

Statistics 5444: Homework 4

For each homework assignment, turn in at the beginning of class on the indicated due date. Late assignments will only be accepted with special permission. Write each problem up *very* neatly (L^AT_EX is preferred). Show all of your work.

Problem 1

In this problem, you will implement a Gibbs sampler for fitting a Cauchy regression model. Recall that a standard regression model has the form $y_i = \beta_0 + \beta_1 x_i + \epsilon$, where $\epsilon \sim N(0, \sigma^2)$. In Homework 3 you found that under the Gamma frailty model, $y_i = \beta_0 + \beta_1 x_i + \epsilon_i$ ($i = 1, \dots, N$), where $\epsilon \sim N(0, \sigma^2/\gamma_i)$ with $\gamma_i \sim \text{Gamma}(\frac{1}{2}, \frac{1}{2})$ specifies a regression model where the errors are no longer Normal, but rather they are Cauchy distributed.

Part 1

Write down the full conditional for $\phi = 1/\sigma^2$, β , and γ_i ($i = 1, \dots, N$). (Yes, you already know this from the previous HW... I just want you to do it again.)

Part 2

Under simulated points (x,y) (use 1,000 pairs), where $(x, y) \sim \text{Cauchy}(0, 1)$, with covariance structure

$$\begin{pmatrix} 1 & 0.8 \\ 0.8 & 1 \end{pmatrix},$$

plot the marginal posterior distributions for $(\beta_0, \beta_1)^T = \beta$. Recall that

$$p(\beta|Y, X) = \int \int p(\beta, \phi, \vec{\gamma}|Y, X) d\phi d\vec{\gamma},$$

where $\vec{\gamma} = \{\gamma_1, \dots, \gamma_N\}$. You may do this via Gibbs sampling. That is Initialize $\beta^{(0)}$ and $\phi^{(0)}$ and sample the joint distribution via
For $t = 1, \dots, T$

- Simulate $\gamma_i^{(t)}$ ($i = 1, \dots, N$) from its full conditional.
- Sample $\beta^{(t)}$ from its full conditional distribution given $\phi^{(t-1)}$ and $\gamma_i^{(t)}$, for ($i = 1, \dots, N$).
- Sample $\phi^{(t)}$ from its full conditional distribution given $\beta^{(t)}$ and $\gamma_i^{(t)}$ ($i = 1, \dots, N$).

And finish by plotting a histogram of the β 's (after burn-in).

Problem 2

In this problem, you will study posterior model probabilities in detail. Consider the problem where

$$x_i \sim N(\theta, \sigma^2),$$

for $i = \{1, \dots, N\}$.

Part a

Let $\sigma^2 = 5$ and simulate

- 400 x 's with $\theta = 10$
- 200 x 's with $\theta = 15$
- 300 x 's with $\theta = 17$
- 100 x 's with $\theta = 20$.

The 1,000 x_i 's form a *mixture* distribution. Plot a histogram of the distribution.

Part b

Consider the hypotheses $\theta = \{10, 15, 17, 20\}$ each with prior probability $\frac{1}{4}$. For example, the first hypothesis is simply

$$\begin{aligned}H_0 &: \theta = 10 \\H_a &: \theta = 15, 17, \text{ or } 20.\end{aligned}$$

Write down the *appropriate* prior distribution for testing this hypothesis.

Part c

Write down the Bayes factor for each hypothesis (there will be 4 of them).

Part d

Randomize the vector of x 's so that there is no particular order in which you sampled them. Plot the x 's. The horizontal axis will go from 1 to 1,000. The vertical axis will have the sample values. There should be no clear pattern in this graph.

Part e

For each hypothesis, compute the posterior probability *in sequence*. That is, look at one data point and compute the posterior probability of the hypothesis. Collect a second data point and compute the posterior using the first two data points. Do this until you have the posterior for all 1,000 data points. Plot the results for each hypothesis. That is, plot the posterior probabilities as a function of the number of data points.

Part f

Test each of the four hypothesis using P values. That is, do the two tailed test and plot the P values in sequence.

Heads up for the next HW

Write down the mixture model that you *should* have used to compute the Bayes Factors. Under this mixture model, show how to compute the *full conditional probability* for each data point belonging to an individual mixture. You do not need to do anything here, but keep this in mind for the future.

Part h

Write your conclusions.

Problem 3

Consider the binomial sampling distribution $x \sim \text{Bin}(N, p)$. Imagine that you observe 100 sample draws from this distribution (call this collection X). Consider testing the hypothesis

$$\begin{aligned}H_0 &: p = 1/2 \\H_a &: p \neq 1/2.\end{aligned}$$

Let each hypothesis have probability $\frac{1}{2}$, and let $p_a(p) \propto p^{\frac{1}{2}}(1-p)^{\frac{1}{2}}$.

Find $p(H_0|X)$.

Problem 4

Prove that under the null hypothesis H_0 , the P-value has a uniform distribution.