Return your written solutions either in person or by email
to vesa.kaarnioja@fu-berlin.de by Monday 4 July, 2022, 12:15

1. Assume that you have a Gaussian posterior distribution

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \sim \pi^y \sim \mathcal{N}\left( \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & p \\ p & 1 \end{pmatrix} \right).$$

   (a) Write a Gibbs sampler for the posterior $\pi^y$. Based on the generated samples, what are the conditional mean estimate of $\pi^y$ and the marginal standard deviations of $x_1$ and $x_2$?

   (b) Repeat part (a) for parameter values $p = 0.5, 0.9, 0.99$, and $0.999$. How does the degree of correlation between $x_1$ and $x_2$ affect the performance of the Gibbs sampler?

2. Suppose we have an inverse problem

$$y = \begin{pmatrix} x_1^2 + x_2^2 \\ x_2 \end{pmatrix} + \eta,$$

where $y \in \mathbb{R}^2$, $x = (x_1, x_2)^{\mathrm{T}} \in \mathbb{R}^2$. Let us set the prior $x = z \cdot \mathbf{1}_{[-4,4]^2}(z)$, where $z \sim \mathcal{N}((0,0)^{\mathrm{T}}, I)$,

$$\mathbf{1}_B(z) = \begin{cases} 1, & z \in B, \\ 0, & \text{otherwise}, \end{cases}$$

and $\eta \sim \mathcal{N}(0, \delta^2 I)$ with $\delta = 0.1$. Suppose we are given the observation $\bar{y} = (7, -2)$. Implement MCMC with Metropolis–Hastings kernel

$$x_{k+1} \sim \sqrt{1 - \beta^2} \cdot x_k + \beta \xi, \quad \xi \sim \mathcal{N}(0, I),$$

for different values of $\beta \in (0,1)$ to sample the posterior density. For each value of $\beta$ produce 10 000 samples and plot them. What do you notice? Also compute for each $\beta$ the *acceptance ratio*, i.e., the ratio between accepted jumps and the total length of the chain. Use the origin as initial value.

Using the best choice of $\beta$, compute the expectation of the posterior, i.e., the conditional mean estimate

$$x_{\mathrm{CM}} = \int_{\mathbb{R}^2} x \pi^y(x) \, \mathrm{d}x.$$