Electrical & Computer ENGINEERING

SPECIAL SEMINAR

WEDNESDAY
OCTOBER 10, 2007

Scaife Hall Auditorium
Room 125

12:30 p.m.
Refreshments—12:00 p.m.

PROF. Fei-Fei Li
PRINCETON UNIVERSITY

Prof. Fei-Fei Li's main research interest is in vision, particularly high-level visual recognition. In computer vision, Fei-Fei's interests span from object and natural scene categorization to human activity categorizations in both videos and still images. In human vision, she has studied the interaction of attention and natural scene and object recognition. In a recent project, she also studies the human brain fMRI activities in natural scene categorization by using pattern recognition algorithms. Fei-Fei graduated from Princeton University in 1999 with a physics degree, and a minor in engineering physics. She received her PhD in electrical engineering from the California Institute of Technology in 2005. Fei-Fei was on faculty in the Electrical and Computer Engineering Dept. at the University of Illinois Urbana-Champaign (UIUC) from Sept 2005 to Dec 2006. Starting Jan 2007, Fei-Fei is an Assistant Professor in the Computer Science Department at Princeton University. She also holds courtesy appointments in the Psychology Department and the Neuroscience Program at Princeton. She is a recipient of the 2006 Microsoft Research New Faculty Fellowship. (Fei-Fei publishes under the name L. Fei-Fei.)

ECE Seminar Hosts:
Radu Marculescu,
radum@ece.cmu.edu
Yi Luo,
y1827@andrew.cmu.edu
Bruno Sinopoli,
brunos@andrew.cmu.edu

Natural Scene Recognition: From Humans to Computers

For both humans and machines, the ability to learn and recognize the semantically meaningful contents of the visual world is an essential and important functionality. In this talk, we will examine the topic of natural scene categorization and recognition in human psychophysical and physiological experiments as well as in computer vision modeling.

I will first present a series of recent human psychophysics studies on natural scene recognition. All these experiments converge to one prominent phenomena of the human visual system: humans are extremely efficient and rapid in capturing the semantic contents of the real-world images. Inspired by these behavioral results, we report a recent fMRI experiment that classifies different types of natural scenes (e.g. beach vs. building vs. forest, etc.) based on the distributed fMRI activity. This is achieved by utilizing a number of pattern recognition algorithms in order to capture the multivariate nature of the complex fMRI data.

In the second half of the talk, we present a generative Bayesian hierarchical model that learns to categorize natural images in a weakly supervised fashion. We represent an image by a collection of local regions, denoted as codewords obtained by unsupervised clustering. Each region is then represented as part of a `theme'. In previous work, such themes were learnt from hand-annotations of experts, while our method learns the theme distribution as well as the codewords distribution over the themes without such supervision. We report excellent categorization performances on a large set of 13 categories of complex scenes.

Carnegie Mellon
Electrical & Computer ENGINEERING