Assignment 1

What you are allowed to do:

- Discuss the problems with your classmates.
- Ask the teacher for help.
- Use a computer for linear operations (e.g. matrix multiplication, matrix vector etc), provided you know how to do it manually.

What you are bound to do:

- Write down the solutions yourself, while being alone.

What you are not allowed to do:

- Copy the written solutions from a classmate.
- Use AI to output the solutions.

Problems

Exercise 0.1. Karl has a boring life. He either spends one hour on instagram, or one hour on TikTok, or watches Youtube videos about Markov Chains for 5 minutes, or sleeps for 25 minutes.

- When he wakes up, he always checks Instagram first.
- When finishing checking Instagram, he either watches Youtube videos about Markov Chains, or checks TikTok with equal probability.
- When he is done with Youtube, he always sleeps since Markov Chains are boring.
- When he is done with TikTok, he will either check Instagram or go to sleep with equal probability.

a. Represent Karl's life as a Markov Chain, with a transition probability matrix (stochastic matrix) and with a weighted graph.
b. Right now, Karl is on Instagram. Find the probability that the tenth app he will open on his phone will be Instagram. E.g. when considering the transitions I-> Y -> S -> I, the second app he opens is Instagram.
c. After many years of such a fascinating life, what percentage of his lifetime will he have spent on Instagram?

Solution

a. The transition probability operator P is (1: Instagram, 2: TikTok, 3: YouTube, 4: Sleep)

$$P = \begin{pmatrix} 0 & 1/2 & 1/2 & 0 \\ 1/2 & 0 & 0 & 1/2 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix}$$

and here the associated graph.

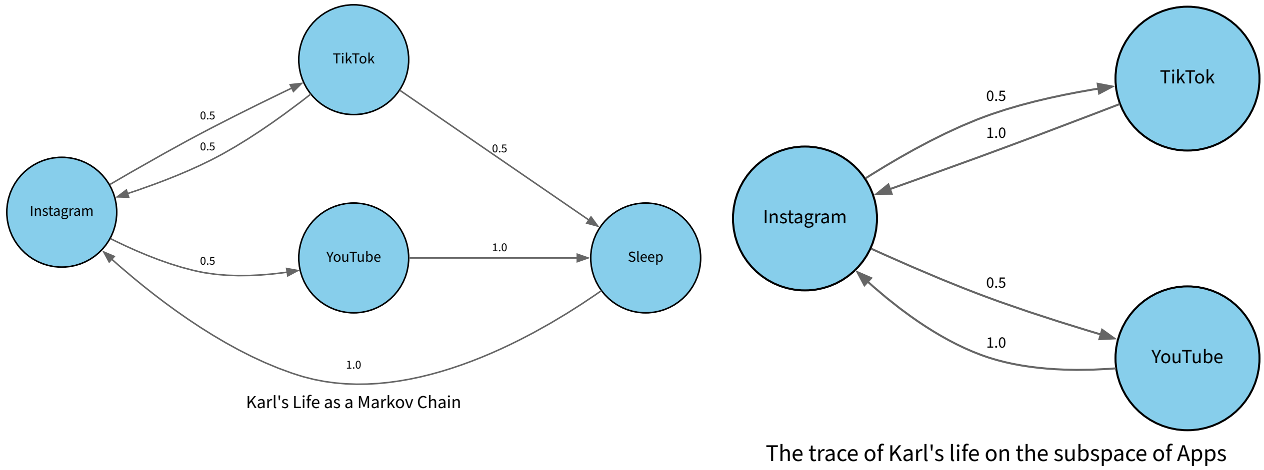

Figure 1: Karl's life and the trace on the subspace of apps. Every second app he uses is Instagram.

b. Every second app he opens is Instagram, so this probability is 1.

c. Let's call $\mathcal{N}_x$ the number of times Karl visits the state x after many years, say after N steps. We understand that the problem asks for the limit $N \to \infty$.

Let $\gamma_x = \mathbb{E}_x[\mathcal{N}_x]/N$ the expected number times he has been in the state x. Then

- The expected number of times he visits YouTube is half the number of times he visits Instagram up to a small (random) error: $\gamma_Y = \gamma_I/2 + \varepsilon_Y$.
- The expected number of times he visits TikTok is also half the number of times he visits Instagram up to a small random error: $\gamma_T = \gamma_I/2 + \varepsilon_T$.
- Similarly: $\gamma_S = \gamma_Y + \gamma_T/2 + \varepsilon_S$.
- Similarly: $\gamma_I = \gamma_S + \gamma_T/2 + \varepsilon_I$.
- To solve one should not forget: $\gamma_I + \gamma_Y + \gamma_T + \gamma_S = 1$.

Here what we mean is that ε_x is some small error $O(1/N)$, which is due to the fact that Karl starts in a specific state (Instagram). So for instance for $N = 1$, we have $\gamma_I = 1$, all the others being 0. So we get $\gamma_I = 4/11 + O(1/N)$, $\gamma_Y = 2/11 + O(1/N)$, $\gamma_T = 2/11 + O(1/N)$, $\gamma_S = 3/11 + O(1/N)$. The fraction of time he spends on Instagram will then be

$$\frac{\gamma_I * 60 + O(1/N)}{\gamma_I * 60 + \gamma_Y * 5 + \gamma_T * 60 + \gamma_S * 25 + O(1/N)} \simeq 0.5393 + O(1/N) \tag{1}$$

So after many years ($N \to \infty$) he will have spent around 54% of his boring life on Instagram.

As a more robust alternative, we can consider

$$\lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} f(X_t) = m(f)$$

where m is the invariant measure of the process. Actually m_x is nothing but the leading term of the aforementioned γ_x, namely $m_I = 4/11$, $m_Y = m_T = 2/11$, $m_S = 3/11$. Now we take $f(x)$ as the amount of time he spends each time on the state x (e.g. $f(I) = 60$ minutes etc), and $g(x) = f(x)\mathbf{1}_{x=I}$. Then we get

$$\lim_{T\to\infty} \frac{\sum_{t=0}^{T-1} g(X_t)}{\sum_{t=0}^{T-1} f(X_t)} = \frac{m(g)}{m(f)} \simeq 0.5393$$

Exercise 0.2. Anna has just received m bikes for her birthday. Each day now, she will pick one bike randomly (each one with probability $1/m$) and return home in the evening. Let X_t be the number of different bikes she will have used after t days (so $X_0 = 0$, $X_1 = 1$, but X_2 may be 1 or 2 and so on).

a. Explain why (X_t) is a Markov chain.
b. Let τ_m be the random time at which she will have tried all the bikes. Is this a stopping time?
c. Find the probability $\mathbb{P}(\tau_m = t)$, recalling that $X_0 = 0$.

Solution

a. The probability that on a given day she will use a never-used bike only depends on how many bikes has she already used, and not on which or which order she used them. So this is a Markov chain with state space $S = \{0, \dots, m\}$, and $p_{x,x} = x/m = 1 - p_{x,x+1}$ (of course $p_{m,m} = 1$ so $p_{m,m+1}$ is just not defined). Each singleton of the state space is a communicating class, with $\{m\}$ being the only closed one.
b. This is a stopping time. Actually it is the hitting time of $\{m\}$.
c. We have $\mathbb{P}_x(\tau_m = t) = \mathbb{P}_x(X_{t-1} = m-1, X_t = m) = p^{(t-1)}_{x,m-1} p_{m-1,m}$. $p_{m-1,m} = 1/m$ so we just need to compute the power $p^{(t-1)}_{0,m-1}$. The matrix P has eigenvalues $\{0, 1/m, 2/m, \dots, 1-1/m, 1\}$. They are all real and simple, and let us call $\lambda_x = x/m$. The fact that $\lambda_x = x/m$ corresponds to the following remark[1]: once Anna has used exactly x bikes (e.g. if the chain starts in the state x), she will not use a never-used bike for t more days with probability λ_x^t. The corresponding left eigenvectors $f^x P = \lambda_x f^x$, and right eigenvectors $P e^x = \lambda_x e^x$ are

$$(e^x)_y = \binom{m-y}{x-y} \mathbf{1}_{y \le x}$$

$$(f^x)_y = (-1)^y \binom{m-x}{y-x} \mathbf{1}_{y \ge x}$$

Thus by linear algebra

$$(p^{(t)})_{x,y} = \sum_{z=0}^{m} \frac{(e^z)_x \, \lambda_z^t \, (f^z)_y}{e^z \cdot f^z} = \sum_{z=x}^{y} (-1)^{y-z} \binom{m-x}{z-x, y-z, m-y} \left(\frac{z}{m}\right)^t$$

In particular

$$\mathbb{P}(\tau_m = t) = p^{(t)}_{0,m-1}/m = \sum_{z=0}^{m-1} (-1)^{m-z+1} \binom{m-1}{z} \left(\frac{z}{m}\right)^{t-1} \tag{2}$$

An alternative method involves the inclusion-exclusion principle. Let A_i be the event that bike i has *not* been used by day t. Let's first compute

$$\mathbb{P}(X_t = m) = 1 - \mathbb{P}\left(\cup_{i=1}^m A_i\right)$$

By the principle of inclusion-exclusion:

$$\mathbb{P}\left(\bigcup_{i=1}^m A_i\right) = \sum_i \mathbb{P}(A_i) - \sum_{i<j} \mathbb{P}(A_i \cap A_j) + \dots + (-1)^{m+1} \mathbb{P}(A_1 \cap \dots \cap A_m)$$

The probability of not picking a specific set of k bikes in t days is: $\mathbb{P}(A_{i_1} \cap \dots \cap A_{i_k}) = (\frac{m-k}{m})^t$ There are

$\binom{m}{k}$ ways to choose which k bikes to exclude. Plugging this into the inclusion-exclusion formula gives

$$\mathbb{P}(\tau_m = t) = \mathbb{P}(X_t = m) - \mathbb{P}(X_{t-1} = m) = \sum_{k=1}^{m} (-1)^k \binom{m}{k} \left((\tfrac{m-k}{m})^t - (\tfrac{m-k}{m})^{t-1} \right)$$

which gives the same answer as Equation 2.

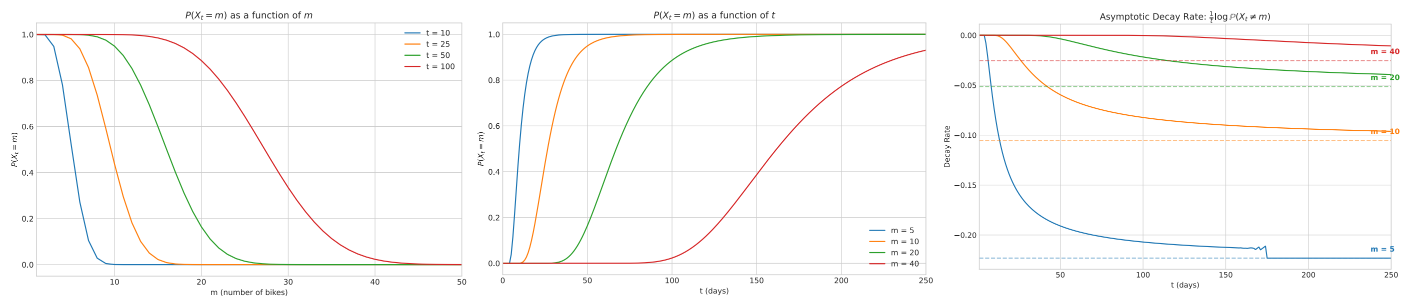

Figure 2: The plots show the $\mathbb{P}(X_t = m)$ as a function of m for fixed t, and more interestingly as a function of t for fixed m. For t large, $\mathbb{P}(X_t \neq m) \simeq (\lambda_{m-1})^t \simeq e^{-t/m}$. We see this fact by computing $\frac{1}{t} \log \mathbb{P}(X_t \neq m)$. For t large, it converges to its asymptote $\log(1 - 1/m)$. For m small, the convergence to 0 is faster: when the probability is thus very small, we see a machine-error issue (the weird behavior in the blue plot), a reminder of the challenging problem of computing exponents numerically.

Exercise 0.3. In a tennis game, two players A and B play points. Player A wins a point with probability p and player B wins a point with probability $1 - p$. The game ends when one of the players has won 4 points in total and at least 2 points more than the opponent.

a. Model the game as a Markov chain on a finite state space, and represent it as a weighted graph.
b. Is this an irreducible Markov chain? If not, determine the communicating classes and the closed ones.
c. What is the probability that the game lasts no more than 5 points?
d. What is the probability that player A wins the game?

Suppose now that player A has a psychological advantage, and if the point decides the game (namely the point may end the game), A wins it with probability $q > p$.

e. What is the probability that player A wins the game now?
f. (Optional) Run a computer simulation to estimate the results of a. and b. numerically for a fixed value of p. Compare the results with the theoretical ones.

Solution

a. While in principle the points count can grow indefinitely, we can make it a finite-state chain by counting only the advantage of the players once both reached 3 points. This is exactly how they count points officially in tennis games, and it can be easily represented by the following directed graph.

[1] One may check that if $p_{x,x}^{(t)} = \lambda^t$ for some $x \in S$, $\lambda \in [0, 1]$ and *all* $t \in \mathbb{N}$, then λ is an eigenvalue of P.

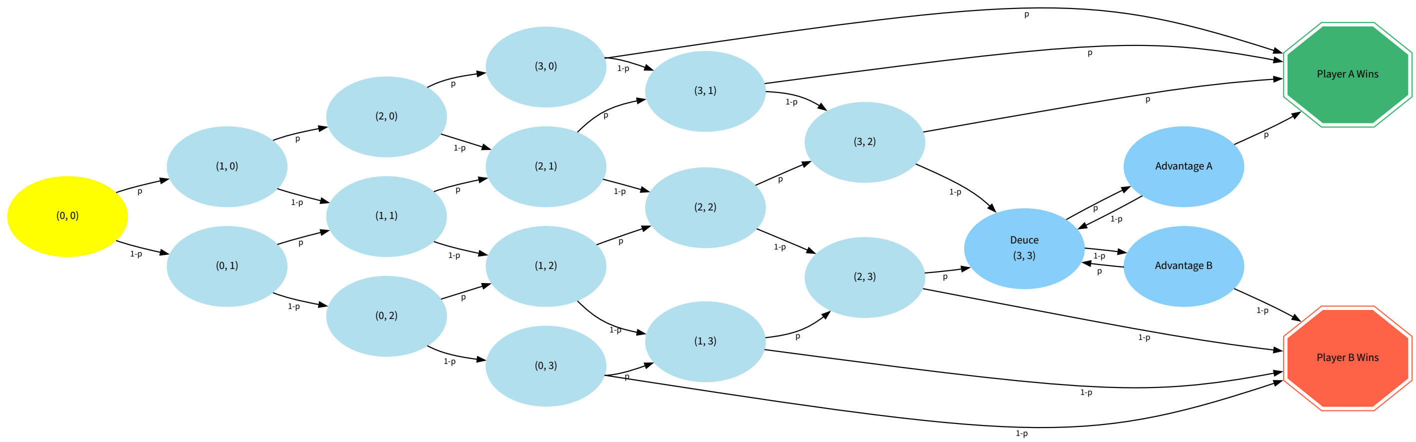

Figure 3

b. Each state where at least one of the two players has no more than 2 points is a communicating class as a singleton. Indeed, this score can never be reached again. Each of the two states where one player wins, is a closed communicating class. Finally the three states, *Deuce*, *Advantage A* and *Advantage B* are all together a communicating class (not a closed one though).

c. The game lasts not more than 5 points if it ends in exactly 4 or 5 points.

- The game ends in 4 points with probability $p^4+(1-p)^4$ (probability of the trajectory where A always wins, plus the one where B always wins).
- The game ends in 5 points exactly on the following trajectories: A wins 3 of the first 4 points, then wins the last one; or B wins 3 of the first 4 points, then wins the last one. The probability in this case is $\binom{4}{1}p^4(1-p)+\binom{4}{1}p(1-p)^4$.

The required probability is then $p^4+(1-p)^4+4p^4(1-p)+4p(1-p)^4$.

d. Let a, b be the states where A wins and B wins respectively, i the initial state with score $(0,0)$, and d the state *Deuce*. We use the notation introduced in the notes: $h_x(y)$ is the probability of hitting the state x when starting in y, and $h_{x,y}(z)$ is the probability of hitting x before y when starting in z. Since $\mathbb{P}(\tau_{\{a,b\}} < \infty) = 1$, in principle we can just define $h_a(x)$ as the probability that A wins when the score is x, solve the linear problem for hitting probabilities and find the required $h_a(i)$.

The state space here is quite large however, so let's compute this probability less systematically. The event that player A wins the game is the union of all the possible trajectories ending in a. There are two types of trajectories: those where B never reaches 3 points (so that *Deuce* is not hit) and those reaching *Deuce*.

- The probability of winning without ever reaching *Deuce* is the sum of the probabilities of all those trajectories: $h_{a,d}(i) = p^4+\binom{4}{1}p^4(1-p)+\binom{5}{3}p^4(1-p)^2$ (there is one trajectory where A wins all points, $\binom{4}{1}$ where B wins 1 point and A wins 4 points including the last one, $\binom{5}{2}$ where B wins 2 points and A wins 4 points including the last one).
- The probability of reaching *Deuce* is $h_d(i) = \binom{6}{3}p^3(1-p)^3$.
- Once we are in *Deuce*, we can forget about the rest and solve the harmonic problem from this point on: if A wins the next two points (probability p^2), A wins the game. If B wins the next two points (probability $(1-p)^2$), B wins the game. If they each win one point (probability $2p(1-p)$), the score returns to deuce. We can write a recursive equation $h_a(d)$:

$$h_a(d) = p^2 + 2p(1-p)h_a(d)$$

or $h_a(d) = \frac{p^2}{p^2+(1-p)^2}$.

- The probability of hitting *Deuce* and winning is then $h_a(d)h_d(i)$ (by the strong Markov property).

Therefore the probability that A wins is

$$h_a(i) = h_{a,d}(i) + h_a(d)h_d(i) = p^4(1 + 4(1-p) + 10(1-p)^2) + 20p^3(1-p)^3 \frac{p^2}{p^2 + (1-p)^2}$$

e. This works with the same approach as before, but hopefully the solution cannot be found online! Let's calculate the probability of A winning under this new condition.

- The probability of winning without ever reaching *Deuce* is $h_{a,d}(i) = [p^3]q + [p^3(1-q) + \binom{3}{1}p^3(1-p)]q + [p^3(1-q)^2 + \binom{3}{1}p^3(1-p)(1-q) + \binom{4}{2}p^3(1-p)^2]q$, where the square bracketed terms, multiplied by q, correspond to the probabilities of winning in 4, 5 and 6 points respectively.
- The probability of reaching *Deuce* can be computed as

$$h_d(i) = (1-q)h_{(3,2)}(i) + qh_{(2,3)} = (1-q)\left[p^3(1-q)^2 + \binom{3}{1}p^3(1-p)(1-q) + \binom{4}{2}p^3(1-p)^2\right] + q\left[(1-p)^3q^2 + \binom{3}{1}(1-p)^3pq + \binom{4}{2}(1-p)^3p^2\right]$$

- The probability of winning from the *Deuce* state solves

$$h_a(d) = p(q + (1-q)h_a(d)) + (1-p)(qh_a(d))$$

or $h_a(d) = (pq)/(1-p-q+2pq)$.

The probability of winning is ultimately

$$h_a(i) = h_{a,d}(i) + h_a(d) * h_d(i) = \frac{pq(-9p^4(q-1) - p^3(q^2-24q+23) + p^2(2q^3-2q^2-15q+15) - 3p(q-1)q^2 + q^3)}{p(2q-1)-q+1}$$

This can be more easily solved just using a linear solver (the harmonic equation), which gives the same solution

```python
import sympy as sp

# Declare sympy symbols
p, q, h_d, h_a, h_b = sp.symbols('p q h_d h_a h_b')

# Solve the deuce sub-problem in-line to get the win probability from deuce.
h_deuce = sp.solve([
 sp.Eq(h_d, p*h_a + (1-p)*h_b),
 sp.Eq(h_a, q + (1-q)*h_d),
 sp.Eq(h_b, q*h_d)
], (h_d, h_a, h_b))[h_d]

# Define a recursive solver for the probability from any score (a,b).
def prob(a, b):
 if a >= 4 and a >= b + 2: return 1
 if b >= 4 and b >= a + 2: return 0
 if a == 3 and b == 3: return h_deuce

 # Use 'q' on game points, otherwise 'p'.
 pt_prob = q if (a == 3 and b < 3) or (b == 3 and a < 3) else p
 return pt_prob * prob(a + 1, b) + (1 - pt_prob) * prob(a, b + 1)

# Print the result from the initial score (0,0)
result = sp.simplify(prob(0, 0))

print(f"""--- Probability of Player A Winning a Game ---

[ LaTeX format ]
{sp.latex(result)}

[ Mathematica format ]
{sp.mathematica_code(result)}

[ Plain text format ]
{result}""")
```

```
--- Probability of Player A Winning a Game ---

[ LaTeX format ]
\frac{p q \left(- 9 p^{4} q + 9 p^{4} - p^{3} q^{2} + 24 p^{3} q - 23 p^{3} + 2 p^{2} q^{3} -

[ Mathematica format ]
p*q*(-9*p^4*q + 9*p^4 - p^3*q^2 + 24*p^3*q - 23*p^3 + 2*p^2*q^3 - 2*p^2*q^2 - 15*p^2*q + 15

[ Plain text format ]
p*q*(-9*p**4*q + 9*p**4 - p**3*q**2 + 24*p**3*q - 23*p**3 + 2*p**2*q**3 - 2*p**2*q**2 - 15
```

f. Here the comparison. Feel free to resize the box and edit the code.

Exercise 0.4. We want to give an analytical characterization of the transition probabilities $q_{x,y}$ of the trace of a recurrent Markov chain. We proved that the trace of **X** on A has transition probabilities, for $x, y \in A$

$$q_{x,y} = \mathbb{P}_x(X_{\tau_A^+} = y)$$

Define the operator $P^{(A^c)}$ with entries

$$p_{x,y}^{(A^c)} = \begin{cases} p_{x,y} & \text{if } y \notin A \\ 0 & \text{if } y \in A \end{cases}$$

Prove that, for each fixed $y \in A$, $q_{x,y}$ is the minimal solution to the problem in the unknown $h \geq 0$

$$((I - P^{(A^c)})h)(x) = p_{x,y}, \qquad x \in S \tag{3}$$

meaning $q_{x,y} - \sum_{z \notin A} p_{x,z} q_{z,y} = p_{x,y}$ for all $x \in S$.

Solution

Fix $y \in A$, and denote $h(x) \equiv q_{x,y} = \mathbb{P}_x(X_{\tau_A^+} = y)$ for $x \in S$. By conditioning on the first step:

$$h(x) = \sum_{z \in S} p_{x,z} \mathbb{P}_x(X_{\tau_A^+} = y | X_1 = z) = \sum_{z \in A} p_{x,z} \mathbf{1}_{y=z} + \sum_{z \notin A} p_{x,z} h(z) = p_{x,y} + \sum_{z \notin A} p_{x,z} h(z)$$

To prove minimality, let now $g \geq 0$ be any non-negative solution to the system Equation 3 for a fixed $y \in A$ (which we omit from the notation). Define $h^{(n)}(x) := \mathbb{P}_x(X_{\tau_A^+} = y, \tau_A^+ \leq n)$. Then $h^{(n)}(x) \leq q_{x,y}$, and by the continuity of probabilities on monotone sequences, $h^{(n)}(x) \uparrow h(x)$. Therefore it is enough to check that $g \geq h^{(n)}$ for all $n \geq 0$. Now we proceed by induction.

- Trivially $g \geq h^{(0)} \equiv 0$.
- $h^{(n)}$ satisfies (reasoning as for $q_{\cdot,y}$ above)

$$h^{(n+1)}(x) = p_{x,y} + \sum_{z \notin A} p_{x,z} h^{(n)}(z)$$

Therefore by the induction hypothesis

$$g(x) = p_{x,y} + (P^{(A^c)} g)(x) \geq p_{x,y} + (P^{(A^c)} h^{(n)})(x) = h^{(n+1)}(x)$$