in the study of natural and formal languages, industrial design, and the way in which we organize almost any kind of knowledge you can think of. Information propagates up the hierarchy as input triggers features in the lower level and their output serves as the input to higher-level features. Information also propagates down the hierarchy as expectations prepare us to see what makes sense in the context of what we've already felt, seen and heard.

David Mumford at Brown and his former student Tai Sing Lee now at CMU have proposed a model of the visual cortex as a hierarchical Bayesian statistical model in which recurrent feedforward and feedback loops serve to integrate top-down contextual priors and bottom-up observations. The Lee and Mumford model offers a proposal for how the cortex is organized, but not how it becomes organized; it seems likely that much of the organization of the cortex is learned. Indeed, it seems we have to be exposed to specific kinds of stimuli at particular times during our development in order to learn certain features of our environment. For example, cats raised in an environment devoid of horizontal lines will grow up unable to see or learn to see horizontal features if later exposed to a richer environment – imagine the poor cat in specially designed room with an Elizabethan collar so it can't even see itself.

"Reactions to this work range from, 'the idea is ludicrous' to the more insightful, 'what's new here will allow us to make progress on problems such as machine vision that have eluded us for decades.'"

Research in machine learning makes it pretty clear that some kinds of learning are hard, unsupervised learning in particular. For instance, learning to classify objects without being shown clearly labeled training data, is devilishly hard. While it seems plausible that we learn some concepts in an unsupervised fashion, for many other learning tasks, it appears we have evolved to take advantage of environmental cues that provide a form of distal supervision. One source of such supervision comes from experiment; we learn that some foods are good to eat and others bad by tasting them. But what about learning to perceive?

Continued on page 14

Ph.D. Profiles

Stefan Roth

Before coming to Brown in the fall of 2001, Stefan received his degree in computer science and engineering from the University of Mannheim in Germany. At Brown he has worked with Michael Black on a number of projects in computer vision and related fields. Inspired by two internships at Intel Research in their Computational Nanovision group, he became interested in modeling the statistics of natural images and the applications of probabilistic image models to image reconstruction problems. Stefan's Ph.D. work also considers the problem of modeling and estimating image motion (optical flow). His paper, "On the Spatial Statistics of Optical Flow" (with Michael Black) received Honorable Mention for the Marr Prize in 2005, which is considered the most prestigious award in the area of computer vision.

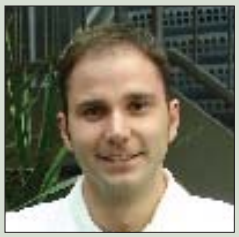

Nikos Triandopoulos

Nikos, a recipient of the Department's Paris Kanellakis fellowship, received his undergraduate diploma in computer engineering and informatics at the University of Patras, Greece, and entered the Ph.D. program at Brown in 2000. His primary research areas are information security, cryptography, and algorithms. His Ph.D. thesis, titled "Efficient Data Authentication" under the supervision of Roberto Tamassia, studies theoretical and practical aspects of the problem of authenticating information in hostile and adversarial computing environments. As a member of the Center for Geometric Computing, and in a collaborative effort with IAM Technologies, Inc., he received the Technological Innovation award from Brown, and his work has led to several patent applications. For the 2006-2007 academic year, he has been awarded an I3P-funded postdoctoral fellowship, and will be working on algorithms for high assurance in cyber-security at the Institute of Security Technology Studies at Dartmouth.

Danfeng Yao

Danfeng Yao received her B.S. degree from Beijing University, China. In 2002, she started her Ph.D. studies at Brown on information security under advisor Roberto Tamassia. She has a number of publications on trust management and user privacy. Danfeng holds a provisional patent on her federated ID management work with Roberto and Professor Michael Goodrich at UC, Irvine. In 2005, she visited CERIAS at Purdue University, where she worked with Professor Mike Atallah and Elisa Bertino on multi-party computation and data integrity. In 2006, she interned in the Trusted Systems Lab at HP Labs, Princeton, NJ. At HP, Danfeng worked closely with Dr. Stuart Haber, and developed security protocols for outsourced systems that have resulted in two HP technical reports. **C!**