


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

Consider random strings of length four made from the 26 lower case English alphabet letters (where each letter is equally likely).

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

Answer:

$$26^4$$

- 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?

Answer:

$$P(\text{English}) = \frac{5,500}{26^4}$$

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

Answer:

Lets consider the probability that no string was an English word. Producing 100,000 unique four letter strings is the same as choosing 100,000 strings from the set of all four letter strings. The equally likely sample space is thus:

$$|S| = \binom{26^4}{100,000}$$

The event space is all the ways you can choose 100,000 strings from the words which are not English:

$$|E| = \binom{26^4 - 5,500}{100,000}$$

$$P(\text{no english}) = 1 - \frac{|E|}{|S|} = 1 - \frac{\binom{26^4 - 5,500}{100,000}}{\binom{26^4}{100,000}}$$

2 Race Condition [20 points]

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

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

Answer:

Let Y be the number of requests in a minute. $Y \sim \text{Poi}(\lambda = 1)$

$$P(Y = 3) = \frac{\lambda^k \cdot e^{-\lambda}}{k!} = \frac{1^3 \cdot e^{-1}}{3!} \approx 0.06$$

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

Answer:

$$1/60$$

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

Answer:

Let X be the number of requests in a second. $X \sim \text{Poi}(\lambda = 1/60)$

$$\begin{aligned} P(\text{crash}) &= 1 - P(X = 0) - P(X = 1) \\ &= 1 - \frac{\lambda^0 \cdot e^{-\lambda}}{0!} - \frac{\lambda^1 \cdot e^{-\lambda}}{1!} \\ &= 1 - e^{-\frac{1}{60}} - \frac{e^{-\frac{1}{60}}}{60} \approx 0.00014 \end{aligned}$$

- 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.

Answer:

$$\begin{aligned} P(\text{crash}) &= 1 - P(\text{no crash}) \\ &= 1 - (1 - p)^{86,400} \\ &= 1 - (1 - 0.00014)^{86,400} \\ &= 0.999994 \end{aligned}$$

Its going to crash...

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 “loan-repay” prediction made by the algorithm (binary)

T : is the true “loan-repay” result (binary)

Recall that cell in the joint table where $(A = i, C = j, Y = k)$ is the probability, $P(A = i, C = j, Y = k)$. For all questions, justify your answer. You may leave your answers with terms that could be input into a calculator.

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

Answer:

Marginalize by adding the mutually exclusive cases where $D = 1$:

$$\begin{aligned}
 P(D = 1) &= \sum_{i,j \in \{0,1\}} P(D = 1, G = i, T = j) \\
 &= 0.01 + 0.01 + 0.02 + 0.08 = 0.12
 \end{aligned}$$

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

Answer:

$$\begin{aligned}
 P(G = 1|D = 1) &= \frac{P(G = 1, D = 1)}{P(D = 1)} \\
 &= \frac{P(G = 1, D = 1, T = 0) + P(G = 1, D = 1, T = 1)}{P(D = 1)} \\
 &= \frac{0.01 + 0.08}{0.12} = 0.75
 \end{aligned}$$

- 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?

Answer:

The algorithm satisfies parity if $P(G = 1|D = 0) = P(G = 1|D = 1)$ OR $P(G = 1|D = 1) = P(G = 1)$

$$P(G = 1) = 0.01 + 0.08 + 0.32 + 0.28 = 0.69$$

No: $P(G = 1|D = 1) \neq P(G = 1)$

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?

Answer:

The algorithm satisfies calibration if $P(G = T|D = 0) = P(G = T|D = 1)$

$$\begin{aligned} P(G = T|D = 0) &= P(G = 1, T = 1|D = 0) + P(G = 0, T = 0|D = 0) \\ &= \frac{0.28 + 0.21}{0.88} \approx 0.56 \\ P(G = T|D = 1) &= P(G = 1, T = 1|D = 1) + P(G = 0, T = 0|D = 1) \\ &= \frac{0.08 + 0.01}{0.12} = 0.75 \end{aligned}$$

No: $P(G = T|D = 0) \neq P(G = T|D = 1)$

e. Fairness definition #3: Equality of odds

An algorithm satisfies “equality of odds” if the probability that the algorithm *predicts a positive outcome* ($G = 1$) is the same regardless of demographics *given* that the outcome will occur ($T = 1$). Does this algorithm satisfy equality of odds?

Answer:

The algorithm satisfies equality of odds if $P(G = 1|D = 0, T = 1) = P(G = 1|D = 1, T = 1)$

$$\begin{aligned} P(G = 1|D = 1, T = 1) &= \frac{P(G = 1, D = 1, T = 1)}{P(D = 1, T = 1)} \\ &= \frac{0.08}{0.08 + 0.02} = 0.8 \\ P(G = 1|D = 0, T = 1) &= \frac{P(G = 1, D = 0, T = 1)}{P(D = 0, T = 1)} \\ &= \frac{0.28}{0.28 + 0.07} = 0.8 \end{aligned}$$

Yes: $P(G = 1|D = 0, T = 1) = P(G = 1|D = 1, T = 1)$

4 Traffic Light [22 points]

Every day you bike to work. But 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 orange 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. Provide a probability density function for the likelihood of non-zero wait times at a light. The integral of your probability density over the range of possible wait times must equal the probability that you hit a red light.

Answer:

I want to see a function:

$$f_X(x) = \frac{1}{8} \text{ if } x \in (0, 4]$$

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

Answer:

This is most easily answered by symmetry:

$$P(W > 2) = \frac{1}{4}$$

- c. What is your expected wait time at one light?

Answer:

Perhaps some students could just figure this out:

$$\begin{aligned} E[W] &= \int_{w=0}^4 \frac{w}{8} dw \\ &= \frac{1}{16} [w^2]_0^4 = 1 \end{aligned}$$

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

Answer:

This one needs some math:

$$\begin{aligned} E[W^2] &= \int_{w=0}^4 \frac{w^2}{8} dw \\ &= \frac{1}{24} [w^3]_0^4 = \frac{8}{3} \\ \text{Var}(W) &= E[W^2] - E[W]^2 = \frac{8}{3} - 1^2 = \frac{5}{3} \end{aligned}$$

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

Answer:

Let T be the total weight time $T = \sum_{i=1}^5 W_i$

$$\begin{aligned} E[T] &= E\left[\sum_{i=1}^5 W_i\right] \\ &= \sum_{i=1}^5 E[W_i] \\ &= 5 \end{aligned}$$

5 Drug Effectiveness

You are working on an algorithm to assist a doctor treating a disease. Research at Stanford has shown that there are two varieties of the disease (A and B). A is more common, occurring in 80% of the population. We don't have a test for the different subtypes, 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 is given the treatment and it is *ineffective*. What is the new probability that the patient has subtype A?

Answer:

This is your classic Bayes:

$$\begin{aligned} P(A|E^C) &= \frac{P(E^C|A)P(A)}{P(E^C|B)P(B) + P(E^C|A)P(A)} \\ &= \frac{(0.4)(0.8)}{(0.8)(0.2) + (0.4)(0.8)} \\ &= \frac{2}{3} \end{aligned}$$

- b. The binary classification "effective" and "ineffective" is a bit reductive. Instead the *effectiveness* 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 is given the treatment and it has an effectiveness of 30. What is the new probability that the patient has subtype A?

Answer:

This is your random variable Bayes:

$$\begin{aligned} P(A|E = 30) &= \frac{f(E = 30|A)P(A)}{f(E = 30|A)P(A) + f(E = 30|B)P(B)} \\ &= \frac{\left(\frac{1}{5\sqrt{2\pi}}e^{-\frac{(x-\mu_a)^2}{50}}\right)(0.8)}{\left(\frac{1}{5\sqrt{2\pi}}e^{-\frac{(x-\mu_a)^2}{50}}\right)(0.8) + \left(\frac{1}{4\sqrt{2\pi}}e^{-\frac{(x-\mu_b)^2}{50}}\right)(0.2)} \\ &= \frac{\frac{1}{5}\left(e^{-\frac{(30-60)^2}{50}}\right)(0.8)}{\frac{1}{5}\left(e^{-\frac{(30-60)^2}{50}}\right)(0.8) + \frac{1}{4}\left(e^{-\frac{(30-20)^2}{50}}\right)(0.2)} \end{aligned}$$

6 Simulation accuracy

You are running a simulation to estimate the probability parameter p of an event. Let X be the number of times the event occurs after 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 and the estimate is less than epsilon: $|\frac{X}{n} - p| < \epsilon$.

Calculate the probability that your estimate will be good enough when event probability $p = 0.7$, $\epsilon = 0.01$ and $n = 1000$ trials. Use an approximation to get a numeric answer. You may leave your answer in terms of the Φ function.

Answer:

This is your random variable Bayes:

$$\begin{aligned} P(|\frac{X}{n} - p| < \epsilon) &= P(-\epsilon < \frac{X}{n} - p < \epsilon) \\ &= P(p - \epsilon < \frac{X}{n} < p + \epsilon) \\ &= P(n(p - \epsilon) < X < n(p + \epsilon)) \\ &= P(1000(0.7 - 0.01) < X < 1000(0.7 + 0.01)) \\ &= P(690 < X < 710) \end{aligned}$$

Note that $X \sim \text{Bin}(n = 1000, p = 0.7)$ but if we want a numeric answer we are going to need to approximate. Let Y be a gaussian approximation of X . $Y \sim N(\mu = 700, \sigma^2 = 210)$.

$$\begin{aligned} P(|\frac{X}{n} - p| < \epsilon) &\approx P(689.5 < Y < 710.5) \\ &\approx \Phi(\frac{710.5 - 700}{\sqrt{210}}) - \Phi(\frac{689.5 - 700}{\sqrt{210}}) \\ &\approx 0.53 \end{aligned}$$