Being an Informed Bayesian: Assessing Prior Informativeness and Prior–Likelihood Conflict

Xiao-Li Meng

Harvard University

Abstract

Dramatically expanded routine adoption of the Bayesian approach has substantially increased the need to assess both the confirmatory and contradictory information in our prior distribution, in reference to the information provided by our likelihood. Our diagnostic approach starts with the familiar posterior matching method; for a given likelihood model, we identify the difference in the sample sizes needed to form two likelihood functions that, when combined respectively with a given prior and a baseline prior, will lead to the same posterior summaries as chosen. This difference can be viewed as a "prior data size" M(k), relative to the likelihood based on k independent, identically distributed observations. The confirmatory information is captured by the M(k) function, which is roughly constant over k when no serious prior-likelihood conflict arises. The contradictory information is detectable in its derivative or finite difference as M(k) tends to decrease with k when contradictory prior specification detracts information from the likelihood. Intriguing findings include a universal low bound, -1, on the derivative of M(k) that represents the most extreme prior-likelihood conflict, and a super-informative phenomenon where the prior effectively gains an extra 50% prior data size relative to the baseline when the prior mean coincides with the truth. We demonstrate our method via several examples, including an application exploring the effects of immunoglobulin levels on lupus nephritis. We also establish theoretical results showing why the derivative of M(k) is a useful indicator for prior-likelihood

conflict. (This is joint work with Matthew Reimherr and Dan Nicolae of The University of Chicago.)