CS 70 Discrete Mathematics and Probability Theory
Fall 2025 Sabin, Hug **Discussion 14A**

1 Markov Chains Intro I

Note 25

A **Markov chain** models an experiment with states, transitioning between states with some probability. A Markov chain is uniquely defined with the following variables:

- $\mathscr{X}$ is the set of possible states in the Markov chain. For this course, we'll only be working with Markov chains with a finite state space.
- X_n is a random variable denoting the state of the Markov chain at timestep n.
- P is the transition matrix. The element row i and column j in the matrix is defined as

  $$P(i,j) = \mathbb{P}[X_{n+1} = j \mid X_n = i].$$

  In particular, this is the probability that we transition from state i to state j.
- π_0 is the initial distribution; it is a row vector, where $\pi_0(i) = \mathbb{P}[X_0 = i]$. (Similarly, π_n is the distribution of states at timestep n; we have $\pi_n(i) = \mathbb{P}[X_n = i]$.)

Markov chains also have the **Markov property**:

$$\mathbb{P}[X_{n+1} = j \mid X_n = i, X_{n-1} = a_{n-1}, \ldots, X_0 = a_0] = \mathbb{P}[X_{n+1} = j \mid X_n = i].$$

That is, the next state depends only on the current state, and not on any prior states (this is also known as the memoryless property of Markov chains).

The **stationary distribution** (or the **invariant distribution**) of a Markov chain is the row vector π such that $\pi P = \pi$. (That is, transitioning does not change the distribution of states.)

A before B: Suppose we want to compute the probability of reaching state A before reaching state B. To compute this quantity, let $\alpha(i) = \mathbb{P}[A \text{ before } B \mid \text{at } i]$. Then, we have:

$$\begin{aligned}
\alpha(A) &= 1 \\
\alpha(B) &= 0 \\
\alpha(i) &= \sum_j P(i,j)\alpha(j)
\end{aligned}$$

Here, we use the law of total probability when computing $\alpha(i)$; we consider all possible transitions *out of* state i. These are called the **first step equations (FSE)**.

Hitting time: Suppose we want to compute the expected number of steps until you reach state A. To compute this quantity, let $\beta(i) = \mathbb{E}[\text{steps until } A \mid \text{at } i]$. Then, the first step equations become:

$$\begin{aligned}
\beta(A) &= 0 \\
\beta(i) &= 1 + \sum_j P(i,j)\beta(j)
\end{aligned}$$

Here, we use the law of total expectation when computing $\beta(i)$; we consider all possible transitions *out of* state i.

(a) Consider the transition matrix P of a Markov chain.

(i) Is it always true that every *row* of P sums to the same value? If so, state this value and briefly explain why this makes sense. If not, briefly explain why.

(ii) Is it always true that every *column* of P sums to the same value? If so, state this value and briefly explain why this makes sense. If not, briefly explain why.

(b) Compute $\mathbb{P}[X_1 = j]$ in terms of π_0 and P. Then, express your answer in matrix notation—that is, give an expression for the row vector π_1, where $\pi_1(j) = \mathbb{P}[X_1 = j]$. Generalize your answer to express π_n in matrix form in terms of n, π_0, and P.

(c) Note that we only need to provide $\mathscr{X}$, P, and π_0 in order to uniquely define a Markov chain; the random variables X_n are implicitly defined.

(i) Explain how you can compute the distributions of the random variables X_n for $n \geq 0$ using only these parameters. (*Hint*: Part (b) can be helpful.)

(ii) The Markov property is also implicit in this definition of a Markov chain. If the Markov property *does not hold*, are $\mathscr{X}$, P, and π_0 sufficient to compute the distributions of X_n for $n \geq 0$? Justify your answer.

2 Skipping Stones

Note 25

We consider a simple Markov chain model for skipping stones on a river, but with a twist: instead of trying to make the stone travel as far as possible, you want the stone to hit a target. Let the set of states be $\mathscr{X} = \{1,2,3,4,5\}$. State 3 represents the target, while states 4 and 5 indicate that you have overshot your target. Assume that from states 1 and 2, the stone is equally likely to skip forward one, two, or three steps forward. If the stone starts from state 1, compute the probability of reaching our target before overshooting, i.e. the probability of $\{3\}$ before $\{4,5\}$.

3 Consecutive Flips

Note 25

Suppose you are flipping a fair coin (one Head and one Tail) until you get the same side 3 times (Heads, Heads, Heads) or (Tails, Tails, Tails) in a row.

(a) Construct an Markov chain that describes the situation with a start state and end state.

(b) Given that you have flipped a (Tails, Heads) so far, what is the expected number of flips to see the same side three times?

(c) What is the expected number of flips to see the same side three times, beginning at the start state?

4 Can it be a Markov Chain?

Note 25

(a) A fly flies in a straight line in unit-length increments. Each second it moves to the left with probability 0.3, right with probability 0.3, and stays put with probability 0.4. There are two spiders at positions 1 and m and if the fly lands in either of those positions it is captured.

Given that the fly starts at state i, where $1 < i < m$, model this process as a Markov Chain. (Don't forget to specify the initial distribution!)

(b) Take the same scenario as in the previous part with $m = 4$. Let $Y_n = 0$ if at time n the fly is in position 1 or 2 and let $Y_n = 1$ if at time n the fly is in position 3 or 4.

Provide the state space for Y_n. Is the process Y_n a Markov chain?