

Chris Piech
CS 109

Oct 27, 2017

Practice Midterm Examination

Problems by Will Monroe, Mehran Sahami and Chris Piech

This is a closed calculator/computer exam. You are, however, allowed to use notes or a textbook 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): _____

1. [15 points] Chris comes to class with 12 fruits in his cooler as follows:

- 5 Apples (Fruit type A)
- 4 Mandarins (Fruit type B)
- 3 Persimmons (Fruit type C)

Note that all fruit of the same type are *indistinguishable*.

a. In how many distinct ways can Chris distribute the fruit to 12 students, where each student only gets one fruit?

b. Say that out of the 12 students, there is a particular pair of students (call them Larry and Sergey) who are only happy if they both receive the same type of fruit. How many distinct ways can fruit be distributed to the 12 students (where each student only gets one fruit) such that Larry and Sergey are happy?

c. Now say that Chris again starts with the 12 fruit described above, but there are only 10 students in the class. He distributes fruit to all 10 students (where each student only gets one fruit, so 2 fruit remain in Will's bag). In how many distinct ways can fruit be distributed under these conditions?

2. [18 points] The Stanford Competitive Robotics Club is building a robot that is able to throw Marshmallow Peeps.¹ A team's robot gets one throw.

- If it throws a distance **less than 4 meters**, the team loses and gets nothing.
- If it throws a distance between **4 and 7 meters**, the team wins \$5.
- If it throws a distance between **7 and 10 meters**, the team wins \$20.
- If it throws a distance **greater than 10 meters**, the team loses and gets nothing.

Stanford's current robot design throws distances that are roughly normally distributed, with $\mu = 7$ meters and $\sigma^2 = 4$ meters² (so $\sigma = 2$ meters). (There's a very small chance the robot throws the Peep backwards: a negative distance. That still counts as a loss; you don't need to handle this case specially.) On average, how much money does the team expect to win?

Give an answer consisting only of numeric values and arithmetic operations (additions, subtractions, multiplications, divisions, and exponents). Your answer *should not include any variables or function calls*.

¹This was the project for an actual final competition for ME 210, Introduction to Mechatronics, a few years back (the theme of the competition was Angry Birds). As it turns out, Peeps' aerodynamic properties change when they get stale, resulting in many missed targets.

3. [20 points] Spotify notices that users do not listen to songs with equal probability. Instead, the probability that a random play (an instance of a user listening to a song) is of the i -th most popular song is distributed as a Zipf random variable. As we (briefly) covered in class, the Zipf random variable is discrete and has PMF:

$$P(X = i) = \frac{\frac{1}{i}}{\sum_{n=1}^N \frac{1}{n}}$$

It is parameterized by N , the total number of songs. Spotify has $N = 30$ million ($3 \cdot 10^7$) songs.

It may be useful in this problem to know that $\sum_{n=1}^{3 \cdot 10^7} \frac{1}{n} \approx 17.8$.

- a. What is the probability that a random play is of the *10th most popular* song?

- b. If there are 1 billion plays on a given day, what is the probability that the most popular song is listened to more than 100 million times that day? Use an approximation. Remember you can include summations in your answers and use variables to represent intermediate values, as long as you define clearly how to compute each one.

4. [20 points] Four 6-sided dice are rolled. The dice are fair, so each one has equal probability of producing a value in $\{1, 2, 3, 4, 5, 6\}$. Let $X =$ the *minimum* of the four values rolled. (It is fine if more than one of the dice has the minimal value.)

(Note: You can define intermediate variables in your answers as long as you clearly state how to compute their values.)

a. What is $P(X \geq k)$ as a function of k ?

b. What is $E[X]$?

c. Let $T =$ the sum of the values rolled on the four dice. Let $S =$ the sum of the largest *three* values on the four dice. In other words, $S = T - X$. What is $E[S]$?

5. [25 points] Say that two different manufacturers (call them A and B) are equally likely to produce screens for laptops. The lifetimes for the screens (measured in hundreds of hours) manufactured by each company are *independently* distributed as follows:

- Manufacturer A: lifetime of screens are normally distributed: $N(20, 4)$
- Manufacturer B: lifetime of screens are exponentially distributed: $\text{Exp}(1/20)$

Say we bought a laptop, have used it for 18 hundred hours so far, and the screen is still working at this point in time.

a. At this point in time, what is the probability that manufacturer A produced the screen?

b. At this point in time, what is the probability that manufacturer B produced the screen?

6. [20 points] In class, we did a number of birthday-related problems, under the assumption that all birthdays in a year are equally likely. In reality, the distribution of birthdays in a roomful of people can be somewhat uneven. Suppose that a room contains n people born in the same year, and each person independently has probability $p_a \approx 0.003$ of being born on any one of the 260 weekdays of that year, and probability $p_b \approx 0.002$ of being born on any one of the 105 weekend days.² Give an expression for the probability that at least one pair of people in that room share a birthday. (Remember your expression may include sums; you do not need to simplify.)

Make an intuitive argument why this probability should be larger or smaller, compared to the probability we derived for the case when all dates were equally likely.

²Real hospitals do deliver more babies on weekdays, although not nearly at the 3:2 ratio we're using here. According to babycenter.com, induced labor and C-sections are responsible for much (but not all) of the difference.

