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 history of technological innovation impacting 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 application of sensor technology to traffic congestion, animal herding, medical monitoring, and climate observations. Robotic locomotion and human robotic interface are being investigated, as well as medical solutions with image-guided surgery and technical applications to aid clinical decisions. Advancements in biological research are underway, including developments in 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 agencies include: the National Science Foundation, the Defense Advanced Research Project Agency, the National Institutes of Health, the Air Force Office of Scientific Research, the National Aeronautics and Space Administration, the Office of Naval Research, and the Department of Defense Research and Engineering. A wide spectrum of US and international companies include: Quanta Computer, Inc., Nokia, Foxconn, Shell, Microsoft Corporation, Toyota, Nippon Telegraph and Telephone Corporation, the Boeing Company, DuPont, SAP, Cisco Systems, Inc, Pfizer, Inc., Other organizations also support CSAIL including the Singapore-MIT Alliance, the Industrial Technology Research Institute, the Defense Science Technology Agency, the Commonwealth Scientific and Industrial Research Organization, the Delta Environmental and Educational Foundation, and the 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.

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. Victor Zue, Srini Devadas, Jim Glass, Frans Kaashoek, Chris Terman, and Steve Ward are the principal investigators from CSAIL.

During this third year of collaboration, 15 investigators pursued projects in 5 main areas:

- Virtualized Computation Platform, focusing on the development of secure and reliable computation and storage.
- T-Net, the development of a direct, secure, authorized, and authenticated access to (mobile) personal devices, including new technologies for robust wireless networking and an IT platform for automotive applications.
- Just Play, distributed systems automatically constructed from ad-hoc collections of disaggregated devices. 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. Also 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.

Mobile Communications EcoSystems

We are two-thirds through this three-year, $7.5 million research project sponsored by Nokia. The goal of the project is to develop new modes of interaction in the use of mobile communication devices and to reduce the development cost of new services and energy-efficient phones. Researchers from Nokia Research Center Cambridge, located within a five minutes' walk from CSAIL, work closely with CSAIL researchers. So far six joint meetings involving 35 to 100 Nokia researchers and technologists have been held at MIT and at Nokia facilities in Helsinki, Tampere, and Oulu in Finland. Arvind, Victor Zue, and Rebecca Henderson (SLOAN) comprise the CSAIL committee that oversees this project.

This initiative started in January 2006 and currently consists of the following projects:

- Virpi, developing techniques to interact with applications on mobile phones primarily using speech and textual dialogs
- ARMO, reducing the development cost of new energy-efficient, high-performance phones
- MyNet/UIA, making it dramatically easier to connect devices to share information and services in a secure manner

- Traz, developing techniques for secure information flow on phones running stock Operating Systems
- Telepresence, providing high quality mobile medical telepresence via radio diversity
- Automatic, scalable summarization of user opinions
- Scalable context awareness via organic location discovery and shared rich maps

Defense Advanced Research Project Agency Urban Challenge

MIT, in partnership with Olin College and Draper Laboratory, competed in the 2007 DARPA Urban Challenge, held in Victorville, CA. Talos, our Land Rover LR3, finished in fourth place and was one of six vehicles to complete an approximately 55 mile urban course in the presence of other robotic vehicles and human driven traffic vehicles. The challenge was the first of its kind to require vehicles to autonomously operate in a dynamic urban environment. Vehicles had to navigate a network of paved suburban and dirt roads among other autonomous cars as well as human-driven vehicles. The cars were tasked with completing a series of mock supply missions while obeying standard California driving laws. Much like people do, they had to deal with many of the challenges inherent to urban driving: navigating among other moving cars, re-planning due to roadblocks, and so on.

CSAIL's vehicle was a Land Rover LR3 equipped with a diverse range of lidar, vision, radar, and navigation sensors connected to a powerful blade cluster computer system. Our vehicle employed novel algorithmic approaches to perception, planning and control for the challenging task of autonomous driving in uncertain, dynamic environments. These innovations provide a strong platform for future research in autonomous driving in GPS-denied and highly dynamic environments with poor a priori information.

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

MIT, in collaboration with Draper Laboratory and BAE Systems, has formed a team of about thirty 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 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, semi-structured, 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 varying 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 that this will be a key element in securing human acceptance of a multi-ton robot operating autonomously in their midst.

Center for Robotics

The CSAIL Robotics Center was formed in September 2005 to provide a mechanism for collaboration and for promoting robotics within CSAIL. The center is codirected by Tomas Lozano-Perez and Daniela Rus. The center has 12 principal investigators from four different departments, 76 graduate students, 11 postdocs, and many MIT undergraduates participating in the Undergraduate Research Opportunities Program.

The CSAIL Robotics Center's mission is to conduct cutting-edge, long-term research and education in robotics. Our research addresses fundamental problems in designing more capable robots and controlling them to interact intelligently with people, the environment and each other. We also explore how increased capability and intelligence can enable new societal benefits through applications in homes, fields, oceans and outer space. Our basic and applied research inspires the development of novel course material aimed both at training students in robotics and at using robots to explore computation, sensing, and control.

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. The W3C is responsible for developing and maintaining the standards that make the web work. 420 member organizations, including most of the world's leading technology companies, are currently working 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 Web 2.0, including technologies behind AJAX, widgets others that support an interesting, client-side experience on the web. 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 upon 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 and new work to address interoperability challenges within health care, life sciences, e-government, multimedia, and energy 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, and these describe how to make the web accessible to people with disabilities, which will also have increasing implications for the world's aging population. A new working group focuses on making the web on mobile phones useful and usable to people in developing countries.

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:

CarTel: A Mobile Sensor Data Management System

CarTel is a software and hardware platform for sensing using moving vehicles that we have been developing since 2005. Data is captured from onboard sensors, including location from GPS sensors, Wi-Fi radios, and onboard diagnostics interfaces standard in most automobiles. This data is stored in a local database on each car and transferred opportunistically, via available Wi-Fi networks or cellular modems to a "portal" where users can browse and visualize their own data as well as contribute data into aggregate views, showing, for example, traffic or the locations of potholes on roads. To allow non-expert users to specify what data they would like to collect from remote vehicles, CarTel includes a simple database-like interface for programming and configuration.

CarTel is currently running on a network of about 40 cars, including cars from two local cab companies and several cars from users in our research group. We have built a range of applications on top of this basic platform. For example, our personal commute management portal allows users to browse historical traffic conditions and retrieve minimum-expected-delay routes between two locations at different times of day, as well as compare their drive times and fuel consumption to other users with whom they choose to share their data. In another application, we have built a road surface quality application that uses car-mounted accelerometers to show city planners and drivers the locations of the most severe potholes in and around Cambridge.

These applications have enabled several new fundamental research results. For example, we have developed networking protocols that allow mobile cars to detect and take advantage of WiFi networks from moving cars by reducing the time establish a connection from about 10 seconds to less than 400 hundred milliseconds, tripling the number of bytes cars can successfully send over wireless networks when compared to unmodified wireless networking software. We have also developed a suite of data management tools to eliminate systematic sensing errors and smooth over gaps of missing data generated by network outages, as well as to query and manage the vast array of data produced by this testbed.

This research is under the direction of associate professor Samuel Madden and professor Hari Balakrishnan.

Automated Initiation of a Neural Stimulator for the Control of Epileptic Seizures

Epilepsy is a disorder of the central nervous system that predisposes individuals to experiencing recurrent seizures. Despite recent advances in the medical and surgical management of epilepsy, more than 20% of individuals with epilepsy never achieve adequate control of their seizures.

Vagus Nerve Stimulation (VNS) is an adjunctive treatment option for medically refractory epilepsy. In some patients, initiating vagus nerve stimulation soon after the onset of a seizure may abort or ameliorate seizure symptoms. Unfortunately, a significant number of patients cannot initiate VNS by themselves following seizure onset. Instead, these patients rely on caregivers to initiate VNS following the start of a seizure. This burdens caretakers and denies patients the therapeutic benefit of vagus nerve stimulation in their absence.

To facilitate the use of VNS, we designed a non-invasive, computerized system that automatically initiates VNS following the computerized detection of seizure onset. The system uses machine learning to build a patient-specific algorithm that detects seizure activity through analysis of non-invasive EEG and ECG signals. In a clinical evaluation of the system on five patients, we demonstrated the feasibility of initiating VNS following the real-time detection of a seizure or epileptiform discharge. We are currently in the process of recruiting patients for a more extensive evaluation of the utility of the system.

This research is being done under the direction of professor John Guttag.

Network Coding

Wireless, in its various forms, is an increasingly dominant communication medium. However, current wireless networks suffer from low throughput as the density and number of users increases. Furthermore, this contention cannot be easily alleviated by using different portions of the spectrum for different applications as RF spectrum is a scarce commodity.

This project is to design and build next-generation wireless networks that support high-throughput dense deployment. Our approach is based on network coding; instead of purely forwarding packets, wireless relays and routers mix (i.e., code) the packets before forwarding. The theory of network coding analytically proves that mixing packets at intermediate nodes achieves multicast capacity and increases throughput. The main challenge, however, is to translate these theoretical results into practical gains. Many of the theoretical assumptions do not apply in practice. For example, most network traffic consists of unicast flows not multicast. Further, traffic is usually bursty, and the sending rate is unknown in advance even to the sender itself, and varies over time. Also, connectivity in a wireless network is highly variable due to changing medium conditions. Thus, in practice, the wireless environment is highly unpredictable and difficult to capture using existing models. Our work addresses this challenge by developing new network coding protocols that are opportunistic, i.e., they make no assumptions about the traffic or the underlying topology, but identify coding opportunities as they arise, and exploit them to increase throughput and reliability.

Over the past couple of years, we have been able to establish network coding as an alternative design for wireless networks that delivers significant improvement in throughput and reliability. We have integrated network coding within the current network stack, extended the approach to both inter-session and intra-session coding, proved robustness against Byzantine adversaries and benefits for privacy, introduced symbol-level network coding, a cross-layer approach that involves cooperation between two layers in the stack: the physical and the network layers, and pushed network coding

all the way to the analog domain, introducing Analog Network Coding, which operates over signals instead of bits. We have also implemented our designs and evaluated them in an actual wireless testbed, showing that they deliver large practical gains.

This research is under the direction of associate professor Dina Katabi.

LabelMe

Currently computers have difficult recognizing objects in images. While practical solutions exist in computer vision for a few simple object classes such as human faces or cars, the more general problem of recognizing all different classes of objects in the world (e.g. guitars, bottles, telephones) remains unsolved. Computer Vision researchers are currently investigating methods that can recognize and localize thousands of different object categories in complex scenes. A key component of these algorithms is the data used to train the computers' model of each object. However, the size and contents of existing datasets, among other factors, limit current methods from scaling to thousands of object categories. Research in object detection and recognition would benefit from large image and video collections with ground truth labels spanning many different object categories in cluttered scenes.

At CSAIL we have developed LabelMe (http://labelme.csail.mit.edu/) a database and an open annotation tool that allows researchers to share and to annotate images. The dataset is dynamic, free to use, and open to public contribution. The application exploits the capacity of the web to concentrate the efforts of a population of users in order to build a database larger than what can be afforded by a single group. The tool has been online since August 2005. As of July 18, 2008, LabelMe had 163,975 images, 44,831 annotated images, and 265,412 labeled objects. The database is continuously growing as users label new objects everyday and contribute pictures. The dataset is being used worldwide, and is also being used as support to teaching in computer vision classes. LabelMe has motivated the emergence of new applications in computer vision and computer graphics.

One important concern when data is collected using web-based tools is quality control. Currently quality control is provided by the users themselves. Polygons can be deleted and the object names can be corrected using the annotation tool online. Despite the lack of a more direct mechanism of control, the annotations are of quite good quality. Another issue is the complexity of the polygons provided by the users—do users provide simple or complex polygon boundaries? Contributors tend to provide good quality annotations sufficient for training and evaluating most object detection and segmentation algorithms.

This research is under the direction of associate professor Antonio Torralba, in collaboration with Bryan Russell at the École Normale Supérieure, France.

Chickenfoot

The migration of applications to the World Wide Web opens up new opportunities for user interface customization. An application that would have been impossible to customize on the desktop suddenly sprouts numerous hooks for customization when

implemented in a web browser, without any effort on the application developer's part. These hooks can be used for automating tasks (clicking links, filling in forms, and extracting data) and for customizing the application (changing appearance, rearranging components, and inserting or removing user interface controls or data). The Chickenfoot project is developing new techniques for exploiting these hooks to make the web easier for users to personalize.

Our research focuses on two barriers users face when customizing an application: *complexity* of the application (particularly if the user is forced to understand its implementation in order to customize it), and *syntax* of the programming language that must be learned. Our approach to the *complexity* barrier is simply stated: a user should never have to view the HTML source of a web page to customize or automate it. Chickenfoot is a scripting environment embedded in the Firefox web browser that provides a platform for automating and customizing web applications through a familiar interface, as web pages rendered in a web browser. It uses novel pattern-matching techniques to allow users to describe components of a web page (targets for interaction, extraction, insertion, or customization) in terms that make sense for the rendered view. For example, the *click* command identifies a button to be clicked using keywords from its visible label, rather than the name it was given by the web page designer.

For the *syntax* barrier, we have found that users can draw from their experience with search engines to write a set of keywords expressing a command, such as "click I'm Feeling Lucky button," "push the Lucky button," or even just "feeling lucky," which an interpreter can convert into an appropriate script command. We call this technique *keyword programming*, since it relies only on keywords, and not on formal syntax or even well-formed natural language. We have explored keyword programming not only in the web automation domain, but also in other domains such as Java development. One surprising result is that programming language syntax often has relatively little information content, and can be inferred automatically from only a handful of keywords—allowing us to design programming systems that reduce the learning and complexity burdens on their users.

The Chickenfoot project is led by associate professor Rob Miller.

Project Angstrom

MIT CSAIL Project Angstrom is a joint research endeavor on the future of multicore computing. With the number of cores per chip scaling to the 1000s by the middle of the next decade, a fundamentally new set of constraints will govern software and hardware design. The programming challenge for multicores is but one major symptom of this issue. Multicore software and applications will require a new kind of reasoning—incorporating the spatial aspect of computing—as well as an increasing focus on energy, reliability, and memory bottlenecks.

Project Angstrom will build a 1,000-processor computer prototype within the next five years to serve as a test bed for multicore research and development initiatives at CSAIL. Project Angstrom will rethink languages, compilers, operating systems, architectures and interfaces for multicore computing, and will explore what this means for emerging

multicore applications in both embedded and general purpose domains, including such areas as computational video, imaging, autonomous vehicle control, web applications, scientific computation, and data center on a chip.

This research is under the direction of professor Anant Agarwal.

Natural Multimodal Interaction

Despite decades of significant advances in both hardware and software, the predominant means of human-computer interaction is mired in technology – the keyboard, mouse and display – that dates back as much as forty years. We'd never be happy with the computers of 40 years ago, so why are we using interaction technology that old?

Our work in natural multimodal interaction seeks to change that. Our vision is to make interacting with computers much more like interacting with other people, by making it possible for people to sketch, gesture, and talk, and have the computer sketch, gesture, and talk back at them.

Our work on sketch understanding has demonstrated systems in a variety of domains, including physics (sketching simple physical devices), electronics (drawing basic analog circuits), and programming (class hierarchies for object-oriented programs). Our systems can understand these hand-drawn sketches, illustrating that understanding by simulating and animating the things drawn, making the sketch appear to come alive.

We are combining sketching and speech in a system that understands hand-drawn chemical structures: The user can draw a structure, then ask questions about it aloud (relying on speech understanding technology built by the SLS group in CSAIL). The user can, for example, draw a compound, then ask "Does this have any known toxic effects?" Our system looks up the answer by sending the interpreted structure to the appropriate web site, then reports back the result by reading aloud the relevant text from that site. The effect is of having an intelligent whiteboard capable of conversing in a limited way about chemistry.

Our work in gesture understanding starts from the premise that computers will be far better at understanding spoken language if theyhave access to the wide range of modalities that accompany speech, notably hand gestures. We have analyzed gestural patterns, discovering parallel structures in the meaning of the associated speech, and have developed novel, structured statistical models for multimodal language processing that enable learning about gesture in its linguistic context, rather than in the abstract. We have used these ideas successfully in a variety of language processing tasks, including resolving ambiguous noun phrases, segmenting speech into topics, and producing keyframe summaries of spoken language. In all three cases, the addition of gestural features—extracted automatically from video—yields significantly improved performanceover state-of-the-art text-only alternatives. This marks the first demonstration that hand gesture improves automatic discourse processing.

Our work on multimodal output generation has studied people describing engineering design, to understand their use of sketching, speech and gesture. We are designing a

system that can generate similar multimodal output, leading eventually to the ability for it to understand and participate in a relatively unstructured multimodal conversation about engineering designs.

One long-term consequence of our work may lie in future electronic textbooks. Imagine a physics text, for example, where every diagram was not only a movie that could be run, but could also be changed by sketching on it by hand. A diagram with masses, cables, and springs, for example, could show the dynamic behavior of the system, then be modified by the user simply by drawing over it, making changes by attaching additional masses, changing the anchor points of springs, etc. The net result would be text in which every diagram was in effect an inviting mini-laboratory that allowed students to explore a variety of scenarios of their own invention

This research is under the direction of professor Randall Davis.

Web 2.0 Computer Games for Language Learning

With the recent rapid emergence of Web 2.0 technologies, and widespread adoption of Voice over IP (VoIP), one can easily imagine a day when students will routinely interact with educational services that depend critically on audio capture and transmission of speech from an ordinary internet browser. With this in mind, we are developing several different web-based systems which allow learners of Chinese to communicate with their computers by speaking to them in Chinese. These systems are typically configured as simple games: card games, translation games, dialogue games, etc. Each game focuses on a restricted vocabulary task, both to provide a limited scope exercise to the student and to help overcome the difficulties of recognizing heavily accented speech. Significant scaffolding provides tutorial help to avoid student frustration. The games can be configured such that the student competes against the computer, or such that two students compete with each other, each from the comfort of their own home.

These systems offer several opportunities for challenging research problems. We have provided an easy-to-use interface to allow a student or a developer to create a set of software flashcards that can be used as playing cards in the games, in order to personalize the vocabulary being learned. We are working towards the goal of being able to enable a non-expert developer to create new games that build on the technology already in place. Some of the games require very high quality spoken language translation, but within a restricted domain. All of the speech from our user base is automatically recorded, and these data will be a valuable resource for research in detecting students' errors in phonetics, tone production, and grammar usage.

This research is under the direction of principal research scientist Stephanie Seneff.

Learning to Summarize Consumer Opinions

Online product reviews are an important source of information for consumers. Websites like Amazon contain hundreds of reviews for almost any imaginable product or service. Reading all relevant reviews would be daunting and time consuming, reading one review may result in biased or incomplete information. To facilitate user's access to review information, we developed Precis, a review compression system. Precis summarizes

hundreds of user opinions into a succinct list of pros and cons that captures common views about the product. Currently, Precis has been applied to generate summaries for 50,000 different products and services, ranging from strollers to cellphones.

In the core of our system is a method for identifying the document-level semantic properties implied by a text. Traditionally, this task has been addressed in a supervised learning framework that requires large quantities of annotated data. We have developed a new method for learning these semantic properties by leveraging unstructured annotations available on the web. As an example, consider the lists of pros and cons that often accompany reviews of products and services. Such end-user annotations are increasingly prevalent online, and they grow organically to keep pace with subjects of interest and socio-cultural trends. The major obstacle to the computational use of such free-text annotations is that they are inherently noisy—there is no fixed vocabulary, no explicit relationship between annotation key phrases, and no guarantee that all relevant semantic properties of a document will be annotated.

To exploit such noisy annotations, we find a hidden paraphrase structure that clusters the key phrases. The paraphrase structure is linked with a latent topic model of the review texts, enabling the system to predict properties of unannotated documents and to effectively aggregate semantic properties of multiple reviews. Our approach is implemented as a hierarchical Bayesian model with joint inference. We find that joint inference increases the robustness of the key phrase clustering and encourages the latent topics to correlate with semantically meaningful properties. Multiple evaluations demonstrate that our model substantially outperforms alternative approaches for summarizing single and multiple documents to a set of semantically salient key phrases.

This research is under the direction of associate professor Regina Barzilay.

Laboratory Sponsored Activities

CSAIL Outreach

Imara

The overall goal of Imara is to find and implement longterm, sustainable solutions to make educational technology and resources available to domestic and international communities.

CommuniTech

CummuniTech 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 centre 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 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 centre particularly targets the youth with high school education and seeks to equip them with basic computing skills that they can apply for personal use, in industry and in education. The centre was started in January 2005 and has so far offered various levels of training to three groups of more than seventy students. The centre takes particular note of the economic difficulties of the many people who would like to benefit from its training and facilities and purposely strives to charge the minimal feasible fees.

CSAIL is currently providing support, with future plans to expand the center. Eric Mibuari, an MIT alumni 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 ISP 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 to direct Web access, email can be much cheaper, more reliable and more convenient in developing areas. The TEK Client. operates as a proxy on the user's machine, enabling users to browse downloaded pages using a standard Web browser. New searches are automatically encoded as emails and sent to the TEK Server, which queries the Web and returns the contents of resulting pages via email. TEK is free software distributed under the GNU LGPL license.

OCW Outreach Initiative

Adnan Esmail, a mechanical engineering graduate student, has combined MIT's OpenCourseWare 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 course information available to those in nations with poor bandwidth infrastructure.

The program has been realized in partnership with Aligarh University in Aligarh, India and Lahore University of Management Sciences, in the capital of Pakistan's Punjab province. It runs on hard drives that have been generously donated by Maxtor for OCW proliferation. This expansion of access 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. They were:

- Professor William Dally, Stanford University, "Stream Processing: Efficiency through Locality"
- Professor Christos Papadimitriou, U.C. Berkeley, "The Algorithmic Lens: How the Computational Perspective is Transforming the Sciences"
- Professor Lenore Blum, Carnegie Mellon University,"(What We Can Learn from the) Canaries in the Computer Science Coal Mine"
- Professor Shree Nayar, T. C. Chang Professor of Computer Science, Columbia University, "Computational Cameras: Redefining the Image"

Organizational Changes

On July 1, 2007 CSAIL underwent a change in its leadership and organizational structure. Victor Zue took over as the director of CSAIL. The director's duties include developing and implementing strategies designed to keep CSAIL growing and evolving; funding raising, determining laboratory policies; and examining promotion cases. Three associate directors, Ananat Agarwal, Daniela Rus and Madhu Sudan were chosen to assist the director with his duties.

CSAIL's four research directorates, Language, Learning, Vision and Graphics (LLVG); Physical, Biological and Social Systems (PBSS); Systems; and Theory were consolidated into three. LLVG and PBSS were combined into a single directorate, Artificial Intelligence. The Systems and Theory directorates remained unchanged. New research directors were also named to steer the research agenda of the laboratory. Tomas Lozano-Perez and William Freeman oversee Artificial Intelligence, Hari Balakrishnan oversees Systems, and Ronitt Rubinfeld oversees Theory.

Additionally, John Guttag was asked to take on the role of space czar, managing allocation of space within CSAIL and implementing improvements to facilities that would increase the quality of the environment for the laboratory's faculty, staff and students.

Jack Costanza continued as the assistant director for infrastructure, overseeing IT 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 was reorganized. This committee meets 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, the space czar, and two additional faculty members, Srini Devadas and Patrick Winston.

Lastly, in May 2008, Elizabeth Bruce joined CSAIL as the director of industry partnership. She will oversee the CSAIL Industry Affiliates Program (CSAIL-IAP) launched in January 2008. CSAIL-IAP is a corporate membership program that offers companies a unique opportunity to access CSAIL's faculty and students through annual conferences, recruiting events and onsite visits. Additionally, Elizabeth is charged with developing and strengthening CSAIL's relationships with both new and existing industry sponsors.

Awards and Honors

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

Arvind	National Academy of Engineering—Member Outstanding Achievement Award—University of Minnesota
Tim Berners-Lee	Fellow IEEE IEEE/RSE Wolfson James Clerk Maxwell Award Pathfinder Award, Kennedy School of Government
Rodney Brooks	AAAI—Classic Paper Award, Honorable Mention
Erik Demaine	Francqui Foundation—International Francqui Chair Carnegie Mellon University and Tokyo University of Technology—Katayanagi Emerging Leadership Prize
Michael Ernst	ACM—ESEC/FSE Distinguished Paper Award

Shafi Goldwasser	ACM—Athena Lecturer, Committee on Women in Computing (ACM-W)
Tommi Jaakkola	Conference on Uncertainty in Artificial Intelligence—Best paper award
Manolis Kellis	Alfred P. Sloan Foundation Award, 2008-2010
Jonathan Kelner	MIT EECS—Sprowls Doctoral Dissertation Award
Tom Leighton	National Academy of Sciences – Member
Charles Lesierson	2008 ACM SIGPLAN Most Influential PLDI Paper Award
Barbara Liskov	ACM SIGPLAN Lifetime Achievement Award for 2008
Nancy Lynch	2007 Dijkstra Prize
Silvio Micali	National Academy of Sciences – Member
Tomaso Poggio	Distinguished Speaker at DARPA-IPTO
Ronald Rivest	Marconi Prize EECS—Jamieson Award for Teaching Excellence
Peter Shor	INFORMS—Computing Society Prize for best paper
Russ Tedrake	Microsoft Research—New Faculty Fellowship Award NSF Career Award
	Jerome Saltzer Award for teaching recitations
Chris Terman	EECS—Jamieson Award for Teaching Excellence
Brian Williams	Science Foundation Ireland—E.T.S. Walton Visitor Award (2007) AAAI—Fellow
Jack Wisdom	National Academy of Sciences—Member
Victor Zue	International Speech Communication Association— Medal of Scientific Achievement

Key Statistics for Academic Year 2008

CSAIL members number 816 in AY2008:

Faculty	81(14% women)
Research Staff	44 (20% women)
Admin, Tech, and Support Staff	75 (57% women)
Postdocs	30 (20% women)
Visitors	124 (15% women)
Paid UROPs	59 (29% women)
MEng Students	38 (13% women)
Graduate Students	365 (22% 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/.