



THE UNIVERSITY OF CHICAGO

Department of Statistics

STATISTICS COLLOQUIUM

DAVID BLEI

Department of Statistics and Computer Science
Columbia University

Scaling and Generalizing Variational Inference

MONDAY, May 18, 2015, at 4:00 PM

Eckhart 133, 5734 S. University Avenue

Refreshments following the seminar in Eckhart 110.

ABSTRACT

Latent variable models have become a key tool for the modern statistician, letting us express complex assumptions about the hidden structures that underlie our data. Latent variable models have been successfully applied in numerous fields.

The central computational problem in latent variable modeling is posterior inference, the problem of approximating the conditional distribution of the latent variables given the observations. Posterior inference is central to both exploratory tasks and predictive tasks. Approximate posterior inference algorithms have revolutionized Bayesian statistics, revealing its potential as a usable and general-purpose language for data analysis.

Bayesian statistics, however, has not yet reached this potential. First, statisticians and scientists regularly encounter massive data sets, but existing approximate inference algorithms do not scale well. Second, most approximate inference algorithms are not generic; each must be adapted to the specific model at hand.

In this talk I will discuss our recent research on addressing these two limitations. I will describe stochastic variational inference, an approximate inference algorithm for handling massive data sets. I will demonstrate its application to probabilistic topic models of text conditioned on millions of articles. Then I will discuss black box variational inference.

For further information and about building access for persons with disabilities, please contact Kirsten Wellman at 773.702.8333 or send email (kwellman@galton.uchicago.edu). If you wish to subscribe to our email list, please visit the following website: <https://lists.uchicago.edu/web/arc/statseminars>.

Black box inference is a generic algorithm for approximating the posterior. We can easily apply it to many models with little model-specific derivation and few restrictions on their properties. I will demonstrate its use on a suite of nonconjugate models of longitudinal healthcare data.

This is joint work based on these two papers:

M. Hoffman, D. Blei, J. Paisley, and C. Wang. Stochastic variational inference. *Journal of Machine Learning Research*, 14:1303-1347.

R. Ranganath, S. Gerrish, and D. Blei. Black box variational inference. *Artificial Intelligence and Statistics*, 2014.