

Markov Chains Handout for Stat 110

Prof. Joe Blitzstein (Harvard Statistics Department)

1 Introduction

Markov chains were first introduced in 1906 by Andrey Markov, with the goal of showing that the Law of Large Numbers does not necessarily require the random variables to be independent. Since then, they have become extremely important in a huge number of fields such as biology, game theory, finance, machine learning, and statistical physics. They are also very widely used for simulations of complex distributions, via algorithms known as MCMC (Markov Chain Monte Carlo).

To see where the Markov model comes from, consider first an i.i.d. sequence of random variables $X_0, X_1, \dots, X_n, \dots$ where we think of n as time. Independence is a very strong assumption: it means that the X_j 's provide no information about each other. At the other extreme, allowing general interactions between the X_j 's makes it very difficult to compute even basic things. Markov chains are a happy medium between complete independence and complete dependence.

The space on which a Markov process “lives” can be either discrete or continuous, and time can be either discrete or continuous. In Stat 110, we will focus on Markov chains X_0, X_1, X_2, \dots in discrete space and time (continuous time would be a process X_t defined for all real $t \geq 0$). Most of the ideas can be extended to the other cases. Specifically, we will assume that X_n takes values in a finite set (the *state space*), which we usually take to be $\{1, 2, \dots, M\}$ (or $\{0, 1, \dots, M\}$ if it is more convenient). In Stat 110, we will always assume that our Markov chains are on finite state spaces.

Definition 1. A sequence of random variables X_0, X_1, X_2, \dots taking values in the *state space* $\{1, \dots, M\}$ is called a *Markov chain* if there is an M by M matrix $Q = (q_{ij})$ such that for any $n \geq 0$,

$$P(X_{n+1} = j | X_n = i, X_{n-1} = i_{n-1}, \dots, X_0 = i_0) = P(X_{n+1} = j | X_n = i) = q_{ij}.$$

The matrix Q is called the *transition matrix* of the chain, and q_{ij} is the *transition probability* from i to j .

This says that given the history $X_0, X_1, X_2, \dots, X_n$, only the *most recent* term, X_n , matters for predicting X_{n+1} . If we think of time n as the present, times before n as the past, and times after n as the future, the Markov property says that given the present, the past and future are conditionally independent.

The Markov assumption greatly simplifies computations of conditional probability: instead of having to condition on the entire past, we only need to condition on the *most recent* value.

2 Transition Matrix

The transition probability q_{ij} specifies the probability of going from state i to state j in one step. The *transition matrix* of the chain is the $M \times M$ matrix $Q = (q_{ij})$. Note that Q is a nonnegative matrix in which each row sums to 1.

Definition 2. Let $q_{ij}^{(n)}$ be the n -step transition probability, i.e., the probability of being at j exactly n steps after being at i :

$$q_{ij}^{(n)} = P(X_n = j | X_0 = i).$$

Note that

$$q_{ij}^{(2)} = \sum_k q_{ik} q_{kj}$$

since to get from i to j in two steps, the chain must go from i to some intermediary state k , and then from k to j (these transitions are independent because of the Markov property). So the matrix Q^2 gives the 2-step transition probabilities. Similarly (by induction), powers of the transition matrix give the n -step transition probabilities:

$q_{ij}^{(n)}$ is the (i, j) entry of Q^n .

Example. Figure 1 shows an example of a Markov chain with 4 states. The chain

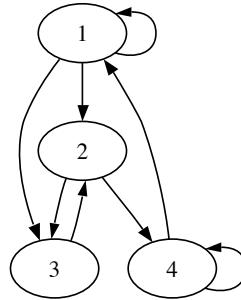


Figure 1: A Markov Chain with 4 Recurrent States

can be visualized by thinking of a particle wandering around from state to state,

randomly choosing which arrow to follow. Here we assume that if there are a arrows originating at state i , then each is chosen with probability $1/a$, but in general each arrow could be given any probability, such that the sum of the probabilities on the arrows leaving i is 1. The transition matrix of the chain shown above is

$$Q = \begin{pmatrix} 1/3 & 1/3 & 1/3 & 0 \\ 0 & 0 & 1/2 & 1/2 \\ 0 & 1 & 0 & 0 \\ 1/2 & 0 & 0 & 1/2 \end{pmatrix}.$$

To compute, say, the probability that the chain is in state 3 after 5 steps, starting at state 1, we would look at the (3,1) entry of Q^5 . Here (using a computer to find Q^5),

$$Q^5 = \begin{pmatrix} 853/3888 & 509/1944 & 52/243 & 395/1296 \\ 173/864 & 85/432 & 31/108 & 91/288 \\ 37/144 & 29/72 & 1/9 & 11/48 \\ 499/2592 & 395/1296 & 71/324 & 245/864 \end{pmatrix},$$

so $q_{13}^{(5)} = 52/243$.

To fully specify the behavior of the chain X_0, X_1, \dots , we also need to give initial conditions. This can be done by setting the initial state X_0 to be a particular state x_0 , or by randomly choosing X_0 according to some distribution. Let (s_1, s_2, \dots, s_M) be a vector with $s_i = P(X_0 = i)$ (think of this as the PMF of X_0 , displayed as a vector). Then the distribution of the chain at any time can be computed using the transition matrix.

Proposition 3. *Define $\mathbf{s} = (s_1, s_2, \dots, s_M)$ by $s_i = P(X_0 = i)$, and view \mathbf{s} as a row vector. Then $\mathbf{s}Q^n$ is the vector which gives the distribution of X_n , i.e., the j th component of $\mathbf{s}Q^n$ is $P(X_n = j)$.*

Proof. Conditioning on X_0 , the probability that the chain is in state j after n steps is $\sum_i s_i q_{ij}^n$, which is the j th component of $\mathbf{s}Q^n$. \square

3 Recurrence and Transience

In the Markov chain shown in Figure 1, a particle moving around between states will continue to spend time in all 4 states in the long run. In contrast, consider the chain shown in Figure 2, and let the particle start at state 1. For a while, the chain may

linger in the triangle formed by states 1, 2, and 3, but eventually it will reach state 4, and from there it can never return to states 1, 2, or 3. It will then wander around between states 4, 5, and 6 forever. In this example, states 1, 2, and 3 are *transient* and states 4, 5, and 6 are *recurrent*. In general, these concepts are defined as follows.

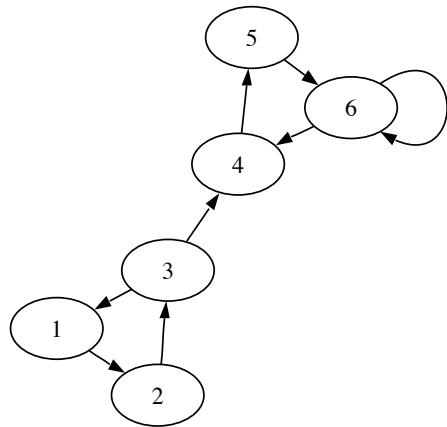


Figure 2: A Markov Chain with States 1, 2, and 3 Transient

Definition 4. State i is *recurrent* if the probability is 1 that the chain will return to i (eventually) if it starts at i . Otherwise, the state is *transient*, which means that if the chain starts out at i , there is a positive probability of never returning to i .

Classifying the states as recurrent or transient is important in understanding the long-run behavior of the chain. Early on in the history, the chain may spend time in transient states. Eventually though, the chain will spend all its time in recurrent states. But what fraction of the time will it spend in each of the recurrent state? To answer this question, we need the concept of a *stationary distribution*.

4 Stationary Distribution

The *stationary distribution* of a Markov chain, also known as the *steady state distribution*, describes the long-run behavior of the chain. Such a distribution is defined as follows.

Definition 5. A row vector $\mathbf{s} = (s_1, \dots, s_M)$ such that $s_i \geq 0$ and $\sum_i s_i = 1$ is a *stationary distribution* for a Markov chain with transition matrix Q if

$$\sum_i s_i q_{ij} = s_j$$

or equivalently,

$$\mathbf{s}Q = \mathbf{s}.$$

Often π is used to denote a stationary distribution, but some prefer to reserve the letter π for some other obscure purpose. The equation $\mathbf{s}Q = \mathbf{s}$ means that if X_0 has distribution given by \mathbf{s} , then X_1 also has distribution \mathbf{s} . But then X_2 also has distribution \mathbf{s} , etc. That is, a Markov chain which starts out with a stationary distribution will stay in the stationary distribution forever.

In terms of linear algebra, the equation $\mathbf{s}Q = \mathbf{s}$ says that \mathbf{s} is a left eigenvector of Q with eigenvalue 1. To get the usual kind of eigenvector (a right eigenvector), take transposes. Writing A' for the transpose of any matrix A , we have $Q'\mathbf{s}' = \mathbf{s}'$.

For example, if $Q = \begin{pmatrix} 1/3 & 2/3 \\ 1/2 & 1/2 \end{pmatrix}$, then $(3/7, 4/7)$ is a stationary distribution since

$$\begin{pmatrix} 3/7 & 4/7 \end{pmatrix} \begin{pmatrix} 1/3 & 2/3 \\ 1/2 & 1/2 \end{pmatrix} = \begin{pmatrix} 3/7 & 4/7 \end{pmatrix}.$$

This was obtained by finding the eigenvector of Q' with eigenvalue 1, and then normalizing (by dividing by the sum of the components).

But does a stationary distribution always exist? Is it unique? It turns out that a stationary distribution always exists. (Recall that we are assuming a finite state space; there may be no stationary distribution if the state space is infinite. For example, it can be shown that if X_n is a random walk on the integers, with $X_n = Y_1 + Y_2 + \dots + Y_n$ where Y_j are i.i.d. with $P(Y_j = -1) = P(Y_j = 1) = 1/2$, then the random walk X_1, X_2, \dots does not have a stationary distribution.)

There are in fact examples where there is not a unique stationary distribution, e.g., in the Gambler's Ruin problem. But under the reasonable assumption that it is possible to get from any state to any other state, then there will be a unique stationary distribution.

Definition 6. A Markov chain with transition matrix Q is *irreducible* if for any two states i and j , it is possible to go from i to j with positive probability (in some number of steps). That is, the (i, j) entry of Q^n is positive for some n . The chain is *reducible* if it is not irreducible.

The chain in Figure 1 is irreducible (in terms of the picture, check that it's possible to go from anywhere to anywhere following the arrows; in terms of the transition matrix, note that all the entries of Q^5 are positive). The chain in Figure 2 is reducible since it is never possible to go from state 4 back to state 3 in any number of steps.

Theorem 7. *Any irreducible Markov chain has a unique stationary distribution. In this distribution, every state has positive probability.*

Definition 8. The *period* of a state i in a Markov chain is the greatest common divisor of the possible numbers of steps it can take to return to i when starting at i . That is, it is the greatest common divisor of numbers n such that the (i, i) entry of Q^n is positive. A state is called *aperiodic* if its period is 1, and the chain itself is called *aperiodic* if all its states are aperiodic, and *periodic* otherwise.

For example, the “clockwork” behavior of states 1, 2, 3 in Figure 2 makes them periodic with period 3. In contrast, all the states in Figure 1 are aperiodic, so that chain is aperiodic.

Theorem 9. *Let X_0, X_1, \dots be an irreducible, aperiodic Markov chain with stationary distribution \mathbf{s} and transition matrix Q . Then $P(X_n = i)$ converges to s_i as $n \rightarrow \infty$. In other words, Q^n converges to a matrix in which each row is \mathbf{s} .*

Therefore, after a large number of steps, the probability that the chain is in state i is close to the stationary probability s_i . We can also use the stationary distribution to find the average time between visits to a state i , and vice versa.

Theorem 10. *Let X_0, X_1, \dots be an irreducible Markov chain with stationary distribution \mathbf{s} . Let r_i be the expected time it takes the chain to return to i , given that it starts at i . Then $s_i = 1/r_i$.*

In the 2×2 example on the previous page, this says that in the long run, the chain will spend $3/7$ of its time in state 1 and $4/7$ of its time in state 2. Starting at state 1, it will take an average of $7/3$ steps to return to state 1. The powers of the transition matrix converge to a matrix where each row is the stationary distribution:

$$\begin{pmatrix} 1/3 & 2/3 \\ 1/2 & 1/2 \end{pmatrix}^n \rightarrow \begin{pmatrix} 3/7 & 4/7 \\ 3/7 & 4/7 \end{pmatrix} \text{ as } n \rightarrow \infty.$$

In general, it may be computationally difficult to find the stationary distribution when the state space is large. Here are two important cases where working with eigenvalue equations for large matrices can be avoided.

Proposition 11. *If each column of the transition matrix Q sums to 1, then the uniform distribution over all states, $(1/M, 1/M, \dots, 1/M)$, is a stationary distribution.*

Proof. Assuming each column sums to 1, the row vector $\mathbf{v} = (1, 1, \dots, 1)$ satisfies $\mathbf{v}Q = \mathbf{v}$. It then follows that $(1/M, 1/M, \dots, 1/M)$ is stationary. \square

For example, if Q is a symmetric matrix (i.e., $q_{ij} = q_{ji}$), then $(1/M, 1/M, \dots, 1/M)$ is stationary. We now define *reversible* Markov chains, which intuitively have the property that they “look the same” if time is reversed (if the chain starts at the stationary distribution), and which are much nicer to work with than general chains.

Definition 12. Let $Q = (q_{ij})$ be the transition matrix of a Markov chain, and that there is $\mathbf{s} = (s_1, \dots, s_M)$ with $s_i \geq 0$, $\sum_i s_i = 1$, such that

$$s_i q_{ij} = s_j q_{ji}$$

for all states i, j . This equation is called the *reversibility* or *detailed balance* condition, and we say that the chain is *reversible* with respect to \mathbf{s} if it holds.

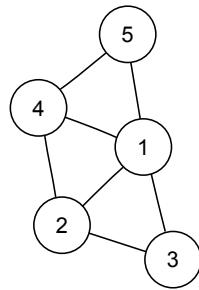
Proposition 13. *Suppose that $Q = (q_{ij})$ is a transition matrix which is reversible with respect to $\mathbf{s} = (s_1, \dots, s_M)$. Then \mathbf{s} is a stationary distribution for the chain with transition matrix Q .*

Proof. We have

$$\sum_i s_i q_{ij} = \sum_i s_j q_{ji} = s_j \sum_i q_{ji} = s_j,$$

where the last equality is because each row sum of Q is 1. Thus, \mathbf{s} is stationary. \square

Example 14. Consider a *random walk on an undirected network*, where a wanderer randomly traverses edges. From a node i , the wanderer randomly picks any of the edges at i , with equal probabilities, and then traverses the chosen edge. For example,



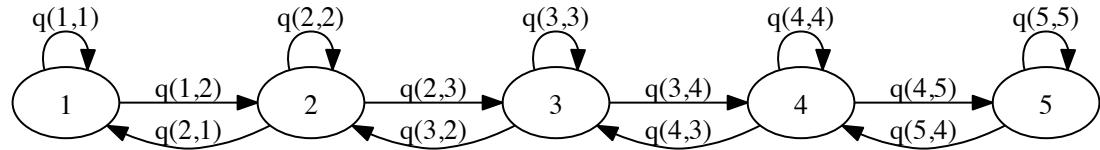
in the network shown above, from node 3 the wanderer goes to node 1 or node 2, with probability $1/2$ each.

The *degree* of a node is the number of edges attached to it, and the *degree sequence* of a network with nodes $1, 2, \dots, n$ is the vector (d_1, \dots, d_n) listing all the degrees, where d_j is the degree of node j . For example, the network above has degree sequence $\mathbf{d} = (4, 3, 2, 3, 2)$. Note that

$$d_i q_{ij} = d_j q_{ji}$$

for all i, j , since q_{ij} is $1/d_i$ if $\{i, j\}$ is an edge and 0 otherwise, for $i \neq j$. By Proposition 13, we have that the stationary distribution is proportional to the degree sequence. In the example above, this says that $\mathbf{s} = (\frac{4}{14}, \frac{3}{14}, \frac{2}{14}, \frac{3}{14}, \frac{2}{14})$ is the stationary distribution for the random walk.

Example 15. A *birth-death chain* on the state space $\{1, 2, \dots, M\}$ is a Markov chain with transition matrix $Q = (q_{ij})$ such that $q_{ij} > 0$ if $|i - j| = 1$ and $q_{ij} = 0$ if $|i - j| \geq 2$. Intuitively, this says it's possible to go "one step to the left" and possible to go "one step to the right" (except at boundaries) but it's impossible to jump further in one step. For example, the chain below is a birth-death chain if the labeled transitions have positive probabilities (except for the "loops" from a state to itself, which are allowed to have 0 probability). We will now show that any birth-death chain is



reversible, and construct the stationary distribution. Let s_1 be a positive number (to be specified later). Since we want $s_1 q_{12} = s_2 q_{21}$, let $s_2 = s_1 q_{12} / q_{21}$. Then since we want $s_2 q_{23} = s_3 q_{32}$, let $s_3 = s_2 q_{23} / q_{32} = s_1 q_{12} q_{23} / (q_{32} q_{21})$. Continuing in this way, let

$$s_j = \frac{s_1 q_{12} q_{23} \dots q_{j-1,j}}{q_{j,j-1} q_{j-1,j-2} \dots q_{21}},$$

for all states j with $2 \leq j \leq M$. Choose s_1 so that the s_j 's sum to 1. Then the chain is reversible with respect to \mathbf{s} , since $q_{ij} = q_{ji} = 0$ if $|i - j| \geq 2$ and by construction $s_i q_{ij} = s_j q_{ji}$ if $|i - j| = 1$. Thus, \mathbf{s} is the stationary distribution.