STK 4021 - Applied Bayesian Analysis and Numerical Methods
Thordis L. Thorarinsdottir
Fall 2014

Problem Set 4

Problem 1 (reflecting random walks). It is often useful in MCMC to have a proposal distribution which is both symmetric and has support only on a certain region. For example, if we know $\theta > 0$, we would like our proposal distribution $J(\theta^*|\theta_k)$ to have support on positive θ values. Consider the following proposal algorithm:

- sample $\tilde{\theta} \sim \mathcal{U}([\theta_k - \delta, \theta_k + \delta])$;
- if $\tilde{\theta} < 0$, set $\theta^* = -\tilde{\theta}$;
- if $\tilde{\theta} \geq 0$, set $\theta^* = \tilde{\theta}$.

In other words, $\theta^* = |\tilde{\theta}|$. Show that the above algorithm draws samples from a symmetric proposal distribution which has support on positive values of θ.

Hint: it may be helpful to write out the associated proposal density $J(\theta^*|\theta_k)$ under the two conditions $\theta_k \leq \delta$ and $\theta_k > \delta$ separately.

Problem 2 (nesting success). Younger male sparrows may or may not nest during a mating season, perhaps depending on their physical characteristics. Researchers have recorded the nesting success of 43 young male sparrows of the same age, as well as their wingspan, and the data appear in the file `msparrownest.dat`. Let Y_i be the binary indicator that sparrow i successfully nests, and let x_i denote their wingspan. Our model for Y_i is

$$\text{logit}\ \mathbb{P}(Y_i = 1|\alpha, \beta, x_i) = \alpha + \beta x_i,$$

where the logit function is given by logit $\theta = \log[\theta/(1-\theta)]$.

(a) Write out the joint sampling distribution $\prod_{i=1}^n p(y_i|\alpha, \beta, x_i)$ and simplify as much as possible.

(b) Formulate a prior probability distribution over α and β by considering the range of $\mathbb{P}(Y = 1|\alpha, \beta, x)$ as x ranges over 10 to 15, the approximate range of the observed wingspans.

(c) Implement a Metropolis algorithm that approximates $p(\alpha, \beta|\boldsymbol{x}, \boldsymbol{y})$. Adjust the proposal distribution to achieve a reasonable acceptance rate, and run the algorithm for a while.

(d) Compare the posterior densities of α and β to their prior densities.

(e) Using output from the Metropolis algorithm, come up with a way to make a confidence band for the following function $f_{\alpha\beta}(x)$ of wingspan:

$$f_{\alpha\beta}(x) = \frac{e^{\alpha+\beta x}}{1 + e^{\alpha+\beta x}},$$

where α and β are the parameters in your sampling model. Make a plot of such a band.

Solutions will be discussed in class on September 26.