Homework 6: Central Limit Theorem + Multiple RVs

For each problem (except those marked as autograded), 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 "HW 6". (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 at 11:59 PM Wednesday August 13th.
You will submit the written problems as a PDF to gradescope.
For the written problems, please put each numbered problem on its own page of the pdf (this will make selecting pages easier when you submit), and ensure that your pdfs are oriented correctly (e.g. not upside-down or sideways).

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. Extra Instructions :o

-1.1. Evaluating Integrals

For calculations that require evaluating integrals (unless we indicate otherwise), you must

(a) Show the integral to evaluate (e.g., $\int_0^2 z \cdot 2dz$)

(b) Show an antiderivative and the values to evaluate at (e.g., $z^2|_0^2$)

(c) Plug in the values and simplify (e.g., $2^2 - 0^2 = 4$)

This is not a problem, so nothing needs to be submitted here.

0. Gradescope Quesitons

Answer the questions on gradescope about CLT, continuity correction, and covariance.

1. Confidence intervals in "real life" [24 points]

In all of the following use the Central Limit Theorem. Use continuity correction if (and only if) you're approximating a discrete random variable.

(a) You are in debt, but luckily you are good at math! You independently owe each of your 30 friends some amount of money by the end of the month, but you aren't sure exactly how much. You do know that you owe each friend an average of $\$135$ with a standard deviation of $\$34$. How much money do you need to be at least 99% sure that you will be able to pay everyone back? Give your answer to the nearest cent. You should treat the amounts of money that you owe as continuous.

(b) After unfortunately failing to pay back your friends, you've decided to turn to other means of earning money: cryptocurrency. After a bit of research, you've narrowed down your focus to 9 different currencies. You're thinking of investing $\$14000$ in every currency. You know you will gross $\$55000$ (for a net return of $\$55000 - \$14000 = \$41000$) with probability p, and with probability $1-p$ this currency will not be a trading at a price you will want to sell (for a net return of $-\$14000$). Since cryptocurrency trading is based primarily on hype, the 9 cryptocurrencies are mutually independent.

Your local bank has agreed to loan you the money as long as your net return from these investments is positive with at least 95% probability. What is the condition on p that needs to be true in order to secure the loan? You should treat the amount of revenue you'll get from these investments as discrete (since the net return will only increase in multiples of $\$41000$ and $-\$14000$).

(c) You've decided to try a less-risky occupation – you were just put in charge of the TA-hiring process and told that we need to hire 312 TA's (that is, you want 312 people, not necessarily for the course CSE 312). You will first choose the number of people, n, to interview. Each of the n interviewees will meet your standards and be hired (independently) with probability 0.75. At least how big should n be to guarantee that you will definitely have at least 312 TA's with probability 95%.

2. Joint Continuous Densities [12 points]

Let X and Y be continuous random variables with the following joint distribution.

$$f_{X,Y}(x,y) = \begin{cases} \frac{1}{243}xy^5 & 0 \le x \le 2, 0 \le y \le 3 \\ 0 & \text{otherwise} \end{cases}$$

(a) Find $\mathbb{P}(X > Y)$ [5 points].

(b) What is the marginal PDF of X? [3 points]

(c) What is the marginal PDF of Y? [3 points]

(d) Are X and Y independent? [1 point]

3. UW Escape Room [15 points]

Kieran has a free afternoon on campus, but can't decide how to spend it. They roll a fair four-sided die to choose between four options:

- If Kieran rolls a 1, they will go to Suzallo to grade some (lovely, wonderful) homework submissions for a number of hours which is Exponential with parameter $\frac{1}{2}$, then go home.
- If Kieran rolls a 2, they will go to Odegaard to study for a number of hours which is Poisson with parameters 5, then go home.
- If Kieran rolls a 3, they will go to Gates to attend various staff meetings for a number of hours which is Geometric with parameter $\frac{1}{3}$, then go home.
- If Kieran rolls a 4, they will get distracted by a butterfly in the quad for 2 hours, then roll again.

Use the law of total expectation to compute the expected number of hours Kieran spends on campus.

4. Doom Scrolling [24 points]

Caleb is going to scroll reels for H hours, where H is a random variable, equally likely to be 1, 2, 3, or 4. The number of mythical-reels M that he sees is random and depends on how long he scrolls for. We are told that

$$\mathbb{P}(M = m \mid H = h) = \frac{c}{h}, \qquad \text{for } m = 1, \ldots, h,$$

for some constant c.

(a) Compute c. You may want to use one of the axioms of probability.

(b) Find the joint distribution of M and H using the chain rule.

(c) Find the marginal distribution of M.

(d) Find the conditional distribution of H given that $M = 1$ (i.e., $Pr(H = h \mid M = 1)$ for each possible h in 1,2,3,4). Use the definition of conditional probability and the results from previous parts.

(e) Suppose that we are told that Caleb saw either 1 or 2 mythical reels. Find the expected number of hours he spent scrolling conditioned on this event. Use the definition of conditional expectation and conditional probability theorems.

(f) Caleb then decides to stream his reaction to each mythical reel he finds. The amount of money Caleb earns for streaming each reaction is a random variable with mean 2 and the money each time is independent of the other times and of the number of mythical reels Caleb sees. What is the expected amount of money Caleb earns?

Warning: you might be tempted to skip some steps and assert that expected amount of money earned is 2 times the expected number of mythical reels seen. Even though the answer is intuitive, its formal derivation is a lot more involved. We expect you to show each step you used to get to that expression. Your work must involve the law of total expectation, conditioning on the number of mythical reels seen. We would encourage you to also use indicator random variables, but other simple random variables may suffice.

5. Real World: ML Fairness

One of the most popular (and influential) applications of the tools of this course to the real world has been in machine learning. You have probably (hopefully!) heard that ML applications have had unintended effects in the real world, especially around fairness. In this assignment, we will use the tools of probability to get a taste of what fairness analysis looks like. Our specific goals are:

- Practice some old concepts (before the final) by applying them in a new context.
- Practice converting between notation and a more intuitive understanding.
- See that probability and conditioning are useful formal representations of what it might mean to be fair.
- Realize that "fair" has more than one reasonable definition, and you can't necessarily have all of them!
- Get a sense that fairness is complicated!

This will **not** be a full introduction to fairness in ML; the topic could easily fill an entire course. This assignment is intended to help you realize that the topic is important, requires careful attention to understand, and that it is a possible application of the tools you learned in this course.

When thinking about ML systems, we need to be careful about what it means to be "fair." A definition that leads to equitable outcomes in one setting may have the opposite effect (setting in stone historic inequality) in another scenario. The actual definition of "fair" to pick is thus critical – we can't know if we're being fair or not if we don't know what fair means! And (as we see next), definitions of fair that might be reasonable in one context or another are sometimes even contradictory! You will see three definitions of "fairness" that have been used in the literature. We'll then consider one potential area of application for these definitions – you'll choose which of the definitions you believe is appropriate for that setting.

There are other questions one must ask about an ML system before using it (should one use an ML system at all? If so what data is appropriate? If so should a human be in-the-loop and how?) As we said right at the start of this assignment these issues are subtle and require careful thought! We can't handle every aspect in a single assignment, so for this one we limit ourselves to just defining fairness.

The basic operation of machine learning is to take in a large set of data, find patterns in that data (using statistics), and make predictions about future data based on the given data.

5.1. Three Definitions of Fairness [6 points]

ML Fairness is usually discussed in the context of disparate outcomes across social groups, and often in terms of effects on protected classes (e.g., based on a person's sex or gender or race—the linked Wikpedia article has a longer list of examples).

To make the examples more concrete, imagine that our ML application is for a bank deciding on making loans. Our ML system will look at historical data of who repaid loans, and use that to predict whether a new person will repay their loan.

Suppose we choose a loan applicant at random: Let L be the indicator that our ML system says our applicant should get a loan, let A be the the indicator that the applicant has the protected-class attribute which we are examining, and let Y be the indicator that the person would truly successfully repay a loan.[1]

A "fairness criterion" is a possible definition of what constitutes "being fair." In this assignment we'll see three definitions of "fair" that have been used in the fairness literature. For example, "precision parity" is defined as $\mathbb{P}(Y=1|L=1,A=1)=\mathbb{P}(Y=1|L=1,A=0)$. That is, if the ML system says our applicant should get a loan, the probability that this applicant will repay the loan remains the same regardless of whether they have our protected attribute or not. Intuitively, the goal of this definition is to ensure that the ML system does not give more "benefit of the doubt" to one group than another (by ensuring that both groups have the same fraction of repaid loans).

(a) Another possible fairness criterion is "true positive parity": $\mathbb{P}(L=1|Y=1,A=1)=\mathbb{P}(L=1|Y=1,A=0)$. Give a (1-2 sentence) intuitive description of what this criterion is trying to achieve. [3 points]

(b) A third possible fairness criterion is "false positive parity." This definition seeks to ensure that among those who do not repay their loan, the probability of our system wanting to give them the loan is equal whether they have the protected attribute or not. Intuitively, the goal is not to give "unfair" benefits (loans unlikely to be repaid) to people without (or with) the protected attribute.

Give the mathematical notation for false positive parity. [3 points]

5.2. You Can't Have it All [6 points]

A theorem[2] says that an ML system cannot achieve all of precision parity, true positive parity, and false positive parity at the same time unless some unrealistic conditions are met[3].

(a) For concreteness, pick an attribute which you want to ensure is treated fairly for this section. You can pick anything from the list of protected classes on wikipedia or another attribute you pick. Tell us which one you picked.

[1] To make it possible to know Y (without using the system to decide on real loans before knowing whether it works) ML systems are usually evaluated on a "test set." A dataset gathered in advance where the true answers (the Y's) are known, but the data was not used to design the model.

[2] proven in the last 10 years!

[3] Specifically, those exceptions are: if the ML system never makes a mistake (which never happens at scale—there are simply too many variables for real people to understand them all perfectly) or if $\mathbb{P}(Y=1|A=0)=\mathbb{P}(Y=1|A=1)$, that is the actual repayment rate of loans is equal regardless of the protected attribute (which, since protected classes are defined for protection because of historic disparate treatment, is uncommon in real data).

(b) For the loan context, suppose you could make the system meet exactly one of the three fairness criteria above. Which of the three definitions of fairness would you want to ensure your ML system meets, and why? Any of the three can be chosen, but we want your answer to show an understanding of the definition and relate to the real-world; we won't grade based on whether we agree with you, just whether you've discussed your reasoning. [3-4 sentences]

(c) Choose one of the other fairness criteria (not the one you chose in the last part). Think of a scenario (something other than "deciding on whether to award loans") and an attribute where you think that idea of fairness would be the most appropriate of the three. Explain why in that scenario your choice is a good one. [3-4 sentences]

5.3. In the news [8 points]

Find an news article about a particular instance of bias in a machine learning (or artificial intelligence) system. We expect the questions in this section can be answered in 1-2 sentences each.

(a) Tell us the title of the article, the source (e.g., "New York Times") and give us a link to it (don't worry about pretty formatting).

(b) What is the ML application trying to learn (e.g. "housing prices" or "recidivism rates")? What group does the article suggest is being treated unfairly?

(c) An ML application can be (roughly) grouped into 3 steps:

  (i) Data acquisition – collecting the data that will be the basis of the system

  (ii) Model selection – choosing which "features" of the data to use in the system and how those data points will be combined (e.g. for a house pricing model, one might decide whether or not to include the number of rooms, square-footage, and time that it has been on the market without selling as 'features' to predict the selling price)

  (iii) Use of prediction – once the model is trained, it will be used to make predictions (the model thinks this house that just went on the market will be worth Y dollars), and those predictions will then be used in the real-world (e.g. a company might decide whether to invest in a property as a result of comparing the model's estimate of its value to its current price)

  Does the article you read try to place the blame on any of these steps? If so, which one(s)?

(d) Does the article state what the author thinks "fairness" would be? If not, is it implied? The definition might match the definitions in this assignment or be something different (there are other fairness criteria we haven't covered).

(e) For this application, what definition of fairness do you think is appropriate (this could be one of the three definitions above, or it could be another definition). Your definition here can be phrased intuitively or in the language of probability (you don't have to list both).

6. 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?