Statistics 5314: Homework 4

For each homework assignment, turn in at the beginning of class on the indicated due date. Late assignments will only be accepted with special permission. Write each problem up *very* neatly (LaTeX is preferred). Show all of your work.

Problem 1

For this problem we will implement the Metropolis algorithm. Consider the model $x_i \sim N(\mu, v = \sigma^2)$. Denote $\boldsymbol{X} = \{x_1, \dots, x_N\}$.

Part a

Under $\mu = 8$ and $\sigma = 2$, simulate $N = 1,000$ samples and plot them.

Part b

Recall that the joint posterior distribution for μ and v follows as

$$p(\mu, \sigma | \boldsymbol{X}) \propto \prod_{i}^{N} p(x_i | \mu, v) p(\mu, v),$$

where $p(\mu, v)$ is some prior distribution on μ and v. For this problem, let the prior be the standard reference choice: $p(\mu, v) \propto \frac{1}{v}$.
Use the Metropolis algorithm, starting from $(\mu^{(0)}, v^{(0)}) = (-10, 20)$ to sample from the joint posterior distribution. Show a plot (2-dimensional) of the iterations for the joint samples. Also, show trace plots for the margins (μ and v). Discuss briefly your proposal distribution and the associated tuning parameters.

Note: multiplying 1,000 numbers together can often times have poor numerical stability (especially if the numbers are extreme). It is often beneficial to operate on a log scale. Of course, for this problem, the likelihood can easily be written down in terms of its sufficient statistics $\bar{x} = \sum x_i / N$ and $s^2 = \sum (x_i - \bar{x})^2 / (N - 1)$, which will also alleviate the problem of computing the product at every iteration.

Problem 2

Consider the AR(1) time series model:

$$x^{(t)} = \alpha x^{(t-1)} + \epsilon^{(t)}.$$

Part a

For $\epsilon^{(t)} \sim N(0, v = \sigma^2)$ answer the following questions:

- For which values of α is the process stationary? What about weakly stationary? Include a explanation of your answer.
- For $\alpha = 0.9$ and $\sigma = 1$, simulate the process. Show a plot of a 1,000 point time series under this model (prior to running the model 10,000 iterations so that it is not sensitive to the starting value (you choose)). Is the process time reversible? Why or why not?

Part b

For $\epsilon^{(t)} \sim \text{Cauchy}(0, 1)$ answer the following questions:

- For which values of α is the process stationary? What about weakly stationary? Include an explanation of your answer.
- For $\alpha = 0.9$, simulate the process. Show a plot of a 1,000 point time series under this model (prior to running the model 10,000 iterations so that it is not sensitive to the starting value (you choose)). Is the process time reversible? Why or why not?

Problem 3

For this problem you will use MCMC to perform inference on the parameters in the models in Problem 2.

Part a

Use the 1,000 samples from part a of problem 2 (or a similar set of samples if you've discarded them already) to construct the joint posterior distribution for (α, v), under the prior distribution $p(\alpha, v) \propto 1/v$.

Part b

Use the 1,000 samples from part b of problem 2 (or a similar set of samples if you've discarded them already) to construct the joint posterior distribution for α, under the prior distribution $p(\alpha) \propto 1$.