Stat 532 **Name:** ____________________________
HW 6
Due 10/9/17

Please prepare your solutions using LaTeX or another word processing software. Note, both of these problems will be presented in class and there may be time to work on them in a lab setting.

1. (60 points) Consider the mixture distribution described on p. 99 (Hoff). This distribution is a joint probability distribution of a discrete variable $\delta = \{1, 2, 3\}$, denoting which mixture component the mass comes from and a continuous variable θ. The target density is $\{Pr(\delta = 1), Pr(\delta = 2), Pr(\delta = 3)\} = (.45, .10, .45)$ and $p(\theta|\delta = i) \sim N(\theta; \mu_i, \sigma_i^2)$ where $\{\mu_1, \mu_2, \mu_3\} = (-3, 0, 3)$ and $\sigma_i^2 = 1/3$ for $i \in \{1, 2, 3\}$.

   1. Generate 1000 samples of θ from this distribution using a Monte Carlo procedure. Hint: first generate $\delta^{(i)}$ from the marginal distribution $p(\delta)$ and then generate $\theta^{(i)}$ from $p(\theta|\delta)$ Plot your samples in a histogram form and superimpose a curve of the density function. Comment on your samples, do they closely match the true distribution?
   2. Next, generate samples from a Gibbs sampler using the full conditional distributions of θ and δ. You already know the form of the full conditional for θ from above. The full conditional distribution for δ is given below:

   $$Pr(\delta = d|\theta) = \frac{Pr(\delta = d) \times p(\theta|\delta = d)}{\sum_{d=1}^{3} Pr(\delta = d) \times p(\theta|\delta = d)}$$

      Hint: for $p(\theta|\delta = d)$ evaluate θ from a normal distribution with parameters $\{\mu_d, \sigma_d^2\}$. Intialize θ at 0.

      (a) Generate 100 samples using this procedure. Plot your samples as a histogram with the true density superimposed on the plot. Also include a plot of your θ value on the y-axis and the iteration number on the x-axis. This is called a trace plot, and allows your to visualize the movement of your MCMC *particle*. Comment on how close your samples match the true density. What does the trace plot reveal about the position of θ over time (the iterations)? Does the proportion of the time the sample spends in each state (δ) match the true probabilities?

      (b) Repeat for 1000 samples.

      (c) Repeat for 10000 samples.
   3. Now repeat part 2, but instead initialize θ at 100. How does this change the results from part 2? Looking at trace plots, do the chains from the three plots seem to be similar?

2. (40 points) Consider a similar scenario to the code for the first Gibbs sampler. Again 100 data points have been generated.

   1. A histogram of the data is shown later as figure (a). What are your thoughts about this data?

2. Now assume you used your MCMC code and came up with figures (b) - (e). Comment on the convergence and the marginal posterior distributions.

3. As a final check you decide to use your MCMC samples to compute the posterior predictive distribution, $p(y^*|y_1, \ldots, y_n)$. Computationally this can be acheived by using each pair $\{\theta^{(i)}, \sigma^{2(i)}\}$ and then simulating from $N(y^*; \theta^{(i)}, \sigma^{2(i)})$. In R this can be done with one line of code:

   ```
   post.pred <- rnorm(num.sims,mean=Phi[,1],sd = sqrt(1/Phi[,2]))
   ```

   Now compare your posterior predictive distribution with the observed data. Are you satisfied that the posterior predictive distribution represents the actual data? Why or why not?

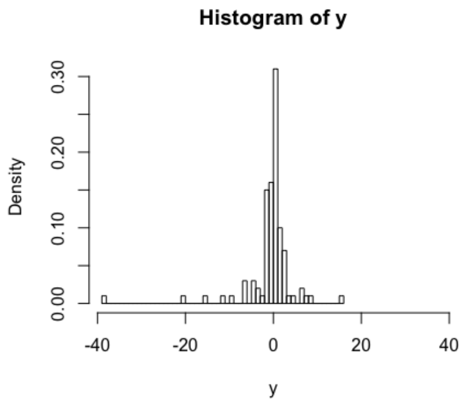

(a) Histogram of the Data

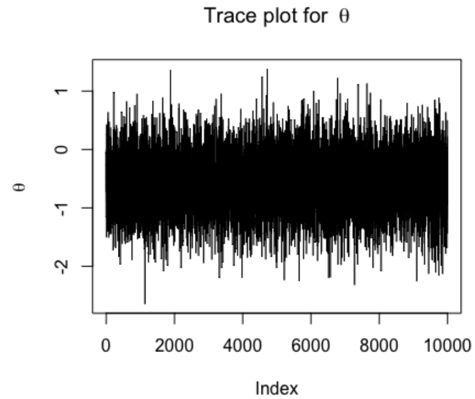

(b) Trace plot for θ

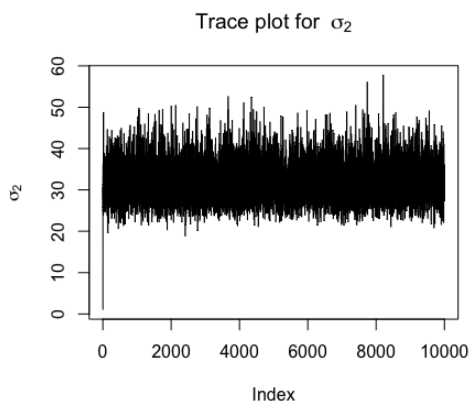

(c) Trace plot for σ^2

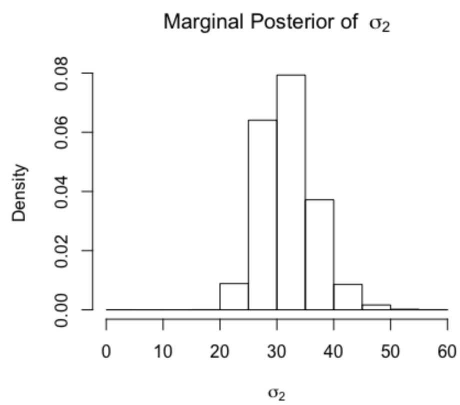

(d) Histogram of marginal posterior for σ^2

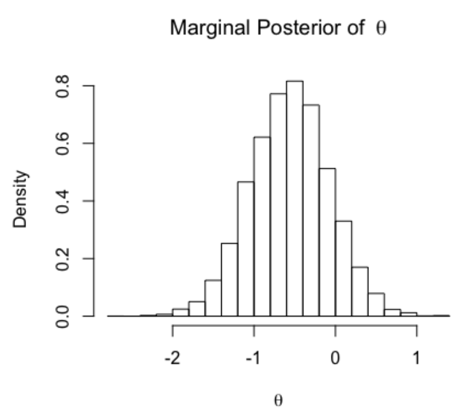

(e) Histogram of marginal posterior for θ