ECE 590.06 Application: Markov Chains

Henry D. Pfister
ECE Department
Duke University

November 1st, 2016

1 A Few Simple Questions

1.1 What is the chance 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 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 linearly 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 a 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://www.math.uiuc.edu/~bishop/monopoly.pdf`

2 What is a Markov chain?

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 homogenous). In the remainder of this note, we assume the FSMC is time invariant and we let $P \in \mathbb{R}^{n\times n}$ denote *transition-probability matrix* with entries $[P]_{i,j} = P_{i,j} = \Pr(X_{t+1} = j \,|\, X_t = i)$. 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

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

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

For a Markov chain starting from state i, the first *hitting time* of state j is a random variable $T_{i,j}$ with distribution

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

If state j is not reachable from state i, then $\Pr(T_{i,j}=\infty)=1$ by convention. If state j is absorbing, then this probability is easy calculated by noticing that

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

Thus, for $m\geq 1$, we find that $\Pr(T_{i,j}=m)=\left[P^m-P^{m-1}\right]_{i,j}$. If state j is not absorbing and $m\geq 1$, then the probabilities must satisfy

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

With a little work, one can show that they are given by the limit of the recursion

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

starting from $\phi_{i,j}^{(1)}=\delta_{i,j}+(1-\delta_{i,j})P_{i,j}$, where $\delta_{i,j}$ is Kronecker delta function.

Exercise 2.1. 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. Use Matlab to compute $\Pr(T_{1,4}=m)$ for $m=1,2,\ldots,100$ and use this to estimate expected number of tosses $\mathbb{E}\left[T_{1,4}\right]$.

A state that will be visited infinitely many times is called *recurrent* and any FSMC without an absorbing state must have recurrent states. If $P_{i,j}>0$ for all $i,j\in[n]$, then the Markov chain can transition to any state at any time and every state is recurrent. In this case, it makes sense to discuss the long-run fraction of time spent in a state and this is given by

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

These probabilities equal the unique solution of the linear system $\pi_j=\sum_{i=1}^{n}\pi_i P_{i,j}$ satisfying $\sum_{j=1}^{n}\pi_j=1$. More generally, this holds for any Markov chain that is irreducible and aperiodic. While we do not define these terms in this note, the advanced reader will see that our final theorem holds also in this case. Finally, one can solve for $\underline{\pi}$ by using row reduction to find a one-dimensional basis for the null space of $(I-P)^T$ and then normalizing the basis vector so that it sums to 1.

Exercise 2.2. 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 probability 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 Contractions on Compact Spaces

First, we recall a modified version of the contraction mapping theorem that applies to compact spaces.

Definition 3.1. Let (X, d) be a compact metric space and let $f\colon X \to X$ be a mapping whose N-th order composition is denoted f^N. We say that f is a *weak contraction* on X if, for some fixed N, f^N satisfies

$$d\left(f^N(\underline{x}), f^N(\underline{y})\right) \le d\left(\underline{x}, \underline{y}\right)$$

for all $\underline{x}, \underline{y} \in X$ with equality if and only if $\underline{x} = \underline{y}$.

Theorem 3.2. *If f is a weak contraction on a compact metric space, then f has a unique fixed point $\underline{x}^*$ and the sequence $\underline{x}_{t+1} = f(\underline{x}_t)$ converges to $\underline{x}^*$ from any $\underline{x}_1 \in X$.*

Sketch of Proof. This can be shown by combining Theorem 3.1 and Theorem 3.5 in "The contraction mapping theorem" by Keith Conrad. Here we only provide a brief outline.

The first step is to establish the conclusion for f^N. To do this, one analyzes $d(f^N(\underline{x}), \underline{x})$ as a function of $\underline{x}$. Since this is a continuous function of $\underline{x}$ on the compact set X, we can assume the minimum value is achieved by some $\underline{a}$. But, by the contraction condition, we find that

$$d\left(f^N(\underline{a}), f^N(f^N(\underline{a}))\right) \le d\left(\underline{a}, f^N(\underline{a})\right)$$

with equality if and only if $f^N(\underline{a}) = f^N(f^N(\underline{a}))$ (i.e., $\underline{a}$ is a fixed point of f^N). Since a strict inequality would contradict the minimality of $d\left(\underline{a}, f^N(\underline{a})\right)$, we conclude that $\underline{a}$ must be a fixed point of f^N.

The convergence of $\underline{x}_{Nt}$ to $\underline{a}$ can be established by considering the sequence $d(\underline{x}_{N(t+1)}, \underline{a})$. Since the contraction property implies that $d(\underline{x}_{N(t+1)}, \underline{a})$ is non-increasing, it must converge to a limit. Also, for any subsequence of $\underline{x}_{Nt}$ that converges to $\underline{x}$, continuity implies that $d(\underline{x}_{N(t+1)}, \underline{a}) \to d(\underline{x}, \underline{a}) = d(f^N(\underline{x}), \underline{a})$. By the contraction propery, this can only happen if $\underline{x} = \underline{a} = f^N(\underline{x})$.

One can extend these results to f by noting that $\underline{a} = f^N(\underline{a})$ implies $f(\underline{a}) = f(f^N(\underline{a})) = f^N(f(\underline{a}))$. From this, we see that $f(\underline{a})$ also equals the unique fixed point of f^N and, thus, $\underline{a} = f(\underline{a})$. Also, $\underline{x}_{s+1} = f(\underline{x}_s)$ for $s = Nt + r$ with $r \in \{0, 1, \ldots, N-1\}$ implies that $\underline{x}_s = f^{Nt+r+1}(\underline{x}_1) = f^{Nt}(f^{r+1}(\underline{x}_1))$. But, since $f^{Nt}(\underline{x}) \to \underline{a}$ as $N \to \infty$ for all $\underline{x} \in X$, this implies that $\underline{x}_s \to \underline{a}$ as $s \to \infty$. □

4 Convergence to Stationarity

Consider the compact metric space (X, d) defined by $X = \{\underline{x} \in \mathbb{R}^n \mid x_i \ge 0, \sum_{i=1}^n x_i = 1\}$ and $d(\underline{x}, \underline{y}) = \left\|\underline{x} - \underline{y}\right\|_1$. Next, let $\pi_i^{(t)} = \Pr(X_t = i)$ and observe that

$$\begin{aligned}
\pi_j^{(t+1)} &= \sum_{i=1}^n \Pr(X_{t+1} = j, X_t = i) \\
&= \sum_{i=1}^n \Pr(X_{t+1} = j | X_t = i) \Pr(X_t = i) \\
&= \sum_{i=1}^n P_{i,j} \pi_i^{(t)}.
\end{aligned}$$

If we let $\underline{\pi}^{(t)}$ denote the row vector $(\pi_1^{(t)}, \ldots, \pi_n^{(t)})$ and $f(\underline{x}) = \underline{x}P$, then the previous equation can be rewritten as

$$\underline{\pi}^{(t+1)} = f\left(\underline{\pi}^{(t)}\right),$$

where $f\colon X \to X$. The following lemmas connects this update to the contraction mapping theorem.

Lemma 4.1. *For $\underline{z} \in \mathbb{R}^n$, $\sum_{i=1}^n z_i \le \sum_{i=1}^n |z_i|$ with equality if and only if all z_i have the same sign.*

Proof. If all z_i have the same sign, then we clearly have equality. If not, this follows from

$$\sum_{i=1}^n z_i = \sum_{i: z_i \ge 0} z_i + \sum_{i: z_i < 0} z_i < \sum_{i: z_i \ge 0} z_i - \sum_{i: z_i < 0} z_i = \sum_{i=1}^n |z_i| .$$

□

Lemma 4.2. *Let Q be an $n \times n$ real matrix whose entries satisfy $Q_{i,j} > 0$. Then, we have*

$$\left\| \underline{x} Q - \underline{y} Q \right\|_1 \le \left\| \underline{x} - \underline{y} \right\|_1$$

for all $\underline{x}, \underline{y} \in X$ with equality if and only if $\underline{x} = \underline{y}$.

Proof. For $\underline{x}, \underline{y} \in X$, we write

$$\begin{aligned}
\left\| \underline{x} Q - \underline{y} Q \right\|_1 &= \sum_{j=1}^n \left| \sum_{i=1}^n (x_i - y_i) Q_{i,j} \right| \\
&\overset{(a)}{\le} \sum_{j=1}^n \left(\sum_{i=1}^n |(x_i - y_i) Q_{i,j}| \right) \\
&= \sum_{i=1}^n |x_i - y_i| \sum_{j=1}^n Q_{i,j} \\
&= \left\| \underline{x} - \underline{y} \right\|_1 ,
\end{aligned}$$

where, by Lemma 4.1, (a) holds with equality if and only if $z_j = \sum_{i=1}^n (x_i - y_i) Q_{i,j}$ has the same sign for all $j \in [n]$. But, if z_j has the same sign for all $j \in [n]$, then

$$\sum_{j=1}^n z_j = \sum_{j=1}^n \sum_{i=1}^n (x_i - y_i) Q_{i,j} = \sum_{i=1}^n (x_i - y_i) = \sum_{i=1}^n x_i - \sum_{i=1}^n y_i = 1 - 1 = 0$$

implies that all $z_j = 0$. Therefore, we conclude that equality occurs if and only if $\underline{x} = \underline{y}$. □

Theorem 4.3 (Perron)**.** *Let P be the transition matrix of a Markov chain (i.e., $P_{i,j} \ge 0$ and $\sum_{j=1}^n P_{i,j} = 1$). If there is a fixed N such that $\left[P^N\right]_{i,j} > 0$ 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 $\underline{\pi}^{(t)} \to \underline{\pi}^*$ 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

$$\pi_i^* \ge \min_{i,j} \left[P^N\right]_{i,j} .$$

Proof. First, we define $f(\underline{x}) \triangleq \underline{x} P$ and observe that $f\colon X \to X$ because $[\underline{x} P]_j = \sum_{i=1}^n x_i P_{i,j} \ge 0$ 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 4.2 to see that f is a weak contraction on (X, d) because

$$d\left(f^N(\underline{x}), f^N(\underline{y})\right) = \left\| \underline{x} P^N - \underline{y} P^N \right\|_1 \le \left\| \underline{x} - \underline{y} \right\|_1 = d\left(\underline{x}, \underline{y}\right) ,$$

with equality if and only if $\underline{x} = \underline{y}$. Then, we apply Theorem 3.2 to see that f has a unique fixed point and that $\underline{\pi}^{(t+1)} = f(\underline{\pi}^{(t)})$ converges to that fixed point. Lastly, we observe that elements of any vector in the range of f^N satisfy the following strictly positive lower bound

$$\left[\underline{x} P^N\right]_j = \sum_{i=1}^n x_i \left[P^N\right]_{i,j} \ge \min_{i,j} \left[P^N\right]_{i,j} \sum_{k=1}^n x_k = \min_{i,j} \left[P^N\right]_{i,j} > 0.$$

This implies that π_i^* is lower bounded by the same quantity. □