Sixth Annual Outstanding Achievement and Advocacy Awards

The accomplishments of this year's Outstanding Achievement and Advocacy (OAA) Award winners and CS undergraduate and graduate student award recipients were celebrated during a banquet held at the Mullins Center on Friday, May 2, 2014.

continued on page 12

(l. to r.): CS Chair Lori Clarke, Zhi-Li Zhang, Donald House, CNS Dean Steve Goodwin, Randy Ellis, Mary-Ellen Prescott, Vanessa Murdock, Prof. Leon Osterweil, and EMC's Cindy Gallerani.

SIXTH OAA AWARDS – – – – – – – – *cont. from page 1*

Prior to the banquet, OAA events included building tours, a reception, and a panel discussion to highlight the varying technical career paths of our award recipients who shared their knowledge and experiences as students at UMass Amherst or as employers of our graduates.

During the banquet, School Chair Lori Clarke and Steve Goodwin, Dean of the College of Natural Sciences, welcomed the attendees. Professor Leon J. Osterweil presented awards to the 2014 OAA award recipients.

The 2014 OAA Award Recipients are:

Outstanding Contributions to Society: **Randy E. Ellis** (UMass Amherst CS Ph.D. '87), Professor in the School of Computing, Department of Mechanical Engineering, and Department of Surgery at Queen's University at Kingston, Ontario. Dr. Ellis's work focuses on using imaging to support and improve the work of surgeons. His work has focused most closely on supporting musculoskeletal surgery. This work has doubtlessly improved the lives of countless thousands of people suffering from problems with their hips, knees, and other joints. He is the Project Leader of a large multidisciplinary group that investigates advanced health-care delivery for the coming decade.

Outstanding Achievement in Education: **Donald H. House** (UMass Amherst CS Ph.D. '84), Professor and Chair, Division of Visual Computing in the School of Computing at Clemson University. Dr. House is well known for his research in the area of physically-based modeling, including approaches to the simulation of cloth, drapery, and other flexible materials. More recently, he has been focusing on perceptual issues in visualization, including texturing volumetric surfaces, eye tracking in stereo display environments, and visualization of uncertainty in data and predictions.

Outstanding Achievement by a Young Alum: **Vanessa Murdock** (UMass Amherst CS M.S. '05, Ph.D. '06), Principal Applied Researcher, Microsoft Bing, Relevance Sciences Group. Her research there focuses on leveraging social media to improve local search. Previously Dr. Murdock worked at Yahoo! Research in Barcelona, Spain, first as a Research Scientist and then as a Senior Research Scientist. At Yahoo! Research she led the Geographic Context and Experience Group, setting the research agenda for topics related to geographic information and user-generated content, and securing nearly a million Euros in funding for her research group.

Outstanding Achievement in Management: **Mary-Ellen Prescott** (UMass Amherst CS B.S. '84), New Technology Program Manager at Bose Corporation. She leads cross-functional teams in the development of high quality-innovative noise cancelling, stereo Bluetooth wireless, and in-ear headphones, while providing leadership, vision, and direction in the execution of strategic goals. Prior to her current position at Bose, Ms. Prescott was a Deputy Program Manager at iRobot Corporation, where she led and managed teams in product development of robotics.

Outstanding Achievement in Research: **Zhi-Li Zhang** (UMass Amherst CS M.S. '92, Ph.D. '97), Qwest Chair Professor and McKnight Distinguished University Professor in the Department of Computer Science and Engineering at the University of Minnesota. Dr. Zhang developed a theory for Internet Quality-of-Service (QoS) and applied it to multimedia applications in the Internet including on-line video streaming. He developed foundational mathematical models for this as well as resource allocation algorithms for this purpose. He also developed a variety of innovative mechanisms such as video smoothing and pre-fetching to support on-line video streaming, and has also contributed to the development of resilient Internet routing algorithms.

Outstanding Support for the School: **EMC Corporation**, a global leader in enabling businesses and service providers to transform their operations and deliver IT as a service. Fundamental to this transformation is cloud computing. Through innovative products and services, EMC accelerates the journey to cloud computing, helping IT departments to store, manage, protect and analyze their most valuable asset — information — in a more agile, trusted and cost-efficient way. EMC has had a longstanding dedication to computer science education within the Commonwealth of Massachusetts and has been a particularly strong supporter of the research and teaching missions of the School of Computer Science at UMass Amherst.

More details on the OAA award recipients, with photos, are posted at **www.cs.umass.edu/oaa2014.**

During the OAA awards celebration, current students and recent alumni were also recognized. The graduate student awards are sponsored by **Yahoo!**, a member of the School's Industrial Affiliates Program (IAP). Professor Sridhar Mahadevan, Graduate Program Director, presented these **Outstanding Graduate Student Awards**:

Outstanding Doctoral Dissertation Award:
Akshat Kumar (Ph.D. '13) and **Chao Li** (Ph.D. '13)
Outstanding Synthesis Project Award:
Niall Emmart and **Fabricio Murai Ferreira**
Outstanding Teaching Assistant Award:
Aditya Sundarrajan and **Kyle Wray**

This year's undergraduate are sponsored by IAP member **Fiksu**. Professor Rod Grupen, Undergraduate Program Director, presented these **Outstanding Undergraduate Awards**:

Overall Academic Achievement: **David B. Lowell**
Overall Achievement in Research: **Daniel M. Stubbs**
Achievement in Networking: **Kyle R. Hughes**
Achievement in Interdisciplinary Studies: **Yael Kaufman**
Achievement in Computing Systems: **Patrick N. Pegus II**
Achievement in Software Engineering: **Elizabeth A. Staruk**
Achievement in Theory: **Sofya Vorotnikova**
Achievement in Robotics: **Jay Ming Wong**

Undergraduate student Fiksu award recipients
(l. to r.): Fiksu's Mark Corner, Prof. Rod Grupen, Yael Kaufman, Kyle Hughes, Patrick Pegus II, Elizabeth Staruk, Daniel Stubbs, Jay Ming Wong, Prof. Andrew McGregor, David Lowell
(not pictured: Sofya Vorotnikova)

Graduate student Yahoo! award recipients
(l. to r.): Aditya Sundarrajan, Kyle Wray, Niall Emmart, and Chao Li
(not pictured: Akshat Kumar and Fabricio Murai Ferreira)

CS alumni enjoy CA meet-up

On July 24, 2014, the School of CS hosted an alumni meet-up for UMass Amherst CS alumni (and student interns) at Google's Facility in Mountain View, CA. CS Chair Lori A. Clarke welcomed everyone to the event with an update on the state of the School, and Professor James Allan gave a presentation on "Who, Where, When, and What? Augmenting things for better search." Professors Leon Osterweil and Robbie Moll also traveled to Mountain View to meet and chat with alumni and interns. It was great to see so many of our alumni at the event: over 70 attended, which is the highest number to attend a CS alumni-organized west coast event! The CS alumni organizing committee consisted of Marc Cartright (Ph.D. '13), Jay Corbett (Ph.D. '93), and Naghi Prasad (Ph.D. '97). Visit **www.cs.umass.edu/alumnievent** for upcoming alumni events (plans are in the works for another CA event next summer). Be sure to join our LinkedIn and Facebook groups for event announcements. See photos page 17.

Alum Connections

Dr. **Carla Brodley** (Ph.D. '94), was appointed Dean of the College of Computer and Information Science at Northeastern University. Previously, she was a Professor of Computer Science at Tufts University and was Chair of Tufts Department of Computer Science from 2010 to 2013.

Alexander L. Wolf (Ph.D. '85) was elected ACM President for a two-year term beginning July 1st. Wolf, a Professor in Computing at Imperial College London, UK, was also awarded the 2014 SIGSOFT Outstanding Research Award "for formative work in software architecture, influential research in distributed event-based systems, and important contributions in software deployment, configuration management, and process."

Akshat Kumar (Ph.D. '13) was selected as the winner of the ICAPS (International Conference on Automated Planning and Scheduling) 2014 Best Dissertation Award. His dissertation "Exploiting Domain Structure in Multiagent Decision-Theoretic Planning and Reasoning" was completed in 2013. The award was presented at the 24th International Conference on Automated Planning and Scheduling that took place in Portsmouth, NH in June. Kumar is currently a faculty member at the School of Information Systems, Singapore Management University.

The Computing Research Association (CRA), in consultation with the National Science Foundation, appointed **Debra Richardson** (Ph.D. '81) as one of six new members of the Computing Community Consortium (CCC) Council. She began a three-year term on July 1st.

Dr. **Michael Franklin** (B.S. '83), Thomas M. Siebel Professor of Computer Science, was named the Chair of the Computer Science Division, Electrical Engineering and Computer Sciences, at the University of California Berkeley.

Bruno Ribeiro (Ph.D. '10) and **Benyuan Liu** (Ph.D. '03), current graduate student **Kun Tu**, Distinguished Professor **Don Towsley**, and co-authors collaborated on a research project involving studying the behavior of users of dating websites. Their paper "Who is Dating Whom: Characterizing User Behaviors of a Large Online Dating Site" was featured in *MIT Technology Review* and a number of other media sites. In another of Ribeiro's projects, partially completed while a Postdoc, a model that he developed predicts the growth and potential death of Facebook and other membership-based websites. Ribeiro notes that his model shows Facebook will not fail anytime soon.

Jennifer (Cotter) Cannon (B.S. '03) is a Software Engineer working at BAE Systems in Hudson, NH for the past 10 years. She received a Master's in Systems Engineering from Johns Hopkins University in 2012. She and her husband are the proud parents of a 10 month old son. "I am grateful for the education I received at UMass Amherst," notes Cannon.

CS is saddened to announce the passing of one of our alumni, **James Tung** (M.S. '69), who died on July 21, 2014 in Poughkeepsie, New York. He was a Systems Analyst for IBM and SAIC.

Jody Daniels accepts campus Distinguished Alumni Service Award

Jody Daniels (M.S. '93, Ph.D. '97) accepted the 2014 UMass Amherst Distinguished Alumni Award during a ceremony held at the Massachusetts State House in April. The Distinguished Alumni Awards are the most prestigious awards conferred by the Alumni Association upon its alumni, faculty, staff and friends.

Daniels is currently serving as a Brigadier General on active duty with the U.S. Africa Command (AFRICOM) in Stuttgart, Germany. She is the Deputy Director, Intelligence and Knowledge Development Directorate (J2), responsible for the intelligence and counterintelligence to facilitate situational awareness, warning and assessments on the threats across the 54 countries in the AFRICOM area of responsibility. We saw it fitting to share her award acceptance speech with our *Significant Bits* audience.

Distinguished guests, faculty, family, friends:

Thank you very much for this honor. I am quite surprised to be standing here today as I hadn't ever considered going to graduate school until my first Army boss asked why I hadn't applied. It had never crossed my mind. A few years later I decided to apply and UMass was an easy choice.

As I transitioned from being a full-time military intelligence officer to being a graduate student in the fall of 1990, the U.S. was just beginning the activities over in the Gulf that would eventually become Operation Desert Storm. As I was the only person that most of my fellow students knew to have served in the military, I was frequently called upon to explain "all things Gulf," "all things Army," and "all things military." I was suddenly thrust into the position of being the Expert on topics about which I knew very little. This caused me to pay a lot more attention to newspapers. So, while I was learning how to read journal and conference papers, I was also learning how to read and interpret the media.

It caused me to take a whole new view of how I looked at the world and my role in it. I realized I was also in a position of influence. I had to be far more careful and conscious with my words. I finally internalized the saying that "Words matter." This came to be ever more true as I reached higher civilian and military positions and more people relied on my words to guide their actions.

(l. to r.): UMass President Robert Caret, UMass Amherst Chancellor Kumble Subbaswamy, Jody Daniels '97, Alumni Association President Andrew MacDougall '95, and Assistant Vice Chancellor for Alumni Relations JC Schnabl.

I left UMass a whole lot wiser in ways I hadn't imagined. It wasn't just a degree or two; it was how to do research, how to frame arguments, how to make an impression, how to create a network. I learned how to listen better.

Those skills have served me well – regardless of whether I'm working with a team of scientists, a battalion of new recruits, or the National Chief of Intelligence for one of our African country partners.

I have had the pleasure of living a dual life – crossing frequently between my civilian and military careers. The overlap has not been as much in technical areas as it has been in communications, management, and leadership skills.

I was fortunate to have the strong academic background that the UMass Amherst Computer Science department gave me that allowed me to succeed at Lockheed Martin. I was fortunate that Lockheed Martin also valued my continuing service to the Army and the Army Reserve. I believe that the combination allowed me to migrate skills between the two careers and be a better leader than otherwise.

I'd like to thank all those faculty who helped me successfully complete the master's and Ph.D. programs, especially, Edwina Rissland, Bruce Croft, James Allan, Nick Belkin, Victor Lesser, Paul Utgoff, and Eliot Moss, but also Lori Clarke, Susan Landau, and Bev Woolf, who never knew it, but I looked upon as role models.

I'd also like to recognize and thank my fellow graduate students who forced me to better articulate my thoughts and arguments while they also supported and encouraged my dual track life. My family deserves much recognition and thanks for their endless support across everything I've done and continue to do.

Finally, I'd like to close with a quote from a leader and statesman who clearly recognized that words matter, and who had multiple careers – serving as the Chairman of the Joint Chiefs of Staff and then as Secretary of State – General Colin Powell: "Tell me what you know. Tell me what you don't know. Then tell me what you think. Always distinguish which is which." As a leader and as an intelligence professional, these are words I try to live by. Once again, thank you for this honor.

Lifetime email addresses for Computer Science alums

We are pleased to announce that CS alum email forwarding addresses are now available for any of our UMass Amherst CS alums. Although you might change employers or Internet providers, your CS alum email address will always stay the same and will forward your email to whatever address you choose. The email address will generally be **firstname.lastname@alum.cs.umass.edu.** To sign up for your email forwarding address, go to **www.cs.umass.edu/lifetime-email-forwarding**.

Elif Aktolga; *Integrating Non-Topical Aspects into Information Retrieval*; (James Allan, Advisor); May 2014; Research Engineer, Apple

When users investigate a topic, they are often interested in results that are not just relevant, but also strongly opinionated or covering a range of times. Often several queries need to be issued with reformulations if initial search results are not satisfactory. In this thesis, we focus on two non-topical dimensions: opinionatedness and time. For improving search results with respect to non-topical dimensions, we use diversification approaches. Results are diversified across a single or multiple non-topical dimensions. The burden of analyzing pre-existing biases for a query and discovering times at which important events happened is fully carried by the system. We show how to combine several dimensions with individual biases for each, while also presenting approaches to time and sentiment diversification. The insights from this work will be very valuable for next generation search engines and retrieval systems.

Jeff Dalton; *Entity-based Enrichment for Information Extraction and Retrieval*; (James Allan, Advisor); May 2014; Software Engineer, Google Inc.

The goal of this work is to leverage knowledge of the world to improve understanding of queries and documents using entities. An entity is a thing or concept that exists in the world, such as a politician, a battle, a film, or a color. Entity-based enrichment (EBE) is a new expansion model for both queries and documents using features from similar entity mentions in the document collection and external knowledge bases, such as Freebase and Wikipedia. With the ultimate goal of improving information retrieval effectiveness, we start from unstructured text and through information extraction, build up rich entity-based representations linked to external knowledge resources. We study the application of entity-based enrichment to improve the effectiveness of each step in the pipeline: 1) Named Entity Recognition, 2) Entity Linking, and 3) Ad hoc document retrieval. The empirical results for EBE in each of these tasks shows significant improvements.

Van Dang; *A Proportionality-based Approach to Search Result Diversification*; (W. Bruce Croft, Advisor); May 2014; Software Engineer, Google Inc.

Search result diversification addresses the problem of queries with unclear information needs by providing a document ranking that covers multiple possible topics for a given query. This increases the likelihood that users will find documents relevant to their specific intent. This dissertation introduces a new perspective on diversity: diversity by proportionality. We consider a result list more diverse, with respect to some set of query topics, when the ratio between the number of documents it provides for each topic matches more closely with the topic popularity distribution. Consequently, we derive a ranking framework for optimizing proportionality and an effectiveness measure. We also show that topical diversity can be achieved by diversifying search results using a set of terms that describe the query topics. This simplifies the task of finding a topic set to finding a term set. We present a technique and several data sources for generating these terms effectively.

Dan Gyllstrom; *Making Networks Robust to Component Failures*; (James Kurose, Advisor); May 2014; Senior Performance Engineer, Akamai Technologies

In this thesis, we consider instances of component failure in the Internet and in networked cyber-physical systems, such as the communication network used by the modern electric power grid (termed the smart grid). We design algorithms that make these networks more robust to various component failures, including failed routers, failures of links connecting routers, and failed sensors. This thesis divides into three parts: recovery from malicious or misconfigured nodes injecting false information into a distributed system (e.g., the Internet), placing smart grid sensors to provide measurement error detection, and fast recovery from link failures in a smart grid communication network.

Andrew Kae; *Incorporating Boltzmann Machine Priors for Semantic Labeling in Images and Videos*; (Erik Learned-Miller, Advisor); May 2104

Semantic labeling is the task of assigning category labels to regions in an image. For example, a scene may consist of regions corresponding to categories such as sky, water, and ground. Labeling regions allows us to better understand the scene itself as well as properties of the objects and their interactions within the scene. Typical approaches for this task include the conditional random field (CRF), which is well-suited to modeling local interactions among adjacent image regions. However the CRF may be limited in dealing with complex, global (long-range) interactions between regions in an image, and between frames in a video. This thesis presents ways to extend the CRF framework and incorporate priors based on the restricted Boltzmann machine (RBM) to model long-range interactions within images and video, for use in semantic labeling.

Tongping Liu; *Reliable and Efficient Multithreading*; (Emery Berger, Advisor); May 2014; Assistant Professor, Univ. of Texas San Antonio

To take advantage of multiple cores, software needs to be written using multithreading. It is notoriously far more challenging to write multithreaded programs correctly and efficiently than sequential ones. developed systems to combat both concurrency errors and performance issues in multithreaded programs. I developed Dthreads, a deterministic threading library that automatically ensures deterministic executions for unmodified C/C++ applications, without requiring programmer intervention or hardware support. Dthreads often matches or even exceeds the performance of standard thread libraries, making deterministic multithreading a practical alternative for the first time. I developed two other systems to attack false sharing, a performance issue that arises when multiple threads access distinct parts of the same cache line simultaneously. The first, Predator, not only precisely identifies but also predicts potential false sharing that does not get manifested. The second system, Sheriff-Protect, automatically eliminates false sharing inside parallel applications without programmer intervention.

Marwan Mattar; *Unsupervised Joint Alignment, Clustering and Feature Learning*; (Allen Hanson and Erik Learned-Miller, Advisors); May 2014; Research Data Scientist, Electronic Arts

Joint alignment is the process of transforming instances in a data set to make them more similar based on a pre-defined measure of joint similarity. This process has great utility in many scientific disciplines including radiology, psychology, and vision. This thesis takes steps towards developing an unsupervised data processing pipeline that includes alignment, clustering and feature learning. We first present an efficient curve alignment algorithm that is effective on many synthetic and real data sets. We show that using the byproducts of joint alignment, the aligned data and transformation parameters, can dramatically improve classification performance. We then incorporate unsupervised feature learning based on convolutional restricted Boltzmann machines to learn a representation that is tuned to the statistics of the data set. We show how these features can be used to improve both the alignment quality and classification performance. Finally, we present a nonparametric Bayesian joint alignment and clustering model which handles data sets arising from multiple modes.

Sameer Singh; *Scaling MCMC Inference and Belief Propagation to Large, Dense Graphical Models*; (Andrew McCallum, Advisor); May 2014; Postdoctoral Research Associate, Department of Computer Science, Univ. of Washington

In the past decade, single-core CPUs have given way to multi-core and distributed computing platforms. At the same time, access to large data collections is progressively becoming commonplace. Inference for probabilistic graphical models, that has been designed to operate sequentially, seems destined to become obsolete in this world of multi-core, multi-node systems. Further, modeling large datasets leads to an escalation in the number of variables, factors, domains, and the density of the models, all of which have a substantial impact on the computational complexity of inference. Motivated by the need to scale inference to large, dense graphical models, in this thesis we explore approximations to Markov chain Monte Carlo (MCMC) and belief propagation (BP) that induce dynamic sparsity in the model to utilize parallelism. These tools for inference enable us to tackle relation extraction, entity resolution, cross-document coreference, and other information extraction tasks over large text corpora.

CS Undergraduate Dean's List – Spring 2014

See full Spring 2014 Dean's List and list of Student Citations (students recognized for their outstanding performance in the classroom) at **www.cs.umass.edu/people/spring2014-deans-list-citations.**

CS graduates celebrated

On Friday, May 9, 2014, the School of Computer Science hosted M.S. and Ph.D. graduates for a post-graduation celebration luncheon. The following day, the undergraduate student graduation celebration luncheon was held. Over 400 CS student graduates, faculty, staff, families, and friends gathered to celebrate the new CS alumni.

UMASSCS SCHOOL OF COMPUTER SCIENCE 50 YEARS