Statistical and computational trade-offs in Bayesian learning*

Tamara Broderick

Abstract

The flexibility, modularity, and coherent uncertainty estimates provided by Bayesian posterior inference have made this approach indispensable in a variety of domains. Since posteriors for many problems of interest cannot be calculated exactly, much work has focused on delivering accurate posterior approximations—though the computational cost of these approximations can sometimes be prohibitive, particularly in a modern, large-data context. Focusing on unsupervised learning problems, we illustrate in a series of examples how we can trade off some typical Bayesian desiderata for computational gains and vice versa. On one end of the spectrum, we sacrifice learning uncertainty to deliver fast, flexible methods for point estimates. In particular, we consider taking limits of Bayesian posteriors to obtain novel K-means-like objective functions as well as scalable, distributed algorithms. On the other end, we consider mean-field variational Bayes (MFVB), a popular and fast posterior approximation method that is known to provide poor estimates of parameter covariance. We develop an augmentation to MFVB that delivers accurate estimates of posterior uncertainty for model parameters.

* Machine Learning External Seminar, Gatsby Unit, December 16, 2015.