Stat 8931, Fall 2005
Homework 2
Due Oct 5, 2005

Q1 Do a Gibbs sampler for the problem described in Section 10.6 (or whatever the "Examples" subsection of "The Gibbs Update" section has turned into) of the lecture notes.

The likelihood is a function of the sufficient statistics

$$\hat{\mu}_n = \frac{1}{n} \sum_{i=1}^{n} x_i$$
$$\hat{\sigma}_n^2 = \frac{1}{n} \sum_{i=1}^{n} (x_i - \hat{\mu}_n)^2$$

Suppose we have observed

$$n = 20$$
$$\hat{\mu}_n = 14.731$$
$$\hat{\sigma}_n^2 = 4.814$$

Suppose that we have independent priors for μ and $\lambda = 1/\sigma^2$ (the "precision") which are

$$\mu \sim \text{Normal}(10, 2^2)$$
$$\lambda \sim \text{Gamma}(3, 1)$$

Although this is the kind of problem that WinBUGS is designed for, write your own Gibbs sampler in R.

Run the Gibbs sampler and get posterior expectation and MCSE for the functions of the state

- μ (the first component)
- $\sigma = \lambda^{-1/2}$ where λ is the second component.

As in the first assignment, run the Markov chain long enough to obtain MCSE smaller than 0.001 for each.

Also as in the first assignment, turn in either an R CMD BATCH file or an Sweave file that contains your complete analysis.