

Fast Robustness Quantification with Variational Bayes

Tamara Broderick, Massachusetts Institute of Technology

In Bayesian analysis, the posterior follows from the data and a choice of a prior and a likelihood. These choices may be somewhat subjective and reasonably vary over some range. Thus, we wish to measure the sensitivity of the posterior 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. An important contributing reason for this lack is the difficulty of calculating robustness measures from MCMC draws, which often lack generality or require additional coding or computation. We demonstrate that, by contrast to MCMC, variational Bayes (VB) techniques are readily amenable to 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, even in very complex models. We use this insight to develop local prior robustness measures for mean-field variational Bayes (MFVB), a particularly popular form of VB.