

Markov Chains II

CS70 Summer 2016 - Lecture 6C

Grace Dinh

27 July 2016

UC Berkeley


Agenda

Classification of MC states

Aperiodicity, irreducibility, ergodicity

Convergence, limiting and stationary distributions

Reference for this lecture: Ch. 7 of
Mitzenmacher and Upfal, "Probability and
Computing"



Markov Chain Properties

Accessibility and Communication


State i is accessible from j if there is some chance that, if I'm at j at some timestep, I'll end up at state i some time later.

Formally: State i is accessible from state j if there exists $n \geq 0$ such that $(P^n)_{i,j} > 0$.

If j is accessible from i and i is accessible from j , then they are said to “communicate”.

Another way of looking at it: directed connectivity. i communicates with j : exists path from i to j in the graph corresponding to the chain.

Accessibility and Communication: Example




Is 1 accessible from 2? **No**. Is 2 accessible from 1? **Yes**. Do 1 and 2 communicate? **No**.

Is 2 accessible from 3? **Yes**. Is 3 accessible from 2? **Yes**. Do 1 and 2 communicate? **Yes**.


Irreducibility

Irreducible Markov chain: every state communicates with every other state.

Or: graph representation is strongly connected.



Irreducible.




Not irreducible.

Recurrent States

Let's say we're at a state i . Do we ever return to it again?

Let $r_{i,j}^t$ denote the probability that we first hit state j in t steps, starting from state i . A state is **recurrent** if $\sum_t r_{i,i}^t = 1$ and **transient** otherwise.



Is state 1 recurrent? **No!**

A Theorem

Suppose we are dealing with a **finite** MC. Then:


- **There is at least one recurrent state.**
- For any recurrent state i , the expected hitting time $h_{i,i}$ if we start from i is finite.

Proof: (first part) Consider a non-recurrent state. If we start at that timestep, there is a nonzero probability that we will never see it again.

Then if we start from that state and do an infinite number of timesteps, the probability that we see that state infinitely many times is zero.

Start anywhere on the MC and do an infinite number of timesteps. Since the MC is finite, some step must appear infinitely many times. So, that step must be recurrent.

Aperiodicity



Intuition: Suppose we're in one of these states at some timestep. Then we can never return to it an odd number of timesteps later.

To capture this intuition: state j is periodic if there exists some integer $\Delta > 1$ such that $P_{j,j}^s = \Pr[X_{t+s} = j | X_t = j] = 0$ unless Δ divides s .

A Markov chain is said to be periodic if any of its states is periodic.

Opposite of periodic: **aperiodic**.

Aperiodicity of Irreducible Chains - Another Definition

Theorem: Assume that the MC is irreducible. Then

$$d(j) := \text{g.c.d.}\{s > 0 \mid P_{j,j}^s > 0\}$$

has the same value for all states i .

Proof: See Lecture note 18.

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

Are the definitions the same? Yes.

If gcd of all the timesteps where $P_{j,j}^s$ is nonzero is greater than 1... On timesteps s that are not multiples of $d(j)$, $P_{j,j}^s$ is zero.

¹gcd = greatest common divisor.



An aperiodic state that is recurrent is called **ergodic**. A Markov chain is said to be ergodic if all its states are ergodic.


“Ludwig Boltzmann needed a word to express the idea that if you took an isolated system at constant energy and let it run, any one trajectory, continued long enough, would be representative of the system as a whole. Being a highly-educated nineteenth century German-speaker, Boltzmann knew far too much ancient Greek, so he called this the “ergodic property”, from *ergon* “energy, work” and *hodos* “way, path.” The name stuck.” (*Advanced Data Analysis from an Elementary Point of View* by Shalizi, pg. 479)

Theorem: A finite, irreducible, aperiodic Markov chain is ergodic.

Stationary and Limiting Distributions

Stationary Distributions: Motivation

Consider the driving exam MC again.



Once we pass the test (state 4), we're done forever. We never leave state 4.

If our distribution is $[0 \ 0 \ 0 \ 1]$: distribution is unchanged over a timestep.

Stationary Distributions: Motivation II

Or how about the two-cycle?



What if our distribution is $[0.5 \ 0.5]$? Does it change with timesteps?

No!

Definition: Stationary Distribution


A distribution π over states in a Markov chain is said to be a *stationary distribution* (a.k.a. an *invariant or equilibrium distribution*) if $\pi = \pi P$.

Basically: not affected by timesteps. If we have this distribution, we have it forever.

Another way of looking at it: π is an eigenvector of P : If we multiply π by P , we get a multiple of π (actually, π itself). Consequence: stochastic matrix always has 1 as an eigenvalue!

To find stationary distribution: solve $\pi P = \pi$ ("balance equations")

An Example




$$\begin{aligned}\pi P &= \pi \Leftrightarrow [\pi_1, \pi_2] \begin{bmatrix} 1-a & a \\ b & 1-b \end{bmatrix} = [\pi_1, \pi_2] \\ &\Leftrightarrow \pi(1)(1-a) + \pi_2 b = \pi_1 \text{ and } \pi_1 a + \pi_2(1-b) = \pi_2 \\ &\Leftrightarrow \pi_1 a = \pi_2 b.\end{aligned}$$

These equations are redundant! Add equation: $\pi_1 + \pi_2 = 1$.
Solves to:

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

Another Example



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

$$\pi P = \pi = [\pi_1, \pi_2] \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = [\pi_1, \pi_2]$$

So:

$$\pi_1 = \pi_1 \text{ and } \pi_2 = \pi_2.$$

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

Main Theorem

Suppose we are given a finite, irreducible, aperiodic Markov chain.
Then:

- There is a unique stationary distribution π .
- For all j, i , the limit $\lim_{t \rightarrow \infty} P_{j,i}^t$ exists and is independent of j .
- $\pi_i = \lim_{t \rightarrow \infty} P_{j,i}^t = 1/h_{i,i}$

Proof: really long and messy, see note 18 or Ch. 7 of MU. (we won't expect you to know this).

Connections between Linear Algebra and Markov Chains

It turns out that the convergence of the limiting distribution to the stationary distribution corresponds to a nice result from linear algebra: if you multiply a random vector by a matrix a lot of times, the result will converge towards an eigenvector (specifically, one corresponding to the highest eigenvalue) w.h.p.

Perron-Frobenius: positive elements \rightarrow single highest eigenvalue (1, here), i.e. one with a unique eigenvector (up to constant factors).

(No, you do not need to know this for the midterms and the homeworks).

The Gambler's Ruin

Suppose you play a game with your friend. Flip a fair coin. Heads: you win a dollar. Tails: you lose a dollar. Repeat.


You win when you get all your friend's money. You lose when your friend gets all of yours.

What is the probability that you win?

If you and your friend have same amount of money: $1/2$ by symmetry.

What if you and your friend are willing to bet different amounts?

The Gambler's Ruin II




Suppose you have l_1 dollars and your friend has l_2 . Express as above Markov chain.

States $-l_1, l_2$ are recurrent; all others are transient. What is the probability that you win (i.e. you hit state l_2 before l_1)?

Let P_i^t be the probability that you're at state i after t timesteps. What's $\lim_{t \rightarrow \infty} P_i^t$ for $i \in [-l_1 + 1, l_2 - 1]$? 0 (since they are transient states).

Want to find: $q := \lim_{t \rightarrow \infty} P_{l_2}^t$: probability that you win (state is absorbed into l_2).

The Gambler's Ruin III



Denote amount of money you have at timestep t as W_t .

What's the expected amount of money you have after a single step?

0. What's the expected gain after t steps, $E[W_t]$? 0, by induction.

So:

$$E[W_t] = \sum_{i \in [-l_1, l_2]} i P_i^t = 0$$

$$\lim_{t \rightarrow \infty} E[W_t] = l_2 q - l_1 (1 - q) = 0$$

Solve: $q = l_1 / (l_1 + l_2)$. The more money you're willing to bet, the more you win!

Random Walks

Motivation

Suppose I give you a connected graph and you walk around on it randomly.

At each vertex you pick a random edge (with uniform probability) to traverse. Probability of choosing a particular edge from vertex i : $1/d(i)$ where $d(i)$ is the degree of i .

This is a Markov chain!

Is it irreducible? Yes, if it's connected.

Aperiodicity of Random Walks

Theorem: A random walk on an undirected, connected graph is aperiodic if and only if the graph is not bipartite.

Proof: Suppose graph is bipartite. Then if I start at a node I can never go back in an odd number of timesteps. So random walk is periodic.

Conversely, suppose graph is not bipartite. Then there's an odd cycle (lecture 6). So we have a path of odd length from any node to itself.

Then there exists an n' such that for all $n \geq n'$, I can go from my start node back to itself in n timesteps. Why?

If n is even: just go to the next node and back $n/2$ times.

If n is odd: Go to some node in cycle (graph is connected). Traverse cycle. Go back. Going to node and back takes even number of timesteps. Traversing cycle takes odd number of timesteps. Total number of timesteps: odd.

So random walk is periodic.

Stationary Distribution of Random Walks

Theorem: A random walk on a graph G converges to a stationary distribution π where $\pi_v = \frac{d(v)}{2|E|}$.

Proof: Is this a distribution at all? $\sum_v d(v) = 2|E|$ so $\sum_v \pi_v = \sum_v d(v)/(2|E|) = 1$. It's a distribution.

Why is it stationary? Let $N(v)$ represent the neighbors of v . Want to show: $\pi = \pi P$. Equivalently:

$$\pi_v = \sum_{u \in N(v)} \frac{d(u)}{2|E|} \frac{1}{d(u)} = \frac{d(v)}{2|E|}$$

.

So π solves the balance equations, so it's stationary.

Immediately follows that for any u , $h_{u,u} = 2 |E| / d(u)$.

Lemma: If $(u, v) \in E$, then $h_{u,v} < 2 |E|$.

Proof:

$$\frac{2 |E|}{d(u)} = h_{u,u} = \frac{1}{d(u)} \sum_{w \in N(u)} (1 + h_{w,u})$$

Cancel:

$$2 |E| = h_{u,u} = \sum_{w \in N(u)} (1 + h_{w,u})$$

Since $v \in N(u)$: $h_{v,u} < 2 |E|$

Cover Time II

Say I start from some vertex and do a random walk. How long does it take me to touch every single node in the graph? **Cover time**: the longest such time (for any starting vertex).

Theorem: Cover time of $G = (V, E)$ is at most $4|V||E|$.

Proof: Choose spanning tree of G .

If we duplicate edges (one going in each direction): there's an Eulerian tour on this tree. Let the vertices traversed by the tour be $v_0, v_1, \dots, v_{2|V|-2} = v_0$.

Expected time to go through vertices in the tour in this order: upper bound on cover time.

Expected time to go from one vertex to the next is at most $2|E|$.

Number of trips we need to do? $2|V| - 2 < 2|V|$.

So: $4|E||V|$ is an upper bound on the cover time.

Application: PageRank

Idea: web search should give you results ordered in such a way that you're more likely to stumble on the top result than the lower results when browsing the web.

Assume you click links on webpages randomly forever. How often are you going to run into a webpage?

Model with a random walk on a **directed** graph! At each webpage: click random link.

Want to find the stationary distribution of this walk. Problem: graph isn't strongly connected.

Solution: with small probability, go to a random website instead of clicking a link.

MC is irreducible and aperiodic, so its limiting distribution must be the unique stationary distribution.

Find the limiting distribution by solving an eigenvalue problem!
(Math 128B, Math 221)

Gig: Random Text