Lecture 16: Markov Chains

CSCI2244-Randomness and Computation

April 25, 2019

This is a summary of the main results, with several examples of how to apply them, but no proofs. Typical presentations involve a heavy load of linear algebra, but here we will just use the very basics of matrix multiplication. A summary of the necessary matrix arithmetic, as well as how to do matrix arithmetic in numpy, is given in an appendix.

1 Example 1. Lily pads.

There are three lily pads on a pond. At regular intervals, a frog gets the urge to jump from one lily pad to another. When he is on either pad 1 or pad 2, he jumps to one of the two others with equal probability. But when he is on pad 3, half the time he will choose to stay there for a spell rather than jump (since it is a sunny spot), one-third of the time jump to pad 1, and one-sixth of the time jump to pad 2. We give two equivalent pictures of this setup, one as a state-transition diagram, the other as a 3×3 matrix M whose ij-entry (entry in row i, column j) gives the probability of jumping from pad i to pad j.

$$M = \begin{bmatrix} 0 & 1/2 & 1/2 \\ 1/2 & 0 & 1/2 \\ 1/3 & 1/6 & 1/2 \end{bmatrix}.$$

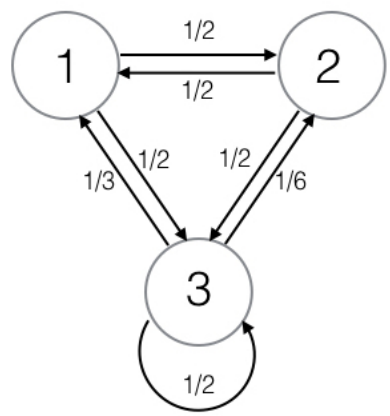

Figure 1: State-transition diagram for the frog jumping between lily pads.

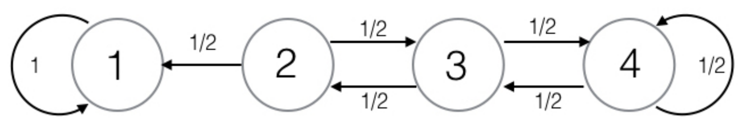

Figure 2: State-transition diagram for drunken professor.

2 Example 2. Drunken professor.

A famous professor of Computer Science attends a conference at a resort hotel. The evening of the day that he presents his paper, he has a bit too much to drink, and tries to navigate the three steps between the hotel bar (4) and his room (1). If he is at one of the two intermediate stops (2,3) he will stagger either to the left or right with equal probability. If he reaches the bar, he will, with equal probability, either stagger left back towards his room or sit in the bar and order another drink before venturing out again. If he reaches his room, he drops into bed and sleeps it off. The state-transition diagram is in Figure 2, and the matrix representation N is shown below:

$$N = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1/2 & 0 & 1/2 & 0 \\ 0 & 1/2 & 0 & 1/2 \\ 0 & 0 & 1/2 & 1/2 \end{bmatrix}.$$

3 Definition.

A *Markov chain* is a pair (S, M), where $S = \{1, \ldots, n\}$ is the set of *states,* and M is an $n \times n$ matrix. We require the matrix to satisfy the following two properties:

For all $i, j \in S$, $M_{ij} \geq 0$. That is, every entry is nonnegative.

For all $i \in S$,

$$\sum_{j=1}^{n} M_{ij} = 1.$$

That is, the sum along every row is 1.

We imagine an experiment that begins by starting at time 0 in some state of S, and then at regular intervals, changing state according to the given probabilities. For example, if we begin the lily pad experiment in state 1, then we flip a coin and transition to state 2 or state 3 according to the outcome. Let's say that we land in state 3. Then we make another random choice, and transition to one of the three states with the given probability, and so on indefinitely. We let X_t denote the random variable giving the state at time t. The basic assumptions underlying the evolution of the chain are that for all times t, and all $i, j \in S$,

$$P(X_{t+1} = j | X_t = i) = M_{ij}.$$

$$P(X_{t+1} = j | X_0 = i_0, \ldots, X_{t-1} = i_{t-1}, X_t = i_t) = P(X_{t+1} = j | X_t = j).$$

The second property means that the new state at each successive time step is independent of all the previous states that the chain was in, except the most recent state.

4 Evolution of the chain and matrix multiplication.

The transition matrix of the Markov chain gives the probabilities of making a transition from one state to any other in a single step. What about the behavior after 2 steps, 3 steps, etc.? As an example, look at the lily pad example. What is the probability of getting from state 3 to state 2 in two steps? There are two ways that this can occur: The frog can remain in state 3 for one time step, and then jump immediately to state 2 in the second; or it can jump to state 1 in the first time step, and then jump to state 2 in the next. The first route [1] will occur with probability

$$\frac{1}{2} \cdot \frac{1}{6}$$

and the second with probability

$$\frac{1}{3} \cdot \frac{1}{2}.$$

Thus the probability of a two-step transition from state 3 to state 2 is

$$\frac{1}{2} \cdot \frac{1}{6} + \frac{1}{3} \cdot \frac{1}{2} = \frac{1}{4}.$$

This is the $(3, 2)$ entry of the matrix product M^2 of M with itself. The full product is given by

$$M^2 = \begin{bmatrix} 5/12 & 1/12 & 1/2 \\ 1/6 & 1/3 & 1/2 \\ 1/4 & 1/4 & 1/2 \end{bmatrix}.$$

---

[1] We are really using our independence assumption here: We are computing

$$P(X_{t+2} = 2 \wedge X_{t+1} = 3 | X_t = 3).$$

It follows from the definition of conditional probability that this is the same as

$$P(X_{t+2} = 2 | X_{t+1} = 3 \wedge X_t = 3) \cdot P(X_{t+1} = 3 | X_t = 3),$$

and from our independence assumption that this is in turn equal to

$$P(X_{t+2} = 2 | X_{t+1} = 3) \cdot P(X_{t+1} = 3 | X_t = 3) = \frac{1}{6} \cdot \frac{1}{2}.$$

This is the situation in general, for every Markov chain, and for any number of steps. That is, the probability of transition from state i to state j in exactly k steps is the ij-entry of M^k. In symbols, for any t,

$$P(X_{t+k} = j | X_t = i) = (M^k)_{ij}.$$

Let's see what happens when we compute high powers of the matrices M and N in our examples. We'll start with the drunk professor:

$$N^{10} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0.81 & 0.04 & 0.06 & 0.09 \\ 0.66 & 0.06 & 0.13 & 0.15 \\ 0.57 & 0.09 & 0.15 & 0.19 \end{bmatrix}.$$

After ten steps, it is more likely than not that the professor is back in his bed, but how likely depends on where he started: If he began at the bar (state 4), there is a 57% chance that he is in bed by time 10. If we prolong the simulation for 10 more steps, we find

$$N^{20} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0.98 & 0.005 & 0.008 & 0.01 \\ 0.96 & 0.008 & 0.015 & 0.019 \\ 0.95 & 0.01 & 0.02 & 0.02 \end{bmatrix}.$$

(Because of rounding, the rows do not exactly add up to 1.) You can see where this is going—as time progresses, it becomes more and more likely that the professor winds up in his bed, with the probability appearing to approach 1, regardless of where he began.

The situation with the lily pads is quite different. We have

$$M^{10} = \begin{bmatrix} 0.2783 & 0.2217 & 0.5 \\ 0.2773 & 0.2227 & 0.5 \\ 0.2777 & 0.2223 & 0.5 \end{bmatrix}.$$

Here we are seeing a different sort of limiting behavior, in which all rows look like they are converging to the same three values. What this means is that after a day of jumping around, the frog will be on lily pad 1 about 28% of the time, lily pad 2 about 22%, and lily pad 3 half the time, *no matter where he started.*

In the next sections we analyze these two sorts of behavior in depth. The drunk's walk is an instance of an *absorbing* Markov chain, and the frog on the lily pad a *regular* Markov chain.

5 Absorbing Markov Chains

A state i in a Markov chain (S, P) is an *absorbing state* if $P_{ii} = 1$.

For example, the state 1 in the drunk professor's walk is an absorbing state. In the lily pad example, there is no absorbing state.

A Markov chain is an *absorbing Markov chain* if it has at least one absorbing state, and if from every state there is a path to an absorbing state.

So, for example, our Markov chain in Example 2 is an absorbing chain. An absorbing chain can have more than one absorbing state. The non-absorbing states of a Markov chain are called *transient states*.

Let's take an absorbing Markov chain and renumber the states so that states $1, \ldots, k$ are transient states, and $k+1, \ldots n$ are absorbing states. A schematic picture of the transition matrix is then

$$P = \left[\begin{array}{c|c} Q & R \\ \hline O & I \end{array}\right].$$

Here, Q is a $k \times k$ square matrix, I the $(n-k) \times (n-k)$ identity matrix, O the $(n-k) \times k$ matrix with all entries 0, and R a $k \times (n-k)$ matrix.

Here are the principal facts about absorbing Markov chains.

1. If i, j are transient states, then

$$\lim_{n \to \infty} (P^t)_{ij} = 0.$$

   In other words, no matter where you start, if you run the chain for a large enough number of steps, the probability of being in the transient state j is close to 0. More simply, you will always wind up in an absorbing state. This being the case, we would like to know how long it takes to reach an absorbing state. This is addressed below.

2. $I - Q$ is invertible. Let $S = (I - Q)^{-1}$. Then the i, j entry of S is the expected number of times that the chain will be in state j if it is started in state i.

3. This being the case, the sum of the entries in the i^{th} row of S is the average number of times the chain will be in a transient state when started in state i. This is the same as the expected time until reaching an absorbing state.

4. $B = SR$ is a $k \times (n - k)$ matrix whose rows are indexed by the transient states and whose columns are indexed by the absorbing states. The i, j entry of B is the probability that the chain, started in transient state i, reaches absorbing state j.

Why these properties hold will be explained in a bit more detail below. Let's first give some examples.

Example 1. We apply this to our drunk professor's walk. We reorder the states so that states 2,3,4 come before the absorbing state 1. The matrix Q giving the transitions among the transient states is

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

The inverse matrix of $I - Q$ is calculated by solving systems of linear equations. We will not carry out the calculation here (and in practice, you would use a computer for it) but the result is

$$S = (I - Q)^{-1} = \begin{bmatrix} 2 & 2 & 2 \\ 2 & 4 & 4 \\ 2 & 4 & 6 \end{bmatrix}.$$

What does this mean? Let's look at the last row, corresponding to state 4, which is the hotel bar. The entries in the corresponding row of S are $2, 4, 6$. this means that the average number of times the professor will be in state 2 when started in state 4 is 2; similarly the average number of times that it will be in state

4 when started in state 4 is 6. The sum $2+4+6=12$ gives the average number of times the professor will be in a transient state, which is the same as the average wait until absorption. In this example, the matrix R is

$$\begin{bmatrix} 1/2 \\ 0 \\ 0 \end{bmatrix},$$

and SR is

$$\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}.$$

All this says is that the probability of winding up in the absorbing state is 1. Since there is only one absorbing state, this gives no additional information—we already knew that.

Example 2. Alice has four dollars and Bob has two dollars. They repeatedly flip coins, betting one dollar on each outcome: Bob wins with heads and loses with tails. The game continues until one of them runs out of money. The game is modeled by the Markov chain shown in Figure 3; the state label represents Alice's winnings. Thus the game ends when it either reaches state 2 (Alice has won two dollars, so Bob is out or money) or state -4 (Bob has won four dollars, so Alice is out of money). In particular, there are now *two* absorbing states.[2]

We order the states as 1,0,-1,-2,-3,2,-4 so that the transient states come first. The matrix Q is then

[2]We could have told another drunk professor story about the chain: In this version he might fall asleep at the bar!

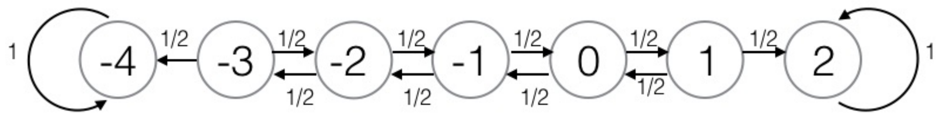

Figure 3: State-transition diagram for the coin-tossing game with Alice and Bob. The state label represents Bob's winnings.

$$Q = \begin{bmatrix} 0 & 1/2 & 0 & 0 & 0 \\ 1/2 & 0 & 1/2 & 0 & 0 \\ 0 & 1/2 & 0 & 1/2 & 0 \\ 0 & 0 & 1/2 & 0 & 1/2 \\ 0 & 0 & 0 & 1/2 & 0 \end{bmatrix},$$

and R is

$$\begin{bmatrix} 1/2 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 1/2 \end{bmatrix}.$$

The matrix $S = (1 - Q)^{-1}$ is given by

$$S = \begin{bmatrix} 5/3 & 4/3 & 1 & 2/3 & 1/3 \\ 4/3 & 8/3 & 2 & 4/3 & 2/3 \\ 1 & 2 & 3 & 2 & 1 \\ 2/3 & 4/3 & 2 & 8/3 & 4/3 \\ 1/3 & 2/3 & 1 & 4/3 & 5/3 \end{bmatrix}.$$

Suppose we start the game in state 0, which corresponds to the second row of the matrix. The row sum is 8, which means that the game will last an average of 8 tosses before someone runs out of money. The matrix SR is

$$\begin{bmatrix} 5/6 & 1/6 \\ 2/3 & 1/3 \\ 1/2 & 1/2 \\ 1/3 & 2/3 \\ 1/6 & 5/6 \end{bmatrix}.$$

The first column corresponds to the absorbing state 2, and the second to the absorbing state -4. The probabilities in the row corresponding to start state 0 are $2/3$ and $1/3$. So that 2/3 of the time when this game is played, the player who brought 4 dollars will win, and 1/3 of the time it will be the player who brought 2 dollars. (This is what you might have guessed!)

6 A bit more by way of explanation

Why do the properties cited above hold? It's sort of obvious that eventually you will wind up in an absorbing state, but let's give something more along the lines of a real proof: Let $q_{i,m}$ be the probability that the chain is *not* absorbed in m steps when we start from the transient state i. Since there is a path from i to the absorbing state, we know there is some m for which $q_{i,m} < 1$ —that is, there is some nonzero probability that the chain will be absorbed. We do this for every state and find the largest of these probabilities and the largest such m. Thus there is some $p < 1$ and $m > 0$ such that the probability of *not* being absorbed in m steps, no matter where you start, is no more than p. This means that the probability of not being absorbed in $2m$ steps is no more than p^2, in $3m$ steps no more than p^3, etc. As $p < 1$, these powers of p approach 0. This is the first property we mentioned.

Now let i, j be transient states of the chain. Let $X_{i,j,t}$ be the Bernoulli random variable whose value is 1 if, when the chain is started in state i, it is in state j at time t, and whose value is 0 otherwise. Then $X_{i,j,t}$ has value 1 with probability $(Q^t)_{ij}$, and $E(X_{i,j,t}) = (Q^t)_{ij}$.

Let $X_{i,j}$ be the number of times that the chain, started in state i, is in state j. Since we always end up in an absorbing state, $X_{i,j}$ is well defined. By the additivity of expected value,

$$E(X_{i,j}) = \sum_{t=0}^{\infty} E(X_{i,j,t}) = \sum_{t=0}^{\infty} (Q^t)_{ij}.$$

We have no *a priori* guarantee that this infinite series is convergent—remember, some random variables do not have a well-defined expected value. However, it *is* convergent, because of the fact cited above, that the entries of Q_t converge exponentially to 0. So we have,

$$E(X_{i,j}) = (I + Q + Q^2 + \cdots)_{i,j}.$$

One last little trick will help us explicitly evaluate the expected time to absorption. Observe that

$$(I - Q)(I + Q + Q^2 + \cdots + Q^k) = I - Q^{k+1},$$

so that passing to the limit, and using the fact that Q^{k+1} converges to the zero matrix,

$$(I - Q)(I + Q + Q^2 + \cdots) = I,$$

thus

$$I + Q + Q^2 + \cdots = (I - Q)^{-1}.$$

This is the second property mentioned above. The third follows from it directly. (I'll leave off explaining the fourth.)

7 Appendix: Matrix multiplication

Just a quick rundown of the few facts you need, examples were given in class:

- If A is an $\ell \times m$ matrix (ℓ rows, m columns) and B is an $m \times n$ matrix, then the $\ell \times n$ matrix AB is defined, and the ij-entry of AB is
$$A_{i1}B_{1j} + A_{i2}B_{2j} + \cdots + A_{im}B_{mj}.$$
- Special case: If A, B are square matrices of the same size (both $n \times n$) then AB is defined. Almost all our matrix multiplication will have this form.
- In general, matrix multiplication is not commutative, *i.e.,* you could have $AB \neq BA$ even if both products are defined. However, it is associative; that is, if either one of $(AB)C$ or $A(BC)$ is defined, then the other is defined, and
$$(AB)C = A(BC).$$
- As a consequence of the associativity, we have $(MM)M = M(MM)$ if M is a square matrix, so we can unambiguously define powers like M^3, M^4, *etc.*
- Matrix addition and subtraction are even easier—it's just component by component: If A, B have the same dimensions (both $m \times n$) then
$$(A + B)_{ij} = A_{ij} + B_{ij}.$$
And likewise subtraction.
- Matrix addition and subtraction obey distributive laws:
$$A(B + C) = AB + AC, (A + B)C = AC + BC,$$
as long as the relevant sums and products are defined.
- Let $n \geq 1$. The $n \times n$ identity matrix, denoted I_n, or just I when we can infer the dimension from the context, is defined by
$$I_{ij} = \begin{cases} 1, & \text{if } i = j \\ 0, & \text{otherwise} \end{cases}.$$
That is, I has 1's on the diagonal, and 0's elsewhere. If M is another $n \times n$ matrix, then
$$IM = MI = M.$$

- *Some* square matrices M are *invertible*: The inverse matrix M^{-1} of M is the unique matrix satisfying

$$M \cdot M^{-1} = M^{-1} \cdot M = I.$$

- Matrices in numpy:
  - Initializing a matrix: Here is our lily pad example.

    ```
    m=array([[0,0.5,0.5],[0.5,0,0.5],[1./3,1./6,0.5]])
    ```

  - Multiplying, adding and subtracting matrices:

    ```
    m.dot(n),m+n,m-m
    ```

    Warning: `m*n` is componentwise multiplication of matrices, NOT the matrix product discussed above.
  - `eye(n)` is the $n \times n$ identity matrix.
  - `inv(m)` returns the inverse of the matrix `m`, if it exists. (Be careful here, because roundoff error can turn a non-invertible matrix into an invertible one.)

Next installment: Analysis of Regular Markov Chains