ECE 6960: Adv. Random Processes & Applications
Lecture Notes, Fall 2010

Lecture 17

Today: (1) State Classification, (2) Limiting probabilities

- Read Bianchi paper for Tuesday's class.

0.1 Transient and Recurrent

Def'n: *Absorbing*
A state is *absorbing* if no other state is accessible from it.

Consider a 5-state MC with states $\{0, 1, 2, 3, 4\}$ and TPM,

$$P = \begin{bmatrix} 0.5 & 0.5 & 0 & 0 & 0 \\ 0.5 & 0.5 & 0 & 0 & 0 \\ 0.25 & 0.25 & 0.25 & 0.25 & 0 \\ 0 & 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 0 & 0.5 & 0.5 \end{bmatrix}$$

Now, are there any absorbing states, and if so, which ones?

In either case, we understand from looking at this example MC that the states $\{0, 1\}$ and states $\{3, 4\}$ are different from state 2. State 2 can occur for a finite period of time, but as time $n \to \infty$, the MC will certainly (with probability 1) end up in states 0 and 1, or 3 and 4, and not return to state 2.

Let f_i be the probability that, starting in state i, the process will ever reenter state i. That is,

$$f_i = P\left[\text{ever returning to state } i\right]$$

Def'n: *Transient*
A state is transient if $f_i < 1$.

We can see that state 2 in the above example is transient, because with probability 0.25, it enters state 3, and then never reenters state 2 after that. With probability 0.5, it enters state 0 or 1, and then never re-enters state 3. So, $f_i = 0.25$.

Def'n: *Recurrent*
A state is recurrent if $f_i = 1$.

We can see that state 3 (and 4) are both recurrent. Assume the process is in state 3 at time 0. At $n = 1$, there is a 0.5 probability of reentering state 3. Otherwise, at time 2, there is a $(0.5)^2$ chance of reentering. Continuing this argument,

$$f_i = \sum_{i=1}^{\infty} 0.5^i = \frac{0.5}{1-0.5} = 1.$$

If state i is recurrent, starting in state i, there is probability one of reentering state i. Once it is back in state i, it will return there again with probability 1. Clearly, it will reenter there infinitely many times.

If state i is transient, starting in state i, there is only probability $f_i < 1$ of reentering state i. That means, there is only a probability f_i^2 that it reenters state i at least twice. Alternatively, there is a $f_i(1-f_i)$ chance that it returns once and then leaves permanently.

In general, the number of returns to state i (when starting at state i) is a geometric random variable with parameter f_i. Note that a geometric r.v. has finite mean. Thus:

- The number of returns to state i for a recurrent state is infinite.
- The number of returns to state i for a transient state is finite.

Theorem: State i is recurrent iff

$$\sum_{n=1}^{\infty} p_{i,i}(n) = \infty.$$

State i is transient iff

$$\sum_{n=1}^{\infty} p_{i,i}(n) < \infty.$$

Proof:
Let I_n be defined as

$$I_n = \begin{cases} 1, & X_n = i \\ 0, & o.w. \end{cases}$$

Then

$$\sum_{n=1}^{\infty} I_n \mid X_0 = i$$

is the total number of returns to state i. We already have that the expected number of returns is infinite for recurrent states and finite for transient states. Converting into a function of $p_{i,i}(n)$,

$$\begin{aligned} E\left[\sum_{n=1}^{\infty} I_n \mid X_0 = i\right] &= \sum_{n=1}^{\infty} E\left[I_n \mid X_0 = i\right] \\ &= \sum_{n=1}^{\infty} P\left[X_n = i | X_0 = i\right] \\ &= \sum_{n=1}^{\infty} p_{i,i}(n) \end{aligned}$$

Notes:

- Transitive property for recurrence: If state i is recurrent, and $i \leftrightarrow j$, then state j is recurrent. (See proof in Leon-Garcia page 664).
- Thus recurrence is a class property, shared by all states in a class.
- Any finite-state MC cannot have all transient states.
- A finite-state irreducible MC must have all recurrent states.

Example: Ross, 4.14
Consider the MC having states $\{0, 1, 2, 3, 4\}$ with TPM

$$P = \begin{bmatrix} 0.5 & 0.5 & 0 & 0 & 0 \\ 0.5 & 0.5 & 0 & 0 & 0 \\ 0 & 0 & 0.5 & 0.5 & 0 \\ 0 & 0 & 0.5 & 0.5 & 0 \\ 0.25 & 0.25 & 0 & 0 & 0.5 \end{bmatrix}$$

What are the classes? What are the recurrent states and transient states?

Example: Random Walk
Consider the MC in Figure 1. A r.p. increments state by one with probability p, and decrements the state by one with probability $1 - p$. This can be seen as the process of gambling, if at each time,

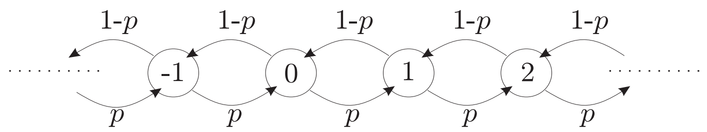

Figure 1: A random walk on an infinite MC.

you either lose or win one coin.

Are the states of this MC recurrent or transient?

Notes:

- All states communicate with each other, so there is one class.
- Thus all states are the same; either they are all transient, or they are all recurrent.
- The MC is irreducible. But, the prior note ("A finite-state irreducible MC must have all recurrent states") does not apply since this is an infinite-state MC.

Solution: Since all states are the same, pick state $i = 0$. We need to find $\sum_{n=1}^{\infty} p_{0,0}(n)$ and see whether it is infinite or finite.

Note that $p_{0,0}(n)$ for n odd is zero. But $p_{0,0}(2n)$, $n = 1, 2, \ldots$, is positive, and it is the probability that out of $2n$ steps, n were to the left, and n were to the right. This is a binomial r.v.,

$$p_{0,0}(2n) = \binom{2n}{n} p^n (1-p)^n = \frac{(2n)!}{n!n!}[p(1-p)]^n$$

We can use Stirling's approximation,

$$n! \sim \sqrt{2\pi} n^{n+1/2} e^{-n}$$

This is exact in the limit as $n \to \infty$, specifically, $a_n \sim b_n$ when $\lim_{n\to\infty} a_n/b_n = 1$. Such an approximation is fine since we only want to know the limiting behavior.

$$\begin{aligned} p_{0,0}(2n) &\sim \frac{\sqrt{2\pi}(2n)^{2n+1/2}e^{-2n}}{2\pi n^{2n+1}e^{-2n}}[p(1-p)]^n \\ &\sim \frac{2^{2n+1/2}}{\sqrt{2\pi}n^{1/2}}[p(1-p)]^n \\ &\sim \frac{[4p(1-p)]^n}{\sqrt{\pi n}} \end{aligned}$$

We can verify that for positive a_n, b_n that if $a_n \sim b_n$ then

$$\sum_{n=1}^{\infty} a_n < \infty \text{ iff } \sum_{n=1}^{\infty} b_n$$

This is due to the definition of the limit applied to the sequence a_n/b_n.

From (1), we have then that $\sum_{n=1}^{\infty} p_{0,0}(n)$ is finite if and only if

$$\sum_{n=1}^{\infty} \frac{[4p(1-p)]^n}{\sqrt{\pi n}}$$

is finite.

Recall that $\sum_n a/\sqrt{n} = \infty$. Thus if $p = 1/2$, and the numerator is $[4p(1-p)]^n = 1$, the sum does not converge. But for $p \neq 1/2$, the numerator is $[4p(1-p)]^n = c^n$ for some constant $0 \leq c < 1$. So,

$$\sum_{n=1}^{\infty} \frac{[4p(1-p)]^n}{\sqrt{\pi n}} < \sum_{n=1}^{\infty} c^n < \infty$$

Thus:

- $p = 1/2$: All states are recurrent. (This is called a symmetric random walk).
- $p \neq 1/2$: All states are transient.

0.2 Periodicity

Def'n: *Period*
A state i in a MC X_n has period d,

$$d = \gcd\{n : P\left[X_n = i | X_0 = i\right] > 0\}$$

where gcd indicates the greatest common divisor.

For example, if

$$\{n : P\left[X_n = i | X_0 = i\right] > 0\} = \{2, 4, 6, 8, \ldots\}$$

Then the period of state i would be 2.

Def'n: *Periodic*
A state i in a MC X_n is periodic if its period $d \geq 2$.

In other words, a period of 1 (the state can reoccur every time) is not called "periodic". Periodicity is a class property, so every state in the same class has the same periodicity.

Def'n: *Aperiodic*
A MC is aperiodic if it contains no periodic states.

A MC is called periodic if if has *any* periodic state.

0.3 Limiting Probabilities

This is in Leon-Garcia 11.3.3.
Let's consider what happens in a MC as $n \to \infty$.

- If the MC has some transient classes and some recurrent classes, then eventually the process enters one of the recurrent classes and remains in that class thereafter.
- Assume we are at recurrent state i at time 0. Let $T_i(k)$ be the number of steps between re-occurrence of state i (the time that elapses between the $k-1$st and kth return). Then the proportion of time spent in state i approaches $1/E\left[T_i\right]$. Let's define π_i as the long-term proportion of time spent in state i,

$$\pi_i = 1/E\left[T_i\right]$$

<u>Def'n:</u> *Positive/Null Reccurrent*
A state i is *positive reccurrent* if $E[T_i] < \infty$. A state i is it null recurrent if $E[T_i] = \infty$.

- Equivalently, a positive recurrent state has $\pi_i > 0$, while a null recurrent state has $\pi_i = 0$.
- Positive / Null recurrence are class properties; all states in the same class are one or the other (or neither if they are transient).

<u>Def'n:</u> *Ergodic MC*
A Markov Chain is an *ergodic* M.C. if it is irreducible, aperiodic, and positive recurrent.

An ergodic MC will revisit states sufficiently frequently that the long-term proportion of time spent in state i will be described by $\pi_i = 1/E[T_i]$ for all states i. Thus

$$\pi_j = \sum_i \pi_i p_{i,j}(1)$$

or in vector/matrix notation,

$$\boldsymbol{\pi}^T = \boldsymbol{\pi}^T P(1)$$

Recall that the state probabilities at time n can be found as

$$\mathbf{p}(n)^T = \mathbf{p}(0)^T [P(1)]^n$$

<u>Theorem:</u> For an ergodic MC, the limit

$$\boldsymbol{\pi} = \lim_{n\to\infty} \mathbf{p}(n)$$

exists, where $\boldsymbol{\pi} = [\pi_1, \ldots, \pi_n]^T$, and satisfies

$$\begin{aligned} \boldsymbol{\pi}^T &= \boldsymbol{\pi}^T P(1) \\ \boldsymbol{\pi} &= P(1)^T \boldsymbol{\pi} \end{aligned}$$

and

$$\sum_i \pi_i = 1.$$

Proof: Not covered.

Equation (1) expresses $\boldsymbol{\pi}$ as an eigen-vector of $P(1)^T$, in particular, an eigen-vector of $P(1)^T$ with eigenvalue equal to 1. In Matlab, use

```
[V, L] = eig(P')
```

which returns eigenvectors in the columns of `V`, and eigenvalues in the diagonal of `L`. The `i`th column of `V` will have `L(i,i)==1`, call that eigenvector $\mathbf{v} = [v_1, \ldots, v_n]^T$. Then

$$\boldsymbol{\pi} = \frac{\mathbf{v}}{\sum_i v_i}.$$

0.4 Examples

Example: Google PageRank
Here's something we use every day: Google's PageRank. Google's PageRank uses the limiting probability of a Markov chain model to figure out the "importance" of all web sites. The Markov chain models each web page as a state, and each link on the page as a potential transition. The transition probabilities are assumed to be uniform to each link. That is, a "robot" would pick randomly from the links in order to generate the next page to surf to. A page is more important if, in the long run, this robot would visit that page more often.

One problem is that there are some pages with no outgoing links. They would be absorbing states. In any case, we would not have an ergodic MC. So, we add in a small probability that the robot transitions randomly to *any* web page on the internet, with equal probability. This makes the MC ergodic.

Another problem is the large size of the transition matrix. Finding the eigenvectors, in general, is an $\mathcal{O}\left(N^3\right)$ algorithm, in general. But we have two advantages – we only need one eigenvector, and the matrix is sparse (not many links exist on a page compared to the number of web pages). Fortunately, there are fast algorithms for this case. Also, updates don't need to occur every instant – Google can update pagerank when it is recalculated.

Example: Epithelial cell division

Cell biologists had wondered why epithelial cells (cells that grow in sheets) are typically 6-sided, but also often 5 or 7-sided. Their assumption was that cells sorted that way for optimal packing. But new work [1] has shown that it is simply a result of randomness, which can be modeled with a Markov chain. The model, shown at left in Figure 2, says that when a cell divides, it picks two points around the cell wall uniformly, and then a new wall forms between those two points, dividing the cell. The number of vertices in the two daughter cells sum to the original number of vertices in the parent cell plus 2 (the new vertices created). Those two new vertices also increase by one the number of sides of two neighboring cells. The results of these two forces; increasing the number of vertices by one, and dividing the number of vertices among two daughter

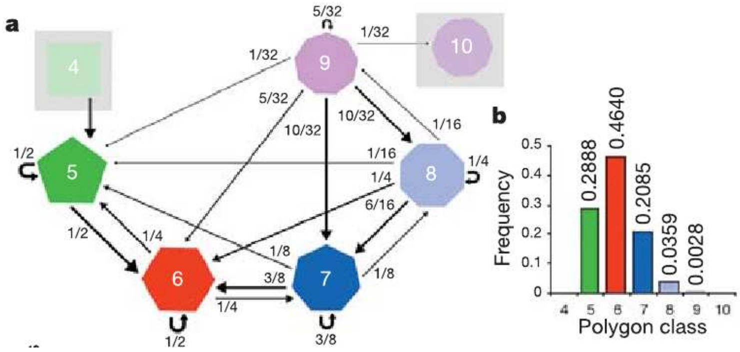

Figure 2: From [1]: at left, the number of sides of an epithelial cell, with transition probabilities due to cell division. At right, the limiting probabilities of each state.

cells, results in an ergodic MC. The stationary distribution, shown in the right of Figure 2, agrees with the distibution of sides from experimental tissue in many different subtypes of epithelial tissue, within a couple percent.

References

[1] M. Gibson, A. Patel, R. Nagpal, and N. Perrimon. The emergence of geometric order in proliferating metazoan epithelia. *Nature*, 442(7106):1038–1041, 2006.