Homework 7: Concentration and MLEs

For each problem, remember you must briefly explain/justify how you obtained your answer, as correct answers without an explanation will not receive full credit. Moreover, in the event of an incorrect answer, we can still try to give you partial credit based on the explanation you provide.

In general, your goal in an explanation is to write enough that a student from class who has attended lecture, but not read the problem yet, could understand your approach, verify your reasoning, and believe your answer is correct. While we do not usually need to see arithmetic, you must include enough work that in principle one could rederive your answer with only a scientific calculator. For each problem, make sure to explicitly define all random variables you use, and be clear about how they are related to each other using proper notation (conditionals, summations, etc.).

Unless a problem states otherwise, you should leave your answer in terms of factorials, combinations, etc., for instance 26^7 or $26!/7!$ or $26 \cdot \binom{26}{7}$ are all good forms for final answers.

Instructions as to how to upload your solutions to gradescope are on the course web page.

Remember that you must tag your written problems on Gradescope.

Submission: You must upload a **pdf** of your written solutions to Gradescope under "Homework 8". (Instructions as to how to upload your solutions to gradescope are on the course web page.) The use of latex is *highly recommended*. (Note that if you want to hand-write your solutions, you'll need to scan them. We will take off points for hand-written solutions that are difficult to read due to poor handwriting and neatness.)

Due Date: This assignment is due Monday August 18th at 11:59 PM.
Late Deadline: The last day to submit will be Tuesday August 19th at 11:59.

We will release solutions to HW8 on Ed on Wednesday August 20th so you can look at them before the final exam. We do not expect to get grades on HW8 back before the final.

Academic Integrity: Please read the full academic integrity policy. If you work with others (and you should!), you must still write up your solution independently and name all of your collaborators in the separate question on gradescope.

1. Sewing Requires Concentration [16 points]

Kieran is attempting to use up their (embarrassingly large) collection of fabric scraps by making a patchwork jacket, made up of 30 individual squares. It takes an average of half an hour for Kieran to attach a square, with a variance of 15 minutes. The time to attach each square is independent.

You should treat time as continuous for this problem.

(a) What is the expectation of the total time to sew the jacket? [2 points]

(b) What is the variance of the total time to sew the jacket? [2 points]

(c) Kieran wants to wear the jacket to lecture, which is in 4 hours. Use Markov's Inequality to bound the probability that Kieran finishes the jacket in time. [6 points]

(d) It's supposed to start raining in 16 hours, and Kieran should really wear a jacket to campus. Use Chebyshev's inequality to bound the probability that Kieran finishes the jacket after at least 14 hours, but within 16 hours. [6 points]

2. Balloons [15 points]

There are 100 city blocks of Seattle in a 10x10 grid (assume for this problem that Seattle is actually laid out in a logical and consistent way). A drone will drop 2500 balloons onto the blocks, with each balloon having equal probability of landing in any of the blocks. No balloons will fall outside the 10x10 grid, and each balloon is independent of the others.

You hope that at the end of the process, each block will have at least 10 balloons. You want to upper bound the probability that you fail to distribute at least 10 balloons to every block.

(a) Use a Chernoff bound from class to bound the probability that the southwest-most block does not get at least 10 balloons. [7 points]

(b) Is the probability that the southwest block gets less than 10 balloons independent of the probability that the northeast block gets less than 10 balloons? Briefly explain (you may give a formal derivation/calculation as an explanation or an informal one). [3 points]

(c) Bound the probability that at least one block gets less than 10 balloons. Give the best bound you can from your answers in (a) and (b). [5 points]

3. Card roulette [15 points]

Zitong, Jeremiah, and Victor love to eat at local restaurants. Each time, they let the waiter decide who pays for the meal.

Every waiter independently chooses to have Zitong pay for the meal with probability θ_Z, Jeremiah pays with probability θ_J, and Victor $1 - \theta_J - \theta_Z$ (Note $0 \leq \theta_J + \theta_Z \leq 1$. The parameters of θ_Z, θ_J are unknown. Suppose that $y_1, ..., y_n$ are n independent, identically distributed samples from this distribution of payer choices. Let x_Z be the number of y_i's equal to Zitong, let x_J be the number of y_i's equal to Jeremiah, and x_V the number equal to Victor. What are the maximum likelihood estimates for θ_J and θ_Z in terms of x_J, x_V, x_Z.

In doing this problem, you do not need to do a second-derivative test (or any other test) to confirm you have a maximizer. You may assume any critical point you find is a maximizer.

For this problem, except where noted, all of your work must be understandable without reference to any calculator (including Wolfram-Alpha). You may check your answers using calculators, but your explanation may not rely on them.

(a) Write the Likelihood function. [3 points]

(b) Write the log-likelihood. [2 points]

(c) Give two equations describing the maximizer of the log-likelihood. Be sure to show work (e.g., taking of partial derivatives) on how you got these equations. [8 points]

(d) You should now have a system of two (linear) equations in two unknowns. Find the solution to this system, and use it to write the MLEs for θ_Z and θ_J. You do **not** have to show your work for this part, and may use online algebra solvers if you wish. [2 points]

4. Unbiased Estimator [5 points]

Let $X_1, X_2, ..., X_n$ be random samples drawn from the continuous distribution Uniform$(1, \theta)$, for $\theta > 0$. We stumble upon a possible estimator $\hat{\theta}$ for θ, where $\hat{\theta} = (2 \cdot \sum_{i=1}^{n} X_i/n) - 2$. Is $\hat{\theta}$ an unbiased estimator of θ?

5. Feedback + Collaboration [1 point]

Answer these questions on the separate Gradescope box for this question.

Please keep track of how much time you spend on this homework and answer the following questions. This can help us calibrate future assignments and future iterations of the course, and can help you identify which areas are most challenging for you.

- Which students did you collaborate with for this homework?
- Is the work that you are submitting your own and does not violate the academic integrity policy outlined in the syllabus?
- How many hours did you spend working on this assignment (excluding any extra credit questions, if applicable)? Report your estimate to the nearest hour.
- Which problem did you spend the most time on?
- Any other feedback for us?