

Lab 6: Gibbs sampling with block updates

Turning in solutions

This lab is part of Homework 5. Solutions to the exercises, as well as the non-lab homework exercises are to be written up and uploaded to Gradescope as a PDF.

Getting started

You will need the following R packages. If you do not already have them installed, please do so first using the `install.packages` function.

```
library(mvtnorm)
library(coda)
```

Gibbs sampling with block updates

One of the problems with Gibbs sampling is that it moves very slowly when posterior variables are highly correlated. This lab explores this issue. Suppose you fit a Bayesian model to a set of data and your posterior includes three variables X , Y , and Z whose joint posterior $\pi(X, Y, Z)$ is multivariate normal with

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} \sim \mathcal{N}_3 \left[\boldsymbol{\theta} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \Sigma = \begin{pmatrix} 1 & 0.9 & 0.1 \\ 0.9 & 1 & 0.1 \\ 0.1 & 0.1 & 1 \end{pmatrix} \right].$$

For deriving the conditional distributions needed in the questions below, you should refer to Chapter 7 of Hoff for the form of conditional normal distributions for any given multivariate normal distribution. See, in particular, the beginning of Section 7.4.

-
1. Given the multivariate normal distribution above, what are the posterior complete conditionals for X , Y , and Z ? That is, derive $\pi(X|Y, Z)$, $\pi(Y|X, Z)$, and $\pi(Z|X, Y)$. Note that you should have three univariate normal distributions.
 2. Write a Gibbs sampler that alternates updating each of the variables. You can set the initial values for all three variables to 0 and the number of mcmc samples to 1,000. Provide a trace plot and an autocorrelation plot of the draws for either X or Y . Comment on the plots.
 3. One option for dealing with this high correlation is doing **block updates**, where multiple variables are updated at once. Give the conditional distributions for $(X, Y)|Z$ and $Z|(X, Y)$. Note that you should have one bivariate normal distribution and one univariate normal distribution.
 4. Write a Gibbs sampler using the conditional distributions in Exercise 3 above, where X and Y are updated together (using a random draw from a bivariate normal), alternating with Z being updated. You can once again set the initial values for all three variables to 0 and the number of mcmc samples to 1,000. Provide a trace plot and an autocorrelation plot of the draws for either X or Y , and comment on the plots.
 5. Comment on the difference between the performance of the two Gibbs samplers. Why is the second more efficient?
