CS 70 Spring 2021 — Discrete Mathematics and Probability Theory — HW 12

Due: Sunday, 04/25 at 10:00 PM
Grace period until Tuesday, 04/27 at 11:59 PM

Sundry

Before you start writing your final homework submission, state briefly how you worked on it. Who else did you work with? List names and email addresses. (In case of homework party, you can just describe the group.)

1 Gambling Woes

Forest proposes a gambling game to you (uh oh!). Every day, you flip two independent fair coins. If both of the coins come up heads, then your fortune triples on that day. If one coin comes up heads and the other coin comes up tails, then your fortune is cut in half. If both of the coins comes up tails, then game over: you lose all of your money! Forest claims that you can get rich quickly with this scheme, but you decide to calculate some probabilities first.

(a) Let M_0 denote your money at the start of the game, and let M_n denote the amount of money you have at the end of the nth day. Compute $\mathbb{E}[M_{n+1} \mid M_n]$.

(b) Use the law of iterated expectation to calculate $\mathbb{E}[M_{n+1}]$ in terms of $\mathbb{E}[M_n]$. Solve your recurrence to obtain an expression for $\mathbb{E}[M_{n+1}]$. Do you think this is a fair game?

(c) Calculate $\mathbb{P}(M_n > 0)$. What is the behavior as $n \to \infty$? Would you still play this game?

2 Iterated Expectation

In this question, we will try to achieve more familiarity with the law of iterated expectation.

1. You lost your phone charger! It will take D days for the new phone charger you ordered to arrive at your house (here, D is a random variable). Suppose that on day i, the amount of battery you lose is B_i, where $\mathbb{E}[B_i] = \beta$. Let $B = \sum_{i=1}^{D} B_i$ be the total amount of battery drained between now and when your new phone charger arrives. Apply the law of iterated expectation to show that $\mathbb{E}[B] = \beta\,\mathbb{E}[D]$. (Here, the law of iterated expectation has a very clear interpretation: the amount of battery you expect to drain is the average number of days it takes for your phone charger to arrive, multiplied by the average amount of battery drained per day.)

2. Consider now the setting of independent Bernoulli trials, each with probability of success p. Let S_i be the number of successes in the first i trials. Compute $\mathbb{E}[S_m \mid S_n]$. (You will need to consider three cases based on whether $m > n$, $m = n$, or $m < n$. Try using your intuition rather than proceeding by calculations.)

3 Strange Dilution

You have a jar of red marbles and blue marbles. At each time step, you draw a marble, and you note the color of the marble. Then, you dilute the proportion of the opposite-colored marbles by a factor of γ, where $0 < \gamma < 1$. (For example: if you pick a red marble, then the proportion of blue marbles is reduced by a factor of γ.) If p is the fraction of marbles that started off as red, what is the expected proportion of red marbles at time n?

4 Oski's Markov Chain

When Oski Bear is studying for CS70, he splits up his time between reading notes and working on practice problems. To do this, every so often he will make a decision about what kind of work to do next.

When Oski is already reading the notes, with probability a he will decide to switch gears and work on a practice problem, and otherwise, he will decide to keep reading more notes. Conversely, when Oski is already working on a practice problem, with probability b he will think of a topic he needs to review, and will decide to switch back over to the notes; otherwise, he will keep working on practice problems.

Assume that (unlike real life, we hope!) Oski never runs out of work to do.

(a) Draw a 2-state Markov chain to model this situation.

(b) In the remainder of this problem, we will learn to work with the definitions of some important terms relating to Markov Chains. These definitions are as follows:

  (a) (Irreducibility) A Markov chain is irreducible if, starting from any state i, the chain can transition to any other state j, possibly in multiple steps.

  (b) (Periodicity) $d(i) := \gcd\{n > 0 \mid P^n(i,i) = \mathbb{P}[X_n = i \mid X_0 = i] > 0\}$, $i \in \mathscr{X}$. If $d(i) = 1\ \forall i \in \mathscr{X}$, then the Markov chain is aperiodic; otherwise it is periodic.

  (c) (Matrix Representation) Define the transition probability matrix P by filling entry (i, j) with probability $P(i, j)$.

  (d) (Invariance) A distribution π is invariant for the transition probability matrix P if it satisfies the following balance equations: $\pi = \pi P$.

  For what values of a and b is the Markov chain irreducible?

(c) For $a = 1$, $b = 1$, prove that the Markov chain is periodic.

(d) For $0 < a < 1$, $0 < b < 1$, prove that the Markov chain is aperiodic.

(e) Construct a transition probability matrix using the Markov chain.

(f) Write down the balance equations for the Markov chain and solve them. Assume that the Markov chain is irreducible.

5 Markov Chains: Prove/Disprove

Prove or disprove the following statements, using the definitions from the previous question.

(a) There exists an irreducible, finite Markov chain for which there exist initial distributions that converge to different distributions.

(b) There exists an irreducible, aperiodic, finite Markov chain for which $\mathbb{P}(X_{n+1} = j \mid X_n = i) = 1$ or 0 for all i, j.

(c) There exists an irreducible, non-aperiodic Markov chain for which $\mathbb{P}(X_{n+1} = j \mid X_n = i) \neq 1$ for all i, j.

(d) For an irreducible, non-aperiodic Markov chain, any initial distribution not equal to the invariant distribution does not converge to any distribution.

6 Markov Property Practice

Let $X_0, X_1, ...$ be a Markov chain with state space S, such that i_j is the value that X_j takes in the j^{th} state. One of the properties that it satisfies is the Markov property:

$$\mathbb{P}(X_n = i_n | X_{n-1} = i_{n-1}, \ldots, X_0 = i_0) = \mathbb{P}(X_n = i_n | X_{n-1} = i_{n-1}), \text{ for all } i_0, i_1, \ldots, i_n \in S, \ n \in \mathbb{Z}_{>0}.$$

Use the Markov property and the total probability theorem to prove the following.

(a) $\mathbb{P}(X_3 = i_3 | X_2 = i_2, X_1 = i_1) = \mathbb{P}(X_3 = i_3 | X_2 = i_2)$, for all $i_1, i_2, i_3 \in S$.

Note: This is not exactly the Markov property because it does not condition on X_0.

(b) $\mathbb{P}(X_3 = i_3 | X_1 = i_1, X_0 = i_0) = \mathbb{P}(X_3 = i_3 | X_1 = i_1)$, for all $i_0, i_1, i_3 \in S$.

(c) $\mathbb{P}(X_1 = i_1 | X_2 = i_2, X_3 = i_3) = \mathbb{P}(X_1 = i_1 | X_2 = i_2)$, for all $i_1, i_2, i_3 \in S$.

7 Knight on a Chessboard (Optional)

This problem is optional and will not be required as part of your submission for this homework assignment. You will receive no additional credit for completing this problem.

(a) An irreducible Markov chain is said to be reversible if there exists a probability distribution π such that

$$\pi(i)P(i,j) = \pi(j)P(j,i) \qquad \forall i,j \in \mathscr{X}. \tag{1}$$

Show that if the chain is reversible, then the distribution π which satisfies (1) is the stationary distribution.

(b) Consider a random walk on a finite undirected connected graph: starting from any vertex, the walk transitions to any of the neighboring vertices with equal probability. The state space $\mathscr{X}$ is the set of vertices of the graph. Show that $\pi(v) = \deg(v)/\sum_{v' \in \mathscr{X}} \deg(v')$, for $v \in \mathscr{X}$, is the stationary distribution.

(c) Let $T_i = \min\{n > 0 : X_n = i\}$ be the first time that the chain reaches state i. Use the result $\mathbb{E}[T_i \mid X_0 = i] = 1/\pi(i)$ to answer the following question:

> A knight begins at a corner of an 8×8 chessboard. At each step, it chooses one of its legal moves uniformly at random and moves to a new square. What is the expected number of steps before the knight returns to the same corner from which it started?

[*Hint*: Think about how this question can be formulated as a random walk on a graph.]

8 Boba in a Straw

Imagine that Jonathan is drinking milk tea and he has a very short straw: it has enough room to fit two boba (see figure).

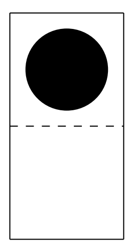

Figure 1: A straw with one boba currently inside. The straw only has enough room to fit two boba.

Here is a formal description of the drinking process: We model the straw as having two "components" (the top component and the bottom component). At any given time, a component can contain nothing, or one boba. As Jonathan drinks from the straw, the following happens every second:

1. The contents of the top component enter Jonathan's mouth.
2. The contents of the bottom component move to the top component.
3. With probability p, a new boba enters the bottom component; otherwise the bottom component is now empty.

Help Jonathan evaluate the consequences of his incessant drinking!

(a) At the very start, the straw starts off completely empty. What is the expected number of seconds that elapse before the straw is completely filled with boba for the first time? [Write down the equations; you do not have to solve them.]

(b) Consider a slight variant of the previous part: now the straw is narrower at the bottom than at the top. This affects the drinking speed: if either (i) a new boba is about to enter the bottom component or (ii) a boba from the bottom component is about to move to the top component, then the action takes two seconds. If both (i) and (ii) are about to happen, then the action takes three seconds. Otherwise, the action takes one second. Under these conditions, answer the previous part again. [Write down the equations; you do not have to solve them.]

(c) Jonathan was annoyed by the straw so he bought a fresh new straw (the straw is no longer narrow at the bottom). What is the long-run average rate of Jonathan's calorie consumption? (Each boba is roughly 10 calories.)

(d) What is the long-run average number of boba which can be found inside the straw? [Maybe you should first think about the long-run distribution of the number of boba.]