Statistical Sciences
UNIVERSITY OF TORONTO

Computational Foundations of Bayesian Inference and Probabilistic Programming

Daniel Roy, University of Cambridge

Thursday February 13, 2014 at 3:30pm

Sidney Smith Hall, Room 1074

Refreshments will be served at 3:15pm

The complexity, scale, and variety of data sets we now have access to have grown enormously, and present exciting opportunities for new applications. Just as high-level programming languages and compilers empowered experts to solve computational problems more quickly, and made it possible for non-experts to solve them at all, a number of high-level probabilistic programming languages with computationally universal inference engines have been developed with the potential to similarly transform the practice of Bayesian statistics. These systems provide formal languages for specifying probabilistic models compositionally, and general algorithms for turning these specifications into efficient algorithms for inference.

In this talk, I will address three key questions at the theoretical and algorithmic foundations of probabilistic programming---and probabilistic modeling more generally---that can be answered using tools from probability theory, computability and complexity theory, and nonparametric Bayesian statistics. Which Bayesian inference problems can be automated, and which cannot? Can probabilistic programming languages represent the stochastic processes at the core of state-of-the-art nonparametric Bayesian models? And if not, can we construct useful approximations? I'll close by relating these questions to other challenges and opportunities ahead at the intersections of computer science, statistics, and probability.