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, computer graphics, and the explorations of 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: 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 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 cyber attacks, rethinking the field of artificial intelligence, and developing the next generation of robots. We are also investigating advanced software-based medical instrumentation and medical informatics systems to aid clinical decision-making. Advancements in biological research are also underway, 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. 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, Department of Defense Research and Engineering, 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 non-federal sponsors include Boeing, BMW of North America, LLC, Ford Motor Co, Foxconn Technology Group, Intel Corporation, Jaguar Land Rover Limited, Lockheed Martin Advanced Technology Laboratories, Microelectronics Advanced Research Corporation, Nissan Motor Co, Ltd., Nippon Telegraph and Telephone Corporation, Northrop Grumman Corporation, Ping An Technology, Qatar Computing Research Institute, Quanta Computer, Inc., Samsung Electronics, Siemens, and Wistron Corporation. Other organizations sponsoring research include Aarhus University, Battelle Memorial Institute, Delta Electronics Foundation, DSO National Laboratories, Epoch Foundation, Industrial Technology Research Institute, Nanyang Technical University, 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:

Quanta/Qmulus Project

In 2005, CSAIL started a ten-year, $45.5 million research collaboration called T-Party with Quanta Computer, Inc., Taiwan. The project was renamed Qmulus in 2010 to reflect the shift in our research emphasis from mobile computing to cloud computing. We describe below progress made during FY 2014–2015, the tenth year of this collaboration, in the four broad areas of research.

Cloud Technologies. The general emphasis of this line of research is in hardware and software security. We are building a secure processor that can be used to maintain privacy in big data applications through the use of Oblivious Random Access Memory, or ORAM, so that memory access patterns do not leak private information (Edwin Sibley Webster Professor of Electrical Engineering and Computer Science Srini Devadas). In another hardware-related project, we developed a first-generation prototype database machine called BlueDBM, a distributed flash-based storage system with hardware acceleration (Johnson Professor of Computer Science and Engineering Arvind Mithal). In software, our research focuses on three aspects: cloud security, multicore, and cloud infrastructure (Charles Piper Professor M. Frans Kaashoek, Assistant Professor Robert Morris, and Associate Professor Nickolai Zeldovich). Finally, we focused on a number of database-related applications that might be a good target for the BlueDBM system (Professor Samuel Madden). Demonstrations of our progress for healthcare-related applications were made in May 2015.

Multimedia. We continue to develop/improve human language technologies for natural human computer interactions and scientific discovery, focusing on health- and wellness-related applications such as logging food intake, interactively exploring an online drug side effect database, and analyzing research literature in order to shed light on the

relationship between wellness and environmental toxins like Glyphosate in genetically modified foods (Drs. James Glass and Stephanie Seneff). We demonstrate low-power chip implementation for speech recognition and for voice activity detection (Professor Anantha Chandrakarsen and Dr. Glass). Finally, we continue to improve our multi-view autostereoscopic systems capable of demonstrating 3D content viewing without the need for users to wear glasses (Associate Professor Wojciech Matusik).

Healthcare. We improved our method of using computational photography and motion magnification for medical applications by introducing a phase-based video processing technique, and by introducing a real-time improvement using a Riesz pyramid extension (Professor Frédo Durand and Thomas and Gerd Perkins Professor of Electrical Engineering and Computer Science William T. Freeman). In cloud-based medicine and healthcare, we continue to develop new tools for visualizing large datasets to support interactive knowledge discovery, and analyzed large datasets for risk stratification (Dugald C. Jackson Professor John Guttag).

Education Technology. Regarding flexible content creation and viewing, we continue to improve Pentimento's interfaces, and we used it in several MIT courses (Professor Durand). We continue to use crowd computing techniques for education, and we demonstrated the concept of "wait learning," which is the use of downtime throughout the day (e.g., waiting for the elevator to arrive, or waiting for computation to complete) to learn a second language (Professor Rob Miller). We demonstrated, through crow-sourcing techniques, that linking various courseware—lectures, slides, textbooks—can lead to better learning in Massive Open Online Course (MOOC) platforms (Delta Electronics Professor of Electrical Engineering Victor Zue). We continue to develop tools to implement the "see one, do one, teach one" pedagogy for online teaching of 6.004x (Professor Steve Ward and Senior Lecturer Dr. Chris Terman).

Qatar Computing Research Institute

In 2012, CSAIL started a seven-year, $35 million research collaboration with the Qatar Computing Research Institute (QCRI) to collaborate on a wide-range of research topics in computer science. Here is a summary of progress made during FY2015 in four broad areas of research.

Cyber Security. The research challenge we address is that of securing computing infrastructure against a broad class of cyber attacks. Our objective is to develop new techniques that can remove many of the vulnerabilities that attackers exploit and that can predict and intercept new (zero-day) attacks that exploit previously unknown vulnerabilities. These objectives are realized through a number of sub-projects: systems that are much more difficult to penetrate, systems that can work through penetrations, and systems that can recover quickly. For example, we are developing mechanisms that monitor a program's execution in order to detect misbehaviors indicative of attacks and to project forward how the attack will evolve (Associate Director and Principal Research Scientist Dr. Howard Shrobe), developing automated techniques for analyzing vulnerability of cyber-physical systems, and synthesizing defense software (Assistant Professor Armando Solar-Lezama), building a strong foundation of cryptographic mechanisms that help us compute on encrypted and authenticated data and programs

(Cryptographer and Associate Professor Vinod Vaikuntanathan), studying the formal specification of reactive programs, with the aim of carrying out rigorous, machine-checked mathematical proofs that satisfy important security properties (Assistant Professor Adam Chlipala), and implementing constructs called security contexts that allow for the user-level implementation of information flow control schemes (Professor Devadas).

Arabic Speech and Language Development. A number of sub-projects are ongoing in human language technologies, including speech recognition error reduction based on diacritized lexical representations (Dr. Glass), improved dialect recognition using Deep Neural Networks (Dr. Glass), development of a sampling-based parser that can easily handle arbitrary global features (Professor Regina Barzilay), and improved semantic parsing for understanding by means of re-ranking models (Drs. Glass and Seneff).

Advanced Analytics and Visualization in Sports. Projects in this area include: video magnification in the presence of large confounding motion (Professors Durand and Freeman), backward-compatible masks for autostereoscopic screens (Associate Professor Matusik), and the use of machine learning techniques to transform soccer event data into features that facilitate such things as learning which sequences of actions on a soccer pitch increase the short-term likelihood of shot on goal (Professor Guttag).

Data Management for Social Computing. There are two research thrusts in this effort. The Humanitarian Technologies team collaborated with the International Community of the Red Cross (ICRC) on their Restoring Family Link app, and developed a new website and tutorials for the framework (Principal Research Scientists Dr. Lalana Kagal). The Social Web Architecture thrust addresses an important and fundamental question related to a decentralized architecture for social networks: How can users get the full benefits of social software even when their friends and colleagues use different vendor's software and different social network providers? (Founders Professor of Engineering Timothy J. Berners-Lee).

We began to extend our data integration framework, starting with interactive data curation. It has since expanded to include various aspects of data management, including elasticity in transactional database management systems (DBMS) and graph processing (Adjunct Professor Michael Stonebraker and Professor Samuel Madden).

Foxconn/Basic Research in Interactive Computing

The Basic Research in Interactive Computing project is a six-year, $5.7 million research collaboration sponsored by Hon Hai/Foxconn Technology Group. This research collaboration completed phase two and began phase three in January 2014. The project is currently funding six major research areas in several areas of computer science, ranging from networking, to human-computer interactions, computer graphics and vision, and theory. The 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 the task of building and using these systems.

During the current year, the collaboration supported five principal investigators in the following research areas:

- Advanced Hand-Tracking and Gesture-Based Interaction (Professor Randall Davis)
- Reducing the Energy Consumption of Cellular Wireless Interfaces (Fujitsu Chair Professor Hari Balakrishnan)
- Bringing Software Defined Networks to Wireless LANs (Andrew and Erna Viterbi Professor of Electrical Engineering and Computer Science Dina Katabi)
- Multicore Software Technology for Video-Conferencing Systems (Professor Charles Leiserson)
- iDiary: Activity Modeling using Audio, Video, and GPS Streams from Phones (Andrew and Erna Viterbi Professor Daniela Rus)

A total of eight Foxconn engineers visited CSAIL during this academic year to facilitate technology transfer and to receive training for advanced research and development in computer science. Each engineer spent six months at CSAIL. Additionally in December 2014, video presentations for the four projects were conducted via video conferencing from CSAIL to Foxconn in Taiwan, summarizing the research results.

InternetPolicy@MIT: Enhancing Public Trust in Our Connected World

In March 2015, MIT launched the new Initiative on Cybersecurity and Internet Policy Research. The goal of the initiative is to create a new field that will help governments and other responsible institutions create public policy frameworks that will increase the trustworthiness of the interconnected digital systems we live in today. The Initiative will reach beyond the confines of what is currently understood by the term "cybersecurity." The field we are creating seeks to understand the interconnected digital systems that will be the foundation of the future flourishing of our societies and on which we are already inexorably dependent. Made up of faculty from across campus and housed in CSAIL, InternetPolicy@MIT has already produced important research results contributing to current debates on the security of new electronic surveillance proposals and launched several efforts to expand the curriculum at MIT.

InternetPolicy@MIT has already begun making a technical contribution to the cybersecurity policy debate. Working with leading security researchers at MIT and colleagues from around the world, we produced a paper ("Keys Under Doormats: Mandating insecurity by requiring government access to all data and communications") analyzing security risks of new wiretapping proposals put forth by law enforcement agencies in the United States and the United Kingdom. This paper was widely cited at several legislative hearings in the US Senate and widely reported in the world press.

As part of the Initiative's mission to train a new generation of technology policy leaders, we are experimenting with a variety of new educational programs. In Spring 2015, initiative leaders Professor Hal Abelson and Daniel Weitzner, Esq., ran an experimental course jointly with Georgetown Law School on privacy technology and legislation. Working with six MIT Course 6 students and 12 law students, we gave a high-intensity

introduction to privacy law and associated computer systems design questions. As a final project, mixed groups of law and EECS students worked together to develop legislative proposals to address specific privacy challenges. Students learned how to make real technology policy contributions. The resulting projects were presented at a seminar in Washington DC to legislators, privacy regulators and advocates.

Intel Science and Technology Center in Big Data

Professors Samuel Madden and Michael Stonebraker continue to head the 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, MIT, Portland State, University of Tennessee, University of California Santa Barbara, and University of Washington) on 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 scale, rate, or sophistication of data processing that existing systems provide. The center is developing new technologies in a number of areas, including: data-intensive scalable computing, machine learning, and computer architecture and domain sciences (genomics, medicine, oceanography, imaging, and remote sensing).

This year, the center was renewed for an additional two years, until September 2017. Over the course of the next two years, the center will focus on a "capstone" project centered around applying big data technologies to medical data processing, including imagery, text, genomics, medical records, labs, and physiological signals. The centerpiece of this capstone project will be a new software architecture for combining multiple data processing systems (each suited to a particular type of data, such as text), together in a new platform called BigDawg.

DARPA Robotics Challenge

Announced in March 2012, the DARPA Robotics Challenge (DRC) was inspired by the Fukushima nuclear disaster and other situations in which human-like mobility and dexterity is required, but which are too hazardous for humans to physically expose themselves to. DARPA aimed to accelerate the development of remotely-operated humanoid robots that can walk and climb over complex terrain within disaster zones (rubble, stairs, ladders), and perform useful tasks that require handling objects (levers, pumps, valves), and using tools (wrenches, drills, saws), and even climbing into and driving ordinary vehicles. In order to encourage autonomy, communications between the robot and the human operator were degraded by a high-latency, low-throughput network link.

CSAIL formed Team MIT in the spring of 2012, and during 2013 finished third then fourth consecutively in the first two rounds of competition, earning the right to compete in the final competition. The final competition took place in June 2015. In this competition the robots had to be power-autonomous, had to operate without a safety belay, and tasks had to be completed in a single run roughly eight times faster than in the previous trials. The MIT student team worked around the clock to field our most advanced research algorithms on this complex humanoid, and arrived at the finals with a system capable of accomplishing all the tasks and with the robot capable of performing

many of the tasks without any input from the human. In the end, the competition tasks and network configuration were too simple—many teams were able to accomplish all of the tasks, often using relatively simple teleoperation-based interfaces. Team MIT performed extremely well. We made a small human operator error on the first day that caused the robot to fall and break its arm, but we were able to showcase our advanced software capabilities by having the robot continue on and perform nearly all of the remaining tasks left-handed. Ultimately the fall prevented us from winning the competition, but we still finished in sixth place and won numerous awards for research and software contributions to the field.

This competition earned CSAIL and MIT recognition as one of the premier labs in the world working on humanoid robots, has produced impressive research results, and has produced incredible software and hardware tools to bootstrap our future research.

World Wide Web Consortium

The World Wide Web Consortium (W3C) was founded at MIT in 1994 by the inventor of the web, Professor 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. Approximately 400 member organizations, including most of the world's leading technology companies, are working to enhance the capabilities of web documents and create an Open Web Platform for application development, available across a wide range of devices, enabling everyone on the planet to collaborate and share data and information.

In recent years, a great many factors (people, devices, bandwidth, policy decisions, etc.) have extended the reach of the web in society. Video, social networking tools, user-generated content, location-based services, and web access from mobile devices are transforming many industries, including mobile, television, publishing, automotive, entertainment, 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.

Core technology focus

W3C standards define an Open Web Platform for application development that has the unprecedented potential to enable developers to build rich interactive experiences, powered by vast data stores that are available on any device. Although the boundaries of the platform continue to evolve, industry leaders speak nearly in unison about how HTML5 will be the cornerstone for this platform. A major milestone was achieved last year with the completion of HTML5. But the full strength of the platform relies on many more technologies that the W3C and its partners are using, 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.

With the completion of HTML5, the web community is also looking at raising the level of abstraction of Web interfaces. Whereas the web began as an information-sharing

mechanism based on simple markup, today's web requires substantial support in a variety of advanced areas. W3C has launched a focus on Application Foundations: eight foundational elements of future development of the web platform, including such foundations as: Media, Performance, Security and Privacy, and Accessibility.

The demand is also driving W3C to expand its agenda and the size of its community. W3C launched Community and Business Groups in 2011. In less than three years after launch, over 4,000 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 pre-standards and standards track work.

Industry impact and broadening the set of participants

In recent years, web technology is not only used by consumers and companies for information sharing, but increasingly the web is the delivery mechanism for companies to deliver their services. Examples of that include telecommunications (where web access is a key service), entertainment (which is increasingly delivered over the web), publishing (EPUB3, the e-book standard builds on the Open Web Platform), and retail and financial services (both impacted by an increase of payments on the web). This has resulted in a diversification of W3C membership, and also has enriched the technical agenda to address new issues that arise.

Centers and Initiatives

CSAIL Alliance Program

The CSAIL Alliance Program (CAP) is a membership-based program that serves as a portal for industry, governmental organizations, and other institutions to engage with CSAIL. CAP's mission is to provide a proactive and comprehensive approach to developing strong connections with all CSAIL has to offer. CAP has 52 members representing a variety of industry sectors and regions, including North America, Asia, South America, and Europe. Key benefits of the program include:

- Keeping abreast of the latest research at CSAIL
- Identifying opportunities to foster and explore new research collaborations, and
- Engaging with CSAIL students for recruiting

Members of CAP are given access to new technologies and ideas as they move from laboratory to marketplace. Throughout the year, CAP hosts events, seminars, meetings, and activities to help connect member companies with the lab. To facilitate collaboration and communication with industry partners, CAP maintains a members-only website, produces research briefings, and publishes a CSAIL Student Profile Book.

CAP's flagship event is the Annual Meeting held in May each year. The event is a two-day conference that showcases CSAIL's research and students as well as serves as a forum for members to network, connect, and learn from each other. Over 200 people from 40 member companies attended the 2015 CAP Annual Meeting. All activities are designed to help members to connect with CSAIL researchers to create value for their organizations.

bigdata@CSAIL

The goal of the MIT Big Data Initiative at CSAIL, a multi-year effort launched in May 2012, is to identify and develop new technologies needed to solve the next-generation data challenges which will require the ability to scale well beyond what today's computing platforms, algorithms, and methods can provide. We want to enable people to truly leverage big data by developing tools and platforms that are reusable, scalable, and easy to deploy across multiple application domains.

Our approach includes two important aspects. First, we work closely with industry and government to provide real-world applications and drive impact. Promoting in-depth interactions between academic researchers, industry, and government is a key goal. Second, we believe the solution to big data is fundamentally multidisciplinary. 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, medicine, smart infrastructure, education, and science. The initiative focuses on four broad research themes: computational platforms; scalable algorithms; privacy and security; and big data applications.

The Big Data Initiative is supported by a group of sponsor companies, including: AIG, Alior Bank, British Telecommunications (BT), EMC, Facebook, Huawei, Intel, Microsoft, Quanta, Samsung, SAP, Shell, and Thomson Reuters.

CyberSecurity@CSAIL

The goal of CyberSecurity@CSAIL is to identify and develop technologies to address the most significant security issues confronting organizations in the next decade. This initiative was launched in March 2015 with five founding partners: Raytheon, BBVA/ Compass Bank, BP, Boeing, and BAE Systems.

Presently, approaches to system security do not give overall security guarantees, but rather attacks are fought individually—"patch and pray" style. CyberSecurity@CSAIL aims to provide an integrated and formal approach to the security of systems, combining design and analysis methods from cryptography, software, and hardware. Our approach includes three key elements: close collaboration with industry for input to shape real-world applications and drive impact; leverage of the breadth and depth of CSAIL security researchers to approach the problem from a multi-disciplinary perspective; and creation of a test-bed for our industry partners to implement and test our tools as well as have our researchers test tools developed by our partners.

The Cybersecurity@CSAIL Lecture Series was launched in April 2015 with a talk by Professor Ben-Israel, chairman of Israel's Space Agency and National R&D Council, on Israel's Triangular Cyber Ecosystem. Additionally, Dr. Howard Shrobe, director of Cybersecurity@CSAIL, has engaged the MIT community by presenting talks to MIT alumni in both Washington, DC and San Francisco.

MIT Center for Wireless Networks and Mobile Computing

The MIT Center for Wireless Networks and Mobile Computing (Wireless@MIT) is an interdisciplinary center created to develop the next generation of wireless networks and mobile devices. Headquartered at CSAIL, the Center creates a focal point for wireless research at MIT and addresses some of the most important challenges facing the wireless and mobile computing fields. The goal of Wireless@MIT is to develop next-generation wireless network technologies and mobile systems. The features of the Center are:

- Access to research and connections with MIT professors and their groups
- Early access to the latest emerging wireless technologies
- A seat at the table on influencing and impacting policies and next-generation technologies, providing neutral ground where companies, academics, and government representatives can discuss the future of the wireless industry and mobile applications.

Over the past year, the center has contributed multiple innovations that enable high network performance and introduce new architectures and services for the Internet of Things, mobile health, low-energy sensors, big data communication and processing, and innovative radio hardware. These contributions came from multiple faculty members and their students, including Professors Dina Katabi, Hari Balkrishnan, Li-Shiuan Peh, Arvind Mithal, and others. The center's work includes contributions from more than 50 MIT faculty, research staff, and graduate students. A key feature of the center is that it encourages collaborations between MIT faculty in various departments and labs. In particular, the center has faculty members from EECS, Mechanical Engineering, and the Sloan School of Management. Its members also span multiple labs including CSAIL, LIDS, MTL, and RLE. The center has eight members, Amazon.com, Cisco, Intel, Mediatek, Microsoft, ST Microelectronics, Telefonica, and Google.

entrepreneurship@CSAIL

During the past year, CSAIL continued its entrepreneurship@CSAIL initiative to help transition our research to start-up-based commercialization while enabling CSAIL students to pursue start-up ideas of their own design. Following an initial successful offering, 6.S078 Entrepreneurship Project was taught in spring 2015 through the EECS Department with 12 registered students and the participation of two local venture capital firms. The students worked on nine projects and completed prototype implementations. Projects included a travel planning web app, an iPhone events app, grocery delivery service implementation, text message communication for African countries, a web interface to existing telephone customer service call centers, automated big data analysis from databases using machine learning, a machine learning system for loan application evaluation, an interactive robotic toy with cloud connectivity, and a 3D printer that forms clay.

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.

The BigDAWG Polystore System

There are two technological factors that drive the need for federations of multiple data storage engines. First, there is an increased interest in disparate data. In the 1990s, enterprises were singularly focused on structured data. Today, there is interest in integrating structured data with text, web pages, semi-structured data, time series, etc. Second, it is now well understood that "one size does not fit all." Specifically, there is no database management system (DBMS) that offers high performance on all of the kinds of data noted above. Hence, it makes sense to deploy applications on multiple DBMS. Obviously, one does not want to force users to learn disparate query languages, so a software layer that supports a simple user-facing abstraction and translates this notation into the various underlying query languages seems like a good idea. Since there is no agreement on a single user-level Esperanto, we conjecture that several such notations will be required, and we define a polystore to be such a software system.

During the last year, we worked out a specific architecture for a polystore system, which we call BigDawg. Moreover, we built a reference implementation of BigDawg that unifies five underlying storage engines and supports novel visualization and browsing interfaces. In addition, we targeted a specific complex medical data set to test out our ideas.

Furthermore, we designed an optimization framework, so BigDawg can be as efficient as possible. Lastly, we prototyped a monitoring framework, so that data objects can be moved between underlying storage systems to balance resource consumption.

This research is under the direction of Professors Michael Stonebraker and Samuel Madden. BigDawg is supported by the Intel Corporation through its Science and Technology Center on big data. There are five universities involved, led by MIT, and approximately 25 students, postdocs, and faculty have contributed to this effort.

Geometric Folding (Computational Oragami)

CSAIL is the leading center for research in the mathematics and algorithms underlying geometric folding—reconfiguring, bending, and creasing of sheets, linkages, robots, and other structures. Professor Erik Demaine and his father Martin Demaine cofounded the area in the late 1990s, and appear in the PBS documentary about origami science and sculpture *Between The Folds*. Erik co-wrote the main book in the field (*Geometric Folding Algorithms*), and teaches the main class on the topic, 6.849 Geometric Folding Algorithms: Linkages, Origami, Polyhedra. Research on geometric folding algorithms has blossomed over the past two decades, including the development of an entire NSF program (Origami Design for Integration of Self-assembling Systems for Engineering Innovation) focused on origami-based engineering. Geometric folding has extensive applications throughout science and engineering, from manufacturing to medicine to robotics. The Demaine Group has consistently co-solved some of the biggest open problems in the field, including:

1. **Carpenter's Rule Problem:** Proved that any 2D chain linkage, with rigid bars and fixed connections, can be universally reconfigured without self-intersection.

2. **Origami Universality:** Proved that any 3D surface can be folded as origami from a large enough sheet of paper, continuously without self-intersection.
3. **Origami Magic:** Proved that any 2D graph drawn on a piece of paper can be aligned by folding the paper flat. One cut then produces exactly the desired cut pattern.
4. **Origamizer Algorithm:** Universal, practical, water-tight foldings of arbitrary 3D surfaces. Experiments confirm feasibility for both paper and sheet metal folding.
5. **Unfolding Polyhedra:** Algorithm for efficiently unfolding the largest solved class of 3D surfaces (orthogonal polyhedra) into 2D "nets."
6. **Hinged Dissections:** Universal algorithm for designing a chain of blocks that can fold into any desired finite set of 2D shapes of equal area, and 3D shapes of equal area and Dehn invariant.
7. **Hinged Robots:** Practical hinged chain patterns that can fold into exponentially many shapes—any shape made out of unit cubes—continuously without self-intersection.
8. **Protein Folding:** Developing geometrically natural algorithms for folding fixed-angle chains, which model protein backbones, providing a hypothesis for how nature folds such structures.
9. **Pleat Folding:** Proved main results on existence and nonexistence of pleat-folded "self-folding" structures, and simulated the equilibrium that results.
10. **Self-Assembly:** Developed structures and algorithms for actuated self-folding of sheets and nets into desired 3D structures, in collaboration with experts in shape-memory alloys and polymers.

As outreach, Erik and Martin Demaine give regular lectures to high school students and exhibit their sculpture that expresses geometric folding to the general public. Notable museums that have exhibited their sculpture includes the Museum of Modern Art (MoMA) and the Renwick Gallery of the Smithsonian American Art Museum. Currently their sculpture is exhibited in about 40 galleries and museums per year.

Spoken Language Understanding for Logging Food Intake

In collaboration with nutritionists at Tufts University, we have created a prototype spoken language understanding system that lets people track the energy and nutritional content of their food intake.

We believe the convenience of speech could boost existing approaches to tracking nutrition, providing a powerful tool for combating health problems such as obesity, particularly for hard-to-reach populations with low literacy.

We are studying whether speech can lower user workload compared to existing self-assessment methods, whether we can accurately quantify caloric and nutrient absorption from spoken descriptions of meals, and whether food properties can efficiently and effectively be ascertained using spoken dialogue-enhanced user interfaces.

Over the past year we have created a prototype speech interface that accesses nutritional database entries mentioned in spoken meal descriptions. The system includes both a web and a touch interface connected to an application server. The language understanding component uses machine learning to automatically process spoken meal descriptions. Our current best methods use a conditional random field model to semantically label the important food concepts (e.g., food, brand, quantity, etc.), and a random forest classifier to associate all concepts belonging to the same food item (e.g., "a cup of skim milk"). On a dataset of 10,000 user-generated meal descriptions, we achieve an average F1 test score of 90.7 for semantic labeling and 96.2 for associating foods with properties (86.3 on predicted labels). The application server includes a component that translates semantically processed food descriptions into nutritional database queries. This component also converts spoken food quantities to the units used for the food in the database.

When the overall system was assessed on a set of 437 meal descriptions containing 975 food concepts, 83% of the semantic tags, 78% of the quantities, and 71% of the USDA hits were judged to be correct. Since the system enables users to correct system understanding by speech or via the interface, we believe these initial results indicate a significant step towards the intelligent conversion of spoken food diaries to actual nutritional feedback.

This research is a collaboration of MIT students Mandy Korpusik, Rachael Naphtal, Patricia Saylor, Calvin Huang, Eann Tuan, and Mari Kobiashvili, Research Scientist Scott Cyphers, and Senior Research Scientist James Glass.

Building Web Applications on Top of Encrypted Data

Using a web application for confidential data requires the user to trust the server to protect the data from unauthorized disclosures. This trust is often misplaced, however, because there are many ways in which confidential data could leak from a server. For example, attackers could exploit a vulnerability in the server software to break in, a curious administrator could peek at the data on the server, or the server operator may be compelled to disclose data by law. How can developers build web applications that protect data confidentiality against attackers with full access to servers?

To address this problem, a group of researchers including students Raluca Ada Popa, Emily Stark, Jonas Helfer, and Steven Valdez, and Professors Hari Balakrishnan, M. Frans Kaashoek, and Nickolai Zeldovich, developed Mylar. Mylar is a new platform for building web applications that stores sensitive data encrypted on the server. The keys that can decrypt the data are stored in some users' web browsers and the data gets decrypted only in those browsers. Even if an attacker fully compromises the server, the attacker gets access to only encrypted data and does not have the necessary decryption keys. Mylar achieves this organization through a new data sharing mechanism, practical ways of computing on encrypted data at the server, and a mechanism for verifying that the application code was not tampered with by a compromised server.

Mylar enables many classes of applications to protect confidential data from compromised servers in a practical way. It leverages the recent shift in web application frameworks towards implementing logic in client-side Javascript code, and sending

data, rather than HTML, over the network, such as in the Meteor framework. This enables a web browser to run most of the application code over plaintext data, on the user's own machine, while ensuring that the server stores only encrypted data. The server, in turn, performs only a limited set of operations on the encrypted data, such as keyword search, which would be impractical to do on the user's machine.

Experimental results show that Mylar imposes modest performance overheads, and requires little developer effort. For example, porting six applications to protect confidential data using Mylar required changing an average of just 36 lines of code per application. The applications included: a system for endometriosis patients, a website for managing homework and grades, a chat application, a forum, a calendar, and a photo-sharing application. The endometriosis application, used to collect data from patients with that medical condition, is in the process of being deployed in collaboration with Newton-Wellesley Hospital.

New Applications of Program Synthesis

All modern computing applications are built from painstakingly handcrafted source code. The difficulty of writing code and the scarcity of people skilled enough to produce it are two of the major limitations to the growth of the computing industry. Our Computer-Aided Programming Group has developed new technologies for software synthesis that can help automate challenging aspects of programming and can even be leveraged to help train new programmers.

One of the main insights behind this work is that while synthesizing large amounts of code from scratch is extremely challenging, synthesizing small code fragments in the context of a larger program is feasible and can already enable a number of important applications. For example, the most straightforward application is to help programmers in cases where some small part of the code is particularly difficult to write and would benefit significantly from automation. For example, we showed that when writing code for super-computers with large-scale parallelism, this approach can help the developer synthesize many of the details of data movement routines that are difficult to get right.

A different application of this form of synthesis involves discovering the high-level intent behind a block of low-level code. For example, we discovered that in data-intensive applications, it is common for programmers to write code for functionality that has already been implemented more efficiently in the database. We have shown that synthesis can help discover the high-level database commands or queries that capture the intent of the low-level code that was written by the programmer. By replacing the code with these high-level queries, it is possible to achieve significant performance improvements over the original code, saving the programmer hours of manual optimization.

Our research has shown that synthesis can even help in teaching new programmers. Specifically, synthesis can assist with the labor-intensive task of providing feedback to students for their programming assignments. Given a student solution to a programming problem containing a small error, synthesis can help discover a new program that is correct with respect to the instructor's solution and that is as similar as possible to the

student's solution. By comparing the differences between these two programs, it is possible to explain to the student exactly what was wrong with his or her solution.

This research has been a collaborative effort between PhD students Rishabh Singh, Alvin Cheung, and Zhilei Xu, with Professor Armando Solar-Lezama and Research Scientist Shoaib Kamil.

Representation by DNN

With the success of new computational architectures for visual processing, such as deep neural networks (DNN) and access to image databases with millions of labeled examples (e.g., ImageNet), the state of the art in computer vision is advancing rapidly. Computer vision is now present among many commercial products, such as digital cameras, web applications, robotics, autonomous driving, and security applications. One important factor for continued progress is to understand the representations that are learned by the inner layers of these deep architectures.

We introduced a new large database of annotated images for scene recognition called the Places database. Building this database is an ongoing multiyear effort and has become an important resource for the computer vision community. Using this database we trained a DNN to do scene classification (e.g., given a picture, the neural network classifies the picture as being an office, a bedroom, or a kitchen—and 200 more possible scene categories). The performances achieved by the DNN are remarkable and constitute the state of the art of scene classification. But what is the nature of the internal representation learned by the network? Why does it work so well? Deep neural networks generally are used as black boxes able to learn complex functions. When they are trained, both the input and output are clearly specified and well understood. But the internal representation built by the network is rarely explored.

To uncover the representation built by the network, which consists of thousands of units, we used crowdsourcing. For each unit we selected the 60 images (out of a set of 200K) that activated each unit more strongly and then we asked non-expert participants to provide a name describing what concept was common among the 60 images. By visualizing stimuli that activate a unit, we have shown that even naïve users have good intuitions about the underlying semantic concept that activates a unit. The results showed that the representation built by the network was interpretable and that it was possible to attach semantics to more than 50% of the units on each layer of the neural network.

The most surprising result was to show that the last layers of the network were building a representation that relied on objects in order to classify scenes. For instance, to decide if a picture corresponded to a bedroom, the network was detecting beds. Indeed, reliable object detectors (people, cars, tables, lamps, etc.) are found inside the network trained to recognize scenes even though no supervision was provided for objects. As scenes are composed of objects, the network for scene classification automatically discovers meaningful object detectors, representative of the learned scene categories. This is remarkable because, despite the fact that we never provided training for objects, the network was learning to localize objects in images.

This research is a collaboration of PhD students Bolei Zhou and Aditya Khosla, Research Scientist Aude Oliva, Visiting Professor Agata Lapedriza, and Professor Antonio Torralba.

Laboratory-Sponsored Activities

CSAIL Outreach

CSAIL's Hour of Code. CSAIL hosted its first annual demo day in conjunction with the global Hour of Code movement, inviting local students to learn more about a wide array of computer science research. Approximately 200 STEM students from schools in the greater Boston area were in attendance.

Reddit "Ask Me Anything." CSAIL encouraged the wider online community to submit questions about computer science and academia in a series of Reddit "Ask Me Anything" (AMA) sessions involving CSAIL researchers. In sum, the AMAs spurred approximately 5,000 comments and questions, as well as more than 160,000 page views.

Middle East Education through Tech. 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 gave presentations during the 2014–2015 Dertouzos Distinguished Lecture Series: Leslie Lamport, Microsoft ("Who Builds a Skyscraper without Drawing Blueprints?") and Geoff Hinton, University of Toronto ("Deep Learning").

CSAIL Research Highlights Competition

On May 18, 2015, CSAIL hosted a Research Highlight Competition for CSAIL graduate students and postdoctoral associates. This event allowed students to create and present refined, accessible demonstrations of their research. Members of the CSAIL community voted on which presentations were the most successful. Winners were then invited to make demonstration videos to add to the CSAIL archive and for use in CSAIL outreach and engagement.

Organizational Changes

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

CSAIL's leadership team includes two associate directors and the 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 director for the 2014-2015 term in October 2014.

Additionally, the CSAIL executive cabinet meets twice per month to review and advise the director on policy, processes, and activities within the laboratory. Members named to the executive cabinet included: Saman Amarsinghe, Regina Barzilay, Jack Costanza, Randall Davis, Srini Devadas, John Fisher, David Gifford, Polina Golladn, Daniel Jackson, Martin Rinard, Karen Shirer, Howard Shrobe, Nickolai Zeldovich, and Victor Zue.

Victor Zue continued in his role of director of international relations, managing the engagment and oversight of various important CSAIL international contracts and international contract negotations. 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. Lori Glover continued in her role as managing director of the CSAIL Alliance Program (CAP).

Srini Devadas remained the 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 the two assistant directors, Jack Costanza and Karen Shirer, and an administrative assistant, Sonya Kovacic.

Awards and Honors

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

Arvind: National Academy of Sciences, India, Foreign Fellow

Hari Balakrishnan: National Academy of Engineering, Member

Bonnie Berger: École Polytechnique Fédérale de Laussane, Honorary Doctorate

Erik Demaine: European Association for Theoretical Computer Science, Nerode Prize

Srini Devadas: ACM, Fellow, IEEE-CEDA, A. Richard Newton Technical Impact Award

Jim Glass, International Speech Communication Association, Fellow

Eric Grimson, ACM, Fellow

Daniel Jackson, ACM SIGSOFT, Distinguished Paper

Charles Leiserson, ACM/IEEE, Ken Kennedy Award, SIAM, Fellow

Robert Morris, ACM, Fellow

Una-May O'Reilly, International Association for Artificial Intelligence and Law, Peter Jackson Best Innovative Application Paper Award, Genetic and Evolutionary Computation Conference, Best Paper in Genetic Programming Track

Li Shiuan Peh, ITS World Congress, Best Paper

Ronitt Rubinfeld, ACM, Fellow

Daniela Rus, IEEE/RSJ International Conference on Intelligent Robots and Systems, Best Paper, Alianta Gala Award, ACM, Fellow, National Academy of Engineering, Member

Julie Shah, Technology Review, Top Innovators under 35, ACM/IEEE International Conference on Human Robot Interaction, Best Paper, NSF, Career Award, HFES, Best Paper Finalist

Michael Stonebraker, ACM, A.M. Turing Award, IEEE ICDE 10-Year Most Influential Paper Award

Madhu Sudan, Infosys, Prize for Mathematical Sciences

Vinod Vaikuntanathan, Microsoft Research, Faculty Fellow

Matei Zaharia, ACM, Doctoral Dissertation Award

Key Statistics for Academic Year 2015

Faculty: 108 (20% women)

Research staff: 31 (16% women)

Administration, technical, and support staff: 81 (68% women)

Postdocs: 84 (10% women)

Visitors: 91 (19% women)

Paid Undergraduate Research Opportunities Program Participants: 156 (37% women)

Master of engineering students: 72 (35% women)

Graduate students: 403 (18% women)

Daniela Rus
Director, Computer Science and Artificial Intelligence Laboratory