

Midterm Classwork

- You have 2 hours and 15 minutes.
- You can check your own notes and books, and the lectures notes on the shared tablet.
- You cannot use your own electronics.

Exercises

Exercise 0.1. Let $\mathbf{X} = (X_t)_{t \geq 0}$ be a homogeneous Markov chain on a finite (or countable) state space S with transition probability $(p_{x,y})_{x,y \in S}$ and initial distribution $\mu \in \mathcal{P}(S)$, meaning $\mathbb{P}(X_0 = x) = \mu_x$. Fix $T \geq 0$ and define $Y_t = X_{T-t}$.

- Prove that $\mathbf{Y} = (Y_t)_{t \in \{0, \dots, T\}}$ has the Markov property. [Compute $\mathbb{P}(Y_t = y_t | Y_{t-1} = y_{t-1}, \dots, Y_0 = y_0)$.]
- Assume that $\mu = m$ is an invariant probability for \mathbf{X} with $m_x > 0$ for all $x \in S$. Check that \mathbf{Y} is a homogeneous Markov chain.

Solution

- To prove the Markov property for \mathbf{Y} , we must show that the probability of the state at time t only depends on the state at time $t-1$, and not on any earlier states. Let y_0, y_1, \dots, y_T be a sequence of states for which $p_{y_{t-i}, y_t}, \mu_{y_T} > 0$ for all $t \leq T$. Fix now $t \leq T$; we compute the conditional probability:

$$\begin{aligned} \mathbb{P}(Y_t = y_t | Y_{t-1} = y_{t-1}, \dots, Y_0 = y_0) &= \frac{\mathbb{P}(Y_t = y_t, Y_{t-1} = y_{t-1}, \dots, Y_0 = y_0)}{\mathbb{P}(Y_{t-1} = y_{t-1}, \dots, Y_0 = y_0)} \\ &= \frac{\mathbb{P}(X_{T-t} = y_t, X_{T-t+1} = y_{t-1}, \dots, X_T = y_0)}{\mathbb{P}(X_{T-t+1} = y_{t-1}, \dots, X_T = y_0)} \\ &= \frac{\mathbb{P}(X_{T-t} = y_t) p_{y_t, y_{t-1}} p_{y_{t-1}, y_{t-2}} \cdots p_{y_1, y_0}}{\mathbb{P}(X_{T-t+1} = y_{t-1}) p_{y_{t-1}, y_{t-2}} \cdots p_{y_1, y_0}} = \frac{\mathbb{P}(X_{T-t} = y_t) p_{y_t, y_{t-1}}}{\mathbb{P}(X_{T-t+1} = y_{t-1})} \end{aligned} \quad (1)$$

This final expression depends only on y_t and y_{t-1} (and the time index t), not on y_{t-2}, \dots, y_0 . Therefore, \mathbf{Y} has the Markov property. The transition probability $q_{x,y}^{(t)}$ for \mathbf{Y} to jump from state x to state y at time t is:

$$q_{x,y}^{t-1} = \mathbb{P}(Y_t = y | Y_{t-1} = x) = p_{y,x} \frac{\mathbb{P}(X_{T-t} = y)}{\mathbb{P}(X_{T-t+1} = x)} \quad (2)$$

Since the transition probabilities depend on t through the distributions of X_{T-t} and X_{T-t+1} , the reversed chain \mathbf{Y} is, in general, **not** homogeneous.

- If the initial distribution μ is an invariant probability m , then the chain is stationary. This means that for any time $t \geq 0$, the distribution of X_t is also m . That is, $\mathbb{P}(X_t = x) = m_x$ for all t . We can substitute

this into the previous formulas (either Equation 1 or Equation 2) to get

$$q_{x,y} = p_{y,x} \frac{m_y}{m_x}$$

This new expression for the transition probabilities, which we can call $p_{x,y}^\dagger$, depends only on the states x and y and the stationary distribution m . It no longer depends on the time index t . Therefore, when the original chain is stationary, the reversed chain \mathbf{Y} is a **homogeneous Markov chain**.

Exercise 0.2. A gambler plays a repeated game. At each turn they win with probability $p \in (0, 1)$ and lose with probability $1 - p$. To encourage playing, the casino gives a free night in their luxury suite, to any player who achieves a streak of 3 consecutive wins. The player decides to play until they win the prize (the free night).

- Model this process as a finite state-space Markov chain $(X_t)_{t \geq 0}$ with an absorbing state representing the win. Define the state space S , draw a graph representing the chain, find the communicating classes, and determine which classes are closed and which are not.
- For a generic Markov chain with Markov operator P , let A be a non-empty subset of the state space S (the “target set”), and let $\tau_A = \inf\{t \geq 0 : X_t \in A\}$ be the first time the chain hits the set A . For any state $x \in S$, let $u(x) = \mathbb{E}_x[\tau_A \mathbf{1}_{\tau_A < \infty}]$ be the expected time to reach A starting from x ¹. Show that u solves the problem for the unknown v :

$$\begin{cases} (I - P)v = 1 & \text{on } A^c \\ v = 0 & \text{on } A \end{cases}$$

- Assume $p = 1/2$. The casino can decide the cost for the gambler to play one turn (say it costs d dollars regardless if gambler wins or loses the turn). A free night costs the casino 100 dollars. For which values of d does the casino, on average, make a profit?

💡 Solution

- The state of the system needs to capture the player’s progress toward the goal, which is the number of consecutive wins. Let X_t be the number of wins since the last loss (the length of the current winning streak) after turn t . So $X_0 = 0$, $X_t = 1$ if the first turn was a win, $X_t = 0$ if the first turn was a loss, and so on. The possible values are the element of the *state space* $S = \{0, 1, 2, 3\}$. $X_t = 0$ if the player just lost, 1 if the player just made a win a win after a loss, and so on. When $X_t = 3$, the player wins the free night and stops playing. This can be visualized with the following graph:




Figure 1

¹You can assume that $\mathbb{P}_x(\tau_A < \infty) = 1$ for all $x \in A$, if you want to simplify the notation. In this case $u(x) = \mathbb{E}_x[\tau_A]$.

$\{0, 1, 2\}$ is a *non-closed communicating class*. $\{3\}$ is a *closed communicating class* which absorbs the chain with probability 1, $\mathbb{P}(\tau_3 < \infty) = 1$.

- b. Let u be the function defined in the problem text. If $x \in A$, the $\tau_A = 0 \mathbb{P}_x$ a.s., so $u = 0$. Thus consider $x \notin A$. Notice that in this case

$$\tau_A = \sum_{t=0}^{\tau_A} \mathbf{1}_{X_t \notin A} = \mathbf{1}_{X_0 \notin A} + \sum_{t=1}^{\tau_A} \mathbf{1}_{X_t \notin A} = 1 + \sum_{t=0}^{\tau_A-1} \mathbf{1}_{X_{t+1} \notin A}$$

Notice that, in the delicate case when $\mathbb{P}(\tau_A < \infty) < 1$, this still holds true, just both sides of the equation maybe be $+\infty$. In particular they coincide on $\{\tau_A < \infty\}$, namely on the paths \mathbf{X} that eventually touch A .

Also, for $X_0 \notin A$, $\tau_A(\mathbf{X}) - 1 = \tau_A(\theta_1 \mathbf{X})$ (which also holds when both sides are $+\infty$), where θ_1 is the time shift. Then conditioning on the first step we have, for $x \notin A$

$$\begin{aligned} u(x) &= \mathbb{E}_x[\tau_A \mathbf{1}_{\tau_A < \infty}] = 1 + \sum_y \mathbb{E}_x \left[\sum_{t=0}^{\tau_A-1} \mathbf{1}_{X_{t+1} \mathbf{1}_{\tau_A < \infty} \notin A} \mid X_1 = y \right] \\ &= 1 + \sum_y p_{x,y} \mathbb{E}_y \left[\sum_{t=0}^{\tau_A} \mathbf{1}_{X_t \mathbf{1}_{\tau_A < \infty} \notin A} \right] = 1 + (Pu)(x) \end{aligned}$$

- c. We need to find the expected number of turns to win, starting from state 0. This is precisely the quantity $u(0)$ where the target set is $A = \{3\}$. We are given $p = 1/2$. The system reads, representing v as a vector (v_0, \dots, v_3) :

$$\begin{cases} \frac{1}{2}v_0 = 1 + \frac{1}{2}v_1 \\ v_1 = 1 + \frac{1}{2}v_0 + \frac{1}{2}v_2 \\ v_2 = 1 + \frac{1}{2}v_0 + \frac{1}{2}v_3 \\ v_3 = 0 \end{cases}$$

$v_0 = 2 + v_1$, thus $v_1 = 4 + v_2$, then $v_2 = 1 + 1 + \frac{1}{2}(1 + v_1) = 4 + v_2/2$. Namely $v_2 = 8$, $v_1 = 12$, $v_0 = 14$. The expected number of games the gambler plays to win the prize is 14. Thus the casino will make profit on average if $14d > 100$ dollars: charging at least \$7.15 per turn will guarantee a positive average profit.

Exercise 0.3. Alice and Bob play a game again their friend Chris. Alice starts with a capital of $x_0 = 100$ dollars, Bob starts with $y_0 = 200$ dollars. Alice starts, then it is Bob's turn, then Alice's again and so on. We assume that Chris accepts debts from Alice and Bob, namely they can continue to play indefinitely even if they are in the red. At each turn when they play, the players can either win or lose one dollar with probability $1/2$. Let X_t (Y_t) be the amount of dollars that Alice (respectively Bob) has after t turns. So for instance $X_0 = 100$, and X_1 can be either 101 or 99, $X_2 = X_1$, while $Y_0 = Y_1 = 200$, Y_2 can be either 201 or 199 and so on. Of course Y_{2t+1} and X_{2t} have a different parity, so it makes sense to confront their distribution only after both players have played an equal number of times: at time $2t$ for $t \in \mathbb{N}$.

Prove that, for $z \in \mathbb{Z}$


$$|\mathbb{P}(X_{2t} = z) - \mathbb{P}(Y_{2t} = z)| \leq q_{2t, y_0 - x_0}$$

where $q_{t,x}$ is the probability that a random walk started at x hits 0 after time t . [Use a coupling argument, as in the proof of the convergence of aperiodic positive-recurrent chains. Be careful with the parity of t .]

Note: Recalling that for t large and $x \neq 0$, $q_{t,x} \approx |x|/\sqrt{\pi t}$ (we did this computation for $x = 0$ which is similar to the

case $|x| = 1$), we thus have that

$$\lim_{t \rightarrow \infty} \sup_z |\mathbb{P}(X_{2t} = z) - \mathbb{P}(Y_{2t} = z)| = 0$$

 Solution

Define $Z_t = Y_t - X_t$. Z_t is a simple symmetric random walk on \mathbb{Z} started at $y_0 - x_0$. Let $\tau := \inf\{t \geq 0 : X_t = Y_t\}$. τ is a stopping time, it is the hitting time of 0 for Z_t . We define a Markov chain \mathbf{Y}' as follows

$$\begin{cases} Y_t & \text{if } t \leq \tau \text{ or } t > \tau \text{ and } t \text{ is odd} \\ X_{t-1} & \text{if } t > \tau \text{ and } t \text{ is even} \end{cases}$$

Then \mathbf{Y}' has the same law of \mathbf{Y} (changes by ± 1 at each even time). Moreover at time $2t$

$$\mathbb{P}(X_{2t} = z, \tau < 2t) = \mathbb{P}(X_{2t-1} = z, \tau < 2t) = \mathbb{P}(Y'_{2t} = z, \tau < 2t) = \mathbb{P}(Y_{2t} = z, \tau < 2t)$$

Therefore

$$|\mathbb{P}(X_{2t} = z) - \mathbb{P}(Y_{2t} = z)| \leq \mathbb{P}(\tau < 2t) = q_{2t, y_0 - x_0}$$