EE 126 — Probability and Random Processes

March 01, 2018 — Discrete-Time Markov Chains II

1 Lecture

Recap

It is the case that $\pi_{n+1} = \pi_n P$. If $\pi = \pi P$, we have an invariant (/stationary) distribution. "If we start in the stationary distribution, we remain in the stationary distribution."

Big Theorem: if the Markov chain is finite and irreducible, then there exists a unique π^* which is an invariant distribution. If the Markov chain is also aperiodic, $\pi_n \xrightarrow{n\to\infty} \pi^*$.

Balance equations: $\pi = \pi P$ is equivalent to the idea that "flow into state i equals flow out of state i." We can use this to solve for the invariant distribution.

Classification of General Markov Chains

We define T_i as $\min(n \geq 1 \mid X_n = i)$. This is a random variable representing the time at which we first return to i, having started at state $X_0 = i$. If the Markov chain is irreducible, then

$$P(T_i < \infty \mid X_0 = i) \quad \begin{cases} = 1 & \text{means that state } i \text{ is recurrent} \\ < 1 & \text{means that state } i \text{ is transient} \end{cases}$$

If a Markov chain is recurrent, then $\mathbb{E}[T_i \mid X_0 = i] < \infty$ signifies **positive recurrence** and $\mathbb{E}[T_i \mid X_0 = i] = \infty$ signifies **null recurrence**. Also, if a Markov chain is irreducible, aperiodic, and positive recurrent, then $P(X_n = j \mid X_0 = i) \xrightarrow{\text{a.s.}} \pi^*(j)$ for all i, j in the state space. In other words, the Markov chain is asymptotically stationary.

For a Markov chain to be recurrent, every state must be recurrent.

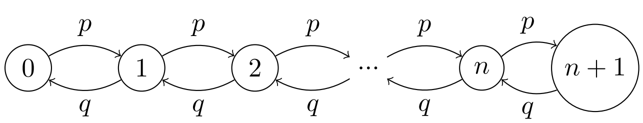

We have $p + q = 1$. If $p < 0.5$, the chain is positive recurrent (there is a pull toward 0). If $p > 0.5$, the chain is transient. If $p = 0.5$, it is null recurrent.

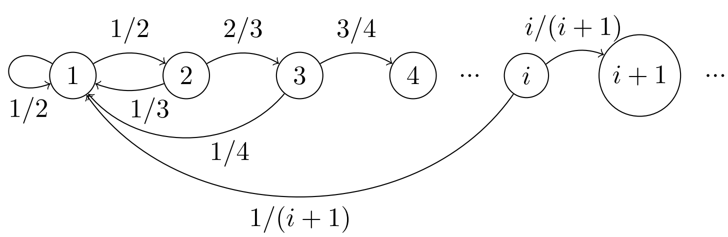

We have $\mathbb{E}[T_1 \mid X_0 = 1] = (1)\left(\frac{1}{2}\right) + (2)\left(\frac{1}{2}\right)\left(\frac{1}{3}\right) + (3)\left(\frac{1}{2}\right)\left(\frac{2}{3}\right)\left(\frac{1}{4}\right) + \ \dots\ + (i)\left(\frac{1}{i(i+1)}\right) + \dots$. This is equivalent to $\sum_{i=1}^{\infty}(i+1)^{-1} \to \infty$. Therefore, the Markov chain is null recurrent (which, once again, says "we'll come back but it'll take forever").

Big Theorem

Markov chains can be either *irreducible* or *reducible*. If a Markov chain is irreducible, it can be either *transient*, *positive recurrent*, or *null recurrent*. (The classification – transient, positive recurrent, or null recurrent – is required for all states.) If a chain is positive recurrent, it can be either *aperiodic* or *periodic*. If the chain is aperiodic, there is convergence to a unique invariant distribution: $\pi \to \pi^*$. If the chain is transient or null recurrent, there is no π^*.

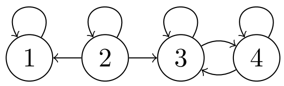

A reducible Markov chain. States 1, 3, and 4 are recurrent. State 2 is transient.

Hitting Times: First-Step Equations

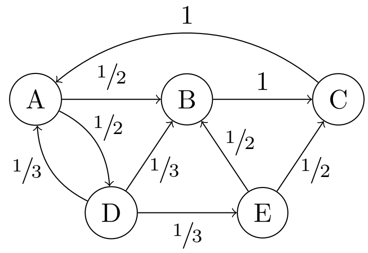

Example: PageRank. Starting at state A, how many steps does it take to reach state E? This time is called the **hitting time** (also known as the first passage time, here the time of first passage into E). We denote it as T_E.

The **mean hitting time**, denoted $\beta_E(A)$, is $\mathbb{E}[T_E \mid X_0 = A]$. We start at A, and we want to know on average how long it takes to hit E. The key to determining $\beta_E(A)$ is to realize that it is related to $\beta_E(i)$ for $i = B, C, D, E$.

$$\begin{aligned}
\beta_E(A) &= 1 + \frac{1}{2}\beta_E(B) + \frac{1}{2}\beta_E(D) \\
\beta_E(B) &= 1 + \beta_E(C) \\
\beta_E(C) &= 1 + \beta_E(A) \\
\beta_E(D) &= 1 + \frac{1}{3}\beta_E(A) + \frac{1}{3}\beta_E(B) + \frac{1}{3}\beta_E(E) \\
\beta_E(E) &= 0
\end{aligned}$$

These are called the **first-step equations**. Solving them gives $\beta_E(A) = 17$, $\beta_E(B) = 19$, $\beta_E(C) = 18$, $\beta_E(D) = 13$, and $\beta_E(E) = 0$.

Example. We flip a fair coin until we get three heads in a row. *How many flips do we perform on average?* In the Markov formulation, a state will correspond to the number of observed consecutive heads. We will have either zero, one, two, or three consecutive heads. The state transition diagram looks like this:

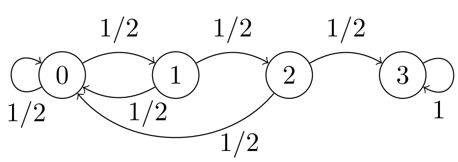

The first step equations are

$$\begin{aligned}
\beta_3(0) &= 1 + \frac{1}{2}\beta_3(0) + \frac{1}{2}\beta_3(1) \\
\beta_3(1) &= 1 + \frac{1}{2}\beta_3(0) + \frac{1}{2}\beta_3(2) \\
\beta_3(2) &= 1 + \frac{1}{2}\beta_3(0) + \frac{1}{2}\beta_3(3) \\
\beta_3(3) &= 0
\end{aligned}$$

and the solution shows that $\beta_3(0) = 14$ flips.