IEOR 3106: Professor Whitt

Lecture Notes, Tuesday, September 24, 2013

More Markov chains

1. Classification of States

See Section 4.3. Discussed in last class.

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: $i \rightsquigarrow j$) [By definition, we allow 0 steps, so that $i \rightsquigarrow i$ for each i.]

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

9. 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.

10. A Markov chain transition matrix P is in *canonical form* if the states are re-labeled (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. It usually looks better if the closed classes are ordered by size, with the smaller ones at the top.

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. [In a finite-state Markov chain, all recurrent states are positive recurrent.]

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. That only can occur in DTMC's with infinitely many states. An example is a symmetric random walk on the nonnegative integers. From every i the walk goes up to $i+1$ with probability $1/2$ and down to $i-1$ with probability $1/2$.

At 0 it goes up to 1, or goes up to 1 with probability 1/2 and stays put at 0 with probability 1/2.

2. Canonical Form for a Probability Transition Matrix

Did example in last class. Find the canonical form of the following Markov chain transition matrix:

(a)

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

(I label the states 1, 2, 3, 4, 5.)

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.

3. More on Absorbing Markov Chains

It is important to see that there are **two basic kinds of markov chains**: (i) **irreducible Markov chains** and (ii) **reducible Markov chains**. In an irreducible Markov chain, it is possible to get to any other state in some finite number of moves; in a reducible Markov chain, it is not. Prominent among the reducible Markov chains are the **absorbing Markov chains**. In absorbing Markov chains there are states that the chain can enter, but once entered, the Markov chain cannot leave these states.

An example of an absorbing Markov chain appeared in the last part of the Markov mouse notes from before. We discussed that to some extent.

The analysis is different for irreducible Markov chains and absorbing Markov chains. For irreducible Markov chains (but not for absorbing Markov chains), we have the fundamental equation $\pi = \pi P$. In contrast, for absorbing Markov chains, we have the fundamental matrix $N = (I - Q)^{-1}$ and the associated matrix of absorption probabilities $B = NR$.

It is important to recognize that these formulas and equations are being expressed concisely in **matrix notation**. For example, the equation $\pi = \pi P$ is a matrix equation; the π appearing there is a probability vector, while P is the square matrix of transition probabilities. If there are n states, then $\pi \equiv (\pi_1, \ldots, \pi_n)$ is a $1 \times n$ matrix (and thus a row vector), while P is an $n \times n$ square matrix with entries $P_{i,j}$. (Thus, the matrix multiplication πP is indeed well defined, and the product must itself be a $1 \times n$ row vector. Thus, the equation $\pi = \pi P$ corresponds to a system of equations that must be satisfied by the steady-state probability mass function, denoted by the probability vector π. In particular, the n equations are

$$\pi_j = \sum_{i=1}^{n} \pi_i P_{i,j}, \quad 1 \le j \le n.$$

To solve this system of equations, we also need the additional equation

$$\sum_{i=1}^{n} \pi_i = 1.$$

Similarly, when we consider an absorbing Markov chain, $N = (I - Q)^{-1}$ is a matrix inverse. We now review Section 3.2.3 of the Markov mouse notes.

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. That means that $B_{i,j}$ is the $(i,j)^{\text{th}}$ element of the matrix NR, which in turn is the number obtained by multiplying the i^{th} row of N by the j^{th} column of R.

4. Gambler's Ruin Problem

The gambler's ruin problem in Section 4.5.1 is an example of an absorbing Markov chain. Indeed, Sections 4.5.1 and 4.6 in the textbook concern absorbing Markov chains and thus relate to the lecture notes on liberating Markov mouse. These cover the same subject from different points of view.

In this problem there are two absorbing states: 0 and N. The remaining states $1, 2, \ldots, N-1$ are transient states. The different treatment in Section 4.5.1 is mostly because that particular model is highly structured, so that there are only a few parameters: the probability of going up one, p (the only other possibility is to go down one, which happens with probability $q = 1-p$), the initial position i and the target N. The goal is to compute the probability P_i of reaching N starting in i before reaching 0. The states 0 and N are absorbing states. The probability P_i can also be found by looking at the correct component of the matrix $B = NR$ in the general absorbing Markov chain theory. In particular, we look at $B_(i, N)$, where the states have their original labels.

However, In this highly structured example, there is an explicit revealing formula for the probability P_i. (See Section 4.5.1.) In particular,

$$P_i = \frac{1-(q/p)^i}{1-(q/p)^N} \quad \text{if} \quad p \neq 1/2,$$

while

$$P_i = \frac{i}{N} \quad \text{if} \quad p = 1/2.$$

In summary, there are three five things to observe here:

1. This is an example of an absorbing Markov chain.
2. Because of the special structure, there is an analytical formula.
3. That formula can be derived by a simple recursive argument.
4. That formula is informative; there is a powerful exponential relation, if $p \neq 1/2$.
5. By symmetry, we can obtain the result for $p > 1/2$ directly from the result for $p < 1/2$.

Elaborating on the fourth point, if $q/p > 1$, then the probability of going down is greater than the probability of going up. Then we should rewrite the expression as

$$P_i = \frac{(q/p)^i - 1}{(q/p)^N - 1}.$$

For larger values of i and n, we have $(q/p)^i >> 1$. Assuming that is the case, we have the approximation

$$P_i = \frac{(q/p)^i - 1}{(q/p)^N - 1} \approx \frac{(q/p)^i}{(q/p)^N} = (p/q)^{N-i} = e^{-\lambda(N-i)},$$

where

$$\lambda \equiv -\log_e(p/q) = \log_e((1-p)/p) > 0.$$

where $(1-p)/p > 1$ because $1-p > p$ since $p < 1/2$. Thus P_i gets exponentially small as $N-i$ increases, and the rate is λ above.

5. Cars and Trucks. We discussed **Exercise 4.30 on cars and trucks**, which was in the homework.

Problem statement. Three out of every four trucks on the road are followed by a car, while only one out of every five cars is followed by a truck. What fraction of vehicles on the road are trucks?

Intended solution. The intent of course is for you to create a Markov chain model. It is easy to set up the matrix P. In particular, we get

$$P = \begin{pmatrix} 4/5 & 1/5 \\ 3/4 & 1/4 \end{pmatrix} ,$$

where the rows and columns are labeled C and T, in that order. Then it is easy to solve $\pi = \pi P$; remember that one equation is always redundant. So that leaves only a single equation. You get an extra equation by noting that the probabilities (the entries of the vector π) must sum to 1. Then you get $\pi = (\pi_C, \pi_T) = (15/19, 4/19)$.

But not so fast! We now observe that the problem wording does not actually imply a Markov chain, but if you use a Markov chain, it turns out that you do get the right answer (whether or not it is a Markov chain). Here is a counterexample showing that the problem formulation does not imply that the system evolves as a Markov chain: The counterexample is a deterministic periodic sequence with the period containing 19 vehicles: 4 trucks and 15 cars. A single period of the deterministic sequence looks like:

$$TTCTCTCCCCCCCCCCCCC$$

The entire sequence repeats that pattern over and over again, with no randomness at all:

$$|TTCTCTCCCCCCCCCCCCC|TTCTCTCCCCCCCCCCCCC|TTCTCTCC\ldots$$

This deterministic pattern satisfies the problem statement, but it is not a Markov chain. Interestingly (but not demonstrated), the Markov chain analysis gives the right answer, even if the vehicles do not appear as a Markov chain.

The solution more generally. How can we solve the problem more generally?

Let C_n be the number of cars among the first n vehicles on the road. We can then relate the number of cars among the first n vehicles to the number of cars among vehicles 2 through $n+1$ by counting the number among following vehicles. Thus, we would directly get the relation

$$C_{n+1} - C_1 = T_n(3/4) + C_n(4/5),$$

where also

$$C_n + T_n = n.$$

Combining these two equations, we get

$$C_{n+1} - C_1 = (n - C_n)(3/4) + C_n(4/5) = n(3/4) + C_n(1/20). \tag{1}$$

Of course, we should not expect equality for all n, but only correctness of proportions asymptotically as $n \to \infty$. We should expect that, as $n \to \infty$, $C_n/n \to C$. Dividing by n in (1) and letting $n \to \infty$, we get

$$C = 3/4 + C(1/20) \quad \text{or} \quad C(19/20) = 3/4,$$

so that

$$C = 15/19,$$

The same as above.

6. The Bonus-Malus Insurance Problem

This is **the insurance example, Example 4.7** in the textbook. This is a slight simplification of a realistic application of Markov chains. We construct the model by constructing the transition matrix, given in the book:

$$P = \begin{pmatrix} a_0 & a_1 & a_2 & 1 - a_0 - a_1 - a_2 \\ a_0 & 0.0 & a_1 & 1 - a_0 - a_1 \\ 0.0 & a_0 & 0.0 & 1 - a_0 \\ 0.0 & 0.0 & a_0 & 1 - a_0 \end{pmatrix}$$

There are four states: $1, 2, 3, 4$, with state 1 being the bonus (good) state. Higher states mean a worse claim record, but in this case we place an upper limit of 4.

The elements of the 4×4 matrix are probabilities: a_k is the probability of having k accidents (actually making k claims) during the year. The row sums necessarily are 1. The probability entries a_k here are in this case in the book are computed as Poisson probabilities; i.e., then

$$a_k = \frac{e^{-\lambda} \lambda^k}{k!} \quad \text{for} \quad k \geq 0 \ .$$

It is natural to use the Poisson distribution to describe the probability a person makes some number of claims. The parameter λ is the mean of that distribution. The parameter λ might well depend on the driver or type of driver. We then use the model to describe the evolving state of the drivers (as a function of λ) and the long-run distribution of premiums, given the premiums charged per state. To appreciate what you can do with the model, you need to learn what you can do to describe performance, after you have built the model. We will be talking about that. You will be reading about that. For example, here you might consider what appropriate premiums in each state should be. This depends on the goals of course, but also on the way the Markov chain evolves.

Suppose the premium in state k is v_k. What is the average premium per year for a customer who has an average of λ accidents per year. We can calculate the steady state probability vector; i.e., we can easily calculate

$$\pi = \pi P$$

with $\pi_1 + \cdots + \pi_4 = 1$, starting from the matrix P with the elements above. (That is easily done on the computer.) Then the average premium per year, which we can represent as the expected premium in steady state, is

$$\text{average premium} \quad = \sum_{k=1}^{4} \pi_k v_k \ .$$

We now discuss a simple approximation, which makes the calculation easy to do by hand. A simple approximation for the case in which accidents are deemed quite rare is to assume that $a_1 = \epsilon$ and $a_k = 0$ for $k > 1$. In other words, we assume that a person has at most one accident each year. We use the notation ϵ because that is common to denote a small positive number. Then necessarily $a_0 = 1 - \epsilon$ and the transition matrix becomes

$$P = \begin{pmatrix} 1 - \epsilon & \epsilon & 0.0 & 0.0 \\ 1 - \epsilon & 0.0 & \epsilon & 0.0 \\ 0.0 & 1 - \epsilon & 0.0 & \epsilon \\ 0.0 & 0.0 & 1 - \epsilon & \epsilon \end{pmatrix}$$

For this special case, it is easy to calculate the steady state probability vector; i.e., we can easily calculate

$$\pi = \pi P$$

with $\pi_1 + \cdots + \pi_4 = 1$. We get

$$\pi = (1/s, c/s, c^2/s, c^3/s) \quad \text{where} \quad s = 1 + c + c^2 + c^3 \quad \text{and} \quad c = \frac{\epsilon}{1-\epsilon} \ .$$

We can now see, via the formula, how the answer depends on the single remaining parameter ϵ.

You might ask: What is the expected number of years starting from the bonus (good) state of 1 until a person next returns to this same state? The answer is

$$1/\pi_1 = s = 1 + c + c^2 + c^3 \quad \text{where} \quad c = \frac{\epsilon}{1-\epsilon}$$

see the Remarks right before Example 4.22 in the book.