University of Massachusetts Amherst

Winter 2012

Significant BITS

Newsletter of the
Department of Computer Science

Human-Machine collaboration for knowledge base construction

Andrew McCallum

Wikipedia's impact has been revolutionary. The collaboratively edited encyclopedia has transformed the way many people learn, browse new interests, share knowledge, and make decisions. Its information is mainly represented in natural language text. However, in many domains—from disaster recovery to bio-medicine—more structured information is useful because it better supports pattern analysis and decision-making.

Such structured information must usually be gathered and assembled from disparate sources. Sometimes this task is performed by humans, but it can be accomplished at much greater scales and speed by information extraction (IE), which automatically populates a database with relevant subsequences of text such as web pages, Twitter™ messages, and research articles. Although both humans and automated IE sometimes make mistakes, they have complementary strengths.

Professor Andrew McCallum sees rising interest in structured knowledge bases with Wikipedia-style breadth and collaborative derivation. His research strives to enable a similar revolution in the creation of such structured knowledge bases by enabling the robust integration of machine-provided information extraction from large text collections with human-provided edits by a diverse population of contributors.

For the past decade he has been working on information extraction and data mining. "Although information extraction and data mining appear together in many applications, their interface in most current systems would better be described as serial juxtaposition than as tight integration," says McCallum. "Information extraction is usually not aware of the emerging patterns and regularities in the database. Data mining methods begin from a populated database, and are often unaware of where the data came from, or its inherent uncertainties. The result is that the accuracy of both suffers," he adds.

McCallum's research has focused on probabilistic models that perform joint inference across multiple components of an information processing pipeline in order to avoid the brittle accumulation of errors. "The need for joint inference appears not only in extraction and data mining, but also in natural language processing, computer vision, robotics, and anywhere we must reason at multiple layers of abstraction," McCallum declares. "I believe that joint inference is one of the most fundamental issues in artificial intelligence."

continued on page 4

Benjamin Marlin joins CS faculty

Benjamin Marlin joined the department this fall as a tenure-track Assistant Professor. His research area is machine learning with a focus on the development and analysis of customized probabilistic models and approximate inference and learning algorithms.

"I'm interested in inference and prediction in all kinds of complex, high-dimensional systems subject to uncertainty and missing data," Marlin says.

continued on page 3

From the Chair

Lori A. Clarke

This is my first newsletter, since I only became department chair this fall (one semester down, 5 more to go—but who is counting?). This past semester has certainly been a learning experience for me. What have I learned? That the department chair receives a tremendous amount of email—it just keeps flooding in. That the UMASS bureaucracy is even more overwhelming than I anticipated—and I expected it to be overwhelming. But there are good things too. I knew that the department was extremely collegial and that the faculty work hard to share the load, but now I have a much better appreciation of how much needs to be done to keep all aspects of the department moving forward. And yes, we always are trying to move forward. Many faculty step up to take on major service jobs — in addition to their research and teaching loads.

And the commitment and professionalism of the staff is amazing. With the departure of Claire Christopherson as our Business Manager (although she is still available to provide much needed advice and guidance) and the retirement of Gwyn Mitchell as a Grant Administrator (although she too provides support as a post-retiree), this has been a difficult semester. But the staff has really pitched in to help cover the load while we look for people who can grow into these positions. Even Sharon Mallory, our past office manager, temporarily came out of retirement to help out. What an amazing place!

The department is embarking on a number of new initiatives. Some are still under wraps, waiting for final approval (more about this in future newsletters). Last year we officially started a new B.A. degree in computer science to complement our B.S. degree. We envision the B.A. degree serving undergraduates who want to combine computer science with studies in non-science disciplines, recognizing the emerging interdisciplinary nature of computing.

We also have started a terminal master's degree program. Previously, we only accepted graduate students who we thought would pursue a Ph.D. degree. With the increasing demand for master's degree students, we now accept students who most likely will be seeking industrial jobs upon graduation, although some are interested in improving their academic credentials before applying to Ph.D. programs in computer science.

We also have been trying to improve the overall environment for our computer science undergraduates. We now have a U-Space in the computer science building where our undergraduates can hang out in between classes. There is usually a monthly First Friday undergraduate event. This last month, there was a Wearable Electronics event where undergraduates could learn how to program wearable processors as well as imaginatively sew them onto their clothing (I hear the sewing was harder than the programming!). The ACM student chapter has also been very active. We had three programming teams compete in the ACM collegiate programming contest, with one team going on to compete in the Northeast North Regional Division.

Finally, let me end with a thank you to Rick Adrion, Jim Kurose, Bruce Croft, and Andy Barto. As "recent" past department chairs, I frequently call on them for advice. Our department has been extremely fortunate to have had such good stewardship over the years. I am serving a department that is in excellent shape. Although external funding has been on the decline nationally, our department funding has been increasing, our undergraduate enrollments have been growing significantly, and our graduate admission applications are rising. And although we had three retirements this year, most of our retirees continue to be research active. Moreover, we are currently searching for four new faculty positions. And next newsletter, I hope to be able to announce some more exciting news.

J.A.N. Lee gives keynote address

Dr. J.A.N. Lee, the department's first chair, returned to campus to give a keynote address in October, 2011 as part of the "50 Years of Computing Celebration" sponsored by the campus' Office of Information Technologies (OIT). He gave the address, "50 Years of Computing: From Humongous to Ubiquitous," to a large crowd at the Student Union Ballroom.

During his talk, Lee spoke about the CS program getting its first computer, a DEC PDP-1, and the difficulties in getting the Ph.D. Program and an Undergraduate Program approved for computer science at UMass Amherst. He discussed much of the history of the department and the field of computing.

Lee was the Head of the Computer Science Program from 1964-1969 (when it was an M.S. Program within the Graduate School) and also a director of the campus' Research Computing Center from 1964-1965.

MARLIN – – – – – – – – – – – *continued from page 1*

"Basing machine learning methods on probabilistic models means that we can explicitly reason about various sources of uncertainty within a consistent, principled framework." Marlin's research draws heavily on areas of both computer science and statistics. "Things get interesting," says Marlin, "when statistical theory collides with computational complexity." According to Marlin, models that attempt to capture the salient features of complex systems nearly always force a trade-off between statistical correctness and computational efficiency, particularly when scaling to massive data sets. Marlin's research addresses these issues through the study and application of several different approximation frameworks including variational algorithms, Markov chain Monte Carlo methods, and non-likelihood-based model estimation.

Marlin's doctoral research was in the area of recommender systems. "In a recommender system like Netflix," says Marlin, "users provide a limited number of ratings for items like movies, and the goal is to leverage rating data from the whole community of users to make personalized predictions and recommendations for each individual user." The underlying system involves thousands of unique items and the preferences of thousands of people. According to Marlin, this makes it extremely complex, high dimensional, uncertain, and incomplete. Marlin's work showed how the properties of missing data in these systems can invalidate the assumptions required for correct estimation of standard statistical models, leading to inaccurate predictions. He proposed an extended modeling framework for mitigating this problem, resulting in significantly improved prediction and ranking performance.

Marlin's recent work spans a range of areas in addition to recommender systems, including unsupervised structure discovery and feature induction, medical informatics, and, most recently, topic modeling. According to Marlin, one of the strengths of machine learning is that it provides a foundation for effective problem solving that easily extends across domains. "There are surprisingly deep connections between the models and algorithms that are applied to seemingly disparate problems like recommending movies, analyzing electronic health care records, and inferring the semantic structure of document collections," Marlin says.

The possibility for collaboration in multiple areas was one of the factors that most impressed Marlin about the Computer Science Department at UMass Amherst. He recently became affiliated with the Center for Intelligent Information Retrieval and is pursuing early-stage collaborations with a number of other groups including the Information Extraction and Synthesis Laboratory, the Sensor Networks Research Group, and the division of Health Informatics and Implementation Science at the UMass Medical School.

Prior to joining the department, Marlin was a Postdoctoral Fellow in the Laboratory for Computational Intelligence at the University of British Columbia (UBC). Marlin received a B.S. degree in Mathematics and Computer Science from McGill University in 2002. He received M.S. and Ph.D. degrees in Computer Science from the University of Toronto in 2004 and 2008, respectively. During graduate school, Marlin spent a year as a Research Intern at Yahoo! Research.

Marlin's undergraduate research project conducted at McGill University won the Best Paper Award at the 2002 Canadian Conference on Computational Geometry. While at the University of Toronto, he held a Canadian Graduate Scholarship from the Natural Sciences and Engineering Research Council of Canada, the top Canadian doctoral scholarship in the sciences. Marlin received the Best Paper prize at the ACM Conference on Recommender Systems in 2009 and an invitation to the Best Papers Track at the International Joint Conference on Artificial Intelligence in 2011. At UBC, Marlin held a Killam Postdoctoral Fellowship as well as a fellowship from the Pacific Institute for the Mathematical Sciences.

Levine receives research award

At the "State of the College of Natural Sciences" celebration held on September 6th, Professor Brian Levine received the 2011 Outstanding Research Faculty Award. He was recognized for his critical contributions to the understanding of networking and security.

Levine's research specialization includes the areas of mobility and forensics. In the area of security, his research is focused on digital forensics and privacy. One example is his investigation of child sexual exploitation on peer-to-peer networks. He and colleagues developed powerful new software, which is being used by 58 out of 61Task Forces on Internet Crimes Against Children (ICAC) around the nation, to collect evidence against people who possess and share illegal images or produce child pornography for the Internet.

In the networking area, his research focuses on mobile networking, including disruption tolerant networks and peer-to-peer networking. Levine's research includes the construction of DieselNet, a long-running mobile network research testbed comprised of computer-equipped buses and outdoor wireless access points.

Levine is the Director of the Commonwealth Center for Forensics & Society, a cooperative between UMass Amherst and the MA State Police Crime Laboratory. He is also the Undergraduate Program Director in the department. Levine was named a UMass Amherst Lilly Teaching Fellow in 2003 and received the College Outstanding Teacher Award in 2007. In 2008, he received the Alumni Award for Excellence in Science and Technology from the University of Albany.

HUMAN-MACHINE COLLABORATION – *cont. from page 1*

His lab has built up a series of successes showing the benefits of joint inference. In 2005 his group won a KDD competition in entity resolution with new research on a conditional random field that jointly accounts for multiple pairwise compatibilities. In 2008 his group obtained new state-of-the-art results on benchmark tasks in ontology- and schema-matching using a joint model trained by their new Metropolis-Hastings-embedded parameter estimation method, SampleRank. In 2011 Postdoctoral Fellow Sebastian Riedel and McCallum won multiple tracks of the BioNLP competition in collaboration with researchers from Stanford who provided features from their parser; the key to this success was the use of dual decomposition for joint inference among multiple constraints on the protein interaction events extracted.

"These results demonstrate joint inference being performed among relatively small sets of variables, but I think the biggest gains will come from joint inference across document boundaries or across entire databases," says McCallum. "Reasoning about data at this scale quickly involves more statistical random variables than can fit in machine memory, however." This problem has led to McCallum's more recent interest in probabilistic databases. "We would like to use database technology not just for storing and querying the results of an IE system, but also for performing IE joint inference itself—managing the many random variables and intermediate results of IE," he adds. In a 2010 VLDB paper, graduate student Michael Wick and McCallum describe just such an approach in which raw textual evidence is presented to the database, and IE inference is performed 'inside the database' using Markov-chain Monte Carlo. "We have taken to calling this an Epistemological Database, indicating that the database doesn't observe the truth; it must infer the truth from available evidence," McCallum explains.

'Truth-discovering' inference continues to run in these database systems as new evidence arrives. New evidence can correctly cause the database to change its mind about its previous conclusions. Since the probabilistic database maintains IE's intermediate results, the system is able to re-visit targeted portions of inference without re-running it from scratch. This is an extremely practical feature since in IE it is common for new documents to arrive in a stream over years of operation. Evidence for the database may include not only additional documents but also new structured records and additional partial databases to be integrated.

This approach also leads directly to a compelling approach for handling human edits. Some traditional approaches use the edits to directly modify the database's notion of 'the truth.' But this is deeply unsatisfying for several reasons: sometimes human edits will be wrong; sometimes humans disagree; and sometimes a correct human edit should be overwritten by IE after the passing of time has made the old edit no longer valid. "A better approach is to model human edits as additional 'mini-documents' (for example 'User X said Y is true on April 2') to be treated as evidence and reasoned about. In this framework we can perform probabilistic reasoning not only about the IE process, and which human edits to incorporate, but also simultaneously about reliability and reputation of the human editors themselves," says McCallum.

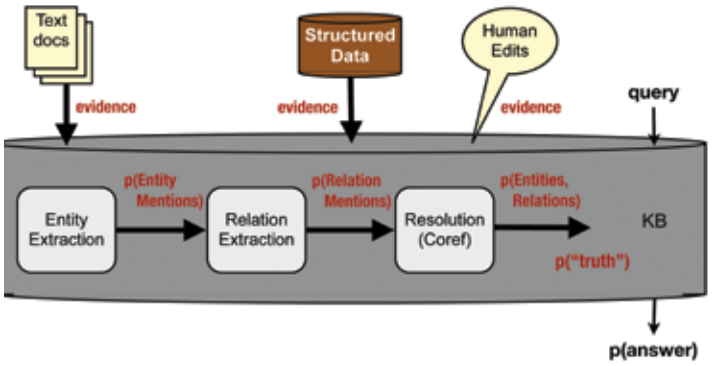

An Epistemological Database doesn't directly receive entities and relations; it infers them internally from available evidence, which may arrive as a stream over time.

McCallum is currently applying these ideas to the construction of knowledge bases of the scientific research literature, with the aim of improving scientific collaboration. "As science becomes more interdisciplinary and complex, it becomes all the more necessary to have tools to manage the intellectual landscape of scientific ideas and collaborations," says McCallum. "By building systems that gather and integrate information about scientists and their work, we can provide tools that will help scientists find collaborators, understand the relations of their work to neighboring scientific fields, translate vocabulary among fields, summarize trends, understand emergence of ideas, know which papers to read, find new scientific sub-areas for fruitful investigation, and identify good candidates for hiring students, postdocs, and faculty. So many of these decisions are currently based on myopic local views and serendipity. We hope to accelerate the rate of scientific progress by providing better tools."

Preliminary results of this work can be found at **rexa.info**. His work on probabilistic programming that supports this research can be found at **factorie.cs.umass.edu**.

Professor McCallum received his Ph.D. from the University of Rochester in 1995. He was a postdoc at Carnegie Mellon University and, later, Vice President of Research and Development at a 170-person start-up company focusing on information extraction from the web. Since 2002 he has directed the Information Extraction and Synthesis Laboratory at the University of Massachusetts Amherst. In 2009 he was named an AAAI Fellow. He has authored over 200 papers in multiple areas of artificial intelligence; his work has received over 22,000 citations.

Wendy Lehnert retires

Professor Emeritus Wendy Lehnert retired in September after a twenty-nine year career as faculty in the UMass Amherst Department of Computer Science.

"I've known Wendy since my days at MIT and hers at Yale. We collaborated on many projects at UMass Amherst, having received one of the department's first DARPA awards to research a system that combined natural language processing (NLP) and case based reasoning (CBR) in an advisory system, followed by other DARPA CBR/Machine Learning (ML) initiatives. We used these ideas in the context of information retrieval as co-PI's with Bruce Croft on the NSF Center for Intelligent Information Retrieval grant," says Professor Edwina Rissland. "Wendy has been one of the most creative researchers in AI and also a wizard-class programmer. She was at the forefront of using ML in NLP. Of her many contributions, one of my favorites is her beautiful work on 'Plot Units.' Long recognized as a leader in AI, she was one of the founding AAAI Fellows."

Lehnert received her B.A. in Mathematics from Portland State University in 1972 and an M.A. in Mathematics from Yeshiva University in 1974. She received her Ph.D. in Computer Science from Yale University in 1977 for her work on a computational model of human question answering. She subsequently joined the faculty at Yale where she held a joint appointment in Computer Science and Psychology. During 1981-82, Lehnert held the position of Vice President for Research at Cognitive Systems, Inc., where she designed commercial natural language processing systems and developed applications for the existing technology.

In 1982, Lehnert joined the tenure-track faculty in the Department of Computer Science at UMass Amherst where she specialized in natural language processing and cognitive models of human thought processes. In the 1990's, Lehnert was a pioneer in the area of information extraction from text and natural language processing performance evaluations. In recent years, Lehnert taught web-related courses and authored a number of books on the topic, including *Light on the Web* and *Web 101*.

In 1984, Lehnert received the Presidential Young Investigator Award from the National Science Foundation in recognition for her work in artificial intelligence, and in 1991 she was elected Fellow of the American Association for Artificial Intelligence. In 1996, she was honored as a UMass Amherst Distinguished Faculty Lecturer. She has served on scientific advisory boards for the National Science Foundation, the National Library of Medicine, and has served as a member of the Information Science and Technology Committee for the Advanced Research Projects Agency of the Department of Defense. Lehnert was elected to the Board of Counselors for the American Association for Artificial Intelligence, as well as the Governing Board for the Cognitive Science Society, and she has also served as a senior editor for *Cognitive Science*. She has published nine books and over 100 journal articles, conference papers and book chapters.

In addition to her CS research, Lehnert has been investigating traditional Chinese medicine as a model of memory-intensive reasoning and problem solving. She is licensed to practice acupuncture and Chinese herbal medicine in the state of Massachusetts.

CAITE expands its role in broadening participation in computing

This year, the Commonwealth Alliance for Information Technology Education (CAITE) has supported more than fifty-five outreach events aimed at recruiting women and underrepresented populations into computing fields. Led by Professor Rick Adrion and based in the department, CAITE is a National Science Foundation-sponsored project aimed at broadening participation in computing.

In October, CAITE was a co-sponsor of the Women in Engineering Diversity Program office's annual Women in Engineering & Computing Career Day, which brought 250 high school students and 50 educators to the campus for a day of hands-on activities, meetings with working engineers, an alum panel, and a keynote address from Cissy George, national engineering director for Verizon Communications. Assistant Professor Hanna Wallach talked to educators about preparing students for the CS program at UMass Amherst, while Ph.D. alumna Rachel Cobleigh and current CS major undergrad Jessica Ray told students about their experiences as women studying and working in the field. CAITE Project Manager Renee Fall also spoke with educators and guidance counselors about the importance of presenting computing careers as promising for young women, and she provided resources CAITE has developed and gathered to help teachers spread this message. Several students reported that the program inspired them to investigate potential careers in STEM (science, technology, engineering, and math) fields.

In another event held this fall, 200 high school students and 50 teachers visited campus for Science Quest, a day of demonstrations and panels organized by UMass Amherst faculty, staff, and graduate students and cosponsored by CAITE. CS Department graduate students Abhinav Parate, John Bowers, and Andres Molina-Markham each held

continued on page 7

Clarke honored with Chancellor's Award

On September 16, 2011, during the Seventh Annual Faculty Convocation ceremony, Chancellor Robert Holub presented Professor Lori Clarke with the Award for Outstanding Accomplishments in Research and Creative Activity. Clarke was one of seven nationally and internationally acclaimed faculty members presented with the award. She was recognized for her pioneering software engineering research and for her leadership in the field of computer science.

Her research focus is software engineering, particularly on software analysis and testing. She was one of the primary developers of symbolic execution, a technique used to reason about the behavior of software systems and to select test data, and she made contributions in the areas of software architecture and object management. Recently she has been applying software engineering techniques to processes in various domains, such as health care and elections, to help find errors or vulnerabilities.

Clarke is the co-director of the Laboratory for Advanced Software Engineering Research. She is currently the chair of the department and a board member of the Computing Research Association's Committee on the Status of Women in Computing Research (CRA-W). She is a Fellow of the ACM and of the IEEE, as well as a former vice chair of the Computing Research Association (CRA) and co-chair of the CRA-W. She received a 2009 College Outstanding Faculty Service Award, a 2004 Distinguished Engineering Alumni Award from the University of Colorado, Boulder, the 2002 ACM SIGSOFT Distinguished Service Award, a 1993 University Faculty Fellowship, and a 1991 UMass Amherst Chancellor's Medal.

Fu named to NIST Advisory Board

Associate Professor Kevin Fu was recently appointed to serve a four-year term on the National Institute of Standards and Technology (NIST) Information Security and Privacy Advisory Board (ISPAB).

"I look forward to serving on the Information Security and Privacy Advisory Board," says Fu. "One of my roles will pertain to security and privacy issues that affect health information technology and medical devices."

ISPAB is a Federal Advisory Committee that brings together senior professionals from industry, government, and academia to help advise NIST, the U.S. Office of Management and Budget (OMB), the Secretary of Commerce, and appropriate committees of the U.S. Congress about information security and privacy issues pertaining to unclassified federal government information systems. The Board reports its findings to the Secretary of Commerce, the Director of OMB, the Director of the National Security Agency, and the appropriate committees of Congress. ISPAB was originally created by the Computer Security Act of 1987 as the Computer System Security and Privacy Advisory Board.

Past computer science members of the Board include Susan Landau of the Department of Computer Science at Harvard University, Howard Schmidt, now the White House Cyber Security Chief, Fred Schneider of the Department of Computer Science at Cornell University, and Ari Schwartz, now in the Office of the Secretary at the Department of Commerce.

Members of this board are known for their past involvement on issues of export control, standardization of cryptographic ciphers, government key escrow, and numerous

CS Women's group selected for national grant

The National Center for Women and Information Technology (NCWIT) has selected the department's Computer Science Women's Group program to receive a Return Path Student Seed Fund grant. The CS Women's group will use the funding to support visits to academic and industrial institutes, to support their Grad-to-Undergrad Mentoring Program, and for activities that increase the awareness of the Computer Science Department for recruiting women to the program.

According to NCWIT, the NCWIT Academic Alliance has partnered with Return Path to offer the NCWIT Student Seed Fund, given to student-run programs and initiatives that promote increased participation of women in computing and IT programs. Return Path is a firm that focuses on boosting the deliverability of email messages for clients.

The award is given to organizations whose initiatives align with NCWIT's mission by proposing activities, events, or other programs that will recruit, retain, and support women in technology- and computing-related majors.

Victor Lesser retires

Distinguished Professor Emeritus Victor Lesser retired in September after thirty-four years of service within the Department of Computer Science at UMass Amherst. While Lesser is now officially retired, he still continues to direct the Multi-Agent Systems (MAS) Laboratory within the department.

"Victor has had an amazing research career. He is widely recognized as the founding father of multi-agent systems," notes Department Chair Lori Clarke. "Our research areas are very different, but the one time we were at the same DARPA meeting, it was clear that Victor was the rock star of the multi-agent systems community. I think everyone is in awe of the number and quality of his Ph.D. students. He has an incredible legacy."

Lesser received an A.B. in Mathematics from Cornell University in 1966, and an M.S. and Ph.D. in Computer Science from Stanford University, in 1970 and 1973, respectively.

Lesser's major research focus is on the control and organization of complex AI systems. He has pioneered work in the development of the blackboard architecture and its control structure, approximate processing for use in control and real-time AI, and a wide variety of techniques for the coordination of, and negotiation among, multiple agents. He was the system architect for the first fully developed blackboard architecture (HEARSAY-II), when he was a research scientist at Carnegie Mellon University from 1972 through 1976. He has also made contributions in the areas of machine learning, signal understanding, diagnostics, plan recognition, and computer-supported cooperative work. He has worked in application areas such as sensor networks for vehicle tracking and weather monitoring, speech and sound understanding, information gathering on the internet, peer-to-peer information retrieval, intelligent user interfaces, distributed task allocation and scheduling, and virtual agent enterprises.

Lesser received the prestigious International Joint Conferences on Artificial Intelligence (IJCAI-09) Award for Research Excellence. He is also a Founding Fellow of AAAI and an IEEE Fellow. He was General Chair of the first International Conference on Multi-Agent Systems (ICMAS) in 1995, and Founding President of the International Foundation of Autonomous Agents and Multi-Agent Systems (IFAAMAS). In 2007, to honor his contributions to the field of multi-agent systems, IFAAMAS established the "Victor Lesser Distinguished Dissertation Award." He also received a Special Recognition Award for his foundational research in generalized coordination technologies from the Information Processing Technology Office at DARPA.

Lesser received the UMass Amherst College of Natural Sciences and Mathematics Outstanding Teaching Award in 2004 and the Outstanding Research Award in 2008. He also received the Chancellor's Award for Outstanding Accomplishments in Research and Creative Activity in 2008.

This past spring, Lesser's former graduate students created an endowed scholarship in his honor to acknowledge his extraordinary career and the mentoring of over 30 MAS graduate students.

CAITE – – – – – – – – – – – *continued from page 5*

Students at Science Quest

short workshops, and CAITE Project Manager Renee Fall spoke with students about opportunities for CS and IT education throughout the day.

CAITE is also working to support community college students, often an older and more diverse population, to transfer and complete bachelors' degrees in computing fields. Pathways Coordinator Cheryl Kiras has been working with admissions staff and faculty from UMass Amherst and Massachusetts community colleges to develop course equivalencies and roadmaps to transfer for students from Bunker Hill, Greenfield, Springfield Technical, and Bristol community colleges. Holyoke Community College is in the process of developing a "Roadmap to transfer to UMass Amherst," and Kiras is working to create similar pathways for Mass-Bay, North Shore, and Quinsigamond community colleges in central Massachusetts. This work by CAITE, along with the campus's Community College Connections program announced last spring, is showing results. Community college transfers into the UMass Amherst CS Department have increased 183% from 2007 to 2011, and 30% of the incoming 2011 students were women and underrepresented minorities.

Salajegheh a top finisher in Innovation Challenge

The team of graduate student Mastooreh (Negin) Salajegheh and her mentors, Professors Kevin Fu and Wayne Burleson, was one of the top four finishers in the Executive Summary and Elevator Pitch stage of the UMass Amherst Innovation Challenge Competition held on December 7th. The team, named SMASH (for **sm**arter f**lash**), introduced techniques to reduce the power consumption of battery-powered devices through smarter flash memory and a zero-power clock (named TARDIS for **T**ime **a**nd **R**emanence **D**ecay **i**s **S**RAM).

Of the 10 teams competing, 4 teams received $1,750 awards; the remaining 6 received $500. The final phase of the academic year-long contest takes place in April, when finalists will compete for more than $50,000 in grand prize money during the Innovation Challenge Final Business Plan portion of the competition.

The SMASH memory uses counterintuitive software techniques to reduce the energy consumption of flash memory twofold by exploiting the electrically cumulative property of floating gate transistors in flash memory cells. The voltage requirements of flash memory typically dictate the overall energy consumption of embedded systems such as smoke detectors and energy harvesting platforms that rely on microcontrollers. The SMASH memory has won the Commercial Ventures and Intellectual Property (CVIP) Technology Fund Award and its patent is pending.

Another problem the team addressed is the need for a zero-power clock. “In a desktop computer, time keeping is easy because one can power a clock. On a smartcard, one must rely on the clock from an untrusted source (e.g., a radio transmission). Without a sense of time, a smartcard remains vulnerable to brute force attacks,” says Salajegheh. “For example, the card would have no idea if an attacker tried a password 1,000,000 times or 1 time. Each power up of the card is stateless and independent.”

Mastooreh (Negin) Salajegheh

The team’s TARDIS technique uses the behavior of volatile, sequential logic to estimate the duration of a power failure. One materialization of the TARDIS is an SRAM clock that estimates the duration of a power-down by measuring the fraction of cells that retain their state across the power failure. The zero-power, approximate time-keeping enabled by the TARDIS promises to expand the use of RFID-scale, transiently-powered computing devices into a variety of applications that require approximate timing information to operate correctly in the face of frequent power failures. The researcher’s initial experiments on a TI MSP430 microcontroller indicate that an SRAM clock can reliably estimate the time of power failures up to 6~minutes in duration. This technology is also patent pending.

Salajegheh is a Ph.D. candidate in the department who works with Fu as a member of the Security and Privacy Research Laboratory (SPQR). Her research focuses on low-power operation of pervasive computers, energy management, probabilistic storage, and trustworthy computing. Mastooreh has received the Outstanding Synthesis Project Award in 2010. She has published or presented papers in ACM *Transactions in Embedded Computing Systems*, USENIX Security, USNEIX FAST, Microsoft Research, and Intel Research. Mastooreh spent the last summer at Microsoft Research working on peer-to-peer communication through Near Field Communications.

Department names Outstanding Dissertation Award recipients

Arruna Balasubramanian

Timothy Wood

Aruna Balasubramanian and Timothy Wood are this year’s UMass Amherst Computer Science Outstanding Dissertation Award Recipients. They will receive their awards at the 2012 awards banquet to be held on May 4, 2012

Balasubramanian, advised by Profs. Arun Venkataramani and Brian Levine, received her Ph.D. in February 2011. She is currently a CIFellow in the Department of Computer Science at the University of Washington. Her Ph.D. thesis title is “Architecting Protocols to Enable Mobile Applications in Diverse Wireless Networks.”

Wood, advised by Prof. Prashant Shenoy, received his Ph.D. in September 2011. He is currently an Assistant Professor in the Department of Computer Science at George Washington University. His Ph.D. thesis title is “Improving Data Center Resource Management Deployment, and Availability with Virtualization.”

David W. Stemple Scholarship in Computing established

Jane Yolen has created a new scholarship in loving memory of her husband, Emeritus Professor David Stemple. Once endowed, the fund will provide an annual David W. Stemple Scholarship in Computing to a first-year graduate student, with priority given to women, to help attract talented Computer Science graduate students in the systems research area.

Stemple, who died in 2006 after a long battle with cancer, was a member of the department's faculty from 1981 until his retirement in 1998. His main research dealt with developing complex database systems. He, along with students Tim Sheard, Subhasish Mazumdar, and Leo Fegaras, concentrated on the specification of database systems and the mechanical reasoning that is required to assure that specifications satisfy their intended properties.

Stemple was involved, starting in 1960, in the development of compilers, operating systems, and database management systems in industry. Then, starting in 1966, he began working at the computing center at the University of Massachusetts Amherst. He was associate director of the computing center from 1971 to 1981, and in 1977 he received his Ph.D. in in Computer Science from UMass Amherst working with Robert W. Taylor. Stemple joined the department as a faculty member in 1981. He served as Chair of the department from 1994 through 1998. He was also an Honorary Professor of Computer Science at the University of St Andrews.

After retiring, Stemple worked as an ornithologist concentrating on the song behaviors of certain montane thrushs from the Andes to the Alps.

To contribute to the fund for the *David W. Stemple Scholarship in Computing* at UMass Amherst, gifts may be made online (**www.cs.umass.edu/about/donate**), by calling 413-545-4721, or by check made out to UMass Amherst with designation to the David W. Stemple Scholarship in Computing and mailed to the Department of Computer Science, 140 Governors Drive, UMass Amherst, Amherst, MA 01003-9264.

Fourth Annual Outstanding Achievement and Advocacy Awards Banquet planned

The fourth annual Outstanding Achievement and Advocacy (OAA) Awards Banquet will be held on the evening of Friday, May 4, 2012 in the UMass Amherst Campus Center. During the banquet, awards will be presented to recognize the remarkable accomplishments of graduates of the department's degree programs and to acknowledge the support of important friends of the department. Current student awards will also be presented. For details and to register for the event, go to: **www.cs.umass.edu/oaa2012**.

Over 85 people joined us in 2011 at the third OAA Awards Banquet to help celebrate the accomplishments of our award recipients and to reconnect with fellow alums.

First place in regional competition

Two UMass Amherst CS undergrad teams placed first and third in the Northeast Regional Preliminary Competition of the annual ACM International Collegiate Programming Competition. Eighteen teams participated in this round. Coached by Professor Erik Learned-Miller, the first place team consists of Yue Wang, Evan Shelhamer, and William (Bill) Crane (shown below, l.to r.). The students on the third place team are Anthony Moh, Khanh Nguyen, and Tung Pham. The first place team traveled to Rochester, NY on November 5, 2011 to participate in the Northeast North American Regional Finals. The team placed fourth, behind MIT, Harvard, and the University of Rochester.

Bruce MacLeod develops cell phone technology for health care delivery in Ghana

Bruce MacLeod (Ph.D.'89)

High quality and timely health information is not readily available to pregnant women and new mothers in the Upper East region of Northern Ghana. Internet access is virtually non-existent, literacy rates are low, and travel distances to clinics can be significant. Yet the need for health information is especially important. Malaria remains the major cause of death among under-5 children, and proper use of intermittent preventative treatment during pregnancy can reduce anemia, fetal loss, and dangerous complications for the unborn child. Local beliefs and superstitions often guide women when information about best practices is not available. In the Upper East, elder female relatives encourage younger pregnant women to labor at home for as long as they can stand before going to a health facility, as a mark of strength and willpower. When women do come to a health clinic, they are given a booklet of basic health information for the pregnancy; unfortunately, most women are unable to access the information because they are illiterate.

The combination of poverty, remote villages, low literacy, and widespread disease creates a health crisis for women and children all over Africa. Lack of information, and bottlenecks in the gathering of data and transfer of information, contribute to the negative health outcomes. Although the Ghana Health System coordinates local volunteer health workers, outreach nurses, clinics, and district health centers to provide care and track the success of interventions, information gaps and challenges pervade the health system.

"Yet cell phone coverage for this remote and rural society is expanding rapidly. In Ghana, the coverage even includes GPRS capabilities which allow communication between cell phones and internet servers. Most of the women in the Upper East of Ghana have mobile phones and our baseline research found that only 16% of women had no access," says Bruce MacLeod, Software Architect on the Mobile Technology for Community Health initiative (MoTeCH), and Associate Professor of Computer Science at the University of Southern Maine.

The team's MoTeCH project attempts a systemic change in Ghana's existing health information system through integrating mobile phones into the rural health system. Pregnant women and new mothers with access to a phone can get relevant, personalized health information from MoTeCH. Nurses are given a low-cost mobile phone for entering data and receive MoTeCH messages about patients overdue for health care. The data entered by nurses is also used to automate the generation of reports for district and regional health administrators.

To enable this functionality, electronic medical records for each women and child in the study area are kept. Basic information about health service delivery and health status are maintained in these records. From this data, MoTeCH calculates the upcoming schedule of recommended care for each patient. For example, schedules of immunization shots are developed for each newborn child and upcoming clinic visit dates are developed for each pregnant woman. When care is due, MoTeCH notifies patients and health workers by mobile phone. As new information comes into the medical record, the schedule of upcoming care for a patient may be updated. While using electronic medical records in this rural setting may seem quite a far reach, it actually solves significant and long standing communication and information transfer problems in a relatively simple and low cost manner.

A pregnant woman who will be receiving MoTeCH messages

The underlying software architecture of MoTeCH evolved from a strategic decision to build on "best-of-class" software applications that have been extensively field tested in Africa, adds MacLeod. The mobile phone data collection system is based on openXData (**www.openxdata.org**). Health information about women and children is maintained in the OpenMRS medical record software application (**www.openmrs.org**). Finally, the task of delivery voice messages to patients was provided by the IntelliIVR (**www.yo.co.ug**) application. Significant project risk was eliminated because these systems have gone through multiple iterations of deployment and revisions in environments similar to those in NorthEast Ghana.

"Developing interoperability logic between these systems was one of the primary software development challenges we faced," says MacLeod. "Data collected by nurses using mobile phones needed to be uploaded into the appropriate places in a medical record. The timing of phone calls to patients using the IVR system is based on data from the electronic medical system. When patients do not attend the clinic within the recommended time frame, which can be calculated from the data in the medical record, text messages are sent to the nurse's mobile phone."

MacLeod's software development work was part of a larger collaboration between Columbia University, the Grameen Foundation, and the Ghana Health Service. The Gates Foundation funded the project with the goal of understanding how mobile phones could be used to achieve measurable health impacts. The system has been functioning since the beginning of the year in the Upper East region of Ghana. Preliminary reactions from patients receiving MoTeCH messages have been generally positive and

A community health officer entering data

nurses are becoming more proficient in using the technology. Ultimately, understanding the full impact and potential of a system like MoTeCH requires a multi-dimensional evaluation, examining the intervention effects on health workers and patients. This evaluation is underway.

"I have had the privilege of working on software systems and health problems in Africa and Asia since I was a graduate student at UMass Amherst," says MacLeod. "My first project sent me to the Sudan in 1984 to work with the Ford Foundation. That was my introduction to the simple joy of working with committed research scientists who are striving to improve health care in these rural communities."

MacLeod has over two decades of experience in developing appropriate software technology for health research in the developing world and in making complex software systems accessible to health researchers who are non-computer scientists. The MoTeCH system is the software standard for INDEPTH network of health research centers (**www.indepth-network.org**). Advised by Prof. Robbie Moll while in the department, MacLeod received an M.S. and Ph.D. in Computer Science from UMass Amherst in 1985 and 1989, respectively.

A Kassena-Nankana extended family compound

In December, The Federal Communications Commission (FCC) appointed **Henning Schulzrinne** (Ph.D. '92) as Chief Technology Officer. He will advise on matters across the agency to ensure that FCC policies are driving technological innovation, including serving as a resource to FCC Commissioners. Schulzrinne is Julian Clarence Levi Professor of Mathematical Methods and Computer Science and Professor of Engineering at the Fu Foundation School of Engineering at Columbia University. He has been an Engineering Fellow at the FCC since 2010. This fall, Schulzrinne also received the IEEE William Terry Lifetime Distinguished Service Award.

A number of CS alums and current faculty had their publications included in the recent list of SIGSOFT Top 10 Most Cited Articles. **Alexander Wolf** (Ph.D. '85) and co-author were listed #1 for "Foundations for the study of software architecture," **James Corbett** (Ph.D. '92), **Matthew Dwyer** (Ph.D. '95), and co-authors were second for "Bandera: extracting finite-state models from Java source code," **Peri Tarr** (Ph.D. '96), former Sr. Research Fellow **Stanley Sutton**, and co-authors were third for "N degrees of separation: multi-dimensional separation of concerns," Prof. **Leon Osterweil** was seventh for "Software processes are software too," and **Matthew B. Dwyer**, Professor **George S. Avrunin**, and **James C. Corbett** were eighth for "Patterns in property specifications for finite-state verification."

CIIR Alum **Giridhar Kumaran** (Ph.D. '08) and his wife Ramya Vaidyanathan announced the birth of their son Pranav Giridhar, born on September 1st.

This fall, **Mario Barrenechea** (B.S. '11) was on the winning team of the Winward Code War Contest held at the University of Colorado Boulder where he is currently a graduate student.

Mimno named 2011 CIFellow

David Mimno (Ph.D. Feb. '12) was named one of the twenty Ph.D. grads to the Computing Community Consortium (CCC) 2011 class of Computing Innovation Fellows (**cifellows.org**). For his fellowship, Mimno will hold a post-doc position at Princeton University working with mentor David Blei.

According to the CCC, "these 20 talented researchers have been competitively awarded postdoctoral positions of up to two years at academic institutions and industrial research laboratories throughout the U.S. Made possible by a $6.5 million National Science Foundation (NSF) grant to CRA - the third grant in as many years - the 2011 CIFellowships are a continuation of the highly successful effort begun in 2009 to forestall a permanent loss of research talent likely to occur as a consequence of the financial crisis and subsequent economic downturn."

Wearable Electronics Night

The first UMass Amherst CS Wearable Electronics Night was held in December for the department's undergrads. Organized by Professors Rui Wang and Deepak Ganesan, and sponsored by Cisco Systems, the event was attended by over thirty students. It was designed to be a fun, artistic, and creative exercise where students could combine computing, electronics, and craft to design electronic textiles or e-textiles. E-textiles are soft electrical circuits created using flexible conductive materials in conjunction with electronic components such as LEDs, sensors, microcontrollers. Thanks to support from Cisco, the department provided students with the basic ingredients for creating the wearable electronics. Prof. Wang gave a tutorial on programming these devices and creating a simple e-textile. Students brought clothing or accessories that they wanted to convert into a wearable electronic. By all accounts, it was a fun and enjoyable night where students got to learn and experiment with building circuits, and quite a few participants created impressive designs. A week later, a follow-up event was organized by undergrad Evan Shelhamer to help students complete their e-textile creations, particularly the programming part. Along with the faculty organizers, graduate students Heather Conboy, Shiri Dori-Hacohen, Jackie Feild, Lisa Friedland, and Bobby Simidchieva were available at the event to help students create the wearable electronics. The event was so successful that another event is planned for the spring semester.

Christopher Amato; *Increasing Scalability in Algorithms for Centralized and Decentralized Partially Observable Markov Decision Processes: Efficient Decision-Making and Coordination in Uncertain Environments*; (Shlomo Zilberstein, Advisor); Sept. 2010; Analytics, Modeling and Simulation Scientist, Aptima, Inc.

As agents are built for ever more complex environments, methods that consider the uncertainty in the system have strong advantages. This uncertainty is common in domains such as robot navigation, medical diagnosis and treatment, inventory management, sensor networks and e-commerce. When a single decision maker is present, the partially observable Markov decision process (POMDP) model is a popular and powerful choice. When choices are made in a decentralized manner by a set of decision makers, the problem can be modeled as a decentralized partially observable Markov decision process (DEC-POMDP). While POMDPs and DEC-POMDPs offer rich frameworks for sequential decision making under uncertainty, the computational complexity of each model presents an important research challenge. As a way to address this high complexity, this thesis develops several solution methods based on utilizing domain structure, memory-bounded representations and sampling. These approaches address some of the major bottlenecks for decision-making in real-world uncertain systems. The methods include a more efficient optimal algorithm for DEC-POMDPs as well as scalable approximate algorithms for POMDPs and DEC-POMDPs. Key contributions include optimizing compact representations as well as automatic structure extraction and exploitation. These approaches increase the scalability of algorithms, while also increasing their solution quality.

Bo An; *Automated Negotiation for Complex Multi-Agent Resource Allocation;* (Victor Lesser, Advisor); Feb. 2011; Postdoctoral Associate, Dept. of Computer Science, University of Southern California.

Automated negotiation (bargaining) is the most widely used approach for multi-agent resource allocation and it has recently received increasing attention. However, information uncertainty, existence of multiple contracting partners and competitors, agents' incentive to maximize individual utilities, and market dynamics make it difficult to calculate agents' rational equilibrium negotiation strategies and to develop successful negotiation agents which behave well in practice. This thesis is concerned with analyzing agents' rational behavior and developing negotiation strategies for a range of complex negotiation contexts. First, we consider the problem of finding agents' rational strategies in bargaining with incomplete information. We provide an algorithm based on the combination of game theoretic analysis and search techniques, which finds agents' equilibrium in pure strategies when they exist. Next, we extend the alternating-offers protocol to handle concurrent negotiations in which each agent has multiple trading opportunities and faces market competition. We provide an algorithm based on backward induction to compute the subgame perfect equilibrium of concurrent negotiation. Third, we present the design and implementation of agents that concurrently negotiate with other entities for acquiring multiple resources. Finally, we consider the problem of allocating networked resources in dynamic environment. We propose a distributed negotiation mechanism where agents negotiate over both a contract price and a decommitment penalty.

Aruna Balasubramanian; *Architecting Protocols to Enable Mobile Applications in Diverse Wireless Networks;* (Arun Venkataramani and Brian Levine, Advisors); Feb. 2011; CI Fellow, Dept. of Computer Science, University of Washington.

The goal of my thesis is to architect robust protocols that overcome disruptions and enable applications in diverse mobile networks. Mobile networks are prone to frequent disruptions and therefore cannot support most Internet applications. In this thesis, I focus on four networks that span the diverse connectivity spectrum to answer the question: What protocol design best overcomes disruptions and enables applications in a given network? I design four utility-driven protocols that tolerate disruptions in an environment by leveraging opportunism. Specifically, I present: 1) RAPID, which uses opportunistic replication to enable bulk transfer in mostly disconnected networks; 2) Thedu, which uses opportunistic prefetching to enable web search in intermittently connected networks; 3) ViFi, which uses opportunistic forwarding to enable Voice over IP (VoIP) in mostly connected mesh networks; and 4) Wiffler, which uses opportunistic augmentation to improve application performance in mostly connected cellular networks. Finally, I present a detailed evaluation of the protocols using implementation and deployment experiments on two large scale vehicular testbeds. Deployment on a real testbed shows that the protocols are practical and can be implemented in realistic usage environments. The evaluations show that the protocols significantly improve performance of applications compared to the state-of-the-art, in their respective environments.

George Bissias; *Bounds on Service Quality for Networks Subject to Augmentation and Attack;* (Brian Levine, Advisor); Sept. 2010; Algorithms Developer, Fluent Mobile.

Assessing a network's vulnerability to attack and random failure is a difficult and important problem that changes with network application and representation. We furnish algorithms that bound the robustness of a network under attack. We utilize both static graph-based and dynamic trace-driven representations to construct solutions appropriate for different scenarios. For static graphs we first introduce a spectral technique for developing a lower bound on the number of connected pairs of vertices in a graph after edge removal, which we apply to random graphs and the power grid of the Philippines. To address the problem of resource availability in networks we develop a second technique for bounding the number of nominally designated client vertices that can be disconnected from all server vertices after either edge or vertex removal (or both). Dynamic networks are modeled as disruption tolerant networks (DTNs). DTNs are composed of mobile nodes that are intermittently connected via short-range wireless radios. In the context of both human and vehicular mobility networks we study both the effect of targeted node removal and the effect of augmentation with stationary relays.

TJ Brunette; *Adaptive Balancing of Exploitation with Exploration to Improve Protein Structure Prediction;* (Oliver Brock, Advisor); May 2011; Senior Fellow, Dept. of Biochemistry, University of Washington.

One of the most significant impediments for protein structure prediction is the inadequacy of conformation space search. Conformation space is too large and the energy landscape too rugged for existing search methods to consistently find near-optimal minima. Conformation space search methods thus have to focus exploration on a small fraction of the search space. The ability to choose appropriate regions, i.e. regions that are highly likely to contain the native state, critically impacts the effectiveness of search. To make the choice of where to explore requires information, with higher-quality information resulting in better choices. Most current search methods are designed to work in as many domains as possible, which leads to less accurate information because of the need for generality. However, most domains provide unique, and accurate information. To best utilize domain-specific information, search needs to be customized for each domain. The first contribution of this thesis customizes search for protein structure prediction, resulting in significantly more accurate protein structure predictions. My results indicate that integrating the information between homologs and fragments significantly improves protein structure prediction accuracy, resulting in several proteins predicted with 1 angstrom RMSD resolution.

Bin Chen; *Improving Processes Using Static Analysis Techniques;* (Lori Clarke and George Avrunin, Advisors); Feb. 2011; Quantitative Developer, Vegasoul Capital.

Real-world processes often undergo improvements to meet certain goals, such as coping with changed requirements, eliminating defects, improving the quality of the products, and reducing costs. Identifying and evaluating the defects or errors in the process, identifying the causes of such defects, and validating proposed improvements all require careful analysis of the process. Human-intensive processes are of particular concern because they can be extremely complex and may be used in critical, including life-critical, situations. To date, the analysis support for such processes is very limited. There has been considerable success lately in using static analysis techniques to analyze hardware systems, software systems, and manufacturing processes. This thesis explores how such analysis techniques can be automated and employed to effectively analyze life-critical, human-intensive processes. We investigated two static analysis techniques: Finite-State Verification (FSV) and Fault Tree Analysis (FTA). We proposed a process analysis framework that is capable of performing both FSV and FTA on rigorously defined processes. We evaluated this framework based on the Little-JIL process definition language and employed it to analyze two real-world, human-intensive processes – an In-Patient Blood Transfusion Process and a Chemo Therapy Process. The results show that the framework can be used effectively to detect defects in such real-world, human-intensive processes.

Michael Hay; *Enabling Accurate Analysis of Private Network Data;* (Gerome Miklau and David Jensen, Advisors); Sept. 2010; Postdoctoral Associate, Dept. of Computer Science, Cornell University.

This dissertation addresses the challenge of enabling accurate analysis of network data while ensuring the protection of network participants' privacy. Massive amounts of data are being collected (facebook activity, email correspondence, cell phone records), and there is huge interest in analyzing the data, but the data is not being shared due to concerns about privacy. Despite much research in privacy-preserving data analysis, existing technologies are inadequate because they were designed for tables, not networks, and cannot be easily adapted to handle the complexities of network data. We develop several technologies to meet this challenge. First, we develop a framework for assessing the risk of publishing a network that has been "anonymized." Using this framework, we show that a small amount of background knowledge about local network structure suffices to re-identify an "anonymous" individual. This motivates our second contribution: an algorithm that transforms network structure to provably lower re-identification risk. In comparison with other algorithms, our approach more accurately preserves important features of the network topology. Finally, we consider an alternative paradigm, in which the analyst accesses data through a regulated query interface. We show that the degree sequence of a network can be accurately estimated while ensuring strong privacy protection.

Vidit Jain; *Using Context to Enhance the Understanding of Face Images;* (Erik Learned-Miller, Advisor); Sept. 2010; Scientist, Yahoo! Labs Bangalore.

Faces are special objects of interest. Developing automated systems for detecting and recognizing faces is useful in a variety of application domains including providing aid to visually-impaired people and managing large-scale collections of images. Humans have a remarkable ability to detect and identify faces in an image, but related automated systems perform poorly in real-world scenarios, particularly on faces that are difficult to detect and recognize. Why are humans so good? There is general agreement in the cognitive science community that the human brain uses the context of the scene shown in an image to solve the difficult cases of detection and recognition. We focus on emulating this approach by using different kinds of contextual information for improving the performance of various approaches for face detection and face recognition: an algorithm that employs the easy-to-detect faces in an image to find the difficult-to-detect faces in the same image and a joint probabilistic model for image-caption pairs. This model solves the difficult cases of face recognition in an image by using the context generated from the caption associated with the same image. Finally, we present an effective solution for classifying the scene shown in an image, which provides useful context for both of the face detection and recognition problems.

George Konidaris; *Autonomous Robot Skill Acquisition;* (Andrew Barto, Advisor); May 2011; Postdoctoral Associate, MIT.

Among the most impressive aspects of human intelligence is skill acquisition–the ability to identify important behavioral components, retain them as skills, refine them through practice, and apply them in new task contexts. Skill acquisition underlies both our ability to choose to spend time and effort to specialize at particular tasks, and our ability to collect and exploit previous experience to become able to solve harder and harder problems over time with less and less cognitive effort. Hierarchical reinforcement learning provides a theoretical basis for skill acquisition, including principled methods for learning new skills and deploying them during problem solving. However, existing work focuses largely on small discrete problems. This thesis identifies the primary obstacles to achieving autonomous skill acquisition in high-dimensional, continuous domains and introduces three methods for overcoming these obstacles: skill chaining, a general skill discovery method for continuous reinforcement learning domains; abstraction selection, an efficient algorithm for selecting a suitable abstraction from a library of available abstractions when creating a new skill; and CST, a method for rapidly building trees of skills (with appropriate abstractions) from sample trajectories obtained via human demonstration, a feedback controller, or a planner. These algorithms are applied to achieve autonomous skill acquisition on the uBot-5.

Ming Li; *Data Management and Wireless Transport for Diverse Sensor Networks;* (Deepak Ganesan and Arun Venkataramani, Advisors); Sept. 2010; Research Staff Member, IBM T.J. Watson.

Today, various sensor networks have emerged and span a wide range of sensing capabilities, computation, energy and communication resources, and user needs. They pose unique design challenges to the distributed system design. We focus on following four challenges in data management and wireless transport. 1) We examine how to explore the resource-rich proxies on the edge of the network to assist the resource-poor sensors. We propose a novel two-tier sensor data management architecture, PRESTO, that proxies model sensed data and predict future data, while sensors check sensed data with model-predicted values. 2) We look at the sensing application regime where a single sensor network has to support diverse users and applications with limited bandwidth and computation resources. We propose a utility-driven approach, MUDS, that maximizes data sharing across users while using limited resources. 3) We seek to improve the matching performance between users' interest and sensed data in large scale sensing applications. We propose BlueDove, a cloud-based publish/subscribe system that takes advantage of the rich resources and flexibility of a computation cloud. 4) We examine how to make the underlying wireless transport between sensor nodes more reliable and efficient. We propose a clean-slate re-design of the network stack, Hop, that uses reliable per-hop block transfer as a building block and builds all other components such as back-pressure congestion control and end-to-end virtual retransmission on top of block transfer.

Daniel Menasche; *Enabling Peer-To-Peer Swarming For Multi-Commodity Dissemination;* (Donald Towsley, Advisor); May 2011; Assistant Professor, Dept. of Computer Science, Federal University of Rio de Janeiro.

Peer-to-peer swarming, as used by BitTorrent, is one of the de facto solutions for content dissemination in today's Internet. Although peer-to-peer swarming has been widely studied for a decade, prior work has focused on the dissemination of one commodity (single file). This thesis focuses on the multi-commodity case. We have discovered through measurements that a vast number of publishers currently disseminate multiple files in a single swarm (bundle). The first contribution of this thesis is a model for content availability. We use the model to show that, when publishers are intermittent, bundling K files increases content availability exponentially as function of K. When there is a stable publisher, we consider content availability among peers. Our second contribution is the estimate of the dependency of peers on the stable publisher (self-sustainability). Then, we investigate reciprocity and the use of barter that occurs among peers. As our third contribution, we prove that the loss of efficiency due to the download of unrequested content to enforce direct reciprocity, as opposed to indirect reciprocity, is at most two in a class of networks without relays. As our fourth contribution, we present conditions for the existence and uniqueness of an equilibrium between publishers and peers.

Albert (Gene) Novark; *Hardening Software Against Memory Errors and Attacks;* (Emery Berger, Advisor); Feb. 2011; Associate, Morgan Stanley.

Programs written in C and C++ are susceptible to a number of memory errors, including buffer overflows and dangling pointers. At best, these errors cause crashes or performance degradation. At worst, they enable security vulnerabilities, allowing denial-of-service or remote code execution. Existing runtime systems provide little protection against these errors. They allow minor errors to crash the program and ensure predictability that allows attackers to consistently exploit vulnerabilities. In this thesis, we introduce a series of runtime systems that detect and tolerate these errors in deployed applications. By design, these systems tolerate minor errors while lowering the probability of successfully exploiting security vulnerabilities. The first such system, Archipelago, protects exceptionally sensitive server applications against severe errors using an object-per-page randomized allocator. It provides near-100% protection against certain common attack vectors. Our second system, DieHarder, combines ideas from Archipelago and previous systems to enable maximal protection against attacks while incurring minimal runtime and memory overhead. Our final system, Exterminator, automatically corrects heap-based buffer overflows and dangling pointers without requiring programmer intervention. Exterminator relies on a low-overhead randomized allocator and statistical inference techniques to isolate and correct errors in deployed applications, deterministically tolerating errors and forcing attackers to adapt to a changing attack surface.

Marek Petrik; *Optimization-based Approximate Dynamic Programming: Reliable Algorithms and Feature Selection;* (Shlomo Zilberstein, Advisor); Sept. 2010; Postoctoral Researcher, IBM Research.

Reinforcement learning algorithms hold promise in many complex domains, such as resource management and planning under uncertainty. Most reinforcement learning algorithms are iterative--they successively approximate the solution based on a set of samples and features. Although these iterative algorithms can achieve impressive results in some domains, they are not sufficiently reliable for wide applicability. Some of the most interesting reinforcement learning algorithms are based on approximate dynamic programming (ADP). This thesis presents new reliable algorithms for ADP that use optimization instead of iterative improvement. We improve on approximate linear programming–an existing method–and derive approximate bilinear programming–a new robust approximate method. The methods presented in this thesis can potentially be used in many practical applications in artificial intelligence, operations research, and engineering. Our experimental results show that optimization-based methods may perform well on resource-management problems and standard benchmark problems and therefore represent an attractive alternative to traditional iterative methods.

Piyanuch (Pla) Silapachote; *Biologically Inspired Vision Systems;* (Allen Hanson, Advisor) May 2011; Faculty, Dept. of Information and Communication Technology, Mahidol University.

One approach to the design of intelligent machines capable of perceiving visual information is to model the fascinating primate vision. Based on discoveries in neuroscience, physiology, and psychology, biologically-plausible models for object recognition and classification can be simple yet achieve high accuracy and generalize well. The proposed system employs unsupervised feature learning, simulating hypercolumns of the primary cortex, a hierarchical feed-forward framework, mimicking simple and complex cells, and neural network classification, a computational model of interconnected neurons. Compared to existing approaches, this system is more biologically inclined as well as more effective. Compared to other state-of-the-art systems it achieves good accuracies with significantly shorter runtimes. Other key features are good generalizability and independence of delicate segmentation procedures employed by many other systems. Experiments are conducted both on natural scenes and challenging realistic underwater marine images. The latter is a rather uncommon data source, no biologically inspired vision systems had previously been applied to it. Despite domain-specific difficulties, such as low image quality, and high diversity of shapes and motions, the potential of the proposed system is shown to be quite promising. This bodes well for future application to other domains currently not considered by mainstream computer vision.

Jacob Sorber; *System Support for Perpetual Mobile Tracking;* (Mark Corner, Advisor); Sept. 2010; Postdoctoral Associate, Dartmouth College.

Advances in low-power electronics, energy harvesting, and sensor technologies will revolutionize mobile and embedded computing, by enabling networks of mobile sensor devices that are long-lived and self-managing. When realized, this new generation of perpetual systems will have a far-reaching and transformative impact, improving scientists' ability to observe natural phenomena, and enabling many ubiquitous computing applications for which regular maintenance is not feasible. Perpetual systems face many challenges. Conditions at runtime are unknown and highly variable. Variations in harvested energy and energy consumption, as well as changes in network connectivity and bandwidth require systems that are able to adapt gracefully at run-time to meet different circumstances. However, when programmers muddle adaptation details with application logic, the resulting code is often difficult to understand as well as maintain. This dissertation demonstrates that perpetual systems can be designed and deployed without sacrificing programming simplicity. We describe two systems. Eon, the first energy-aware language, allows programmers to simply express energy policies and delegate the complexities of energy management to the underlying system. The second, Tula, automatically balances the inherently dependent activities of data collection and data delivery, while also ensuring that devices have fair access to network resources.

Siddharth Srivastava; *Foundations and Applications of Generalized Planning;* (Neil Immerman and Shlomo Zilberstein, Advisors); Sept. 2010; Postdoctoral Research Associate, Dept. of Computer Science, UMass Amherst.

Research in AI Planning is largely focused on the problem of finding plans or sequences of actions for going from a specific initial state to a goal state. The inherent complexity of this task and uncertainty in real-world situations make it desirable to find "generalized" plans which can be used to solve classes of similar problem instances. However, such generalized solutions typically require loops of actions, which are impossible even to verify as correct in the general case because of the undecidability of the halting problem for Turing Machines. In my thesis, I addressed both theoretical and applied aspects of this problem. I first develop an efficient method for determining the correctness and utility of a general class of loops of actions. Next, I utilize this method to develop original algorithms for computing generalized plans by identifying loops in sample classical plans and merging classical plans together; by using classical planners to automate this process and thereby solve a given generalized problem from scratch; and also by conducting a direct search in the space of abstract states.

John Sweeney; *A Teleological Approach to Robot Programming by Demonstration;* (Roderic Gupen, Advisor); Feb. 2011; Research Associate, FDO Partners, LLC.

This dissertation presents an approach to robot programming by demonstration based on two key concepts: demonstrator intent is the most meaningful signal that the robot can observe, and the robot should have a basic level of behavioral competency from which to interpret observed actions. I argue that programming by demonstration can be organized into declarative and procedural components. The declarative component represents a reusable outline of underlying behavior that can be applied to many different contexts. The robot can use the knowledge encapsulated in sensorimotor schemas to interpret the demonstration. The procedural component represents the dynamic portion of the task that is based on features observed at run time. I describe how statistical models, Bayesian methods in particular, can be used to model these components. These models have many features that are beneficial for learning in this domain, such as tolerance for uncertainty, and the ability to incorporate prior knowledge into inferences. I demonstrate this architecture through experiments on a bimanual humanoid robot using tasks from the pick and place domain. Additionally, I develop and experimentally validate a model for generating grasp candidates using visual features that is learned from demonstration data.

Louis Theran; *Problems in Generic Combinatorial Rigidity: Sparsity, Sliders, and Emergence of Components;* (Ileana Streinu, Advisor); Sept. 2010; Research Assistant Professor, Mathematics Dept., Temple University.

A bar-joint framework is made of fixed-length bars connected by universal joints with full rotational degrees of freedom; the allowed motions preserve the lengths and connectivity of the bars, and a framework is rigid if the only allowed motions are trivial motions of Euclidean space. The remarkable Maxwell-Laman Theorem says that rigidity of generic bar-joint frameworks depends only on the graph that has as its edges the bars and as its vertices the joints. We generalize the "degree of freedom counts" that appear in the Maxwell-Laman theorem to the very general setting of (k,l)-sparse and (k,l)-graded sparse hypergraphs, giving graph-theoretic, matroidal, and algorithmic characterization of them. We then introduce a new rigidity model: slider-pinning rigidity. This is an elaboration of the planar bar-joint model to include sliders, which constrain a vertex to move on a specific line. We prove the analogue of the Maxwell-Laman Theorem for slider pinning, using, as a lemma, a new proof of Whiteley's Parallel Redrawing Theorem. We then study the emergence of rigid substructures in a generic framework with the combinatorics of a sparse Erdös-Renyi random graph, proving the existence of a sharp threshold for them to exist and a linear-sized lower bound when they emerge.

Chang Wang; *A Geometric Framework for Transfer Learning Using Manifold Alignment;* (Sridhar Mahadevan, Advisor); Sept. 2010; Research Scientist, IBM TJ Watson Research.

Many machine learning problems involve dealing with a large amount of high-dimensional data across diverse domains. Annotating or labeling the data is also expensive as it involves significant human effort. This work explores a solution to these problems by exploiting the property that high-dimensional data in real-world application domains often lies on a lower-dimensional structure, whose geometry can be modeled as a graph or manifold. We propose a set of novel manifold-alignment based approaches for transfer learning. The proposed approaches transfer knowledge across different domains by finding low-dimensional embeddings of the datasets to a common latent space, which simultaneously match corresponding instances while preserving local or global geometry of each input dataset. We develop several extensions of a manifold alignment framework to more challenging situations, including (1) when no correspondences across domains are given; (2) when the global geometry of each input domain needs to be respected; (3) when label information rather than correspondence information is available. Another contribution of this thesis is the study of multiscale methods for manifold alignment. Multiscale alignment automatically generates alignment results at different levels by discovering the shared intrinsic multilevel structures of the given datasets, providing a common representation across all input datasets.

Philipp Weis; *Expressiveness and succinctness of first-order logic on finite words;* (Neil Immerman, Advisor); May 2011; Software Engineer, Google, Inc.

Expressiveness, and more recently succinctness, are two central concerns of finite model theory and descriptive complexity theory. We develop new bounds on the expressiveness and succinctness of first-order logic with two variables on finite words, present a related result about the complexity of the satisfiability problem for this logic, and explore a new approach to the generalized star-height problem from the perspective of logical expressiveness. Using our complete characterization of the expressive power of first-order logic with two variables on finite words, we prove that the quantifier alternation hierarchy for this logic is strict, settling the main remaining open question about the expressiveness of this logic. We also prove a polynomial-size small-model property for this logic, leading to an NP algorithm and thus proving that the satisfiability problem for this logic is NP-complete. Finally, we investigate the generalized star-height problem. As of today, we do not even know whether there exists a regular language that has generalized star-height larger than 1. We show that this problem can be phrased as an expressiveness question for first-order logic with a restricted transitive closure operator, and use established tools from finite model theory to gain new insights into the generalized star-height problem.

Faculty News






This fall, Professor **David Mix Barrington** took over **Robbie's Moll**'s position as Associate Chair for Academic Affairs. Professor **Rod Grupen** took on Barrington's former role as Chief Undergraduate Advisor. Associate Professor **Sridhar Mahadevan** replaced Professor **James Allan** as Graduate Program Director, and Lecturer **Timothy Richards** stepped in as the Honors Program Director.

Associate Professor **Kevin Fu** and graduate student **Ben Ransford**, along with collaborators from MIT, received the SIGCOMM Best Paper Award for "They Can Hear Your Heartbeats: Non-Invasive Security for Implantable Medical Devices."

Associate Professor **Emery Berger** gave an invited talk at USENIX WOOT 2011 (Workshop on Offensive Technologies). In March, he will chair the 2012 Workshop on Determinism and Correctness in Parallel Programming (WoDET 2012). Berger also started a new graduate course in Systems that will become a core class in the curriculum. In December, he received a Google Research Award for his proposal "Causal Profiling: Low-Overhead, High Precision Profiling."

Berger and **Fu** were each elevated to the level of ACM Senior Member.

In October, Distinguished Professor **Bruce Croft** gave the talk "Social Search: Information Retrieval with Ephemeral Data" as part of the 2011 Samuel D. Conte Distinguished Lecture Series at Purdue University Department of Computer Science.

Barto feted

The department showed its appreciation for Professor Andrew Barto's years of service as department chair with a celebration this fall. Barto was chair from 2007 - 2011. He is shown opening his gift during the party.

In December, Professor **Leon Osterweil** gave a keynote address, "What I Have Learned From Defining, Analyzing, and Executing Processes," at the 2011 Asia-Pacific Software Engineering Conference (APSEC 2011).

Professor **Andrew McCallum** will be the General Chair of 2012 International Conference on Machine Learning in Edinburgh, Scotland. He received new research awards from Google and Oracle for work on information extraction, probabilistic programming and probabilistic databases. He also designed and taught new graduate-level course on Graphical Models last spring.

Assistant Professor **Hanna Wallach** was a co-organizer of the Second Workshop on Computational Social Science and the Wisdom of Crowds, held at the Neural Information Processing Systems Foundation Conference (NIPS 2011) in December. She was also one of four invited participants in two video interviews for The Science Network on the history and future of the Neural Information Processing Systems (NIPS) Conference.

Associate Professor **Hava Siegelmann** was named a Senior Member of IEEE and also a Senior Member of the International Neural Networks Society. She was also the Program Chair of the 2011 International Joint Conference on Neural Networks held in San Jose, California.

Professor **Brian Levine** was the Program co-Chair of the 2011 ACM MobiCom Conference, and the Technical Program Vice Chair of the 2011 DFRWS Digital Forensics Research Conference.

In November, Distinguished Professor Emeritus **Victor Lesser** gave a lecture, "A Practical Application of Cooperative Multi-Agent Technology: The DCAS Severe Weather Radar Detection and Tracking System," at North Carolina State University as part of the Triangle Computer Science Distinguished Lecturer Series that includes Duke University, University of North Carolina and North Carolina State University.

Professor **Eliot Moss**, alum **Antony Hosking** (Ph.D. '95), and Richard Jones recently published the book *The Garbage Collection Handbook: The Art of Automatic Memory Management*.

Assistant Professor **Rui Wang** received funding from Robert Bosch Research for a project on "High-Definition 4D Rendering for Future Navigation Systems."

With National Science Foundation (NSF) funding, Professor Emeritus **Rick Adrion** and MIT geoscientist Christopher Hill are leading the effort to use new technology to teach middle school science students in Springfield and Holyoke.

Professor Emeritus **Allen Hanson**, CS graduate student **Adam Williams**, along with colleagues from Chemistry, Microbiology, and Computer Science, are working on a project that displays large-scale interactive molecules to prod individuals to personally explore a vast array of molecular structures. The "Molecular Playground", supported by a Camille and Henry Dreyfuss Foundation grant, is now on display at the Springfield Science Museum.

Professor **Edwina Rissland** has returned to the NSF for a second term as a Program Director in the Robust Intelligence cluster in the Division of Information and Intelligent Systems within the Directorate for Computer & Information Science & Engineering (CISE). She is also on the management team for the CISE Computing Research Infrastructure program.

Adjunct Professor **Jane Fountain** gave a keynote address at GovCamp Singapore. The event focuses on positioning government as a platform to engage key sectors of a country to come together in one conference and to jointly collaborate on how to improve citizen engagement and services using technology.

Adjunct Professor **Lee Spector** is a co-PI on the Four Colleges BioMathematics Consortium NSF grant and PI on an NSF grant titled "Evolution of Robustly Intelligent Computational Systems." He co-directs a Hampshire College "Design, Art, and Technology" program, supported by the Sherman Fairchild Foundation grant, and is involved in launching Hampshire's new Creativity Center, supported by funding form Eileen Fisher. Spector is also Editor-in-Chief of the Springer journal Genetic Programming and Evolvable Machines.

Associate Professor **Gerome Miklau** and his wife Johanna Callard are the proud parents of Owen Callard Miklau, born on September 2nd.

Samuel Dunn, **Zachary Dunn**, and **Timothy Richards** joined the department as Lecturers this fall.

Researcher News

Jianqing Ma joined the Networks group as a Visiting Professor from Fudan University, and **Liyuan Sun** joined the group as a Visiting Scholar from Tsinghua University.

David Soergel joined the Information Extraction and Synthesis Lab (IESL) as a Postdoctoral Research Associate.

Chongjie Zhang joined the Multi-Agent Systems Lab as a Postdoctoral Research Assistant.

Student News

Vinay Shah was selected for Honorable Mention in the Computing Research Association's Outstanding Undergraduate Researcher Award competition for 2012.

Undergrad **Evan Shelhamer** is one of four to receive the 2010-2011 Upsilon Pi Epsilon ACM Student Chapter Scholarship Award.

Thanks, Claire

After 18 years in the department as our Director of Administration and Finance, Claire Christopherson has accepted a new position as Director of Administration and Finance for the Massachusetts Green High Performance Computer Center (MGHPCC) in Holyoke. "Claire has been the glue the kept the department together and running smoothly for all of these years. We greatly appreciate her dedication to the department and to the campus. She will definitely be missed," says department chair Lori Clarke.

At the 2011 ACM Mobicom Student Workshop held in Las Vegas, Nevada in September, **Yung-Chih Chen** received an award for the Best Theory Session Talk for his paper, "A Simple Queueing Network Model of Mobility in a Campus Wireless Network," and **Anand Seetharam** won the "Best Startup Idea" competition.

The UMass CS ACM Student Chapter and Yahoo! hosted a Hackathon in September. The Kinect Hacking winner was **Paul Edwards** and the Rom Hacking winner was **Nathan Deren**.

CS graduate students **Scott Kuindersma** and **George Konidaris** were presented with the Best Student Video award during the Twenty-Fifth Conference on Artificial Intelligence (AAAI-11) for their video, "Autonomous Robot Skill Acquisition."

CIIR grad students **Marc Cartright** and **Henry Field** received the Best Paper Award at the BooksOnline'11 Workshop for their paper "Evidence Finding using a Collection of Books."

Staff News

Debora Comeau joined the Computer Science Computing Facility as their Administrative Assistant/Bookkeeping Staff.

Matthew Cornell (M.S. '91) returned to the department as a Research Fellow in the Knowledge Discovery Lab. Previously, he contributed to the recently published Wiley book, *Mindhacker: 60 Tips, Tricks, and Games to Take Your Mind to the Next Level*.

Rachel Leach joined the Commonwealth Alliance for Information Technology Education as a Communications Assistant.

Brian Martin joined the the Information Extraction and Synthesis Lab as an Associate Software Engineer.

Diane Muller is the new Receptionist in the department.

Susan Overstreet joined the department as an Accountant, replacing **Gwyn Mitchell**, who retired this fall.

Significant Bits

Newsletter of the
Department of Computer Science
College of Natural Sciences
at the University of Massachusetts Amherst

140 Governors Drive
University of Massachusetts Amherst
Amherst, MA 01003-9264

"Significant Bits" is published twice a year by the Department of Computer Science, University of Massachusetts Amherst **(www.cs.umass.edu)**.
Your news, suggestions, comments, and contributions are welcome. Please mail them to the address above or send them electronically to **bits@cs.umass.edu**.

Department Chair	Lori A. Clarke
Editor	Jean Joyce
Art Direction	North Haven Design
Graduate Student Liaisons	Filip Jagodzinski, Borislava Simidchieva
Contributors	Rachel Leach, Bruce MacLeod, Andrew McCallum, UMass Amherst News Office staff

NONPROFIT ORG.
U.S. Postage
PAID
Permit No. 2
AMHERST, MA

Thanks for your support

The following alumni and friends have actively supported the Department of Computer Science from May 2011 through December 2011. Such financial support is greatly appreciated and helps maintain a world-class instructional and research program. Contributions from alums and friends help to fund important special activities that are not supported through the state budget.

Those interested in helping the department should visit **www.cs.umass.edu/about/donate** for online donations or send a check made out to *UMass Amherst* to:
Attn: Jean Joyce, Department of Computer Science
University of Massachusetts Amherst,
140 Governors Drive, Amherst, MA 01003-9264
Please state that your gift is restricted to Computer Science.

Mrs. Christine M. Ahrens
Mr. Daniel R. Amirault ('10)
J.T. and Patricia Amirault ('82)
Dr. Krishnamoorthy Arvind ('91)
Dr. Kevin D. Ashley ('88)
Mr. Mariappan Asokan ('86)
Mr. Armen P. Babikyan ('02)
Aruna and Niranjan Balasubramanian ('11)
Mr. James J. Barber, Jr. ('08)
Dr. Daniel J. Barrett ('98)
Mr. John V. Bellissimo
Prashant and Hema Bhatt
George and Kim Bowker ('83)
Dr. Eric W. Brown ('96)
Mr. Daniel F. Burke ('82)
Jamie Callan and Adele Weitz ('93)
Ms. Jennifer M. Cannan ('03)
Dr. Aaron G. Cass ('05)
Mr. Joseph F. Catalano ('02)
Dr. Shenze Chen ('92)
Mr. Aaron Cho ('04)
Dr. Yuan-Chieh R. Chow ('77)
Mr. Eric Christophersen ('05)
Mr. Jacky Cheuk Kei Chu ('01)
Mr. Paul A. Chukiu ('07)
Mr. Stephen D. Clare ('01)
Dr. Daniel D. Corkill ('83)
Ms. Kristen M. Day
Ms. Donna E. Deeley ('70)
Mr. Charles F. Denison ('03)
Edward and Mary Depierro
Mr. Richard T. Donlon
Mr. Wayne W. Duso ('85)
Maury Falkoff and Luanne Hendricks ('82)
Dr. Zhengzhu Feng ('05)
Dr. Claude L. Fennema, Jr. ('91)
Victor and Laura Firoiu ('98)
Mr. William Fisher ('75)
In Memory of Michael J. Frongillo
Mr. Thomas A. Galanis ('99)
Mr. Jiuhe Gan ('95)
Paul and Valerie Gilman ('82)
Lisa and Jon Glass ('85)
Margaret and Mark Godin
Mr. John Goshdigian ('73)
Mr. Charles B. Greely ('85)
Mr. Thomas E. Greene ('07)
Ms. Judith L. Greenspan ('82)
Cheryl and Edmund Harrington ('75)
Mr. Michael J. Hartman ('82)
Sharon and William Hingley
Mr. Richard L. Housel ('80)
Ms. Jin Huang ('96)
Mr. Michael P. Johnson ('97)
Mrs. Colleen B. Kettle ('84)
Mr. Hanumantha R. Kodavalla ('88)
Mr. Gordon M. Koizumi ('72)
Mr. E. Micah Kornfield ('02)
Mr. Anthony J. Leonardi ('05)
Mr. Jones S. Leung ('04)
Mr. Austin P. Maher ('87)
Dr. Victoria U. Manfredi ('05)
Mrs. Ruth C. McGlothlin ('62)
Mr. Jacob D. Mitchell ('10)
Dr. Michael C. Monks ('87)
Mr. Michael T. Morganti ('85)
Mr. Brian J. Morris ('91)
Dr. Daniel E. Neiman ('92)
Dr. Jitendra D. Padhye ('00)
Mr. Charles N. Paliocha ('83)
Dr. Achilleas Papakostas ('93)
Ms. Areti N. Papanastasiou ('87)
Dr. Ron Papka ('99)
Dr. Jasmina D. Pavlin ('85)
Ms. Desislava I. Petkova ('08)
Mr. Robert M. Pilachowski ('88)
Yao Ren and Qin Yao ('00)
Ms. Pat Rogers
Ms. Lilian Ruiz
Mr. David R. Schmitt ('78)
Mr. Michael J. Shaughnessy ('88)
Ms. Kimberly D. Silverman ('83)
Mr. Michael J. Sullivan ('71)
Keith and Amy Swan ('90)
Miss Nina Tsao ('75)
Vijay Turlapati & Vijayalakshmi Merla
Mrs. Renee B. Venne
Mr. John M. Vervaert ('73)
Dr. Rukmini Vijaykumar ('88)
Dr. Stephen T. Vinter ('85)
Mr. Timothy J. Volpe ('01)
Mr. Hao-Jin Wang ('88)
Yunqing and Changting Wang ('02)
Mr. Harteg S. Wariyar ('07)
Mrs. Nancy W. Wasiuk ('79)
Prof. Charles C. Weems, Jr. ('84)
Dr. John L. Woods ('70)
Timothy and Heidi Wright ('86)
Mr. Yan Xiao ('00)
Dr. Zhongfei Zhang ('96)

Victor Lesser Graduate Scholarship
Dr. Malini K. Bhandaru ('98)
Dr. Bryan C. Horling ('98)
Dr. Tuomas W. Sandholm and Dr. Christina M. Fong ('96)
Dr. Thomas A. Wagner ('00)

David W. Stemple Memorial Scholarship
Prof. Lori A. Clarke and Prof. Leon J. Osterweil
Dr. Jane Yolen

Paul Utgoff Memorial Scholarship
Mrs. Karen Utgoff
Dr. Victor A. Utgoff

Corporate and Matching Gifts
Bosch Research and Technology Center
Cisco Systems Inc.
EMC Corporation
Microsoft Corporation
TripAdvisor
VistaPrint
Yahoo! Inc.