

Sundry

Before you start your homework, write down your team. Who else did you work with on this homework? List names and email addresses. (In case of homework party, you can also just describe the group.) How did you work on this homework? Working in groups of 3-5 will earn credit for your "Sundry" grade.

Please copy the following statement and sign next to it:

I certify that all solutions are entirely in my words and that I have not looked at another student's solutions. I have credited all external sources in this write up.

1 Law of Large Numbers

Recall that the *Law of Large Numbers* holds if, for every $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\left| \frac{1}{n} S_n - \mathbb{E} \left[\frac{1}{n} S_n \right] \right| > \varepsilon \right) = 0.$$

In class, we saw that the Law of Large Numbers holds for $S_n = X_1 + \dots + X_n$, where the X_i 's are i.i.d. random variables. This problem explores if the Law of Large Numbers holds under other circumstances.

Packets are sent from a source to a destination node over the Internet. Each packet is sent on a certain route, and the routes are disjoint. Each route has a failure probability of p and different routes fail independently. If a route fails, all packets sent along that route are lost. You can assume that the routing protocol has no knowledge of which route fails.

For each of the following routing protocols, determine whether the Law of Large Numbers holds when S_n is defined as the total number of received packets out of n packets sent. Answer **Yes** if the Law of Large Number holds, or **No** if not, and give a brief justification of your answer. (Whenever convenient, you can assume that n is even.)

1. **Yes** or **No**: Each packet is sent on a completely different route.

2. **Yes or No:** The packets are split into $n/2$ pairs of packets. Each pair is sent together on its own route (i.e., different pairs are sent on different routes).
3. **Yes or No:** The packets are split into 2 groups of $n/2$ packets. All the packets in each group are sent on the same route, and the two groups are sent on different routes.
4. **Yes or No:** All the packets are sent on one route.

2 Guessing a Fair Coin Flip

The purpose of this question is to provide further evidence of the vague idea that the entropy measures the amount of “information” contained in a random variable.

- (a) Let $X \sim \text{Geometric}(1/2)$. Compute $H(X)$.
- (b) Now, you want to guess the value of X . Consider the following strategy: First, you ask if $X = 1$. If the answer is “no”, then you ask if $X = 2$. If the answer is “no” again, then you ask if $X = 3$, and so on. In other words, for each $k \in \mathbb{Z}_+$, you successively ask if $X = k$ until you receive the answer “yes”. What is the expected number of questions you must ask before you find out the value of X ? Show that it equals the entropy.

3 Trust No One

Your friend has a coin that they insist is fair. You ask them to toss the coin n times in front of you so you can trust what they say. You’ll consider the coin biased if the coin comes up heads or tails more than 90% of the time.

- (a) Using Chebyshev’s Inequality, how many samples do you need to collect before you can say the coin is fair with 95% confidence?
- (b) What if you use the Chernoff bound?
- (c) For this part, you will go through a proof of a special case of Chernoff bound.

Let Y_i be random variables bounded by the interval $[-1, 1]$ with 0 mean, and let $Y = \sum_{i=1}^n Y_i$.

- (i) We’ll begin similarly to the proof of Chernoff bound. Let $\hat{\mu} = \mu + \varepsilon$, for $\varepsilon > 0$. Show that the following inequality holds:

$$\mathbb{P}\left(\frac{Y}{n} \geq \hat{\mu}\right) \leq \frac{\mathbb{E}[e^{\lambda Y_1}]^n}{e^{\lambda \varepsilon n}}$$

- (ii) Now, we’ll use the convexity of $e^{\lambda Y_i}$, for $\lambda > 0$, which implies for a random variable bounded by the interval $[a, b]$:

$$e^{\lambda Y_i} \leq \frac{b - Y_i}{b - a} e^{\lambda a} + \frac{Y_i - a}{b - a} e^{\lambda b}$$

Now using this inequality, show that

$$\mathbb{E}[e^{\lambda Y_i}] \leq \frac{e^\lambda + e^{-\lambda}}{2}.$$

- (iii) Now, we will extend this inequality further. Remember that the Taylor series expansion of e^x is $\sum_{k=0}^{\infty} x^k/k!$. Using this, show that

$$\mathbb{E}[e^{\lambda Y_i}] \leq \exp\left(\frac{\lambda^2}{2}\right).$$

- (iv) Now you can plug this inequality into the expression bounding $\mathbb{P}(Y/n \geq \hat{\mu})$. After this, show that the tightest bound you can get is:

$$\mathbb{P}\left(\frac{Y}{n} \geq \hat{\mu}\right) \leq \exp\left(-n\frac{\epsilon^2}{2}\right)$$

- (d) Like the general case for Chernoff bound, this bound also holds on both sides. How many samples do you need if you use the bound from the previous part? [*Hint*: Use the random variable $Y_i = 2X_i - 1$.]

4 Pokemon Craze

You and your friend are both trying to catch a Dratini. Unfortunately, you each can only attempt to catch one Dratini per day. Once you or your friend catch a Dratini, that person stops while the other person continues to try to catch a Dratini if they haven't already caught one. The probability an attempt at catching a Dratini is successful is p . What is the expected number of days this process takes, if you two both try to catch a Dratini every day? Solve this using a Markov chain with three states. (What is this in terms of two geometric random variables?)

(Also, for fun, consider how this problem would change if instead of stopping, the first person kept on trying to catch a Dratini in case he could donate it to the other person.)

5 Markov's Coupon Collecting

Courtney is home for Thanksgiving and needs to make some trips to the Traitor Goes grocery store to prepare for the big turkey feast. Each time she goes to the store before the holiday, she receives one of the n different coupons that are being given away. You may recall that we studied how to calculate the expected number of trips to the store needed to collect at least one of each coupon. Using geometric distributions and indicator variables, we determined that expected number of trips to be $n(\ln n + \gamma)$.

Let's re-derive that, this time with a Markov chain model and first-step equations.

- (a) Define the states and transition probabilities for each state (explain what states can be transitioned to, and what probabilities those transitions occur with).

- (b) Now set up first-step equations and solve for the expected number of grocery store trips Courtney needs to make before Thanksgiving so that she can have at least one of each of the n distinct coupons.

6 To Infinity (But Not Beyond)

[In this problem, we consider a very simple infinite Markov chain. This problem is intended to show you that even when we're in a situation where the theorems from lecture and notes don't necessarily apply, we can still use our intuition to find the stationary distribution, though we still need to give a formal proof that our intuition led us to the correct answer.]

Consider a Markov chain whose state space is \mathbb{N} with the following transition probabilities:

$$\begin{aligned} P(0,0) &= \frac{3}{4}, \\ P(0,1) &= \frac{1}{4}, \\ P(i,i-1) &= \frac{1}{2}, & \forall i > 0, \\ P(i,i) &= P(i,i+1) = \frac{1}{4}, & \forall i > 0, \\ P(i,j) &= 0, & \text{otherwise.} \end{aligned}$$

- (a) For the moment, let's ignore the constraint that $\sum_{i \in \mathbb{N}} \pi(i) = 1$ and just assume that $\pi(0) = 1$. What value must $\pi(1)$ take on to ensure that $\pi(0)$ does not change when we move forward one time step?
- (b) If we assume that $\pi(0) = 1$ and $\pi(1)$ takes on the value you found in the last part, what value must $\pi(2)$ take on to ensure that $\pi(1)$ remains unchanged after one time step? And if $\pi(2)$ takes on that value, what must $\pi(3)$ be?
- (c) What pattern do you see emerging? If π followed this pattern and $\pi(0) = 1$, what would π be?
- (d) Now we bring back the constraint that $\sum_{i \in \mathbb{N}} \pi(i) = 1$. What can π be such that it fulfills this constraint and still follows the pattern?
- (e) Prove that the π from the previous part is the stationary distribution for the Markov chain.

7 Allen's Umbrellas

Every morning, Allen walks from his home to Soda, and every evening, Allen walks from Soda to his home. Suppose that Allen has two umbrellas in his possession, but he sometimes leaves his umbrellas behind. Specifically, before leaving from his home or Soda, he checks the weather. If it is raining outside, he will bring his umbrella (that is, if there is an umbrella where he currently is). If it is not raining outside, he will forget to bring his umbrella. Assume that the probability of rain is p .

We will model this as a Markov chain. Let $\mathcal{X} = \{0, 1, 2\}$ be the set of states, where the state i represents the number of umbrellas in his current location. Write down the transition matrix, determine if the distribution of X_n converges to the invariant distribution, and compute the invariant distribution. Determine the long-term fraction of time that Allen will walk through rain with no umbrella.