# ECE 586 Application: Markov Chains

Henry D. Pfister
ECE Department
Duke University

October 12th, 2020

## 1 A Few Simple Questions

### 1.1 What is the chance that a game of Candyland will last $m$ moves?

Candyland is an American boardgame where players draw cards to move and the goal is to reach the candy castle first. It can be played by very young children because it requires neither reading nor counting. Players draw cards randomly and, if a colored card is drawn, they move their piece to the next position of that color. If the card has a picture, they move to the position with that picture. There are also spaces that allow shortcuts or cause delays. A picture of board can be found at:
`https://kim.scarborough.chicago.il.us/images/cl-2010`

### 1.2 What is the chance that a game of Chutes and Ladders lasts $m$ moves?

Chutes and Ladders (aka Snakes and Ladders outside of the US) is a boardgame where a single die is rolled to determine how far you move on a gameboard defined by a grid. Some locations contain ladders that let you skip ahead while others contains chutes (aka snakes) that you move you backwards. Historically, it is based on an ancient game from India that teaches morality by associating ladders with virtues and snakes with vices. For more information, see:
`https://en.wikipedia.org/wiki/Snakes_and_Ladders`
`http://www.datagenetics.com/blog/november12011/`

### 1.3 What are the best properties to buy in Monopoly?

Monopoly is an American boardgame where players move around the gameboard buying, selling, and developing properties. Rent is collected from other players who land on your properties. Properties differ both in their expense (e.g., Park Place is valued much more highly than Baltic Avenue) and the chance that players will land on them. Markov chains can be used to estimate how often players will land on each property, which can be used to estimate their value. For way too much information, see:
`http://pfister.ee.duke.edu/courses/ece586/monopoly.pdf`

## 2 What is a Markov chain?

### 2.1 Introduction

A finite-state Markov chain (FSMC) with $n$ states is a sequence of random variables $X_1, X_2, X_3, \ldots$ where each $X_i \in [n] \triangleq \{1, 2, \ldots, n\}$ and

$$\Pr\left(X_{t+1} = j \,|\, X_t = i, X_1, X_2, \ldots, X_{t-1}\right) = \Pr\left(X_{t+1} = j \,|\, X_t = i\right).$$

If $\Pr(X_{t+1} = j \,|\, X_t = i)$ does not depend on $t$, then the Markov chain is called time invariant (or homogeneous). In the remainder of this note, we assume the FSMC is time invariant and we let $P \in \mathbb{R}^{n\times n}$ denote the *transition-probability matrix* with entries $[P]_{i,j} \triangleq P_{i,j} = \Pr(X_{t+1} = j \,|\, X_t = i)$. Since each row of $P$ represents a probability distribution, we see that $P_{i,j} \geq 0$ and $\sum_{j=1}^{n} P_{i,j} = 1$. The Markov property also implies that

$$\begin{aligned}\Pr\left(X_{t+2} = j \,|\, X_t = i\right) &= \sum_{k=1}^{n} \Pr\left(X_{t+2} = j \,|\, X_{t+1} = k\right) \Pr\left(X_{t+1} = k \,|\, X_t = i\right) \\ &= \sum_{k=1}^{n} P_{k,j} P_{i,k} = \sum_{k=1}^{n} P_{i,k} P_{k,j} = \left[P^2\right]_{i,j}.\end{aligned}$$

By induction, one can also show that $\Pr(X_{t+m} = j \,|\, X_t = i) = [P^m]_{i,j}$. Thus, given a fixed starting state, one can calculate the probability of being in state $i$ after $m$ steps by computing the $m$-th power of a matrix. Using the notation $\underline{\pi}^{(t)} = (\pi_1^{(t)}, \ldots, \pi_n^{(t)})$ with $\pi_i^{(t)} \triangleq \Pr(X_t = i)$, we also see that

$$\begin{aligned}\pi_j^{(t+1)} &= \sum_{i=1}^{n} \Pr(X_{t+1} = j, X_t = i) \\ &= \sum_{i=1}^{n} \Pr(X_{t+1} = j | X_t = i) \Pr(X_t = i) \\ &= \sum_{i=1}^{n} P_{i,j} \pi_i^{(t)} = \left[\underline{\pi}^{(1)} P^t\right]_j,\end{aligned}$$

where $\pi_i^{(1)}$ is the probability that the process starts in state $i$.

Let the random variable $U$ be uniformly distributed on the interval $[0,1)$ and let $u$ be a realization of $U$. For a discrete random variable $X \in [n]$, one can use $U$ to simulate $X$ by assigning subintervals of $[0,1)$ to each of the $n$ possibilities for $X$. Let $F_X(x) = \sum_{i=1}^{x} \Pr(X = i)$ for $x \in [n]$ be the cumulative distribution function of $X$. Then, we can set $X = x$ if $u \in [F_X(x-1), F_X(x))$ (with $F_X(0) = 0$ by convention). This works because

$$\Pr\left(U \in [F_X(x-1), F_X(x))\right) = F_X(x) - F_X(x-1) = \Pr(X = x).$$

For built-in sampling functions in Matlab and Python, see `randsrc` and `np.random.choice`, respectively.

Similarly, one can simulate a Markov chain by using pseudo-random numbers to generate realizations of the process. Let $u_1, u_2, \ldots$ be a realization of a sequence of independent and identically distributed (i.i.d.) copies of $U$. Then, a realization $x_1, x_2, \ldots$ of $X_1, X_2, \ldots$ is generated by choosing $x_t$ to be the unique value satisfying

$$\sum_{i=1}^{x_t - 1} \Pr(X_t = i | X_{t-1} = x_{t-1}) \leq u_t < \sum_{i=1}^{x_t} \Pr(X_t = i | X_{t-1} = x_{t-1}).$$

## 2.2 Absorbing States and Hitting Times

State $j$ is called *reachable* from state $i$ if $[P^m]_{i,j} > 0$ for some $m \in \mathbb{N}$. For a Markov chain starting from state $i$, the first *hitting time* of state $j$ is a random variable $T_{i,j}$ with distribution

$$\Pr\left(T_{i,j} = m\right) = \Pr\left(X_{m+1} = j, X_m \neq j, X_{m-1} \neq j, \ldots, X_2 \neq j \,|\, X_1 = i\right),$$

where, by convention, $\Pr\left(T_{i,j} = 0\right) = \delta_{i,j}$ and $\delta_{i,j}$ is Kronecker delta function. If state $j$ is not reachable from state $i$, then $T_{i,j} = \infty$ (i.e., $\Pr(T_{i,j} = \infty) = 1$) by convention. Note that it is also possible to construct examples where $0 < \Pr(T_{i,j} = \infty) < 1$.

If the process can become stuck in a single state (e.g., let $j$ be the state at the end of a game), then that state is called *absorbing* and $P_{j,j} = 1$. If state $j$ is absorbing, then the distribution of the hitting time $T_{i,j}$ satisfies

$$\Pr\left(T_{i,j} \leq m\right) = \Pr\left(X_{m+1} = j \,|\, X_1 = i\right) = [P^m]_{i,j}.$$

This follows from initializing $X_1 = i$ and observing the equality of the two events "$T_{i,j} \leq m$" and "$X_{m+1} = j$". Since $X_1 = i$, if $X_{m+1} = j$, then we clearly have $T_{i,j} \leq m$. On the other hand, if $T_{i,j} \leq m$, then $X_t = j$ for some $t \leq m+1$ and, thus, $X_{m+1} = j$ because state $j$ is absorbing. Hence, for $m \geq 1$, we find that

$$\Pr(T_{i,j} = m) = \Pr\left(T_{i,j} \leq m\right) - \Pr\left(T_{i,j} \leq m-1\right) = \left[P^m - P^{m-1}\right]_{i,j}.$$


Figure 1: Miniature chutes and ladders game

More generally, if state $j$ is not absorbing and $m \geq 0$, then extending all paths by one shows that

$$\Pr\left(T_{i,j} \leq m+1\right) = \begin{cases} 1 & \text{if } i = j \\ \sum_{k=1}^{n} P_{i,k} \Pr\left(T_{k,j} \leq m\right) & \text{otherwise.} \end{cases}$$

This implies that $\phi_{i,j}^{(m)} \triangleq \Pr\big(T_{i,j} \leq m\big)$ satisfies the recursion, starting from $\phi_{i,j}^{(0)} = \delta_{i,j}$, given by

$$\phi_{i,j}^{(m+1)} = \delta_{i,j} + (1 - \delta_{i,j}) \sum_{k=1}^{n} P_{i,k} \phi_{k,j}^{(m)}.$$

**Exercise 1.** (15 pts for compute_Phi_ET function + 10 pts for $\mathbb{E}\left[T_{1,4}\right]$ + 15 pts for simulate_hitting_time function + 10 pts for histogram/mean) What is the distribution of the number of fair coin tosses before one observes 3 heads in a row? To solve this, consider a 4-state Markov chain with transition probability matrix

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

where $X_t = 1$ if the previous toss was tails, $X_t = 2$ if the last two tosses were tails then heads, $X_t = 3$ if the last three tosses were tails then heads twice, and $X_t = 4$ is an absorbing state that is reached when the last three tosses are heads. Write a computer program (e.g., in Python, Matlab, ...) to compute $\Pr(T_{1,4} = m)$ for $m = 1, 2, \ldots, 100$ and use this to compute and print an estimate of the expected number of tosses $\mathbb{E}\left[T_{1,4}\right]$. Write a computer program that generates 500 realizations from this Markov chain. Then, use them to plot a histogram of $T_{1,4}$ and compute/print an estimate of the expected number of tosses $\mathbb{E}\left[T_{1,4}\right]$.

**Exercise 2.** (10 pts for Markov chain + 10 pts for CDF/mean + 10 pts for histogram/mean) Consider the miniature chutes and ladders game shown in Figure 1. Assume a player starts on the space labeled 1 and plays by rolling a fair four-sided die and then moves that number of spaces. If a player lands on the bottom of a ladder, then they automatically climb to the top. If a player lands at the top of a slide, then they automatically slide to the bottom. This process can be modeled by a Markov chain with $n = 16$ states where each state is associated with a square where players can start their turn (i.e., players never start at the bottom of a ladder or the top of a slide). To finish the game, players must land exactly on space 20 (i.e., if your roll would take you beyond 20, then no move is made). Compute the transition probability matrix $P$ of the implied Markov chain. For this Markov chain, use the program from Exercise 1 to compute and plot the cumulative distribution of the number turns a player takes to finish (i.e., the probability $\Pr\big(T_{1,20} \leq m\big)$ where $T_{1,20}$ is the hitting time from state 1 to state 20). Compute and print the mean $\mathbb{E}[T_{1,20}]$. Use the program from Exercise 1 to generate 500 realizations from this Markov chain. Then, use them to plot a histogram of $T_{1,20}$ and compute/print an estimate of the expected number of tosses $\mathbb{E}\left[T_{1,20}\right]$.

**Optional Challenge Question:** If the first player rolls 3 and climbs the ladder to square 8, then what is the probability that the second player will win.

### 2.3 Recurrent States and Stationary Probabilities

A state is called *recurrent* if it is expected to return to itself infinitely many times. Two states reachable from each other are called *communicating*. If all pairs of states are communicating, the Markov chain is called *irreducible* and all states are recurrent. A state distribution $\underline{\pi}$ is called *stationary* if it satisfies $\underline{\pi}P = \underline{\pi}$. A matrix $A$ is called *positive* if $A_{i,j} > 0$ for all $i, j \in [n]$.

If the transition-probability matrix is positive (i.e., the process can transition to any state in one step), then the Markov chain has a unique stationary distribution. More generally, this holds for any Markov chain that is irreducible [1]. If the $n$-step transition-probability matrix $P^n$ is positive, then the stationary distribution $\underline{\pi}$ also satisfies the steady-state limit

$$\pi_i = \lim_{t\to\infty} \Pr(X_t = i),$$

which equals the expected fraction of time that the process spends in state $i$.

One can find the stationary distribution by first rewriting $\underline{\pi}P = \underline{\pi}$ as

$$(I - P)^T \underline{\pi}^T = \underline{0}.$$

Then, one can solve for $\underline{\pi}^T$ (up to normalization) by applying row reduction to $(I - P)^T$ and computing the one-dimensional basis of the null space. This can be computed in Matlab and Python using the functions `null` and `scipy.linalg.null_space`. Lastly, one must normalize the resulting basis vector by enforcing the condition $\sum_{j=1}^{n} \pi_j = 1$ (e.g., by dividing by the sum of its entries).

**Example 1.** In a certain city, it is said that the weather is rainy with a 90% probability if it was rainy the previous day and with a 50% probability if it not rainy the previous day. If we assume that only the previous day's weather matters, then we can model the weather of this city by a Markov chain with $n = 2$ states whose transitions are governed by

$$P = \begin{bmatrix} 0.9 & 0.1 \\ 0.5 & 0.5 \end{bmatrix}.$$

Under this model, what is the steady-state probability of rainy weather?

To find this, we solve for the stationary distribution. As described above, we write

$$(I - P)^T = \begin{bmatrix} 0.1 & -0.5 \\ -0.1 & 0.5 \end{bmatrix} \implies \begin{bmatrix} 1 & -5 \\ 0 & 0 \end{bmatrix},$$

where $\implies$ denotes row reduction. Thus, $\pi_1 - 5\pi_2 = 0$ and $\pi_1 + \pi_2 = 1$ imply that $\pi_1 = 1 - \pi_2 = 5/6$ is the steady state-probability of rainy weather.

**Exercise 3.** (10 pts for stationary_distribution function + 5 pts for MC1 + 5 pts for MC2) Write a program to compute the stationary distribution of a Markov chain when it is unique. Consider a game where the gameboard has 8 different spaces arranged in a circle. During each turn, a player rolls two 4-sided dice and moves clockwise by a number of spaces equal to their sum. Define the transition matrix for this 8-state Markov chain and compute its stationary distribution.

Next, suppose that one space is special (e.g., state-1 of the Markov chain) and a player can only leave this space by rolling doubles (i.e., when both dice show the same value). Again, the player moves clockwise by a number of spaces equal to their sum. Define the transition matrix for this 8-state Markov chain and compute its stationary probability distribution.

# 3 Convergence to the Stationary Distribution

Consider the vector space $V = \mathbb{R}^n$ of row vectors equipped with the norm $\|\underline{v}\|_1 = \sum_{i=1}^{n} |v_i|$ and let $X = \{\underline{x} \in V \mid x_i \geq 0, \sum_{i=1}^{n} x_i = 1\}$ be the subset of probability vectors.

**Lemma 1.** *Let $P$ be the $n \times n$ transition matrix of a Markov chain (i.e., $P_{i,j} \geq 0$ and $\sum_{j=1}^{n} P_{i,j} = 1$). If $P_{i,j} \geq \alpha \geq 0$ for all $i, j \in [n]$, then, for all $\underline{x}, \underline{y} \in X$, one finds that*

$$\left\|\underline{x}P - \underline{y}P\right\|_1 \leq (1 - n\alpha) \left\|\underline{x} - \underline{y}\right\|_1$$

*Proof.* For $\underline{x}, \underline{y} \in X$, observe that

$$
\begin{aligned}
\left\|\underline{x}P - \underline{y}P\right\|_1 &= \sum_{j=1}^{n} \left| \sum_{i=1}^{n} (x_i - y_i)(P_{i,j} - \alpha + \alpha) \right| \\
&= \sum_{j=1}^{n} \left| \sum_{i=1}^{n} (x_i - y_i)(P_{i,j} - \alpha) + \sum_{i=1}^{n} (x_i - y_i)\alpha \right| \\
&\overset{(a)}{=} \sum_{j=1}^{n} \left| \sum_{i=1}^{n} (x_i - y_i)(P_{i,j} - \alpha) \right| \\
&\overset{(b)}{\leq} \sum_{j=1}^{n} \sum_{i=1}^{n} |x_i - y_i| \, (P_{i,j} - \alpha) \\
&= \sum_{i=1}^{n} |x_i - y_i| \sum_{j=1}^{n} (P_{i,j} - \alpha) \\
&\leq \sum_{i=1}^{n} |x_i - y_i| \, (1 - n\alpha) \\
&= (1 - n\alpha) \left\|\underline{x} - \underline{y}\right\|_1 ,
\end{aligned}
$$

where $(a)$ follows from $\sum_{i=1}^{n} (x_i - y_i) = 0$ and $(b)$ holds because $P_{i,j} - \alpha \geq 0$ and $|\sum_i z_i| \leq \sum_i |z_i|$. □

**Theorem 1** (Perron)**.** *Let $P$ be the transition matrix of a Markov chain (i.e., $P_{i,j} \geq 0$ and $\sum_{j=1}^{n} P_{i,j} = 1$). If there is a fixed $N$ such that $\left[P^N\right]_{i,j} \geq \alpha > 0$ for all $i, j \in [n]$, then the iteration*

$$\underline{\pi}^{(t+1)} = P \underline{\pi}^{(t)}$$

*has a unique fixed-point $\underline{\pi}^*$ and $\left\|\underline{\pi}^{(t)} - \underline{\pi}^*\right\|_1 \leq 2(1 - \alpha n)^{\lfloor (t-1)/N \rfloor}$ from any starting point $\underline{\pi}^{(1)} \in X$. A Markov chain satisfying this condition is called irreducible and aperiodic. In addition, the fixed-point vector is strictly positive and satisfies $\pi_i^* \geq \alpha$ for all $i \in [n]$.*

*Proof.* First, we define $f(\underline{x}) \triangleq \underline{x}P$ and observe that $f \colon X \to X$ because $[\underline{x}P]_j = \sum_{i=1}^{n} x_i P_{i,j} \geq 0$ and

$$\sum_{j=1}^{n} [\underline{x}P]_j = \sum_{j=1}^{n} \sum_{i=1}^{n} x_i P_{i,j} = \sum_{i=1}^{n} x_i \sum_{j=1}^{n} P_{i,j} = 1.$$

Next, we apply Lemma 1 to see that $f^N(\underline{x}) \triangleq \underline{x}P^N$ is a contraction on $(X, d)$ with $d\left(\underline{x}, \underline{y}\right) = \left\|\underline{x} - \underline{y}\right\|_1$ because

$$d\left(f^N(\underline{x}), f^N(\underline{y})\right) = \left\|\underline{x}P^N - \underline{y}P^N\right\|_1 \leq (1 - \alpha n) \left\|\underline{x} - \underline{y}\right\|_1 = (1 - \alpha n) d\left(\underline{x}, \underline{y}\right).$$

Thus, the contraction mapping theorem shows that $f^N$ has a unique fixed-point $\underline{\pi}^* \in X$ and that $\underline{\pi}^{(t+N)} = f^N(\underline{\pi}^{(t)})$ converges (along the subsequence $t_k = kN + 1$) to $\underline{\pi}^*$ for any $\underline{\pi}^{(1)} \in X$ . For $t = t_k + s$ with $s \in [N-1]$, we note that $\underline{\pi}^{(t_k+s)} = f^{kN}\left(f^s(\underline{\pi}^{(1)})\right)$ also converges to to $\underline{\pi}^*$ as $k \to \infty$.

The rate of convergence follows from $\|\underline{\pi} - \underline{\pi}^*\|_1 \leq 2$ and

$$\left\|f^{kN}\left(\underline{\pi}\right) - \underline{\pi}^*\right\|_1 = \left\|f^{kN}\left(\underline{\pi} - \underline{\pi}^*\right)\right\|_1 \leq (1 - \alpha n)^k \left\|\underline{\pi} - \underline{\pi}^*\right\|_1 \leq 2(1 - \alpha n)^k,$$

which both hold for $\underline{\pi} \in X$. This construction leads to $k = \lfloor (t-1)/N \rfloor$ and gives the stated bound. Lastly, we observe that the elements of any vector in the range of $f^N$ satisfy the following strictly positive lower bound

$$\left[\underline{x}P^N\right]_j = \sum_{i=1}^{n} x_i \left[P^N\right]_{i,j} \geq \min_{i,j} \left[P^N\right]_{i,j} \sum_{k=1}^{n} x_k = \min_{i,j} \left[P^N\right]_{i,j} \geq \alpha.$$

This implies that $\pi_i^*$ is lower bounded by the same quantity for all $i \in [n]$. □

# References


[1] R. Durrett, *Elementary probability for applications.* Cambridge University Press, 2009.