Math 312 Lecture Notes
Markov Chains

Warren Weckesser
Department of Mathematics
Colgate University

Updated, 30 April 2005

Markov Chains

A (finite) Markov chain is a process with a finite number of states (or outcomes, or events) in which the probability of being in a particular state at step $n+1$ depends only on the state occupied at step n.

Let $S = \{S_1, S_2, \ldots, S_r\}$ be the possible states. Let

$$\vec{\mathbf{p}}_n = \begin{bmatrix} p_1 \\ p_2 \\ \vdots \\ p_r \end{bmatrix} \tag{1}$$

be the vector of probabilities of each state at step n. That is, the ith entry of $\vec{\mathbf{p}}_n$ is the probability that the process is in state S_i at step n. For such a probability vector, $p_1 + p_2 + \cdots + p_r = 1$.

Let

$$p_{ij} = \text{Prob(State } n+1 \text{ is } S_i \mid \text{State } n \text{ is } S_j), \tag{2}$$

and let

$$P = \begin{bmatrix} p_{11} & p_{12} & \cdots & p_{1r} \\ p_{21} & p_{22} & \cdots & p_{2r} \\ \vdots & & \ddots & \\ p_{r1} & & & p_{rr} \end{bmatrix} \tag{3}$$

That is, p_{ij} is the (conditional) probability of being in state S_i at step $n+1$ given that the process was in state S_j at step n. P is called the *transition matrix*. P contains all the conditional probabilities of the Markov chain. It can be useful to label the rows and columns of P with the states, as in this example with three states:

$$\begin{array}{cc} & \overbrace{\begin{array}{ccc} S_1 & S_2 & S_3 \end{array}}^{\text{State } n} \\ \text{State } n+1 \left\{ \begin{array}{c} S_1 \\ S_2 \\ S_3 \end{array} \right. & \begin{bmatrix} p_{11} & p_{12} & p_{13} \\ p_{21} & p_{22} & p_{23} \\ p_{31} & p_{32} & p_{33} \end{bmatrix} \end{array}$$

The fundamental property of a Markov chain is that

$$\vec{\mathbf{p}}_{n+1} = P\vec{\mathbf{p}}_n. \tag{4}$$

Given an initial probability vector $\vec{\mathbf{p}}_0$, we can determine the probability vector at any step n by computing the iterates of a linear map.

The information contained in the transition matrix can also be represented in a *transition diagram.* In a transition diagram, the states are arranged in a diagram, typically with a "bubble" around each state. If $p_{ij} > 0$, then an arrow is drawn from state j to state i, and the arrow is labeled with p_{ij}. Examples are given in the following discussions.

In these notes, we will consider two special cases of Markov chains: regular Markov chains and absorbing Markov chains. Generalizations of Markov chains, including continuous time Markov processes and infinite dimensional Markov processes, are widely studied, but we will not discuss them in these notes.

Regular Markov Chains

Definition. A Markov chain is a *regular* Markov chain if some power of the transition matrix has only positive entries. That is, if we define the (i, j) entry of P^n to be $p_{ij}^{(n)}$, then the Markov chain is regular if there is some n such that $p_{ij}^{(n)} > 0$ for all (i, j).

If a Markov chain is regular, then no matter what the initial state, in n steps there is a positive probability that the process is in *any* of the states.

Essential facts about regular Markov chains.

1. $P^n \to W$ as $n \to \infty$, where W is a constant matrix and all the columns of W are the same.
2. There is a unique probability vector $\vec{\mathbf{w}}$ such that $P\vec{\mathbf{w}} = \vec{\mathbf{w}}$.

   Note:

   (a) $\vec{\mathbf{w}}$ is a fixed point of the linear map $\vec{\mathbf{x}}_{i+1} = P\vec{\mathbf{x}}_i$.

   (b) $\vec{\mathbf{w}}$ is an eigenvector associated with the eigenvalue $\lambda = 1$. (The claim implies that the transition matrix P of a regular Markov chain *must* have the eigenvalue $\lambda = 1$. Then $\vec{\mathbf{w}}$ is the eigenvector whose entries add up to 1.)

   (c) The matrix W is $W = \begin{bmatrix} \vec{\mathbf{w}} & \vec{\mathbf{w}} & \cdots & \vec{\mathbf{w}} \end{bmatrix}$.
3. $P^n\vec{\mathbf{p}}_0 \to \vec{\mathbf{w}}$ as $n \to \infty$ for *any* initial probability vector $\vec{\mathbf{p}}_0$. Thus $\vec{\mathbf{w}}$ gives the long-term probability distribution of the states of the Markov chain.

Example: Sunny or Cloudy. A meteorologist studying the weather in a region decides to classify each day as simply *sunny* or *cloudy.* After analyzing several years of weather records, he finds:

- the day after a sunny day is sunny 80% of the time, and cloudy 20% of the time; and
- the day after a cloudy day is sunny 60% of the time, and cloudy 40% of the time.

We can setup up a Markov chain to model this process. There are just two states: S_1 = *sunny*, and S_2 = *cloudy*. The transition diagram is

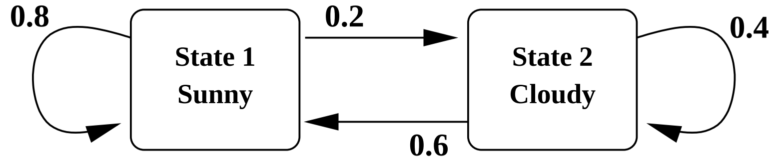

and the transition matrix is

$$P = \begin{bmatrix} 0.8 & 0.6 \\ 0.2 & 0.4 \end{bmatrix}. \tag{5}$$

We see that all entries of P are positive, so the Markov chain is regular. (The conditions of the definition are satisfied when $n = 1$.)

To find the long-term probabilities of sunny and cloudy days, we must find the eigenvector of P associated with the eigenvalue $\lambda = 1$. We know from Linear Algebra that if $\vec{\mathbf{v}}$ is an eigenvector, then so is $c\vec{\mathbf{v}}$ for any constant $c \neq 0$. The probability vector $\vec{\mathbf{w}}$ is the eigenvector that is also a probability vector. That is, the sum of the entries of the vector $\vec{\mathbf{w}}$ must be 1.

We solve

$$\begin{aligned} P\vec{\mathbf{w}} &= \vec{\mathbf{w}} \\ (P - I)\vec{\mathbf{w}} &= \vec{\mathbf{0}} \end{aligned} \tag{6}$$

Now

$$P - I = \begin{bmatrix} -0.2 & 0.6 \\ 0.2 & -0.6 \end{bmatrix} \tag{7}$$

If you have recently studied Linear Algebra, you could probably write the answer down with no further work, but we will show the details here. We form the augmented matrix and use Gaussian elimination:

$$\begin{bmatrix} -0.2 & 0.6 & \vdots & 0 \\ 0.2 & -0.6 & \vdots & 0 \end{bmatrix} \rightarrow \begin{bmatrix} 1 & -3 & \vdots & 0 \\ 0 & 0 & \vdots & 0 \end{bmatrix} \tag{8}$$

which tells us $w_1 = 3w_2$, or $w_1 = 3s$, $w_2 = s$, where s is arbitrary, or

$$\vec{\mathbf{w}} = s \begin{bmatrix} 3 \\ 1 \end{bmatrix} \tag{9}$$

The vector $\vec{\mathbf{w}}$ must be a probability vector, so $w_1 + w_2 = 1$. This implies $4s = 1$ or $s = 1/4$. Thus

$$\vec{\mathbf{w}} = \begin{bmatrix} 3/4 \\ 1/4 \end{bmatrix}. \tag{10}$$

This vector tells us that in the long run, the probability is 3/4 that the process will be in state 1, and 1/4 that the process will be in state 2. In other words, in the long run 75% of the days are sunny and 25% of the days are cloudy.

Examples: Regular or not? Here are a few examples of determining whether or not a Markov chain is regular.

1. Suppose the transition matrix is
$$P = \begin{bmatrix} 1/3 & 0 \\ 2/3 & 1 \end{bmatrix}. \tag{11}$$
We find
$$P^2 = \begin{bmatrix} (1/3)^2 & 0 \\ (2/3)(1+1/3) & 1 \end{bmatrix}, \quad P^3 = \begin{bmatrix} (1/3)^3 & 0 \\ (2/3)(1+1/3+(1/3)^2) & 1 \end{bmatrix}, \tag{12}$$
and, in general,
$$P^n = \begin{bmatrix} (1/3)^n & 0 \\ (2/3)(1+1/3+\cdots+(1/3)^{n-1}) & 1 \end{bmatrix}. \tag{13}$$
The upper right entry in P^n is 0 for all n, so the Markov chain is *not* regular.

2. Here's a simple example that is not regular.
$$P = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \tag{14}$$
Then
$$P^2 = I, \quad P^3 = P, \quad \text{etc.} \tag{15}$$
Since $P^n = I$ if n is even and $P^n = P$ if n is odd, P always has two entries that are zero. Therefore the Markov chain is not regular.

3. Let
$$P = \begin{bmatrix} 1/5 & 1/5 & 2/5 \\ 0 & 2/5 & 3/5 \\ 4/5 & 2/5 & 0 \end{bmatrix} \tag{16}$$
The transition matrix has two entries that are zero, but
$$P^2 = \begin{bmatrix} 9/25 & 7/25 & 5/25 \\ 12/25 & 10/25 & 6/25 \\ 4/25 & 8/25 & 14/25 \end{bmatrix}. \tag{17}$$
Since all the entries of P^2 are positive, the Markov chain is regular.

Absorbing Markov Chains

We consider another important class of Markov chains. A state S_k of a Markov chain is called an *absorbing state* if, once the Markov chains enters the state, it remains there forever. In other words, the probability of leaving the state is zero. This means $p_{kk} = 1$, and $p_{jk} = 0$ for $j \neq k$.

A Markov chain is called an *absorbing chain* if

(i) it has at least one absorbing state; and

(ii) for every state in the chain, the probability of reaching an absorbing state in a finite number of steps is nonzero.

An essential observation for an absorbing Markov chain is that it will eventually enter an absorbing state. (This is a consequence of the fact that if a random event has a probability $p > 0$ of occurring, then the probability that it does not occur is $1 - p$, and the probability that it does not occur in n trials is $(1 - p)^n$. As $n \to \infty$, the probability that the event does not occur goes to zero.) The non-absorbing states in an absorbing Markov chain are called *transient states.*

Suppose an absorbing Markov chain has k absorbing states and ℓ transient states. If, in our set of states, we list the absorbing states first, we see that the transition matrix has the form

$$
\begin{array}{c}
\\ \\ S_1 \\ S_2 \\ \vdots \\ S_k \\ S_{k+1} \\ \vdots \\ S_{k+\ell}
\end{array}
\begin{array}{c}
\overbrace{\begin{array}{cccc} & & & \end{array}}^{\text{Absorbing States}} \quad \overbrace{\begin{array}{ccc} & & \end{array}}^{\text{Transient States}} \\
\begin{array}{ccccccc} S_1 & S_2 & \cdots & S_k & S_{k+1} & \cdots & S_{k+\ell} \end{array} \\
\begin{bmatrix}
1 & 0 & \cdots & 0 & p_{1,k+1} & \cdots & p_{1,k+\ell} \\
0 & 1 & & \vdots & \vdots & & \vdots \\
\vdots & & \ddots & 0 & & & \\
0 & \cdots & 0 & 1 & p_{k,k+1} & \cdots & p_{k,k+\ell} \\
0 & \cdots & \cdots & 0 & p_{k+1,k+1} & \cdots & p_{k+1,k+\ell} \\
\vdots & & & \vdots & \vdots & & \vdots \\
0 & \cdots & \cdots & 0 & p_{k+\ell,k+1} & \cdots & p_{k+\ell,k+\ell}
\end{bmatrix}
\end{array}
$$

That is, we may partition P as

$$
P = \begin{bmatrix} I & R \\ \mathbf{0} & Q \end{bmatrix} \tag{18}
$$

where I is $k \times k$, R is $k \times \ell$, $\mathbf{0}$ is $\ell \times k$ and Q is $\ell \times \ell$. R gives the probabilities of transitions from transient states to absorbing states, while Q gives the probabilities of transitions from transient states to transient states.

Consider the powers of P:

$$
P^2 = \begin{bmatrix} I & R(I + Q) \\ \mathbf{0} & Q^2 \end{bmatrix}, \quad P^3 = \begin{bmatrix} I & R(I + Q + Q^2) \\ \mathbf{0} & Q^3 \end{bmatrix}, \tag{19}
$$

and, in general,

$$
P^n = \begin{bmatrix} I & R(I + Q + Q^2 + \cdots + Q^{n-1}) \\ \mathbf{0} & Q^n \end{bmatrix} = \begin{bmatrix} I & R\sum_{i=0}^{n-1} Q^i \\ \mathbf{0} & Q^n \end{bmatrix}, \tag{20}
$$

Now I claim that[1]

$$\lim_{n\to\infty} P^n = \begin{bmatrix} I & R(I-Q)^{-1} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \tag{21}$$

That is, we have

1. $Q^n \to \mathbf{0}$ as $n \to \infty$, and
2. $\displaystyle\sum_{i=0}^{\infty} Q^i = (I-Q)^{-1}$.

The first claim, $Q^n \to \mathbf{0}$, means that in the long run, the probability is 0 that the process will be in a transient state. In other words, the probability is 1 that the process will eventually enter an absorbing state. We can derive the second claim as follows. Let

$$U = \sum_{i=0}^{\infty} Q^i = I + Q + Q^2 + \cdots \tag{22}$$

Then

$$QU = Q\sum_{i=0}^{\infty} Q^i = Q + Q^2 + Q^3 + \cdots = (I + Q + Q^2 + Q^3 + \cdots) - I = U - I. \tag{23}$$

Then $QU = U - I$ implies

$$\begin{aligned} U - UQ &= I \\ U(I-Q) &= I \\ U &= (I-Q)^{-1}, \end{aligned} \tag{24}$$

which is the second claim. (The claims can be rigorously justified, but for our purposes, the above arguments will suffice.)

The matrix $R(I-Q)^{-1}$ has the following meaning. The column i of $R(I-Q)^{-1}$ gives the probabilities of ending up in each of the absorbing states, given that the process started in the i^{th} transient states.

There is more information that we can glean from $(I-Q)^{-1}$. For convenience, call the transient states T_1, T_2, $\ldots$, T_ℓ. (So $T_j = S_{k+j}$.) Let $V(T_i, T_j)$ be the expected number of times that the process is in state T_i, given that it started in T_j. (V stands for the number of "**v**isits".) Also recall that Q gives the probabilities of transitions from transient states to transient states, so

$$q_{ij} = \text{Prob(State } n+1 \text{ is } T_i \mid \text{State } n \text{ is } T_j) \tag{25}$$

I claim that $V(T_i, T_j)$ satisfies the following equation:

$$V(T_i,T_j) = e_{ij} + q_{i1}V(T_1,T_j) + q_{i2}V(T_2,T_j) + \cdots + q_{i\ell}V(T_\ell,T_j) \tag{26}$$

where

$$e_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases} \tag{27}$$

[1] There is a slight abuse of notation in the formula given. I use the symbol $\mathbf{0}$ to mean "a matrix of zeros of the appropriate size". The two $\mathbf{0}$'s in the formula are not necessarily the same size. The $\mathbf{0}$ in the lower left is $\ell \times k$, while the $\mathbf{0}$ in the lower right is $\ell \times \ell$.

Why? Consider just the term $q_{i1}V(T_1, T_j)$. Given that the process started in T_j, $V(T_1, T_j)$ gives the expected number of times that the process will be in T_1. The number q_{i1} gives the probability of making a transition from T_1 to T_i. Therefore, the product $q_{i1}V(T_1, T_j)$ gives the *expected number of transitions from* T_1 *to* T_i, given that the process started in T_j. Similarly, $q_{i2}V(T_2, T_j)$ gives the expected number of transitions from T_2 to T_i, and so on. Therefore the total number of expected transition to T_i is $q_{i1}V(T_1, T_j) + q_{i2}V(T_2, T_j) + \cdots + q_{i\ell}V(T_\ell, T_j)$.

The expected number of transitions into a state is the same as the expected number of times that the process is in a state, except in one case. Counting the *transitions* misses the state in which the process started, since there is no "transition" into the initial state. This is why the term e_{ij} is included in (26). When we consider $V(T_i, T_i)$, we have to add 1 to the expected number of transitions into T_i to get the correct expected number of times that the process was in T_i.

Equation (26) is actually a set of ℓ^2 equations, one for each possible (i, j). In fact, it is just one component of a matrix equation. Let

$$N = \begin{bmatrix} V(T_1,T_1) & V(T_1,T_2) & \cdots & V(T_1,T_\ell) \\ V(T_2,T_1) & V(T_2,T_2) & & \\ \vdots & & \ddots & \\ V(T_\ell,T_1) & & & V(T_\ell,T_\ell) \end{bmatrix} \tag{28}$$

Then equation (26) is the (i, j) entry in the matrix equation

$$N = I + QN. \tag{29}$$

(You should check this yourself!) Solving (29) gives

$$\begin{aligned} N - QN &= I \\ (I - Q)N &= I \\ N &= (I - Q)^{-1} \end{aligned} \tag{30}$$

Thus the (i, j) entry of $(I - Q)^{-1}$ gives the expected number of times that the process is in the i^{th} transient state, given that it started in the j^{th} transient state. It follows that the sum of the i^{th} column of N gives the expected number of times that the process will be in some transient state, given that the process started in the j^{th} transient state.

Example: The Coin and Die Game. In this game there are two players, *Coin* and *Die*. *Coin* has a coin, and *Die* has a single six-sided die. The rules are:

- When it is *Coin*'s turn, he or she flips the coin. If the coin turns up **heads**, *Coin* wins the game. If the coin turns up **tails**, it is *Die*'s turn.
- When it is *Die*'s turn, he or she rolls the die. If the die turns up **1**, *Die* wins. If the die turns up **6**, it is *Coin*'s turn. Otherwise, *Die* rolls again.

When it is *Coin*'s turn, the probability is 1/2 that *Coin* will win and 1/2 that it will become *Die*'s turn. When it is *Die*'s turn, the probabilities are

- 1/6 that *Die* will roll a **1** and win,

- 1/6 that *Die* will roll a **6** and it will become *Coin*'s turn, and
- 2/3 that *Die* will roll a **2**, **3**, **4**, or **5** and have another turn.

To describe this process as a Markov chain, we define four *states* of the process:

- *State 1*: *Coin* has won the game.
- *State 2*: *Die* has won the game.
- *State 3*: It is *Coin*'s turn.
- *State 4*: It is *Die*'s turn.

We represent the possible outcomes in the following transition diagram:

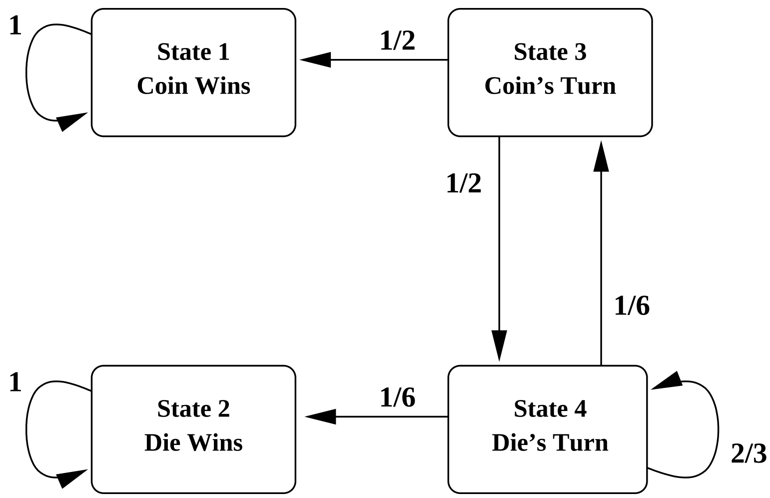

This is an absorbing Markov chain. The absorbing states are State 1 and State 2, in which one of the players has won the game, and the transient states are State 3 and State 4.

The transition matrix is

$$P = \begin{bmatrix} 1 & 0 & 1/2 & 0 \\ 0 & 1 & 0 & 1/6 \\ 0 & 0 & 0 & 1/6 \\ 0 & 0 & 1/2 & 2/3 \end{bmatrix} = \left[\begin{array}{cc:cc} 1 & 0 & 1/2 & 0 \\ 0 & 1 & 0 & 1/6 \\ \hdashline 0 & 0 & 0 & 1/6 \\ 0 & 0 & 1/2 & 2/3 \end{array}\right] = \begin{bmatrix} I & R \\ \mathbf{0} & Q \end{bmatrix} \tag{31}$$

where

$$R = \begin{bmatrix} 1/2 & 0 \\ 0 & 1/6 \end{bmatrix} \quad \text{and} \quad Q = \begin{bmatrix} 0 & 1/6 \\ 1/2 & 2/3 \end{bmatrix}. \tag{32}$$

We find

$$I - Q = \begin{bmatrix} 1 & -1/6 \\ -1/2 & 1/3 \end{bmatrix}, \tag{33}$$

so

$$N = (I - Q)^{-1} = \begin{bmatrix} 4/3 & 2/3 \\ 2 & 4 \end{bmatrix}, \tag{34}$$

and

$$R(I - Q)^{-1} = \begin{bmatrix} 2/3 & 1/3 \\ 1/3 & 2/3 \end{bmatrix} \tag{35}$$

Recall that the first column of $R(I - Q)^{-1}$ gives the probabilities of entering State 1 or State 2 if the process starts in State 3. "Starts in State 3" means *Coin* goes first, and "State 1" means *Coin* wins, so this matrix tells us that if *Coin* goes first, the probability that *Coin* will win is 2/3, and the probability that *Die* will win is 1/3. Similarly, if *Die* goes first, the probability that *Coin* will win is 1/3, and the probability that *Die* will win is 2/3.

From (34), we can also conclude the following:

1. If *Coin* goes first, then the expected number of turns for *Coin* is 4/3, and the expected number of turns for *Die* is 2. Thus the expected total number of turns is $10/3 \approx 3.33$.

2. If *Die* goes first, then the expected number of turns for *Coin* is 2/3, and the expected number of turns for *Die* is 4. Thus the expected total number of turns is $14/3 \approx 4.67$.

The following table gives the results of our "experiment" with the Coin and Die Game along with the predictions of the theory. In class, a total of 220 games were played in which *Coin* went first. *Coin* won 138 times, and the total number of turns was 821, for an average of 3.73 turns per game.

Quantity	*Predicted*	*Experiment*
Percentage Won by Coin	66.7	62.7
Average Number of Turns per Game	3.33	3.73

It appears that in our experiment, *Die* won more often than predicted by the theory. Presumably if we played the game a lot more, the experimental results would approach the predicted results.