

Statistics 5444: Homework 1

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

Let $x_i \sim N(\mu, \sigma^2)$, for $i = 1, \dots, N$. Assume that σ^2 is a known parameter.

In this problem, we will infer from the observed x_i 's the posterior sampling distribution $p(\mu|\mathbf{X})$, where $\mathbf{X} = \{x_1, \dots, x_N\}$.

Part 1

What is the likelihood function $L(\mu|\mathbf{X})$?

Part 2

Under the *reference* prior $p(\mu) \propto 1$, find the posterior distribution for μ .

Problem 2

This problem is very similar in spirit to Problem 1. Let $X \sim \text{Bin}(N, p)$, so that $p(X = x) = \binom{N}{x} p^x (1 - p)^{N-x}$.

Consider the *reference* prior $p \sim \text{Beta}(\frac{1}{2}, \frac{1}{2})$ (We will motivate this prior distribution later). Recall that the pdf for $z \sim \text{Beta}(\alpha, \beta)$ follows as

$$p(z) = \frac{1}{B(\alpha, \beta)} z^{\alpha-1} (1 - z)^{\beta-1},$$

where $B(\alpha, \beta) = \int_0^1 z^{\alpha-1} (1 - z)^{\beta-1} dz$.

Specify the posterior distribution $p(p|x)$. That is, explicitly name the type of distribution that the posterior follows and find the parameters of the distribution.

Problem 3

This problem is again very similar to problem 1. This time we will be doing posterior inference on the ordinary least squares (simple regression) problem. Recall the basic set up of the regression problem:

$$y_i = \sum_{i=1}^p x_i b_i + \epsilon_i \quad \text{for } i = 1, \dots, N,$$

where $\epsilon_i \sim N(0, \sigma^2)$. Generally ($x_1 = 1$), so that b_1 is interpreted as an intercept term.

In matrix notation, we can write down the system of equations as

$$Y = \mathbf{X}\beta + \boldsymbol{\epsilon},$$

where $Y = (y_1, \dots, y_n)^T$ is a vector of y 's, $\beta = (b_1, \dots, b_p)^T$ is a vector of (unknown) coefficients, and \mathbf{X} is the corresponding ($N \times p$) matrix of regressors. The error term follows the multivariate normal distribution:

$$\boldsymbol{\epsilon} \sim N(\mathbf{0}, \Sigma),$$

where $\mathbf{0}$ is the ($N \times 1$) zero vector, and $\Sigma = \sigma^2 I_{N \times N}$.

Recall that the density for the multivariate normal distribution (with mean vector μ and covariance matrix Σ) is written

$$p(x|\mu, \Sigma) = \frac{1}{(2\pi)^{N/2} |\Sigma|^{1/2}} e^{-\frac{1}{2}(x-\mu)^T \Sigma^{-1} (x-\mu)},$$

where $|\cdot| = \text{Det}(\cdot)$.

Part 1

Under the reference prior $p(\beta) \propto 1$, find the posterior distribution for β .

Part 2

Recall that the MLE estimate is $\hat{\beta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T Y$. Find $\text{Var}(\hat{\beta})$ and compare the posterior estimate to MLE estimate.

Problem 4

See P.M. Lee, page 10. The top paragraph illustrates the *prosecutor's fallacy*. After reading, the Author states that the two (under the scenario illustrated in the paragraph) conditional distributions $p(E|I)$ and $P(I|E)$ are “equal if and on if the prior probability $P(I)$ of innocence is taken to be $\frac{1}{2}$ ”. Justify his claim (i.e. show that the prior must equal 1/2 (approximately??)).

Problem 5

Consider $X \sim \text{Bin}(N, p)$, So that $p(X = x) = \binom{N}{x} p^x (1 - p)^{N-x}$. Consider the uniform prior $p(p) = 1$, where $0 \leq p \leq 1$. The posterior sampling distribution for this should be clear from Problem 2 (note that this prior is a $\text{Beta}(1, 1)$).

The prior that you have specified places an equal amount of weight on every possible value of p .

However, some people like to work with the log-odds, which we write as

$$\Lambda = \log \left(\frac{p}{1 - p} \right).$$

Part 1

Find $p(\Lambda)$. That is, find the pdf for Λ (this is just a simple transformation problem).

Part 2

Now, do the problem the other way around. Consider placing a uniform prior on (Λ) . That is, let $p(\Lambda) \propto 1$. What is the implied prior distribution on p ?

Problem 6

Consider the binomial sampling distribution $p(x) = \binom{N}{x} p^x (1 - p)^{N-x}$. In classical inference, once you observe data (x) , you might write down the (95%) confidence interval for p as

$$\hat{p} \pm 1.96 \times \sqrt{\frac{\hat{p}(1 - \hat{p})}{N}}, \quad (1)$$

where $\hat{p} = \frac{x}{N}$ and $\hat{\sigma} = \sqrt{\frac{x}{N}(1 - \frac{x}{N})}$. This of course follows from the asymptotic normality assumption of the estimator \hat{p} , and holds for sample sizes of about 30 or larger (t-intervals are used for smaller sample sizes).

Generate a uniform spacing of p 's from .01 to .99, with a step size of .01. Under each value of p , generate a Binomial random number with $N = 30$. Now construct the confidence interval, given by (??). Since you know the actual value of p , you can confirm if this interval does in fact cover the *true* parameter.

Repeat this 10,000 times for each value of p , and keep track of the frequency of times each of 10,000 repetitions cover the known parameter. Show a plot of the behavior of the "true" coverage probabilities for the full range of p 's (0.01, 0.02, 0.03, ..., 0.99) (The y-axis will have frequency on it and the x-axis will be over p).

Repeat the exercise with $N = \{50, 100\}$.

Conclude with your thoughts on the experiment. Are you surprised?