StatMathAppli2017 - September 2017

Frequentist coverage in Bayesian and post-selection restricted space problems

Kevin Duisters & Johannes Schmidt-Hieber
Mathematical Institute, Leiden

The frequentist coverage of Bayesian credible sets for restricted parameter space problems of the type $\theta > \lambda \in \mathbb{R}$, endowed with a flat prior, has been widely studied in the sequence model. In particular for symmetric and log-concave error densities g, $1 - 3\alpha/2$ is known to be a sharp lower bound for the frequentist coverage of the $(1-\alpha)$ HPD credible set. We describe the non-trivial behaviour of this coverage in the case of restrictions of the type $|\theta| > \lambda > 0$ and prove a $1 - 3\alpha/2$ lower bound for a richer class of densities g as corollary. We apply this machinery in two settings. Firstly, we relate these HPD credible sets to post-selection confidence sets in (high-dimensional) $l1$-regularized sequence models, which are derived by conditioning g on observing $|Y| > \lambda > 0$. We prove in general that the frequentist coverage of any θ_0 by the $1-\alpha$ HPD credible set equals the (conditional $|Y| > \lambda$) probability of the random variable Y falling in the post-selection confidence set constructed when observing $y = \theta_0$. From a more practical point of view, a numerical study of the post-selection coverage is made for several choices of g,λ and α. Despite the different targets of inference, we may compare this coverage to that of the debiased LASSO and Bayesian HPD with restricted flat prior in this way. Secondly, in ongoing work, we investigate how the frequentist behaviour of Bayesian uncertainty quantification with sparse priors can be described using these tools.

This is joint work with Johannes Schmidt-Hieber.