

This homework is due November 24, 2014, at 12:00 noon.

1. Section Rollcall!

In your self-grading for this question, give yourself a 10, and write down what you wrote for parts (a) and (b) below as a comment. Please put the answers in your written homework as well.

- (a) What discussion did you attend on Monday last week? If you did not attend section on that day, please tell us why.
- (b) What discussion did you attend on Wednesday last week? If you did not attend section on that day, please tell us why.

2. Hashing & Drunk Man Lab

(Note): This Virtual Lab will be pushed to Homework 13. If you want to get started early on it, that's great! If not, just hand it in with Homework 13.

The Birthday Paradox, which we explored in last week's lab, is a specific form of a more general concept known as hashing. In hashing, we are interested in minimizing the probability that two or more keys are hashed to the same table location. This also can be looked at as a balls and bins problem, where we are trying to minimize the probability of balls being thrown into the same bin.

Let A be the event that there is no collision (i.e. no two balls are thrown into the same bin), and $P(A)$ be the probability that there is no collision. For the questions below, let's consider n bins and m balls (for the Birthday Paradox, $n = 365$).

For each part, students who want to can choose to completely rewrite the question. Basically, you can come up with your own formulation of how to do a series of experiments that result in the same discoveries. Then, write up the results nicely using plots as appropriate to show what you observed. You can also rewrite the entire lab to take a different path through as long as they convey the key insights aimed at in each part.

Please download the IPython starter code from Piazza or the course webpage, and answer the following questions.

- (a) For $n = 100, 365, 500, 1000$, and for $m = 0.1n, 0.3n, \sqrt{n}$, simulate throwing m balls into n bins. For every (m, n) pair, do this for 1000 trials. For each value of n , plot the fraction of trials for which there is no collision vs. m . What do you observe as m gets larger?
- (b) Let's take a detour and look at a very useful approximation for $\ln(1 - x)$. Plot $\ln(1 - x)$ for 1000 values of x from $0.01 - 0.1$. What is the approximate slope of this graph and why? (think about Taylor expansion.)
- (c) Repeat simulating throwing m balls in n bins for $n = 100, 365, 500, 1000$, and 20 values of m between 1 and $2\sqrt{n}$. Now for each n , plot the negative logarithm of the fraction of trials for which there were no collisions, for different values of m . Does this plot look linear in m ? Quadratic? Exponential?

- (d) Now overlay the plots in part (c) with $\frac{m^2}{2n}$. What do you notice?
 From these graphs, and for each value of n , find the approximate m value such that $P(A) \approx 0.5$ for both the simulated data and $\frac{m^2}{2n}$. How much does this differ?
- (e) The drunk man is back, and he wants your help to plot 100 sample paths that he could take from time $t = 0$ to $t = 999$ (1000 timesteps). Assume the same probabilities from Homework 11's Question 5, that is, the man moves forward with probability 0.5, backward with probability 0.3, and stays exactly where he is with probability 0.2. What do you observe about his paths? Where do you think he should end up at after 1000 timesteps, on average?
Hint: Implement the function `drunk_man`, which returns a list of elements that starts at 0, and every element thereafter is one more, one less, or equal to the previous one, with the correct probability for each possibility.
- (f) Redo the previous part, but this time, the man moves forward with probability 0.3, backward with probability 0.5, and stays exactly where he is with probability 0.2. What do you observe about his paths? Where do you think he should end up at after 1000 timesteps, on average?
 Finally, assume the man moves forward and backward with probability 0.5 each, i.e. he never stays still. Plot his 100 sample paths. What do you observe? Does this graph remind you of something we have done in a previous virtual lab?
- (g) One simple way of solving counting problems is to enumerate all possibilities, and then count the ones that we are looking for. In this question, let's check your answer to Question 5 from the last homework. Implement the function `count_paths(t)`, which generates all the paths the Drunk Man can take in t timesteps, then counts the number of paths in which he returns to 0 at time t and it is his first return. Remember that we no longer care about probabilities when counting paths.
Hint: One way to generate all possible paths is to use `itertools.product`. You might also want to implement the function `catalan(n)`, which computes the n^{th} Catalan number.
- (h) The Bernoulli distribution is the probability distribution of a random variable which takes value 1 with success probability p and value 0 with failure probability $1 - p$. It can be used, for example, to represent a coin toss, where 1 is defined to mean "heads" and 0 is defined to mean "tails".
 In other words:

$$f(k; p) = \begin{cases} p & \text{if } k = 1 \\ 1 - p & \text{if } k = 0 \end{cases}$$

Related values:

$$E[X] = p$$

$$\text{Var}(X) = p(1 - p)$$

Plot the probability mass functions (pmf) for Bernoulli random variables with success probabilities of 0.1, 0.3, 0.5, 0.7, and 0.9, respectively.

Hint: `scipy.stats.bernoulli.pmf` will be useful for this question.

- (i) The binomial distribution gives us the probability of observing k successes, each with a probability p , out of N attempts.

$$f(k; N; p) = \binom{N}{k} p^k (1 - p)^{N-k}$$

where

$$\binom{N}{k} = \frac{N!}{k!(N-k)!}$$

with $k = 0, 1, 2, \dots, N$

Related values:

$$E[X] = Np$$

$$\text{Var}(X) = Np(1-p)$$

If I toss a 0.6-biased coin 20 times, what is the approximate probability of getting exactly 15 heads? Answer this question graphically by plotting the PMF of a binomial random variable with parameters (20, 0.6).

Looking at the plot, at which value of n is the probability of getting n heads maximized?

Hint: As before, you will find `scipy.stats.binom.pmf` useful.

(j) Question 9, part (e)

(k) Question 9, part (f)

Reminder: When you finish, don't forget to convert the notebook to pdf and merge it with your written homework. Please also zip the `ipynb` file and submit it as `hw12.zip`.

3. Fundamentals

True or false? For the following statements, provide either a proof or a simple counterexample. Let X, Y, Z be arbitrary random variables.

- (a) If (X, Y) are independent and (Y, Z) are independent, then (X, Z) are independent.
- (b) If (X, Y) are dependent and (Y, Z) are dependent, then (X, Z) are dependent.
- (c) Assume X is discrete. If $\text{Var}(X) = 0$, then X is a constant.
- (d) $E[X]^4 \leq E[X^4]$

4. Markov Conversation

Alice is hosting a party. As she's talking to her guests, she notices that conversations naturally transition between casual and more interesting topics. Consider the following simple model of conversations: Each type of topic takes a certain amount of time, and can transition to different topics as specified.


1. Casual Topics: These take 5 minutes. At the end, they can transition into an Interesting topic (w.p. 25%), or the conversation can terminate (w.p. 75%).
2. Interesting Topics: These take 16 minutes. At the end, they are always followed by a Casual Topic.

The following diagram illustrates the conversation flow, where " $\emptyset$ " means the conversation has terminated, and "1", "2" correspond to the conversation topics.


Using the above model:

- (a) What is the expected length of a conversation that starts on a Casual topic?
(*Hint: Let the random variable X be the length of a conversation that starts on a Casual topic. Try to express $E[X]$ in terms of conditional expectations...*)

As the party progresses, Alice revises her model of conversation to include three types of topics:



1. Casual Topics: These take 5 minutes. At the end, they can transition into an Interesting topic (w.p. 25%), or the conversation can terminate (w.p. 75%).
2. Interesting Topics: These take 15 minutes. At the end, they can be followed by a Deep topic (w.p. 50%), or can go back to a Casual topic (w.p. 50%).
3. Deep Topics: These take 25 minutes. They are always followed by an Interesting topic.



Alice starts to wonder how the expected length of her conversations depend on who she talks to. Assume the following model: If she talks to acquaintances, they start on a Casual topic. With close friends, they start on an Interesting topic (w.p. 50%), or on a Deep topic (w.p. 50%). Using the revised model:

- (b) Alice starts talking to her acquaintance Bob. What is the expected length of their conversation?
Hint: Let X_1, X_2, X_3 be the length of her conversation if she starts on topic 1, 2, 3 respectively. Try to relate the expected values of the three R.Vs.
- (c) Alice starts talking to her close friend Charlie. What is the expected length of their conversation?
- (d) Assume people at the party are equally likely to be close friends or acquaintances. But Eve noticed that Alice and Dave talked for 45 minutes! What is the probability that Dave is a close friend of Alice?

5. Round the Clock

As your new year's resolution you decided to have three meals a day, with no exceptions. However you do not like putting names on things, and so you won't be calling your meals breakfast, lunch, and/or dinner. You simply love them equally and hate labeling them. Furthermore, towards your goal of being well-organized you decide that you will have each meal at the same time everyday.

For each one of your meals, you decide that you would have it at an exact but completely random hour during the whole day. i.e. for each meal you pick a time from the set of 24 exact hours $\{12\text{am}, 1\text{am}, \dots, 10\text{pm}, 11\text{pm}\}$ and then you have that meal at that exact time every day. Note that you might even have two meals at the same time (i.e. you sample with replacement)! But if you schedule two or more meals at the same time, you decide on a completely random ordering over the meals scheduled for that time, and then you'll have your meals according to that ordering every day. Furthermore, for simplicity assume that it takes no time for you to finish a meal.

- (a) For each one of your meals, what is the expected amount of time you have to wait after finishing it, until you get to have your next meal?

- (b) Now suppose that your friend calls you at an exact random hour during the day (24 different possibilities). S/he asks you to go to his/her place and have your next meal with him/her. You do not want to ruin your schedule, so you tell your friend that you'll be at his/her place in X hours to have your next meal, where $X \geq 0$ is the number of hours until your next meal is scheduled. **What is the expected value of X ?** (If a phone call falls on the exact hour of a meal, break ties randomly as in part (a)).

6. Linearity

Solve each of the following problems using linearity of expectation. Explain your methods clearly.

- (a) In an arcade, you play game A 10 times and game B 20 times. Each time you play game A, you win with probability $1/3$ (independently of the other times), and if you win you get 3 tickets (redeemable for prizes), and if you lose you get 0 tickets. Game B is similar, but you win with probability $1/5$, and if you win you get 4 tickets. What is the expected total number of tickets you receive?
- (b) A monkey types at a 26-letter keyboard with one key corresponding to each of the lower-case English letters. Each keystroke is chosen independently and uniformly at random from the 26 possibilities. If the monkey types 1 million letters, what is the expected number of times the sequence "book" appears?
- (c) A building has n floors numbered $1, 2, \dots, n$, plus a ground floor G. At the ground floor, m people get on the elevator together, and each gets off at a uniformly random one of the n floors (independently of everybody else). What is the expected number of floors the elevator stops at (not counting the ground floor)?
- (d) A coin with Heads probability p is flipped n times. A "run" is a maximal sequence of consecutive flips that are all the same. (Thus, for example, the sequence HTHHHTTH with $n = 8$ has five runs.) Show that the expected number of runs is $1 + 2(n-1)p(1-p)$. Justify your calculation carefully.

7. Company Selection

Company A produces a particular device consisting of 10 components. Company A can either buy all the components from Company S or Company T, and then uses them to produce the devices without testing every individual component. After that, each device will be tested before leaving the factory. The device works only if every component works properly. Each working device can be sold for x dollars, but each non-working device must be thrown away. Products from Company S have a failure probability of $q = 0.01$ while Company T has a failure probability of $q/2$. However, every component from Company S costs \$10 while it costs \$30 from Company T. Should Company A build the device with components from Company S or Company T in order to maximize its expected profit per device? (Your answer will depend on x .)

8. (Extra Credit) Would You Bet It?

In addition to the oral exam, suppose EECS70 offers an extra credit challenge as follows.

- Bet: 4 homework points.
- Rules:
 - Initially, there are 52 cards with 4 suits ($\spadesuit, \heartsuit, \diamondsuit, \clubsuit$) and 13 values (A, 2, 3, 4, 5, 6, 7, 8, 9, 10, J, Q, K).
 - In each round, you draw a card and put it in front of you. If two cards in front of you have the same value, the game is over; otherwise, you survive this round and earn 1 homework point.
 - All cards you drew stay in front of you until the game is over.

Would it be wise to play the game? You might want to use a spreadsheet program to help your calculation.


9. Hosting a party

You are hosting a party, and as the host it is rude to talk to only some of your guests. In reality you should socialize with all your guests at least once. For the parts that are marked with [VL], please complete your answer in the Virtual Lab skeleton.

- Your first strategy is to continually pick a person at random and go chat with him/her. Assume that in the party, there are n guests plus you. If you do the chatting strategy m times ($m \leq n$), what is the probability that you talk with everyone at most once?
- Now suppose that instead of randomly picking a single guest, you instead pick two people at random and chat with the one you've chatted with the least number of times (it does not matter how you break the ties). Now what is the chance that after doing this m times ($m \leq n$), you haven't talked to anyone more than once? (Assume $n \geq 2$.)
- Suppose you use the first strategy. What is the expected number of times you should chat with people before having chatted with every guest at least once?
- Suppose you use the second strategy. What is the expected number of times you should chat with people before having chatted with every guest at least once? (Assume $n \geq 2$.)
- (*Extra Credit*) [VL] Let $n = 100$. Plot the probabilities in (a) and (b) by varying m from 1 to n .
- (*Extra Credit*) [VL] Plot the expected numbers of times in (c) and (d) by varying n from 2 to 100.

10. (*Extra Credit*) Fill in the Blanks!

This is a flowchart for counting how many ways to put k objects into n bins. Fill in the blanks with the correct formula. Remember, each box corresponds to one subproblem.



11. Write your own problem

Write your own problem related to this week's material and solve it. You may still work in groups to brainstorm problems, but each student should submit a unique problem. What is the problem? How to formulate it? How to solve it? What is the solution?