Dimension Dependence of Sampling Algorithms in Hierarchical Bayesian Inverse Problems

Sergios Agapiou, University of Warwick
S.Agapiou@warwick.ac.uk

We consider the inverse problem of recovering u from noisy observations of $\mathcal{A}^{-1}u$, where $\mathcal{A}^{-1} : \mathcal{X} \to \mathcal{Y}$ is an injective bounded operator and $\mathcal{X}, \mathcal{Y}$ separable Hilbert spaces. We model this as $y = \mathcal{A}^{-1}u + \xi$, where ξ is Gaussian noise, $\xi \sim \mathcal{N}(0, \lambda^{-1}\mathcal{C}_1)$, for some positive scaling λ^{-1} modelling the noise level. We adopt a widely used Hierarchical Bayesian approach assuming a Gaussian prior on the unknown, $u \sim \mathcal{N}(0, \delta^{-1}\mathcal{C}_0)$ for some positive scaling δ^{-1}, and Inverse Gamma hyperpriors, $\lambda \sim \Gamma(\alpha_\lambda, \beta_\lambda),\ \delta \sim \Gamma(\alpha_\delta, \beta_\delta)$. We are interested in the case where $\mathcal{X}, \mathcal{Y}$ are infinite-dimensional, however, in such cases in practice the statistical inference is implemented in $\mathbb{R}^N$. Working in $\mathbb{R}^N$, we can apply the Bayes rule to get the *posterior distribution* $\mathbb{P}(u, \lambda, \delta|y)$ which is the object of interest in the Bayesian setting. The choice of distributions is conjugate to the model giving an explicit form for the conditional distributions $\mathbb{P}(u|\lambda, \delta, y)$, $\mathbb{P}(\lambda|u, \delta, y)$, $\mathbb{P}(\delta|u, \lambda, y)$, and making natural the use of the Gibbs sampler to sample the posterior. The goal of this talk is to explain the dimension dependence of the Gibbs sampler. In particular, we show that as $N \to \infty$ the chain mixes rapidly in the scale of the inverse noise covariance, but poorly with respect to the scale of the inverse prior covariance. We propose a reparametrization of the problem which is stable with respect to the increase in dimension. This is joint work with John Bardsley (Montana), Omiros Papaspiliopoulos (Barcelona) and Andrew M. Stuart (Warwick).