

14 From a Conditional Lindley’s Paradox to Block Hyper-g Priors

Chris Hans (The Ohio State University)

Thick-tailed mixtures of g priors have gained traction as a default choice of prior distribution in Bayesian regression settings. The motivation for these priors usually focuses on properties of model comparison and variable selection (e.g., asymptotic consistency of model posterior probabilities), as well as computational considerations. Standard mixtures of g priors mix over a single, common scale parameter that shrinks all regression coefficients in the same manner, and the particular form of the mixture distribution determines the model comparison properties. In this paper we focus on the effect of the mono-shrinkage induced by mixing over a single scale parameter and propose new mixtures of g priors that allow for differential shrinkage across collections of regression coefficients. We introduce a new “conditional information asymptotic” that is motivated by the common data analysis setting where at least one regression coefficient is much larger than others. We analyze existing mixtures of g priors under this limit and reveal two new behaviors, “Essentially Least Squares (ELS)” estimation and the “Conditional Lindley’s Paradox (CLP)”, and argue that these behaviors are, in general, undesirable. As the driver behind both of these behaviors is the use of a single, latent scale parameter that is common to all coefficients, we propose a block hyper- g prior, defined by first partitioning the covariates into groups and then placing independent hyper- g priors on the corresponding blocks of coefficients. We provide conditions under which ELS and the CLP are avoided by the new class of priors, and provide consistency results under traditional sample size asymptotics. This is joint work with Agniva Som (Duke University) and Steven MacEachern (Ohio State University).

The discussant can find a tech report related to this talk at <http://arxiv.org/abs/1406.6419>.