Statistics Workshop Part II

Markov Chains

Silvia Heubach

Department of Mathematics

California State University Los Angeles

A **random variable** X is a function $X : S \longrightarrow \mathbb{R}$, i.e., each outcome in the sample space is associated with a numerical value.

Example 1: X = number of heads in three coin tosses

S	HHH	HHT HTH THH	TTH THT HTT	TTT
x	3	2	1	0
$p(x)$	$\frac{1}{8}$	$\frac{3}{8}$	$\frac{3}{8}$	$\frac{1}{8}$

A **stochastic process** $\{X_t\}$ is a random variable observed over time.

- **Discrete** or **continuous** time processes.
- A **Markov process** is a process where the next state of the system depends only on the current state (or more generally, on a finite number of immediate past states)
- A discrete Markov process is called a **Markov chain**

Ingredients for Markov Chain

- $X_0, X_1, X_2, \ldots, X_n, \ldots$ describes the **state** of the chain
- **State space** $= \{1, 2, 3, \ldots, N\}$
- (Stationary) **transition probabilities** $p_{i,j} = P(X_{n+1} = j | X_n = i)$
- Graphical display - states are nodes, transitions are edges
- **Transition matrix** $P = (p_{i,j})_{i,j=1}^{N}$. Note that all rows in P have to add to 1, i.e., P is a **stochastic** matrix.

Example 2: A gambler plays on one of four slot machines, each of which pays off a reward with probability 1/10. If the player wins on any machine, he continues to play the same machine; otherwise, he switches at random to one of the adjacent machines (where machine 1 and machine 4 are considered adjacent).

$S = \{1, 2, 3, 4\}$

$P(\text{stay}) = p_{i,i} = \mathbf{1/10}$ for all $i = 1, \ldots, 4$

$P(\text{switch to adjacent}) = p_{i,i+1(\text{mod}4)} = p_{i,i-1(\text{mod}4)} = (9/10)(1/2) = \mathbf{9/20}$

Transition Matrix and Graphical Representation

$$P = \begin{pmatrix} \frac{1}{10} & \frac{9}{20} & 0 & \frac{9}{20} \\ \frac{9}{20} & \frac{1}{10} & \frac{9}{20} & 0 \\ 0 & \frac{9}{20} & \frac{1}{10} & \frac{9}{20} \\ \frac{9}{20} & 0 & \frac{9}{20} & \frac{1}{10} \end{pmatrix}$$

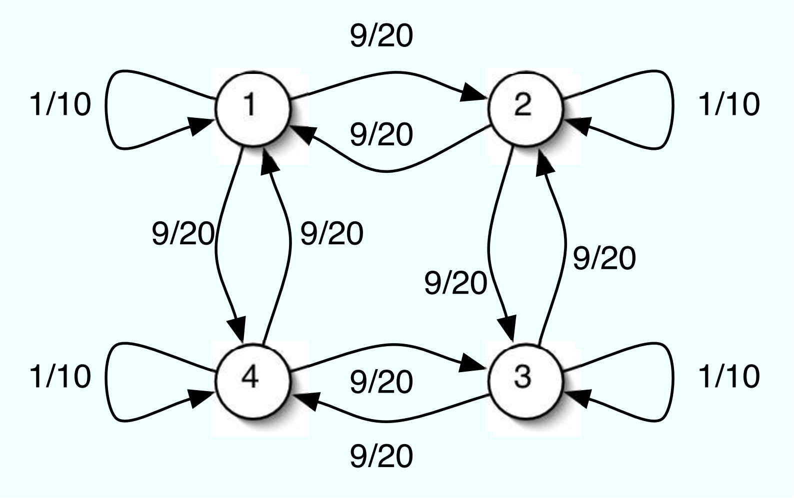

Questions

- What is the probability for a specific **path**, i.e., sequence of particular states visited?
- What is the probability for moving from state i to j in n steps?
- What is the probability to be in state i at time n?
- What is the percentage of time that state i occurs in the long run?

For Question 1, look at the sequence of states visited
$1, 2, 3, 3, 4, 1, 4$

We usually represent such a sequence in this form:

$$1 \to 2 \to 3 \to 3 \to 4 \to 1 \to 4$$

$$P(1 \to 2 \to 3 \to 3 \to 4 \to 1 \to 4)$$

$$= P(X_1 = 2|X_0 = 1)P(X_2 = 3|X_1 = 2)P(X_3 = 3|X_2 = 3)$$

$$\cdot P(X_4 = 4|X_3 = 3)P(X_5 = 1|X_4 = 4)P(X_6 = 4|X_5 = 1)$$

$$= p_{1,2} \cdot p_{2,3} \cdot p_{3,3} \cdot p_{3,4} \cdot p_{4,1} \cdot p_{1,4}$$

For Question 2, look at $P(X_2 = 1 | X_0 = 1) = p_{1,1}^{(2)}$

Possibilities:

$1 \rightarrow 1 \rightarrow 1$		$p_{1,1} \cdot p_{1,1}$
$1 \rightarrow 2 \rightarrow 1$		$p_{1,2} \cdot p_{2,1}$
$1 \rightarrow 3 \rightarrow 1$		$p_{1,3} \cdot p_{3,1}$
$1 \rightarrow 4 \rightarrow 1$		$p_{1,4} \cdot p_{4,1}$

$$\begin{aligned} \Rightarrow p_{1,1}^{(2)} &= p_{1,1} \cdot p_{1,1} + p_{1,2} \cdot p_{2,1} + p_{1,3} \cdot p_{3,1} + p_{1,4} \cdot p_{4,1} \\ &= \left(\frac{1}{10}\right)^2 + \left(\frac{9}{20}\right)^2 + 0 + \left(\frac{9}{20}\right)^2 \approx \mathbf{.415} \end{aligned}$$

Note: $p^{(2)}_{1,1} = \sum_{l=1}^{4} p_{1,l} \cdot p_{l,1} = (P \cdot P)_{1,1} = (P^2)_{1,1}$

$$\begin{pmatrix} \mathbf{p_{1,1}} & \mathbf{p_{1,2}} & \mathbf{p_{1,3}} & \mathbf{p_{1,4}} \\ p_{2,1} & p_{2,2} & p_{2,3} & p_{2,4} \\ p_{3,1} & p_{3,2} & p_{3,3} & p_{3,4} \\ p_{4,1} & p_{4,2} & p_{4,3} & p_{4,4} \end{pmatrix} \cdot \begin{pmatrix} \mathbf{p_{1,1}} & p_{1,2} & p_{1,3} & p_{1,4} \\ \mathbf{p_{2,1}} & p_{2,2} & p_{2,3} & p_{2,4} \\ \mathbf{p_{3,1}} & p_{3,2} & p_{3,3} & p_{3,4} \\ \mathbf{p_{4,1}} & p_{4,2} & p_{4,3} & p_{4,4} \end{pmatrix}$$

In general, the matrix P^k contains the probabilities for the **k-step transitions probabilities**

$$p^{(k)}_{i,j} = (P^k)_{i,j}$$

For Question 3, we use the Law of Total Probability, conditioning on the initial state. Let

$$\vec{p}^{(n)} = (P(X_n = 1), P(X_n = 2), \ldots, P(X_n = N))$$

be the vector of **state probabilities at time** n, and $\vec{p}^{(0)}$ be the vector of **initial state probabilities**. Then

$$P(X_n = i) = \sum_{l=1}^{N} P(X_n = i | X_0 = l) \cdot P(X_0 = l)$$

and we can compute the vector of state probabilities as $\vec{p}^{(n)} = \vec{p}^{(0)} \cdot P^n$.

$$P^1 = \begin{pmatrix} \frac{1}{10} & \frac{9}{20} & 0 & \frac{9}{20} \\ \frac{9}{20} & \frac{1}{10} & \frac{9}{20} & 0 \\ 0 & \frac{9}{20} & \frac{1}{10} & \frac{9}{20} \\ \frac{9}{20} & 0 & \frac{9}{20} & \frac{1}{10} \end{pmatrix}$$

$$P^5 = \begin{pmatrix} 0.17 & 0.33 & 0.17 & 0.33 \\ 0.33 & 0.17 & 0.33 & 0.17 \\ 0.17 & 0.33 & 0.17 & 0.33 \\ 0.33 & 0.17 & 0.33 & 0.17 \end{pmatrix}$$

$$P^{10} = \begin{pmatrix} 0.28 & 0.22 & 0.28 & 0.22 \\ 0.22 & 0.28 & 0.22 & 0.28 \\ 0.28 & 0.22 & 0.28 & 0.22 \\ 0.22 & 0.28 & 0.22 & 0.28 \end{pmatrix}$$

$$P^{15} = \begin{pmatrix} 0.24 & 0.26 & 0.24 & 0.26 \\ 0.26 & 0.24 & 0.26 & 0.24 \\ 0.24 & 0.26 & 0.24 & 0.26 \\ 0.26 & 0.24 & 0.26 & 0.24 \end{pmatrix}$$

$$P^{20} = \begin{pmatrix} 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \end{pmatrix}$$

Example 4: Assume the gambler in Example 3 always starts at slot machine 1. What is the probability that the player will move to slot machine 3 after the 10th game? (Note that X_i = machine played in game $i + 1$.)

$\vec{p}^{(0)} = (1, 0, 0, 0)$. Need to compute $\vec{p}^{(10)} = \vec{p}^{(0)} \cdot P^{10}$.

$$\vec{p}^{(0)} \cdot \begin{pmatrix} 0.28 & 0.22 & 0.28 & 0.22 \\ 0.22 & 0.28 & 0.22 & 0.28 \\ 0.28 & 0.22 & 0.28 & 0.22 \\ 0.22 & 0.28 & 0.22 & 0.28 \end{pmatrix} = (0.28\ 0.22\ \mathbf{0.28}\ 0.22)$$

For Question 4, we need the concept of a stationary distribution and a regular matrix.

A state probability vector is called a **stationary distribution** π if it satisfies $\pi \cdot \mathbf{P} = \pi$. The transition matrix P is **regular**, if, for some n, the entries of the matrix P^n are all positive (no zeros allowed).

Theorem: If P is a **regular** transition matrix, then the powers P^n approach a matrix all of whose rows are the same, and this row vector is the stationary distribution π.

$$P^1 = \begin{pmatrix} \frac{1}{10} & \frac{9}{20} & 0 & \frac{9}{20} \\ \frac{9}{20} & \frac{1}{10} & \frac{9}{20} & 0 \\ 0 & \frac{9}{20} & \frac{1}{10} & \frac{9}{20} \\ \frac{9}{20} & 0 & \frac{9}{20} & \frac{1}{10} \end{pmatrix}$$

$$P^5 = \begin{pmatrix} 0.17 & 0.33 & 0.17 & 0.33 \\ 0.33 & 0.17 & 0.33 & 0.17 \\ 0.17 & 0.33 & 0.17 & 0.33 \\ 0.33 & 0.17 & 0.33 & 0.17 \end{pmatrix}$$

$$P^{10} = \begin{pmatrix} 0.28 & 0.22 & 0.28 & 0.22 \\ 0.22 & 0.28 & 0.22 & 0.28 \\ 0.28 & 0.22 & 0.28 & 0.22 \\ 0.22 & 0.28 & 0.22 & 0.28 \end{pmatrix}$$

$$P^{15} = \begin{pmatrix} 0.24 & 0.26 & 0.24 & 0.26 \\ 0.26 & 0.24 & 0.26 & 0.24 \\ 0.24 & 0.26 & 0.24 & 0.26 \\ 0.26 & 0.24 & 0.26 & 0.24 \end{pmatrix}$$

$$P^{20} = \begin{pmatrix} 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \end{pmatrix}$$

- P is a regular transition matrix (P^5, for example, has all positive entries) $\Rightarrow$ stationary distribution exists
- P^{20} suggests that $\pi = (1/4, 1/4, 1/4, 1/4)$
- Check: $\pi \cdot P = \pi$?

Note: If no computer is available to compute matrix powers of P, one can solve for $\pi = (\pi_1, \pi_2, \ldots, \pi_N)$ by setting up the system of N equations corresponding to $\pi \cdot P = \pi$ and solve for π_i, $i = 1, \ldots N$. Due to redundancy in this system, we also need to use that $\sum_{i=1}^{N} \pi_i = 1$.

$$(\pi_1\ \pi_2\ \pi_3\ \pi_4) \cdot \begin{pmatrix} \frac{1}{10} & \frac{9}{20} & 0 & \frac{9}{20} \\ \frac{9}{20} & \frac{1}{10} & \frac{9}{20} & 0 \\ 0 & \frac{9}{20} & \frac{1}{10} & \frac{9}{20} \\ \frac{9}{20} & 0 & \frac{9}{20} & \frac{1}{10} \end{pmatrix} = (\pi_1\ \pi_2\ \pi_3\ \pi_4)$$

This gives the following system of equations:

$$\begin{aligned}
\frac{1}{10}\pi_1 + \frac{9}{20}\pi_2 + \frac{9}{20}\pi_4 &= \pi_1 \quad (1)\\
\frac{9}{20}\pi_1 + \frac{1}{10}\pi_2 + \frac{9}{20}\pi_3 &= \pi_2 \quad (2)\\
\frac{9}{20}\pi_2 + \frac{1}{10}\pi_3 + \frac{9}{20}\pi_4 &= \pi_3 \quad (3)\\
\frac{9}{20}\pi_1 + \frac{9}{20}\pi_3 + \frac{1}{10}\pi_4 &= \pi_4 \quad (4)\\
\pi_1 + \pi_2 + \pi_3 + \pi_4 &= 1 \quad (5)
\end{aligned}$$

Summary

The stationary distribution can be found in two ways:

1. Use software to compute high powers of the transition matrix. When all rows are identical, read off the stationary distribution. Verify by computing $\pi \cdot P$ and checking that the result equals π.

2. Solve the system of equations given by $\pi \cdot P = \pi$, $\sum_{i=1}^{N} \pi_i = 1$ for π.

The stationary distribution gives the long term probabilities for the Markov chain to be in a given state, and does NOT depend on the initial distribution $\vec{p}^{(0)}$ (see exercises).