

IEOR 3106: Introduction to Operations Research: Stochastic Models

Professor Whitt

SOLUTIONS to Numerical Part of Homework Assignment 3

Markov Chains

Part I. Computational Exercises with Matlab

1. Markov Mouse

Refer to the nine-state model of Markov Mouse moving around within the closed maze presented in class.

(a) Exploit the structure of the transition matrix (without doing elaborate computation) to deduce what the following transition probabilities must be (exactly or approximately):

These probabilities are now given:

$$\begin{array}{llll} P_{1,1}^3 = 0 & P_{1,1}^{17} = 0 & P_{4,4}^{17} = 0 & P_{1,2}^{18} = 0 \\ P_{1,2}^{21} \approx 1/4 & P_{1,6}^{21} \approx 1/4 & P_{3,4}^{21} \approx 1/4 & P_{1,1}^{22} \approx 1/6 \\ P_{1,9}^{22} \approx 1/6 & P_{9,1}^{22} \approx 1/6 & P_{1,5}^{22} \approx 1/3 & P_{5,5}^{22} \approx 1/3 \end{array}$$

Explanation: The first four multi-step transition probabilities above are all 0, because the chain is periodic with period 2: You can go from an odd-numbered state to an odd-numbered state only in an even number of steps. Similarly, you can go from an even-numbered state to an even-numbered state only in an even number of steps.

By symmetry the long run probability of being in each of the even-numbered states must be identical. Thus the first three non-zero transition probabilities above should be approximately 1/4.

We can compute the long run probability of going from an odd-numbered state to the state 1 (and also the states 3, 7 and 9) in an even number of steps in the long run (e.g., in 14 – 22 steps), using the determined long-run probabilities for being in the neighboring even-numbered state the step before:

$$\begin{aligned} P_{1,1}^{22} &= P_{1,2}^{21} * P_{2,1}^1 + P_{1,4}^{21} * P_{4,1}^1 \\ &\approx \frac{1}{4} \times \frac{1}{3} + \frac{1}{4} \times \frac{1}{3} \\ &\approx 1/6 . \end{aligned}$$

That same argument works for $P_{9,1}^{22}$ and $P_{i,j}^{22}$ where i and j each can be one of the state 1, 3, 7 or 9. It also applies for those j if $i = 5$.

Finally, $P_{i,5}^{22} \approx 1/3$ when i is one of 1, 3, 5, 7 or 9, because we subtract 4 times the probability of being in each of the other odd-numbered states: $1 - 4 \times (1/6) = 1/3$.

We thus can uncover the long-run behavior of the Markov chain in this case by doing only elementary calculations.

(b) Using MATLAB or another computer program, calculate the following matrices:

$$P^{14} \quad P^{16} \quad P^{17} \quad (P^{16} + P^{17})/2$$

(You can turn in computer printout or just write down a few entries.)

Look at the MATLAB output displayed for the (stationary) MarkovMouse example (PDF) with link from the Lectures link on the course web page. The matrix data are there too.

You see the periodic structure when you look at P^{14} , P^{16} and P^{17} : The matrices P^{14} and P^{16} have zeros but look the same. The matrix $(P^{16} + P^{17})/2$ has all rows equal to the stationary vector. We take the average of two successive powers of the matrix P because the Markov chain has period 2.

(c) Using one of the MATLAB programs **stat.m** or **stationary.m** (which can be downloaded from the course web page), or your own program, calculate the stationary probability vector of P . Compare your answer to the calculations in part (b).

The displayed MATLAB output on the web page show what happens.

(d) What is the long-run proportion of time (steps) that the mouse spends in Room 5?

Find the π such that $\pi = \pi P$. You get $\pi_2 = \pi_4 = \pi_6 = \pi_8 = 1/8 = 0.1250$, $\pi_1 = \pi_3 = \pi_7 = \pi_9 = 1/12 = 0.08333$, and $\pi_5 = 1/6 = 0.16667$. So the answer is $\pi_5 = 0.1667$

(e) Is there a unique stationary probability vector? If so, what is the stationary probability that the mouse spends in Room 5?

Yes, we found it in part (d) in order to answer the question there.

(f) What is the limit (as $n \rightarrow \infty$) of the probability that the mouse is in Room 5 at step n , given that the mouse starts in Room 1?

This Markov chain is periodic with period 2, so there is no simple limit. For large n , $P_{1,5}^{2n+1} = 0$ and $P_{1,5}^{2n} \simeq 2\pi_5 = 0.3334$. The probabilities approach these alternating values.

(g) What is the expected number of transitions between successive visits to Room 5?

$1/\pi_5 = 6$. See Remark (ii) on page 212.

2. Escaping Markov Mouse

Now allow Markov mouse to escape from the nine-room maze. In this example we allow escape from rooms 3, 7 and 9 through doors to the outside. Once the mouse escapes, it never returns. Let the probability of leaving each room through any one of the available doors be equally likely. Thus the probability of escaping from each of the rooms 3, 7 and 9, after a visit to that room, is $1/3$ in each case. Suppose that we want to keep track of the room from which the mouse eventually leaves the maze; i.e., exiting from Room 3 is to be distinguished from exiting from Room 7 or Room 9.

(a) What is the overall probability transition matrix to model the movement of the mouse now?

This example is directly an absorbing Markov chain. The matrix can be written in the block-matrix form

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

where I is a 3×3 identity matrix, Q is a 9×9 matrix and R is a 9×3 matrix. The matrices Q and R are provided as MATLAB data files on the web page.

(b) Using the MATLAB program **absorbing.m** (which can be downloaded from the course web page), or your own program, calculate the three matrices N , m and B describing the behavior of this absorbing Markov chain.

The analysis with the program **absorbing.m** is displayed on the web page.

(c) Starting in room 1, what is the probability of the mouse eventually leaving the maze from room 7?

$$B_{1,7} = 0.3684$$

Starting in room 1, what is the probability of the mouse eventually leaving the maze from room 9?

$$B_{1,9} = 0.2632$$

Starting in room 9, what is the probability of the mouse eventually leaving the maze from room 3?

$$B_{9,3} = 0.2105$$

(d) Starting in room 1, what is the expected number of transitions before the mouse leaves the maze?

$$m_1 = 14.263$$

(e) Starting in room 2, what is the expected number of visits to room 4 before the mouse leaves the maze?

$$N_{2,4} = 1.56$$

3. The Canonical Form for a Markov Transition Matrix

The canonical form of a Markov chain transition matrix has the states ordered so that the states in the same communicating class appear together. Moreover, the ergodic (closed communicating) classes appear first, with the transient states appearing below.

The canonical form looks like:

$$P = \begin{pmatrix} P_1 & & & \\ & P_2 & & \\ & & P_3 & \\ R_1 & R_2 & R_3 & Q \end{pmatrix}$$

where the entries of this matrix P are submatrices, with Q and P_i for all i being square matrices. The empty spaces are understood to be matrices of 0's. If Q is $m \times m$ and P_i is $k \times k$, then R_i is $m \times k$. The matrices P_i are understood to be irreducible Markov transition matrices. The states associated with the submatrix Q are the transient states. From any state associated with the submatrix Q , the chain will visit that state only finitely often.

Find the canonical form of each of the following Markov transition matrices:

(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

(a)

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

(b)

$$P = \begin{pmatrix} 0.0 & 0.2 & 0.0 & 0.0 & 0.8 \\ 0.0 & 0.0 & 0.0 & 1.0 & 0.0 \\ 1.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 1.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 1.0 & 0.0 \end{pmatrix}$$

Note that the set $\{1, 2, 4, 5\}$ is a closed communicating class, containing recurrent states, while $\{3\}$ is an open communicating class, containing a single transient state. Thus, an appropriate order of the states (one of several) is 1, 2, 4, 5, 3.

With that ordering of the states, you get

$$P = \begin{pmatrix} 0.0 & 0.2 & 0.0 & 0.8 & 0.0 \\ 0.0 & 0.0 & 0.1 & 0.0 & 0.0 \\ 1.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 1.0 & 0.0 & 0.0 \\ 1.0 & 0.0 & 0.0 & 0.0 & 0.0 \end{pmatrix}$$

(c)

$$P = \begin{pmatrix} 0.1 & 0.2 & 0.1 & 0.0 & 0.2 & 0.2 & 0.2 & 0.0 \\ 0.0 & 0.2 & 0.0 & 0.0 & 0.0 & 0.8 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.1 & 0.0 & 0.0 & 0.0 & 0.9 & 0.0 \\ 0.0 & 1.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.1 & 0.2 & 0.1 & 0.6 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 1.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.7 & 0.0 & 0.0 & 0.0 & 0.3 & 0.0 \\ 0.1 & 0.0 & 0.0 & 0.4 & 0.0 & 0.0 & 0.0 & 0.5 \end{pmatrix}$$

Note that $\{2, 6\}$ and $\{3, 7\}$ are two closed communication classes, containing recurrent states, while $\{4\}$, $\{1, 5\}$ and $\{8\}$ are open communicating classes, containing transient states. Since you can go from $\{8\}$ to $\{1, 5\}$ and $\{4\}$ and from $\{1, 5\}$ to $\{4\}$, it is natural to order the transient states as 4, 1, 5, 8 (with state 4 closest to the recurrent states). Thus it is natural to order all the states by: 2, 6, 3, 7, 4, 1, 5, 8. The detailed ordering of the 4 transient states among themselves has not been emphasized, but what we have done includes such an ordering.

With that ordering, you get the canonical form:

$$P = \begin{pmatrix} 0.2 & 0.8 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 1.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.1 & 0.9 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.7 & 0.3 & 0.0 & 0.0 & 0.0 & 0.0 \\ 1.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.2 & 0.2 & 0.1 & 0.2 & 0.0 & 0.1 & 0.2 & 0.0 \\ 0.2 & 0.0 & 0.1 & 0.0 & 0.6 & 0.1 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.4 & 0.1 & 0.0 & 0.5 \end{pmatrix}$$

In block-matrix form, it looks like

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

where P_1 is 2×2 , P_2 is 2×2 , Q is 4×4 , and both R_1 and R_2 are 4×2 .

4. More on Stationary Vectors

Calculate the stationary vector of the following Markov transition matrix. Compare to high powers of the matrix P .

$$P = \begin{pmatrix} 0 & 0 & .25 & .75 & 0 & 0 \\ 0 & 0 & .5 & .5 & 0 & 0 \\ 0 & 0 & 0 & 0 & .875 & .125 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

The Markov chain here is periodic with period 3. Thus, as with other periodic chains, high powers will always contain zeros. We see limiting behavior if we look at P^{3k+i} for successive k , where i is fixed at 0, 1 or 2. We see the stationary probability vector if we look at $(P^{21} + P^{22} + P^{23})/3$.

In this case, the stationary probability vector is

$$(0.11970.21370.13680.19660.11970.2137) .$$

See the six-state periodic chain example (PDF) on the web page. The data for the matrix P is given there as well.
