CS70: Lecture 25.

Markov Chains: Distributions; Continuous Probability

1. Review
2. Distribution (Cont'd)
3. Irreducibility
4. Convergence
5. Continuous Probability: Introduction

Review

▶ Markov Chain:
  ▶ Finite set $\mathscr{X}$; π_0; $P = \{P(i,j), i,j \in \mathscr{X}\}$;
  ▶ $Pr[X_0 = i] = \pi_0(i), i \in \mathscr{X}$
  ▶ $Pr[X_{n+1} = j \mid X_0, \ldots, X_n = i] = P(i,j), i,j \in \mathscr{X}, n \geq 0$.
  ▶ Note:
    $Pr[X_0 = i_0, X_1 = i_1, \ldots, X_n = i_n] = \pi_0(i_0)P(i_0,i_1)\cdots P(i_{n-1},i_n)$.
▶ First Passage Time:
  ▶ $A \cap B = \emptyset; \beta(i) = E[T_A | X_0 = i]; \alpha(i) = P[T_A < T_B | X_0 = i]$
  ▶ $\beta(i) = 1 + \sum_j P(i,j)\beta(j); \alpha(i) = \sum_j P(i,j)\alpha(j)$.

Balance Equations

Question: Is there some π_0 such that $\pi_m = \pi_0, \forall m$?

Definition A distribution π_0 such that $\pi_m = \pi_0, \forall m$ is said to be an invariant distribution.

Theorem A distribution π_0 is invariant iff $\pi_0 P = \pi_0$. These equations are called the balance equations.

Proof: $\pi_n = \pi_0 P^n$, so that $\pi_n = \pi_0, \forall n$ iff $\pi_0 P = \pi_0$. □

Thus, if π_0 is invariant, the distribution of X_n is always the same as that of X_0.

Of course, this does not mean that X_n does not move. It means that the probability that it leaves a state i is equal to the probability that it enters state i.

The balance equations say that $\sum_j \pi(j)P(j,i) = \pi(i)$.
That is,

$$\sum_{j \neq i} \pi(j)P(j,i) = \pi(i)(1 - P(i,i)) = \pi(i)\sum_{j \neq i} P(i,j).$$

Thus, $Pr[\text{enter } i] = Pr[\text{leave } i]$.

Poll

Suppose $\pi(j)P(j,i) = \pi(i)P(i,j), \forall i,j$. Select all true statements.
▶ π is certainly an invariant distribution.
▶ π might not be an invariant distribution.
▶ π can't be an invariant distribution.

Balance Equations

Theorem A distribution π_0 is invariant iff $\pi_0 P = \pi_0$. These equations are called the balance equations.
Example 1:

$\pi P = \pi \quad \Leftrightarrow \quad [\pi(1), \pi(2)] \begin{bmatrix} 1-a & a \\ b & 1-b \end{bmatrix} = [\pi(1), \pi(2)]$
$\Leftrightarrow \quad \pi(1)(1-a) + \pi(2)b = \pi(1)$ and $\pi(1)a + \pi(2)(1-b) = \pi(2)$
$\Leftrightarrow \quad \pi(1)a = \pi(2)b$.

These equations are redundant! We have to add an equation:
$\pi(1) + \pi(2) = 1$. Then we find

$$\pi = [\frac{b}{a+b}, \frac{a}{a+b}].$$

Balance Equations

Theorem A distribution π_0 is invariant iff $\pi_0 P = \pi_0$. These equations are called the balance equations.
**Example 2:**

$\pi P = \pi \Leftrightarrow [\pi(1), \pi(2)] \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = [\pi(1), \pi(2)] \Leftrightarrow \pi(1) = \pi(1)$ and $\pi(2) = \pi(2)$.

Every distribution is invariant for this Markov chain. This is obvious, since $X_n = X_0$ for all n. Hence, $Pr[X_n = i] = Pr[X_0 = i], \forall(i,n)$.

Irreducibility

Definition A Markov chain is irreducible if it can go from every state i to every state j (possibly in multiple steps).

Examples:

[A]   [B]   [C]

[A] is not irreducible. It cannot go from (2) to (1).

[B] is not irreducible. It cannot go from (2) to (1).

[C] is irreducible. It can go from every i to every j.

If you consider the graph with arrows when $P(i,j) > 0$, irreducible means that there is a single connected component.

Existence and uniqueness of Invariant Distribution

Theorem A finite irreducible Markov chain has one and only one invariant distribution.

That is, there is a unique positive vector $\pi = [\pi(1), \ldots, \pi(K)]$ such that $\pi P = \pi$ and $\sum_k \pi(k) = 1$.

Proof: See EE126, or MC Note. (Proof is not in scope.)
**Note:** We know already that some reducible Markov chains have multiple invariant distributions.

Fact: If a Markov chain has two different invariant distributions π and ν, then it has infinitely many invariant distributions. Indeed, $p\pi + (1-p)\nu$ is then invariant since

$$[p\pi + (1-p)\nu]P = p\pi P + (1-p)\nu P = p\pi + (1-p)\nu.$$

Long Term Fraction of Time in States

Theorem Let X_n be an irreducible Markov chain with invariant distribution π.

Then, for all i,

$$\frac{1}{n}\sum_{m=0}^{n-1} 1\{X_m = i\} \to \pi(i), \text{ as } n \to \infty.$$

The left-hand side is the fraction of time that $X_m = i$ during steps $0, 1, \ldots, n-1$. Thus, this fraction of time approaches $\pi(i)$.

Proof: See EE126. MC Note gives a plausibility argument.

□

Long Term Fraction of Time in States

Theorem Let X_n be an irreducible Markov chain with invariant distribution π. Then, for all i, $\frac{1}{n}\sum_{m=0}^{n-1} 1\{X_m = i\} \to \pi(i)$, as $n \to \infty$.

Example 1:

$$P = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \qquad \pi P = \pi \Rightarrow \pi = [1/2, 1/2]$$

The fraction of time in state 1 converges to $1/2$, which is $\pi(1)$.

Long Term Fraction of Time in States

Theorem Let X_n be an irreducible Markov chain with invariant distribution π. Then, for all i, $\frac{1}{n}\sum_{m=0}^{n-1} 1\{X_m = i\} \to \pi(i)$, as $n \to \infty$.

Example 2:

$$\pi = [0.4, 0.6]$$

Convergence to Invariant Distribution

Question: Assume that the MC is irreducible. Does π_n approach the unique invariant distribution π?

Answer: Not necessarily. Here is an example:

$$P = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \qquad \pi P = \pi \Rightarrow \pi = [1/2, 1/2]$$

Assume $X_0 = 1$. Then $X_1 = 2, X_2 = 1, X_3 = 2, \ldots$.

Thus, if $\pi_0 = [1, 0]$, $\pi_1 = [0, 1], \pi_2 = [1, 0], \pi_3 = [0, 1]$, etc.

Hence, π_n does not converge to $\pi = [1/2, 1/2]$.

Periodicity

Theorem Assume that the MC is irreducible. Then

$$d(i) := \text{g.c.d.}\{n > 0 \mid Pr[X_n = i \mid X_0 = i] > 0\}$$

has the same value for all states i.

Proof: See MC Note. □

Definition If $d(i) = 1$, the Markov chain is said to be aperiodic. Otherwise, it is periodic with period $d(i)$.

Example

[A]
[B]

[A]: $\{n > 0 \mid Pr[X_n = 1 | X_0 = 1] > 0\} = \{3, 6, 7, 9, 10, \ldots\} \Rightarrow d(1) = 1$.
$\{n > 0 \mid Pr[X_n = 2 | X_0 = 2] > 0\} = \{3, 4, \ldots\} \Rightarrow d(2) = 1$.
[B]: $\{n > 0 \mid Pr[X_n = 1 | X_0 = 1] > 0\} = \{3, 6, 9, \ldots\} \Rightarrow d(i) = 3$.
$\{n > 0 \mid Pr[X_n = 5 | X_0 = 5] > 0\} = \{6, 9, \ldots\} \Rightarrow d(5) = 3$.

Poll

X_n is a MC with the transition diagram as shown. Select all true statements.

- ▶ X_n is irreducible.
- ▶ X_n is reducible.
- ▶ X_n is periodic.
- ▶ X_n is aperiodic.

Convergence of π_n

Theorem Let X_n be an irreducible and aperiodic Markov chain with invariant distribution π. Then, for all $i \in \mathscr{X}$,

$$\pi_n(i) \to \pi(i), \text{ as } n \to \infty.$$

Proof See EE126, or MC Note. □

Example

Convergence of π_n

Theorem Let X_n be an irreducible and aperiodic Markov chain with invariant distribution π. Then, for all $i \in \mathscr{X}$,

$$\pi_n(i) \to \pi(i), \text{ as } n \to \infty.$$

Proof See EE126, or MC Note. □

Example

Convergence of π_n

Theorem Let X_n be an irreducible and aperiodic Markov chain with invariant distribution π. Then, for all $i \in \mathscr{X}$,
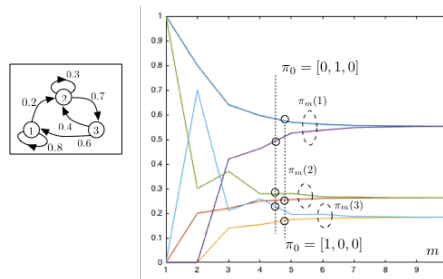
$$\pi_n(i) \to \pi(i), \text{ as } n \to \infty.$$

Proof See EE126, or MC Note. □

Example

Poll

X_n is a MC with the transition diagram as shown. Fix $b = 1$, and let a have any value such that $0 < a < 1$. Select all true statements.

- ▶ MC is always irreducible.
- ▶ MC is always aperiodic.
- ▶ MC has a unique invariant distribution π.
- ▶ The long-term fraction of time spent in each state converges to the invariant probability for that state.
- ▶ π_n always converges to the invariant distribution π.

Calculating π

Let P be irreducible. How do we find π?

Example: $P = \begin{bmatrix} 0.8 & 0.2 & 0 \\ 0 & 0.3 & 0.7 \\ 0.6 & 0.4 & 0 \end{bmatrix}$.

One has $\pi P = \pi$, i.e., $\pi[P - I] = \mathbf{0}$ where I is the identity matrix:

$$\pi \begin{bmatrix} 0.8-1 & 0.2 & 0 \\ 0 & 0.3-1 & 0.7 \\ 0.6 & 0.4 & 0-1 \end{bmatrix} = [0,0,0].$$

However, the sum of the columns of $P - I$ is $\mathbf{0}$. This shows that these equations are redundant: If all but the last one hold, so does the last one. Let us replace the last equation by $\pi\mathbf{1} = 1$, i.e., $\sum_j \pi(j) = 1$:

$$\pi \begin{bmatrix} 0.8-1 & 0.2 & 1 \\ 0 & 0.3-1 & 1 \\ 0.6 & 0.4 & 1 \end{bmatrix} = [0,0,1].$$

Hence,

$$\pi = [0,0,1] \begin{bmatrix} 0.8-1 & 0.2 & 1 \\ 0 & 0.3-1 & 1 \\ 0.6 & 0.4 & 1 \end{bmatrix}^{-1} \approx [0.55, 0.26, 0.19]$$

Interesting Example: How to Gamble, if You Must

Dubins and Savage, How to Gamble if You Must: Inequalities for Stochastic Processes. Dover Books on Mathematics. Paperback - July 23, 2014. (Original Edition, 1965.)

Recall the 'heads or tails game':
At each step, you win 1 w.p. p and loose 1 w.p. $q = 1 - p$.
You start with 10 and you want to maximize the probability of getting to 100 before you get to 0.

In their celebrated masterpiece, Dubins and Savage proved that the optimal strategy, if $p \leq 1/2$, is the bold one, always betting the maximum, and if $p \geq 1/2$, then an optimal strategy is the timid one, always betting the minimum.

There are relatively few problems for which one can prove such a clean result. However, there is a systematic approach to calculate the optimal strategy for many problems. We explain that approach next on this problem.

Original Strategy

Recall the original strategy: bet 1 each time. Then,

Let $\alpha(n)$ be the probability of reaching 100 before 0, starting from n, for $n = 0, 1, \ldots, 100$.

$$\alpha(0) = 0; \alpha(100) = 1.$$
$$\alpha(n) = p\alpha(n+1) + q\alpha(n-1), 0 < n < 100.$$

Solving, we find

$$\alpha(n) = \frac{1-\rho^n}{1-\rho^{100}} \text{ with } \rho = qp^{-1}.$$

For $p = 0.46$, we get $\alpha(10) \approx 3.5 \times 10^{-6}$.

Bold: Estimate

We can do better. Let us bet all we have. Then, with probability p^4 we have

$$10 \to 20 \to 40 \to 80 \to 160.$$

With $p = 0.46$, we see that we get to 100, at least with probability $(0.46)^4 = 0.0448$. This is much better than 3.5×10^{-6}.

Thus, the probability of winning the game (i.e., getting to 100 before 0) is at least 0.0448 when playing bold.

Bold: Analysis

What is the exact probability of winning when playing bold? Here is the corresponding MC:

The FSE for $\alpha(n) = Pr[T_{100} < T_0 \mid X_0 = n]$ are

$$\alpha(10) = p\alpha(20) + q0; \alpha(20) = p\alpha(40) + q0; \alpha(40) = p\alpha(80) + q0$$
$$\alpha(80) = p1 + q\alpha(60); \alpha(60) = p1 + q\alpha(20)$$

To solve, let $\alpha(10) = x$. Then, we find

$$\alpha(20) = p^{-1}x; \alpha(40) = p^{-1}\alpha(20) = p^{-2}x$$
$$\alpha(80) = p^{-1}\alpha(40) = p^{-3}x; p^{-3}x = p + q\alpha(60); \alpha(60) = p + qp^{-1}x.$$

We solve the last two equations for x.
We find $x = p^2(1+q)/(p^{-2} - q^2) \approx 0.0735$.

Summary

Markov Chains

▶ Markov Chain: $Pr[X_{n+1} = j | X_0, \ldots, X_n = i] = P(i,j)$
▶ FSE: $\beta(i) = 1 + \sum_j P(i,j)\beta(j); \alpha(i) = \sum_j P(i,j)\alpha(j)$.
▶ $\pi_n = \pi_0 P^n$
▶ π is invariant iff $\pi P = \pi$
▶ Irreducible $\Rightarrow$ one and only one invariant distribution π
▶ Irreducible $\Rightarrow$ fraction of time in state i approaches $\pi(i)$
▶ Irreducible + Aperiodic $\Rightarrow \pi_n \to \pi$.
▶ Calculating π: One finds $\pi = [0, 0, \ldots, 1]Q^{-1}$ where $Q = \cdots$.