

IEOR 6711: Stochastic Models I

Fall 2013, Professor Whitt

Class Lecture Notes: Thursday, October 24.

Introduction to Markov Chains: The Big Picture

1 Basic Structure

A stochastic process is a **Markov process** if the conditional probability of a future event given present and past states depends only on the present state. Let $\{X_n : n \geq 0\}$ be a **discrete-time Markov chain (DTMC)**, with “chain” meaning it takes values in a discrete state space, which we may take to be the integers. Let

$$P_{i,j} \equiv P(X_{k+1} = j | X_k = i)$$

be a **transition probability**, which we take to be independent of the initial time k . (We assume that the DTMC has stationary transition probabilities.)

Let the n -step transition probabilities be denoted by

$$P_{i,j}^{(n)} \equiv P(X_{k+n} = j | X_k = i).$$

For a finite-state DTMC, the model can be taken to be specified by the finite transition matrix $P \equiv (P_{i,j})$ (plus an initial probability vector). Then the associated matrix of n -step transition probabilities is just the n -fold product of the matrix P with itself; i.e., we have the important link to linear algebra

$$P^{(n)} = P^n, \quad n \geq 1,$$

which is based on the property

$$P_{i,j}^{(2)} = \sum_k P_{i,k} P_{k,j}.$$

1.1 Reaching

We say that the DTMC can **reach** state j from state i or that state j is **accessible** from state i , which we denote by $i \rightsquigarrow j$, if

$$P_{i,j}^{(n)} > 0 \quad \text{for some } n \geq 0.$$

Note that we do not require that this take place in a single step. By convention, i is always accessible from i , because that can occur in 0 steps. That is, by convention, $P_{i,i}^{(0)} \equiv 1 > 0$.

We say that the states i and j **communicate** if both $i \rightsquigarrow j$ and $j \rightsquigarrow i$. Communication is an **equivalence relation**, and we write $i \leftrightarrow j$ if i and j communicate (Prop 4.2.1), e.g., see Wikipedia. As such, it induces a partition of the set of states into **equivalence classes**. These are classified as **open** or **closed** according to their returning properties; see below. (A class is closed if the chain cannot leave the class; otherwise it is open.) The DTMC is said to be **irreducible** if there is a single communication equivalence class; otherwise it is said to be **reducible**. It is often convenient to put a reducible DTMC into **canonical form**; see below.

1.2 Returning

A state is said to be **recurrent** if the DTMC returns to that state with probability 1. A state is recurrent if and only if the expected number of returns to that state is infinite, Prop 4.2.3. A state that is not recurrent is said to be **transient**. The expected number of returns to a transient state is necessarily finite. We will show how to compute that expected value.

A recurrent state is said to be **positive recurrent** if the expected time to return (starting there) is finite; otherwise the (recurrent) state is said to be **null recurrent**. All recurrent states in finite-state DTMC's are positive recurrent. An example of a null recurrent DTMC is a simple random walk on the integers, with probability 1/2 of going up and probability 1/2 of going down. Recurrence becomes an interesting issue for simple random walks on the integer lattice of \mathbb{R}^m .

A state i is said to have **period** d if $P_{i,i}^{(n)} = 0$ whenever n is not divisible by d , and if d is the largest such integer. If a state has period 1 it is said to be **aperiodic**.

1.3 Class Properties

The notions of periodicity, recurrence and positive recurrence (and thus transience) are all **class properties**. That is, if $i \leftrightarrow j$, then j has the same period as i ; see Prop 4.2.2. Moreover, one is recurrent (positive recurrent, or transient) if and only if the other is; see Cor 4.2.4 and Prop 4.3.2.

1.4 Canonical Form

It is often helpful to reorder the states of a reducible DTMC so that the structure is more clearly visible. We illustrate by example.

Find the canonical form of the following Markov chain transition matrix:

(a)

$$P = \begin{pmatrix} 0.1 & 0.0 & 0.0 & 0.9 & 0.0 \\ 0.0 & 0.4 & 0.0 & 0.0 & 0.6 \\ 0.3 & 0.3 & 0.0 & 0.4 & 0.0 \\ 0.3 & 0.0 & 0.0 & 0.7 & 0.0 \\ 0.0 & 0.7 & 0.0 & 0.0 & 0.3 \end{pmatrix}$$

(I label the states 1, 2, 3, 4, 5.)

Notice that the sets $\{1, 4\}$ and $\{2, 5\}$ are closed communicating classes containing recurrent states, while $\{3\}$ is an open communicating class containing a transient state.

So you should reorder the states according to the order: 1, 4, 2, 5, 3. The order 2, 5, 1, 4, 3 would be OK too, as would 5, 2, 4, 1, 3. We put the recurrent states first and the transient states last. We group the recurrent states together according to their communicating class. Using the first order - 1, 4, 2, 5, 3 - you get

$$P = \begin{pmatrix} 0.1 & 0.9 & 0.0 & 0.0 & 0.0 \\ 0.3 & 0.7 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.4 & 0.6 & 0.0 \\ 0.0 & 0.0 & 0.7 & 0.3 & 0.0 \\ 0.3 & 0.4 & 0.3 & 0.0 & 0.0 \end{pmatrix}$$

Notice that the canonical form here has the structure:

$$P = \begin{pmatrix} P_1 & 0 & 0 \\ 0 & P_2 & 0 \\ R_1 & R_2 & Q \end{pmatrix},$$

where P_1 and P_2 are 2×2 Markov chain transition matrices in their own right, whereas R_i is the one-step transition probabilities from the single transient state to the i^{th} closed set. In this case, $Q \equiv (0)$ is the 1×1 sub-matrix representing the transition probabilities among the transient states. Here there is only a single transient state and the transition probability from that state to itself is 0. The chain leaves that transient state immediately, never to return..

2 Irreducible Chains: Apply Renewal Theory

We now focus on an irreducible DTMC. Equivalently, we focus on one communication class within a general DTMC. The key results are given in two theorems of Ross.

Theorem 4.3.1. Existence of a Limiting Distribution. In that setting, we can apply renewal theory to conclude that the obvious limits are valid. We use the fact that successive visits to any fixed state constitute a delayed renewal process. Thus we apply Proposition 3.5.1 as well as the results in §3.3 and §3.4. All cases are covered by allowing the mean time between renewals to be infinite as well as finite. Note that the theorems in Chapter 3 also apply to the case in which the mean time is infinite as well as finite.

Theorem 4.3.3. Existence of a Unique Stationary Distribution. We can deduce that a strictly positive limit exists (allowing for adjustment of the statement to account for periodicity, if it is present) if and only if there exists a proper stationary distribution, i.e., if and only if the matrix equation

$$\pi = \pi P$$

has a positive solution whose elements sum to 1. There is at most one such probability vector solution. Stationary distributions are valid for periodic chains, whereas limiting distributions are not. We call π a **stationary distribution**, because the right side of the equation can be interpreted as describing the probability vector that results if we start with an initial probability vector π and make a single transition according to the probability matrix P . Thus πP is the probability vector giving the distribution of X_1 assuming that X_0 is distributed according to π . The equation says that if we have an initial distribution of π , then the distribution after one transition (and thus after any number of transitions) is again π ; so that π is indeed a “stationary probability vector.”

Theorem 4.3.1 establishes the existence of a limiting distribution (in various senses), while Theorem 4.3.3 establishes the existence of a unique stationary distribution. In this case, for aperiodic irreducible DTMC’s these two necessarily coincide, but in general these two notions

are different. When they coincide, we may regard both as corresponding to the notion of a **steady-state distribution**.

These two results give **two ways** to determine the steady-state behavior. We either can determine the mean time to go from state j to state j , $E[T_{j,j}]$, or we can solve $\pi = \pi P$. These two ways are connected by the simple relation

$$E[T_{j,j}] = 1/\pi_j.$$

The validity of this relation is easy to see via the elementary renewal theorem, applied to this case. Let $N_{i,j}(n)$ be the number of visits to state j , starting in state i , among the first n steps. Then

$$E[N_{i,j}(n)] = \sum_{k=1}^n P_{i,j}^{(k)}.$$

First, from the elementary renewal theorem, we have $E[N_{i,j}(n)]/n \rightarrow 1/E[T_{j,j}]$. Second, we can extend this limit to an arbitrary initial probability vector

$$n^{-1} \sum_{i=0}^{\infty} P(X_0 = i) E[N_{i,j}(n)] \rightarrow 1/E[T_{j,j}] \quad \text{as } n \rightarrow \infty.$$

Proof. To prove the assertion immediately above, first note that

$$\begin{aligned} & \left| \left(n^{-1} \sum_{i=0}^{\infty} P(X_0 = i) E[N_{i,j}(n)] \right) - 1/E[T_{j,j}] \right| = \left| \sum_{i=0}^{\infty} P(X_0 = i) \left((E[N_{i,j}(n)]/n) - (1/E[T_{j,j}]) \right) \right| \\ & \leq \sum_{i=0}^{\infty} P(X_0 = i) \left| (E[N_{i,j}(n)]/n) - (1/E[T_{j,j}]) \right|, \end{aligned}$$

where $0 \leq E[N_{i,j}(n)]/n \leq 1$ and $0 \leq 1/E[T_{j,j}] \leq 1$, so that $|(E[N_{i,j}(n)]/n) - (1/E[T_{j,j}])| \leq 1$. For any $\epsilon > 0$, choose $m \equiv m(\epsilon)$ such that $\sum_{i=m+1}^{\infty} P(X_0 = i) < \epsilon/2$ and choose $n_0 \equiv n_0(\epsilon)$ such that $|E[N_{i,j}(n)]/n - 1/E[T_{j,j}]| < \epsilon/2$ for all i , $0 \leq i \leq m$ and all n , $n \geq n_0$. Then

$$\sum_{i=0}^{\infty} P(X_0 = i) \left| (E[N_{i,j}(n)]/n) - (1/E[T_{j,j}]) \right| \leq (1 \times (\epsilon/2) + (\epsilon/2) \times 1) = \epsilon.$$

Second, using the stationary distribution provided by Theorem 4.3.3, we can deduce that, if we use initial probability vector π ,

$$n^{-1} \sum_{i=0}^{\infty} P(X_0 = i) E[N_{i,j}(n)] = n^{-1} \sum_{i=0}^{\infty} \pi_i E[N_{i,j}(n)] = \pi_j.$$

Hence, indeed, we must have $\pi_j = 1/E[T_{j,j}]$.

3 Absorbing Chains: Apply Linear Algebra

We now consider a common case of reducible DTMC's in which all closed sets contain only one state. These special states are called **absorbing** states. The other states are the transient states. Basic linear algebra applies to the case P is finite.

For absorbing DTMC's the canonical form produces the transition matrix P with the **block matrix form**

$$P = \begin{pmatrix} I & 0 \\ R & Q \end{pmatrix},$$

where I is an identity matrix (1's on the diagonal and 0's elsewhere) and 0 (zero) is a matrix of zeros. The square submatrix Q is a transition matrix among the transient states, but some of its row sums are necessarily less than 1 (substochastic). This is achieved by putting the absorbing states first, and then the transient states.

Use the **fundamental matrix** $N = (I - Q)^{-1}$. See §4.4 and second part of Markov mouse notes. The expected total number of visits to transient state j starting in transient state i is $N_{i,j}$. The probability of being absorbed in absorbing state k starting in transient state i is $B_{i,k}$, where $B \equiv NR$.

We remark that there is a fundamental matrix for irreducible chains too, often denoted by Z , having the form $Z \equiv (I - P + A)^{-1}$, where A is a matrix all of whose rows are the stationary probability vector π , e.g., see p. 31 of Asmussen (2003).