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 brings together faculty, researchers, and students in a wide-ranging program of study, research, and experimentation. The laboratory's membership stands at 813, comprising 166 faculty members and research staff, 454 graduate students, 61 undergraduates, 96 visitors, affiliates, and postdoctoral associates and fellows, and 36 support staff. The majority are affiliated academically with the Department of Electrical Engineering and Computer Science, but CSAIL also includes faculty and students from Mathematics, Brain and Cognitive Sciences, Aeronautics and Astronautics, Mechanical Engineering, the Biological Engineering Division, and the Harvard–MIT Division of Health Sciences and Technology.

CSAIL research is sponsored by US government agencies, including AFOSR, ARDA, CIA, DARPA, NASA, NIH, and NSF; by companies like ABB, Acer, BAE Systems, Boeing, CISCO, Delta Electronics, Hewlett-Packard, Honda, IBM, Microsoft, NTT, Nokia, Phillips, Quanta, Shell, Sun, and Toyota; and by large-scale institutional collaborations such as the Cambridge-MIT Institute, the Commonwealth Scientific and Industrial Research Organisation (Australia), the Industrial Technology Research Institute (Taiwan), and the Singapore-MIT Alliance. Additionally, the World Wide Web Consortium (W3C) consists of nearly 400 organizations that help set standards for a continuously evolving World Wide Web.

CSAIL is organized into four broad research areas:

- *Architecture, Systems, and Networks* (ASN) covers all aspects of the building of both hardware and software computational systems. Srini Devadas is ASN's research director.
- *Theory* looks at the fundamental mathematical underpinnings of all aspects of computer science and artificial intelligence. David Karger is the research director.
- *Language, Learning, Vision, and Graphics* (LLVG) 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* (PBSS) 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 CSAIL 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, Tim Berners-Lee, Erik Demaine, William Freeman, Daniel Jackson, Thomas Knight Jr., Lissa Natkin, assistant director, and Karen Shirer, assistant director.

Research Projects

Project Oxygen

In late 1999, CSAIL launched the five-year Project Oxygen. Oxygen sought to bring about revolutionary changes in the ways computers and humans interact through a vision of pervasive, human-centered computing. The Oxygen Alliance, between MIT and six world-class companies—Acer, Delta Electronics, Hewlett-Packard, Nokia, NTT, and Philips—was formed in June 2000. Since then, researchers at MIT have worked actively on all aspects of Oxygen, sometimes in close collaboration with our Oxygen partners.

In June 2005, Project Oxygen reached its five-year milestone. Numerous accomplishments in diverse research areas included human-centered interfaces that combine multiple modalities (speech, vision, gesture, etc.) to make it easier to interact with machines; new ways to pinpoint people and devices indoors that can provide location-based services; a new computing architecture and chips to enable devices to morph into various functions on the fly via software; a flexible software infrastructure that can reconfigure itself as people and devices move about; new security methods that can protect the integrity of mobile devices and networks; and new collaboration technologies that can create new ways for people to work together. Victor Zue, Rodney Brooks, and Anant Agarwal comprise the committee that oversaw this project.

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 online at *http://cisp.csail.mit.edu/*.

T-Party

In March 2005, a team of CSAIL principal investigators led by Rodney Brooks and Victor Zue held a four-day workshop for Quanta Computer, Inc., at their headquarters outside Taipei, Taiwan. Quanta's chairman, Barry Lam, had asked CSAIL to think about the future of mobile computing, and the workshop presented the results of three months of brainstorming by the CSAIL community. At the conclusion of the workshop, Quanta and

CSAIL entered into a five-year, $20 million research agreement to look into hardware and software technologies that would be key to future computation and communication platforms. Rodney Brooks, Victor Zue, Chris Terman, Randy Davis, and Srini Devadas comprise the committee that oversees this project.

World Wide Web Consortium

The World Wide Web Consortium (W3C) is where the framework for today's web was developed (including HTML and XML) and where the framework for tomorrow's web is being designed. The consortium is cohosted by MIT, the European Research Consortium for Informatics and Mathematics, and Keio University and has an additional 15 outreach offices around the world. To further increase global participation, W3C opened an office in India, is exploring the potential for new offices in China and Latin America, and has lowered membership fees for organizations in developing countries. Membership in W3C has grown to almost 400 companies and organizations this year. W3C has issued 90 web standards in total, including, over the past year, standards to support XML key management, speech synthesis markup, efficient web services, and the use of international character sets. The "Architecture of the World Wide Web, Volume One" presents for the first time in one document core components, constraints, and good practices for web standardization and development. W3C's Patent Policy has now been phased in successfully, ensuring that standards will be developed under the most comprehensive patent policy in the internet standards industry and be implemented without the barrier of licensing fees. Efforts recently started will further expand interoperability between programs, data, devices, and people across the expanding world of things on the internet. For example, the new Mobile Web Initiative is forging standards and guidelines to facilitate adaptation of web content to devices with different display capabilities and access capacities. W3C's early leadership in Semantic Web research has resulted in completion of the basic Semantic Web standards, motivated development of a new query standard, and increased momentum for new efforts to develop a rules standard and to address interoperability challenges within the health and life sciences.

Research Highlights

Lecture Retrieval and Access

In the past decade, there has been a dramatic increase in the availability of online academic lecture material. Low-cost media and fast networks have opened the door to new and exciting ways to disseminate knowledge in media formats ranging from audio recordings to streaming video. These educational resources can potentially change the way people learn—students with disabilities can enhance their educational experience, professionals can keep up with recent advancements in their field, and people of all ages can satisfy their thirst for knowledge. It is conspicuous, however, that in contrast to many other communicative activities, lecture processing has until now enjoyed little benefit from the development of human language technology.

The goal of CSAIL research in this area is to provide efficient access to large online multimodal collections of educational material through development of the Lecture Retrieval and Access (LERA) system. The exploration began by analyzing the

performance of commercial speech recognition engines on lecture recordings. An unexpectedly high word error rate (WER)—66.8 percent—motivated research on new recognition techniques tailored for the lecture genre. Preliminary investigations have improved the WER to 40.6 percent through acoustic and language model adaptation. The new method depends on effective use of a domain-specific written corpus (e.g. textbooks) to accurately predict lecture vocabulary. In parallel, methods that can accurately recognize lecture fragments, which are crucial for successful navigation, are being explored. Specifically, an unsupervised learning method that can automatically identify technical terms that are spoken multiple times in an audio-only source (i.e. no text) is being developed. These technical terms can subsequently be used for browsing and for retrieval of lecture material. Future plans include exploration of synergies between text-based and audio-based keyword discovery mechanisms.

As a first step toward automatic summarization and content retrieval, we focused on the problem of topical segmentation of lecture material and developed a segmentation algorithm based on spectral clustering methods and random walks that has been successfully applied to image segmentation. These methods were able to pinpoint the most salient topic boundaries in spontaneous speech transcripts, but at a finer granularity the performance degraded. Currently we are exploring ways to modify existing spectral clustering methods to take into account the constraints imposed by lecture discourse.

Finally, we developed a tool for automatic terminology extraction from a corpus of lecture transcripts. We are currently in the process of customizing this tool to the needs of the MIT OpenCourseware Initiative, where the tool will be used to automatically create metadata for a large repository of MIT course materials. This research is headed by professor Regina Barzilay.

Machine Learning for Natural Language Processing

Natural language processing (NLP), or computational linguistics, deals with the application of computational methods to problems involving linguistic data. Examples of application areas within NLP include automatic (machine) translation between languages; dialogue systems, which allow a human to interact with a machine using natural language; and information extraction, where the goal is to transform unstructured text into structured (database) representations that can be searched and browsed in flexible ways. In terms of technical and scientific challenges, computational linguistics involves fundamental questions concerning how to structure formal models (e.g. statistical models) of natural language phenomena and the design of algorithms that implement these models.

The research focus has been on machine-learning approaches to NLP problems. A first subarea of research has concentrated on the development of theory and algorithms underlying machine-learning approaches to *structured* problems—problems that involve complex, discrete structures such as strings, labeled sequences, or trees. Structured problems of this kind are extremely prevalent in NLP, as well as in related research areas such as computational biology, speech recognition, and machine vision. One contribution over the last year has been the development of a new family of

algorithms for discriminative parameter estimation. A second contribution has been the development of a new formalism—case-factor diagrams—that subsumes two central frameworks in probabilistic models for structured problems: Markov random fields (MRFs) and weighted context-free grammars (WCFGs).

A second subarea of research has concentrated on applying machine-learning methods to NLP problems such as machine translation, dialogue systems, and speech recognition. One approach learns to map sentences to logical forms that represent their meaning. For example, in building a natural language interface to a database of geographical information, the goal would be to map strings such as "what states border Texas?" to logical forms that can be used to query the database. This research has developed a radically new approach to the problem, one that simultaneously learns the syntactic and semantic representations required for the mapping. A second application is machine translation. Recent work has shown how syntactic information can be used to improve the quality of machine translation systems through a linguistically motivated account of systematic differences in word order between languages such as German and English.

In the past year, this research has resulted in best paper awards at three major conferences in the field: Empirical Methods in Natural Language Processing 2004 and Uncertainty in Artificial Intelligence 2004 and 2005. The research is headed by professor Michael Collins.

Realistic Graphics and Computational Photography

In the areas of realistic graphics and computational photography, a new theoretical framework has been developed that studies light transport from a signal-processing perspective. The frequency spectrum of the light array is modeled and analyzed to discover how it is modified by phenomena such as propagation in free space, occlusion, or glossy reflection. This approach has important potential applications in light simulation and inverse problems. New techniques have been developed to model the appearance of materials—that is, their texture and the way they reflect light. Unique high-resolution measurements of reflectance functions have been performed and made available to stimulate research in this area. In the domain of computational photography and video, a videoprocessing technique has been introduced that analyzes and magnifies small motions. A hardware and software solution has been designed that can extract a foreground object from a video sequence based on defocus effects. This research is headed by professor Fredo Durand.

Computational Modeling of Biological Shape

The study of biological shape and its variability has always played an important role in the life sciences. In human brain mapping, recent advances in magnetic resonance imaging have enabled noninvasive studies of human neuroanatomy at high resolution, yielding important insights into normal development and the pathologies of anatomical structures. Our research focuses on computational shape modeling and analysis that provide detailed, accurate characterization of shape variations and their interpretation in biologically meaningful terms of development and deformations.

In spite of the recent advances in computational approaches to modeling shape, many biological structures—including the exceedingly complex folding patterns of the cortical surface—still present a significant challenge for statistical shape analysis. We are developing novel shape representations that will lead to robust and efficient learning of shape variability from images. This research will enable statistical analysis of complex shapes in many different areas of biomedical imaging.

Together with our collaborators from the Harvard Medical School, we are validating the novel methods for shape analysis in the neuroimaging studies of normal aging and diseases, such as Alzheimer's and schizophrenia. In a more recent collaboration with a group from the Whitehead Institute, we are building on our techniques to characterize the shape and appearance of cells from high-throughput imaging data in the gene knockdown experiments. This project aims to identify the relevant pathways and the functionality of previously uncatalogued genes based on the changes in the cellular phenotype caused by the gene knockdown. Beyond biomedical imaging, our research offers improvements in modeling shape in many other fields that require understanding and characterization of 3-D shape from a set of examples, including computer vision, geology, and biochemistry. This research is headed by professor Polina Golland.

Databases

The research focus of this group is on sensor-related database applications, with a particular focus on building interfaces that facilitate collection of large amounts of data from widely distributed networks of sensors, particularly those connected via wireless and intermittent network connections. For example, in the CarTel project we are building a distributed wireless platform for automotive applications. One primary research focus is on providing a secure, easy-to-program environment that allows a wide range of automotive diagnostic and traffic-related applications to be rapidly built. Another research thrust involves integrating probabilistic and statistical models into databases of sensor information to detect unusual sensory events and predict future behavior. In CarTel, such models can be used to predict traffic at particular times on particular routes using historical data. This research is headed by professor Sam Madden.

Laboratory Sponsored Activities

K–12 Outreach

With the move to the Stata Center, CSAIL has become more visible to the local communities as a place to visit. We focus our attention on high schools, giving tours where 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 Tennessee visit to learn about the research we do. This past year we began teleconferencing with student groups who are unable to visit MIT. Our first was with a Canadian high school in Nova Scotia.

Middle East Education through Technology

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's primary goal is to provide an environment for Israeli and Palestinian youth in which societal barriers can become less rigid and friendships can be formed by sharing classes and learning a common computer language. Many of our students volunteer to teach MEET summer courses at the Hebrew University in Jerusalem.

MIT–Africa Internet Technology Initiative

CSAIL donated decommissioned computer equipment to the MIT–Africa Internet Technology Initiative (MIT–AITI), a student-run organization that gives students the opportunity to develop methods to bring IT to developing nations. These computers were installed in the Alliance High School and Laare Computer Center in rural Kenya. This effort was headed by Jack Costanza and Frank Tilley.

Museum of Science Discussion Series

On Saturday afternoons during the spring semester, CSAIL and the Museum of Science presented a series of popular audience talks geared specifically to children and parents visiting the museum. This presented an opportunity for our faculty to explain their research to very young children. Professors Rod Brooks, Polina Golland, and Erik Demaine and Dr. Stephanie Seneff participated in the series.

Professional Development

CSAIL sponsored a management training seminar entitled "Leadership for an Effective Research Environment." This seminar, which was offered to 12 junior faculty and research staff in CSAIL and the Research Laboratory for Electronics, was taught jointly by management consultant Chuck McVinney and professor Charles E. Leiserson. The seminar was designed to foster the nontechnical (human-centered) skills needed to develop and manage an effective research-group culture.

During the seminar, participants engaged in case studies of academic life, including student advising, rejection of papers from journals, and dealing with research sponsors. Activities ranged from role-playing games to various management "instruments" that give feedback on the effectiveness of situational decision making. In addition, group discussions addressed a variety of topics of concern to junior faculty, such as tenure, grants, consulting, staff supervision, and how to develop an effective professional statement.

Seminar Series

Four distinguished speakers gave presentations during this year's Dertouzos Lecture Series. They were Jon Kleinberg, professor of computer science, Cornell University; Robert Newton, dean of the College of Engineering at the University of California–Berkeley; Deborah Estrin, professor of computer science, University of California–Los Angeles; and Daphne Koller, associate professor of computer science, Stanford.

Awards and Honors

Our faculty and staff won many awards this year, including the following: Tim Berners-Lee was awarded the Great Briton Award in Science and Technology; Regina Barzilay, Erik Demaine, Fredo Durand, Dina Katabi, Sam Madden, and Robert Morris received NSF Career Awards; Alan Edelman received the 2005 Mathematical Association of America's Lester R. Ford Award (with Gil Strang); Shafi Goldwasser became a member of the National Academy of Engineering; John Guttag was admitted to the American Academy of Arts and Sciences; Frans Kaashoek became a fellow of the Association for Computing Machinery; Butler Lampson was elected to the National Academy of Sciences; Una-May O'Reilly was elected a fellow of the International Society of Genetic and Evolutionary Computation; Ron Rivest received the Massachusetts Innovation and Technology Exchange (MITX) Lifetime Achievement Award; Michael Stonebraker was awarded the IEEE John von Neumann Medal; and Gerald Sussman was elected a fellow of the American Association for the Advancement of Science.

Affirmative Action

CSAIL supports the affirmative action goals of the Institute.

Rodney Brooks
Director
Matsushita Professor of Robotics

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