Chris Piech
CS109

Fall 2018
Oct 30th

CS109 Midterm Exam

This is a closed calculator/computer exam. You are, however, allowed to use notes in the exam. The last page of the exam is a Standard Normal Table, in case you need it. You have 2 hours (120 minutes) to take the exam. The exam is 120 points, meant to roughly correspond to one point per minute of the exam. You may want to use the point allocation for each problem as an indicator for pacing yourself on the exam.

In the event of an incorrect answer, any explanation you provide of how you obtained your answer can potentially allow us to give you partial credit for a problem. For example, describe the distributions and parameter values you used, where appropriate. It is fine for your answers to include summations, products, factorials, exponentials, and combinations, unless the question specifically asks for a numeric quantity or closed form. Where numeric answers are required, the use of fractions is fine.

I acknowledge and accept the letter and spirit of the honor code. I pledge to write more neatly than I have in my entire life:

Signature: ____

Family Name (print): ____

Given Name (print): ____

Email (preferably your gradescope email): ______________________________

1 Random Writer [16 points]

Consider random strings of length four made from the 26 lower case English alphabet letters (where each letter is equally likely and letters can occur in the string more than once).

a. How many unique strings of length four are there?

b. There are 5,500 four letter words in English. What is the probability that a randomly generated string of length four is an English word?

c. You generate 100,000 *unique* four letter strings. What is the probability that at least one is an English word?

2 Race Condition [16 points]

You are reviewing code for a server and realize that it has a "race condition" meaning that if more than one "request" occurs in the same second, the server will crash. Requests occur at an average rate of 1 a minute.

a. What is the probability of exactly 3 requests in a given minute?

b. What is the average rate of requests per *second*?

c. What is the probability the server will crash in a given second (ie it receives more than one request)?

d. Let p be the answer to part (c). Give an expression for the probability that the server will crash on a given day. Note: There are 86,400 seconds in a day.

3 Algorithmic Fairness [24 points]

An artificial intelligence algorithm is being used to make a binary prediction (G for guess) for whether a person will repay a microloan. The question has come up: is the algorithm "fair" with respect to a binary demographic (D for demographic)? To answer this question we are going to analyze the historical predictions of the algorithm and compare the predictions to the true outcome (T for truth). Consider the following joint probability table from the history of the algorithm's predictions:

	D = 0		**D = 1**	
	G = 0	G = 1	G = 0	G = 1
T = 0	0.21	0.32	0.01	0.01
T = 1	0.07	0.28	0.02	0.08

D: is the demographic of an individual (binary).
G: is the "repay" prediction made by the algorithm. 1 means predicted repay.
T: is the true "repay" result. 1 means did repay.

Recall that cell $(D = i, G = j, T = k)$ is the probability $P(D = i, G = j, T = k)$. For all questions, justify your answer. You may leave your answers with terms that could be input into a calculator.

a. What is $P(D = 1)$?

b. What is $P(G = 1|D = 1)$?

c. Fairness definition #1: Parity
An algorithm satisfies "parity" if the probability that the algorithm makes a *positive prediction* ($G = 1$) is the same regardless of the demographic variable. Does this algorithm satisfy parity?

This is the same table, for your convenience:

	D = 0		D = 1	
	G = 0	G = 1	G = 0	G = 1
T = 0	0.21	0.32	0.01	0.01
T = 1	0.07	0.28	0.02	0.08

d. Fairness definition #2: Calibration
An algorithm satisfies "calibration" if the probability that the algorithm is *correct* ($G = T$) is the same regardless of demographics. Does this algorithm satisfy calibration?

e. Fairness definition #3: Equality of odds
An algorithm satisfies "equality of odds" if the probability that the algorithm predicts a positive outcome *given* given that the true outcome is positive ($G = 1|T = 1$) is the same regardless of demographics. Does this algorithm satisfy equality of odds?

4 Traffic Light [20 points]

Every day you bike to work and your commute is impacted by traffic lights. You are determined to figure out the probability distribution of how long you have to wait.

For all lights on your commute: when you arrive at the light, there is a 50% chance that the light is green and a 50% chance that the light is red (we treat yellow as green). If the light is green, your wait time is 0. If the light is red, your wait time is equally likely to be any value in the range 0 to 4 mins.

a. What is the probability you wait more than 2 minutes at one light?

b. What is your expected wait time at one light?

c. What is the variance of your wait time for one light?

d. You pass through 5 lights on your way to work (each light works in the same way). What is your expected total wait time?

5 Drug Effectiveness [24 points]

You are working on an algorithm to assist a doctor treating a disease. Research at Stanford has shown that there are two subtypes of the disease (A and B). Subtype A is more common; 80% of people with the disease have subtype A. We don't have a test for which subtype a patient has, but a patient's subtype changes how effective the standard "treatment" is.

a. If the patient has subtype A, there is a 0.6 probability that the treatment is effective. If the patient has subtype B, there is a 0.2 probability that the treatment is effective. A patient with the disease (unknown subtype) is given the treatment and it is *in*effective. What is the new probability that the patient has subtype A?

b. The classification “effective” and “ineffective” is a bit reductive. Instead the effective*ness* of the treatment can be measured as a number:

The effectiveness of the treatment:
Given that the patient has subtype A, is gaussian: $X_A \sim N(\mu = 60, \sigma^2 = 25)$
Given that the patient has subtype B, is gaussian: $X_B \sim N(\mu = 20, \sigma^2 = 16)$

A patient with the disease (unknown subtype) is given the treatment and it has an effectiveness of 30. What is the new probability that the patient has subtype A? Reminder: you do not have to simplify exponents.

6 Simulation Accuracy [20 points]

You are running simulations to estimate the probability of an event. Let X be the number of times the event occurs in n independent simulations. You estimate the event probability to be $\approx \frac{X}{n}$.

We say that an estimate is "good enough" if the difference between the true probability (p) and the estimate is less than epsilon: $-\varepsilon < \frac{X}{n} - p < \varepsilon$.

Calculate the probability that your estimate will be good enough if: $p = 0.7$, $\varepsilon = 0.01$ and $n = 1000$. Use an approximation to get a closed form expression. You may leave your answer in terms of the Φ function. It is ok for your answer to include square roots.

That's the last question of the exam! We hope you had fun. Algorithmic fairness is an important (and growing) field of CS. There are over 21 mathematical definitions of fairness, several of which were derived here at Stanford!