Lesson 7. Introduction to Markov Chains

1 Overview

- In the Poisson process, the occurrences of new arrivals do not depend on the past history of the process
- This lesson: a stochastic process model
  - where the future depends on and can be predicted to some extent by past history
  - the effect of the past on the future is summarized by a **state** that changes probabilistically over time

2 The Case of the Random Behavior

Jungle.com is an online retailer that sells everything from books to toothbrushes. Their data analytics group is currently evaluating changes to Jungle.com's computer architecture, and needs a model that describes customer behavior. The group has identified four key types of customer transactions:

(1) visit the Jungle.com home page to start shopping ("log on"),
(2) fetch the main page of a product,
(3) fetch and read the reviews of a product, and
(4) finish shopping by checking out or closing the browser ("log off").

The data analytics group believes that the next transaction a customer requests is strongly influenced by the last (most recent) transaction requested, and not significantly influenced by anything else. For a given customer, let

- N be a random variable representing the next transaction a customer requests, and
- L be a random variable representing the last transaction requested.

Based on its substantial historical data, it has determined the following conditional pmfs:

a	1	2	3	4
$p_{N\|L=1}(a)$	0	0.95	0.01	0.04
$p_{N\|L=2}(a)$	0	0.27	0.63	0.10
$p_{N\|L=3}(a)$	0	0.36	0.40	0.24
$p_{N\|L=4}(a)$	0	0	0	1

3 Markov chains

- Discrete-time, discrete-state stochastic process $\{S_0, S_1, S_2, \dots\}$ where

$$S_n = \textbf{state} \text{ at time step } n$$
$$\mathcal{M} = \{1, \dots, m\} = \textbf{state space}\text{, or set of possible states}$$

- States evolve according to **one-step transition probabilities**:

- The initial state S_0 is determined by the **initial-state probabilities**:

- $\{S_0, S_1, S_2, \dots\}$ is a **Markov chain** if:

  - In other words, $\{S_0, S_1, S_2, \dots\}$ satisfies **the Markov property**: the conditional probability of the next state given the history of past states only depends on the last state
  - As a consequence:

Example 1. Recall that the performance-modeling group at Jungle.com believes that the next transaction a customer requests is essentially solely influenced by the last transaction requested. Compute the probability of the sequence of transactions 1, 2, 4.

- The sample paths of a Markov chain are completely characterized by a corresponding sequence of one-step transition probabilities and initial-state probabilities

- Note that a Markov chain is **time-stationary** because

  - In other words, the conditional probability of the next state given the last one does not depend on the number of time steps taken so far
  - As a consequence:

- Sometimes it is useful to study a variation of the Markov chain that is not time-stationary, where the one-step transition probabilities may change over time

  - These kinds of Markov chains are beyond the scope of this course

4 Representations of Markov chains

- We can organize the one-step transition probabilities into a **one-step transition matrix**:

$$\mathbf{P} = \begin{bmatrix} p_{11} & p_{12} & \cdots & p_{1m} \\ p_{21} & p_{22} & \cdots & p_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ p_{m1} & p_{m2} & \cdots & p_{mm} \end{bmatrix}$$

- We can also organize the initial-state probabilities into a **initial-state vector**:

$$\mathbf{p} = \begin{bmatrix} p_1 \\ p_2 \\ \vdots \\ p_m \end{bmatrix}$$

Example 2. Write the one-step transition matrix and initial-state vector for the Jungle.com Markov chain. Why do the rows of the one-step transition matrix sum up to 1?

- We can also draw a **transition probability diagram** where
  - each node represents a state of the system
  - a directed arc connects state i to state j if a one-step transition from i to j is possible
  - the one-step transition probability p_{ij} is written next to the arc from i to j

Example 3. Draw the transition probability diagram for the Jungle.com Markov chain.

5 Next time...

- How can we use the one-step transition matrix $\mathbf{P}$ and initial-state vector $\mathbf{p}$ to answer questions like:
  - Given that we are in state i right now, what is the probability we will be in state j after n time steps?
  - What is the unconditional probability we will be in state j after n time steps?