Advanced Bayesian Methods: Theory and Applications in R

02-MCMC - Exercises

This exercise is about implementing a slice sampler for a general log-posterior.

Use the `cars` dataset:

```
R> head(cars)

  speed dist
1     4    2
2     4   10
3     7    4
4     7   22
5     8   16
6     9   10
```

Estimate the following log-linear model:

$$\log(\texttt{dist}) = \beta_0 + \beta_1 \log(\texttt{speed}) + \varepsilon,$$

where $\varepsilon \sim N(0, \sigma^2)$.

The full log-posterior for $\boldsymbol{\beta} = (\beta_0, \beta_1)$ and σ^2 is the sum of the log-likelihood and log-prior.

1. **Log-likelihood**:

$$\log p(\mathbf{y}|\mathbf{X}, \boldsymbol{\beta}, \sigma^2) = -\frac{n}{2}\log(2\pi\sigma^2) - \frac{1}{2\sigma^2}(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})^\top(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})$$

2. **Log-prior for** $\beta|\sigma^2$ (Normal prior):

$$\log p(\beta|\sigma^2) = -\frac{p}{2}\log(2\pi\sigma^2) - \frac{1}{2\sigma^2}(\boldsymbol{\beta} - \mathbf{m}_0)^\top \mathbf{M}_0^{-1}(\boldsymbol{\beta} - \mathbf{m}_0)$$

3. **Log-prior for** σ^2 (Inverse Gamma prior):

$$\log p(\sigma^2) = a\log(b) - \Gamma(a) - (a+1)\log(\sigma^2) - \frac{b}{\sigma^2}$$

The full log-posterior is:

$$\log p(\boldsymbol{\beta}, \sigma^2|\mathbf{y}, \mathbf{X}) = -\frac{n}{2}\log(2\pi\sigma^2) - \frac{1}{2\sigma^2}(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})^\top(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})$$
$$-\frac{p}{2}\log(2\pi\sigma^2) - \frac{1}{2\sigma^2}(\boldsymbol{\beta} - \mathbf{m}_0)^\top \mathbf{M}_0^{-1}(\boldsymbol{\beta} - \mathbf{m}_0)$$
$$+a\log(b) - \Gamma(a) - (a+1)\log(\sigma^2) - \frac{b}{\sigma^2}$$

1. **Implement the Slice Sampler**: Write a function in R to sample from this log-posterior. Hint instead of sampling uniform in $(0, 1)$, remove height from the log-posterior using `rexp(1)`.
2. **Estimate the Model**: Use the `cars` dataset to estimate $\boldsymbol{\beta}$ and σ^2. Visualize the estimated function on the orignal scale of the data.