

Fourth Annual Outstanding Achievement and Advocacy Award ceremony held

The department celebrated the accomplishments of this year's Outstanding Achievement and Advocacy (OAA) Award winners along with undergraduate and graduate student award recipients during a banquet held at the UMass Amherst Campus Center on Friday, May 4, 2012.

Department Chair Lori Clarke and Steve Goodwin, Dean of the College of Natural Sciences, welcomed the attendees. Robert Sproull, Chair of the Computer Science and Telecommunications Board of the U.S. National Academies, was the featured address speaker of the evening. Professor Leon J. Osterweil presented the 2012 OAA awards.

The 2012 OAA Award Recipients are:



- Outstanding Achievement by a Young Alum: **Justin Borgman** (B.S. '02) - Chief Executive Officer and Co-Founder of Hadapt



- Outstanding Support for the Department: **Fredrick W. Byron, Jr.** - retired Vice-Chancellor for Research of UMass Amherst; former Dean of the College of Natural Sciences and Mathematics



- Outstanding Achievement in Education: **Claire Cardie** (Ph.D. '94) - Professor in the Departments of Computer Science and Information Science at Cornell University and Co-Founder and Chief Scientist of Appinions

- Outstanding Achievement in Research: **Edmund Durfee** (Ph.D. '87) - Professor in the Department of Electrical Engineering and Computer Science, and in the School of Information, at the University of Michigan



- Outstanding Achievement in Entrepreneurship: **Andrew Merlino** (B.S. '85) - President and Founder of Pixel Forensics



- Outstanding Support for the Department: **Ken Schmidt** - Academic Relations Director for Yahoo!

- Outstanding Achievement in Research: **Henning Schulzrinne** (Ph.D. '92) - Julian Clarence Levi Professor of Computer Science at Columbia University and Chief Technology Officer of the Federal Communications Commission

- Outstanding Achievement in Technology Development: **Ted Selker** (M.S. '81) - Associate Director of Mobility Research and Visiting Professor at Carnegie Mellon University Silicon Valley



- Outstanding Contributions to Society: **Elliot Soloway** (Ph.D. '78) - Arthur F. Thurnau Professor in the College of Engineering, the School of Information and the School of Education at the University of Michigan

More details on the careers of the OAA award recipients, along with photos of the event, are posted at www.cs.umass.edu/oa2012.

During the evening's celebration, current students and recent alums were also recognized. The undergraduate and graduate awards were generously sponsored by Yahoo!. Professor Sridhar Mahadevan, Graduate Program Director, presented the following **Outstanding Graduate Student Awards**:

- Outstanding Dissertation Award: **Aruna Balasubramanian** (Ph.D. '11) - CIFellow, Department of Computer Science, University of Washington
- Outstanding Dissertation Award: **Timothy Wood** (Ph.D. '11) - Assistant Professor, Department of Computer Science, George Washington University
- Outstanding Synthesis Award: **Andres Molina-Markham**
- Outstanding Synthesis Award: **Robert Walls**
- Outstanding Teaching Assistant Award: **Brandon McPhail**
- Outstanding Teaching Assistant Award: **Philip Thomas**



(l. to r.) Andres Molina-Markham, Robert Walls, Philip Thomas, Aruna Balasubramanian, and Tim Wood with Yahoo!'s Ken Schmidt and Profs. Lori Clarke and Sridhar Mahadevan

Professor David Mix Barrington, Department Associate Chair, presented the following **Outstanding Undergraduate Awards** to students in this year's graduating class:

- Overall Achievement in Computer Science: **Evan Shelhamer**
- Achievement in Graphics: **Stephen Giguere**
- Achievement in Interdisciplinary Study: **Vinay Shah**
- Achievement in Interdisciplinary Study: **Derek Wood**
- Achievement in Machine Learning: **Tristan Warneke**
- Achievement in Security: **Jessica Ray**
- Achievement in Systems: **Marek Blat**
- Achievement in Theory: **Matthew Kelly**



(l. to r.): Evan Shelhamer, Matthew Kelly, Marek Blat, Jessica Ray, Derek Wood, and Stephen Giguere

Sawhney and Zhang named IEEE Fellows

Two CS alums, Harpreet Sawhney (Ph.D. '92) and Zhi-Li Zhang (Ph.D. '97), were named 2012 Fellows of the Institute of Electrical and Electronics Engineers (IEEE).

Sawhney, the Technical Director of the Vision and Learning Laboratory within SRI International's Vision Technology Group, was recognized for contributions to video algorithms. While at UMass Amherst, he was advised by Profs. Allen Hanson and Edward Riseman. Sawhney has extended the scope of vision technology at SRI International from basic sensor processing to recognition and reasoning about objects and activities in complex scenes. He led the conceptualization and realization of innovative algorithms in 2D and 3D processing. He introduced the idea of video object fingerprinting to tackle the difficult problem of object tracking across non-overlapping cameras. Sawhney has published more than 90 papers, holds 34 patents, and was named an SRI Fellow in 2011.

Zhang, Qwest Chair Professor in the Department of Computer Science and Engineering at the University of Minnesota, was recognized for his contributions to Internet quality-of-service theory and applications. Zhang's UMass Amherst advisors were Distinguished Professors Jim Kurose and Don Towsley. His research interests lie broadly in computer communication and networks, Internet technology, multimedia and emerging applications. While his past research was centered on the analysis, design and development of scalable Internet QoS solutions to support performance-demanding multimedia applications, his current research thrusts focus primarily on i) building *highly scalable, resilient and secure* Internet infrastructure and mechanisms to enhance Internet service availability, reliability and security; and on ii) developing *next-generation, service-oriented, manageable* Internet architectures to provide better support for creation, deployment, operations and management of value-added Internet services and underlying networks.

The IEEE Board of Directors confers the IEEE Grade of Fellow upon a person with an outstanding record of accomplishments in any of the IEEE fields of interest. The total number selected in any one year cannot exceed one-tenth of one percent of the total voting membership.

SAVE THE DATE – September 27

CS Alum Social in Cambridge, MA

Computer Science alums: You are cordially invited to a social gathering of CS alums and faculty to be held on Thursday, September 27, 2012 from 6:30 – 9:00 p.m. at Google in Kendall Square, Cambridge, MA. Check out www.cs.umass.edu/alumsocial2012 for more details as they become available. We will post to our LinkedIn and Facebook groups as well.



Victor Bahl (UMass Amherst Ph.D. '97), Director of the Mobile Computing Research Center in Microsoft Research, received a University of Massachusetts 2012 Distinguished Alumni Award at the Massachusetts State House on April 17. The Distinguished Alumni Award is the highest honor bestowed by the UMass Amherst Alumni Association on alumni, faculty, and friends; it honors those who have built

on their UMass Amherst experience to attain notable achievements in business, public or community service realms. On April 18, he gave a presentation in the CS department. He discussed "Putting the Cloud in the Palm of Your Hand."

Alexander L. Wolf (Ph.D. '85) was elected as Vice President of the Association for Computing Machinery (ACM). Wolf is a Professor in the Department of Computing at Imperial College London. He received our department's Outstanding Achievement in Research Award in 2010. More details: www.acm.org/news/featured/new-acm-officers.

Computing for Ordinary Mortals, by **Rob St. Amant** (Ph.D. '96), will be published by Oxford University Press in October, 2012. *Ordinary Mortals* is a popular science book for lay readers about important concepts in computing. St. Amant is an Associate Professor in the Computer Science Department at North Carolina State University.

Aruna Balasubramanian (Ph.D. '11) is the runner-up award winner for the ACM SIGCOMM Doctoral Dissertation award for 2011. Her thesis is entitled *Architecting Protocols to Enable Mobile Applications in Diverse Wireless Networks*. Balasubramanian is currently a Postdoctoral CIFellow Research Associate at the University of Washington.

Joseph Kiniry (M.S. '95), Associate Professor at the IT University of Copenhagen, is one of the PIs on "DemTech: Trustworthy Democratic Technology" (www.demtech.dk), one of the three major strategic grants awarded in Denmark in 2011. The grant will fund research that Kiniry started eight years ago focusing on educating governments and citizenry in the problems of electronic voting and showing how e-voting technology might be used and constructed, treating elections as critical systems. His previous work in e-voting led, in part, to the nationwide ban on the use of computers in elections in the Netherlands and Ireland. In September 2012 Kiniry will begin a new position as a Full Professor and Head of the Software Engineering Section at the Technical University of Denmark.

This spring, **Steve Willis** (B.S./BDIC '78), a consultant to financial, management, and technical startups, discussed his career in startups and technology over lunch with students in the department. Willis' discussion, "The Road from LISP and Ethernet's Promiscuous Mode to the NYSE," was organized by the UMass Amherst ACM student chapter.



Niranjan Balasubramanian; *Query Dependent Selection of Retrieval Alternatives*; (James Allan, Advisor); Sept. 2011; Post-Doctoral Researcher, Dept. of Computer Science & Engineering, University of Washington

This thesis investigates query-dependent selection of retrieval alternatives for information retrieval systems. Information retrieval research focuses on developing query representations and retrieval models for ranking documents. However, no single query representation or retrieval model performs the best for all queries. Selecting the best representation or retrieval model for each query can yield improved performance. This thesis addresses an important question in dynamically combining retrieval alternatives: *How to estimate the relative performance of the different alternatives for each query?* We treat query-dependent selection as a general problem of selecting between the result sets of different alternatives. We develop a relative effectiveness estimation technique that uses retrieval-based features in a learning formulation to directly predict differences between the results sets. We apply this general technique to select between alternative reduced versions for long queries and to combine multiple ranking algorithms. We extend the selection problem to include efficiency constraints and develop easy-to-compute features for efficient combination of retrieval models. Finally, we provide an extensive analysis of selection performance to better understand when query-dependent selection can be useful.



Alan Carlin; *Decision-Theoretic Meta-reasoning in Partially Observable and Decentralized Settings*; (Shlomo Zilberstein, Advisor); Feb. 2012; Research Engineer, Aptima

This thesis examines decentralized meta-reasoning. For a single agent or multiple agents, it may not be enough for agents to compute correct decisions if they do not do so in a timely or resource efficient fashion. By the time a decision is computed, the conditions leading up to it may have changed. The reasoning about one's computation process is referred to as meta-reasoning. Aspects of meta-reasoning considered in this thesis include the reasoning about how to allocate computational resources, including when to stop one type of computation and begin another, and when to stop all computation and report an answer. Given a computational model, this translates into computing how to schedule the basic computations that solve a problem.

This thesis constructs meta-reasoning strategies for the purposes of monitoring and control in multi-agent settings, specifically settings that can be modeled by the Decentralized Markov Decision Process (Dec-POMDP). It uses decision theory to optimize computation for efficiency in time and space in communicative and non-communicative decentralized settings. Whereas base-level reasoning describes the optimization of actual agent behaviors, the meta-reasoning strategies produced by this thesis dynamically optimize the computational resources which lead to the selection of base-level behaviors.



David Cooper; *Computational Affect Detection for Education and Health*; (Hava Siegelmann and Beverly Woolf, Advisors); Sept. 2011; Postdoctoral Researcher, Section of Biomedical Image Analysis, University of Pennsylvania

Emotional intelligence is a clear factor in education, health care, and day-to-day interaction. With the increasing use of computer technology, computers are interacting with more and more individuals. This interaction provides an opportunity to increase knowledge about human emotion for human consumption, well-being, and improved computer adaptation. This thesis explores the efficacy of using up to four different sensors in three domains for computational affect detection. We first consider computer-based education: four sensors are used to detect student emotions relevant to learning, such as frustration, confidence, excitement and interest during a computer tutoring session. The best emotion classifier accuracies range from 78% to 87.5%. We use voice data collected in a clinical setting to differentiate both gender and culture of the speaker. We produce classifier accuracies between 84% and 94% for gender, and between 58% and 70% for American vs. not American culture. Finally, we use video and audio in a health care education scenario to detect students' emotions during a clinical simulation evaluation. The video data provide classifiers with accuracies between 63% and 88% for learning relevant emotions. In total, this work is a step forward in the automatic computational detection of affect in realistic settings.



Gregory Druck; *Generalized Expectation Criteria for Lightly Supervised Learning*; (Andrew McCallum, Advisor); Sept. 2011; Research Scientist, Yummlly

Machine learning has facilitated many recent advances in natural language processing and information extraction. Unfortunately, most machine learning methods rely on costly labeled data, which impedes their application to new problems. Even in the absence of labeled data we often have a wealth of prior knowledge about these problems. For example, we may know which labels particular words are likely to indicate for a sequence labeling task, or we may have linguistic knowledge suggesting probable dependencies for syntactic analysis. This thesis focuses on incorporating such prior knowledge into learning, with the goal of reducing annotation effort for information extraction and natural language processing tasks. We advocate constraints on expectations as a flexible and interpretable language for encoding prior knowledge. We focus on the development of Generalized Expectation (GE), a method for learning with expectation constraints and unlabeled data. We explore the various flexibilities afforded by GE criteria, derive efficient algorithms for GE training, and relate GE to other methods for incorporating prior knowledge into learning. We then use GE to develop lightly supervised approaches to text classification, dependency parsing, sequence labeling, and entity resolution that yield accurate models for these tasks with minimal human effort. We also consider the incorporation of GE into interactive training systems that actively solicit prior knowledge from the user and assist the user in evaluating and analyzing model predictions.



Gary B. Huang; *Weakly Supervised Learning for Unconstrained Face Processing;* (Erik Learned-Miller, Advisor); May 2012; Scientist, Howard Hughes Medical Institute

Machine face recognition has traditionally been studied under the assumption of a carefully controlled image acquisition process. By controlling image acquisition, variation due to factors such as pose, lighting, and background can be largely eliminated. Applications of face recognition have had mixed success when deployed in conditions where the assumption of controlled image acquisition no longer holds. This dissertation focuses on this unconstrained face recognition problem, where face images exhibit the same amount of variability that one would encounter in everyday life.

We formalize unconstrained face recognition as a binary pair matching problem (verification), and present a data set for benchmarking performance. We observe that it is comparatively much easier to obtain many examples of unlabeled face images than face images that have been labeled with identity or other higher level information. We thus focus on improving face verification by leveraging the information present in this source of weakly labeled data. We show how unlabeled face images can be used to perform unsupervised face alignment and unsupervised feature discovery, both leading to improvements in recognition accuracy. We additionally combine the two methods, leading to an unsupervised alignment system with verification accuracy matching supervised alignment.



Pallika Kanani; *Resource-bounded Information Acquisition and Learning;* (Andrew McCallum, Advisor); May 2012; Senior Research Staff Member, Oracle Labs

In many scenarios it is desirable to augment existing data with information acquired from an external source. For example, information from the Web can be used to fill missing values in a database or to correct errors. In many machine learning problems, acquiring additional feature values can lead to improved data quality and accuracy. However, there is often a cost associated with information acquisition, and we typically have resource limitations. In this thesis, I explore Resource-bounded Information Acquisition and Learning. The process of acquiring information from an external source involves multiple steps, such as deciding what subset of information to obtain, locating the documents that contain the required information, acquiring relevant documents, extracting the specific piece of information, and combining it with existing information to make useful decisions. Resource-bounded Information Acquisition (RBIA) involves saving resources at each stage of the information acquisition process. I explore four special cases of the RBIA problem in real-world domains. Finally, I propose a general framework for RBIA, that considers the state of the database at each time point, dynamically adapts to the results of previous acquisition steps and properties of the next ones, and perform them to acquire most information with least resources.



Dov Katz; *Interactive Perception of Articulated Objects for Autonomous Manipulation;* (Oliver Brock, Advisor); Sept. 2011; Postdoctoral Associate, The Robotics Institute, Carnegie Mellon University

This thesis develops robotic skills for manipulating novel articulated objects. Examples of everyday articulated objects include scissors, pliers, door handles, books, and drawers. The degrees of freedom of an articulated object describe the relationship among its rigid bodies, and are often relevant to the object's intended function. Autonomous manipulation of articulated objects is therefore a prerequisite for many robotic applications in our everyday environments. In structured environments robots accomplish manipulation tasks with impressive accuracy and speed. In contrast, when object models are not available, in unstructured environments such as our homes and offices, manipulation remains largely unsolved. We assume that to enable autonomous manipulation of objects in our everyday environments, robots must be able to acquire information about these objects. Acquiring information about the world from sensor data is a challenging problem. There is simply too much information that could be perceived. Solving such a general perception problem is infeasible. Instead, we propose to leverage our understanding of the task, in order to determine what information is relevant for manipulation. The main contribution of this thesis is in introducing Interactive Perception, a new perceptual framework, which exploits the synergies that arise when crossing the boundary between action and perception.



David Mimno; *Topic Regression;* (Andrew McCallum, Advisor); Feb. 2012; Postdoctoral Research Associate, Department of Computer Science, Princeton University

Text documents are generally accompanied by non-textual information, such as authors, dates, publication sources, and, increasingly, automatically recognized named entities. Work in text analysis has often involved predicting these non-text values based on text data for tasks such as document classification and author identification. This thesis considers the opposite problem: predicting the textual content of documents based on non-text data. In this work, I study several regression-based methods for estimating the influence of specific metadata elements in determining the content of text documents. Such topic regression methods allow users of document collections to test hypotheses about the underlying environments that produced those documents.



Hala Mostafa; *Exploiting Structure in Coordinating Multiple Decision Makers;* (Victor Lesser, Advisor); Sept. 2011; Research Scientist, BBN Technologies

This thesis is concerned with sequential decision making by multiple agents, whether they are cooperative or self-interested. The practical intractability of the general problem led to efforts in identifying special cases that admit efficient computation, yet still represent a wide enough range of problems. In our work, we identify the class of problems with structured interactions, where actions of one agent can have non-local effects on the transitions and/or rewards of another agent.

We addressed the following research questions: 1) How can we compactly represent this class of problems?; 2) How can we efficiently calculate agent policies that maximize team reward (for cooperative agents) or achieve equilibrium (self-interested agents)?; and 3) How can we exploit structured interactions to make reasoning about communication offline tractable? We developed a new decision-theoretic model, Event-Driven Interactions with Complex Rewards, that explicitly represents structured interactions. We looked at the issue of communication in both cooperative and self-interested settings. For self-interested agents, we studied an example setting where communication is an integral part of problem solving, but where the self-interested agents have a reason to be reticent (e.g. privacy concerns). We formulate this problem as a game of incomplete information and present a general algorithm for calculating approximate equilibrium profile in this class of games.



Megan Olsen; *Variations on Stigmergic Communication to Improve Artificial Intelligence and Biological Modeling;* (Hava Siegelmann, Advisor); Sept. 2011; Assistant Professor, Dept. of Computer Science, Loyola University Maryland

Stigmergy refers to indirect communication that was originally found in biological systems. It is used for self-organization by ants, bees, and flocks of birds, by allowing individuals to focus on local information. Through local communication among individuals, larger patterns are formed without centralized communication. This self-organization is just one type of system studied within complex systems, and is generally referred to as emergent behavior. My thesis examines emergent behavior in artificial intelligence and biology: the growth of cancer, population dynamics, emotions, multi-agent fault tolerance, and real-time strategic AI for games. I design variations on stigmergy to accomplish my two goals: a) to develop novel computational models of complex biological systems, and b) to tackle key AI research questions by proposing new algorithms and techniques that are inspired by those complex biological systems. My contributions are a new agent-based cancer growth model, a proposed use of localized communication for removing cancer cells, improved multi-agent fault tolerance through localized messaging, a new approach to modeling predator-prey dynamics using computational emotions, and improved strategic game AI through computational emotions.



Jangwon Seo; *Search Using Social Media Structures;* (W. Bruce Croft, Advisor); Sept. 2011; Software Engineer, Google

Social applications on the Web have appeared as communication spaces for sharing knowledge and information. In particular, social applications can be considered valuable information sources. We address methods for finding relevant information in social media applications that use unique properties of these applications. Specifically, we focus on three unique structures in social media: hierarchical structure, conversational structure, and social structure. These structures are designed to organize information and encourage people to participate in discussions in social applications. To exploit these structures in retrieval frameworks, we introduce an effective representation framework for

multiple contexts. We then discuss how to discover or define each structure and how to extract relevant contexts from the structure. Using the representation framework, these relevant contexts are integrated into retrieval algorithms. In addition, we address two other challenges related to social media search--identification of relevant substructures in retrieval objects and analysis of text reuse structures.



Timothy Wood; *Improving Data Center Resource Management, Deployment, and Availability with Virtualization;* (Prashant Shenoy, Advisor); Sept. 2011; Assistant Professor, Dept. of Computer Science, The George Washington University

The increasing demand for storage and computation has driven the growth of large data centers--the massive server farms that run many of today's Internet and business applications. A data center can comprise many thousands of servers and can use as much energy as a small city. The massive amounts of computation power contained in these systems results in many interesting distributed-systems and resource-management problems. This thesis investigates challenges related to data centers, with a particular emphasis on how new virtualization technologies can be used to simplify deployment, improve resource efficiency, and reduce the cost of reliability, all in application-agnostic ways. We focus on three of the key challenges faced when building and running data centers: 1) simplifying the deployment of applications by providing models of virtualization overheads and new placement strategies that allow for greater consolidation; 2) automating performance management techniques that exploit the virtualization abstraction layer to dynamically allocate resources; and 3) providing high reliability services to applications even in the face of site-wide disasters.



Xing Yi; *Discovering and Using Implicit Data for Information Retrieval;* (James Allan, Advisor); Sept. 2011; Scientist, Yahoo! Labs

In real-world information retrieval (IR) tasks, the searched items and/or the users' queries often have implicit information associated with them--information that describes unspecified aspects of the items or queries. This indirectly available information has been shown to improve search effectiveness. However, in many real-world IR challenges this information is incomplete or missing in a large portion of the data. We present a language-modeling based general perspective for discovering implicit information and demonstrate how to use the discovered data for search. We investigate four specific IR challenges: (1) finding relevant records in semi-structured databases where many records contain incomplete or empty fields; (2) searching web pages that have little or no associated anchor text; (3) using click-through records in web query logs to help search pages that have no or very few clicks; and (4) discovering plausible geographic locations for web queries that contain no explicit geographic information. Intuitively, we use contextual similarity among data for discovering implicit information for search. We empirically demonstrate the effectiveness of our approaches using different IR challenges. Our research shows that supporting information discovery tailored to different search tasks can enhance IR systems' search performance.



Chongjie Zhang; *Scaling Multi-Agent Learning in Complex Environments*; (Victor Lesser, Advisor); Sept. 2011; Postdoctoral Research Associate, Dept. of Computer Science, UMass Amherst

Cooperative multi-agent systems (MAS) are finding applications in a wide variety of domains, including sensor networks, robotics, distributed control, collaborative decision support systems, and data mining. A cooperative MAS consists of a group of autonomous agents that interact with one another in order to optimize a global performance measure. A central challenge in cooperative MAS research is to design distributed coordination policies. Existing techniques are inadequate to address this challenge in large-scale complex applications with uncertain

environments. This dissertation develops a new multi-agent reinforcement learning (MARL) paradigm for tackling this challenging coordination problem. To scale up the learning, this learning paradigm exploits interaction locality and non-local information in a coherent way. Using this paradigm, agents concurrently learn their policies based on local observations, and, meanwhile, their learning processes are coordinated by a non-local control mechanism to ensure the global learning performance. This dissertation presents both efficient algorithms for multi-agent learning (MAL) with limited observability and a scalable control framework for coordinating MAL. I have applied and evaluated this learning paradigm in diverse problem domains, including distributed task allocation, network routing, and sensor networks.

CS Undergraduate Dean's List – Spring 2012

Alpert, Warren	Giguere, Stephen J.	Muehlberg, Jacob	Simms, Jeremy R.
Avery, Sean T.	Gonda, Peter S.	Murphy, Brendan Edward	Simon, Timothy
Bach, Michael Curtis	Gordon, Rachel Clare	Nakashian, Louis E.	Sims, Robert J.
Briggs, Priscilla E.	Hamza, Mazen Samer	Naparstek, Nicolas	Singer, Henry Oskar
Bristol, Alexander Lawrence	Hardy, Dustin Power	Neilson, Tyler J.	Stapleton, Brian Fitzpatrick
Bronk, Harold R.	Hayes, James Patrick	Nguyen, Anh Tung	Staruk, Elizabeth A.
Burkatovskiy, Aleksandr Valeriy	Hebert, Mitchell Ryan	Nolin, Brian J.	Stearns, Brett C.
Cadwell-Frost, Amanda Jean	Hingley, Jesse A.	Noran, Sean Stephen	Stewart, Ian Ajameian
Chadowitz, Chaniel	Hoffman, Scott A.	Olejnik, Katarzyna Lucja	Stubbs, Daniel M.
Chan, Andrew	Hong, John	Otsuka, Maxwell Sho	Sullivan, Dylan Patrick
Cheang, Wai	Hughes, Kyle Robert	Ottaviano, Kenneth Robert	Szeto, Ryan Bing-Shue
Chen, Chengyang	Ioannou, Nicolas	Pegus, Patrick	Tang, Roger Z.
Chudy, Joshua Roman	Itkin, David Jonas	Pelletier, Jesse N.	Trask, Joshua H.
Ciollaro, Christopher James	Izotov, Andrey Kirillovich	Peris, Sahil Loy Joe	Truong, Kevin
Conlon, Nicholas	Kaufman, Yael	Perodeau, Timothy Lawrence	Tsang, Ivan
Connell, Ryan J.	Lee, Anna	Perreault, Jason Edward	Turlapati, Hridya
Daole-Wellman, George W.	Lewis, Patrick Joseph	Pezzzone, Jeffrey Michael	Vorotnikova, Sofya
Deroche, Michael A.	Li, Lily	Pham, Tung Thanh	Warneke, Tristan Tamas
DeRoy, Michael K.	Ling, Yan	Powers, Kenneth Wendell	Williams, Albert Beckman
Droeske, Trevor Michael	Lowell, David B.	Raitto, Caleb L.	Wong, Ann L.
Duggan, Daniel Joseph	Masi, Daniel Patrick	Rivais, Elaina	Wong, Jay Ming
Edwards, Gen Michael	May, Nathaniel David	Rudovol, Ivan	Wong, Lp Tommy
Eid, Jacob Imad	McAndrew, Jeffrey Douglas	Rutter, James J.	Wongwajarachot, Victor
Filipek, Alexander William	McMahon, Samuel Robert	Rynderman, Mordechai	Young, Cody J.
Fox, John B.	Moh, Heng Li Vahid Anthony	Schuler, Joachim D.	Yuen, Lit Wa

CS undergrads awarded course citations

Course citations recognize outstanding performance of undergraduate students who have significantly exceeded the standards for an A grade in a particular course. The list

of spring 2012 course citation award recipients is below. The course for which the student received the citation is noted after his/her name.

Ayoub, Richard (187)	Gonda, Peter (187)	Muehlberg, Jacob (240)	Stubbs, Daniel (401)
Bjorge, Michael (187)	Higgins, Lauren (121)	Noran, Sean (187 and 250)	Szeto, Ryan (250)
Blat, Marek (445)	Hughes, Kyle (311 and 320)	Otsuka, Maxwell (121)	Tenaglia, Steven (240)
Burkatovskiy, Aleksandr (187)	Izotov, Andrey (445)	Ottaviano, Kenneth (121)	Vorotnikova, Sofya (401)
Chen, Chengyang (187)	Jacek, Nicholas (401)	Pegus, Patrick (240)	Whittemore, Max (187)
Crespi, Tyler (187)	Kelly, Matthew (445)	Powers, Kenneth (305)	Williams, Albert (250)
Deroche, Michael (320)	Lee, Anna (240)	Richardson, Samuel (121)	
Droeske, Trevor (187)	Lewis, Patrick (311)	Sprague, Brandon (187)	
Dunford, Andrew (187)	Lowell, David (240)	Stapleton, Brian (240)	