Chapter 9

Gibbs Sampling for the Normal Model. Lab to turn in at the end of session. Must print before leaving. Do problems 8.2.1 and 8.2.2 and 8.2.3

9.1 Introduction

We introduced in the last notes the Gibbs sampler for a simple case, a bivariate posterior distribution for (θ_1, θ_2). In this lesson, we will try to reproduce the posterior distributions presented in Hoff's book, page 94. The program here is slightly different, and we also do the plots differently. The results, however, are the same.

Remember, Gibbs sampling approximates the posterior distribution $P(\theta_1, \theta_2)$ by drawing random numbers iteratively from the conditionals $p(\theta_1 \mid \theta_2)$ and $p(\theta_2 \mid \theta_1)$. The distributions of the θ_1 and the θ_2 generated this way are the posterior marginal distributions of θ_1 and of θ_2, respectively.

9.2 Obtaining the posterior distribution of $(\theta, \sigma^2 \mid y_1, ..., y_n)$ in the normal model with unknown mean and variance

The joint posterior distribution of $(\theta, \sigma^2) \mid y_1, ..., y_n$ can be obtained by drawing from these two full conditional distributions:

The conditional distribution of θ given the σ^2 and the data (see chapter 5 of Hoff's and your class notes for this one):

$$\theta \mid \sigma^2, y_1, ..., y_n) \sim Normal(\mu_n, \tau_n^2)$$

with

$$\mu_n = \frac{\mu_0/\tau_0^2 + n\bar{y}/\sigma^2}{1/\tau_0^2 + n/sigma^2}$$

and

$$\tau_n^2 = \left(1/\tau_0^2 + n/\sigma^2\right)^{-1}$$

The conditional distribution of σ^2 given θ and the data (see chapter 6 of Hoff's):

$$\sigma^2 \mid \theta, y_1, ..., y_n \sim Igamma(\nu_n, \sigma_n^2)$$

where

$$\nu_n = \nu_0 + n$$

and

$$\sigma_n^2 = \frac{1}{\nu_n}[\nu_0\sigma_0^2 + n \quad s_n^2]$$

Notice that $ns_n^2 = (n-1)s^2 + n(\bar{y} - \theta)^2$. We use this in the code below.

Notice also that we will be drawing $1/\sigma^2$ which is gamma distributed instead of σ^2 directly. This is no problem. Once we simulate the precision $1/\sigma^2$, we obtain σ^2 by inverting the precision.

We will do the simulation for the midget data given in Chapter 5 of Hoff's. Thus, we use the values given for everything in that chapter. Read the chapter for their rationale.

9.2.1 EXERCISE 1 TO TURN IN

Open a file and insert the results you get running this program, with some explanations of your findings. We are trying to estimate the marginal posterior mean and standard deviation of the Midge wing length seen in Chapter 5.

```
###############################################
# Midge wing length. Simplifying the program on page 94 of Hoff's:
# Simulating the posterior distribution of mu and sigma with Gibbs sampling
# Notice that to do the Gibbs for the posterior standard deviation we
# make use of the precision. Also, plots are not as in Hoff, because we
# will be plotting the posterior standard deviation instead of the precision
# or the variance. But if you do the calculation, precision=1/variance
# and sd = sqrt(1/precision)
###############################################

#########  Data and data summary ##########

mu0=1.9   # mean of prior distribution of theta  (mu_0)
t20 = 0.95^2     # variance of prior distribution of theta  (tau_0^2)
s20=0.01   #  variance of prior distribution of sigma squared (sigma_0^2)
nu0=1  # number of prior observations   (nu_0)

y=c(1.64,1.7,1.72,1.74,1.82,1.82,1.9,2.08)  # midge wing length data

mean.y = mean(y)    # mean of the data (ybar)
var.y= var(y)  # variance of the data (s^2)  - s2 in program
  n= length(y)  # sample size of observed data

  #initial values ybar and s2

########## Initial values for the chains #############

S=1000
mun=rep(0,S)  # open space for the mun
t2n = rep(0,S)  # open space for the tau squaren
s2n=rep(0,S)
mun[1] = mean.y  ;   t2n[1] = 1/var.y     #initial values

post.mean=rep(0,S)  # room for draws of posterior mean
post.mean[1] = mean.y
```

```
post.tau2 =rep(0,S)  # room for draws of the posterior precision
post.tau2[1]= t2n[1]

########### Gibbs sampling ###################

for(i in 2:S) {

# Generate a new value for the posterior mean from its full conditional

mun[i] = (mu0/t20 + n*mean.y*post.tau2[i-1])/(1/t20 + n*post.tau2[i-1])
t2n[i]  = 1/(1/t20 + n*post.tau2[i-1] )
post.mean[i] = rnorm(1, mun[i], sqrt(t2n[i]))

# Generate a new 1/sigma^2 value from its full conditional
nun=nu0+n
s2n[i]=(nu0*s20 + (n-1)*var.y + n*(mean.y -post.mean[i])^2 )/ nun
post.tau2[i] =  rgamma(1, nun/2, nun*s2n[i] /2 )

}

# plot traces of your gibbs sampling of mean and standard deviation
# save to your files

plot(post.mean[500:S], type="l", lty=1)  # plot markov chain for post mean
plot(sqrt(1/post.tau2[500:S]),type="l",lty=2) # plot markov chain for post sd

# Plot joint traces   and save to your file
plot(post.mean[500:S],sqrt(1/post.tau2[500:S]),type="l")    # line plot to see the traces

# Plot scatter plot of draws from joint posterior distribution
plot(post.mean[500:S],sqrt(1/post.tau2[500:S]), type="pt")  # point plot

# Plot marginal distributions of the posterior mean and posterior standard deviation

hist(post.mean, main="Marginal distribution of posterior mean")
hist(sqrt(1/post.tau2), main="Marginal distribution of posterior standard deviation")


######## Add the following on your own:
###### Mean, median, standard deviation of the posterior mean
###### Mean, median, standard deviation of the posterior standard deviation
####### 95\% posterior intervals for the posterior mean and posterior standard deviation
### Interpret
```

9.2.2 EXERCISE 2 TO TURN IN

Repeat the work done above, but instead of starting with the sample mean and sample precision as the initial values for the Gibbs chains, start with initial values 0 for the posterior mean and 1 for the posterior precision. Copy paste the commands that you will need to change. How different are your conclusions? Support your answer with the plots and summary statistics of the posterior marginal distributions.

9.2.3 EXERCISE 3 TO TURN IN

Repeat exercise 2, but start with initial value for the posterior mean of 10 and initial value for the posterior precision of 3.