5.4 EXERCISES

BMED6420
Brani Vidakovic; Fall 2018
Consult the class slides, hints, and cited literature for the solution of exercise problems.

1. Laplace's Method for Binomial-Beta. Approximate the posterior by a normal model using Laplece's method. The likelihood is binomial $X|p \sim \mathcal{B}in(n, p)$ with and the prior is $p \sim \mathcal{B}e(\alpha, \beta)$ with parameters:

(a) $\alpha = \beta = 1$ (flat prior)
(b) $\alpha = \beta = 1/2$ (Jeffreys' prior)
(c) $\alpha = \beta = 2$.

How well the Laplace's method approximates 95% CS for p? Compare exact equi-tailed CS's with the approximations. Use $n = 20$ and $X = 8$.

2. Coin and Probability of Heads. A coin is flipped unknown number of times, say n, and X heads are observed. Suppose we are interested in probability of heads p and assume model

$$P(X = k|n, p) \sim \mathcal{B}in(n, p).$$

The model is completed with priors on n and p. We assume that n is Poisson with mean λ and p is beta $\mathcal{B}e(\alpha, \beta)$. The priors on n and p are independent.

Assume $X = 6$ is observed. Sample form the posterior using
(a) Metropolis algorithm
(b) Gibbs sampler

Assume that $\alpha = \beta = 20$ and $\lambda = 18$.

3. Amanita muscaria. With its bright red, sometimes dinner-plate-sized caps, the fly agaric (*Amanita muscaria*) is one of the most striking of all mushrooms (Fig. 1a). The white warts that adorn the cap, the white gills, a well-developed ring, and the distinctive volva of concentric rings distinguish the fly agaric from all other red mushrooms. The spores of the mushroom print white, are elliptical, and have a larger axis in the range of 7 to 13 μm (Fig. 1b).

Measurements of the diameter X of spores for $n = 51$ mushrooms are given in the following table:

Figure 1: *Amanita muscaria* and its spores. (a) Fly agaric or *Amanita muscaria*. (b) Spores of *Amanita muscaria*.

10	11	12	9	10	11	13	12	10	11
11	13	9	10	9	10	8	12	10	11
9	10	7	11	8	9	11	11	10	12
10	8	7	11	12	10	9	10	11	10
8	10	10	8	9	10	13	9	12	9
9									

Assume that the measurements are normally distributed with mean μ and variance σ^2, but both parameters are unknown and of interest.

Suppose that the prior on μ is normal $\mathcal{N}(12, 2^2)$ and the prior on $\tau = 1/\sigma^2$ is gamma $\mathcal{G}a(2, 4)$.

(a) Develop Gibbs sampling algorithm and find Bayes estimators of μ and τ.

(b) If an inverse gamma $\mathcal{IG}a(4, 2)$ is placed on σ^2, would the solution be different? Develop a Gibbs sampling for this case keeping the prior on μ as in (a).

(c) Develop a Metropolis algorithm for the priors as in (a). Choice of proposal distributions is up to you.

4. Jeremy via Metropolis. The Jeremy example was solved exactly (as a conjugate problem) and by using WinBUGS. Recall that Jeremy's IQ test score X was modeled as normal $\mathcal{N}(\theta, 80)$, while the location parameter θ had prior $\mathcal{N}(110, 120)$. We could think of θ as Jeremy's intrinsic IQ level. On his IQ test Jeremy scored $X = 98$.

Develop Metropolis MCMC scheme for sampling from the posterior.

(i) Sample the proposal θ_p from Cauchy distribution centered at the current state of the chain θ.

```
thetaprop = theta  + 10*tan(pi*rand(1,1)-pi/2);
```

Since the proposal distribution depends on $(\theta_p - \theta)^2$ it does not factor into the acceptance ratio ρ.

Your target is the product of likelihood and prior and does not need to be normalized to be a distribution. Start with $\theta = 100$ and burn in 500 simulations. Save 50,000 simulations (not including the burn in).

(a) Estimate posterior mean, posterior standard deviation, fine the median and 95% equaly-tailed credible set.

(b) Plot the normalized histogram and superimpose the theoretical posterior density.

5. Jeremy via Gibbs Sampling.

Translate the following WinBUGS code into custom made Gibbs Sampling program in MATLAB.

```
 model{
 x ~ dnorm(mu, tau)
pprec <- 1/120
mu ~ dnorm(110, pprec)
tau ~ dgamma(0.01, 1)
sig2  <- 1/tau
}

DATA
list(  x = 98  )

INITS
list(mu=100, tau=0.01)
```

Find full conditionals for μ and τ.

Compare Bayes summaries (posterior mean, median, 2.5 and 97.5 percentiles, histograms) between WinBUGS and Gibbs outputs.

Gibbs with Metropolis Step. When in Gibbs sampling scheme, one or more full conditionals are not tractable, their function can be replaced by a so-called Metropolis step. An example with Metropolis step is given below.

Georgia Deaths from Kidney Cancer 1985-1989 by Counties. Data set contains death counts from Kidney Cancer for 159 Georgia counties as well as the county population size. The data are from years 1985-1989.

It is of interest to estimate the death rate (per 100,000) for each county, as well as the all-Georgia death rate.

The model is Poisson $\mathcal{P}oi(\lambda_i n_i), i = 1, \ldots, k$ where n_i is the population size divided by 100,000 for county i; here $k = 169$. Also, $n = n_1 + \ldots + n_k$ is population of Georgia in units of 100,000.

The model is as follows:

$$y_i \sim \mathcal{P}oi(\lambda_i n_i), i = 1, \ldots, k$$

$$
\begin{aligned}
&\lambda_i \sim \mathcal{G}a(\alpha, \beta), i = 1, \ldots, k \\
&\alpha \sim \mathcal{U}(0, A) \\
&\beta \sim \mathcal{U}(0, B)
\end{aligned}
$$

for specified constants A and B.

The joint distribution is:

$$
f(\boldsymbol{y}, \boldsymbol{\lambda}, \alpha, \beta) = \prod_{i=1}^{k} \left[\frac{\lambda_i^{y_i}}{y_i!} e^{-\lambda_i n_i} \times \frac{\beta^{\alpha} {\lambda_i}^{\alpha-1}}{\Gamma(\alpha)} e^{-\beta\lambda_i} \right] \times \frac{1}{A} \mathbf{1}(0 \le \alpha \le A) \times \frac{1}{B} \mathbf{1}(0 \le \alpha \le B).
$$

$$
\begin{aligned}
\pi(\lambda_i | \boldsymbol{y}, \boldsymbol{\lambda}_{\neq i}, \alpha, \beta) &\propto f(\boldsymbol{y}, \boldsymbol{\lambda}, \alpha, \beta) \\
&\propto \lambda_i^{y_i} e^{-\lambda_i n_i} \lambda_i^{\alpha-1} e^{-\beta\lambda_i}, i = 1, \ldots, k.
\end{aligned}
$$

Therefore,

$$
[\lambda_i | \boldsymbol{y}, \boldsymbol{\lambda}_{\neq i}, \alpha, \beta] \sim \mathcal{G}a(y_i + \alpha, n_i + \beta), \; i = 1, \ldots, k.
$$

$$
\begin{aligned}
\pi(\alpha | \boldsymbol{y}, \boldsymbol{\lambda}, \beta) &\propto f(\boldsymbol{y}, \boldsymbol{\lambda}, \alpha, \beta) \\
&\propto \frac{\beta^{k\alpha}}{(\Gamma(\alpha))^k} \left(\prod_{i=1}^{k} \lambda_i \right)^{\alpha-1} \frac{1}{A} \mathbf{1}(0 \le \alpha \le A) \\
&\propto \frac{\left(\beta^k \prod_{i=1}^{k} \lambda_i \right)^{\alpha}}{(\Gamma(\alpha))^k} \mathbf{1}(0 \le \alpha \le A).
\end{aligned}
$$

$$
\begin{aligned}
\pi(\beta | \boldsymbol{y}, \boldsymbol{\lambda}, \alpha) &\propto f(\boldsymbol{y}, \boldsymbol{\lambda}, \alpha, \beta) \\
&\propto \beta^{k\alpha} e^{-\beta \sum_{i=1}^{k} y_i} \times \frac{1}{B} \mathbf{1}(0 \le \beta \le B).
\end{aligned}
$$

Therefore,

$$
[\beta | \boldsymbol{y}, \boldsymbol{\lambda}, \alpha] \sim \mathcal{G}a \left(k\alpha + 1, \sum_{i=1}^{k} y_i \right) \mathbf{1}(0 \le \beta \le B).
$$

The Metropolis Step for generating α:

Proposal $\alpha' \sim \mathcal{N}(\alpha, \sigma^2)$. The Metropolis ratio depends on target only,

$$
r = \left(\frac{\Gamma(\alpha)}{\Gamma(\alpha')} \right)^k \times C^{\alpha' - \alpha} \times \mathbf{1}(0 \le \alpha' \le A),
$$

for $C = \beta^k \prod_{i=1}^{k} \lambda_i$.

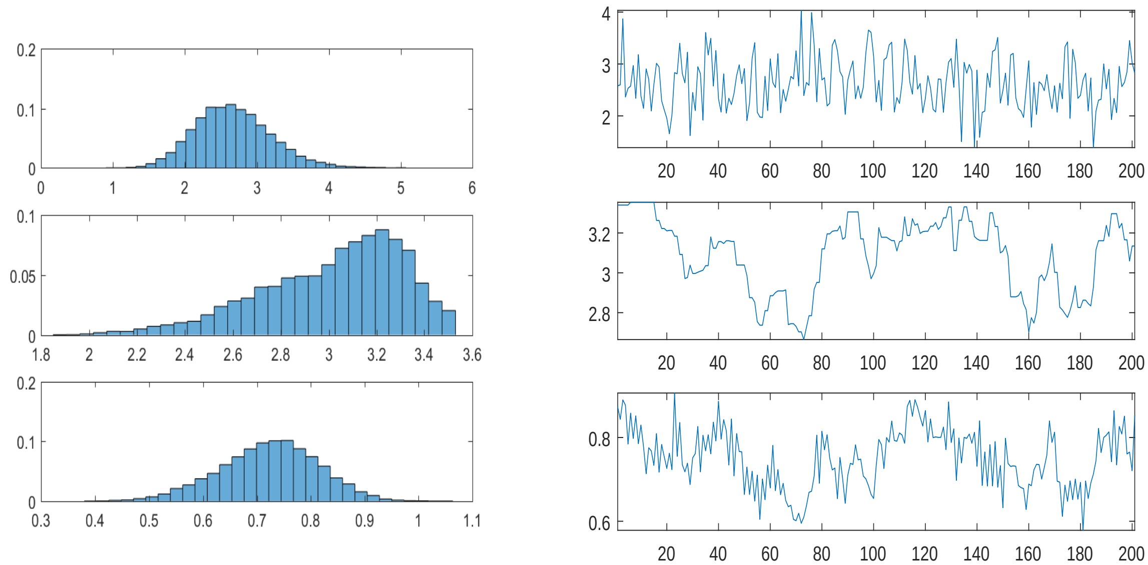

From `gibbsmetstep.m`:

Exponential Survival Times. Survival times of n post-surgery patients assigned to a treatment group are given as $T_1, \dots, T_n$. In a control group of m post-surgery patients the survival times are $C_1, \dots, C_m$.

The proposed model is

$$\begin{aligned} T_i &\sim \mathcal{E}(\lambda\theta),\ i = 1, \dots, n \\ C_i &\sim \mathcal{E}(\lambda),\ i = 1, \dots, m \\ &\lambda > 0, \theta > 0. \end{aligned}$$

where T_i's and C_i's are assumed independent.

To complete the model assume the following priors on λ and θ,

$$\begin{aligned} \pi(\lambda) &= \frac{1}{\lambda}, \quad \text{[Jeffreys' choice]} \\ \theta &\sim \mathcal{G}a(1/2, 1/2), \end{aligned}$$

where gamma is parameterized via the rate parameter.

(a) Show that the likelihood and posterior depend on $T = \sum_{i=1}^n T_i$ and $C = \sum_{i=1}^m C_i$.

(b) Identify full conditionals for $[\lambda|\theta, T, C]$ and $[\theta|\lambda, T, C]$.

(c) For $n = 10, m = 12, T = 18.26$, and $C = 26.78$ set Gibbs' sampling scheme.

(d) From a Gibbs sample of size 10000 estimate θ and find 95% equitailed credible set. Is the credible set containing 1? Discuss.

(e) How would you deal with the analysis if some times are censored? Repeat the analysis if $T_{10} = 2.12, C_{11} = 1.54$ and $C_{12} = 1.98$ are in fact censored times.

Testing the Effectiveness of a Seasonal Flu Shot. Assume 30 individuals are given a flu shot at the start of winter. At the end of winter, follow up to see whether they contracted flu. Let

$$x_i = \begin{cases} 1, & \text{no flu (shot effective)} \\ 0, & \text{flu (shot not effective)} \end{cases}$$

Suppose the 30th individual was unavailable for followup. Define $y = \sum_{i=1}^{29} x_i$. If p is the probability the shot is effective, then

$$f(y|p) = \binom{29}{y} p^y (1-p)^{29-y}, \; y = 0, 1, \ldots, 29.$$

With the complete data (y plus x_{30}),

$$f(y, x_{30}|p) = \binom{30}{y + x_{30}} p^{y+x_{30}} (1-p)^{30-y-x_{30}}, \; y + x_{30} = 0, 1, \ldots, 30.$$

With uniform prior on p and $[x_{30}|p] \sim \mathcal{B}er(p)$ the joint distribution of $[p, x_{30}|y]$ is proportional to $f(y, x_{30}|p)$. Thus, the full conditionals are

$$[p|x_{30}, y] \sim \mathcal{B}e(y + x_{30} + 1, 30 - y - x_{30} + 1),$$

and $[x_{30}|y, p] \sim \mathcal{B}er(p)$.

If $y = 21$, estimate p using Gibbs sampler that uses these two full conditionals.

Hints/Results/Solutions to Some of the Exercises

1. Laplace's Method for Binomial-Beta.

2. Coin and Probability of Heads.

Hint for (a). Show first that $[X|p] \sim \mathcal{P}oi(\lambda p)$. To do this you will need to integrate n from the $X \sim \mathcal{B}in(n,p)$.

$$P(X = k|p) = \sum_{n=k}^{\infty} P(X = k|n,p) \frac{\lambda^n}{n!} e^{-\lambda}.$$

Since $P(X = k|n,p) = \binom{n}{k} p^k (1-p)^{n-k}$ after change of variables $u = n-k$ and rearrangement of the sum you will arrive to Poisson $\mathcal{P}oi(p\lambda)$ distribution. Here k is constant, and you will need to use $\sum_{u=0}^{\infty} [(1-p)\lambda]^u/u! = \exp\{(1-p)\lambda\}$.

Now, the posterior is proportional to the product of beta prior and marginal likelihood for $[X|p]$ which is the Metropolis target function.

Take as proposal beta $\mathcal{B}e(X + \alpha, \beta)$ distribution. Note that this is an independence proposal which conveniently simplifies acceptance probability in Metropolis algorithm.

Pick any p between 0 and 1 and start the simulation.

Hint for (b). For Gibbs, you will need full conditionals. Argue that:

$$[n|X,p] \sim X + \mathcal{P}oi((1-p)\lambda).$$

and

$$[p|n,X] \sim \mathcal{B}e(\alpha + X, \beta + n - X).$$

The first one is simply number of heads (observed) plus number of tails (unobserved) and by argument similar to one in (a) distributed as $\mathcal{P}oi((1-p)\lambda)$. The second follows from a conjugate Binomial-Beta setup.

3. Amanita muscaria.

Hint: Bayes' inference depends on the sufficient statistics $\overline{X} \sim \mathcal{N}(\mu, \frac{\sigma^2}{n})$.

4. Jeremy via Metropolis.

5. Jeremy via Gibbs Sampling.

6. Gibbs with Metropolis Step.

```
%
% Kidney Cancer in Georgia 1985-1989 by county
%
close all
clear all

%rand('seed',2);
%randn('seed',2);
filename = 'C:\Brani\Courses\isyestatg\ISyE6420Spring2018\HWs\georgiakcd.xlsx';


data = xlsread(filename,  'B:D');

 y = data(:,2);
 ni = data(:,3)/100000;
 [n ~]=size(data);
 A=10;
 B=10;
 %
lambda = ones(n,1); % initial values
alpha = 1;
beta=1;
sigma=0.1;
%
lambdas =[lambda]; %save all lambdas
alphas =[alpha];
betas= [beta];
%
%
tic
for i = 1:20000  %beware slow...
   prodlambdas = prod(lambda);
   sumlambdas = sum(lambda);
   % lambda_k ~ Gamma(y_k + alpha, ni_k + beta), k=1,...,n
   lambda=zeros(n,1);
   for k = 1:n
       lambda(k) = gamrnd(y(k)+alpha, 1/(ni(k)+beta));
   end
   lambdas=[lambdas  lambda];

   % METROPOLIS STEP for alpha
   alpha_prop =  alpha + sigma*randn;
```

```
   c= prodlambdas * beta^n;
   if   (alpha_prop < A) .* (alpha_prop > 0)
          r = (gamma(alpha).^n * c^alpha_prop)/(gamma(alpha_prop).^n * c^alpha) ;
   else
          r = 0;
   end
   if rand < r
       alpha = alpha_prop;
   end
   alphas=[alphas alpha];
   % RESTRICTED GAMMA for beta

   beta_new = B+1;
   while beta_new > B
        beta_new = gamrnd(n*alpha + 1, 1/sumlambdas);
   end
   beta = beta_new;
   betas = [betas beta];
  end
toc
%Burn in 500
burn=500;
 lambdas = lambdas(:, burn+1:end);
 alphas = alphas(burn+1:  end);
  betas = betas(burn+1: end);
figure(1)
nbins=30;
subplot(3,1,1)
histogram(lambdas(33,:),nbins,'Normalization','Probability')
subplot(3,1,2)
histogram(alphas,nbins,'Normalization','Probability')
subplot(3,1,3)
histogram(betas,nbins,'Normalization','Probability')
print -depsc 'gacancer1.eps')

figure(2)
subplot(3,1,1)
plot(lambdas(33,2000:2200),'-')
axis tight
subplot(3,1,2)
plot(alphas(2000:2200),'-')
axis tight
```

```
subplot(3,1,3)
plot(betas(2000:2200),'-')
axis tight
print -depsc 'gacancer2.eps')

ma=mean(alphas)
mb=mean(betas)
mla = mean(lambdas')
ma/mb
sum(y)/sum(ni)
```

Exponential Survival Times.

Hint: The posterior is proportional to

$$\prod_{i=1}^{n}(\lambda\theta\exp\{-\lambda\theta T_i\}) \times \prod_{i=1}^{m}(\lambda\exp\{-\lambda C_i\}) \times \frac{1}{\lambda} \times \theta^{1/2-1}\exp\{-\theta/2\}$$

The above product is proportional to

$$\lambda^{m+n-1}\exp\{-\lambda(\theta T + C)\},$$

when θ is constant, and to

$$\theta^{n-1/2}\exp\{-(\lambda T + 1/2)\theta)\},$$

when λ is constant. Thus, $[\lambda|\theta, T, C] \sim \mathcal{G}a(m+n, \theta T+C)$ and $[\theta|\lambda, T, C] \sim \mathcal{G}a(n+1/2, \lambda T+1/2)$.