid Undergraduate Research Opportunities Program Participants: 104 (34 percent women)
- Master of engineering students: 30 (27 percent women)
- Graduate students: 333 (21 percent women)

Daniela Rus
Director
Professor of Computer Science and Engineering