Computer Science and Artificial Intelligence Laboratory

The principal mission of the Computer Science and Artificial Intelligence Laboratory (CSAIL) is to invent new information technologies and discover new computing paradigms that better the lives of people everywhere, while creating an internal culture of joyous innovation. Over the last four decades, CSAIL has been instrumental in the development of many computer science and artificial intelligence innovations, including time-shared computing, public key encryption, computer chess, web standards, GNU, TCP/IP, ARPANet, and many more.

CSAIL research is sponsored by US government agencies including the US Army, the Air Force Research Laboratory, the CIA, the Defense Advanced Research Projects Agency, the Departments of Defense, Justice, Energy, Education, and Information, NASA, the National Geospatial-Intelligence Agency, the National Institutes of Health, the National Science Foundation, and the Office of Naval Research; by companies such as ABB, Acer, BAE Systems, Boeing, CISCO, Delta Electronics, Dupont, Ford, Hewlett-Packard, Honda, Intel, Microsoft, Mitre, Mitsubishi, NTT, Nokia, Phillips, Pfizer, Quanta, Shell, Sun, Toyota, and Verizon; and by large-scale institutional collaborations such as the Cambridge–MIT Institute, the Commonwealth Scientific and Industrial Research Organisation (Australia), the Electronics and Telecommunications Research Institute (Korea), the Industrial Technology Research Institute (Taiwan), and the Singapore–MIT Alliance; and by other institutions including DSO National Laboratories (Singapore), the Epoch Foundation, the Deshpande Center, the Packard Foundation, and the Sloan Foundation. Additionally, the World Wide Web Consortium (W3C) consists of over 400 organizations that help set standards for a continuously evolving web.

CSAIL is organized into four broad research areas:

- Systems covers all aspects of the building of both hardware and software computational systems. Steve Ward is the research director.
- Theory looks at the fundamental mathematical underpinnings of all aspects of computer science and artificial intelligence. Madhu Sudan is the research director.
- Language, Learning, Vision, and Graphics includes work on the sorts of things that all people manage to do effortlessly, both emulating those abilities and simulating their appearance. Leslie Kaelbling is the research director.
- Physical, Biological, and Social Systems might also be called complex adaptive systems. It covers work from robotics to molecular biology and from semantic systems to computational models of politics. Randy Davis is the research director.

The CSAIL Executive Committee determines policies, examines promotion cases, and discusses strategies designed to keep the Laboratory at its peak. Members are drawn from each research area and include, in addition to the four research directors mentioned above, Rodney Brooks (director), Victor Zue (codirector), Arvind, Bonnie Berger, Tim Berners-Lee, Erik Demaine, Srini Devadas, William Freeman, Daniel Jackson, Thomas Knight Jr., Daniela Rus, Jack Costanza (assistant director), and Karen Shirer (assistant director).

Research Projects

Within CSAIL we have many single- and multi-investigator projects, but we also have a number of virtual centers and large-scale mega projects.

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. It has 12 principal investigators from four different departments, 76 graduate students, 11 postdocs, and many MIT UROPs.

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 research, basic and applied, inspires the development of novel course material aimed both at training students in robotics and at using robots to explore computation, sensing, and control.

Center for Information Security and Privacy

The goal of the Center for Information Security and Privacy (CISP) is to develop both the theoretical foundation for secure systems and to engineer practical systems. CISP addresses the full range of security concerns and technologies, from hardware and architectural support for secure computing to applications with significant societal impact, such as electronic voting and email authentication. CISP has particular strength in cryptography, both theoretical and applied—indeed, many of the major advances and themes in cryptographic research originated in CISP. Much of the center's applied research focuses on problems emerging from the growth of the internet, as millions of embedded networked devices come online. CISP also develops techniques for promoting individual privacy in the face of pervasive sensors and networks. CISP is headed by Professor Ron Rivest, and Professors Srini Devadas and Frans Kaashoek act as associate heads. More information can be found at http://cisp.csail.mit.edu/.

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. Rodney Brooks, Randy Davis, Srini Devadas, Chris Terman, and Victor Zue comprise the CSAIL side of the committee that oversees this project.

During the first year of this collaboration, the following four projects were being pursued by some 20 principal investigators:

- 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
- "Just Play," distributed systems automatically constructed from ad-hoc collections of disaggregated devices
- Natural Interactions, emphasizing the use of human language as a central ingredient in a multimodal interface (combining speech, vision, and gesture) for naive users and in real-world environments where a traditional graphical user interface is not practical.

Mobile Communications EcoSystems

This is a three-year, $7.5 million research project sponsored by Nokia. The goal of the project is to develop new ways for people to use mobile communication devices in smart spaces, using a tightly coupled model of university/industry relations in which Nokia researchers from their research center in Cambridge, situated within a five-minute walk from the Laboratory, work closely with CSAIL researchers. Arvind, Rodney Brooks, and Victor Zue comprise the CSAIL committee overseeing this project. The project started in January 2006, and it currently consists of the following projects:

- Simone: Develop advanced speech-based interaction capabilities for mobile phones
- Startmobile: Make natural language application development easier
- Mynet/UIA: Make it easy to connect devices and services together in a secure manner
- Asbestos: Control information flow to prevent inadvertent or malicious disclosure of personal information
- Swapme: Make it possible to develop applications that are policy, preference, and context aware
- Composeme: Discover composability of web services, phone data, and software components
- Armo: Make it easy to develop high-performance, energy-efficient hardware for new mobile phones.

World Wide Web Consortium

Hosted through a partnership between MIT, the European Research Consortium for Informatics and Mathematics, and Keio University, the World Wide Web Consortium (W3C) has standardized the foundation of today's web (including HTML and XML) and is developing the foundation for tomorrow's web. The consortium has experienced significant growth in technical work and participation over the past year. W3C opened its 18th and 19th offices in India and China and saw its membership climb by over 10 percent to 415 companies and organizations. Members are drawn by new work that is transforming the current web of linked documents to an expanded web of data and services across a wide range of devices, and that aims to enable everyone on the planet to collaborate and share information. Think "Web 2.0." The Rich Web Clients Activity is standardizing AJAX and other technologies to support improved 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 larger computers and screens. Future work will focus on web ubiquity, leveraging web technologies to support interoperable data exchange and multimodal interaction (including graphical, voice, and gesture) across devices as diverse as appliances, entertainment systems, printers, projectors, cars, planes, specialized hand-held devices, and so on. The foundational specifications for web services are developed at W3C, with work on policy and semantic annotations just recently underway. W3C's early leadership in semantic web research has resulted in completion of the basic semantic web standards and is supporting current work on query and rules interchange languages, as well as new work to address interoperability challenges within health care, life sciences, multimedia, and geospatial domains. New web content accessibility guidelines describe how to make the web accessible to people with disabilities, which will also have increasing implications for the world's aging population. Though XML became a standard in 1998, important additions such as XML Query, Voice XML, and binary XML are under development. Work continues on these and other technologies, with a total of 60 W3C working and interest groups.

Research Highlights

Over 300 CSAIL research projects can be found at http://publications.csail.mit.edu/abstracts/abstracts06. Highlights of a few are included here.

Biped Locomotion in Uncertain Environments

Robotic biped locomotion is a difficult problem for several reasons. Walking is underactuated: the stance leg acts as a pivot about which the walker falls forward at each step. In contrast, a typical robotic arm is fully actuated and can freely control each degree of freedom independently. Environmental uncertainties for a real-world walker (such as the upcoming profile and compliance of terrain) are more unpredictable and can have greater impact on the dynamics as compared to the disturbances a robotic arm typically experiences in a factory setting, and efficiency becomes more critical for a machine that must carry its own power source as it travels.

The underactuation of the process makes implementing traditional multilink robot control in uncertain real-world environments quite challenging. This research focuses on the assumption that efficient biped walking will best be achieved by exploiting the inherent coupling between degrees of freedom rather than "canceling" them. This is a biomimetic strategy: there is strong evidence that human walking also exploits the coupled pendulum dynamics of the stance and swing legs, and mechanical walking machines that have been built on these passive dynamic principles often exhibit very natural (humanlike) gaits. Our biped research attacks this problem of developing efficient and robust biped walking from several directions: (1) theory, (2) control (using reinforcement learning), and (3) experimental verification.

We seek to establish a better analytical understanding of the passive (unactuated) dynamics of the system, particularly when subjected to disturbance forces and/or uncertain terrain. The basic model we are using is the “compass gait” walker, so named because it resembles a compass. It has two straight legs with distributed masses and a single pivot. Given appropriate initial conditions, such a device can walk down a small, even slope forever: it reaches a steady limit cycle where the energy lost at each foot impact is exactly recovered by the change in potential energy (height) of each step. However, motions on a rough surface will be less predictable. This is a 2-D theoretical model of walking that assumes all motion is constrained to a single plane; a real walker must also maintain lateral stability to avoid falling.

Another piece of this work is to appropriately quantify what “stability” means for walking in uncertain environments. For instance, we suggest that measuring performance in terms of actual expected number of steps taken (or “mean first passage time” until falling) provides a more meaningful metric than less direct methods now used, which depend on instantaneous (geometric) measures (e.g., “center of mass” location or the “zero moment point”) that do not incorporate the inherent stability (or its lack) of the coupled-dynamic system as a whole in judging the stability of a walking gait. We view successful walking as a “metastable” state: a state that is expected to persist for a very long time but that is also expected to fail occasionally, given large enough disturbances from the environment (and cannot therefore be defined as strictly stable). While any one walker design may not be guaranteed to continue walking forever given any eventuality, some walkers can be expected to perform much better than others. We wish to quantify the baseline stability of a passive walker both to provide a benchmark against which we can judge the performance of an actuated and actively controlled walker, and also to identify how particular (passive) mechanical designs can enhance the stability of the compass gait, and we wish to develop these results into principles for creating inherently stable mechanical biped design.

Creating control policies for these nonlinear systems is a challenge of its own. Linearization about a particular path of expected operating points only suffices if disturbance forces or uncertainties (in sensor information or upcoming terrain) can be guaranteed to be sufficiently small. But a real-world walker must expect the unexpected. This requires knowledge of a policy to deal with the types of uncertainties (tripping, slipping, etc.) one expects with walking, rather than a less flexible plan that depends on staying closely on track to one particular path of motion. Reinforcement learning can be employed to learn a policy for selecting control actions. As a walker is exposed to a variety of disturbances (walking on carpet instead of tile flooring, recovering from a push, etc.), it can test and evaluate the performance of various actions to compile a policy for acting that does not rely on sticking to one specific plan of action for success.

This research is headed by Professor Russ Tedrake.

Ortholog and Paralog Detection with Multiple Complete Genomes

The first challenge in comparative genomics is to reliably determine orthologous genes across multiple species. Current approaches work well with pairs of species but accumulate inaccuracies as more relatives are considered. Hence, new methods are necessary to achieve accurate and reliable alignments of the upcoming mammalian, fly, and fungal species. We have developed and prototyped a new phylogenetic reconstruction method: species informed distance-based reconstruction (SINDIR). SINDIR finds the maximum likelihood gene tree using the known species and pairwise distance estimates between gene sequences. The algorithm is designed specifically for the phylogenomic problem—namely, the use of phylogeny to reconstruct the ancestry of all genes within several complete genomes. The resulting phylogeny can be used to make reliable ortholog assignments across several species, which is useful in comparative studies and also sheds light on the evolution of the species.

With a distance-based phylogenetic approach, it is possible to efficiently reconstruct phylogenies of thousands of gene families, some as large as hundreds of genes, in a reasonable amount of time. However, existing distance-based approaches, such as neighbor joining, often construct erroneous gene trees due to long branch attraction. Often these errors produce gene trees that disagree with the known species tree and thus erroneously infer gene duplications and losses. The key insight of our algorithm is to use the known species tree during the gene tree construction in order to identify unlikely inferred gene duplications and losses.

The algorithm is given as input for a rooted species tree topology and a small subset of genes known to be orthologous within the species of interest. Given these inputs, the algorithm computes a density estimation of the distribution of phylogenetic trees for true orthologs. Once the distribution parameters are learned, the algorithm takes a distance matrix for a family of genes with unknown orthology. In a search constrained by the distance matrix, the algorithm finds the maximum likelihood tree according to the learned distribution. Using reconciliation, the orthology of the gene family can be inferred.

This research is headed by Professor Manolis Kellis.

Quantitative Information-Flow Tracking for Real Systems

Because digital information can be easily copied, it can easily go astray. Information that is intended to be confidential can accidentally be transferred to subjects who should not have it, and information that should not be trusted (perhaps because it comes from an unknown party on a network) can affect the operation of critical processes. To prevent such failures, computer systems should track the provenance of the information they work with, so that it can be given the correct levels of protection and trust. We are exploring a new, more practical technique for such information-flow tracking.

The key technical challenge in performing such tracking is to follow information not just as it is stored and transferred, but as it is used and modified by potentially complex calculations. For instance, in most operating systems it is up to a program to mark its outputs as confidential if necessary. However, if the program makes this determination incorrectly (perhaps because it contains a bug), the information is not protected. An alternative strategy called mandatory access control makes the assumption that every output of a secret-using program must be secret. While safe, this policy makes many programs impossible to run. What is needed instead—and is missing from existing systems—is a way to track information flow at a fine-grained level through the operations performed on data. Previous language-based approaches, such as type systems, can track information at a fine-grained level but are not applicable to programs written in languages such as C and C++. Operating system techniques are potentially applicable to any program but require substantial changes in system architecture.

Our approach to this challenge has two innovative aspects. First, the information tracking is performed on preexisting software by tracking its operation at the instruction level. This makes the technique effective even for memory-unsafe languages such as C. Because the level of analysis exactly matches a program's actual execution, the results are both comprehensive and precise. Second, the measurement of information flow is a numeric estimate of how many bits of information could have been transferred at any point in a particular program execution. Most previous information-flow approaches use a binary notion of security: either *any* of amount of information might flow or *none.* However, more important in practice is to distinguish between executions that reveal a limited quantity of information versus those that reveal too much—for instance, between an execution of a password-checking program that reveals whether a password is correct and one that releases the entire password database.

To define the number of bits of information that a program transmits from its secret inputs to its public outputs, consider the graph representation of the program's computation as a flow network in which each intermediate value is a channel whose capacity is its size in bits. Then the number of bits that the program might leak is bounded by the maximum flow in this network. Rather than computing this maximum flow globally, however, a more efficient local approach is to annotate the boundaries between portions of the program that operate primarily on secret data and those that operate primarily on public data. The encapsulation strategy corresponds to finding a minimum cut in the flow network and can also be used to bound the amount of information leaked by secret-dependent control-flow decisions (so-called indirect flows).

This research is headed by Professor Michael Ernst.

A Fast Approximation of the Bilateral Filter Using a Signal-Processing Approach

The bilateral filter is a nonlinear filter that smooths a signal while preserving strong edges. It has demonstrated great effectiveness for a variety of problems in computer vision and computer graphics. Unfortunately, little was known about the accuracy of this acceleration. We propose a new signal-processing analysis of the bilateral filter, which complements the recent studies that analyzed it as a preliminary data element or as a robust statistics estimator. Importantly, this signal-processing perspective allows us to develop a novel bilateral filtering acceleration using a downsampling in space and intensity. This affords a principled expression of the accuracy in terms of bandwidth and sampling. The key to our analysis is to express the filter in a higher-dimensional space where the signal intensity is added to the original domain dimensions. The bilateral filter can then be expressed as simple linear convolutions in this augmented space followed by two simple nonlinearities.

This new formulation allows us to derive simple criteria for downsampling the key operations and to achieve important acceleration of the bilateral filter. We have shown that, for the same running time, our method is significantly more accurate than previous acceleration techniques. From a practical point of view, we are able to process large images in a very short time. For instance, a 2-megapixel picture is filtered in less than a second, and the result is visually identical to the exact computation that requires several minutes, even with optimized code. In comparison, the previous technique shows important visual differences.

This work introduces a novel interpretation of the bilateral filter. It contributes to a better understanding of the nonlinear filters by unveiling a strong link between linear and nonlinear techniques. We believe this work can yield further theoretical progress in this domain.

From a practical point of view, the achieved speedup and faithfulness pave the way for high-quality and interactive image manipulations. Being able to process in a second picture at screen resolution enables a direct user interaction where only an offline computation was available before. Such a tool opens new possibilities in the field of computational photography and video processing.

This research is headed by Professor Fredo Durand.

Geometric Folding Algorithms

The area of geometric folding and unfolding is attractive in that problems and even results can be easily understood with little knowledge of mathematics or computer science, yet the solutions are difficult and involve many sophisticated techniques. The general form of a folding problem is to determine how a particular geometric object (e.g., linkage, a piece of paper, a polyhedron, or protein) can be reconfigured or folded according to a few constraints, which depend on the object being folded and the problem of interest. In particular, we are interested in characterizing the existence of foldings and efficient algorithms for finding foldings—or in proving that such algorithms are impossible.

Most folding and unfolding problems are attractive from a pure mathematical standpoint for the beauty of the problems themselves. But additionally, most of the problems have close connections to important industrial applications. Linkage folding has applications in robotics and hydraulic tube bending and has connections to protein folding. Paper folding has applications in sheet-metal bending, packaging, and air-bag folding. Unfolding polyhedra has applications in manufacturing, particularly sheet-metal bending.

In various collaborations, we have solved a variety of folding problems. Here we describe a few of the more recent results:

- Energy-driven approach to linkage unfolding. We present a new algorithm for unfolding planar polygonal linkages without self-intersection based on following the gradient flow of a “repulsive” energy function. This algorithm has several advantages over previous methods.

- Flat-state connectivity of fixed-angle linkages. We explore which classes of linkages in 3-D have the property that each pair of their flat states—that is, their embeddings in 2-D without self-intersection—can be connected by a continuous dihedral motion that avoids self-intersection throughout.
- Infinitesimally locked linkages. We propose a new algorithmic approach for analyzing whether planar linkages are locked in the sense that they cannot be moved into some other configuration while preserving the bar lengths and not crossing any bars. The idea is to examine self-touching or degenerate frameworks in which multiple edges converge to geometrically overlapping configurations.
- Short interlocked linkages. We study collections of linkages in 3-space that are interlocked in the sense that the linkages cannot be separated without one bar crossing through another. We explore pairs of linkages—one open chain and one closed chain, each with a small number of joints—and determine which can be interlocked.
- Producible protein chains. Fixed-angle polygonal chains in 3-D serve as an interesting model of protein backbones. Here we consider such chains produced by a ribosome, modeled crudely as a cone, and examine the constraints this model places on the producible chains.

This research is headed by Professor Erik Demaine.

For all of the above research projects, see http://publications.csail.mit.edu/abstracts/abstracts06 for full citations.

Laboratory Sponsored Activities

K–12 Outreach

In the past year, CSAIL has hosted high school students who participate in a volunteer internship giving the students hands-on experience with complex systems management, network, and web technology. Students from Commonwealth High School, Framingham High School, and the Massachusetts Academy of Science at Worcester Polytechnic Institute took advantage of this program.

CSAIL is highly visible to the local communities as a place to visit because it has an architectural focus and because it is the MIT home of robots and computers. We focus our outreach attention on high schools, giving tours during which we strive to generate enthusiasm in the fields of computer science and artificial intelligence, combining talks with film clips of some of the devices that have been conceived and built at the Lab. Most of the groups are from Massachusetts and the surrounding states, but we have had high school children from as far away as Nevada visit to learn about the research we do. This past year we increased our tour capabilities by adding a virtual tour of the building. Our goal is increasing awareness and encouraging applicants, especially young women and minorities.

Collaborations with the MIT Museum

The combination of a visit to CSAIL and a visit to the MIT Museum is a natural activity for people traveling to the campus. The two organizations have worked closely this year to coordinate visits and to make presentations at each other's sites when appropriate. This included a special evening program lecture by CSAIL's director, Rodney Brooks.

Outreach Programs

CSAIL support of existing programs continues to grow. The MIT Summer Research Program increased the minority summer students at the Laboratory from 4 in 2004 to 12 in 2005. Several other existing outreach programs use visits to CSAIL as part of their normal activity. The Lab's voluntary outreach programs overseas in Kenya, Fiji, and Israel were aided this year by funds from outside sources.

Kenya

CSAIL members traveled to Laare, Kenya, to provide technical expertise and educational resources to the Laare Training Center. The Laare Training Center received some of our decommissioned computers last year and is preparing to receive another shipment this year. Lab member Anthony Zolnik, along with Eric Mibuari, founder of the Laare Training Center while chairman of the Africa Internet Technology Initiative, held a training session that addressed educational technology and management model options for the training center. After the trip, we learned that the Laare Training Center founders have been given a three-acre parcel of land on which to build their own training center. CSAIL is assisting in raising awareness and funds to help build this computer training center.

Fiji

This year the CSAIL Fiji outreach program provided new laptops, digital cameras, and printers to multiple islands located in Fiji. Professor Daniela Russ and Jack Costanza traveled to Taveuni, where they held training sessions instructing teachers and students in the use of their new laptops, including the K–12 educational software installed on the computers. Google provided funding for the laptops, and the Infrastructure Group@ CSAIL provided the technical expertise to configure and program the hardware. CSAIL prepared to send three more laptops by the end of summer 2006. We are also assisting in raising awareness and funds to implement the first K–12 computer lab on Taveuni.

Israel

In conjunction with the Sloan School of Management and the Office of the President, CSAIL continued to support Middle East Education through Technology (MEET) this past year. MEET is an innovative educational initiative aimed at creating a common professional language between young Israeli and Palestinian 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 meet.csail.mit.edu and provide technical support to the MEET program.

Professional Development

CSAIL hosted the Center for Women in Enterprise/Springboard Enterprise Bootcamp event in October 2005, a "workshop designed to provide entrepreneurs with the tools they need to successfully develop and execute a venture capital presentation." For women thinking about starting their own business, this information is invaluable. Established in 1995, the Center for Women & Enterprise (CWE) is a nonprofit organization dedicated to helping women start and grow their own businesses. Springboard Enterprise forums are designed to increase investment opportunities for women-led firms and to help women entrepreneurs navigate the equity markets. CWE's mission is "to empower women to become economically self-sufficient and prosperous through entrepreneurship."

CSAIL continues to expand our outreach efforts to broaden computer science and artificial intelligence diversity in society across gender and racial boundaries.

Seminar Series

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

- Professor Avi Wigderson, Institute for Advanced Study: "The Power and Weakness of Randomness in Computation"
- Professor Randal E. Bryant, Carnegie Mellon University: "Formal Verification of Infinite State Systems using Boolean Methods"
- Professor Michael Black, Brown University: "Repairing the Damaged Brain with Computation: The Development of a Neural Motor Prosthesis"
- Dr. David E. Shaw, D.E. Shaw Research and Development and Center for Computational Biology and Informatics, Columbia University: "New Architectures for a New Biology"

Awards and Honors

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

Ted Adelson became a member of the National Academy of Science and received the IEEE Longuet-Higgins Award.

Tim Berners-Lee was awarded the Sigma Xi Common Wealth Award for Science and Invention; the Financial Times Lifetime Achievement Award; the Die Quadriga Award; and the President's Medal of the Institute of Physics.

Regina Barzilay was awarded the Ross career development chair; *Technology Review*'s "Top 35 Innovators under the Age of 35"; the Institute of Electrical and Electronics Engineers (IEEE) Intelligent Systems' "AI Ten to Watch"; and Microsoft Research's New Faculty Fellowship Award.

Rod Brooks became an Association for Computing Machinery (ACM) fellow and a corresponding member of the Australian Academy of Science.

Jack Costanza received the Steven Wade Neiterman Award.

Randy Davis received the Inaugural Influential Paper Award of the International Foundation of Autonomous Agents and Multi-Agent Systems (with R. Smith).

Fern DeOliviera was awarded the MIT–DAPER Instructor of the Year.

Erik Demaine was awarded the Sloan Foundation Research Fellowship and was appointed Esther and Harold E. Edgerton professor.

Jim Glass was appointed as lecturer in Health Sciences and Technology.

Tom Greene was named the Academic Resource Center's Freshman Advisor of the Year.

Frans Kaashoek became a member of the National Academy of Engineering.

Dina Katabi was appointed the NBX career development chair and was awarded a Sloan Foundation Research Fellowship.

Butler Lampson received the ACM Symposium on Operating Systems Principles Hall of Fame Award for "Hints for Computer System Design" and the International Federation for Information Processing TC11 Kristian Beckman Award for Information Security.

Tom Leighton received the ACM Special Interest Group on Algorithms and Computation Theory Distinguished Service Prize.

Barbara Liskov was awarded an honorary doctorate from the Eidgenössische Technische Hochschule (ETH), Switzerland.

Nancy Lynch was awarded the Van Wijngaarden Prize from Centrum voor Wiskunde en Informatica.

Sam Madden was named as one of *Technology Review*'s "Top 35 Innovators under the Age of 35."

Silvio Micali was awarded UC–Berkeley's Distinguished Alumnus of the Year, the University of Pennylvania's Rademacher Lecturer, and the Information Security Executive Award.

Rob Miller received User Interface Software and Technology's Best Paper Award.

Joel Moses was appointed acting director of the Engineering Systems Division.

Tomaso Poggio was the keynote speaker at SigGraph, ETH's 150-year anniversary, the International Conference on Mathematical Harmonic Analysis, and at the Institute for Pure and Applied Mathematics; he was also appointed a member of the Scientific Advisory Board of the Institute for Scientific Interchange Foundation (Turin) and a member of the Comitato di Esperti dell'ISICT Genoa.

Nick Roy received an NSF Career Award.

Madhu Sudan was awarded Danny Lewin Outstanding Professor.

Peter Szolovitz became a member of the Institute of Medicine of the National Academies and a fellow of the American Institute for Medical and Biological Engineering.

Patrick Winston received the MIT Graduate Student Council Graduate Teaching Award.

Scott Ananian, Krste Asanovic, Bradley C. Kuszmaul, Charles E. Leiserson, and Sean Lie received the 2006 IEEE Micro "Top Picks" Award as one of the 10 most significant papers of the year in computer architecture for their paper "Unbounded Transactional Memory."

Key Statistics for AY2006	
Faculty	78 (13% women)
Research staff	42 (19% women)
Admin, tech & support staff	80 (37% women)
Postdocs	26 (31% women)
Visiting faculty	68 (1% women)
Paid UROPs	41 (27% women)
MEng students	61 (26% women)
PhD & MS students	402 (20% women)
Total	**798 CSAIL members**

Affirmative Action

CSAIL supports the affirmative action goals of the Institute.

Rodney Brooks
Director
Panasonic Professor of Robotics

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