Fast Quantification of Uncertainty and Robustness with Variational Bayes

Tamara Broderick, Massachusetts Institute of Technology

In Bayesian analysis, the posterior follows from the data and choices of both prior and likelihood. These choices may be somewhat subjective and reasonably vary over some range. Thus, we wish to measure the sensitivity of posterior estimates to variation in these choices. While the field of robust Bayes has been formed to address this problem, its tools are not commonly used in practice. We demonstrate that variational Bayes (VB) techniques are readily amenable to fast robustness analysis. Since VB casts posterior inference as an optimization problem, its methodology is built on the ability to calculate derivatives of posterior quantities with respect to model parameters. We use this insight to develop local prior robustness measures for mean-field variational Bayes (MFVB), a particularly popular form of VB due to its fast runtime on large data sets. A potential problem with MFVB is that it has a well-known major failing: it can severely underestimate uncertainty and provides noinformation about covariance. We generalize linear response methods from statistical physics to deliver accurate uncertainty estimates for MFVB---both for individual variables and coherently across variables. We call our method linear response variational Bayes (LRVB).