Computational Statistics
7. exercise session, 9.4.2010

1. Let the target distribution be $N(0, 1)$ and suppose we use the random walk Metropolis algorithm, where the distribution of the increment w in the updating step $\theta' \leftarrow \theta + w$ is $N(0, 1)$. Suppose we initialize the chain with the value $\theta^{(0)} = -40$ and that the proposed value is $\theta' = -40 + 1.1$.

a) Calculate $\pi(\theta)$, $\pi(\theta')$, $q(\theta' \mid \theta)$ and $q(\theta \mid \theta')$ using double precision floating point numbers (which is the accuracy used by R). Calculate the M–H ratio by plugging the four values in eq. (7.2).

b) This time, calculate first the logarithm of the M–H ratio (by calculating values of log-densities) and then exponentiate. What is the the value of the next state $\theta^{(1)}$ in a properly implemented Metropolis–Hastings sampler?

2. (Boundary reflection in random walk M–H). Suppose the parameter space is $(0, \infty)$ and that the proposal θ' is calculated as follows when the current state is θ. Here g is some pdf defined on the whole real line.

1. Draw increment w from pdf g.
2. Set $\theta' \leftarrow |\theta + w|$.

Find an expression for the M–H ratio. Show that this expression simplifies, when g is a symmetric density. (Hint: in order to find the density $q(\theta' \mid \theta)$, first find the cdf of the proposal distribution, and then calculate the pdf by differentiating.)

3. Suppose the parameter space is $(0, \infty)$ and that the proposal θ' is calculated as follows when the current state is θ.

1. Set $\phi \leftarrow \ln(\theta)$.
2. Draw v from $N(0, \sigma^2)$, and set $\phi' \leftarrow \phi + v$.
3. Set $\theta' \leftarrow \exp(\phi')$.

Calculate a formula for the proposal density $q(\theta' \mid \theta)$, and write down the formula for the M–H ratio.

4. Suppose that we have a target density $\pi(\theta)$ of a continuous distribution, which is non-zero only on $(0, \infty)$. We decide to reparametrize the statistical model using $\phi = \ln(\theta)$, and to use random walk proposal for ϕ using $N(0, \sigma^2)$ as the distribution of the increment.

Calculate a formula for the M–H ratio. What is the relationship of your result with the M–H ratio that you calculated in the previous problem?

5. Derive the Gibbs sampler for the two-variate normal distribution $N(\mu, \Sigma)$ truncated to the positive quadrant, which has the unnormalized density

$$p(x_1, x_2) \propto N\left(x \mid \mu, \Sigma\right) 1_+(x_1)\, 1_+(x_2),$$

where $x = (x_1, x_2)$, and 1_+ is the indicator function of the set $(0, \infty)$, i.e.,

$$1_+(u) = \begin{cases} 1, & \text{if } u > 0 \\ 0, & \text{otherwise.} \end{cases}$$

Hint: the covariance matrix can be written as

$$\Sigma = \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix}, \qquad -1 < \rho < 1, \quad \sigma_1, \sigma_2 > 0.$$

and the mean vector μ as (μ_1, μ_2). The non-truncated bivariate normal density can factorized as

$$\begin{aligned} N((x_1, x_2) \mid \mu, \Sigma) &= N(x_1 \mid \mu_1, \sigma_1^2)\; f_{2|1}(x_2 \mid x_1) \\ &= N(x_2 \mid \mu_2, \sigma_2^2)\; f_{1|2}(x_1 \mid x_2), \end{aligned}$$

and you can look up the formulas for these conditional densities in some book.