In the name of God.

Sharif University of Technology

Stochastic Process

Fall 2025

Hamid R. Rabiee

Homework 5 | Markov Chains, Hidden Markov Model | Deadline : 1404/10/22

- **Blind copying answers from LLMs is not acceptable. Any sign of blind usage will result in losing the entire grade for the homework.**
- **Provide full justifications and explanations for every step of your solution. You are required to show that you fully understand what you have written!**
- **You may want to use a computer in solving some of the questions. If so, upload a Jupyter notebook containing the code and the results in addition to your solution file. In this case, you must also explain the algorithm in the solution file.**

1. [40] (Expand Your Dungeon!) There is a dungeon with multiple rooms, part of which is depicted in Fig. 1. Each room has a one-way door to one or more other rooms, and there is a uniform probability of choosing among the available doors for someone present in a room. Please answer the following:

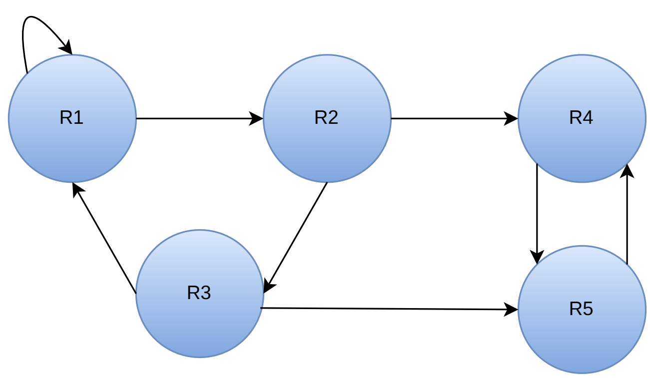

   (a) Is the process of traversing the rooms a Markov chain? Why or why not? If not, what should be added to make it one?
   (b) How many communicating classes exist in the current version of the dungeon? Which ones are transient and which ones are recurrent?
   (c) Please expand the dungeon by adding 11 new rooms and 11–40 new doors, resulting in three new communicating classes: two recurrent and one transient.
   (d) Change the doors so that traversing the dungeon can lead to three different final stationary distributions.
   (e) Starting from room R1, what is the probability of ending up in each of the possible stationary distributions?

(f) Select one transient room (named R^T) and one recurrent room with no transient neighbors, such that there exists a path from R^T to it (named R^R). Calculate the expected number of steps needed to reach R^R starting from R^T, conditioned on the fact that the traverser follows a path that eventually reaches R^R.

Solution:

(a) Yes, it is a Markov chain as it is. Presence in the rooms defines the states of the Markov chain, and the probability of being in each room depends on the last room the traverser was in.

(b) 2, One recurrent (R4, R5), and one transient (R1, R2, R3)

(c) Please see Fig. 1c.

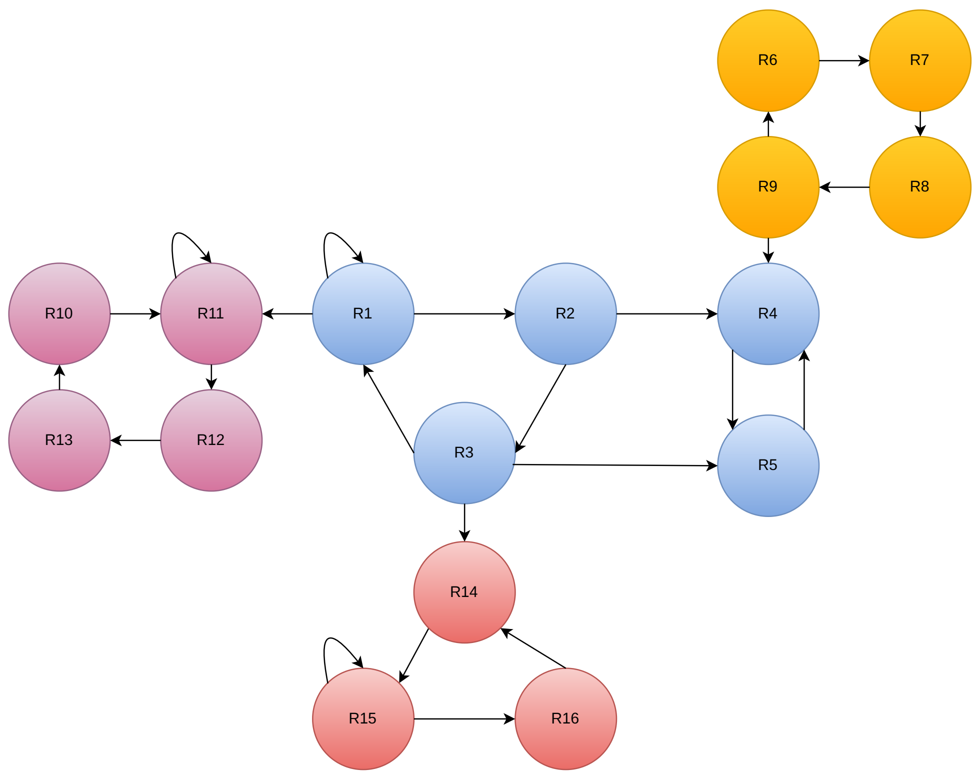

(d) We can end up in each of the communicating recurrent classes after traversing the rooms, and there are three of them. To have three possible stationary distributions—one unique distribution for each class—all rooms in each communicating recurrent class need to be aperiodic and irreducible. We can add a self-loop door to R5 to satisfy the aperiodicity condition.

(e) We can aggregate the rooms related to each recurrent class into a single state for this question, as shown in Fig. 1e:

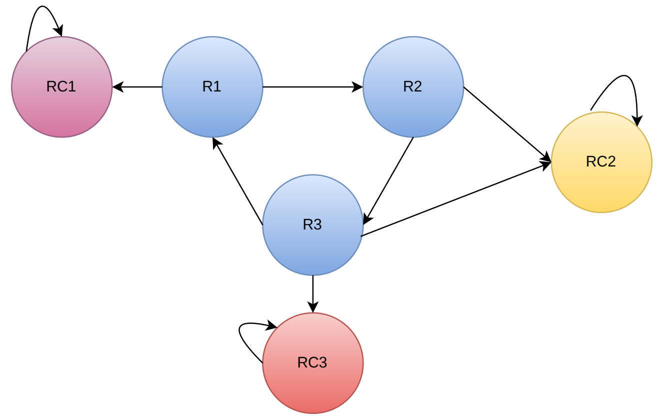

We define $q_{i,j}$ as the probability of eventually reaching recurrent class j, starting from room i. Following the probability rule $q_{i,j} = \sum_{k=1}^{3} p_i k \times q_{k,j}$:

$$\begin{cases} q_{1,1} = \frac{1}{3} \times 1 + \frac{1}{3} \times q_{1,1} + \frac{1}{3} q_{2,1} \\ q_{2,1} = \frac{1}{2} \times 0 + \frac{1}{2} \times q_{3,1} \\ q_{3,1} = \frac{1}{3} \times 0 + \frac{1}{3} \times 0 + \frac{1}{3} \times q_{1,1} \end{cases} => q_{1,1} = \frac{6}{11} \tag{1}$$

$$\begin{cases} q_{1,2} = \frac{1}{3} \times 0 + \frac{1}{3} \times q_{1,2} + \frac{1}{3} q_{2,2} \\ q_{2,2} = \frac{1}{2} \times 1 + \frac{1}{2} \times q_{3,2} \\ q_{3,2} = \frac{1}{3} \times 1 + \frac{1}{3} \times 0 + \frac{1}{3} \times q_{1,2} \end{cases} => q_{1,2} = \frac{4}{11} \tag{2}$$

$$\begin{cases} q_{1,3} = \frac{1}{3} \times 0 + \frac{1}{3} \times q_{1,3} + \frac{1}{3} q_{2,3} \\ q_{2,3} = \frac{1}{2} \times 0 + \frac{1}{2} \times q_{3,3} \\ q_{3,3} = \frac{1}{3} \times 0 + \frac{1}{3} \times 1 + \frac{1}{3} \times q_{1,3} \end{cases} => q_{1,3} = \frac{1}{11} \tag{3}$$

As expected, $q_{1,1} + q_{1,2} + q_{1,3} = 1$.

(f) Assuming R3 and R16, we define $S_{i,j}$ as the expected number of steps that it takes to each state j, starting from i, We would have $S_{i,j} = \sum_k 1 + (p_i k \times S_{k,j})$ and $\forall i : S_{i,i} = 0$. Note that given the conditional assumption, the traverser would never pass the doors $2-4$, $3-5$, and $1-11$.

$$\begin{cases} S_{3,16} = 1 + (\frac{1}{2} \times S_{14,16} + \frac{1}{2} S_{1,16}) \\ S_{1,16} = 1 + (\frac{1}{2} \times S_{1,16} + \frac{1}{2} S_{2,16}) \\ S_{2,16} = 1 + S_{3,16} \\ S_{14,16} = 1 + S_{15,16} \\ S_{15,16} = 1 + (\frac{1}{2} \times S_{15,16} + \frac{1}{2} S_{16,16}) \end{cases} => \begin{cases} e_5 => S_{15,16} = 2 \\ e_4 => S_{14,16} = 3 \\ e_2 \& e_3 => S_{1,16} = S_{3,16} + 3 \\ e_1 => S_{3,16} = 5 + S_{14,16} \\ => S_{3,16} = 8 \end{cases} \tag{4}$$

2. [20]

- Show that an ergodic Markov chain with M states must contain a cycle with $\tau <$ M states.
- Let X be a fixed state on this cycle of length τ. Let T(m) be the set of states accessible from X in m steps. Show that:

$$\forall m \geq 1; T(m) \subseteq T(m+\tau) \tag{5}$$

Solution:

Solution: The states in any cycle are distinct and thus a cycle contains at most M states. An ergodic chain must contain cycles, since for each pair of states $\ell \neq j$, there is a walk from ℓ to j and then back to ℓ; if any state i other than ℓ is repeated in this walk, the first i and all subsequent states before the second i can be eliminated. This can be done repeatedly until a cycle remains.

Finally, suppose a cycle contains M states. If there is any transition $P_{im} > 0$ for which (i, m) is not a transition on that cycle, then that transition can be added to the cycle and all the transitions between i and m on the existing cycle can be omitted, thus creating a cycle with fewer than M states. If there are no nonzero transitions other than those in the original cycle with M states, then the Markov is periodic with period M and thus not ergodic.

b) Let ℓ be a fixed state on a cycle of length $\tau < \mathsf{M}$. Let $\mathcal{T}(m)$ be the set of states accessible from ℓ in m steps. Show that for each $m \geq 1$, $\mathcal{T}(m) \subseteq \mathcal{T}(m+\tau)$. Hint: For any given state $j \in \mathcal{T}(m)$, show how to construct a walk of $m + \tau$ steps from ℓ to j from the assumed walk of m steps.

Solution: Let j be any state in $\mathcal{T}(m)$. Then there is an m-step walk from ℓ to j. There is also a cycle of length τ from state ℓ to ℓ. Concatenate this cycle (as a walk) with the above m step walk from ℓ to j, yielding a walk of length $\tau + m$ from ℓ to j. Thus $j \in \mathcal{T}(m+\tau)$ and it follows that $\mathcal{T}(m) \subseteq \mathcal{T}(m+\tau)$.

3. [10] What is the order of complexity of the Viterbi algorithm in a discrete-state Markov chain with K different states for the observation sequence $Y_1, \ldots, Y_N$? Does it produce the optimal solution? Prove your answer either by a counterexample or by mathematical justification.

Solution:

(a)
- $O(NK^2)$
- Yes, it is a kind of dynamic programming resulting in the optimum solution.
- The goal is:

$$\begin{aligned}
&\arg\max_{X_{1:N}}(P(X_1, ..., X_N, Y_1, ..., Y_N)) = \\
&\arg\max_{X_{1:N}}(\pi_1 \prod_{i=2}^{N}[P(X_i|X_{i-1})] \prod_{i=1}^{N}[P(Y_i|X_i)]) = \\
&\arg\max_{X_N} P(Y_N|X_N) \arg\max_{X_{1:N-1}}(\pi_1 \prod_{i=2}^{N}[P(X_i|X_{i-1})] \prod_{i=1}^{N-1}[P(Y_i|X_i)])
\end{aligned} \tag{6}$$

Now if we find $\arg\max_{X_{1:N-2}}(\pi_1 \prod_{i=2}^{N-1}[P(X_i|X_{i-1})] \prod_{i=1}^{N-1}[P(Y_i|X_i)])$ as a function of X_{N-1} and call it $\mu(X_{N-1})$, we would have:

$$\begin{aligned}
&\arg\max_{X_{1:N-1}}(\pi_1 \prod_{i=2}^{N}[P(X_i|X_{i-1})] \prod_{i=1}^{N-1}[P(Y_i|X_i)]) = \\
&\arg\max_{X_{N-1}}(\mu(X_{N-1})P(X_N|X_{N-1}))
\end{aligned} \tag{7}$$

, which is what Viterbi does. At each step, it calculates the probability of the chain of the hidden states with maximum joint probability with the observations as a function of the value of the state at the step ($\mu(X_t)$), and then calculates the same function for the next step using $P(Y_{t+1}|X_{t+1}) \times \max_{X_t}[\mu(X_t) \times P(X_{t+1}|X_t)]$. For the last step ($N$), the algorithm takes maximum over all possible values for X_N.

4. [15] We have two baskets, each containing m white balls and m black balls. At each step, one random ball is swapped with another random ball from the other basket. Consider the number of black balls as the state descriptor for the described Markov chain. Is this Markov chain time-reversible? Prove your answer either by a counterexample or by mathematical justification.

Solution:

Clearly, there is a path from each state to every other state with a positive probability (it just needs to select the differing balls and change their baskets), and all states are aperiodic because we can state in each state with a self loop by only swapping white balls. For the state of full black/white balls, we have a path of going to the neighbor and returning (a path of length 2), and going to the neighbor, looping over the neighbor with one step, and returning (a path of length 3), which gives a period of 1. Therefore the chain is also aperiodic. So we have a stationary probability and it is unique.

In the next step, we intuitively propose $\pi_i = \frac{\binom{m}{i}\binom{m}{m-i}}{\binom{2m}{m}}$. Why this one? Because it is like the possible number of selections of i black balls and $m-i$ white balls over all possible selections of $2m$ balls. Now we need to prove that it is the stationary probability by showing it satisfies both conditions:

- $\sum_i \pi_i = 1$

$$\sum_i \pi_i = \sum_{i=0}^{i=m} \frac{\binom{m}{i}\binom{m}{m-i}}{\binom{2m}{m}} = \frac{\sum_{i=0}^{i=m}\binom{m}{i}\binom{m}{m-i}}{\binom{2m}{m}} \tag{8}$$

  Now using a binomial expansion, we can show $\sum_{i=0}^{i=m}\binom{m}{i}\binom{m}{m-i} = \binom{2m}{m}$
  We know $(x+y)^z = \sum_{k=0}^{z}\binom{z}{k}x^k y^(z-k)$, therefore the coefficient of x^m in expansion of $(x+1)^{2m}$ would be $\binom{m}{2m}$. Another way to write $(x+1)^{2m}$ is $(x+1)^m \times (x+1)^m$, and the factors that their multiplication results in x^m, is a x^k from the first expansion and $x^(m-k)$ from the second one and obviously, the coefficient is the multiplication of theirs. Now we have options from 0 to m for k. The first coefficient is $\binom{m}{k}$ and the second one is $\binom{m}{k-m}$. Therefore, $\sum_{i=0}^{i=m}\binom{m}{i}\binom{m}{m-i} = \binom{2m}{m}$.
- $\pi_i = \sum_j \pi_j p_{ji}$.
  We only have three states that can go to the state of i with one step:
  - Being at state i (π_i), swapping either two back balls or two white balls $p_{ii} = \frac{i}{m}\frac{m-i}{m} + \frac{m-i}{m}\frac{i}{m}$.
  - Being at state $i+1$ (π_{i+1}), swapping a back ball with a white ball $p_{i+1,i} = \frac{i+1}{m}\frac{i+1}{m}$
  - Being at state $i-1$ (π_{i-1}), swapping a white ball with a black ball $p_{i-1,i} = \frac{m-i+1}{m}\frac{m-i+1}{m}$

  Therefore we would have:

$$\begin{aligned}\sum_j \pi_j p_{ji} &= \frac{\binom{m}{i}\binom{m}{m-i}}{\binom{2m}{m}}\frac{2m(m-i)}{m^2} + \frac{\binom{m}{i+1}\binom{m}{m-i-1}}{\binom{2m}{m}}\frac{(i+1)^2}{m^2} + \frac{\binom{m}{i-1}\binom{m}{m-i+1}}{\binom{2m}{m}}\frac{(m-i+1)^2}{m^2} \\ &= \frac{(m!)^2}{m^2(i!)^2((m-i)!)^2} \times \left[2i(m-i) + (m-i)^2 + i^2\right] = \frac{m^2(m!)^2}{m^2(i!)^2((m-i)!)^2} = \pi_i\end{aligned} \tag{9}$$

  To prove time reversibility we need to show that $\pi_i p_{ij} = \pi_j p_{ji}$, which we need to show between only two successive states here (as the other p_{ij}s are 0), i and i + 1:

$$
\begin{aligned}
&p_{i,i+1} = \frac{m-i}{m}\frac{m-i}{m}, p_{i+1,i} = \frac{i+1}{m}\frac{i+1}{m} \\
&\pi_i p_{i,i+1} = \pi_{i+1} p_{i+1,i} => \frac{\binom{m}{i}\binom{m}{m-i}}{\binom{2m}{m}}(\frac{m-i}{m})^2 = \frac{\binom{m}{i+1}\binom{m}{m-i-1}}{\binom{2m}{m}}(\frac{i+1}{m})^2 \\
&=> (\frac{m!}{i!(m-i)!}\frac{m-i}{m})^2 = (\frac{m!}{(i+1)!(m-i-1)!}\frac{i+1}{m})^2 \\
&=> (\frac{1}{(m-i)}(m-i))^2 = (\frac{1}{(i+1)}(i+1))^2, \text{which holds as 1=1!}
\end{aligned} \tag{10}
$$

5. [15] We have a three-armed bandit machine with arms A, B, and C. When an arm is pulled, it sends a character from the set A, B, C to the machine's calculation unit. Due to noise in the system, each arm sends its corresponding character with probability 0.8 and each of the other two characters with probability 0.1. If the machine receives the continuous pattern $ACAC$, it rewards the puller with probability 0.9; in all other states, it rewards the puller with probability 0.0001. Calculate the following:

   - If people tend to pull the middle arm (B) with a weight of 2 and each of the other arms with a weight of 1, what is the expected number of pulls between two rewards of the machine?
   - If the machine costs 1$ per try and gives a reward of 100$, what would be the profit rate of the owner?

Solution:

(a) The Markov Chain is as follows:

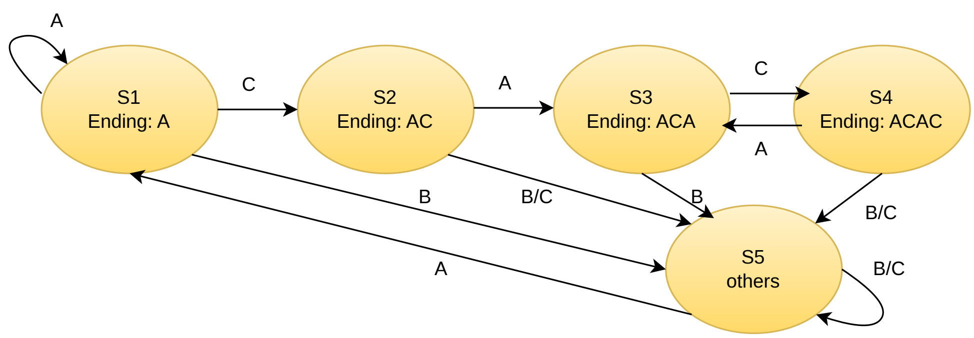

If we show the act of pulling an arm with $A_{A/B/C}$, and the act of the character being inserted by the character it self:

$$\begin{cases} p(A) = p(A_A)p(A|A_A) + p(A_B)p(A|A_B) + p(A_C)p(A|A_C) = 0.275 \\ p(B) = p(A_A)p(B|A_A) + p(A_B)p(B|A_B) + p(A_C)p(B|A_C) = 0.45 \\ p(C) = p(A_A)p(C|A_A) + p(A_B)p(C|A_B) + p(A_C)p(C|A_C) = 0.275 \end{cases} \tag{11}$$

So the transition matrix would be:

$$\begin{bmatrix} 0.275 & 0.275 & 0 & 0 & 0.45 \\ 0 & 0 & 0.275 & 0 & 0.725 \\ 0.275 & 0 & 0 & 0.275 & 0.45 \\ 0 & 0 & 0.275 & 0 & 0.725 \\ 0.275 & 0 & 0 & 0 & 0.725 \end{bmatrix}$$

for stationary probability we have $\pi \times T = \pi$, solving the equation in python ($T^T \times \pi^T = \pi^T$ so π is a right eigenvector of T^T with an eigenvalue of 1)
$=> \pi = [0.254, 0.070, 0.021, 0.006, 0.649]$
The probability of winning would be $\pi_4 \times 0.9 + (1 - \pi_4) \times 0.0001 = 0.0055$. The expected number of pulls between two wins would be $\frac{1}{p_{\text{win}}} = 181.8181$

(b) Profit rate is :cost $- p_{\text{win}} \times$ prize $= 1 - 0.0055 \times 100 = 0.45\$$ per pull.