

Matrix Applications: Markov Chains and Game Theory

Christopher Carl Heckman

Department of Mathematics and Statistics, Arizona State University

checkman@math.asu.edu

Two important applications of matrices which are discussed in MAT 119 are Markov Chains and Game Theory. Here, we present a brief summary of what the textbook covers, as well as how to solve certain problems in these applications.

Markov Chains

1. The Setup.

Markov Chains model a situation, where there are a certain number of *states* (which will unimaginatively be called $1, 2, \dots, n$), and whether the state changes from state i to state j is a constant probability. In particular, it does not matter what happened, for the state to be in state i in the first place.

In general, the state that the situation is in will only be known probabilistically; thus, there is a probability of p_1 that it is in state 1, a probability of p_2 that it is in state 2, etc. A row vector v with n entries represents the probabilistic knowledge of the state if each entry is nonnegative, and the sum of the entries in v is 1; in that case, v is called a *probability vector*.

The change in state from one point in time to another point in time is determined by an $n \times n$ *transition matrix* P , which has the properties that:

- (a) $P_{i,j} \geq 0$ for all i, j .¹
- (b) $P_{i,1} + P_{i,2} + \dots + P_{i,n} = 1$, for all i .

The entry $P_{i,j}$ represents the probability that you will change from state i to state j .

The probability vector v_0 is used for the “initial set-up.”

Moving Around A Square. For instance, if you imagine 1, 2, 3, and 4 being corners of a square (where 1 and 3 are non-adjacent, and 2 and 4 are non-adjacent), you start at corner #1, and in each step you decide randomly and fairly which of the following to do:

- (a) Stay where you are;
- (b) Move from n to $n + 1$ (Where you go to 1, if you’re at 4); and
- (c) Move from n to $n - 1$ (where you go to 4, if you’re at 1);

then $v_0 = [1, 0, 0, 0]$ and the transition matrix is $P = \begin{pmatrix} 1/3 & 1/3 & 0 & 1/3 \\ 1/3 & 1/3 & 1/3 & 0 \\ 0 & 1/3 & 1/3 & 1/3 \\ 1/3 & 0 & 1/3 & 1/3 \end{pmatrix}$.

2. Terminology.

Now for some vocabulary:

A Markov Chain is called *regular* if there is some positive integer $k > 0$ such that $(P^k)_{i,j} > 0$ for all i, j .² This means you can potentially get from any state to any other state in k steps. The example above (“Moving Around A Square”) is regular, since every entry of P^2 is positive.

A probability vector t is a *fixed probability vector* if $t = tP$. $[0.25, 0.25, 0.25, 0.25]$ is a fixed probability vector of “Moving Around A Square” (and is in fact the only one); $[1, 0, 0, 0, 0]$ and $[0, 0, 0, 0, 1]$ are fixed probability vectors of “Drunken Walk” (covered in the “Behavior of Absorbing Markov Chains” section).

¹ $A_{i,j}$ is my notation for the entry in the i th row and the j th column of A . Some books use $a_{i,j}$ for this number. Also, “for all i ” should be interpreted to mean “for all i for which this makes sense.”

² P^n is defined in the sense of matrix multiplication.

3. Finding probability vectors after k iterations.

The probability of being in state j after one iteration can be broken in to the sum of the probability of being in state i , times the probability that you go from state i to state j , summed over all states i . Thus, if v_0 is the initial state, the probability vector after one iteration is v_0P . Iterating, we find that the probability vector after k state is

$$v^{(k)} = v_0P^k.$$

For instance, in “Moving Around A Square”, the probability vector after one iteration is

$$v^{(1)} = v_0P = [1, 0, 0, 0] \cdot \begin{pmatrix} 1/3 & 1/3 & 0 & 1/3 \\ 1/3 & 1/3 & 1/3 & 0 \\ 0 & 1/3 & 1/3 & 1/3 \\ 1/3 & 0 & 1/3 & 1/3 \end{pmatrix} = \left[\frac{1}{3}, \frac{1}{3}, 0, \frac{1}{3} \right],$$

the probability vector after two iterations is

$$v^{(2)} = v_0P^2 = [1, 0, 0, 0] \cdot \begin{pmatrix} 1/3 & 1/3 & 0 & 1/3 \\ 1/3 & 1/3 & 1/3 & 0 \\ 0 & 1/3 & 1/3 & 1/3 \\ 1/3 & 0 & 1/3 & 1/3 \end{pmatrix}^2 = \left[\frac{1}{3}, \frac{2}{9}, \frac{2}{9}, \frac{2}{9} \right],$$

and so on.

4. Finding all fixed probability vectors.

A fixed probability vector satisfies the condition $t = tP$, so $t(P - I) = 0$. Because we are used to solving systems of linear equations where the vector is a column vector, we will take the transpose.³ Since the transpose reverses the order of the factors, we end up solving

$$(P - I)^T \cdot t^T = 0. \tag{1}$$

But a probability vector also has to have entries which add up to 1; we can incorporate this into the system of linear equations by writing this condition as

$$[1, 1, \dots, 1] \cdot t^T = 0. \tag{2}$$

Since P is a transition matrix, the rows of $(P - I)^T$ add up to $[0, 0, \dots, 0]$. This means that the n th equation in (1) is redundant, and we can incorporate (2) into the system by removing the last row of $(P - I)^T$ and replacing it by $[1, 1, \dots, 1]$. If we call this new matrix R , then we need to find solutions to

$$R \cdot \begin{bmatrix} t_1 \\ t_2 \\ \vdots \\ t_n \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}. \tag{3}$$

(Note that we have not insisted that the entries of t be nonnegative. After we find all solutions, we will put restrictions on the free variables (if there are any) to make these entries nonnegative.)

For instance, for “Moving Around A Square”, our system of equations in (3) becomes

$$\begin{pmatrix} -2/3 & 1/3 & 0 & 1/3 \\ 1/3 & -2/3 & 1/3 & 0 \\ 0 & 1/3 & -2/3 & 1/3 \\ 1 & 1 & 1 & 1 \end{pmatrix} \cdot \begin{bmatrix} t_1 \\ t_2 \\ t_3 \\ t_4 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}.$$

³ If A is an $m \times n$ matrix, then the *transpose* of A (denoted A^T) is the $n \times m$ matrix B such that $B_{i,j} = A_{j,i}$ for all i, j .

This system has a unique solution, namely $t = [0.25, 0.25, 0.25, 0.25]$.⁴ For an example of a Markov Chain with more than one fixed probability vector, see the “Drunken Walk” example below.

5. Behavior of absorbing Markov Chains.

A state i is an *absorbing state* if $P_{i,i} = 1$; it is one from which you cannot change to another state. (“Moving Around A Square” has no absorbing states.) A Markov Chain with at least one absorbing state, and for which all states potentially lead to an absorbing state, is called an *absorbing Markov Chain*.

Drunken Walk.⁵ There is a street in a town with a De-tox center, three bars in a row, and a Jail, all in a row. A man is in the bar in the middle. Every hour, whenever he is in a bar, he takes a drink, goes outside, chooses a direction at random with equal probability, and then walks one building in that direction. If he ends up at the De-tox center or the Jail, he is taken in for the rest of the night. If we use 1, 2, ..., 5 to denote the five buildings (where the De-tox center is 1), then we end up with a Markov Chain with the following transition matrix:

$$P = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0.5 & 0 & 0.5 & 0 & 0 \\ 0 & 0.5 & 0 & 0.5 & 0 \\ 0 & 0 & 0.5 & 0 & 0.5 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

and initial probability vector $v_0 = [0, 0, 1, 0, 0]$.

“Drunken Walk” is an absorbing Markov Chain, since 1 and 5 are absorbing states.⁶ If we look for all fixed probability vectors, we find that we need to solve the system

$$\begin{pmatrix} 0 & 0.5 & 0 & 0 & 0 \\ 0 & -1 & 0.5 & 0 & 0 \\ 0 & 0.5 & -1 & 0.5 & 0 \\ 0 & 0 & 0.5 & -1 & 0 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix} \cdot \begin{bmatrix} t_1 \\ t_2 \\ t_3 \\ t_4 \\ t_5 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}.$$

There are an infinite number of solutions:

$$t = [t_1, t_2, t_3, t_4, t_5] = [a, 0, 0, 0, 1 - a].$$

Since every entry must be nonnegative, we must have $a \geq 0$ and $1 - a \geq 0$, which means $0 \leq a \leq 1$.

With an absorbing Markov Chain, we can also ask the following questions: How many iterations should we expect before entering an absorbing state? How many times do we expect to be in a non-absorbing state i ? What is the probability of reaching absorbing state j , assuming we start at state i ?⁷ We answer these questions below.

To do this, we take the transition matrix P and shuffle the rows and columns so that all of the absorbing states are together, and all the absorbing states are together. If $P_{i,i} = 1$, we can swap row i of P with row 1, and then swap column i of P with column 1, to get the 1 in the first row and first column. After that, we can take another absorbing state (say, $P_{j,j} = 1$, where $j \neq 1$) and place the 1 in the second row and second column, etc. By doing so, we can obtain a sort of “canonical form” for P , which is as follows:

$$P = \begin{bmatrix} I_r & 0 \\ S & Q \end{bmatrix}, \tag{4}$$

⁴ If you have a regular Markov Chain, then it can be proven that there is exactly one fixed probability vector, and the long-term behavior of that Markov Chain is that fixed probability vector.

⁵ “Drunken Walk” is based on the “Gambler’s Ruin” problem.

⁶ This Markov Chain is not regular. No absorbing Markov Chain is regular.

⁷ The last question, in the context of an interrupted dice game, was what led Blaise Pascal to start investigating the topic of probability.

where r is the number of absorbing states; we will also let s be the number of non-absorbing states.⁸ For the example of “Drunken Walk”, the matrix P becomes

$$P = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.5 & 0.5 \\ 0 & 0.5 & 0.5 & 0 & 0 \\ 0.5 & 0 & 0.5 & 0 & 0 \end{pmatrix},$$

so that $r = 2$, $s = 3$, $S = \begin{pmatrix} 0 & 0 \\ 0 & 0.5 \\ 0.5 & 0 \end{pmatrix}$, and $Q = \begin{pmatrix} 0 & 0.5 & 0.5 \\ 0.5 & 0 & 0 \\ 0.5 & 0 & 0 \end{pmatrix}$. Note that the rows of S , the rows of Q , and the columns of Q correspond to the states 3, 4, and 2, in that order, and the columns of S correspond to the absorbing states 1, 5 (in that order), so that S is $s \times r$ and Q is $s \times s$.

Another matrix associated with an absorbing Markov Chain is the so-called *fundamental matrix* T , given by:

$$T = I_s + Q + Q^2 + Q^3 + \cdots = (I_s - Q)^{-1}. \quad (5)$$

The probability of getting from state i to state j in k steps is $(Q^k)_{i,j}$. Thus $T_{i,j}$ is the probability of getting from state i to state j in 0 steps, plus the probability in 1 step, plus the probability in 2 steps, etc. This sum is the expected number of times we go from state i to state j .⁹ If we add the entries in the i th row of T , we obtain the expected number of times that we are in a non-absorbing state.

For “Drunken Walk”,

$$T = \left[\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} - \begin{pmatrix} 0 & 0.5 & 0.5 \\ 0.5 & 0 & 0 \\ 0.5 & 0 & 0 \end{pmatrix} \right]^{-1} = \begin{pmatrix} 2 & 1 & 1 \\ 1 & 1.5 & 0.5 \\ 1 & 0.5 & 1.5 \end{pmatrix}.$$

In terms of the problem, where the man starts in bar #3 (row 1), this states that the man will expect to be in bar #3 two times, in bar #2 once, and in bar #4 once. The expected number of drinks he will get is $2 + 1 + 1 = 4$. (If he started in bar #2 or bar #4, the expected number of drinks is $1 + 1.5 + 0.5 = 3$.)

To find the probability that the man ends up in states 1 or 5 (De-tox or Jail), we calculate the matrix $T \cdot S$; then $(T \cdot S)_{i,j}$ is the probability that the Markov Chain ends up in absorbing state j , assuming it starts in (the non-absorbing¹⁰) state i . For “Drunken Walk”,

$$T \cdot S = \begin{pmatrix} 2 & 1 & 1 \\ 1 & 1.5 & 0.5 \\ 1 & 0.5 & 1.5 \end{pmatrix} \cdot \begin{pmatrix} 0 & 0 \\ 0 & 0.5 \\ 0.5 & 0 \end{pmatrix} = \begin{pmatrix} 0.5 & 0.5 \\ 0.25 & 0.75 \\ 0.75 & 0.25 \end{pmatrix}.$$

So the probability of the man ending up in De-tox is 0.5 and in Jail is 0.5. (We could have guessed this from the symmetry of the problem.) If he starts in bar #2, the probability he ends up in De-tox is 0.75, and the probability he ends up in Jail is 0.25.¹¹

⁸ If every state is an absorbing state, then the whole matrix $P = I_n$; but then the questions asked above are not interesting.

⁹ This sum converges because all the eigenvalues of Q turn out to be less than 1 in absolute value, the matrix equivalent of the geometric ratio test.

¹⁰ If the initial state is absorbing, the probability of ending up in j is easy to calculate.

¹¹ Remember, the row for bar #2 is the bottom row of T and $T \cdot S$.

Game Theory

1. *The Setup.*

A *game* refers to a situation where there are two or more players, who each have a set of options. If these players choose particular options, the rules of the game determine which of the players get rewarded. For MAT 119, all games are assumed to have two players, and the reward is determined by a matrix A called the *payoff matrix*, and each player will have only a finite number of options, which will be numbered. The number A_{ij} determines how many “points” Player I receives from Player II if Player I chooses option i and Player II chooses option j .¹² In addition, it will be assumed that both players have received the payoff matrix in advance.¹³

An example of a game is “Rock–Scissors–Paper”, where we will let 1 denote Rock, 2 denote Scissors, and 3 denote Paper. In this game, Rock breaks (beats) Scissors, Scissors cut (beat) Paper, and Paper covers (beats) Rock. The payoff matrix is

$$A = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}.$$

2. *Finding optimal strategies: Saddle Points.*

Some payoff matrices have one or more *saddle points*. A saddle point is a pair of strategies (i, j) such that $A_{i,m} \leq A_{i,j}$ for all m , and $A_{m,j} \geq A_{i,j}$ for all m . For instance, if the payoff matrix is $A = \begin{pmatrix} 2 & 1 \\ 3 & 4 \end{pmatrix}$, then $(2, 1)$ is a saddle point.

Saddle points represent optimum solutions. Remember that Player I is trying to maximize his payoff, and Player II is trying to minimize it. By choosing option 2 in the example in the previous paragraph, Player I is guaranteed of receiving 3 points, possibly more. By choosing option 1, Player II is guaranteed to have to give up at most 3 points. These two strategies are *optimal*, in the sense that Player I cannot find a strategy where he/she receives more than 3 points, and Player II cannot find a strategy where he/she gives up fewer than 3 points, even if randomized strategies are allowed (see below). The *value* of the game is the number of points Player I receives from Player II when both use optimum strategies.

If a payoff matrix has two or more saddle points, the values are the same.¹⁴

3. *Finding optimal strategies: “Dominated” Rows and Columns.*

Sometimes, it can be determined that a certain player will never use a certain option. For instance, if the payoff matrix is

$$A = \begin{pmatrix} 2 & 5 & 8 \\ 3 & 2 & 7 \end{pmatrix}, \tag{6}$$

Player II will never use option 3, because by playing option 1 instead, Player II will decrease the expected payoff, no matter what Player I does.

In general, Player I will never play option i if there is another option i' such that $A_{i,j} \leq A_{i',j}$ for all j ; this allows you to remove a row from the payoff matrix. Player II will never play option j if there is an option j' such that $A_{i,j} \geq A_{i,j'}$ for all i ; this allows you to remove a column from the payoff matrix.

Sometimes, if you remove a row or column, you will be able to remove another, and/or find a saddle point.

¹² Since one player wins at the expense of another, these type of games are sometimes called *zero sum* games. Games which are not zero-sum are harder to analyze; one particularly notorious example is called the Prisoner’s Dilemma.

¹³ The technical term for this situation is *perfect information*. Chess is an example of another game with perfect information; Poker is a game without perfect information, since you can’t see your opponent’s cards.

¹⁴ Left as an exercise for the reader.

4. Finding optimal strategies: Choosing between two options.

Not every game has a saddle point; Rock–Scissors–Paper does not, for instance. In this case, both players use what are called *randomized strategies*; each chooses option 1 with probability p_1 , option 2 with probability p_2 , etc., up to p_k . Note that $[p_1, p_2, \dots, p_k]$ is a probability vector,¹⁵ and if Player I uses the randomized strategy P_I , and Player II uses the randomized strategy P_{II} , then the expected payoff is $P_I A P_{II}$.

For the purposes of the class, you may assume that each player has exactly two options.¹⁶ Thus $P_I = [p, 1 - p]$ and $P_{II} = [q, 1 - q]$, for some real numbers p and q between 0 and 1. If we use the payoff matrix shown in (6) above (keeping in mind that Player II will never choose option 3), then in the “average case”, Player II will see two choices, with payoffs of

$$P_I A = [p, 1 - p] \cdot \begin{pmatrix} 2 & 5 \\ 3 & 2 \end{pmatrix} = [2p + 3(1 - p), 5p + 2(1 - p)] = [3 - p, 2 + 3p].$$

Player II will now choose whichever option gives the smaller number; so the payoff will be the minimum of $3 - p$ and $2 + 3p$.

Player I is always trying to maximize the number of points he/she receives, so he/she will choose the value of p that maximizes $\min\{3 - p, 2 + 3p\}$. It turns out that this value of p is the one that makes $3 - p = 2 + 3p$,¹⁷ namely $p = 0.25$. If Player I uses the strategy $[0.25, 0.75]$, he/she will expect to win $3 - p = 2 + 3p = 2.75$ points every time the game is played.

We can also determine Player II's strategy. If Player II uses the strategy $[q, 1 - q]$, then Player I is faced with

$$A P_{II}^\top = \begin{pmatrix} 2 & 5 \\ 3 & 2 \end{pmatrix} \begin{bmatrix} q \\ 1 - q \end{bmatrix} = \begin{bmatrix} 2q + 5(1 - q) \\ 3q + 2(1 - q) \end{bmatrix} = \begin{bmatrix} 5 - 3q \\ 2 + q \end{bmatrix}.$$

Player II's strategy is optimized when $5 - 3q = 2 + q$, or when $q = 0.75$. If Player II uses the strategy $[0.75, 0.25]$, he/she will expect to lose $5 - 3q = 2 + q = 2.75$ points every time the game is played.

The two strategies are thus optimal, and the value of the game (how much Player I is expected to win every time the game is played) is 2.75.

Note that, for Rock–Scissors–Paper, the strategy $[1/3, 1/3, 1/3]$ is optimal for both players, and the value of the game is 0.¹⁸

¹⁵ Like we discussed in the “Markov Chains” section.

¹⁶ Finding optimal solutions with more than two options can be done using Linear Programming.

¹⁷ The proof left for the reader.

¹⁸ This can, and should, be verified by the reader.