

Lesson 14. Introduction to Markov Chains

1 Overview

- We characterized the sample path of a stochastic process as a sequence of state changes (the S_n 's) occurring at random points in time (the T_n 's)
- In some settings, we care more about the state changes than the time at which the changes occur
- Today: a stochastic process model that focuses on the transitions between states

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").

A **session** is a sequence of transactions that begins with a log on (1) and ends with a log off (4).

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

Denote the corresponding conditional cdfs as $F_{N|L=1}$, $F_{N|L=2}$, $F_{N|L=3}$, $F_{N|L=4}$.

- Let's model the transitions between customer transaction types as a stochastic process
- State variables:

- System events:

- In this model, $n =$

- We don't keep track of event epochs, so our algorithm Simulation can be simplified to:

algorithm Simulation:

- | | |
|----------------------|-----------------------------------|
| 1: $n \leftarrow 0$ | (initialize system-event counter) |
| $e_0()$ | (execute initial system event) |
| 2: $e_1()$ | (update state of the system) |
| $n \leftarrow n + 1$ | (update system-event counter) |
| 3: go to line 2 | |

- Since $p_{N|L=4}(4) = 1$, once a sample path reaches state 4, it stays there
- The stochastic process model above – with the property that the probability distribution of the next state only depends on the last state – is called a **Markov chain**
- We can generalize this model so that the initial state is allowed to be a random variable with cdf F_{S_0} :

- In the Jungle.com case, we can think of S_0 as a degenerate random variable with $\Pr\{S_0 = 1\} = 1$ to model that a session always begins with a log on

3 Markov chains

- Discrete-time, discrete-state stochastic process $\{S_n; n = 0, 1, 2, \dots\}$
- State space $\mathcal{M} = \{1, \dots, m\}$
- States evolve according to the algorithmic model above
- $\{S_n; n = 0, 1, 2, \dots\}$ is a **Markov chain** if:

- In other words, $\{S_n; n = 0, 1, 2, \dots\}$ 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, 2, 4.

- A Markov chain is **time-stationary** if:

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

- In this course, we assume that Markov chains are time-stationary unless told otherwise
- For a time-stationary Markov chain, the **one-step probabilities** p_{ij} are defined as:

- The **initial-state probabilities** p_i are defined as:

Example 2. Assuming time-stationarity, express the probability of the sequence of transactions 1, 2, 2, 4 in terms of the one-step transition and initial-state probabilities.

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

4 Representations of Markov chains

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

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

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

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

- 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.

- a. Write the one-step transition matrix and initial-state vector for the Jungle.com Markov chain.
- b. Draw the transition probability diagram for the Jungle.com Markov chain.

5 Next time...

- Using 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?