

University of Chicago and Toyota Technological Institute at Chicago
Machine Learning Seminar Series

PRESENTS:



Yali Amit
UChicago Statistics

Title: Optimization of latent variables in deep network applications

Abstract: The primary use of deep networks is to provide direct prediction of response variables, whether discrete or continuous, following a forward pass through the network. In some cases however, we know that the values of certain well defined but unobserved latent variables are key to successful prediction. For example the scale or rotation of an object in an image. In deep learning the typical solution is to provide extensive (also known as augmented) training sets where the expected range of values of the latent variables is well represented. In generative modeling it is more common to define a statistical model with a distribution of the observed data conditional on some unobserved latent variables, and online optimization or simulation w.r.t to the latent variables is required when computing likelihoods. I will describe some experiments, with 'deformable classifiers' where we train a deep network together with an optimization step over predefined latent variables, and require this optimization to also be performed online during test time. I will show that this enables learning with much smaller data sets, at the cost of more intensive computation, and provides as output not just the class label but the optimal instantiation of the latent variables for the test example. I will also show some application of these ideas for training generator networks.

Bio: Yali Amit is Professor of Statistics and Computer Science at the University of Chicago. He fields of interest are computer vision, machine learning and computational neuroscience.

Host: Rebecca Willett

Wednesday, May 15, 1:00 – 2:00 pm
Harper Center (Booth) Room 219
Pizza provided by UChicago CS Department