From Probabilistic Models to Decision Theory and Back Again*

Sanmi Koyejo

Abstract: Decision making is a fundamental aspect of probabilistic inference, and involves selecting a point estimate from a probabilistic model e.g. a prediction or a parameter estimate, which optimizes some measure of utility. In the first part of the talk, I will outline some a new results on binary and multilabel prediction with complex performance measures, such as F-measure and Jaccard measure. Perhaps surprisingly, the prediction which maximizes the utility for a range of such measures takes a particularly simple form as the thresholding of the probability of the positive class. This result motivates simple but effective classifiers for "brain reading" i.e. predicting cognitive processes from fMRI data.

In the second part of the talk, I will discuss a "dual" of the decision making problem i.e. estimating a probabilistic model that satisfies the user determined utility. This approach to probabilistic inference may be employed for incorporating potentially useful prior information encoded in the user-devised risk function. One useful application of the resulting procedure is for incorporating structural assumptions on the support of latent variables. I will outline some new results which show that the support constrained information projection is submodular, thus can be efficiently optimized using standard techniques. Of particular interest is where the the support constraints can be represented as a matroid, resulting in new techniques for probabilistic (variable/group/tree) sparse regression. I will discuss a recent application of these techniques to develop a novel sparse probabilistic canonical correlation analysis for the joint analysis of fMRI statistical maps and behavioural measurements.

Bio: Sanmi (Oluwasanmi) Koyejo an engineering research associate in the Poldrack Lab at Stanford University and an Assistant Professor (starting Fall 2016) in the Department of Computer Science at the University of Illinois at Urbana-Champaign. Koyejos research involves the development and analysis of probabilistic and statistical machine learning techniques motivated by, and applied to various modern big data problems, with a particular focus on analysis of large scale biological data such as neuroimaging and genetics data. Koyejo completed his Ph.D in Electrical Engineering at the University of Texas at Austin under the supervision of Joydeep Ghosh and was a postdoc with Russell Poldrack and Pradeep Ravikumar. He has been the recipient of several awards including the outstanding NCE/ECE student award, a best student paper award from the conference on uncertainty in artificial intelligence (UAI) and a trainee award from the Organization for Human Brain Mapping (OHBM).

*Machine Learning External Seminar, Gatsby Unit, July 5, 2016.