

2 Bayesian Estimation of Sparse Signals with a Continuous Spike-and-Slab Prior

Veronika Rocková (University of Pennsylvania)

We introduce a new framework for estimation of normal means, bridging the gap between popular frequentist strategies (LASSO) and popular Bayesian strategies (spike-and-slab). The main thrust is to introduce the family of Spike-and-Slab LASSO (SS-LASSO) priors, which form a continuum between the Laplace prior and the point-mass spike-and-slab prior. We establish several appealing frequentist properties of SS-LASSO priors, contrasting them with these two limiting cases. To this end, we adopt the penalized likelihood perspective on Bayesian modal estimation and introduce the framework of “Bayesian penalty mixing” with spike-and-slab priors. Of primary interest to us are very non-concave penalties, yielding potentially very multimodal posteriors. Similarly as the LASSO mode, we show that the SS-LASSO global posterior mode is near-minimax rate-optimal under squared error loss (with suitable penalties). Going further, we show that the whole posterior “keeps pace with the global mode” and concentrates at the near-minimax rate, a property that is known not to hold for the single Laplace prior. Furthermore, minimax-rate optimality is obtained with a suitable class of independent product priors (for known levels of sparsity) as well as with dependent mixing priors (adapting to the unknown levels of sparsity). Up to now, the rate-optimal posterior concentration has been established only for spike-and-slab priors with a point mass at zero. Thus, the SS-LASSO priors, despite being continuous, possess similar optimality properties as the “methodologically ideal” point-mass mixtures. These results provide valuable theoretical justification for our proposed class of priors, underpinning their intuitive appeal and practical potential.