Bayesian Theory and Computation, Problem Set 1

March 26, 2021 Due 04/09/2021

Problem 1.

A retail company monitoring the sales of a new product records the numbers sold at each of I similar retail outlets in each of n consecutive weeks. Sales are modeled using a Poisson distribution for each outlet with a constant weekly rate:

$$x_{ij}|\theta_i \overset{iid}{\sim} \text{Poisson}(\theta_i)$$

for $1 \le i \le I,\ 1 \le j \le n$, where x_{ij} is the sales at outlet i in week j. Possible variation between outlets is modeled by assuming the means θ_i are independently drawn from the gamma prior with uncertain mean $1/\mu$ and specified shape $a > 0$, i.e.,

$$\theta_i|\mu \sim \text{Gamma}(a, a\mu).$$

The x_{ij} are assumed conditionally independent of μ given θ. For notation, set $\Theta = \theta_{1:I} = \{\theta_i\}_{i=1}^I$ and $X = x_{1:I,1:n} = \{x_{ij}, 1 \le i \le I, 1 \le j \le n\}$.

(1) Write down the complete form of $p(X|\theta, \mu)$.

(2) Show that, conditional on X and μ, the θ_i are independent and Gamma distributed, namely $\theta_i|X, \mu \sim \text{Gamma}(a + ny_i, a\mu + n)$, independently, where y_i is a function of the data X. Identify y_i.

(3) By integrating over Θ, find the marginal likelihood $p(X|\mu)$. Be sure to identify all terms involving μ, but you may ignore other "constants".

(4) If μ has the Gamma prior $\mu \sim \text{Gamma}(r, s)$, show that

$$p(\mu|X) \propto \mu^{r+aI-1}e^{-s\mu}(n + a\mu)^{-q}, \quad \mu > 0$$

where $q = I(a + n\bar{x})$ and $\bar{x}$ is the overall sample mean.

(5) Assuming the result of (4), let r and s tend to zero. Show that the resulting posterior for μ is such that $\mu = \phi/\bar{x}$, where $\phi \sim F_{k,h}$ with $k = 2aI$ and $h = 2nI\bar{x}$.

Some facts related to the F distribution. If $\phi \sim F_{k,h}$, then

$$p(\phi) \propto \phi^{k/2-1}/(h + k\phi)^{(k+h)/2}, \quad \phi^{-1} \sim F_{h,k}.$$

Problem 2.

Consider the following Markov chain taking values on the nonnegative integers. The one-step transition probabilities are given by $(i, j \geq 0)$

$$\Pr(X_{t+1} = j | X_t = i) \equiv p_{ij} = \begin{cases} \frac{1}{2j} & j = i+1, i \geq 0 \\ \frac{j}{2(j+1)} & j = i \geq 1 \\ \frac{1}{2} & j = i = 0 \\ \frac{1}{2} & j = i-1, i \geq 1 \\ 0 & |j-i| > 1 \end{cases}$$

(1) Verify that these transition probabilities are legitimate.

(2) Argue that this Markov chain is irreducible.

(3) Is this chain aperiodic? Why or why not?

(4) Verify that the Poisson(1) distribution is a stationary distribution.

(5) Starting with $X_0 = 1$, simulate this Markov chain for $10,000$ iterations. Compute the mean and variance of these $10,000$ draws and compare with the theoretical mean and variance of the stationary distribution.

Problem 3.
Consider the following model

$$\begin{aligned} \phi_1 | \mu, \sigma^2 &\sim \mathcal{N}(\mu, \sigma^2/(1-\rho^2)), \\ \phi_{j+1} | \phi_1, \ldots, \phi_j, \mu, \sigma^2 &\sim \mathcal{N}(\mu + \rho(\phi_j - \mu), \sigma^2), \quad j = 1, \ldots, J-1 \\ y_j | \phi_1, \ldots, \phi_J, \tau^2 &\sim \mathcal{N}(\phi_j, \tau^2), \quad j = 1, \ldots, J, \end{aligned}$$

where μ, σ^2, τ^2 are unknown parameters of interest, ρ is a known constant between -1 and 1, $\phi = (\phi_1, \ldots, \phi_J)$ is a sequence of latent variables, and $y = (y_1, \ldots, y_J)$ is the observed data. (This is known as a *Gaussian state-space model.*)

(1) Assume the prior $p(\mu, \sigma^2, \tau^2) \propto \tau^{-4} e^{-1/(2\tau^2)}$. That is, the prior on (μ, σ^2) is a non-informative constant prior, and independently, the prior on $1/\tau^2$ is an exponential distribution (check this). Write down the joint posterior of all parameters and latent variables given the observed data y.

(2) Derive the full conditionals. That is, compute the distribution of each of $\mu, \sigma^2, \tau^2, \phi_1, \ldots, \phi_J$, given all others and the observed data y.

(3) Implement a Gibbs sampler for simulating from the joint posterior.

(4) Simulate a data set with $J = 100, \rho = 0.8, \mu = 0, \tau^2 = \sigma^2 = 1$. Use you Gibbs sampler to simulate from the posterior. Comment on its convergence after examining the trajectories of the draws of μ, σ^2, τ^2 and the autocorrelations.

(5) Based on the Gibbs draws, plot the marginal posteriors of μ, σ^2 and τ^2 and compare with the true values. Do the true values lie within or outside of their respective 95% credible intervals?