Statistics 651: Introduction to Applied Bayesian Methods
Mini-Project # 3
Due: November 6, 2009, 4pm

1. Use the data from the website `http://madison.byu.edu/bayes/faculty.dat` as in HW # 2 to apply the bootstrap importance sampling algorithm. Generate 5000 observations from the posterior and do the following:

   (a) Choose an appropriate importance function and tell me what it is and why you think it is an appropriate importance function.

   (b) Calculate the posterior distribution (joint, if applicable) of the parameters of your model. Present the joint distribution in its **bivariate form**.

   (c) Calculate $E[\Theta|\mathbf{Y}]$, the posterior mean.

   (d) Calculate $V[\Theta|\mathbf{Y}]$, the posterior variance.

   (e) Calculate $\sqrt{V[\Theta|\mathbf{Y}]}$, the posterior standard deviation.

   (f) Calculate the predictive distribution of the "next" average faculty evaluation.

   (g) In general, it is considered "good" to score a 5 or better. Calculate the probability that a randomly selected professor gets a 5 or better in her next evaluation.

2. Demonstrate that the Metropolis-Hastings (M-H) algorithm works. Generate 10000 observations from the distribution of X, where

$$f^*(x) \propto (1 + (x - 10)^2/3)^{-2}, \qquad -\infty < x < \infty.$$

   Demonstrate that the M-H algorithm works by plotting a density plot of your draws compared with the actual distribution. Plot them both on the same axes with a different line type to distinguish the two.

3. A famous reliability experiment was performed and $n = 23$ ball bearings were placed on test and the failure times were recorded. They can be found at `http://madison.byu.edu/bayes/ballbearing2.dat` and the observations are the number of revolutions (recorded in 10^6 units). It is well known that these data have a Weibull failure time. Using Gibbs sampling as your computational tool, fit the ball bearing data. Be very careful in your choice of likelihood (there are at least 3 different ways to pose a Weibull model). Experts know that ball bearings should average about 50-70 (10^6 revolutions), and that there should absolutely no bearing under 10. Choose priors that make sense to you. It is a tough prior specification! Make sure that your priors agree with the parameter space you determine. After you have your prior specification, fit the model. Report the following:

   (a) Posterior means and variances for your chosen parameterization.

   (b) Posterior distributions (joint, *and* marginal), with your priors included on the same plot as your marginals.

   (c) Posterior predictive distribution for the next ball bearing failure time.

   (d) 97% HPD for each of your parameters.

   (e) Superior ball bearings will last 120 (10^6 revolutions). Find the posterior probability that our next ball bearing will be a superior ball bearing. That is, find the $Pr(Y_{n+1} > 100 | y_1, ..., y_n)$.

   (f) Report the final settings for your MCMC, (BURN, LENGTH, σ_{cand}, candidate density, etc.).

   (g) Include a copy of your code in an appendix to your homework assignment.