Reversible Markov Chains

EECS 126 (UC Berkeley)

Spring 2019

1 Reversibility

Consider an irreducible Markov chain $(X_n)_{n\in\mathbb{N}}$ on the finite state space $\mathcal{X}$ with transition probability matrix P. When is it the case that the Markov chain "looks the same" regardless of whether we run it forwards in time or backwards in time? Formally, fix a positive integer N and define the **reversed chain** $Y_n := X_{N-n}$ for $n = 0, 1, \dots, N$. Then, $(Y_0, \dots, Y_N) = (X_N, \dots, X_0)$, so $(Y_n)_{n=0}^N$ is the sequence of states we observe if, starting at time N, we run the original Markov chain "backwards". To justify the name "reversed chain", we prove:

Theorem 1. *If the irreducible Markov chain $(X_n)_{n\in\mathbb{N}}$ is started from the stationary distribution π, then the reversed chain $(Y_n)_{n=0}^N$ is an irreducible Markov chain with transition probabilities $\hat{P}(x,y) = \pi(y)P(y,x)/\pi(x)$ for $x, y \in \mathcal{X}$. The stationary distribution for the reversed chain is also π.*

Proof. For a positive integer $k < N$ and a feasible sequence of states $x_0, x_1, \dots, x_k, x_{k+1} \in \mathcal{X}$,

$$\begin{aligned}
&\mathbb{P}(Y_{k+1} = x_{k+1} \mid Y_k = x_k, \dots, Y_1 = x_1, Y_0 = x_0) \\
&\qquad = \mathbb{P}(X_{N-k-1} = x_{k+1} \mid X_{N-k} = x_k, \dots, X_{N-1} = x_1, X_N = x_0)
\end{aligned}$$

(now use the "backwards Markov property")

$$\begin{aligned}
&= \mathbb{P}(X_{N-k-1} = x_{k+1} \mid X_{N-k} = x_k) = \frac{\mathbb{P}(X_{N-k-1} = x_{k+1}, X_{N-k} = x_k)}{\mathbb{P}(X_{N-k} = x_k)} \\
&= \frac{\pi(x_{k+1})P(x_{k+1}, x_k)}{\pi(x_k)}.
\end{aligned}$$

Therefore, $(Y_n)_{n=0}^N$ is a Markov chain and the transition probabilities are $\hat{P}(x,y) := \pi(y)P(y,x)/\pi(x)$ for $x, y \in \mathcal{X}$. Irreducibility of $(Y_n)_{n=0}^N$ follows from irreducibility of the original chain (you are encouraged to think about why this is true). Finally, since $\mathbb{P}(X_n = x) = \pi(x)$ for all $n \in \mathbb{N}$ (since the original chain is started from stationarity), then $\mathbb{P}(Y_n = x) = \pi(x)$ for all $n = 0, 1, \dots, N$, which implies that π is the stationary distribution for the reversed chain. $\square$

Now, to answer the question we posed above: when does the reversed chain look the same as the original chain? We need the transition probabilities to be the same in both chains, i.e., $\hat{P}(x, y) = P(x, y)$ for all $x, y \in \mathcal{X}$, or equivalently,

$$\pi(x)P(x, y) = \pi(y)P(y, x) \qquad \text{for all } x, y \in \mathcal{X}. \tag{1}$$

A Markov chain whose stationary distribution π and transition probability matrix P satisfy (1) is called **reversible**.

Perhaps surprisingly, the notion of reversibility is more than just a mathematical curiosity. Reversibility finds applications in the design of queueing systems and sampling algorithms; we will soon discuss the latter in more detail. We give an example of the former:

Example 1. Consider a discrete-time queue: at each $n \in \mathbb{N}$, a customer arrives with probability p (this is called a *Bernoulli process*), and if there are any customers in the queue, then one of the customers is served with probability q. The arrivals and services are all assumed to be mutually independent. Then, the length of the queue is a Markov chain, and in fact it turns out to be reversible. Notice though that when we look at the queue backwards in time, then the customer *departures* become *arrivals*, which implies that the departure process is also a Bernoulli process!

1.1 Detailed Balance

The equations (1) are also called the **detailed balance equations**. In general, the condition for stationarity is $\pi(y) = \sum_{x \in \mathcal{X}} \pi(x)P(x, y)$ for all $y \in \mathcal{X}$. Imagine that a total mass of $\pi(y)$ sits at state y, and when the Markov chain undergoes a transition, the total mass of $\pi(y)$ leaves state y and is divided up among the neighbors of y according to the transition probabilities $P(y, \cdot)$. The balance equations then express the *global* condition that the total mass leaving state y equals the sum of the mass entering y. On the other hand, the detailed balance equations reflect a *local* condition that the mass exchanged along each edge (x, y) is balanced; this is a *stronger* condition than global balance.

Exercise 1. Consider an irreducible Markov chain on the finite state space $\mathcal{X}$ with transition probability matrix P.

1. Suppose that $\tilde{P}$ is another transition probability matrix of the same dimensions as P (the rows sum to 1). Show that if π is a probability distribution on $\mathcal{X}$ which satisfies $\pi(x)P(x,y) = \pi(y)\tilde{P}(y,x)$ for all $x, y \in \mathcal{X}$, then π is the stationary distribution and $\tilde{P}$ is the transition probability matrix of the reversed chain.

2. In particular, if π satisfies the detailed balance equations (1), then π is the stationary distribution.

To summarize the situation, the detailed balance equations are *sufficient* for stationarity but not *necessary*; there exist Markov chains whose stationary distributions do not satisfy detailed balance. [1] However, it is useful to try detailed balance first, since solving (1) is usually easier than solving the general balance equations. To take full advantage of the power of this tool, it is helpful to know the following result:

Exercise 2. The graph associated with a Markov chain is formed by taking the transition diagram of the chain, forgetting the directions of all of the edges, removing multiple edges, and removing self-loops. If the graph of a finite-state irreducible Markov chain is a tree, then the stationary distribution of the Markov chain satisfies detailed balance.

In particular, Markov chains which look like a *line* satisfy detailed balance: for example, a random walk which can only move to the immediate left or to the immediate right. This includes the queue in Example 1.

Example 2 (Metropolis-Hastings)**.** Often in applications, we encounter probability distributions that are computationally intractable, but if we can design algorithms which produce *samples* from the desired distribution, then we can infer important properties about the distribution (such as the expectation or the mode). **Monte Carlo Markov Chain (MCMC)** refers to class of algorithms which produce samples from a distribution by designing a Markov chain where the given distribution is stationary for the chain. After performing transitions until the chain reaches stationarity, we can then take samples from the chain to approximate the given distribution.

[1] If you are interested, necessary and sufficient conditions for detailed balance are given by the *Kolmogorov cycle condition*.

An important MCMC algorithm, called the **Metropolis-Hastings algorithm**, only requires knowledge about the desired distribution *up to a constant factor*. (The constant factor is important—it is often the reason why the probability distribution is computationally intractable!) Crucially, the Metropolis-Hastings chain satisfies detailed balance.