Finite sample bound for the Bernstein–von Mises theorem

The University of Tokyo, Keisuke Yano
The University of Tokyo, Kengo Kato

1. Abstract

This talk brings a contribution to the Bernstein–von Mises theorem and its application. The Bernstein–von Mises theorem states that the posterior distribution is approximated by the Gaussian distribution of which the mean is the maximum likelihood estimate and the covariance matrix is Fisher information matrix. Focusing on linear regression model in high dimensions, a finite sample bound of the total variation distance between a posterior distribution and the Gaussian distribution is derived. Our analysis has the following two features: The model is possibly misspecified; The variance of the error distribution is possibly unknown. Applying the finite sample bound, frequentist coverage probabilities of Bayesian credible sets are evaluated. Fast convergence of the coverage probabilities is discovered.

2. Problem setting

We consider a linear regression model

$$Y = X\beta_0 + \varepsilon,$$

where $Y = (Y_1, \dots, Y_n)^\top \in \mathbb{R}^n$ is a vector of response variables, $X \in \mathbb{R}^{n\times p}$ is a design matrix, $\beta_0 \in \mathbb{R}^p$ is an unknown coefficient vector, and $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)^\top \in \mathbb{R}^n$ is a vector of i.i.d. error variables with mean zero and variance $\sigma_0^2 \in (0, \infty)$. In our analysis, we do not impose the assumption that the true distribution of ε is Gaussian.

Fitting a Gaussian distribution on the error ε, we obtain the likelihood

$$L(\beta, \sigma^2; Y) = \frac{1}{(2\pi\sigma^2)^{n/2}} \mathrm{e}^{-\|Y - X\beta\|^2/(2\sigma^2)}.$$

Putting a prior distribution on β and σ^2, we obtain the marginal posterior distribution of β. We derive the finite sample bound of the total variation distance between the marginal distribution and the Gaussian distribution and evaluate the coverage probabilities of Bayesian credible rectangles.