Last updated on March 11, 2026

BM5163 Bayesian Inference in Bioengineering

Problem Set 3

Instructions

1. You are expected to work on these problems on your own and not submit the solutions.

Questions

1. For the following data from experiments, estimate distribution parameters using the maximum likelihood estimate (MLE)
   (a) data follows an exponential distribution
   (b) data follows a uniform distribution
   (c) data follows a normal distribution
2. For problems 1-3 in problem set 2, calculate 95% equal-tailed and HPD credible intervals.
3. Consider the following problem from the first exam

   "Suppose one student misses this exam for some genuine reason. Once all the exams are over at the end of the course, can you design a systematic approach to award the marks to the student for the missed exam? You can assume that only one student has missed one exam. All other exams were attended by all the students. Give all the details of your calculation. "

   Currently, BM5163 has 15 students, and there are 5 exams scheduled (including the project). If the marks scored by the jth student in the ith exam are y_{ij}, estimate the marks for one student who missed one exam. You can also write a `Python` program that takes scores of all the students as input and gives the average value and 95%-HPD interval for the marks.
4. You lose your keys. You estimate an 80% chance they are on the counter and a 20% chance they are in your jacket pocket. You glance at the counter. If they are there, a quick glance has a 90% chance of spotting them. If they are not there, you obviously won't spot them (0%). You do not spot the keys. What is the updated probability that the keys are in your jacket pocket?

   Note: the following questions are not necessarily well defined. You have to analyze the situations and also identify the information that is relevant for making a decision.

5. Your close friend, who usually writes long, enthusiastic replies to your messages, responds to one of your WhatsApp messages with a single letter: "K". You know from past experience that on the rare occasions they are genuinely furious with you, they almost always send short, one-word texts. Because of this, you immediately panic. Is your fear justified?
6. You receive an email from an unknown sender with the subject line "URGENT: Your account has been compromised." You recently read a cybersecurity report stating that scammers use the word "URGENT" in roughly 50% of all their phishing emails. Your friend looks at your screen and says, "Don't click that! If half of all scams use that word, there's a 50% chance you're being scammed right now." Is your friend right?
7. The Mumbai police is investigating a case with zero suspects and zero leads. They find a tiny speck of blood at the scene and run the DNA profile through a national criminal database of 30 million people. The computer spits out a single match which is a man living in a non-neighboring state. The DNA lab testifies that the test has a false-positive rate of only 1 in 1 million. At trial, the prosecutor points to the defendant and says, "The test is 99.9999% accurate. Therefore, there is only a 1 in 1 million chance this man is innocent." Is the prosecutor right? How does the fact that the detective searched a database of 30 million people fundamentally change the weight of the evidence compared to if they had a single prime suspect to begin with?
8. You are a science journalist reviewing two newly published nutritional studies.

- Study A shows that eating 200 fewer calories a day leads to weight loss.
- Study B shows that listening to extreme heavy metal music for 10 minutes a day leads to weight loss.

Both studies are conducted by independent university labs, both used sample sizes of 100 people, and both achieved the exact same statistical significance ($p < 0.05$, meaning there is less than a 5% chance of seeing these results if the intervention actually did nothing). A strict frequentist statistician might argue that both studies provide the exact same "strength" of evidence since their p-values are identical. How would you report these two studies?

9. You receive an unsolicited email on Monday from a financial analyst predicting a specific stock will go up this week. It goes up. Next Monday, he emails predicting it will go down. It goes down. This happens perfectly for six consecutive weeks. On the seventh week, he emails you saying, "I have proven my system works flawlessly. Pay me INR 100,000 for my prediction for week 7." Will you do it?

10. Now suppose you are the scamster/financial analyst from the last question. On the first Monday, you had sent 1024 emails predicting a particular stock will go up and 1024 emails claiming it will go down. It went up, but not because of any particular reason to your knowledge. So you selected the email addresses of those to whom you had predicted it would go up, and divided the whole group into two of 512 each. To one group of 512 you sent another email that it would go up and the other 512 it would go down. Of course, you turned out to be correct for one of those groups of 512. You divide them again and keep doing it for 6 weeks. For the seventh week, you demanded INR 100,000 for your prediction. What is your maximum likelihood estimate of the amount you will get through this scam?