

Chapter 15

Introduction to Markov Chains

Abstract The theory of Markov chains is rooted in the work of Russian mathematician Andrey Markov and has an extensive body of literature to establish its mathematical foundations. The availability of computing resources has recently made it possible to use Markov chains to analyze a variety of scientific data. Monte Carlo Markov chains are now one of the most popular methods of data analysis. This chapter presents the key mathematical properties of Markov chains, necessary to understand its implementation as Monte Carlo Markov chains.

15.1 Stochastic Processes and Markov Chains

This section presents key mathematical properties of Markov chains. The treatment is somewhat theoretical, but necessary to ensure that the applications we make to the analysis of data are consistent with the mathematics of Markov chains, which can be very complex. The goal is therefore that of defining and understanding a basic set of definitions and properties necessary to use Markov chains for the analysis of data, especially via the Monte Carlo simulations.

Markov chains are a specific type of stochastic processes, or sequence of random variables. A typical example of Markov chain is the so-called random walk, in which at each time step a person randomly takes a step either to the left, or to the right. As time progresses, the location of the person is the random variable of interest, and the collection of such random variables forms a Markov chain. The ultimate goal of a Markov chain is to determine the *stationary distribution* of the random variable. For the random walk, where we are interested in knowing the probability that at a given time the person is located n steps to the right or to the left of the starting point.

In the typical case of interest for the analysis of data, a dataset Z is fit to a parametric model. The goal is to create a Markov chain for each parameter of the model, in such a way that the stationary distribution for each parameter is the distribution function of the parameter. The chain will therefore result in the knowledge of the best-fit value of the parameter, and of confidence intervals, making use of the information provided by the dataset.

15.2 Mathematical Properties of Markov Chains

A stochastic process is defined as a sequence of variables X_t ,

$$\{X_t, \text{ for } t \in T\} \quad (15.1)$$

where t labels the sequence. The domain for the index t is indicated as T to signify “time.” The domain is usually a subset of the real numbers ($T \subset \mathbb{R}$) or of the natural numbers ($T \subset \mathbb{N}$). As time progresses, the random variables X_t change value, and the stochastic process describes this evolution.

A Markov chain is a particular stochastic process that satisfies the following properties:

1. The time domain is the natural numbers ($T \subset \mathbb{N}$), and each random variable X_t can have values in a countable set, e.g., the natural numbers or even an n -dimensional space (\mathbb{N}^n), but not real numbers (\mathbb{R}^n). A typical example of a Markov chain is one in which $X_i = n$, where both i (the time index) and n (the value of the random variable) are natural numbers. Therefore a Markov chain takes the form of

$$X_1 \rightarrow X_2 \rightarrow X_3 \rightarrow \dots \rightarrow X_n \rightarrow \dots$$

The random variable X_i describes the state of the system at time $t = i$. The fact that Markov chains must be defined by way of countable sets may appear an insurmountable restriction, since it would appear that the natural domain for an n -parameter space is \mathbb{R}^n . While a formal extension of Markov chains to \mathbb{R}^n is also possible, this is not a complication for any practical application, since any parameter space can be somehow “binned” into a finite number of states. For example, the position of the person in a random walk was “binned” into a number of finite (or infinite but countable) positions, and a similar process can be applied to virtually any parameter of interest for a given model. This means that the variable under consideration can occupy one of a countable multitude of states $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_n, \dots$, and the random variable X_i identifies the state of the system at time step i , $X_i = \varepsilon_n$.

2. A far more important property that makes a stochastic process a Markov chain is the fact that subsequent steps in the chain are only dependent on the current state of the chain, and not on any of its previous history. This “short memory” property is known as the *Markovian property*, and it is the key into the construction of Markov chains for the purpose of data analysis. In mathematical terms, given the present time $t = n$, the future state of the chain at $t = n + 1$ (X_{n+1}) depends only on the present time (X_n), but not on past history. Much of the efforts in the construction of a Monte Carlo Markov chain lies in the identification

of a *transition probability* from state ε_i to state ε_j between consecutive time steps,

$$p_{ij} = P(X_{n+1} = \varepsilon_j / X_n = \varepsilon_i). \quad (15.2)$$

A Markov chain requires that this probability be time-independent, and therefore a Markov chain has the property of time homogeneity. In Chap. 16 we will see how the transition probability takes into account the likelihood of the data Z with the model.

The two properties described above result in the fact that Markov chain is a sequence of states determined by transition probabilities p_{ij} (also referred to as *transition kernel*) that are fixed in time. The ultimate goal is to determine the probability to find the system in each of the allowed states. With an eye towards future applications for the analysis of data, each state may represent values of one or many parameters, and therefore a Markov chain makes it possible to reconstruct the probability distribution of the parameters.

Example 15.1 (Random Walk) The *random walk* is a Markov chain that represents the location of a person who randomly takes a step of unit length forward with probability p , or a step backward with probability $q = 1 - p$ (typically $p = q = 1/2$). The state of the system is defined by the location i at which the person find itself at time $t = n$,

$$X_n = \{\text{Location } i \text{ along the } \mathbb{N}^+ \text{ axis}\}$$

where \mathbb{N}^+ indicates all positive and negative integers. For this chain, the time domain is the set of positive numbers ($T = \mathbb{N}$), and the position can be any negative or positive integer (\mathbb{N}^+). The transition probability describes the fact that the person can only take either a step forward or backward:

$$p_{ij} = \begin{cases} p & \text{if } j = i + 1, \text{ or move forward} \\ q & \text{if } j = i - 1, \text{ or move backward} \\ 0 & \text{otherwise.} \end{cases} \quad (15.3)$$

◇

The chain satisfies the Markovian property, since the transition probability depends only on its present position, and not on previous history.

Example 15.2 Another case of a Markov chain is a simple model of diffusion, known as the *Ehrenfest chain*. Consider two boxes with a total of m balls. At each time step, one selects a ball at random from either box, and replaces it in the other box. The state of the system can be defined via the random variable

$$X_n = \{\text{Number of balls in the first box}\}.$$

The random variable can have only a finite number of values $(0, 1, \dots, m)$. At each time step, the transition probability is

$$P_{ij} = \begin{cases} \frac{m-i}{m} & \text{if } j = i + 1 \text{ (box had } i \text{ balls, now has } i + 1) \\ \frac{i}{m} & \text{if } j = i - 1 \text{ (box had } i \text{ balls, now has } i - 1). \end{cases} \quad (15.4)$$

For example, in the first case it means that we chose one of $m - i$ balls from the second box. The transition probabilities depend only on the number of balls in the first box at any given time, and are completely independent of how the box came to have that many balls. This chain therefore satisfies the Markovian property. \diamond

15.3 Recurrent and Transient States

We are interested in knowing how often a state is visited by the chain and, in particular, whether a given state can be visited infinitely often. Assume that the system is initially in state ε_i . We define u_k the probability that the system returns to the initial state in *exactly* k time steps, and v_n the probability that the system returns to the initial state at time n , with the possibility that it may have returned there other times prior to n . Clearly, it is true that $v_n \geq u_n$.

To determine whether a state is recurrent or transient, we define

$$u \equiv \sum_{n=1}^{\infty} u_n \quad (15.5)$$

as the probability of the system returning the initial state ε_i for the first time at some time n . The state can be classified as *recurrent* or *transient* according to the probability of returning to that state:

$$\begin{cases} u = 1 & \text{state is recurrent;} \\ u < 1 & \text{state is transient.} \end{cases} \quad (15.6)$$

Therefore a recurrent state is one that will certainly be visited again by the chain. Notice that no indication is given as to the time at which the system will return to the initial state.

We also state a few theorems that are relevant to the understanding of recurrent states. Proofs of these theorems can be found, for example, in the textbook by Ross [38] or other books on stochastic processes, and are not reported here.

Theorem 15.1 *With v_n the probability that the system returns to a state ε_i at time n ,*

$$\text{state } \varepsilon_i \text{ is recurrent} \iff \sum_{n=1}^{\infty} v_n = \infty. \quad (15.7)$$

This theorem states that, if the system does return to a given state, then it will do so infinitely often. Also, since this is a necessary and sufficient condition, any transient state will not be visited by the chain an infinite number of times. This means that transient states will not be visited any more after a given time, i.e., they are only visited during an initial period. The fact that recurrent states are visited infinitely often means that it is possible to construct a sample distribution function for recurrent states with a precision that is function of the length of the chain. No information is, however, provided on the timing of the return to a recurrent state.

We also introduce the definition of *accessible* states: a state ε_j is said to be accessible from state ε_i if $p_{ij}(m) > 0$ for some natural number m , meaning that there is a non-zero probability of reaching this state from another state in m time steps. The following theorems establish properties of accessible states, and how the property of accessibility relates to that of recurrence.

Theorem 15.2 *If a state ε_j is accessible from a recurrent state ε_i , then ε_j is also recurrent, and ε_i is accessible from ε_j .*

This theorem states that once the system reaches a recurrent state, the states visited previously by the chain must also be recurrent, and therefore will be visited again infinitely often. This means that recurrent states form a network, or class, of states that share the property of recurrence, and these are the states that the chain will sample over and over again as function of time.

Theorem 15.3 *If a Markov chain has a finite number of states, then each state is accessible from any other state, and all states are recurrent.*

This theorem ensures that all states in a finite chain will be visited infinitely often, and therefore the chain will sample all states as function of time. This property is of special relevance for Monte Carlo Markov chain methods in which the states of the chain are possible values of the parameters. As the chain progresses, all values of the parameters are accessible, and will be visited in proportion of the posterior distribution of the parameters.

Example 15.3 (Recurrence of States of the Random Walk) Consider the random walk with transition probabilities given by (15.3). We want to determine whether the initial state of the chain is a recurrent or a transient state for the chain. The probability of returning to the initial state in k steps is clearly given by the binomial distribution,

$$p_{ii}(k) = \begin{cases} 0 & \text{if } k \text{ is odd} \\ C(n, k)p^n q^n & \text{if } k = 2n \text{ is even} \end{cases} \quad (15.8)$$

where

$$C(n, k) = \binom{k}{n} = \frac{k!}{(k-n)!n!} \quad (15.9)$$

is the number of combinations [of n successes out of $k = 2n$ tries, see (3.3)]. Using Stirling's approximation for the factorial function in the binomial coefficient,

$$n! \simeq \sqrt{2\pi n} n^n e^{-n},$$

the probability to return at time $k = 2n$ to the initial state becomes

$$v_k = p_{ii}(k) = \binom{2n}{n} p^n q^n = \frac{(2n)!}{(n!)^2} p^n q^n \simeq \frac{\sqrt{4\pi n} (2n)^{2n} e^{-2n}}{2\pi n n^{2n} e^{-2n}} p^n q^n = \frac{(4pq)^n}{\sqrt{\pi n}}$$

which holds only for k even.

This equation can be used in conjunction with Theorem 15.1 to see if the initial state is transient or recurrent. Consider the series

$$\sum_{n=1}^{\infty} v_n = \sum_{n=1}^{\infty} \frac{1}{\sqrt{\pi n}} (4pq)^n.$$

According to Theorem 15.1, the divergence of this series is a necessary and sufficient condition to prove that the initial state is recurrent.

(a) $p \neq q$. In this case, $x = 4pq < 1$ and

$$\sum_{n=1}^{\infty} \frac{1}{\sqrt{\pi n}} (4pq)^n < \sum_{n=1}^{\infty} x^n = \frac{x}{1-x};$$

since $x < 1$, the series converges and therefore the state is transient. This means that the system may return to the initial state, but only for a finite number of times, even after an infinite time. Notice that as time progresses the state of the system will drift in the direction that has a probability $> 1/2$.

(b) $p = q = 1/2$, thus $4pq = 1$. The series becomes

$$\frac{1}{\sqrt{\pi}} \sum_{n=1}^{\infty} \frac{1}{n^{1/2}}. \quad (15.10)$$

It can be shown (see Problem 15.2) that this series diverges, and therefore a random walk with the same probability of taking a step to the left or to the right will return to the origin infinitely often. \diamond

15.4 Limiting Probabilities and Stationary Distribution

The ultimate goal of a Markov chain is to calculate the probability that a system occupies a given state ε_i after a large number n of steps. This probability is called the *limiting probability*. According to the frequentist approach defined in (1.2), it is given by

$$p_j^* = \lim_{n \rightarrow \infty} p_j(n), \quad (15.11)$$

where $p_j(n)$ is the probability of the system to be found in state ε_j at time $t = n$. With the aid of the total probability theorem, the probability of the system to be in state ε_j at time $t = n$ is

$$p_j(n) = \sum_k p_k(n-1)p_{kj}. \quad (15.12)$$

In fact $p_k(n-1)$ represents the probability of being in state ε_k at time $n-1$, and the set of probabilities $p_k(n-1)$ forms a set of mutually exclusive events which encompasses all possible outcomes, with the index k running over all possible states. This formula can be used to calculate recursively the probability $p_j(n)$ using the probability at the previous step and the transition probabilities p_{kj} , which do not vary with time.

Equation (15.12) can be written in a different form if the system is known to be in state ε_i at an initial time $t = 0$:

$$p_{ij}(n) = P(X_n = \varepsilon_j) = \sum_k p_{ik}(n-1)p_{kj} \quad (15.13)$$

where $p_{ij}(n)$ is the probability of the system going from state ε_i to ε_j in n time steps.

The probabilities $p_j(n)$ and $p_{ij}(n)$ change as the chain progresses. The limiting probabilities p_j^* , on the other hand, are independent of time, and they form the *stationary* distribution of the chain. General properties for the stationary distribution can be given for Markov chains that have certain specific properties. In the following we introduce additional definitions that are useful to characterize Markov chains, and to determine the stationary distribution of the chain.

A number of states that are accessible from each other, meaning there is a non-zero probability to reach one state from the other ($p_{ij} > 0$), are said to *communicate*, and all states that communicate are part of the same *class*. The property of communication (\leftrightarrow) is an equivalence relation, meaning that it obeys the following three properties:

- (a) The reflexive property: $i \leftrightarrow i$;
- (b) The symmetric property: if $i \leftrightarrow j$, then $j \leftrightarrow i$; and
- (c) The transitive property: if $i \leftrightarrow j$ and $j \leftrightarrow k$, then $i \leftrightarrow k$. Therefore, each class is separate from any other class of the same chain. A chain is said to be *irreducible* if it has only one class, and thus all states communicate with each other.

Another property of Markov chains is periodicity. A state is said to be *periodic* with period T if $p_{ii}(n) = 0$ when n is not divisible by T , and T is the largest such integer with this property. This means that the return to a given state must occur in multiples of T time steps. A chain is said to be *aperiodic* if $T = 1$, and return to a given state can occur at any time. It can be shown that all states in a class share the same period.

The uniqueness of the stationary distribution and an equation that can be used to determine it are established by the following theorems.

Theorem 15.4 *An irreducible aperiodic Markov chain belongs to either of the following two classes:*

1. All states are positive recurrent. In this case, $p_i^* = \pi_i$ is the stationary distribution, and this distribution is unique.
2. All states are transient or null recurrent; in this case, there is no stationary distribution.

This theorem establishes that a “well behaved” Markov chain, i.e., one with positive recurrent states, does have a stationary distribution, and that this distribution is unique. Positive recurrent states, defined in Sect. 15.3, are those for which the expected time to return to the same state is finite, while the time to return to a transient or null recurrent state is infinite. This theorem also ensures that, regardless of the starting point of the chain, the same stationary distribution will eventually be reached.

Theorem 15.5 *The limiting probabilities are the solution of the system of linear equations*

$$p_j^* = \sum_{i=1}^N p_i^* p_{ij}. \quad (15.14)$$

Proof According to the recursion formula (15.12),

$$p_j(n) = \sum_{i=1}^N p_i(n-1)p_{ij}. \quad (15.15)$$

Therefore the result follows by taking the limit $n \rightarrow \infty$ of the above equation. \square

If we consider a chain with a probability distribution at a time t_0 that satisfies

$$p_j(t_0) = \sum_{i=1}^N p_i(t_0)p_{ij}, \tag{15.16}$$

then the probability distribution of the states satisfies (15.14), and the chain has reached the stationary distribution $p_i(n) = p_i^*$. Theorem 15.5 guarantees that, from that point on, the chain will maintain its stationary distribution. The importance of a stationary distribution is that, as time elapses, the chain *samples* this distribution. The sample distribution of the chain, e.g., a histogram plot of the occurrence of each state, can therefore be used as an approximation of the posterior distribution.

Example 15.4 (Stationary Distribution of the Ehrenfest Chain) We want to find a distribution function p_j^* that is the stationary distribution of the Ehrenfest chain. This case is of interest because the finite number of states makes the calculation of the stationary distribution easier to achieve analytically. The condition for a stationary distribution is

$$p_j^* = \sum_{i=1}^N p_i^* p_{ij}$$

where N is the number of states of the chain. The condition can also be written in matrix notation. Recall that the transition probabilities for the Ehrenfest chain are

$$p_{ij} = \begin{cases} \frac{m-i}{m} & \text{if } j = i + 1 \\ \frac{i}{m} & \text{if } j = i - 1, \end{cases}$$

and they can be written as a transition matrix \mathbf{P}

$$\mathbf{P} = [p_{ij}] = \begin{bmatrix} 0 & 1 & 0 & 0 & \dots & 0 & 0 \\ \frac{1}{m} & 0 & \frac{m-1}{m} & 0 & \dots & 0 & 0 \\ 0 & \frac{2}{m} & 0 & \frac{m-2}{m} & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & 0 & \dots & 1 & 0 \end{bmatrix}. \tag{15.17}$$

Notice that the sum of each line is one, since

$$\sum_j p_{ij} = 1$$

is the probability of going from state ε_i to any state ε_j . In (15.17) you can regard the vertical index to be $i = 0, \dots, m$, and the horizontal index $j = 0, \dots, m$.

The way in which we typically use (15.14) is simply to verify whether a distribution is the stationary distribution of the chain. In the case of the Ehrenfest chain, we try the binomial distribution as the stationary distribution,

$$p_i = \binom{m}{i} p^i q^{m-i} \quad i = 0, \dots, m,$$

in which p and q represent the probability of finding a ball in either box. At equilibrium one expects $p = q = 1/2$, since even an initially uneven distribution of balls between the two boxes should result in an even distribution at later times. It is therefore reasonable to expect that the probability of having i balls in the first box, out of a total of m , is equivalent to that of i positive outcomes in a binary experiment.

To prove this hypothesis, consider $p = [p_0, p_1, \dots, p_m]$ as a row vector of dimension $m + 1$, and verify the equation

$$p = pP, \tag{15.18}$$

which is the matrix notation for the condition of a stationary distribution. For the Ehrenfest chain, this condition is

$$[p_0, p_1, \dots, p_m] = [p_0, p_1, \dots, p_m] \begin{bmatrix} 0 & 1 & 0 & 0 & \dots & 0 & 0 \\ \frac{1}{m} & 0 & \frac{m-1}{m} & 0 & \dots & 0 & 0 \\ 0 & \frac{2}{m} & 0 & \frac{m-2}{m} & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & 0 & \dots & 1 & 0 \end{bmatrix}.$$

For a given state i , only two terms (at most) contribute to the sum,

$$p_i = p_{i-1}p_{i-1,i} + p_{i+1}p_{i+1,i}. \tag{15.19}$$

From this we can prove that the $p = q = 1/2$ binomial is the stationary distribution of the Ehrenfest chain (see Problem 15.1). \diamond

Summary of Key Concepts for this Chapter

- *Markov chain*: A stochastic process or sequence of random variables as function of an integer time variable.
- *Markovian property*: It is the key property of Markov chains, stating that the state of the system at a given time depends only on the state at the previous time step.
- *Recurrent and transient state*: A recurrent state occurs infinitely often while a transient state only occurs a finite number of times in the Markov chain.
- *Stationary distribution*: It is the asymptotic distribution of each variable, obtained after a large number of time steps of the Markov chain. When the variable represents a model parameter, the stationary distribution is the posterior distribution of the parameter.

Problems

15.1 Consider the Markov chain for the Ehrenfest chain described in Example 15.4. Show that the stationary distribution is the binomial with $p = q = 1/2$.

15.2 Show that the random walk with $p = q = 1/2$ (15.10) returns to the origin infinitely often, and therefore the origin is a recurrent state of the chain.

15.3 For the random walk with $p \neq q$, show that the origin is a transient state.

15.4 Assume that the diffusion model of Example 15.2 is modified in such a way that at each time step one has the option to choose one box at random from which to replace a ball to the other box.

- (a) Determine the transition probabilities p_{ij} for this process.
- (b) Determine whether this process is a Markov chain.

15.5 Using the model of diffusion of Problem 15.4, determine if the binomial distribution with $p = q = 1/2$ is the stationary distribution.