Review (Basics)

Def A Markov chain is a collection of random variables $X_0, X_1, \dots$ over the same probability space with a common state space S such that

$$\mathbb{P}(X_{n+1} = i_{n+1} \mid X_0 = i_0, X_1 = i_1, \dots, X_{n-1} = i_{n-1}, X_n = i_n) = \mathbb{P}(X_{n+1} = i_{n+1} \mid X_n = i_n)$$

for any $i_0, i_1, \dots, i_{n+1} \in S$

In this class, a Markov chain is specified by:

① A finite state space S: $|S| < \infty$

② A time-homogeneous transition matrix $\mathbb{P}$:

$$\mathbb{P}(i,j) = \mathbb{P}(X_1 = j \mid X_0 = i) = \mathbb{P}(X_{n+1} = j \mid X_n = i) \text{ for any } n \geq 0.$$

③ An initial distribution π_0:

$$\sum_{k \in S} \pi_0(k) = 1$$

Thm Let $\mathbb{P}_n$ be the n-step transition matrix $\mathbb{P}_n(i,j) = \mathbb{P}(X_n = j \mid X_0 = i)$

π_n be the distribution over the states after n steps.

Then

$$\mathbb{P}_n = \mathbb{P}^n$$

$$\pi_n = \pi_0 \mathbb{P}^n$$

Review (Conditioning)

<u>Thm</u> Let $A, B \subset S$ be mutually exclusive. Let T_A, T_B be the times until we visit a state in A and B, respectively.
Define $\alpha(i) = \mathbb{P}(T_A < T_B \mid X_0 = i)$, $\beta(i) = \mathbb{E}[T_A \mid X_0 = i]$
Then

$\alpha(i) = 1$ if $i \in A$

$\alpha(i) = 0$ if $i \in B$

$$\begin{aligned}\alpha(i) &= \sum_{j \in S} \mathbb{P}(T_A < T_B, X_1 = j \mid X_0 = i) \\ &= \sum_{j \in S} \mathbb{P}(T_A < T_B \mid X_0 = i, X_1 = j) \cdot \mathbb{P}(X_1 = j \mid X_0 = i) \\ &= \sum_{j \in S} \mathbb{P}(T_A < T_B \mid X_0 = j)\, \underline{\mathbb{P}}(i,j) \quad \text{by homogeneity} \\ &= \sum_{j \in S} \alpha(j)\, \underline{\mathbb{P}}(i,j) \quad \text{if } i \notin A,\ i \notin B\end{aligned}$$

$\beta(i) = 0$ if $i \in A$

$$\begin{aligned}\beta(i) &= \sum_{j \in S} \mathbb{E}[T_A \mid X_0 = i, X_1 = j] \cdot \mathbb{P}(X_1 = j \mid X_0 = i) \\ &= \sum_{j \in S} \left(1 + \mathbb{E}[T_A \mid X_0 = j]\right) \cdot \underline{\mathbb{P}}(i,j) \quad \text{by homogeneity} \\ &= \sum_{j \in S} \underline{\mathbb{P}}(i,j) + \sum_{j \in S} \beta(j)\, \underline{\mathbb{P}}(i,j) \\ &= 1 + \sum_{j \in S} \beta(j)\, \underline{\mathbb{P}}(i,j) \quad \text{if } i \notin A\end{aligned}$$

<u>Thm</u> Let V_A be the number of times a state in A is visited.
Define $\gamma(i) = \mathbb{E}[V_A \mid X_0 = i]$.
Then

$$\gamma(i) = 1 + \sum_{j \in S} \gamma(j)\, \underline{\mathbb{P}}(i,j) \quad \text{if } i \in A$$

$$\gamma(i) = \sum_{j \in S} \gamma(j)\, \underline{\mathbb{P}}(i,j) \quad \text{if } i \notin A$$

Note: $\mathbb{E}[V_A]$ may be infinite.

Compartmental Model

Ex Suppose that for some disease, the population can be split into three groups: susceptible, infected, and removed.

Each day,

- of the susceptible, 20% become infected;
  80% stay susceptible.
- of the infected, 20% become susceptible;
  10% become removed;
  70% stay infected
- of the removed, p become susceptible
  $1-p$ stay removed.

In the population prior to exposure, 95% are susceptible and 5% are removed.

We can model this as a Markov process:

$S = \{S, I, R\}$

$$\mathbb{P} = \begin{array}{c c} & \begin{array}{ccc} S & I & R \end{array} \\ \begin{array}{c} S \\ I \\ R \end{array} & \begin{bmatrix} 0.8 & 0.2 & 0 \\ 0.2 & 0.7 & 0.1 \\ p & 0 & 1-p \end{bmatrix} \end{array}$$

$$\pi_0 = [\, 0.95 \quad 0 \quad 0.05 \,]$$

[Demo]

Invariant Distributions

Def A distribution π is invariant for transition matrix $\mathbb{P}$ if
$$\pi = \pi \mathbb{P}$$
These equations are called the balance equations.

Thm $\pi_n = \pi_0$ for all $n \geq 0$ if and only if π_0 is invariant

Pf Suppose $\pi_n = \pi_0$ for all $n \geq 0$. Then
$$\pi_1 = \pi_0 \mathbb{P} = \pi_0,$$
so π_0 is invariant.

Suppose $\pi_0 = \pi_0 \mathbb{P}$. Then
$$\pi_1 = \pi_0 \mathbb{P} = \pi_0$$
$$\pi_{n+1} = \pi_n \mathbb{P} = \pi_0$$
By induction, $\pi_n = \pi_0$ for all $n \geq 0$

Note A Markov chain may have many invariant distributions. For example, $\mathbb{P} = \mathbb{I}$ has infinitely many.

Note Invariance means that the net flow in and out of states is equal.

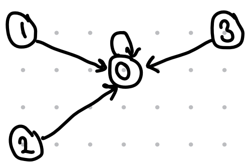

$\pi(0)$ of the particles are leaving state 0

$\pi(0) \cdot \mathbb{P}(0,0) + \pi(1) \cdot \mathbb{P}(1,0) + \pi(2) \cdot \mathbb{P}(2,0) + \pi(3) \cdot \mathbb{P}(3,0)$ are entering state 0

The balance equations say these flows are equal:
$$\underset{\substack{\uparrow \\ \text{leaving } i}}{\pi(i)} = \underbrace{\sum_{j \in S} \pi(j) \cdot \mathbb{P}(j,i)}_{\text{entering } i}$$

Irreducibility

Def We say i can reach j ($i \to j$) if there is a path with positive probability from i to j.

A Markov chain is irreducible if
$\forall i, j$ $\quad i \to j$ and $j \to i$ $\quad$ (you can get from any state to any state)

Ex Which of the following chains is irreducible?

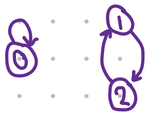

$0 \not\to 1$

Reducible

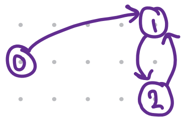

$1 \not\to 0$

Reducible

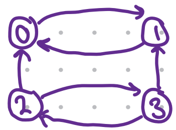

$0 \not\to 2$

Reducible

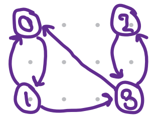

Irreducible

To show a chain is irreducible, construct a path that
- starts at any state
- goes through all of the other states
- ends at the starting state
- has positive probability

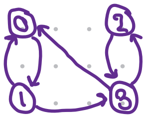

$0 \to 1 \to 3 \to 2 \to 3 \to 0$

Long Run Proportion of Time

Thm If a finite Markov chain is irreducible, then, for any initial distribution π_0,

$$\lim_{n\to\infty} \frac{1}{n}\sum_{m=0}^{n-1} \mathbb{1}\{X_m = i\} = \pi(i) \quad \text{for all } i \in S$$

That is, the fraction of time spent in each state is given by π.
Also, π is invariant; therefore the invariant distribution exists and is unique

Pf Because the chain is irreducible, each state will be visited infinitely many times.
For each state i, let T_i be the number of steps to return to i starting with $X_0 = i$. Then let

$\pi(i) = \dfrac{1}{\mathbb{E}[T_i]}$, the fraction of time spent in state i

So the fraction of time the chain spends in state i is $\pi(i)$, as desired.
Now, π is invariant. Note that in n steps, the chain is in state i for $n\pi(i)$ steps.
Over large n, the chain transitions from any state j to i

$$\underbrace{n\pi(j)}_{\text{in state } j} \cdot \underbrace{\mathbb{P}(j,i)}_{\text{transition to state } i}$$

times. The total visits to state i is the sum of all visits from the states:

$$n\pi(i) = \sum_{j\in S} n\pi(j)\mathbb{P}(j,i)$$

This can be written as $\pi = \pi\mathbb{P}$.
Proof of uniqueness is more complicated.

Note An irreducible Markov chain's distribution does not necessarily converge to π:

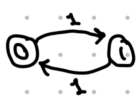

The chain spends half the time in each state, but every initial distribution does not converge to $[1/2, 1/2]$, e.g.

$[0 \;\; 1] \to [1 \;\; 0] \to [0 \;\; 1] \to \dots$

This is because of the periodic behavior of the chain.

Periodicity

Def The period of a state i, denoted $d(i)$, is

$$d(i) = \gcd\left\{ n : \mathbb{P}_n(i,i) > 0 \right\};$$

that is, the period of a state is the greatest common divisor of the lengths of paths from i to i.

In this class we only consider the period of states in irreducible chains.

Ex Find $d(1)$ for each of the following chains

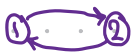

$1 \to 2 \to 1$; $1 \to 2 \to 1 \to 2 \to 1$

$d(1) = \gcd\{2, 4, 6, \dots\} = 2$

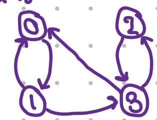

$1 \to 0 \to 1$; $1 \to 3 \to 0 \to 1$;

$d(1) = \gcd\{2, 3, \dots\} = 1$

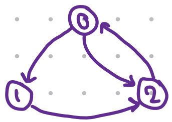

$1 \to 2 \to 0 \to 1$; $1 \to 2 \to 0 \to 2 \to 1$

$d(1) = \gcd\{3, 4, \dots\} = 1$

Clm If a Markov chain is irreducible, the periods of all states are the same:

$d(i) = d(j)$ for all $i, j \in S$

Pf As an exercise

Def If $d(i) = 1$ for any state in an irreducible Markov chain, we say the chain is aperiodic

Ex Which of the following chains is aperiodic?

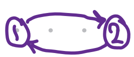

periodic, since $d(1) = 2$

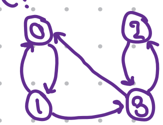

aperiodic, since $d(1) = 1$

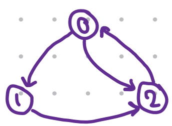

aperiodic, since $d(1) = 1$

To show a chain is aperiodic, find two loops with coprime lengths

Markov Chain Convergence Theorem

Thm (Finite Markov Chain Convergence) If $X_0, X_1, \dots$ is a Markov chain on S with time-homogeneous transition matrix $\mathbb{P}$ and

1. $|S| < \infty$
2. $\mathbb{P}$ is irreducible: $\forall i,j \in S,\ \exists n\ \mathbb{P}_n(i,j) > 0$
3. $\mathbb{P}$ is aperiodic: $\forall i \in S,\ d(i) = 1$,

then $X_0, X_1, \dots$ has an invariant distribution $\pi = \pi \mathbb{P}$ such that

$$\lim_{n\to\infty} \mathbb{P}_n(i,j) = \pi(j)$$

i.e. the n-step transition probabilities converge to π.

Moreover,

$$\lim_{n\to\infty} \frac{1}{\mathbb{E}[T_j]} = \pi(j),$$

i.e. $\pi(j)$ is the long-run proportion of time spent in state j.

Cor $\pi_n(i) = \mathbb{P}(X_n = i) \to \pi(i)$ as $n \to \infty$

Pf

$$\begin{aligned}
\pi_n(i) &= \sum_{j\in S} \pi_0(j)\, \mathbb{P}_n(j,i) \\
&\to \sum_{j \in S} \pi_0(j)\, \pi(i) \quad \text{as } n\to\infty \\
&= \pi(i) \sum_{j\in S} \pi_0(j) \\
&= \pi(i)
\end{aligned}$$

Note We can find the long-run probability of an event by conditioning on π.

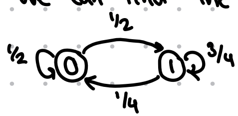

$\pi = [\frac{1}{3}, \frac{2}{3}]$

What is the long-run probability that we stay in the same state?

$$\begin{aligned}
\mathbb{P}(\text{stay}) &= \mathbb{P}(\text{stay} \mid \text{state } 0)\cdot \mathbb{P}(\text{state } 0) + \mathbb{P}(\text{stay} \mid \text{state } 1)\cdot \mathbb{P}(\text{state } 1) \\
&= \frac{1}{2}\cdot\frac{1}{3} + \frac{3}{4}\cdot\frac{2}{3} \\
&= \frac{2}{3}
\end{aligned}$$

Ehrenfest Chain 1

Q In the Ehrenfest model, there are two containers, containing a total of N particles.
At each step
- a container is selected uniformly at random
- a particle is selected uniformly at random, independently of the container

and the selected particle is placed in the selected container; if the particle was already in the container, it remains in place.
Let X_n be the number of particles in the first container at time n.
What are the transition probabilities of the chain?

The number of particles can either increase by 1, decrease by 1, or remain the same.

$$\mathbb{P}(X_{n+1} = i+1 \mid X_n = i) = \frac{1}{2} \cdot \frac{N-i}{N} = \frac{N-i}{2N}$$

$$\mathbb{P}(X_{n+1} = i-1 \mid X_n = i-1) = \frac{1}{2} \cdot \frac{i}{N} = \frac{i}{2N}$$

$$\mathbb{P}(X_{n+1} = i \mid X_n = i) = 1 - \frac{i}{2N} - \frac{N-i}{2N} = \frac{1}{2}$$

So

$$\mathbb{P}(i,j) = \begin{cases} \frac{N-i}{2N} & \text{if } j = i+1 \\ \frac{1}{2} & \text{if } j = i \\ \frac{i}{2N} & \text{if } j = i-1 \\ 0 & \text{otherwise} \end{cases}$$

Prove that any distribution over the states converges to some distribution π.

We must show that the chain is irreducible and aperiodic.

Irreducible: Consider the path $0 \to 1 \to 2 \to \dots \to N \to N-1 \to \dots \to 0$.

Aperiodic: note that $d(0) = \gcd\{1, \dots\} = 1$. So all states have period 1.

Ehrenfest Chain II

Q In the Ehrenfest model, there are two containers, containing a total of N particles.
At each step
- a container is selected uniformly at random
- a particle is selected uniformly at random, independently of the container

and the selected particle is placed in the selected container; if the particle was already in the container, it remains in place.
Let X_n be the number of particles in the first container at time n.
Find the stationary distribution of the chain.

The balance equations are

$$\pi(0) = \frac{1}{2}\pi(0) + \frac{1}{2N}\pi(1)$$

$$\pi(j) = \frac{N-(j-1)}{2N}\pi(j-1) + \frac{1}{2}\pi(j) + \frac{j+1}{2N}\pi(j+1) \quad \text{for } 1 \le j \le N-1$$

$$\pi(N) = \frac{1}{2N}\pi(N-1) + \frac{1}{2}\pi(N)$$

Rewrite the first few equations in terms of $\pi(0)$

$$\pi(0) = \frac{1}{2}\pi(0) + \frac{1}{2N}\pi(1) \Rightarrow \pi(1) = N\pi(0) = \binom{N}{1}\pi(0)$$

$$\pi(1) = \frac{N}{2N}\pi(0) + \frac{1}{2}\pi(1) + \frac{2}{2N}\pi(2) \Rightarrow \pi(2) = \frac{N}{2}\left(\pi(1) - \pi(0)\right) = \frac{N(N-1)}{2}\pi(0) = \binom{N}{2}\pi(0)$$

By induction, $\pi(j) = \binom{N}{j}\pi(0)$.
Then

$$\sum_{j=0}^{N}\pi(j) = \sum_{j=0}^{N}\binom{N}{j}\pi(0) = \pi(0)\sum_{j=0}^{N}\binom{N}{j}1^j1^{N-j} = \pi(0)2^N = 1 \Rightarrow \pi(0) = \frac{1}{2^N}$$

So

$$\pi(k) = P(X_n = k) = \binom{N}{k}\frac{1}{2^N},$$

i.e. $\pi(k) \sim \text{Binomial}(N, 1/2)$.