Statistical Computing: Bayesian Inference

June 12

Markov chain Monte Carlo

We will simulate a sample $\mu^{[1]}, \mu^{[2]}\mu^{[3]}, \dots, \mu^{[M]} \sim p(\mu|\boldsymbol{y})$. The kernel of the target density is

$$p(\mu|\boldsymbol{y}) \propto \left[\sum_{i=1}^{n}(y_i - \mu)^2\right]^{-n/2} exp\left[-\frac{(\mu-\eta)^2}{2\tau^2}\right] \tag{1}$$

where

- $M = 25000$
- $n = 10$
- $\sum y_i = 17.34$
- $\sum y_i^2 = 32$
- $\eta = 1.8$
- $\tau^2 = 0.25$

Use the Metropolis algorithm to simulate from this. The steps are

1. Initialise $\mu^{[0]} = 1.8$
2. Initialise an $M \times 1$ vector for storing values of $\mu^{[j]}$
3. Find $\mu^{[1]}$ using the following steps
   - Set μ^{old}=$\mu^{[0]}$
   - Generate $\mu^{new} \sim N(\mu^{old}, \nu^2)$ where ν^2 is the variance of the proposal. For now, set $\nu^2 = 0.02$.
   - Compute the acceptance probability

     $$\alpha = min\left(1, \frac{p(\mu^{new}|\boldsymbol{y})}{p(\mu^{old}|\boldsymbol{y})}\right) \tag{2}$$

   - Generate an 0/1 indicator called *accept* where $\Pr(accept = 1) = \alpha$
   - If $accept = 1$ then set $\mu^{[1]} = \mu^{[new]}$, otherwise $\mu^{[1]} = \mu^{[old]}$
4. Use a loop for **j** *in* 2:M. Inside the loop:
   - Do the same as in step 3, but instead of updating $\mu^{[0]} \to \mu^{[1]}$, update $\mu^{[j-1]} \to \mu^{[j]}$

Bayes Inference

Using your sample from the first part, find an approximation of:

1. The Posterior Mean $E[\mu|\boldsymbol{y}]$
2. The Posterior Median of $\mu|\boldsymbol{y}$
3. The Posterior Mode of $\mu|\boldsymbol{y}$ (*difficult*)
4. A 95% credible interval for $\mu|\boldsymbol{y}$
5. The posterior probability $\Pr(\mu > 1.7|\boldsymbol{y})$

Don't forget to remove the first few iterates in the sample (burn-in).

Better Code

Make the following changes to your code

- Write code to give a trace plot of your iterates.
- Write code that allows you to compute the acceptance rate (i.e. the percentage of times a proposed value of μ^{new} is accepted).

Assess scheme

Instead of $\nu^2 = 0.02$, try half ($\nu^2 = 0.01$) this value.

1. Do you expect the acceptance rate to change? Does it? Why?
2. Do you expect the posterior mean to change? Does it? Why?

Do the same for ($\nu^2 = 0.04$).

Preparation for Tuesday's Lecture

Download the package CODA. Try to use the functions *geweke.diag* and *EffectiveSize* on your Monte Carlo sample.