

# ALUM Matters

*A newsletter for alums of the Department of Computer Science*

## Jacobs studies biological and artificial intelligence

**M**ultidisciplinary research seems to be the rage these days; a search of the Internet with Google yields 698,000 hits for the keywords “National Science Foundation” and “multidisciplinary.” University of Rochester Professor Robert Jacobs (UMass Amherst CS Ph.D. ’90) has been pursuing multidisciplinary research on biological and artificial intelligence for many years. He has developed new machine learning algorithms and new computational models of human cognition in a variety of domains, including perception, motor control, and memory organization.

While at UMass Amherst, Jacobs was advised by Professor Andrew Barto. Afterwards, he served as a postdoctoral fellow in the lab of Professor Michael Jordan at the Department of Brain & Cognitive Sciences at the Massachusetts Institute of Technology. He then was a postdoctoral fellow in the lab of Professor Stephen Kosslyn at the Department of Psychology at Harvard University. He currently is a faculty member at the University of Rochester where he serves as Professor of Brain

& Cognitive Sciences, of Computer Science, and of Vision Science.

His research activities have been shaped by the belief that biological and artificial intelligence should be studied in tandem. “No one discipline has all the answers regarding the nature of cognition and intelligence, though different disciplines have yielded different insights and perspectives,” says Jacobs. “To obtain a complete understanding of intelligence, we’ll need to work across traditional disciplinary boundaries. The payoff from this work will be both technological innovations that improve our lives and also new ways of thinking about our own minds.”

One question that has always fascinated Jacobs is whether human cognition is “optimal” in some computational or mathematical sense. He has addressed this question in several research projects.

For his Ph.D. dissertation, he considered whether it is preferable to have a highly modular computational system in which different modules perform different tasks, or

if it is preferable to have a single monolithic system that performs all tasks. He studied this issue through the development of a “mixtures of experts” system, consisting of multiple learning devices, such as multiple artificial neural networks. The system’s learning algorithm uses a competitive learning scheme to adaptively partition the training

data so that different devices learn different subsets of data items. His simulation results showed that there are significant advantages

to highly modular systems in terms of learning speed. At the same time, however, there are also disadvantages because such systems require special mechanisms to integrate information to produce robust performances in novel circumstances.

This finding led him to think about how people integrate information from different information sources. A convenient area to study this issue is the domain of perception where we perceive our environments through multiple sensory modalities, such as vision, audition, and touch. In one set of experiments, Jacobs and his colleagues examined whether people combine visual and auditory information in a statistically optimal manner for the purpose of spatial localization. Previous research suggested that people exhibit “visual capture” —if visual and auditory information indicate slightly different locations for an event, people will localize the event to the location indicated by the visual modality, a phenomenon known as the “ventriloquist effect.” This outcome reflects the fact that people are more accurate at making spatial judgments based on visual information than on auditory information. What would happen if people were placed in a virtual-reality environment in which



A subject using a visual-haptic virtual reality environment. The subject sees virtual objects rendered in the head-mounted display, and grasps these objects using the force-feedback devices connected to his thumb and index fingers.

visual information was highly “noisy”? The experimental data collected by Jacobs and his colleagues indicates that people start averaging the locations indicated by visual and auditory modalities, with the auditory location being assigned a larger weight when the visual information is corrupted with noise. These data can be accounted for by an optimal statistical model known as a Kalman filter, indicating that people in these studies integrated the information provided by their visual and auditory modalities in an optimal manner.

More recently, Jacobs has been examining whether human perceptual learning is optimal. Imagine that each time you saw and grasped an object, your (unconscious) estimate of the object’s depth based on haptic (touch) information was always 10% larger than your estimate based on visual information. If you had reason to believe that your haptic estimates were more reliable, would you adapt your visual system so that, in the future, its interpretations were more consistent with those of your haptic system? Jacobs and his colleagues used a novel virtual reality environment to address this question (see photo on opposite page). People in their experiments wore a head-mounted display which allowed them to see “virtual” objects. In addition, people’s thumbs and index fingers were connected to robot arms comprising a haptic force-feedback device which allowed them to grasp these objects. The experimental data indicate that when there is a discrepancy between people’s visual and haptic percepts, and when the visual information is noisy, people recalibrate their visual systems so that these systems’ interpretations are more consistent with those of their haptic systems. These data support the hypothesis that people continuously monitor the interpretations of their sensory systems, and keep them calibrated by adapting each system whenever its interpretations fail to match the interpretations of other systems.

Jacobs continues to examine the optimality of people’s perceptual learning. “When quantitatively comparing people’s learning performances with those of state-of-the-art machine learning devices using the same data sets, it often seems as if people do not learn as much from each data item as they theoretically could,” adds Jacobs. “Why do people show sub-optimal learning?”

Jacobs is now examining different training procedures to see which ones elicit the largest learning effects. Finding effective training procedures for perceptual skills has enormous importance for science education. For example, radiologists need to learn to visually distinguish tumors

Ph.D. alums **Peter Desnoyers** and **Charles Sutton** received the 2007-2008 Awards for an Outstanding Doctoral Dissertation, sponsored by Yahoo!. Desnoyer’s dissertation was titled “Distributed Sensing Networks: Archiving, Indexing, and Querying,” and Sutton’s was “Local Training for Conditional Random Fields.” Both received their Ph.D.s in 2008.

*IEEE Intelligent Systems* selected **Jennifer Neville** (Ph. D. ’06) as one of “AI’s 10 to Watch.” The feature on the 10 top researchers was in the May/June issue. Neville is an Assistant Professor of Computer Science and Statistics at Purdue University. Her research interests are in machine learning and data mining.



**Amy McGovern** (Ph.D. ’02) received a five year NSF CAREER grant for her project “Developing Dynamic Relational Models to Anticipate Tornado Formation.” McGovern is an Assistant Professor in the School of Computer Science at the University of Oklahoma.

**Balaraman (“Ravi”) Ravindran** (Ph.D. ’04) was awarded tenure and a promotion to Associate Professor in the Department of Computer Science and Engineering, Indian Institute of Technology Madras.

Center for Intelligent Information Retrieval alums **Jeremy Pickens** (Ph.D. ’04) and **Chirag Shah** (M.S. ’06) along with co-authors Gene Golovchinsky, Pernilla Qvarfordt, and Maribeth Back received the Best Paper Award at the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR ’08) for their paper “Algorithmic Mediation for Collaborative Exploratory Search.” Pickens is a Research Scientist at FX Palo Alto Labs and Shah is a Ph.D. student at the University of North Carolina, Chapel Hill.

from other tissue by viewing mammograms, geologists need to learn to visually recognize different types of rock samples, astronomers need to learn to visually interpret stellar spectrograms, and biologists need to learn to visually identify different cell structures. Although much is known about how to train people to learn new facts, defining effective training procedures for teaching people new perceptual skills is an open area of research.

## Outstanding CS alum awards

Department Chair Andy Barto has announced a new Awards Program aimed at recognizing outstanding achievements by alums of the department. Awards will be presented at an annual UMass Amherst Computer Science Awards Banquet; the first is planned for the weekend of May 1, 2009. The awards will recognize achievement in such areas as entrepreneurship, scientific research, and education.

Details such as award categories, the nomination process, and the date, place, and time of the awards banquet, will be forthcoming. CS Alums are encouraged to stay in touch and

participate by making a nomination and/or attending the awards banquet.

Barto also announced steps to make it easier for CS Alums to keep up with department news. In addition to the biannual Significant Bits newsletter, and the web site, [www.cs.umass.edu](http://www.cs.umass.edu), the department has also set up a LinkedIn group (UMass Amherst Dept. of Computer Science), and a Facebook group (UMass CS) that all CS Alums are encouraged to join. The LinkedIn and Facebook groups are expected to be a forum for discussions and news of particular interest to CS Alums. Be sure to check them out!