

IEOR 4106: Professor Whitt

Lecture Notes

Introduction to Markov Chains

1. Markov Mouse: The Closed Maze

We start 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, as shown in the figure below

The Maze

1	2	3
4	5	6
7	8	9

There are doors leading to adjacent rooms, vertically and horizontally. In particular, there are doors

from 1 to 2, 4
from 2 to 1, 3, 5
from 3 to 2, 6
from 4 to 1, 5, 7
from 5 to 2, 4, 6, 8
from 6 to 3, 5, 9
from 7 to 4, 8
from 8 to 5, 7, 9

from 9 to 6,8

We assume that the mouse is a Markov mouse; i.e., the mouse moves randomly from room to room, with the probability distribution of the next room depending only on the current room, not on the history of how it got to the current room. (This is the **Markov property**.) Moreover, we assume that the mouse is equally likely to choose each of the available doors in the room it occupies. (That is a special property beyond the Markov property.)

We now model the movement of the mouse as a Markov chain. The **state** of the Markov chain is the room occupied by the mouse. We let the time index n refer to the n^{th} room visited by the mouse. So we make a discrete-time Markov chain. Specifically, we let X_n be the state (room) occupied by the mouse on step (or time or transition) n . The initial room is X_0 . The room after the first transition is X_1 , and so forth. Then $\{X_n : n \geq 0\}$ is the **discrete-time Markov chain (DTMC)**; it is a discrete-time discrete-state stochastic process. The mouse is in room X_n (a random variable) after making n moves, after having started in room X_0 , which could also be a random variable.

We specify the evolution of the Markov chain by specifying the one-step transition probabilities. We specify these transition probabilities by specifying the **transition matrix**. For our example, making a discrete-time Markov chain model means that we define a 9×9 Markov transition matrix consistent with the specification above. For the most part, **specifying the transition matrix is specifying the model**. (We also must say how we start. The starting point could be random, in which case the initial position would be specified by a probability vector.)

Notation. It is common to denote the transition matrix by P and its elements by $P_{i,j}$; i.e., $P_{i,j}$ denotes the probability of going to state j next when currently in state i . (When the state space has m states (is finite), P is a square $m \times m$ matrix.) ■

For example, for our Markov mouse model, we have $P_{1,2} = 1/2$ and $P_{1,4} = 1/2$, with $P_{1,j} = 0$ for all other j , $1 \leq j \leq 9$. And we have $P_{2,1} = P_{2,3} = P_{2,5} = 1/3$ with $P_{2,j} = 0$ for all other j . And so forth.

Here is the total transition matrix:

$$P = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \\ 9 \end{matrix} & \begin{pmatrix} 0 & 1/2 & 0 & 1/2 & 0 & 0 & 0 & 0 & 0 \\ 1/3 & 0 & 1/3 & 0 & 1/3 & 0 & 0 & 0 & 0 \\ 0 & 1/2 & 0 & 0 & 0 & 1/2 & 0 & 0 & 0 \\ 1/3 & 0 & 0 & 0 & 1/3 & 0 & 1/3 & 0 & 0 \\ 0 & 1/4 & 0 & 1/4 & 0 & 1/4 & 0 & 1/4 & 0 \\ 0 & 0 & 1/3 & & 1/3 & 0 & 0 & 0 & 1/3 \\ 0 & 0 & 0 & 1/2 & 0 & 0 & 0 & 1/2 & 0 \\ 0 & 0 & 0 & 0 & 1/3 & 0 & 1/3 & 0 & 1/3 \\ 0 & 0 & 0 & 0 & 0 & 1/2 & 0 & 1/2 & 0 \end{pmatrix} \end{matrix} .$$

It is common to label the columns, but we did not above. They are numbered in the same way as the rows. Here the columns are numbered from 1 to 9 starting at the left. Hence, the matrix element in the upper left corner is $P_{1,1}$, while the matrix element in the lower right corner is $P_{9,9}$. The matrix element $P_{2,3}$ appears in the second row and third column. The rows represent the starting state, while the column represents the next step; i.e., $P_{2,3}$ is the probability of going next to state 3 given that you are starting in state 2. The Markov property implies that the probability does not depend on the earlier history. Each time the Markov chain is in state 2, its transition probabilities are the same, independent of how it happened to get there. (We are assuming that the transition probabilities do not depend on either time or the previous states visited.)

We can analyze the evolution of the Markov chain over time (successive transitions) by looking at powers of the matrix P . In particular, **the k -step transition matrix is just the k^{th} power of the matrix P** . Thus, $P^{(k)}$ denotes the k -step transition matrix, while P^k denotes the k^{th} power of the transition matrix P . A main theorem is that $P^{(k)} = P^k$.

You should look at the powers of the matrix P in MATLAB (available, for example, in the IEOR Computer Lab in 301 Mudd) or some other suitable programming language. (Even the TI89 calculator can do matrix multiplication.) You can also access MATLAB from Butler Library and the Mathematics Building. Sample MATLAB code is given in my web-page examples; see the Computational Tools link. Use the program `stat.m` or the program `stationary.m` to analyze this example.

It is important here to know how to do **matrix multiplication**. Suppose that $A \equiv (A_{i,j})$ is a $k \times m$ matrix, while $B \equiv (B_{i,j})$ is an $m \times n$ matrix. Then the matrix product AB is the $k \times n$ matrix with matrix elements

$$(AB)_{i,j} \equiv \sum_{p=1}^m A_{i,p} B_{p,j}$$

Hence, if P is the transition matrix of an m -state DTMC, then P and P^k are $m \times m$ matrices, with

$$P_{i,j}^{(2)} = P_{i,j}^2 \equiv \sum_{p=1}^m P_{i,p} P_{p,j}$$

and

$$P_{i,j}^{(k)} = P_{i,j}^k \equiv \sum_{p=1}^m P_{i,p}^{k-1} P_{p,j} \quad \text{for } k \geq 2 .$$

It is important to see why this matrix multiplication is the correct thing to do. In words, the probability of moving from state i to state j in k steps can be calculated by considering the probability of going from state i to state p in $k-1$ steps and then going next from state p to state j in one step (which is computed by the product $P_{i,p}^{k-1} P_{p,j}$), summed over all possible intermediate states p . This is sometimes called “the law of total probabilities.” The different intermediate states are distinct alternatives; you have to be somewhere at step $k-1$, and you can only be in one state. Summing over all alternatives thus gives the answer.

Much can be done by thinking. Here are some probabilities to estimate:

$$\begin{aligned} P_{1,1}^3 &= ? \\ P_{1,1}^7 &= ? \\ P_{1,9}^7 &= ? \\ P_{2,4}^7 &= ? \\ P_{2,8}^{17} &= ? \\ P_{1,1}^{17} &= ? \\ P_{2,2}^{18} &= ? \\ P_{1,2}^{17} &= ? \\ P_{1,2}^{19} &= ? \\ P_{4,2}^{18} &= ? \\ P_{2,2}^{18} &= ? \end{aligned}$$

$$\begin{aligned}
P_{1,1}^{18} &= ? \\
P_{1,7}^{18} &= ? \\
P_{3,9}^{18} &= ? \\
P_{1,5}^{18} &= ? \\
P_{5,5}^{18} &= ?
\end{aligned}$$

The simple formulas that result can be seen as a consequence of periodicity and symmetry. Note that the state necessarily alternates between even and odd: You go from even to odd to even to odd, and so forth. Hence, 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 in an even number of steps. This is a **periodic** DTMC with **period 2**.

One of the key ideas is that there is **statistical regularity in the long run**. This regularity can be seen by calculating high powers of the matrix P . Because of the periodicity, in this example we do not have P^n converge as $n \rightarrow \infty$. We would have that without the periodicity, but we have periodicity here. The matrix powers converge if we look at the powers P^{2k} or P^{2k+1} . We also obtain convergence if we look at $(P^{2k} + P^{2k+1})/2$ (average of the entries). This convergence is remarkably fast. You should convince yourself by performing calculations.

In order to determine what the long-run probabilities are for this example, you can exploit **symmetry**. By symmetry, the long-run probabilities of being in the states 2, 4, 6 and 8 should be identical. Hence, for large k , we should have

$$P_{2,2}^{2k} \approx P_{2,8}^{2k} \approx P_{4,2}^{2k} \approx P_{8,8}^{2k} \approx \frac{1}{4}$$

In fact, it may be somewhat surprising, but k need not actually be too large. You should verify this with Matlab.

Similarly, the long-run probability of being in the odd states except state 5 should be identical. The odd-numbered states are a bit tricky, but we can work from the even-numbered states:

$$P_{1,1}^{2k} = P_{1,2}^{2k-1}P_{2,1} + P_{1,4}^{2k-1}P_{4,1} \approx \left(\frac{1}{4}\right)\left(\frac{1}{3}\right) + \left(\frac{1}{4}\right)\left(\frac{1}{3}\right) = \frac{1}{6}$$

With Markov chains, we are often interested in the long-run transition probabilities. These transition probabilities often converge as the number of transitions increases, but in this example they do not, because of the periodicity. Because of the periodicity, the probabilities do not simply converge. They alternate between positive values (which converge) and 0.

These properties that can be deduced from detailed analysis can be confirmed by looking at powers of the transition matrix P . So the probabilities above are:

$$\begin{aligned}
P_{1,1}^3 &= 0 \\
P_{1,1}^7 &= 0 \\
P_{1,9}^7 &= 0 \\
P_{2,4}^7 &= 0 \\
P_{2,8}^{17} &= 0 \\
P_{1,1}^{17} &= 0 \\
P_{2,2}^{18} &\approx 1/4
\end{aligned}$$

$$\begin{aligned}
P_{1,2}^{17} &\approx 1/4 \\
P_{1,2}^{19} &\approx 1/4 \\
P_{4,2}^{18} &\approx 1/4 \\
P_{2,2}^{18} &\approx 1/4 \\
P_{1,1}^{18} &\approx 1/6 \\
P_{1,7}^{18} &\approx 1/6 \\
P_{3,9}^{18} &\approx 1/6 \\
P_{1,5}^{18} &\approx 1/3 \\
P_{5,5}^{18} &\approx 1/3
\end{aligned}$$

For any initial probability vector $\alpha \equiv (\alpha_1, \dots, \alpha_m)$, αP gives the new probability vector after one transition. An important concept is the notion of a **stationary probability vector**. An important theorem about **irreducible** finite-state DTMC's states that there exists a unique probability vector $\pi \equiv (\pi_1, \dots, \pi_m)$ such that

$$\pi = \pi P .$$

That is a matrix equation. For the corresponding matrix elements, we have

$$\pi_j = \sum_{k=1}^m \pi_k P_{k,j} \quad \text{for all } j, \quad 1 \leq j \leq m .$$

(A DTMC is irreducible if it is possible to go from any state to any other state in some number of transitions.)

We can relate the positive limiting probability vectors to the stationary probabilities π_j . These positive approximate limiting probabilities are exactly twice the elements of the stationary probability vector π , because of the periodicity. For example, $\pi_2 = 1/8$, while $P_{2,2}^{18} \approx 1/4$.

Looking ahead to future classes: As we will discuss later, the Markov maze here, as well as much larger mazes, can be analyzed very quickly by exploiting the structure of a **random walk on a graph**. The idea is to let the rooms be nodes in the graph. The doors from room to room then become arcs in the graph. There is an arc in the graph between node i and node j whenever there is a positive probability of going from room i to room j in the maze (in one step). Because of the structure of a random walk on a graph, the stationary probability of being in room i turns out to be the number of doors out of room i divided by the sum over the rooms of the number of doors out of each room. That simple formula is a consequence of the detailed balance that holds for this example. Detailed balance is equivalent to the Markov chain being **time reversible**. See Section 4.8 in the textbook. We will get there in due course.

Summary So Far

1. For a finite-state DTMC the model is mostly just a matrix P .

(There is also the initial distribution, a probability vector.)

2. Multi-step transition probabilities are given by powers of the transition matrix P .
3. The long-run behavior can be seen from looking at high powers of the matrix P .

4. We let X_n be the state of the Markov chain after the n^{th} transition; then $\{X_n : n \geq 0\}$ is the stochastic process (Markov chain). The transition probability is

$$P_{i,j} \equiv P(X_{n+1} = j | X_n = i)$$

5. The Markov property asserts that the conditional probability of future states, given the present states and information about past states, depends only on the present state. As a formula, the Markov property states that

$$P(X_{n+1} = j | X_n = i, X_0 = i_0, X_1 = i_1, \dots, X_{n-1} = i_{n-1}) = P(X_{n+1} = j | X_n = i)$$

for all possible past events $\{X_0 = i_0, X_1 = i_1, \dots, X_{n-1} = i_{n-1}\}$.

Additional Points:

6. The long-run behavior can also be found from the stationary probabilities, assuming the Markov chain is irreducible, which it is. (We are also using the fact that the state space is finite. When the state space is infinite, we need to make extra assumptions.) The stationary probability can be found by solving $\pi = \pi P$. The stationary probability vector π is defined as being a probability vector such that, if we then transition according to P , the distribution after the transition is again π . That is, π is a stationary probability vector under the operator P . (In general, the matrix P is thought of as an operator because it maps probability vectors into probability vectors.

(Solving $\pi = \pi P$ is just solving a system of linear equations. Subtle point: there is one redundant equation. The extra equation needed is obtained by the condition that the probabilities should add to 1.)

7. Often, but not always, the stationary probabilities (of an irreducible Markov chain) coincide with the long-run limiting probabilities. They do when the DTMC is aperiodic (period 1). They do not when the Markov chain is periodic. The closed maze is an example of a periodic chain (period 2).

8. If a chain is periodic with period d , then the stationary probabilities can be obtained approximately from the average of d successive high matrix powers; e.g., from $(P^{n+1} + P^{n+2} + \dots + P^{n+d})/d$ for large n .

9. If a chain is periodic with period d , then the long-run limiting probabilities are obtained from the stationary probabilities, e.g., $P_{i,i}^{nd} = d\pi_i$. The probabilities $P_{i,j}^n$ assume the values $d\pi_j$ once and 0 $d - 1$ times as n increases.

10. Sample MATLAB output and programs are given on the computational tools web page for this course.

2. Liberating Markov Mouse: The Open Maze

2.1 The First Absorbing Chain

Suppose that we add a door from room 9 to the outside, and that we allow the mouse to leave room 9 from any of the doors available to it. Now there are three doors out of room 9, one of which takes it outside. Suppose also that the mouse does not reenter the maze after it has left.

This modification significantly changes the problem. Now the mouse will eventually leave the maze with probability one. There is no interesting long-run distribution of being in the different rooms. Eventually the mouse will leave the maze.

However, we can still model the process as a Markov chain. We can add an extra state for the “outside.” We can call the outside state 10. We thus have a 10×10 Markov transition matrix. Now the probabilities of row 9 change: Now we have $P_{9,6} = P_{9,8} = P_{9,10} = 1/3$. And there is a 10th row with entries $P_{10,10} = 1.0$ and $P_{10,j} = 0$ for all other j . We also have $P_{i,10} = 0$ for all $i < 9$.

New questions become interesting for this absorbing Markov chain. Now we want to know about the time it takes until the mouse leaves the maze. Now we want to know about the expected number of visits to the various rooms before the mouse leaves the maze.

We now show how to compute the expected number of steps from each starting state until the mouse leaves the maze. We obtain a system of 9 equations in 9 unknowns by conditioning on what happens in the first step.

Let T_i be the time (number of steps) until the mouse leaves the maze (first enters state 10) starting in state i . We cannot find $E[T_i]$ for one i directly, but we can solve for all of the $E[T_i]$ in a system of equations. For example, the first equation is

$$E[T_1] = 1 + (1/2)E[T_2] + (1/2)E[T_4] ,$$

while the second equation is

$$E[T_2] = 1 + (1/3)E[T_1] + (1/3)E[T_3] + (1/3)E[T_5] .$$

We can use MATLAB again to solve this system of equations.


2.2 The Second Absorbing Markov Chain

We now consider a larger example with more possible absorbing states. Now we put a door leading out of the maze from rooms 3, 7 and 9. The probabilities of leaving through each of these doors is $1/3$. As before, we let the mouse choose each of the available doors with equal probability. Wherever the mouse leaves the original 9 rooms, we assume that the mouse cannot return. He stays outside. Here is the new picture:

We now pay attention to the door from which the mouse leaves the maze. We thus obtain the 12-state absorbing Markov chain with 3 absorbing states. We say that the Markov chain enters state 10 if the mouse leaves through the door out of room 3; we say that the Markov chain enters state 11 if the mouse leaves through the door out of room 7; we say that the Markov chain enters state 12 if the mouse leaves through the door out of room 9. We let the mouse stay outside the maze when it leaves, so the Markov chain is absorbing: we have $P_{10,10} = P_{11,11} = P_{12,12} = 1$. Here is the new transition matrix:

$$P = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \\ 9 \\ 10 \\ 11 \\ 12 \end{matrix} & \left(\begin{array}{cccccccccccc} 0 & 1/2 & 0 & 1/2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1/3 & 0 & 1/3 & 0 & 1/3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1/3 & 0 & 0 & 0 & 1/3 & 0 & 0 & 0 & 1/3 & 0 & 0 \\ 1/3 & 0 & 0 & 0 & 1/3 & 0 & 1/3 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1/4 & 0 & 1/4 & 0 & 1/4 & 0 & 1/4 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1/3 & 0 & 1/3 & 0 & 0 & 0 & 1/3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1/3 & 0 & 0 & 0 & 1/3 & 0 & 0 & 1/3 & 0 \\ 0 & 0 & 0 & 0 & 1/3 & 0 & 1/3 & 0 & 1/3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1/3 & 0 & 1/3 & 0 & 0 & 0 & 1/3 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{array} \right) . \end{matrix}$$

The Maze with Outside Rooms



This is an absorbing Markov chain, which is reducible. We analyze it in a different way than we analyze an irreducible Markov chain. Use the program `absorbing.m` on this example.

2.3. 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 .

2.3.1 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

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

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

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

To summarize,

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

In matrix form, we have

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

where the identity matrix I here has the same dimension m as Q . (Since Q is the submatrix corresponding to the transient states, $Q^n \rightarrow 0$ as $n \rightarrow \infty$, 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 $N = \text{inv}(I - Q)$. 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

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

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 $n \rightarrow \infty$. Hence, $(I - Q)N_n \rightarrow I$ as $n \rightarrow \infty$. 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.

2.3.2 Other Quantities of Interest: M and B

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.

2.3.3. 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, is an irreducible Markov chain. It has a unique stationary distribution, obtained by solving $\pi = \pi P$. We have considered a DTMC with a finite state space. There are issues about what happens with an infinite state space. With an infinite state space, such as the nonnegative integers, it is possible that the DTMC is transient or null-recurrent instead of positive recurrent. With a finite state space, the irreducible DTMC is always positive recurrent, in which case the stationary probability vector π is proper (sums to 1). In the other two cases, it is the 0 vector.
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 $1/2$.