Stochastic Processes
&
Markov Chain

Introduction to Stochastic Processes

What is a Stochastic Process?

We already know.

Simple Analogy:

Imagine repeatedly flipping a coin. The outcome of each flip is a random variable, and the sequence of these outcomes over time is a stochastic process.

Modeling a Coin Flip Sequence as a Stochastic Process

Let's define a sequence of random variables $\{X_n\}_{n \in T}$, where n represents the number of the coin flip (i.e., time).

Index Set: Currently our index set is $T = \mathbb{N} = \{1,2,3,...\}$, which represents the flip number.

State Space: The possible outcomes of a single coin flip are "Heads" or "Tails."
We can assign numerical values to these outcomes for mathematical convenience.

- Let's define the state space as $S = \{0,1\}$.
- $X_n = 1$ if the nth flip is Heads.
- $X_n = 0$ if the nth flip is Tails.

Random Variables: Each X_n is a random variable with a parameter p, where p is the probability of getting Heads (for a fair coin, $p = 0.5$).

- $P(X_n = 1) = p$ (the probability of getting Heads).
- $P(X_n = 0) = 1 - p$ (the probability of getting Tails).

The entire stochastic process is the sequence of these random variables,

$$\{X_1 , X_2 , X_3 , \ldots\} \;\; or \; \{X_n, n = 1,2,3, \ldots\}$$

For example, the sequence of outcomes "Heads, Tails, Heads, Heads, ..." corresponds to the sequence of values 1,0,1,1, ...

On a more abstract level... (only for modeling a coin flip sequence)

Let $\Omega = \{H, T\}$ be sample space for one coin flip and

$$X_n : \Omega^\infty \to \{0,1\}, \qquad n = 1,2,3, \dots,$$

where

$$Xn = \begin{cases} 1, & if\ the\ nth\ flip\ is\ Head, \\ 0, & if\ the\ nth\ flip\ is\ Tail. \end{cases}$$

The collection

$$\{X_n : n \in \mathbb{N}\}$$

is a stochastic process.

Markov Chain: Introduction & Motivation:

Scenario 1: The Predictable Library Student

Imagine a student, let's call him **Sikandar**. Every evening, Sikandar does one of two things:

(i) Study in the library or

(ii) Watch a movie at home.

Now, what does Sikandar do?

It's completely random! Sometimes he studies, sometimes he watches a movie.

There's no pattern. It's like flipping a coin every night.

How would we predict what Sikandar does tomorrow?"

What is the probability that Sikandar will STUDY?

What is the probability that Sikandar will WATCH THE MOVIE?

Scenario 2: The Habit-Driven Student

Now, let's meet **Anmol**. He also only does two things:

(i) Study in the library or

(ii) Watch a movie.

But Anmol is a creature of *habit*.

His decision for *tomorrow* depends *only* on what he did *today*.

Here's Anmol's rule:

- If Anmol studied today, there's a 90% chance he'll feel good and take a break to watch a movie tomorrow. Only a 10% chance he'll study again.
- If Anmol watched a movie today, he might feel a little guilty. So, there's an 80% chance he'll study tomorrow, and only a 20% chance he'll watch another movie.

New question:

Anmol watched a movie *today*. What is the *most likely* thing he will do *tomorrow*?

If he studies tomorrow, what is the *most likely* thing he will do the *day after tomorrow*?"

Now, we can make a prediction! Not with certainty, but with probability.
And the key insight is that we only needed to know one thing:
What Anmol did today.

What Anmol did last week or last month doesn't matter for predicting tomorrow.
The past is irrelevant, only the present matters.

Mathematical Model for Anmol's Weekly Habit

State Space: The set of all possible states of the system. We define:

State 0 = Study: "Studying in the library"

State 1 = Movie: "Watching a movie"

The state space is the set $S = \{0,1\}$.

We define the random variable X_n to represent Anmol's activity on day n:

$$X_n = \begin{cases} 0 & if\ studying\ on\ day\ \ n \\ 1 & if\ watching\ movie\ on\ day\ n \end{cases}$$

Transition Probabilities (P_{ij})

The probability of moving from state i today to state j tomorrow is denoted by P_{ij}.

From the problem description:

- $P_{00} = \ P(Study\ tomorrow \mid Study\ today) = 0.10$
- $P_{01} = \ P(Movie\ tomorrow \mid Study\ today) = 0.90$
- $P_{10} = \ P(Study\ tomorrow \mid Movie\ today) = 0.80$
- $P_{11} = \ P(Movie\ tomorrow \mid Movie\ today) = 0.20$

Transition Probability Matrix (P)

This matrix collects all transition probabilities.

The entry in the $(i+1)$th row and $(j+1)$th column is P_{ij}.

Important: Rows represent the current state. Columns represent the next state. Each row must sum to 1.

$$P = \begin{bmatrix} P_{00} & P_{01} \\ P_{10} & P_{11} \end{bmatrix} = \begin{bmatrix} 0.1 & 0.9 \\ 0.8 & 0.2 \end{bmatrix}$$

Interpretation of the matrix:

- **Row 1 (Currently Studying):** 10% chance to study again, 90% chance to watch a movie.
- **Row 2 (Currently Watching a Movie):** 80% chance to study, 20% chance to watch another movie.

State Transition Diagram

A visual representation of the Markov chain is often helpful. Arcs are labeled with transition probabilities.

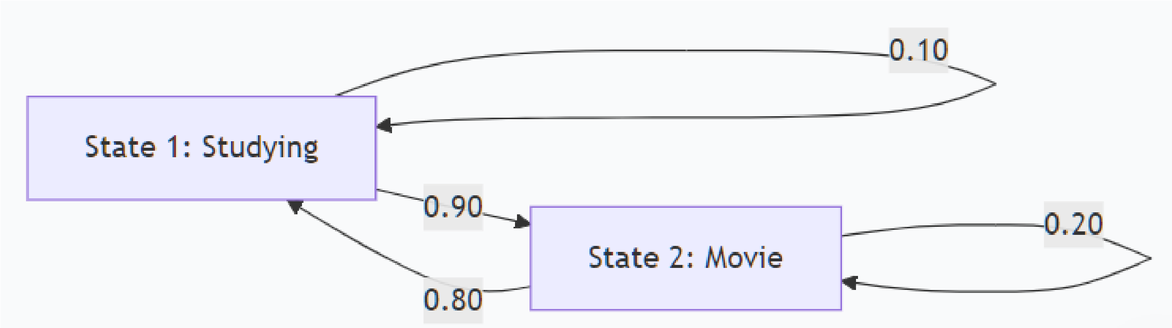

Example: Predicting the Future

Question: Suppose Anmol watched a movie on Monday ($n = 0$). What is the probability he is studying on Wednesday ($n = 2$)?

We need to find the 2-step transition probability $\boldsymbol{P_{10}^{(2)}}$.

Method 1: Probability Tree (Enumeration)

There are two paths from Movie (Mon) to Study (Wed):

1. **Movie (Mon) -> Study (Tue) -> Study (Wed):**
$$P_{10} \times P_{00} = (0.8) \times (0.1) = 0.08$$
2. **Movie (Mon) -> Movie (Tue) -> Study (Wed):**
$$P_{11} \times P_{10} = (0.2) \times (0.8) = 0.16$$

Total probability:

$$P_{10}^{(2)} = 0.08 + 0.16 = 0.24$$

Method 2: Matrix Multiplication (More powerful)

The 2-step transition matrix is $P^{(2)} = P^2 = P \times P$ (by Chapman Kolmogorov Theorem)

$$P^2 = \begin{bmatrix} 0.1 & 0.9 \\ 0.8 & 0.2 \end{bmatrix} \begin{bmatrix} 0.1 & 0.9 \\ 0.8 & 0.2 \end{bmatrix}$$

$$= \begin{bmatrix} 0.01 + 0.72 & 0.09 + 0.18 \\ 0.08 + 0.16 & 0.72 + 0.04 \end{bmatrix}$$

$$= \begin{bmatrix} 0.73 & 0.27 \\ 0.24 & 0.76 \end{bmatrix}$$

The entry in **Row 2, Column 1** of P^2 is 0.24. This confirms our calculation:

$$P_{10}^{(2)} = 0.24$$

Markov Chain: Formal Definition (Page 192)

Let $\{X_n : n = 0,1,2,\dots\}$ be a stochastic process taking values in a finite or countably infinite state space $\mathcal{S}$. If $X_n = i$, the process is said to be in state i at time n.

The process is called a **Markov chain** if it satisfies the **Markov property**:

For every $n \geq 0$ and for all states $i_0, i_1, \dots, i_{n-1}, i, j \in \mathcal{S}$,

$$P(X_{n+1} = j \mid X_n = i, X_{n-1} = i_{n-1}, \dots, X_1 = i_1, X_0 = i_0) = P(X_{n+1} = j \mid X_n = i).$$

That is, the conditional probability of the future state X_{n+1} depends only on the present state X_n and not on the past states $X_0, \dots, X_{n-1}$.

Assume

$$P_{ij} = P\{X_{n+1} = j \mid X_n = i\} \text{ for all } n \geq 0,$$

The quantities P_{ij} are called **transition probabilities**, and they satisfy

$$P_{ij} \geq 0 \text{ for all } i, j \in \mathcal{S}, \text{and} \sum_{j \in \mathcal{S}} P_{ij} = 1 \text{ for each } i \in \mathcal{S}.$$

Let $\mathbf{P}$ denote the matrix of one-step transition probabilities P_{ij}, so that

$$\mathbf{P} = \begin{bmatrix} P_{00} & P_{01} & P_{02} & \cdots \\ P_{10} & P_{11} & P_{12} & \cdots \\ \vdots & \vdots & \vdots & \\ P_{i0} & P_{i1} & P_{i2} & \cdots \\ \vdots & \vdots & \vdots & \end{bmatrix}$$

The matrix P is called **Transition Probability Matrix .**

Example (Forecasting the Weather)

Suppose that the chance of rain tomorrow depends on previous weather conditions only through whether or not it is raining today and not on past weather conditions. Suppose also that if it rains today, then it will rain tomorrow with probability α; and if it does not rain today, then it will rain tomorrow with probability β.

If we say that the process is in state 0 when it rains and state 1 when it does not rain, then the preceding is a two-state Markov chain whose transition probabilities are given by

$$\mathbf{P} = \begin{bmatrix} \alpha & 1-\alpha \\ \beta & 1-\beta \end{bmatrix}.$$

Ref: Introduction to probability models, Ross, S. M., Academic press, 2014.