IEOR 3106: Professor Whitt

Lecture Notes, Thursday, September 28, 2006

More on Markov Chains

Liberating Markov Mouse: The Open Maze

We started a week ago by considering how to model a mouse moving around in a maze. The maze is a closed space containing nine rooms. The space is arranged in a three-by-three array of rooms, with doorways connecting the rooms. See the first Markov mouse notes.

Suppose that we let the mouse escape. Suppose that we put a door leading out of the maze from rooms 3, 7 and 9. Let the mouse still choose one of the available doors out of each room with equal probability. Then the probabilities of leaving these corner rooms 3, 7 and 9 through one of these new doors is 1/3. As before, we let the mouse choose each of the available doors with equal probability, independently of how the mouse got to its present room. Assume that the mouse doe not re-enter the maze once it has left.

We now pay attention to the door from which the mouse leaves the maze. To do so, we add three new states: 10, 11 and 12. These states 10, 11 and 12 mean, respectively that the mouse left the maze from the door out of room 3, 7 and 9. We thus obtain the 12-state absorbing Markov chain with 3 absorbing states that appears on the course Matlab web page. The non-absorbing states are called *transient* states. To analyze this problem with matlab, use the program absorbing.m on this example.

3. Analyzing an Absorbing Chain

We now indicate how to analyze an absorbing Markov chain. This analysis applies to the absorbing Markov chain we have just defined, but also to other absorbing Markov chains. We first label the states so that all the absorbing states appear first, and then afterwards we put the transient states (the states that we will eventually leave, never to return. The transition matrix then has the block matrix form

$$P = \begin{pmatrix} I & 0 \\ R & Q \end{pmatrix} ,$$

where I is an identity matrix (1's on the diagonal and 0's elsewhere) and 0 (zero) is a matrix of zeros. In this case, I would be 3×3, R is 9×3 and Q is 9×9). The matrix Q describes the probabilities of motion among the transient states, while the matrix R gives the probabilities of absorption in one step (going from one of the transient states to one of the absorbing states in a single step). In general Q would be square, say m by m, while R would be m by k, and I would be k by k.

3.1 The Fundamental Matrix N

First suppose that we want to calculate the expected number of times the chain spends in transient state j starting in transient state i. Let $T_{i,j}$ be the total number times and let $N_{i,j} \equiv E[T_{i,j}]$ be the expected number of times. It is convenient to write

$$T_{i,j} = T_{i,j}^{(0)} + T_{i,j}^{(1)} + T_{i,j}^{(2)} + T_{i,j}^{(3)} + T_{i,j}^{(4)} + T_{i,j}^{(5)} + \cdots$$

where $T_{i,j}^{(k)}$ is the number of times at the k^{th} transition. Clearly, $T_{i,j}^{(k)}$ is a random variable that is either 1 (if the chain is in transient state j on the k^{th} transition) or 0 (otherwise). By definition, we say that $T_{i,j}^{(0)} = 1$ if $i = j$, but $= 0$ otherwise. Since these random variables assume only the values 0 and 1, we have

$$\begin{aligned} N_{i,j} \equiv E[T_{i,j}] &= E[T_{i,j}^{(0)} + T_{i,j}^{(1)} + T_{i,j}^{(2)} + T_{i,j}^{(3)} + T_{i,j}^{(4)} + T_{i,j}^{(5)} + \cdots] \\ &= E[T_{i,j}^{(0)}] + E[T_{i,j}^{(1)}] + E[T_{i,j}^{(2)}] + E[T_{i,j}^{(3)}] + E[T_{i,j}^{(4)}] + E[T_{i,j}^{(5)}] + \cdots \\ &= P(T_{i,j}^{(0)} = 1) + P(T_{i,j}^{(1)} = 1) + P(T_{i,j}^{(2)} = 1) + P(T_{i,j}^{(3)} = 1) + P(T_{i,j}^{(4)} = 1) + \cdots \\ &= Q_{i,j}^0 + Q_{i,j}^1 + Q_{i,j}^2 + Q_{i,j}^3 + Q_{i,j}^4 + Q_{i,j}^5 + \cdots \end{aligned}$$

To summarize,

$$N_{i,j} \equiv Q_{i,j}^{(0)} + Q_{i,j}^{(1)} + Q_{i,j}^{(2)} + Q_{i,j}^{(3)} + \cdots .$$

In matrix form, we have

$$\begin{aligned} N &= Q^{(0)} + Q^{(1)} + Q^{(2)} + Q^{(3)} + \cdots \\ &= I + Q + Q^2 + Q^3 + \cdots \end{aligned}$$

where the identity matrix I here has the same dimension m as Q. (Since Q is the submatrix corresponding to the transient states, $Q^n \to 0$ as $n \to \infty$, where here 0 is understood to be a matrix of zeros.)

Multiplying by $(I - Q)$ on both sides, we get a simple formula, because there is cancellation on the righthand side. In particular, we get

$$(I - Q) * N = I ,$$

so that, multiplying on the left by the inverse $(I - Q)^{-1}$, which can be shown to exist, yields

$$N = (I - Q)^{-1}$$

In MATLAB, we would write $N = inv(I - Q)$. The matrix $N = (I - Q)^{-1}$ is called the *fundamental matrix of an absorbing Markov chain.*

We can be a little more careful and write

$$N_n \equiv I + Q + Q^2 + Q^3 + \cdots + Q^n .$$

Then the cancellation yields

$$(I - Q)N_n = I - Q^{n+1} .$$

We use the fact that Q is the part of P corresponding to the transient states, so that Q^n converges to a matrix of zeros as $n \to \infty$. Hence, $(I - Q)N_n \to I$ as $n \to \infty$. Writing

$$(I - Q)N = I ,$$

we see that both N and $I - Q$ are nonsingular, and thus invertible, yielding $N = (I - Q)^{-1}$, as stated above.

3.2 Other Quantities of Interest

Other quantities of interest can be computed using the fundamental matrix N, as we now show.

Let M_i be the expected number of steps until absorption starting in transient state i and let M be the $m \times 1$ column vector with elements M_i. The total number of steps until absorption is the sum of the numbers of steps spent in each of the transient states before absorption (always starting in transient state i). Hence,

$$M_i = N_{i,1} + N_{i,2} + \cdots + N_{i,m} \ ,$$

assuming, as before, that there are m transient states. In matrix form,

$$M = N * w \ ,$$

where w is a $m \times 1$ column vector of ones.

Let $B_{i,l}$ be the probability of being absorbed in absorbing state l starting in transient state i. Breaking up the overall probability into the sum of the probabilities of being absorbed in state l in each of the possible steps, we get

$$B_{i,l} = R_{i,l} + (Q * R)_{i,l} + (Q^2 * R)_{i,l} + \cdots$$

so that

$$\begin{aligned} B &= R + Q * R + Q^2 * R + Q^3 * R + \cdots \\ &= (I + Q + Q^2 + \cdots) * R \\ &= N * R \ . \end{aligned}$$

Hence, $B = NR$, where N is the fundamental matrix above.

In summary, it is easy to compute the matrices N, M and B describing the evolution of the absorbing Markov chain, given the key model elements - the matrices Q and R.

You should do the escaping Markov mouse example using MATLAB. The MATLAB program absorbing.m does that for you. The data and the program are on my web page on the computational-tools page.

For some related material, see Sections 4.5.1 and 4.6 of Ross.

4. Summary

1. There are two kinds of basic Markov chains, those considered in so far in these notes.

2. The first kind, illustrated by the closed maze, has a unique stationary distribution, obtained by solving $\pi = \pi P$.

3. The second kind, illustrated by the open maze (the escaping mouse), is an absorbing Markov chain.

4. You ask different questions for the two kinds of chains.
(And so, there are different tools and different answers.)

5. There are even more complicated Markov chains, but we usually decompose them and analyze component chains of the two types above.

6. Finite matrices can be applied only when the state space is finite. However, the main story (and results) go over to infinite-state Markov chains. Indeed, Ross discusses the more general case from the beginning. A simple example is a simple random walk, in which you go to the right or left at each step, each with probability 1/2.

5. Classification of States

See Section 4.3 on p. 189.

Concepts:

1. State j is *accessible* from state i if it is possible to get to j from i in some finite number of steps. (notation: $i \rightsquigarrow j$)

2. States i and j *communicate* if both j is accessible from i and i is accessible from j. (notation: $i \leftrightsquigarrow j$)

3. A subset A of states in the Markov chain is *a communication class* if every pair of states in the subset communicate.

4. A communication class A of states in the Markov chain is *closed* if no state outside the class is accessible from a state in the class.

5. A communication class A of states in the Markov chain is *open* if it is not closed; i.e., if it is possible for the Markov chain to leave that communicating class.

6. A Markov chain is *irreducible* if the entire chain is a single communicating class.

7. A Markov chain is *reducible* if there are two or more communication classes in the chain; i.e., if it is not irreducible.

8. A Markov chain transition matrix P is in *canonical form* if the states are re-labelled (re-ordered) so that the states within closed communication classes appear together first, and then afterwards the states in open communicating classes appear together. The recurrent states appear at the top; the transient states appear below. The states within a communication class appear next to each other.

9. State j is a *recurrent state* if, starting in state j, the Markov chain returns to state j with probability 1

10. State j is a *transient state* if, starting in state j, the Markov chain returns to state j with probability < 1; i.e., if the state is *not* recurrent.

11. State j is a *positive-recurrent state* if the state is recurrent and if, starting in state j, the expected time to return to that state is finite.

12. State j is a *null-recurrent state* if the state is recurrent but, , starting in state j, the expected time to return to state j is infinite.

6. Canonical Form for a Probability Transition Matrix

Find the canonical form of the following Markov chain transition matrix:

(a)

$$P = \begin{pmatrix} 0.1 & 0.0 & 0.0 & 0.9 & 0.0 \\ 0.0 & 0.4 & 0.0 & 0.0 & 0.6 \\ 0.3 & 0.3 & 0.0 & 0.4 & 0.0 \\ 0.3 & 0.0 & 0.0 & 0.7 & 0.0 \\ 0.0 & 0.7 & 0.0 & 0.0 & 0.3 \end{pmatrix}$$

Notice that the sets $\{1,4\}$ and $\{2,5\}$ are closed communicating classes containing recurrent states, while $\{3\}$ is an open communicating class containing a transient state.

So you should reorder the states according to the order: $1,4,2,5,3$. The order $2,5,1,4,3$ would be OK too, as would $5,2,4,1,3$. We put the recurrent states first and the transient states last. We group the recurrent states together according to their communicating class. Using the first order - $1,4,2,5,3$ - you get

$$P = \begin{pmatrix} 0.1 & 0.9 & 0.0 & 0.0 & 0.0 \\ 0.3 & 0.7 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.4 & 0.6 & 0.0 \\ 0.0 & 0.0 & 0.7 & 0.3 & 0.0 \\ 0.3 & 0.4 & 0.3 & 0.0 & 0.0 \end{pmatrix}$$

Notice that the canonical form here has the structure:

$$P = \begin{pmatrix} P_1 & 0 & 0 \\ 0 & P_2 & 0 \\ R_1 & R_2 & Q \end{pmatrix},$$

where P_1 and P_2 are 2×2 Markov chain transition matrices in their own right, whereas R_i is the one-step transition probabilities from the single transient state to the i^{th} closed set. In this case, $Q \equiv (0)$ is the 1×1 sub-matrix representing the transition probabilities among the transient states. Here there is only a single transient state and the transition probability from that state to itself is 0. The chain leaves that transient state immediately, never to return..