

Solutions

Name: _____

Bayesian Statistics, 22S:138
PRACTICE PROBLEMS for Midterm 2, 2008
Midterm 2, 2007

A neurologist has four patients with epilepsy, all of whom are taking the same drug to control their seizures. They have been on the drug long enough that their pattern of seizures has stabilized.

The neurologist asks each patient to keep track of the number of seizures he or she has in each of three months. Let $x_{i,j}$ represent the number of seizures that patient i has in month j , where i goes from 1 to 4 and j goes from 1 to 3.

The neurologist believes that each patient has a true mean number of seizures per month θ_i , and that the number of seizures that each patient has in a given month is a draw from a Poisson distribution with parameter θ_i . That is,

$$x_{i,j} | \theta_i \sim \text{Poisson}(\theta_i), \quad i = 1, \dots, 4, \quad j = 1, \dots, 3.$$

and the $x_{i,j}$ s are considered conditionally independent given θ_i .

Furthermore, she believes that, although the θ_i s are different for different patients, they are similar, and that that similarity can be expressed by saying that all the θ_i s are like draws from a common Gamma density. That is:

$$\theta_i | \alpha, \beta \sim \text{Gamma}(\alpha, \beta), \quad i = 1, \dots, 4$$


Finally, she believes that the following prior densities are appropriate for the parameters α and β :

$$\alpha \sim \text{Exponential}(0.1)$$

$$\beta \sim \text{Gamma}(0.1, 1)$$

The neurologist carries out a Bayesian analysis to learn about each individual patient's mean parameter θ_i and about the density from which the θ_i s are drawn. She uses WinBUGS to run three MCMC chains using different sets of initial values. Code, plots, and the table of node statistics from the WinBUGS analysis are attached.

1. Draw a directed graph of the model.



2. According to your directed graph, which model quantities will appear in the full conditional density of α ? You do *not* need to derive the form of the full conditional density; just name the quantities that will appear in it.

β , all the θ_i 's, and the 0.1 from prior on α .

3. Circle all of the true statements in the list below.

- (a) All of the $x_{i,j}$ s for all patients and all months are treated as exchangeable in this model.
- (b) $x_{2,1}$, $x_{2,2}$ and $x_{2,3}$ are considered exchangeable in this model.
- (c) All the θ_j s are treated as exchangeable in this model.
- (d) None of the above.

4. For the parameter α , the autocorrelation between values drawn from the MCMC samplers 5 iterations apart is (circle the one best answer):

- (a) close to 0
- (b) close to 0.5
- (c) close to 1
- (d) impossible to estimate from output provided

5. How many iterations do the Gelman-Rubin plot and table for α suggest should be discarded as burn-in? Give a rough numeric answer and explain in one or two sentences how you chose it.

There are many possible good answers here. We want the red line to get close to 1 (<1.2) and stay there. This happens at about 100 iterations. We want the blue and green lines to come together and stabilize. This could be said to happen at any iteration after ~~between~~ 400.

6. What is the 95% credible set for θ_1 (numeric answer)?

(1.442, 4.871)

7. Circle all of the following statements that are true regarding the credible set in the previous question

- (a) Using different prior densities for α and β could have changed the endpoints of the credible set for θ_2 .
- (b) For the neurologist, there is 95% probability that patient number 2's true mean rate of seizures per month lies in this interval.
- (c) In 95% of all possible months, the number of seizures that patient 2 has will lie in this interval.
- (d) If the neurologist ran her WinBUGS samplers for an addition 5000 iterations, the endpoints of the credible set would be likely to change a little bit but not substantially.

8. Fill in the table below. Identify each of the expressions in the first column as one of the following types: prior, joint posterior, marginal posterior, or full conditional density. Also indicate with "yes" or "no" whether useful Bayesian inference might be based on it.

Density	Density Type	Useful for inference?
$p(\theta_1 \theta_2, \dots, \theta_4, \alpha, \beta, x)$	full condit	no
$p(\theta_1, \dots, \theta_4, \alpha, \beta x)$	joint post	may be
$p(\alpha)$	prior	no
$p(\theta_1 x)$	marginal post	yes

9. Running more than one MCMC chain from different initial values (circle all that apply):

- (a) helps locate all modes in the posterior density
- (b) reduces autocorrelation in the output from each chain
- (c) enables the use of the BGR diagnostic in WinBUGS
- (d) makes trace plots more helpful in indicating convergence failure
- (e) shrinks the estimates of posterior standard deviations

10. Suppose the neurologist had included only patients 1, 2, and 3 in her study, but had kept all other aspects of the analysis the same. Would you expect her to obtain the same posterior means and standard deviations for θ_1 , θ_2 , and θ_3 as she obtained when all 4 patients were in the analysis? Explain in one or two sentences why or why not.

No. Because in hierarchical models, estimation of group specific parameters borrows strength from observations in other groups, estimation of θ_1 , θ_2 , and θ_3 is affected by the inclusion or exclusion of patient 4's data.