

CS109 Midterm Examination

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.

Problem	Score
1 (17 pts)	
2 (16 pts)	
3 (26 pts)	
4 (20 pts)	
5 (16 pts)	
6 (25 pts)	
Total (120 pts)	

THE STANFORD UNIVERSITY HONOR CODE

- A. The Honor Code is an undertaking of the students, individually and collectively:
- (1) that they will not give or receive aid in examinations; that they will not give or receive unpermitted aid in class work, in the preparation of reports, or in any other work that is to be used by the instructor as the basis of grading;
 - (2) that they will do their share and take an active part in seeing to it that others as well as themselves uphold the spirit and letter of the Honor Code.
- B. The faculty on its part manifests its confidence in the honor of its students by refraining from proctoring examinations and from taking unusual and unreasonable precautions to prevent the forms of dishonesty mentioned above. The faculty will also avoid as far as practicable, academic procedures that create temptations to violate the Honor Code.
- C. While the faculty alone has the right and obligation to set academic requirements, the students and faculty will work together to create optimal conditions for honorable academic work.

I acknowledge and accept the letter and spirit of the honor code:


Signature: _____

Name (print): _____

- c. (5 points) What is the probability that a random ordering of the songs has two songs from artist A in a row or two songs from artist B in a row?

- d. (4 points) Spotify tries an algorithm where they play the 10 songs in two rounds:
- In the first round, they play one song from each artist.
 - In the second round, they play the other song from each artist, preserving the **same artist ordering** as in the first round.

Here is an example of one ordering produced by the algorithm. A_2 is the second song by artist A:



How many ways are there to order songs under this new algorithm?

Spotify uses a shuffling algorithm like the one in part (d) for the reason described in (b)

- c) (8 points) To slow down the hacker you make your website wait 2^n seconds before responding to a login request, where n is the number of times a visitor has incorrectly guessed their password. If the hacker has a $1/1000$ chance of correctly guessing a password on each attempt, what is the expected amount of time it will take the hacker to guess a user's password?

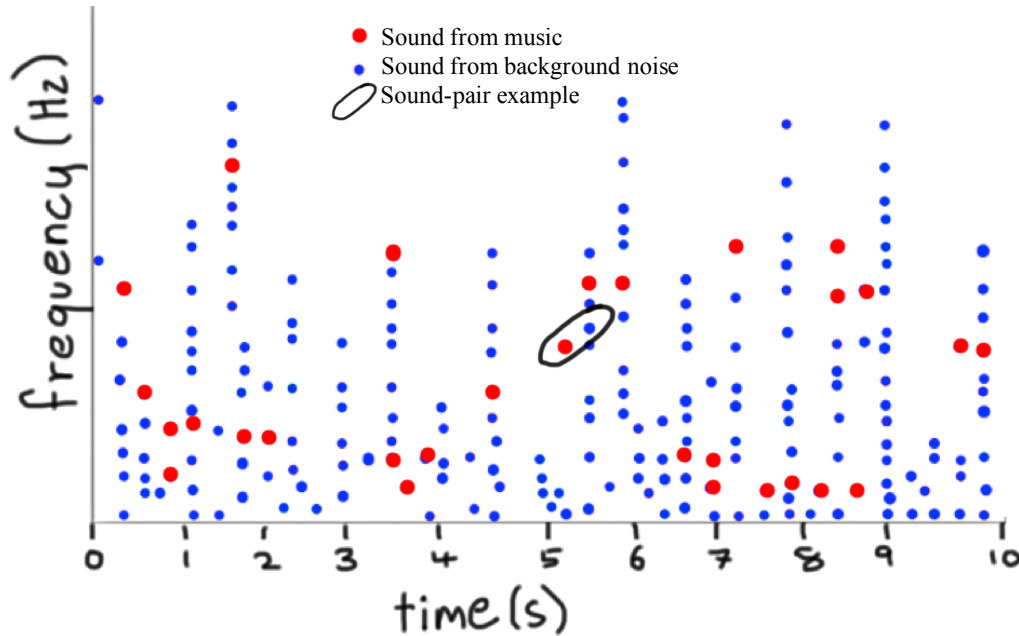
The strategy used in part (d) is a great password protection algorithm.

3. Shazam (26 points)

Shazam is an app that listens to a 10 second sample of background music playing (for example in a restaurant) and guesses the song. A user has just sent in a sample which we will use to explore how Shazam works. Within the 10 second sample there are:

- 50 sounds heard from the background *music* and
- 2000 sounds heard from background *noise*

A “sound” is a frequency heard at a particular time. Here is a visualization of our 10 second sample (a few sounds are omitted for visual clarity):



- a. (3 points) There are 2050 total sounds in the sample. How many distinct *pairs* of two sounds are there? Sounds can not be paired with themselves. Sound-pairs with the same two sounds are *not* distinct.

- b. (4 points) How many distinct sound pairs exist such that *both* sounds in the sound-pair are from the music (as opposed to being from background noise)?

- c. (5 points) Every distinct pair of sounds from the 10 second sample casts a vote as to what song the pair thinks is playing. If both sounds in a sound-pair are from the music, the pair always casts a vote for the correct song. Otherwise, since the pair contains background noise, the pair casts a vote uniformly at random from a set of 5 songs (always the same five songs, including the correct song).

We want **more than 1/5 of the total number of votes** to go to the correct song. How many of the pairs containing background noise must vote for the correct song in order for the correct song to get 1/5 of the votes?

- d. (6 points) Let d be the number of sound-pairs that you calculated in part (c). What is the probability that the correct song receives more than 1/5 of all the votes?

- e. (8 points) You are at the Tree House and hear music playing. You believe that there is an:

80% chance the song is: Hold Up by Beyonce (event X_1)

20% chance the song is: Can't Get Used by Andy Williams (event X_2)

You run "Shazam" and the app returns that it predicts the Andy Williams song is playing. Let q be the probability that the correct song is returned by Shazam. Let $(1 - q)$ be the probability that an incorrect song is returned. What is your new probability for X_1 and X_2 ? Express your answer in terms of q .

This is the basic version of Shazam. Wisdom of the crowds is a phenomenon where a crowd (in this case a crowd of note-pairs) is correct consistently and wrong randomly.

- c. (8 points) The visitor clicks in the box at pixel (3,3) which has a distance of 2 pixels from the center of the box. What is Google's **new** belief that the visitor is a robot?

Recaptcha uses more sophisticated statistics of natural human mouse gestures and clicks, but this problem covers the central idea behind the new click based recaptchas.

5. Exponentials (16 points)

- a. (8 points) What is the probability that an exponential random variable $X \sim \text{Exp}(\lambda)$ takes on a value that is within one standard deviation of its mean?
- b. (8 points) Let $X_1 \sim \text{Exp}(\lambda_1)$ and $X_2 \sim \text{Exp}(\lambda_2)$ be independent exponential random variables. Let $M = \max(X_1, X_2)$ where \max is a function which returns the larger of the two values. Give an expression for the cumulative density function of M .

- c. (7 points) A Poisson random variable with a parameter of $\lambda \geq 1000$, can be approximated by a normal. The approximating normal should match the mean and the variance of the Poisson. Let $X \sim \text{Poi}(1000)$. What is an approximation for $P(990 < X < 1000)$. Give your answer to two decimal places.
- d. (6 points) The probit function, $\Phi^{-1}(x)$ is the inverse of the CDF of a standard normal. It maps from probabilities to the standard normal CDF input that would produce said probability. For example, you can confirm using the Standard Normal Table that $\Phi(0.1) = 0.5398$, so $\Phi^{-1}(0.5398) = 0.1$. Give a closed form expression, using the probit function, for an approximation of K.

Websites solve this problem constantly.

