

Poster Snapshot

Posterior Contraction Rate and Asymptotic Bayes Optimality for One Group Global-Local Shrinkage Priors in Sparse Normal Means Problem

Sayantana Paul (ISI)

Consider the normal means model under sparsity, where the level of sparsity is unknown. We investigate some optimality properties of inference on the mean vector using a broad class of one-group global-local shrinkage priors including the horseshoe. Such a class of priors was considered earlier in Ghosh and Chakrabarti (2017). We show that both the empirical Bayes and the full Bayes posterior distribution of the mean vector contracts around the true mean vector as well as the empirical Bayes estimate at a near minimax rate with respect to the squared L_2 loss. These results generalize those of van der Pas et al. (2014, 2017) for a broad class of priors.

We next consider simultaneous hypothesis testing for the coordinates of the mean vector to infer which ones are actually non-zero. We investigate asymptotic decision theoretic optimality of a testing rule based on a full Bayes approach using our chosen class of one-group priors when the prior on the global shrinkage parameter is allowed to be chosen arbitrarily on a suitable support. Under the asymptotic framework of Bogdan et al. (2011), we prove that the Bayes risk of our full-Bayes decision rule asymptotically attains the risk of the Bayes Oracle defined in Bogdan et al. (2011), up to some multiplicative constant. This is the first optimality result of its kind in this problem in the full Bayes approach. Our result is adaptive in nature and reinforces the argument that appropriately designed one-group shrinkage priors can be very reasonable alternatives to two-groups priors for inference in sparse problems.

This is a joint work with Dr. Arijit Chakrabarti (Indian Statistical Institute, Kolkata)