CS 70 Spring 2024 | Discrete Mathematics and Probability Theory Seshia, Sinclair | HW 14

1 Rahil's Dilemma

Note 22

Youngmin and Rahil decided to play a game: A fair coin is flipped until either the last two flips were all heads - then Youngmin wins, or the last three flips were all tails - then Rahil wins. Compute the probability that Rahil wins.

Solution: The corresponding Markov chain is: states are $\mathscr{X} = \{\emptyset, H, HH, T, TT, TTT\}$ and the transition probability matrix is

$$\begin{bmatrix} 0 & 1/2 & 0 & 1/2 & 0 & 0 \\ 0 & 0 & 1/2 & 1/2 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1/2 & 0 & 0 & 1/2 & 0 \\ 0 & 1/2 & 0 & 0 & 0 & 1/2 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Let $\alpha(i)$ denote the probability of Rahil winning, that is, reaching state TTT before HH. Then, the first-step equations for α are

$$\begin{aligned} \alpha(TTT) &= 1 \\ \alpha(HH) &= 0 \\ \alpha(TT) &= \frac{1}{2}(\alpha(TTT) + \alpha(H)) = \frac{1}{2}(1 + \alpha(H)) \\ \alpha(T) &= \frac{1}{2}(\alpha(TT) + \alpha(H)) \\ \alpha(H) &= \frac{1}{2}(\alpha(T) + \alpha(HH)) = \frac{1}{2}\alpha(T) \\ \alpha(\emptyset) &= \frac{1}{2}(\alpha(T) + \alpha(H)) \end{aligned}$$

Solving these equations, we get

$$\alpha(TT) = \frac{3}{5}$$
$$\alpha(T) = \frac{2}{5}$$
$$\alpha(H) = \frac{1}{5}$$
$$\alpha(\emptyset) = \frac{3}{10}$$

Hence, Rahil wins with probability $\frac{3}{10}$.

2 A Bit of Everything

Suppose that $X_0, X_1, \ldots$ is a Markov chain with finite state space $S = \{1, 2, \ldots, n\}$, where $n > 2$, and transition matrix P. Suppose further that

$$P(1,i) = \frac{1}{n} \quad \text{for all states } i \text{ and}$$
$$P(j, j-1) = 1 \quad \text{for all states } j \neq 1,$$

with $P(i, j) = 0$ everywhere else.

(a) Prove that this Markov chain is irreducible and aperiodic.

(b) Suppose you start at state 1. What is the distribution of T, where T is the number of transitions until you leave state 1 for the first time?

(c) Again starting from state 1, what is the expected number of transitions until you reach state n for the first time?

(d) Again starting from state 1, what is the probability you reach state n before you reach state 2?

(e) Compute the stationary distribution of this Markov chain.

Solution:

(a) For any two states i and j, we can consider the path $(i, i-1, \ldots, 2, 1, j)$, which has nonzero probability of occurring. Thus, this chain is irreducible. To see that it is aperiodic, observe that $d(1) = 1$, as we have self-loop from state 1 to itself.

(b) At any given transition, we leave state 1 with probability with probability $\frac{n-1}{n}$, independently of any previous transition. Thus, the distribution is Geometric, with parameter $\frac{n-1}{n}$.

(c) Suppose that $\beta(i)$ is the expected number of transitions necessary to reach state n for the first time, starting from state i. We have the following first step equations:

$$\begin{aligned}\beta(1) &= 1 + \sum_{j=1}^{n} \frac{1}{n}\beta(j),\\ \beta(i) &= 1 + \beta(i-1) \quad \text{for } 1 < i < n, \text{ and}\\ \beta(n) &= 0.\end{aligned}$$

We can simplify the second recurrence to

$$\beta(i) = i - 1 + \beta(1) \quad \text{for } 1 < i < n.$$

Substituting this simplified recurrence into the first equation, we get that

$$\beta(1) = 1 + \frac{1}{n}\sum_{i=1}^{n-1}(i-1+\beta(1)) = 1 + \frac{1}{n}\sum_{i=1}^{n-1}(i-1) + \frac{1}{n}\sum_{i=1}^{n-1}\beta(1) = 1 + \frac{(n-2)(n-1)}{2n} + \frac{n-1}{n}\beta(1),$$

which we can solve to get that

$$\beta(1) = \boxed{n + \frac{1}{2}(n-1)(n-2)}.$$

(d) Suppose that $\alpha(i)$ is the probability that we reach state n before we reach state 2, starting from state i. One immediate observation we can make is that from any state i in $\{2, \ldots, n-1\}$, we are guaranteed to see state 2 before state n, as we can only take the path $(i, i-1, \ldots, 2, 1)$. Hence, $\alpha(i) = 0$ if $i \in \{2, \ldots, n-1\}$. Moreover, $\alpha(n) = 1$, so

$$\alpha(1) = \sum_{i=1}^{n} \frac{1}{n}\alpha(i) = \frac{1}{n}\alpha(1) + \frac{1}{n},$$

hence $\alpha(1) = \boxed{\frac{1}{n-1}}$.

(e) We have the balance equations

$$\begin{aligned}\pi(i) &= \frac{1}{n}\pi(1) + \pi(i+1) \quad \text{if } i \neq n, \text{ and}\\ \pi(n) &= \frac{1}{n}\pi(1).\end{aligned}$$

We can collapse the first recurrence to

$$\pi(i) = \frac{n-i}{n}\pi(1) + \pi(n) = \frac{n-i+1}{n}\pi(1),$$

so we can express each stationary probability in terms of the stationary probability of state 1. We can finish by using the normalization equation:

$$\pi(1) + \pi(2) + \cdots + \pi(n) = 1 \implies \frac{1}{n}\pi(1)\sum_{i=1}^{n} n - i + 1 = 1.$$

The last sum can be rearranged to be the sum of the integers from 1 up to n, so we get that

$$\pi(1) = \frac{2}{n+1} \implies \pi = \boxed{\frac{2}{n(n+1)} \begin{bmatrix} n & n-1 & \cdots & 1 \end{bmatrix}}.$$

3 Playing Blackjack

Note 22 Suppose you start with \$1, and at each turn, you win \$1 with probability p, or lose \$1 with probability $1-p$. You will continually play games of Blackjack until you either lose all your money, or you have a total of n dollars.

(a) Formulate this problem as a Markov chain.

(b) Let $\alpha(i)$ denote the probability that you end the game with n dollars, given that you started with i dollars.

Notice that for $0 < i < n$, we can write $\alpha(i+1) - \alpha(i) = k(\alpha(i) - \alpha(i-1))$. Find k.

(c) Using part (b), find $\alpha(i)$, where $0 \le i \le n$. (You will need to split into two cases: $p = \frac{1}{2}$ or $p \neq \frac{1}{2}$.)

Hint: Try to apply part (b) iteratively, and look at a telescoping sum to write $\alpha(i)$ in terms of $\alpha(1)$. The formula for the sum of a finite geometric series may be helpful when looking at the case where $p \neq \frac{1}{2}$:

$$\sum_{k=0}^{m} a^k = \frac{1-a^{m+1}}{1-a}.$$

Lastly, it may help to use the value of $\alpha(n)$ to find $\alpha(1)$ for the last few steps of the calculation.

(d) As $n \to \infty$, what happens to the probability of ending the game with n dollars, given that you start with i dollars, with the following values of p?

(i) $p > \frac{1}{2}$
(ii) $p = \frac{1}{2}$
(iii) $p < \frac{1}{2}$

Solution:

(a) We have the following state transition diagram:

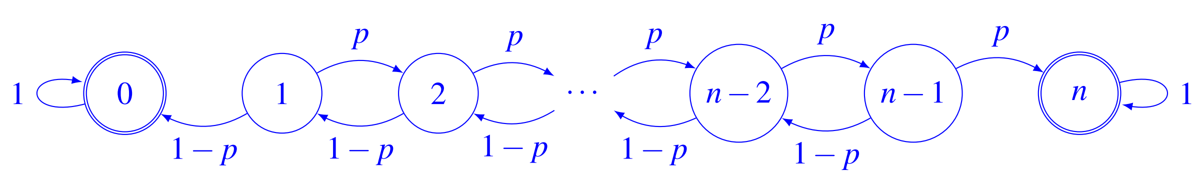

In particular, we have $n+1$ states, $\{0,1,2,\ldots,n\}$, where the transition probability from i to $i+1$ is p, and the transition probability from i to $i-1$ is $1-p$. The transition probabilities for $i=0$ and $i=n$ are edge cases, where we stay in place with probability 1.

(b) If we start with i dollars, this means that we start at state i. The next transition can either be to state $i+1$ with probability p, or to state $i-1$ with probability $1-p$. This means that we have

$$\alpha(i) = p\alpha(i+1) + (1-p)\alpha(i-1).$$

Here, a trick is to expand $\alpha(i) = p\alpha(i) + (1-p)\alpha(i)$. Substituting this in, we can rewrite

$$\begin{aligned}
p\alpha(i) + (1-p)\alpha(i) &= p\alpha(i+1) + (1-p)\alpha(i-1) \\
(1-p)\left(\alpha(i) - \alpha(i-1)\right) &= p\left(\alpha(i+1) - \alpha(i)\right) \\
\alpha(i+1) - \alpha(i) &= \frac{1-p}{p}\left(\alpha(i) - \alpha(i-1)\right)
\end{aligned}$$

(c) Now that we have a relationship between $\alpha(i+1)-\alpha(i)$ and $\alpha(i)-\alpha(i-1)$, notice that we can iteratively apply the recurrence to get

$$\begin{aligned}
\alpha(i+1) - \alpha(i) &= \frac{1-p}{p}\left(\alpha(i) - \alpha(i-1)\right) \\
&= \left(\frac{1-p}{p}\right)^2 \left(\alpha(i-1) - \alpha(i-2)\right) \\
&\vdots \\
&= \left(\frac{1-p}{p}\right)^i \left(\alpha(1) - \alpha(0)\right) \\
&= \left(\frac{1-p}{p}\right)^i \alpha(1)
\end{aligned}$$

since $\alpha(0) = 0$ (once we lose all our money, we stop and can never reach n).

Further, notice that we have the telescoping sum

$$[\alpha(i) - \alpha(i-1)] + [\alpha(i-1) - \alpha(i-2)] + \cdots + [\alpha(1) - \alpha(0)] = \alpha(i) - \alpha(0) = \alpha(i).$$

This means that we have the summation

$$\begin{aligned}
\alpha(i) &= \sum_{k=0}^{i-1} \left(\alpha(k+1) - \alpha(k)\right) \\
&= \sum_{k=0}^{i-1} \left(\frac{1-p}{p}\right)^k \alpha(1) \\
&= \alpha(1) \sum_{k=0}^{i-1} \left(\frac{1-p}{p}\right)^k \\
&= \alpha(1) \cdot \frac{1 - \left(\frac{1-p}{p}\right)^i}{1 - \frac{1-p}{p}}
\end{aligned}$$

[Note that if $p = \frac{1}{2}$, the last step is not valid; in fact, since $\frac{1-p}{p} = 1$, this means that $\alpha(i) = i\alpha(1)$. We'll come back to this case later.]

The previous formula applies for all $0 < i \leq n$, so we can let $i = n$ and simplify to find $\alpha(1)$:

$$1 = \alpha(n) = \alpha(1) \cdot \frac{1 - \left(\frac{1-p}{p}\right)^n}{1 - \frac{1-p}{p}}$$

$$\frac{1 - \frac{1-p}{p}}{1 - \left(\frac{1-p}{p}\right)^n} = \alpha(1)$$

Plugging this back in for $\alpha(i)$, we have

$$\alpha(i) = \frac{1 - \frac{1-p}{p}}{1 - \left(\frac{1-p}{p}\right)^n} \cdot \frac{1 - \left(\frac{1-p}{p}\right)^i}{1 - \frac{1-p}{p}} = \frac{1 - \left(\frac{1-p}{p}\right)^i}{1 - \left(\frac{1-p}{p}\right)^n}.$$

Going back to the case where $p = \frac{1}{2}$, we saw that the summation simplifies to $\alpha(i) = i\alpha(1)$. Since $\alpha(n) = 1$, this means that $1 = n\alpha(1)$, or $\alpha(1) = \frac{1}{n}$. This means that we have

$$\alpha(i) = i\alpha(1) = \frac{i}{n}.$$

Together, we have the following formula for any $0 \leq i \leq n$:

$$\alpha(i) = \begin{cases} \frac{1-\left(\frac{1-p}{p}\right)^i}{1-\left(\frac{1-p}{p}\right)^n} & p \neq \frac{1}{2} \\ \frac{i}{n} & p = \frac{1}{2} \end{cases}.$$

(d) (i) If $p > \frac{1}{2}$, then $\frac{1-p}{p} < 1$, and as $n \to \infty$, the $\left(\frac{1-p}{p}\right)^n$ term in the denominator vanishes. This means that all we're left with is the numerator, and as such

$$\lim_{n\to\infty} \alpha(i) = 1 - \left(\frac{1-p}{p}\right)^i.$$

(ii) If $p = \frac{1}{2}$, then we know that $\alpha(i) = \frac{i}{n}$. As $n \to \infty$, this fraction goes to 0, and we have

$$\lim_{n\to\infty} \alpha(i) = 0.$$

(iii) If $p < \frac{1}{2}$, then $\frac{1-p}{p} > 1$, and as $n \to \infty$, the $\left(\frac{1-p}{p}\right)^n$ term in the denominator blows up. This means that the denominator tends to $-\infty$, while the numerator remains bounded for any fixed i. This means that the entire fraction tends to 0, i.e,

$$\lim_{n\to\infty} \alpha(i) = 0.$$

Note that this problem shows that, even in the case of a fair game (i.e., $p = \frac{1}{2}$), the probability that a gambler wins \$n before going broke tends to zero as $n \to \infty$. This is one version of the so-called "Gambler's Ruin" problem. Only in the case where $p > \frac{1}{2}$, i.e., when the game is strictly in the gambler's favor, does the gambler come out on top with positive probability.