Title: A Probabilistic Theory of Deep Learning

Speaker: Richard Baraniuk, Rice University

Abstract
A grand challenge in machine learning is the development of computational algorithms that match or outperform humans in perceptual inference tasks that are complicated by nuisance variation. For instance, visual object recognition involves the unknown object position, orientation, and scale in object recognition while speech recognition involves the unknown voice pronunciation, pitch, and speed. Recently, a new breed of deep learning algorithms have emerged for high-nuisance inference tasks that routinely yield pattern recognition systems with near- or super-human capabilities. But a fundamental question remains: Why do they work? Intuitions abound, but a coherent framework for understanding, analyzing, and synthesizing deep learning architectures has remained elusive. We answer this question by developing a new probabilistic framework for deep learning based on the Deep Rendering Model: a generative probabilistic model that explicitly captures latent nuisance variation. By relaxing the generative model to a discriminative one, we can recover two of the current leading deep learning systems, deep convolutional neural networks and random decision forests, providing insights into their successes and shortcomings, a principled route to their improvement, and new avenues for exploration.

BIO

Richard G. Baraniuk is the Victor E. Cameron Professor of Electrical and Computer Engineering at Rice University. His research interests lie in new theory, algorithms, and hardware for sensing, signal processing, and machine learning. He is a Fellow of the IEEE and AAAS and has received national young investigator awards from the US NSF and ONR, the Rosenbaum Fellowship from the Isaac Newton Institute of Cambridge University, the ECE Young Alumni Achievement Award from the University of Illinois, the Wavelet Pioneer and Compressive Sampling Pioneer Awards from SPIE, the IEEE Signal Processing Society Best Paper Award, and the IEEE Signal Processing Society Technical Achievement Award. His work on the Rice single-pixel compressive camera has been widely reported in the popular press and was selected by MIT Technology Review as a TR10 Top 10 Emerging Technology. For his teaching and education projects, including Connexions (cnx.org) and OpenStax College (openstaxcollege.org), he has received the C. Holmes MacDonald National Outstanding Teaching Award from Eta Kappa Nu, the Tech Museum of Innovation Laureate Award, the Internet Pioneer Award from the Berkman Center for Internet and Society at Harvard Law School, the World Technology Award for Education, the IEEE-SPS Education Award, the WISE Education Award, and the IEEE James H. Mulligan, Jr. Medal for Education.