Computer Science and Artificial Intelligence Laboratory

The Computer Science and Artificial Intelligence Laboratory (CSAIL) takes a computational approach 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, and explore human–machine interactions.

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, 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 innovation that has affected how people interact and do business. Previous innovations from laboratory members include time-shared computing, public key encryption, computer chess, web standards, GNU, TCP/IP, and ARPANet. Current research explores mobile computing, the next generation of laptops, and the application of sensor technology to traffic congestion, animal herding, medical monitoring, and climate observations. Robotic locomotion and human robotic interface are under investigation as well as medical solutions with image-guided surgery and technical applications to aid clinical decisions. Advancements in biological research are also under way, including developments in the field of computational biology and the application of machine learning to the interpretation of complete genomes and understanding gene regulation.

CSAIL research is sponsored by a large number of diverse sources, from US government contracts to the private sector. US government sponsors include the National Science Foundation (NSF), Air Force (including Air Force Research Laboratory and Air Force Office of Scientific Research), National Institutes of Health, Army Research Office, Defense Advanced Research Project Agency, Department of Defense Research and Engineering, Navy (including Office of Naval Research and Naval Air Systems Command), and the US Department of Education. US and international industrial sponsors include Quanta Computer, Inc.; Nokia; Foxconn; Boeing; Ford; Shell; Nippon Telegraph and Telephone Corporation; Cisco Systems; Pfizer, Inc.; Star Alliance; Toshiba; Eli Lilly; and Du Pont. Other organizations sponsoring research include Industrial Technology Research Institute, Singapore–MIT Alliance, Defense Science and Technology Agency, Delta Electronics Foundation, DSO National Laboratories, and Epoch Foundation.

Research Projects

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

T-Party

T-Party is a five-year, $20 million research project sponsored by Quanta Computer, Inc. The goal of the project is to (1) develop the next generation of platforms for computing and communication beyond personal computers, (2) create new systems for the development and seamless delivery of information services in a world of smart devices and sensors, and (3) move from a device-centric perspective to a human-centric one. Anant Agarwal, Hari Balakrishnan, Regina Barzilay, Srini Devadas, Fredo Durand, William Freeman, Jim Glass, John Guttag, Frans Kaashoek, Dina Katabi, Sam Madden, Rob Miller, Robert Morris, Stephanie Seneff, Christopher Terman, Steve Ward, and Victor Zue are the principal investigators from CSAIL.

During this fourth year of the collaboration, 17 investigators pursued projects in eight main areas:

- Virtualized computation platform, focusing on the development of secure and reliable computation and storage
- Unmanaged internet architecture, the development of a direct, secure, authorized, and authenticated access to (mobile) personal devices, including new technologies for robust wireless networking and an information technology platform for automotive applications
- Robust wireless communication using clever coding techniques to increase the throughput of wireless networks
- Just Play, distributed systems automatically constructed from ad hoc collections of disaggregated devices, which includes work on automation using web-based interfaces and the use of “soft hardware” for smart device prototypes
- Natural interactions, emphasizing the use of human language as a central ingredient in a multimodal interface (combining speech and gesture) for naive users and in real-world environments where a traditional graphical user interface is not practical
- Natural language summarization, techniques for automatically creating summaries from natural language descriptions (e.g., those found on web pages)
- Multicore and high-performance computing applications and architectures applied to tasks in the multimedia and medical computing domains
- Computational image manipulation leading to higher-quality images requiring fewer bits to transmit and store

Agile Robotics for Logistics: Developing a Voice-Commandable Robotic Forklift

MIT, in collaboration with Draper Laboratory and BAE Systems, formed a team of about 30 faculty, staff, and students with expertise in a variety of areas including drive-by-wire actuation, sensing and perception for situational awareness, planning under uncertainty, agile control for mobility and manipulation, and natural human–computer interfaces and interaction to support military logistics operations such as supply-chain management.

With funding from the US Army, we are developing a prototype robotic forklift capable of locating, engaging, lifting, transporting, and depositing pallets throughout an outdoor warehouse facility, all while under voice and gestural command of a human supervisor. Key technical challenges and innovations in the project include achieving effective situational awareness in a changing, semistructured, outdoor environment with nearby humans and other vehicles; safe interoperation with humans, including seamless autonomy handoff and return; effective task-level multimodal voice and gesture interface; and safe autonomous engagement, transport, and placement of various pallet loads over uneven outdoor terrain. The forklift embodies a novel technique called “apparent intent,” in which it announces its imminent motion through a variety of visible and audible annunciators before actually moving. We think this will be a key element in securing human acceptance of a multiton robot operating autonomously in their midst.

This research is under the direction of professor Seth Teller.

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, and theory. Our 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 year one, this effort supported seven projects as shown below:

- Bit-switched wireless networks (professors Hari Balakrishnan and Robert Morris)
- Speech understanding in multimedia environments (Drs. Jim Glass and Stephanie Seneff)
- Computational imaging for future displays and cameras (professors Fredo Durand, Edward Adelson, and Bill Freeman)
- Scene and object recognition (professors Antonio Torralba and Bill Freeman)
- Real-time hand tracking as a user interface (professors Jovan Popovic and Randall Davis)
- Research in theoretical computer science (professors Madhu Sudan and Ronitt Rubinfeld)

- Active personalization through NLP (professor Regina Barzilay)

In year two, the last two projects were replaced by a new project, unsupervised audiovisual person identification, by Dr. Glass and Professors Torralba and Davis.

A total of 12 Foxconn engineers in three shifts have visited CSAIL over an 18-month period to facilitate technology transfer and to receive training for advanced research and development in computer science. In addition, seven principal investigators visited Foxconn in Taiwan in fall 2008 to present our preliminary research results and to interact with Foxconn senior management to better understand how our research can achieve maximum impact for our sponsor.

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 ensure the long-term growth of the web. Currently, 360 member organizations, including most of the world's leading technology companies, are working to enhance the capabilities used within web documents and to transform today's web of linked documents to an expanded web of data and services across a wide range of devices, enabling everyone on the planet to collaborate and share data and information.

W3C is completing standards key to the interactivity features referred to as Web 2.0, including technologies behind AJAX, widgets, and other technologies that support a variety of richer document content and browser-based applications. The Mobile Web Initiative is making web browsing on mobile phones and other small devices as useful and usable as it is on desktop and notebook computers. W3C is building on this work to enable a web that is ubiquitous and interoperable across a wider range of devices as diverse as appliances, entertainment systems, printers, projectors, cars, planes, and specialized hand-held devices.

W3C's early leadership in semantic web (Web 3.0) research has resulted in completion of the basic semantic web standards and is supporting current work on rules interchange languages, ontology description language extensions, and continuing work to address interoperability challenges within health care, life sciences, e-government, multimedia, financial data, and social networking domains.

There is an increasing focus on ensuring that the web is both usable and useful to people who up to now experience challenges in these areas. New web content accessibility guidelines are expected to be completed this year. They describe how to make web content and web applications accessible to people with disabilities. These guidelines will have increasing implications for the world's aging population. Other work in W3C focuses on making the web on mobile phones useful and usable to people in developing countries and on designing web sites that meet guidelines for both accessibility and mobile access.

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.

Autonomous Micro-Air Vehicles

Unmanned air vehicles (UAVs) have been enormously beneficial in a number of environments. While UAV technology has been very successful in autonomous operation at high altitudes, supporting the military, UAVs have not been deployed to the same extent in urban environments or indoors—for example, operating in close proximity to troops or working with first responders performing house-to-house searches in a Hurricane Katrina scenario. The operational limits are largely due to the fact that most existing UAVs rely heavily on a baseline communications infrastructure for their operation, such as the Global Positioning System (GPS) satellite network and a connection with human operators. This reliance on infrastructure and human support limits the areas of operation; UAVs cannot easily fly in GPS-denied environments, in crowded or cluttered environments such as indoors or urban environments, or in close proximity to humans.

In contrast, small-scale or micro-UAVs (MAVs), have the promise of operating as teammates to humans in a variety of domains. MAVs are generally capable of flight in cluttered environments and can fly safely near people. Additionally, MAVs can easily be carried by soldiers or first responders and deployed when needed, rather than launched from a remote location. For example, imagine a first responder such as a firefighter carrying a quad-rotor MAV. Before entering an unknown building, the firefighter deploys the MAV, which circles the building, identifies an ingress point (e.g., an open window), and enters the building. The MAV uses sensors to build a map of the building and annotates the map with human-centered labels, such as "apartment," "open door," or "stairs." The firefighter chooses an entry point to the building and is met by the MAV on entry. The MAV flies above throughout the building, watching the firefighter's back as a "wingman" or flying ahead to act as a remote sensor when instructed.

The MAV applications have enabled new fundamental research results in probabilistic inference and planning under uncertainty. For example, we have developed a motion planning algorithm that allows the vehicle to deliberately plan exploration strategies that keep the vehicle safe. Additionally, we have developed new task planning algorithms based on probabilistic sensing that allow the vehicle to sense the environment in a number of modalities, such as range sensing, computer vision, and nontraditional sensors such as magnetic fields. We have demonstrated the technology in a number of environments; for example, we have shown that our MAV can fly up to a building, recognize an open window, enter the building, and autonomously explore it. The most recent research resulted in complete two-dimensional and partial three-dimensional models of different floors of the Stata Center, including the "student street."

In July, the MAV team competed in the 5th International Aerial Robotics Competition and was the only team able to complete the mission. In addition to winning the competition, the MAV team was awarded a prize for the best systems integration.

This research is under the direction of professor Nicholas Roy.

Programmable Matter: Creating Objects by Programming Smart Sheets and Smart Sand

Programmable matter is achieved when a collection of small robotic modules that are physically connected have the ability to respond to the request of creating a goal shape autonomously. Our project aims to create programmable matter by fabricating smart materials that integrate sensing, actuation, communication, computation, and connections to create "smart sand" and "smart paper." These light, compact, and smart materials pack compactly and create objects with desired functionalities (e.g., a wrench, an antenna, a cup, a plane) on demand. When use is completed, the same material is reconfigured to make new parts with different shapes and functionalities. The ability to make generically what is needed, when it is needed, will enable everyday objects in our world to become smart objects.

Our research program in programmable matter combines the state of the art in microrobotics, materials science, robotic planning and control, information theory, ad hoc networking, topology, and advanced materials and fabrication methods. We are developing devices and planning and control algorithms for two programmable matter concepts: making shapes by folding and making shapes by subtraction. We are designing and building two types of novel devices: smart sheets that can fold themselves into desired shapes by a process analogous to folding origami and the bag of smart (centimeter scale) sand capable of forming shapes by a process analogous to sculpting. We are also developing the planning and control algorithms required for programming and achieving desired shapes using these devices. So far, we have demonstrated (1) a smart sheet capable of folding itself into a tray, a table, and a star and (2) a bag of smart sand capable of sculpting a stylized flower, a humanoid, and a dog.

This work has contributed folding and subtraction as concepts for achieving programmable matter and novel devices capable of folding and subtraction. This research led to fundamental results in understanding the nature of self-organization and what can be achieved by engineered systems capable of self-folding and subtraction. This work aims to demonstrate the convergence between materials and machines; materials are gradually becoming smarter and machines are gradually becoming more flexible and capable.

This research is under the direction of professors Daniela Rus and Erik Demaine.

Super Maneuverable Flying Robots: Learning to Fly Like a Bird

Birds routinely execute aerial maneuvers far beyond capabilities of our best aircraft control systems. To better understand birds, to advance capabilities of aerial vehicles, and to address fundamental problems in nonlinear control, the CSAIL Robot Locomotion Group is attempting to engineer a robotic bird that can reliably land on a perch.

During a perching maneuver, birds rotate their wings and bodies so that they are almost perpendicular to the direction of travel and oncoming airflow. This maneuver increases the aerodynamic drag on the bird both by increasing the surface area exposed to the flow and by creating a low-pressure pocket of air behind the wing. Viscous and pressure forces combine for a rapid deceleration, but the maneuver has important consequences: the wings become "stalled," meaning they experience a dramatic loss of lift and of control authority, and the dynamics become unsteady and turbulent, making them incredibly difficult to model and predict accurately. The task is further complicated by uncertain wind dynamics (which affect both the bird and the perch) and the partial observability of the airflow. Yet birds perch with apparent ease.

We hypothesize that designing a high-performance control system for perching is easier than accurately modeling the dynamics ("birds don't solve Navier-Stokes"). In a departure from the dominant model-centric control approaches, we have been developing tools for machine learning to build coarse, approximate models of the aerodynamics and to improve the performance of the vehicle through trial-and-error motor learning. So far, we demonstrated aggressive perching maneuvers with a small fixed-wing aircraft, built a small wind tunnel for flow visualization and modeling, and engineered a two-meter wingspan robotic bird capable of autonomous steady, level flight.

This research is under the direction of professor Russ Tedrake. Multimedia content is available at http://groups.csail.mit.edu/locomotion/videos.html.

Systems Support for Multicore Processors

Multicore architectures, which have several processing cores on a die, have been adopted by most chip manufacturers. Hardware trends suggest that large-scale multicore machines with tens to hundreds of cores will appear within five years. Operating systems historically have had problems providing good performance on multiprocessors; the prevalence of multicore processors will make these problems much more urgent.

One source of poorly scaling operating system services is use of data structures modified by multiple cores; such data structures do not benefit from caching and thus are slow to access. The semantics of existing operating system application program interfaces (APIs) often require this kind of sharing; for example, the POSIX standard requires each of a process's file descriptors to be usable by any of the process's threads (i.e., on any core). The most common approach to increasing scalability is to redesign data structures with fine-grained locking or wait-free primitives; this approach is programmer intensive and reduces contention but does not eliminate it.

The approach we are exploring exploits the fact that some instances of a given resource type need to be shared, while others do not. If the operating system were aware of an application's sharing requirements, it could choose resource implementations suited to those requirements. This approach, however, requires rethinking operating system APIs to make sharing requirements explicit. Our prototype operating system, Corey, includes three major new APIs designed in this way. Address trees allow applications to control which parts of the address space are private per core and which are shared; kernel cores allow applications to dedicate cores to run specific kernel functions, avoiding contention over the data those functions use; and shares are lookup tables for kernel objects that allow applications to control which object identifiers are visible to other cores. These abstractions are implementable without intercore sharing by default but allow sharing among cores as directed by applications. Corey's use of these abstractions allows it to outperform Linux for a number of system-intensive applications.

This research is under the direction of professors Frans Kaashoek and Robert Morris.

Conversational Automotive Human–Machine Interaction

Over the past few years, researchers in CSAIL's Spoken Language Systems (SLS) group have been developing speech and language technology that can operate in mobile, multimedia, and multimodal environments. One of their prototype conversational speech interfaces, called City Browser, allows users to find addresses on a map, search for points of interest like restaurants and hotels, and obtain driving directions.

Over the past year, a new version of City Browser specifically targeting automobile applications has been developed. The automotive version of City Browser is currently deployed in a BMW 530xi sedan, which has been loaned to SLS by BMW in support of their speech research. The interface makes use of the car's built-in display, sound system, and iDrive controller. Speech is captured through an array microphone positioned on the driver's sun visor. The interface stands out from both commercially available interfaces and other research prototypes because it combines a capable graphical user interface with a conversational spoken language interface in an actual automobile.

Over the last six months, a preliminary user study involving 127 subjects solving 10 tasks was conducted by collaborators in the MIT AgeLab to study the effectiveness and acceptance of this technology across different age groups. Overall responses were quite positive, with more than three-quarters of the users reporting enjoying their interaction with the system.

This activity on vehicle-based conversational interfaces has helped motivate several important research topics, including noise robust speech recognition, dynamic vocabulary speech understanding, context-dependent help, information summarization, and personalization.

This research is under the direction of Drs. James Glass and Stephanie Seneff.

Robust Robotic Manipulation

Robot manipulators are very widespread in factory applications. Factory robots operate at high speed and with great precision, but the cost of doing so is that their operating environments must be very carefully engineered so that there is little or no variability in the task they must perform. If parts are not appropriately placed and aligned, the industrial robot has no way to compensate: it will simply grasp in the same way it always does.

Humans, in contrast, are very robust at manipulation tasks: if you reach out for something, as soon as you touch it, your fingers give you important information about where the object is and how, exactly, your fingers are placed with respect to the object. You can then, very flexibly, adjust the way you are grasping the object, to achieve a stable grasp with the object in the desired orientation in your hand (it's important, for example, that you wind up with a pencil or a spoon oriented the right way).

Our research is focused on integrating sensory information and reasoning about uncertainty into robot manipulation to improve its robustness.

We begin with an approach to visual object detection that can locate an object in an image, recognize its type, and construct a very coarse three-dimensional model of the object's location and shape. We can derive a probability distribution over the position and shape of the object from the output of the vision system: we call this distribution the robot's "belief" about where the object is.

In the next phase, our robot, which is equipped with fingertips that can sense the location and surface orientation of objects that it comes in contact with, must determine how to approach the object. It makes use of the entire belief about the object, taking into account the degree of uncertainty. If the uncertainty is very high, it might decide to perform some groping actions that give more information and narrow down its belief about the object's location. If the uncertainty is lower, it might decide to attempt to grasp an object by its handle. If, in the course of attempting to grasp the handle, the robot misses or makes contacts it does not expect, it uses that information to update its belief about the object's location and tries again.

This initial work will serve as the basis for constructing a flexible, robust, robotic manipulation system that can, for example, find a desired object in a cluttered cupboard by removing some objects to be able to detect and then reach the object of interest. Ultimately, such flexible robots will have wide application, from assistants for elderly people to flexible factory workers.

This research is under the direction of professors Tomas Lozano-Perez and Leslie Pack Kaelbling.

Computational Photography

The principle of image formation has remained largely unchanged since the invention of photography: a lens focuses light rays from the scene onto a two-dimensional sensor that records this information directly into a picture. The final image is a simple copy of the optical image reaching the sensor and image-quality enhancement is usually obtained through improvement in the optics. The emerging field of computational photography challenges this view and proposes to leverage computation between the optical images and the final picture to alleviate physical limitation, enable flexible postcapture editing, record new types of information such as depth, and enable a novel visual experience.

The addition of computation means that the optical image does not need to be similar to the final image, which greatly expands the possible imaging strategies. Computation can be more than a simple postprocessing that takes as input a traditionally formed image; it also deeply changes the rules of the game for the optical side of imaging. New optics must be designed together with the computation to optimize the whole imaging process. Until now, optics has been the key to enhancing our ability to view and image the world, but digital processing provides us with a new tool that vastly expands our ability to form and enhance images.

This research is under the direction of professors Fredo Durand, William T. Freeman, and Edward H. Adelson.

Cilk: Solving the Multicore-Programming Problem

Parallel programming has a reputation for being notoriously difficult. The widespread availability of multicore processors has therefore created a strong demand for tools that make it easier for application programmers to harness a multitude of processing cores. Research on the Cilk (pronounced "silk") multithreaded programming language, developed at MIT over the last two decades, simplifies the parallel-programming task by making it a seamless extension of conventional serial programming.

Cilk minimally extends the C programming language with just five keywords that allow programmers to specify interactions among computational threads in a simple and high-level fashion. Cilk's runtime system contains a "work-stealing" scheduler that dynamically maps the computation onto available processing cores in a provably efficient fashion. The performance of an application written in Cilk is mathematically guaranteed to scale up linearly with the number of processing cores as long as the application has sufficient parallelism and the architecture has sufficient bandwidth. In addition, Cilk provides a race detector that guarantees finding and localizing pernicious parallel-programming bugs, called determinacy races.

Cilk has won numerous awards and has influenced many modern industrial concurrency platforms including Threading Building Blocks (Intel), Parallel Patterns Library (Microsoft), Fortress (Sun Microsystems), and X10 (IBM). In 2006, MIT licensed Cilk to a start-up, called Cilk Arts, which is producing a C++-based concurrency platform called Cilk++. All these systems contain work-stealing schedulers modeled after the original MIT Cilk system.

This research is under the direction of professsor Charles E. Leiserson and Dr. Bradley C. Kuszmaul.

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 the Institute, we hope our work will enable us to help bridge the digital divide.

CommuniTech

CommuniTech was founded in 2000 by a former CSAIL graduate student and is managed by the Public Service Center. The program seeks to bridge the digital divide in our local communities by providing adults in low-income families with the tools they need to gain access to valuable information they can use to better their lives and the lives of their families.

The two main foci of the program are providing skills and accessibility for technology management. The first portion provides a six-week training course in basic computer applications, which is taught by MIT student volunteers. At the end of the program, used computers that have been donated to the Institute and refurbished by students are given to participants for home use. An alliance with a local internet provider gives a discounted rate to participants in the program.

The skills learned in the first portion of the program and subsequent refinement due to the donated computers facilitates greater computer proficiency, marketability to prospective employers, and connection to the world at large.

Lacotec Laare, Kenya

The Laare Community Technology Centre in Kenya was founded by Eric Mibuari ’06, 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. This center is in an area of the country where nearly no one owns a computer and very few people have ever seen one. At present, the center particularly targets youth with high school educations and seeks to equip them with basic computing skills that they can apply for personal use, in industry, and in education. The center was started in January 2005 and has so far offered various levels of training to three groups of more than 70 students. The center takes particular note of the economic difficulties of the many people who would like to benefit from its training and facilities and strives to charge the minimal feasible fees.

CSAIL is currently providing support, with future plans to expand the center. Eric Mibuari, an MIT alumnus and the center’s founder, has secured four acres of public land donated by the community to build a new, larger center. The plan is to improve the center’s learning facilities, expand its curriculum, access online resources, and increase its capacity to serve more people in the community.

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 Israeli and Palestinian young leaders. MEET enables its participants to acquire advanced technological and leadership tools while empowering them to create positive social change within their own communities. Many of our students volunteer to teach MEET summer courses at the Hebrew University in Jerusalem. CSAIL continues to host http://meet.csail.mit.edu/ and provide technical support to the MEET program.

TEK

While 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 and internet service provider 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 (e.g., a single phone in a school), it is often difficult to time-share between internet and voice. Furthermore, unreliable network and power infrastructures can sometimes block internet access altogether. Several CSAIL members have been supporting the TEK project. TEK stands for time equals knowledge, and the project empowers low-connectivity communities by providing a full internet experience using email as the transport mechanism.

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 with 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 CSAIL's support 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 generously donated by Maxtor for OCW proliferation. This expansion of access will give talented students without the technological resources they need a chance to reach their full educational potential.

Seminar Series

Four distinguished speakers gave presentations during this year's Dertouzos Lecture Series:

- Cynthia Dwork, Microsoft Research, "Privacy: A Natural Resource to Be Conserved"
- Marc Raibert, Boston Dynamics, "BigDog: Is There Light at the End of the Tunnel?"
- David DeWitt, University of Wisconsin and Microsoft Research, "Clustera: A Data-Centric Approach to Scalable Cluster Management"
- Mendel Rosenblum, Stanford University, "The Impact of Virtualization on Modern Computing Environments"

Organizational Changes

The leadership and organizational structure of CSAIL remained the same in 2009. Victor Zue continued to serve as director of CSAIL. The director's duties include developing and implementing strategies designed to keep CSAIL growing and evolving, fund raising, determining laboratory policies, and examining promotion cases. Three associate directors—Anant Agarwal, Daniela Rus, and Madhu Sudan—remained in their positions, assisting the director with his duties.

CSAIL's principal investigators are organized into three research directorates: artificial intelligence, systems, and theory. These directorates met weekly during the academic year. Tomas Lozano-Perez and William Freeman continued to oversee artificial intelligence, and Hari Balakrishnan continued to oversee the systems directorate. Nancy Lynch took on oversight of the theory directorate in 2009.

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.

CSAIL's executive committee met monthly to advise the director on policies and processes within the laboratory. The executive committee consists of the director, associate directors, research directors, assistant directors, space czar, and two additional faculty members—Srini Devadas and Patrick Winston.

Elizabeth Bruce continued as director of industry partnership. She oversees the CSAIL Industry Affiliates Program (CSAIL-IAP), which grew to 28 members in 2009. CSAIL-IAP is a corporate membership program that offers companies the opportunity to access CSAIL's faculty and students through annual conferences, recruiting events, and onsite visits. Additionally, Elizabeth and Dr. Chris Terman lead CSAIL's industry strategy working group, which explored policies and processes to improve engagement with our industry partners.

Awards and Honors

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

Hari Balakrishnan: ACM Fellow

Tim Berners-Lee: Foreign associate of the US National Academy of Sciences

Konstantinos Daskalakis: Game Theory and Computer Science Prize, Game Theory Society; ACM Doctoral Dissertation Award

Randall Davis: 2009 GSA Graduate Counselor Award, Electrical Engineering and Computer Science

Fredo Durand: Best Paper Honorable Mention with Bill Freeman, Anat Levin, and Yar Weiss, Computer Vision and Pattern Recognition (CVPR) 2009

Bill Freeman: IEEE Fellow; Best Paper Honorable Mention with Fredo Durand, Anat Levin, and Yar Weiss, CVPR 2009

Michel Goemans: ACM Fellow

Polina Golland: Medical Image Computing and Computer-Assisted Intervention International Conference, two Best Paper Runner Ups with Boon Thye Thomas Yeo, Mert R. Sabuncu, and Bruce Fischl and with Wanmei Ou and Matti S. Hamalainen

Daniel Jackson: 2009 MacVicar Teaching Fellow

Dina Katabi: Sigcomm ACM, Best Paper Award with Shyamnath Gollakota; William R. Bennett Prize in the Field of Communications Networking with Sachin Katti

Tom Leighton: 1st Class SIAM Fellow; National Academy of Sciences, elected member

Barbara Liskov: Institute Professor; 2008 Turing Award

Nancy Lynch: PODC/CONCUR 60th Birthday Celebration Symposium Honor; Brooklyn College 2009 Distinguished Alumna Award

Joel Moses: ACM Fellow

Tomaso Poggio: NSF Distinguished Lecturer August 2008; Werner Reichardt Center for Integrative Neuroscience Inauguration, University of Toronto, main speaker, December 2008; NSF computer and information science and engineering/biological sciences/social behavioral and economic sciences/mathematics and physical sciences/engineering, distinguished lecturer, April 2009; EEE Symposium on Computational Intelligence for Financial Engineering 2009, elected honorary member

Nicholas Roy: 2008 International Conference on Robotics and Automation Best Paper Award with Ruijie He and Sam Prentice; 2008 MAV Best Overall Mission Performance, Best Rotorcraft Performance; 2009 Institute of Electrical and Electronics Engineers Robotics and Automation Society Early Career Award

Daniela Rus: Association for the Advancement of Artificial Intelligence, elected member

Michael Sipser: American Academy of Arts and Sciences, elected member

Madhu Sudan: ACM Fellow

Peter Szolovits: National Research Council Computer Science and Telecommunications Board, elected member; 2009 Harvard–MIT Health Sciences and Technology Thomas A. McMahon Mentoring Award

Antonio Torralba: CVPR 2009, Best Student Paper with Jenny Yuen

Victor Zue: Academia Sinica, Taiwan, Academician

Key Statistics for Academic Year 2009

CSAIL members in AY2009: 814
Faculty: 83 (13% women)
Research staff: 36 (19% women)
Administration, technical, and support staff: 70 (60% women)
Postdocs: 36 (22% women)
Visitors: 122 (15% women)
Paid Undergraduate Research Opportunities Program: 90 (31% women)
Master of engineering students: 46 (24% women)
Graduate students: 331 (23% women)

Victor Zue
Director
Delta Electronics Professor of Electrical Engineering and Computer Science

More information about the Computer Science and Artificial Intelligence Laboratory can be found at http://www.csail.mit.edu/.