Absorbing Markov Chains

Allan Yashinski

July 21, 2021

15.1 Introduction

We consider Markov chains with absorbing states, i.e. states from which it is impossible to leave. As a motivating example, we consider a tied game of tennis.

15.2 A Tied Tennis Game

To win a game of tennis, you have to score at least 4 points, but you must also have at least 2 more points than your opponent. When the game is tied 3–3, it is called a “deuce”. From here, the game continues until one player has scored 2 more points than the other. We will model a tied game using a Markov chain (we will ignore the portion of the game before the tie is reached.) We will refer to the two players as A and B. There are five states of this Markov chain:

1 Deuce (tie)

2 Advantage A (A has 1 more point than B)

3 Advantage B (B has 1 more point than A)

4 A wins.

5 B wins.

Our main assumption is that at any given point, there is a probability of p that A will score (and consequently a probability of $1 - p$ that B will score). The Markov chain looks like:

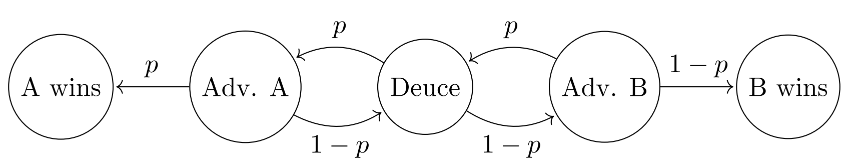

The new thing about this type of Markov chain is that the process we are modeling can end. We model this by declaring that once we enter the state where A (or B) has won, we remain in that state with probability 1. Let’s rewrite it, taking this into account and replacing the state names with the numbers we assigned to them:

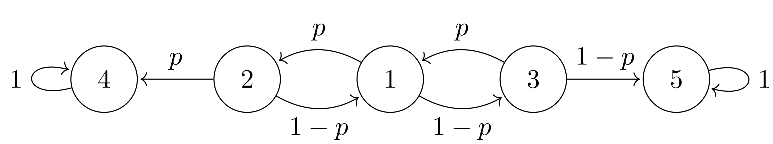

The transition matrix is

$$T = \begin{bmatrix} 0 & 1-p & p & 0 & 0 \\ p & 0 & 0 & 0 & 0 \\ 1-p & 0 & 0 & 0 & 0 \\ 0 & p & 0 & 1 & 0 \\ 0 & 0 & 1-p & 0 & 1 \end{bmatrix}.$$

States 4 and 5 are called *absorbing states.* These are the states from which it is impossible to leave. We refer to all other states $1, 2, 3$ as *transient states.* Note that it is impossible to go from an absorbing state to a transient state, so T is not regular.

Let's suppose that $p = 0.6$, so that

$$T = \begin{bmatrix} 0 & 0.4 & 0.6 & 0 & 0 \\ 0.6 & 0 & 0 & 0 & 0 \\ 0.4 & 0 & 0 & 0 & 0 \\ 0 & 0.6 & 0 & 1 & 0 \\ 0 & 0 & 0.4 & 0 & 1 \end{bmatrix}.$$

It is natural to ask what the probability of player A or B winning, given that the game starts in deuce. Observe that we will not get any answers by simply searching for the steady-state vector (as in the regular case). Here, there are two linearly independent eigenvectors:

$$\bar{v}_1 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \qquad \bar{v}_2 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}.$$

Those alone do not seem to carry any useful information. If the game starts in deuce, then we consider the Markov chain

$$\bar{x}_{k+1} = T\bar{x}_k, \qquad \bar{x}_0 = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

The vector $\bar{x}_k = T^k \bar{x}_0$ tells us the probability that we are in any of the five states after k steps. Let's investigate with some computations:

$$\bar{x}_1 = \begin{bmatrix} 0 \\ 0.6 \\ 0.4 \\ 0 \\ 0 \end{bmatrix}, \quad \bar{x}_2 = \begin{bmatrix} 0.48 \\ 0 \\ 0 \\ 0.36 \\ 0.16 \end{bmatrix}, \quad \bar{x}_3 = \begin{bmatrix} 0 \\ 0.288 \\ 0.192 \\ 0.36 \\ 0.16 \end{bmatrix}, \quad \bar{x}_4 = \begin{bmatrix} 0.2304 \\ 0 \\ 0 \\ 0.5328 \\ 0.2368 \end{bmatrix}.$$

For example, $\bar{x}_4$ tells us that after four iterations (scores), there is a 23.04% chance that the game is back in deuce, so it is still going on. Also, there is a 53.28% chance that A has already won and a 23.68% chance that B has already won. We continue with more computations:

$$\bar{x}_6 = \begin{bmatrix} 0.1106 \\ 0 \\ 0 \\ 0.6157 \\ bar0.2737 \end{bmatrix}, \quad \bar{x}_{10} = \begin{bmatrix} 0.0255 \\ 0 \\ 0 \\ 0.6747 \\ 0.2999 \end{bmatrix}, \quad \bar{x}_{30} = \begin{bmatrix} 0.0000 \\ 0 \\ 0 \\ 0.6923 \\ 0.3077 \end{bmatrix}.$$

We see that after 30 steps, there is a nearly 0% chance that the game is still going on in deuce, but there is a 69.23% chance that A has won and a 30.77% chance that B has won. These numbers indicate are what we are really looking for!

Formally, the probabilities we seek should be found in the limit

$$\lim_{k\to\infty} \bar{x}_k = \lim_{k\to\infty} T^k \bar{x}_0.$$

In contrast with the regular case, **this limit will depend on the initial condition** $\bar{x}_0$. In what follows, we will think a little more carefully about how to compute this limit.

15.3 Absorbing Markov Chains

In a Markov chain, an *absorbing state* is one in which you get stuck forever (like A wins/B wins above). By an *absorbing Markov chain*, we mean a Markov chain which has absorbing states and it is possible to go from any transient state to some absorbing state in a finite number of steps.

Suppose we have an absorbing Markov chain with r transient states $t_1, \ldots, t_r$ and s absorbing states $a_1, \ldots, a_s$. The whole Markov chain has $r+s$ states, and we order them as follows:

$$t_1, \ldots, t_r, a_1, \ldots, a_s.$$

The transition matrix will have the following "block" form

$$T = \begin{bmatrix} Q & \bar{0}_{r\times s} \\ R & I_s \end{bmatrix},$$

where

- Q is an $r \times r$ matrix which holds the probabilities of moving from a transient state to another transient state.
- R is an $s\times r$ matrix which holds the probabilities of moving from a transient state to an absorbing state.

- $\bar{0}_{r\times s}$ is the $r \times s$ matrix of all 0's. These 0's represent the probabilities of moving from an absorbing state to a transient state (which is impossible).
- I_s is the $s \times s$ identity matrix, which holds the probabilities of transitioning between absorbing states (which is impossible as we just get stuck in the same absorbing state.)

In the tennis match example, the transition matrix is

$$T = \begin{bmatrix} 0 & 0.4 & 0.6 & 0 & 0 \\ 0.6 & 0 & 0 & 0 & 0 \\ 0.4 & 0 & 0 & 0 & 0 \\ 0 & 0.6 & 0 & 1 & 0 \\ 0 & 0 & 0.4 & 0 & 1 \end{bmatrix},$$

so we have

$$Q = \begin{bmatrix} 0 & 0.4 & 0.6 \\ 0.6 & 0 & 0 \\ 0.4 & 0 & 0 \end{bmatrix}, \qquad R = \begin{bmatrix} 0 & 0.6 & 0 \\ 0 & 0 & 0.4 \end{bmatrix}.$$

We are interested in $\lim_{k\to\infty} T^k \bar{x}_0$, so it is natural to investigate powers of T. The block form of T allows us to compute:

$$\begin{aligned} T^2 &= \begin{bmatrix} Q & \bar{0}_{r\times s} \\ R & I_s \end{bmatrix} \begin{bmatrix} Q & \bar{0}_{r\times s} \\ R & I_s \end{bmatrix} = \begin{bmatrix} Q^2 & \bar{0}_{r\times s} \\ RQ+R & I_s \end{bmatrix} \\ T^3 &= \begin{bmatrix} Q & \bar{0}_{r\times s} \\ R & I_s \end{bmatrix} \begin{bmatrix} Q^2 & \bar{0}_{r\times s} \\ RQ+R & I_s \end{bmatrix} = \begin{bmatrix} Q^3 & \bar{0}_{r\times s} \\ RQ^2+RQ+R & I_s \end{bmatrix} \\ T^4 &= \begin{bmatrix} Q & \bar{0}_{r\times s} \\ R & I_s \end{bmatrix} \begin{bmatrix} Q^3 & \bar{0}_{r\times s} \\ RQ^2+RQ+R & I_s \end{bmatrix} = \begin{bmatrix} Q^4 & \bar{0}_{r\times s} \\ RQ^3+RQ^2+RQ+R & I_s \end{bmatrix}. \end{aligned}$$

In general, the pattern is

$$T^k = \begin{bmatrix} Q^k & \bar{0}_{r\times s} \\ R+RQ+\ldots+RQ^{k-1} & I_s \end{bmatrix} = \begin{bmatrix} Q^k & \bar{0}_{r\times s} \\ R\sum_{i=0}^{k-1} Q^i & I_s \end{bmatrix}.$$

When we take a limit as $k \to \infty$, we will encounter the series

$$\sum_{i=0}^{\infty} Q^i = I + Q + Q^2 + Q^3 + \ldots$$

We've seen in the past that this series converges to $(I-Q)^{-1}$ under the assumption that the column sums of Q are all less than 1. For our tennis example,

$$Q = \begin{bmatrix} 0 & 0.4 & 0.6 \\ 0.6 & 0 & 0 \\ 0.4 & 0 & 0 \end{bmatrix}$$

the column sums are $1, 0.4, 0.6$; unfortunately they are not all less than 1. However,

$$Q^2 = \begin{bmatrix} 0.48 & 0 & 0 \\ 0 & 0.24 & 0.36 \\ 0 & 0.16 & 0.24 \end{bmatrix}$$

does have column sums less than 1 (and the same is true for $Q^3, Q^4, \ldots$) This turns out to be enough to guarantee convergence. For a general absorbing Markov chain, some power Q^k of Q must have column sums less than 1 because the column sums of T^k are exactly 1. It then follows by considering our formula above for T^k, in which Q^k is the upper left block. The definition of absorbing Markov chain will imply that the lower left block will have nonzero entries in each column for large enough k.

Theorem 15.3.0.1. For an absorbing Markov chain, some power Q^k will have column sums less than 1. Consequently,

$$\sum_{i=0}^{\infty} Q^i = I + Q + Q^2 + Q^3 + \ldots = (I - Q)^{-1}.$$

Since this series converges, its terms must go to zero. So we also obtain

$$\lim_{k\to\infty} Q^k = \bar{0}_{r\times r}.$$

Putting it together gives the following important result.

Theorem 15.3.0.2. For an absorbing Markov chain with $T = \begin{bmatrix} Q & \bar{0}_{r\times s} \\ R & I_s \end{bmatrix}$, we have

$$\lim_{k\to\infty} T^k = \begin{bmatrix} \bar{0}_{r\times r} & \bar{0}_{r\times s} \\ R(I_r - Q)^{-1} & I_s \end{bmatrix}.$$

In our tennis example, we have

$$T = \begin{bmatrix} 0 & 0.4 & 0.6 & 0 & 0 \\ 0.6 & 0 & 0 & 0 & 0 \\ 0.4 & 0 & 0 & 0 & 0 \\ 0 & 0.6 & 0 & 1 & 0 \\ 0 & 0 & 0.4 & 0 & 1 \end{bmatrix}, \qquad Q = \begin{bmatrix} 0 & 0.4 & 0.6 \\ 0.6 & 0 & 0 \\ 0.4 & 0 & 0 \end{bmatrix}, \qquad R = \begin{bmatrix} 0 & 0.6 & 0 \\ 0 & 0 & 0.4 \end{bmatrix}.$$

We can compute directly (using MATLAB) that

$$(I_3 - Q)^{-1} = \begin{bmatrix} 1.9231 & 0.7692 & 1.1538 \\ 1.1538 & 1.4615 & 0.6923 \\ 0.7692 & 0.3077 & 1.4615 \end{bmatrix}$$

and

$$R(I_3 - Q)^{-1} = \begin{bmatrix} 0.6923 & 0.8769 & 0.4154 \\ 0.3077 & 0.1231 & 0.5846 \end{bmatrix}.$$

Consequently,

$$\lim_{k\to\infty} T^k = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0.6923 & 0.8769 & 0.4154 & 1 & 0 \\ 0.3077 & 0.1231 & 0.5846 & 0 & 1 \end{bmatrix}.$$

Now we are in a position to properly answer questions about the tennis match. Suppose the game starts in deuce, so that $\bar{x}_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \end{bmatrix}^T$. Then to find the probabilities of what state we are in "after a long time", we need to compute

$$\lim_{k\to\infty} T^k \bar{x}_0 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0.6923 & 0.8769 & 0.4154 & 1 & 0 \\ 0.3077 & 0.1231 & 0.5846 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0.6923 \\ 0.3077 \end{bmatrix}.$$

This says that the probability that A wins is 0.6923 and the probability that B wins is 0.3077. There is a probability of 0 that the game is still going on in Deuce or Advantage A/B. Note that multiplying by this $\bar{x}_0$ just extracted the first column. Further, the probabilities we were interested in were just the entries of the first column of $R(I_3 - Q)^{-1}$.

Theorem 15.3.0.3. For an absorbing Markov chain, the (i, j)-entry of $R(I_r - Q)^{-1}$ is the probability of ending in absorbing state a_i given that we started in transient state t_j.

For example, the second column says that if the game is in Advantage A, then A has a 0.8769 probability of winning, whereas B has a 0.1231 probability of coming back and winning.

The matrix $(I_r - Q)^{-1}$ contains more useful information.

Theorem 15.3.0.4. For an absorbing Markov chain,

1. The (i, j)-entry of $(I_r - Q)^{-1}$ contains the expected number of visits to transient state t_i given that we started in transient state t_j.
2. The j-th column sum of $(I_r - Q)^{-1}$ contains the expected number of time steps before absorption, given that we started in transient state t_j.

The second statement follows from the first. To understand the first, we need to view $(I_r - Q)^{-1}$ as the infinite series

$$(I_r - Q)^{-1} = I_r + Q + Q^2 + Q^3 + \dots$$

The (i,j)-entry of $(I_r - Q)^{-1}$ is the sum of the (i,j)-entries of each of the powers Q^k. Recall that the (i,j)-entry of Q^k contains the probability of starting in t_j and ending in t_i after exactly k steps. The infinite sum of all these probabilities is the expected value described in the theorem.

For our tennis example,

$$(I_3 - Q)^{-1} = \begin{bmatrix} 1.9231 & 0.7692 & 1.1538 \\ 1.1538 & 1.4615 & 0.6923 \\ 0.7692 & 0.3077 & 1.4615 \end{bmatrix}.$$

The first entry says that if the game starts in deuce, then the expected number of times it will be in deuce is 1.9231 (Note this counts the starting deuce. Consequently, all diagonal entries should be at least 1.) The expected number is like an average. If we played a bazillion games of tennis starting in deuce, we would find that, on average, we had 1.9231 deuces per game.

The first column sum is 3.8462. The interpretation is that (assuming the game starts in deuce), it will take, on average, 3.8462 scores until the game ends. As an example, if A and B play 30 games of tennis, and we assume there is a score every minute, then we should expect that they played for a total of

$$(30 \text{ games})(3.8462 \text{ scores/game})(1 \text{ minute/score}) = 115.386 \text{ minutes}.$$

(Again, all games start in deuce. They are not playing the beginning portion of a traditional tennis game.)

If we return to the original

$$T = \begin{bmatrix} 0 & 1-p & p & 0 & 0 \\ p & 0 & 0 & 0 & 0 \\ 1-p & 0 & 0 & 0 & 0 \\ 0 & p & 0 & 1 & 0 \\ 0 & 0 & 1-p & 0 & 1 \end{bmatrix}$$

before we assigned a value to p, we have

$$Q = \begin{bmatrix} 0 & 1-p & p \\ p & 0 & 0 \\ 1-p & 0 & 0 \end{bmatrix}, \qquad R = \begin{bmatrix} 0 & p & 0 \\ 0 & 0 & 1-p \end{bmatrix}.$$

Using MATLAB, we can invert and get

$$(I_3 - Q)^{-1} = \frac{1}{2p^2 - 2p + 1} \begin{bmatrix} 1 & 1-p & p \\ p & p^2 - p + 1 & p^2 \\ 1-p & (1-p)^2 & p^2 - p + 1 \end{bmatrix}$$

and

$$R(I_3 - Q)^{-1} = \frac{1}{2p^2 - 2p + 1} \begin{bmatrix} p^2 & p(p^2 - p + 1) & p^3 \\ (1-p)^2 & (1-p)^3 & (1-p)(p^2 - p + 1) \end{bmatrix}.$$

Looking at these allows us to answer questions in terms of p, the probability that A scores. For example, the expected number of scores before the game ends is the first column sum of $(I_3 - Q)^{-1}$, which is

$$\frac{1}{2p^2 - 2p + 1}(1 + p + (1-p)) = \frac{2}{2p^2 - 2p + 1}.$$

By looking at the graph of this function, we see that its maximum is when $p = 0.5$, and the expected number of scores is 4. (Does it make sense why the maximum should be at $p = 0.5$?) We also see that if $p = 0$ or $p = 1$, the expected number of scores is exactly 2. (Again, does it make sense?)

Exercise 15.1. Find the smallest value for p such that A is more likely to win when the game is in Advantage B.