Chapter 5: Markov Chains

it only matters where you are, not where you've been...

5.1 Introduction

A.A.Markov
1856-1922

So far, we have examined several stochastic processes using transition diagrams and First-Step Analysis.
The processes can be written as $\{X_0, X_1, X_2, \ldots\}$,
where X_t is the .

On the transition diagram, X_t corresponds to

In the Gambler's Ruin (Section **??**), X_t is the amount of money the gambler possesses after toss t. In the model for gene spread (Section **??**), X_t is the number of animals possessing the harmful allele A in generation t.

The processes that we have looked at via the transition diagram have a crucial property in common:

It does <u>not</u> depend upon $X_0, X_1, \ldots, X_{t-1}$.

Processes like this are called

Example: Random Walk (see Chapter 7)

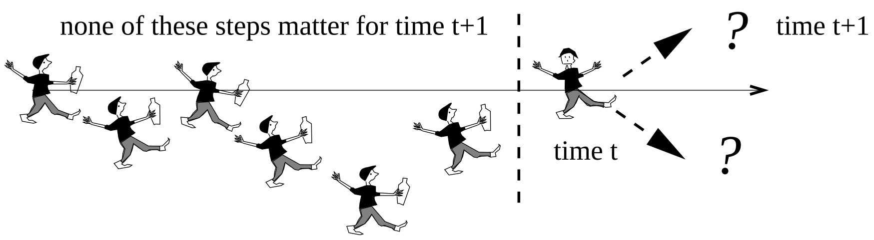

The text-book image of a Markov chain has a flea hopping about at random on the vertices of the transition diagram, according to the probabilities shown.

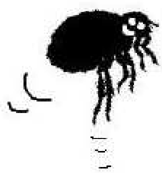

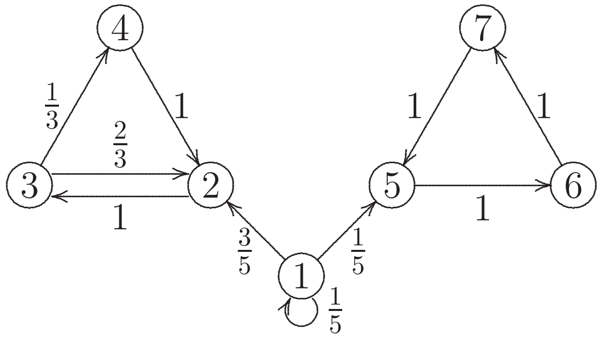

The transition diagram above shows a system with 7 possible states:

Questions of interest

- Starting from state 1, what is the probability of ever reaching state 7?
- Starting from state 2, what is the expected time taken to reach state 4?
- Starting from state 2, what is the long-run proportion of time spent in state 3?
- Starting from state 1, what is the probability of being in state 2 at time t? Does the probability converge as $t \to \infty$, and if so, to what?

We have been answering questions like the first two using first-step analysis since the start of STATS 325. In this chapter we develop a unified approach to all these questions using the matrix of transition probabilities, called the

5.2 Definitions

The Markov chain is the process

Definition: The **state** of a Markov chain at time t is the

For example, if $X_t = 6$, we say

Definition: The **state space** of a Markov chain, S, is the set of values that each X_t can take. For example, $S = \{1, 2, 3, 4, 5, 6, 7\}$.

Let S have size N (possibly infinite).

Definition: A **trajectory** of a Markov chain is

For example, if $X_0 = 1$, $X_1 = 5$, and $X_2 = 6$, then the trajectory up to time $t = 2$ is

More generally, if we refer to the trajectory $s_0, s_1, s_2, s_3, \ldots$, we mean that

'Trajectory' is just a word meaning

Markov Property

The basic property of a Markov chain is that

This is called the
It means that

We formulate the Markov Property in mathematical notation as follows:

$$\boxed{\mathbb{P}(X_{t+1} = s \,|\, X_t = s_t, X_{t-1} = s_{t-1}, \ldots, X_0 = s_0) \;=\; \mathbb{P}(X_{t+1} = s \,|\, X_t = s_t),}$$

for all $t = 1, 2, 3, \ldots$ and for all states $s_0, s_1, \ldots, s_t, s$.

Explanation:

$$\mathbb{P}(\underset{\uparrow}{X_{t+1} = s} \quad | \quad \underset{\uparrow}{X_t = s_t}, \quad \underbrace{\cancel{X_{t-1} = s_{t-1}, X_{t-2} = s_{t-2}, \ldots, X_1 = s_1, X_0 = s_0}}_{\uparrow})$$

Definition: Let $\{X_0, X_1, X_2, \ldots\}$ be a sequence of discrete random variables. Then $\{X_0, X_1, X_2, \ldots\}$ is a **Markov chain** if

---

5.3 The Transition Matrix

We have seen many examples of **transition diagrams** to describe Markov chains. The transition diagram is so-called because it shows the transitions between different states.

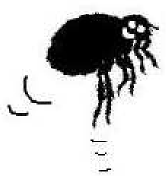

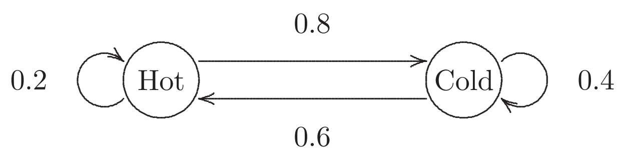

We can also summarize the probabilities in a **matrix:**

The matrix describing the Markov chain is called the
It is the most important tool for analysing Markov chains.

Transition Matrix

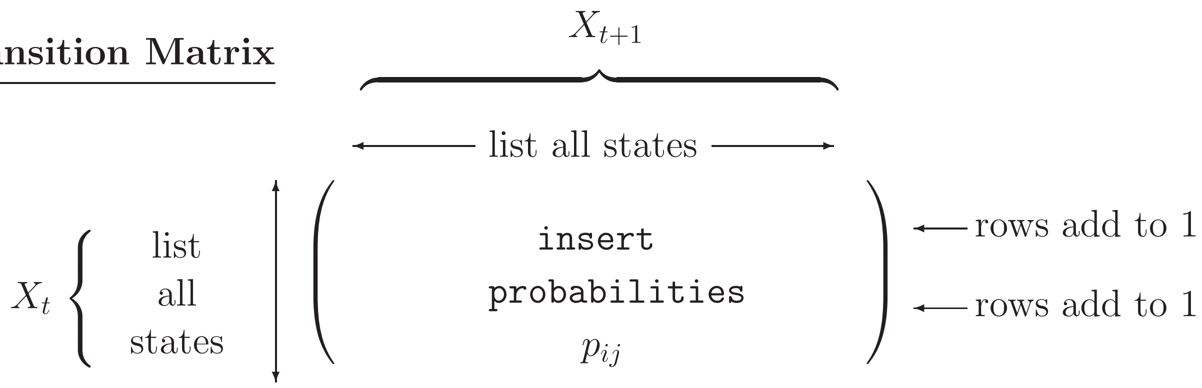

The transition matrix is usually given the symbol

In the transition matrix P:

-
-
-

Notes: 1. The transition matrix P must list *all* possible states in the state space S.

2. P is a *square matrix* $(N \times N)$, because X_{t+1} and X_t both take values in the same state space S (of size N).

3. The **rows** of P should each

$$\sum_{j=1}^{N} p_{ij} = \sum_{j=1}^{N} \mathbb{P}(X_{t+1} = j \mid X_t = i) = \sum_{j=1}^{N} \mathbb{P}_{\{X_t = i\}}(X_{t+1} = j) = 1.$$

This simply states that X_{t+1} *must* take one of the listed values.

4. The **columns** of P do **not** in general sum to 1.

Definition: Let $\{X_0, X_1, X_2, \ldots\}$ be a Markov chain with state space S, where S has size N (possibly infinite). The **transition probabilities** of the Markov chain are

Definition: The **transition matrix** of the Markov chain is

5.4 Example: setting up the transition matrix

We can create a transition matrix for any of the transition diagrams we have seen in problems throughout the course. For example, check the matrix below.

Example: Tennis game at Deuce.

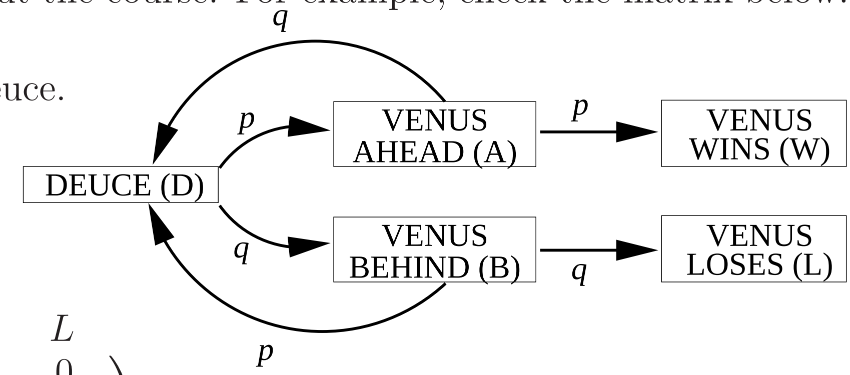

	D	A	B	W	L
D	0	p	q	0	0
A	q	0	0	p	0
B	p	0	0	0	q
W	0	0	0	1	0
L	0	0	0	0	1

5.5 Matrix Revision

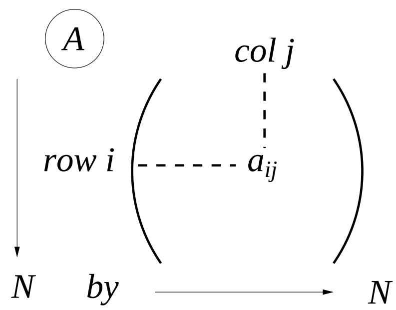

Notation

Let A be an $N \times N$ matrix.

We write $A = (a_{ij})$,
i.e. A comprises elements a_{ij}.

The (i, j) element of A is written both as a_{ij} and $(A)_{ij}$:
e.g. for matrix A^2 we might write $(A^2)_{ij}$.

Matrix multiplication

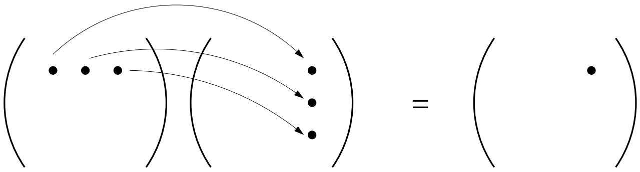

Let $A = (a_{ij})$ and $B = (b_{ij})$ be $N \times N$ matrices.

The product matrix is $A \times B = AB$, with elements $(AB)_{ij} = \sum_{k=1}^{N} a_{ik}b_{kj}$.

Summation notation for a matrix squared

Let A be an $N \times N$ matrix. Then

$$(A^2)_{ij} = \sum_{k=1}^{N} (A)_{ik}(A)_{kj} = \sum_{k=1}^{N} a_{ik}a_{kj}.$$

Pre-multiplication of a matrix by a vector

Let A be an $N \times N$ matrix, and let $\boldsymbol{\pi}$ be an $N \times 1$ column vector: $\boldsymbol{\pi} = \begin{pmatrix} \pi_1 \\ \vdots \\ \pi_N \end{pmatrix}$.

We can pre-multiply A by $\boldsymbol{\pi}^T$ to get a $1 \times N$ row vector, $\boldsymbol{\pi}^T A = \big((\boldsymbol{\pi}^T A)_1, \ldots, (\boldsymbol{\pi}^T A)_N\big)$, with elements

$$(\boldsymbol{\pi}^T A)_j = \sum_{i=1}^{N} \pi_i a_{ij}.$$

5.6 The t-step transition probabilities

Let $\{X_0, X_1, X_2, \ldots\}$ be a Markov chain with state space $S = \{1, 2, \ldots, N\}$.

Recall that the elements of the transition matrix P are defined as:

$$(P)_{ij} = p_{ij} = \mathbb{P}(X_1 = j \mid X_0 = i) = \mathbb{P}(X_{n+1} = j \mid X_n = i) \quad \text{for any } n.$$

p_{ij} is the probability of making a transition FROM state i TO state j in a SINGLE step.

Question: what is the probability of making a transition from state i to state j over **two** steps?

We are seeking $\mathbb{P}(X_2 = j \mid X_0 = i)$. Use the

The two-step transition probabilities are therefore given by

3-step transitions: We can find $\mathbb{P}(X_3 = j \mid X_0 = i)$ similarly, but conditioning on the state at time 2:

$$\begin{aligned}
\mathbb{P}(X_3 = j \mid X_0 = i) &= \sum_{k=1}^{N} \mathbb{P}(X_3 = j \mid X_2 = k)\mathbb{P}(X_2 = k \mid X_0 = i) \\
&= \sum_{k=1}^{N} p_{kj} \left(P^2\right)_{ik} \\
&= (P^3)_{ij}.
\end{aligned}$$

The three-step transition probabilities are therefore given by the matrix P^3:

$$\mathbb{P}(X_3 = j \mid X_0 = i) = \mathbb{P}(X_{n+3} = j \mid X_n = i) = \left(P^3\right)_{ij} \quad \text{for any } n.$$

General case: t-step transitions

The above working extends to show that the t-step transition probabilities are given by the matrix P^t for any t:

We have proved the following Theorem.

Theorem 5.6: Let $\{X_0, X_1, X_2, \ldots\}$ be a Markov chain with $N \times N$ transition matrix P. Then the t-step transition probabilities are given by the matrix P^t. That is,

$$\mathbb{P}(X_t = j \mid X_0 = i) = \left(P^t\right)_{ij}.$$

It also follows that

$$\mathbb{P}(X_{n+t} = j \mid X_n = i) = \left(P^t\right)_{ij} \text{ for any } n. \qquad \square$$

5.7 Distribution of X_t

Let $\{X_0, X_1, X_2, \ldots\}$ be a Markov chain with state space $S = \{1, 2, \ldots, N\}$.

Now each X_t is a random variable, so it has a

We can write the probability distribution of X_t as an

For example, consider X_0. Let $\boldsymbol{\pi}$ be an $N \times 1$ vector denoting the probability distribution of X_0:

In the flea model, this corresponds to

Notation: we will write to denote that the row vector of probabilities is given by the row vector $\boldsymbol{\pi}^T$.

Probability distribution of $\boldsymbol{X_1}$

Use the Partition Rule, conditioning on X_0:

This shows that
The row vector $\boldsymbol{\pi}^T P$ is therefore

Probability distribution of $\boldsymbol{X_2}$

Using the Partition Rule as before, conditioning again on X_0:

$$\mathbb{P}(X_2 = j) = \sum_{i=1}^{N} \mathbb{P}(X_2 = j \mid X_0 = i)\mathbb{P}(X_0 = i) = \sum_{i=1}^{N} \left(P^2\right)_{ij} \pi_i = \left(\boldsymbol{\pi}^T P^2\right)_j .$$

The row vector $\boldsymbol{\pi}^T P^2$ is therefore the probability distribution of X_2:

$$\boxed{\begin{array}{rcl} X_0 & \sim & \boldsymbol{\pi}^T \\ X_1 & \sim & \boldsymbol{\pi}^T P \\ X_2 & \sim & \boldsymbol{\pi}^T P^2 \\ & \vdots & \\ X_t & \sim & \boldsymbol{\pi}^T P^t. \end{array}}$$

These results are summarized in the following Theorem.

Theorem 5.7: Let $\{X_0, X_1, X_2, \ldots\}$ be a Markov chain with $N \times N$ transition matrix P. If the probability distribution of X_0 is given by the $1 \times N$ row vector $\boldsymbol{\pi}^T$, then the probability distribution of X_t is given by the $1 \times N$ row vector $\boldsymbol{\pi}^T P^t$. That is,

Note: The distribution of X_t is
The distribution of X_{t+1} is
Taking one step in the Markov chain corresponds to

Note: The t-step transition matrix is .
The $(t+1)$-step transition matrix is
Again, taking one step in the Markov chain corresponds to

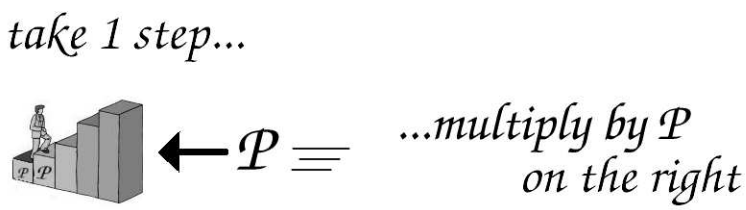

5.8 Trajectory Probability

Recall that a trajectory is a sequence of values for $X_0, X_1, \ldots, X_t$.

Because of the Markov Property, we can find the probability of any trajectory by multiplying together the starting probability and all subsequent single-step probabilities.

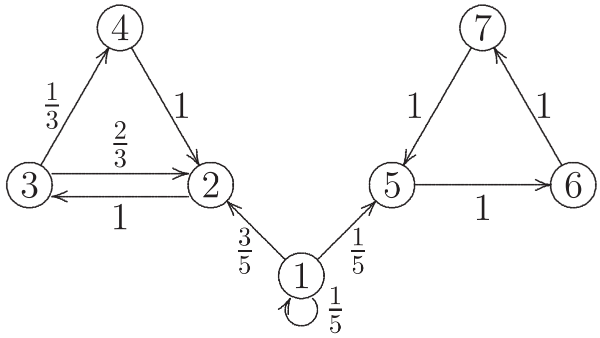

Example: Let $X_0 \sim (\frac{3}{4}, 0, \frac{1}{4}, 0, 0, 0, 0)$. What is the probability of the trajectory 1, 2, 3, 2, 3, 4?

Proof in formal notation using the Markov Property:

Let $X_0 \sim \boldsymbol{\pi}^T$. We wish to find the probability of the trajectory $s_0, s_1, s_2, \ldots, s_t$.

$$\begin{aligned}
&\mathbb{P}(X_0 = s_0, X_1 = s_1, \ldots, X_t = s_t) \\
&= \mathbb{P}(X_t = s_t \,|\, X_{t-1} = s_{t-1}, \ldots, X_0 = s_0) \times \mathbb{P}(X_{t-1} = s_{t-1}, \ldots, X_0 = s_0) \\
&= \mathbb{P}(X_t = s_t \,|\, X_{t-1} = s_{t-1}) \times \mathbb{P}(X_{t-1} = s_{t-1}, \ldots, X_0 = s_0) \quad \text{(Markov Property)} \\
&= p_{s_{t-1},s_t}\mathbb{P}(X_{t-1} = s_{t-1} \,|\, X_{t-2} = s_{t-2}, \ldots, X_0 = s_0) \times \mathbb{P}(X_{t-2} = s_{t-2}, \ldots, X_0 = s_0) \\
&\quad \vdots \\
&= p_{s_{t-1},s_t} \times p_{s_{t-2},s_{t-1}} \times \ldots \times p_{s_0,s_1} \times \mathbb{P}(X_0 = s_0) \\
&= p_{s_{t-1},s_t} \times p_{s_{t-2},s_{t-1}} \times \ldots \times p_{s_0,s_1} \times \pi_{s_0}.
\end{aligned}$$

5.9 Worked Example: distribution of X_t and trajectory probabilities

Purpose-flea zooms around the vertices of the transition diagram opposite. Let X_t be Purpose-flea's state at time t $(t = 0, 1, \ldots)$.

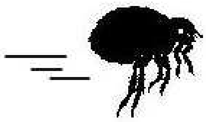

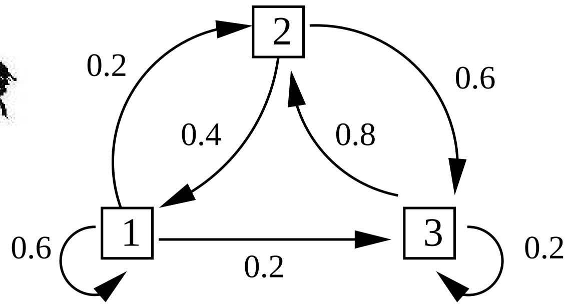

(a) Find the transition matrix, P.

Answer: $P = \begin{pmatrix} 0.6 & 0.2 & 0.2 \\ 0.4 & 0 & 0.6 \\ 0 & 0.8 & 0.2 \end{pmatrix}$

(b) Find $\mathbb{P}(X_2 = 3 \mid X_0 = 1)$.

$$\mathbb{P}(X_2 = 3 \mid X_0 = 1) = \left(P^2\right)_{13} = \begin{pmatrix} 0.6 & 0.2 & 0.2 \\ \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \end{pmatrix} \begin{pmatrix} \cdot & \cdot & 0.2 \\ \cdot & \cdot & 0.6 \\ \cdot & \cdot & 0.2 \end{pmatrix}$$

$$= 0.6 \times 0.2 + 0.2 \times 0.6 + 0.2 \times 0.2$$

$$= 0.28.$$

Note: we only need one element of the matrix P^2, so don't lose exam time by finding the whole matrix.

(c) Suppose that Purpose-flea is equally likely to start on any vertex at time 0. Find the probability distribution of X_1.

From this info, the distribution of X_0 is $\boldsymbol{\pi}^T = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$. We need $X_1 \sim \boldsymbol{\pi}^T P$.

$$\boldsymbol{\pi}^T P = \left(\tfrac{1}{3} \quad \tfrac{1}{3} \quad \tfrac{1}{3}\right) \begin{pmatrix} 0.6 & 0.2 & 0.2 \\ 0.4 & 0 & 0.6 \\ 0 & 0.8 & 0.2 \end{pmatrix} = \left(\tfrac{1}{3} \quad \tfrac{1}{3} \quad \tfrac{1}{3}\right).$$

Thus $X_1 \sim \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$ and therefore X_1 is also equally likely to be 1, 2, or 3.

(d) Suppose that Purpose-flea begins at vertex 1 at time 0. Find the probability distribution of X_2.

The distribution of X_0 *is now* $\boldsymbol{\pi}^T = (1, 0, 0)$. *We need* $X_2 \sim \boldsymbol{\pi}^T P^2$.

$$\boldsymbol{\pi}^T P^2 = (1 \;\; 0 \;\; 0) \begin{pmatrix} 0.6 & 0.2 & 0.2 \\ 0.4 & 0 & 0.6 \\ 0 & 0.8 & 0.2 \end{pmatrix} \begin{pmatrix} 0.6 & 0.2 & 0.2 \\ 0.4 & 0 & 0.6 \\ 0 & 0.8 & 0.2 \end{pmatrix}$$

$$= (0.6 \;\; 0.2 \;\; 0.2) \begin{pmatrix} 0.6 & 0.2 & 0.2 \\ 0.4 & 0 & 0.6 \\ 0 & 0.8 & 0.2 \end{pmatrix}$$

$$= (0.44 \quad 0.28 \quad 0.28) \,.$$

Thus $\mathbb{P}(X_2 = 1) = 0.44$, $\mathbb{P}(X_2 = 2) = 0.28$, $\mathbb{P}(X_2 = 3) = 0.28$.

Note that it is quickest to multiply the vector by the matrix first: we don't need to compute P^2 *in entirety.*

(e) Suppose that Purpose-flea is equally likely to start on any vertex at time 0. Find the probability of obtaining the trajectory (3, 2, 1, 1, 3).

$$\begin{aligned} \mathbb{P}(3, 2, 1, 1, 3) &= \mathbb{P}(X_0 = 3) \times p_{32} \times p_{21} \times p_{11} \times p_{13} \quad \textit{(Section 5.8)} \\ &= \tfrac{1}{3} \times 0.8 \times 0.4 \times 0.6 \times 0.2 \\ &= 0.0128. \end{aligned}$$

5.10 Class Structure

The state space of a Markov chain can be partitioned into a set of non-overlapping

States i and j are in the same communicating class if there is some way of getting from state i to state j, AND there is some way of getting from state j to state i. It needn't be possible to get between i and j in a ***single*** step, but it must be possible over some number of steps to travel between them both ways.

We write

Definition: Consider a Markov chain with state space S and transition matrix P, and consider states $i, j \in S$. Then **state *i* communicates with state *j*** if:

1. there exists some t such that $(P^t)_{ij} > 0$, AND
2. there exists some u such that $(P^u)_{ji} > 0$.

Mathematically, it is easy to show that the communicating relation $\leftrightarrow$ is an **equivalence relation,** which means that it partitions the sample space S into non-overlapping equivalence classes.

Definition: States i and j are in the same **communicating class** if

Every state is a member of

Example: Find the communicating classes associated with the transition diagram shown.

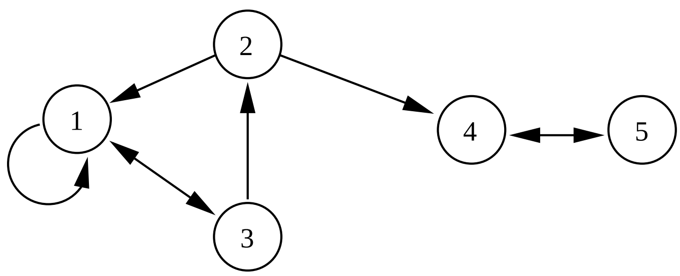

Solution:

Definition: A communicating class of states is **closed** if

That is, the communicating class C is closed if $p_{ij} = 0$ whenever $i \in C$ and $j \notin C$.

Example: In the transition diagram above:

- Class $\{1, 2, 3\}$ is
- Class $\{4, 5\}$ is

Definition: A state i is said to be **absorbing** if

Definition: A Markov chain or transition matrix P is said to be **irreducible** if

---

5.11 Hitting Probabilities

We have been calculating hitting probabilities for Markov chains since Chapter 2, using First-Step Analysis. The hitting probability describes the probability that the Markov chain will *ever* reach some state or set of states.

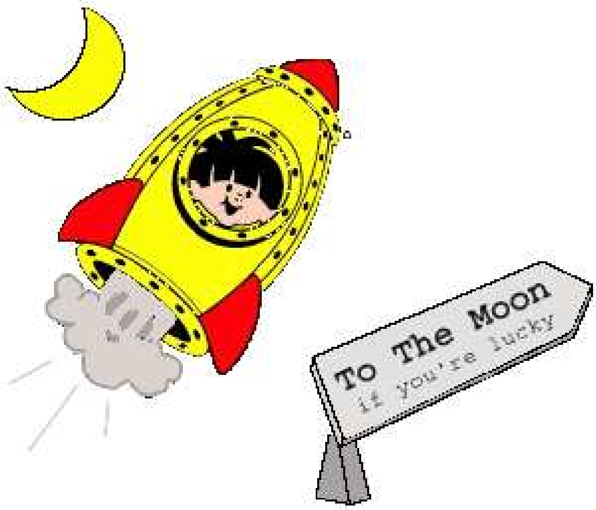

In this section we show how hitting probabilities can be written in a single vector. We also see a general formula for calculating the hitting probabilities. In general it is easier to continue using our own common sense, but occasionally the formula becomes more necessary.

Vector of hitting probabilities

Let A be some subset of the state space S. (A need not be a communicating class: it can be any subset required, including the subset consisting of a single state: e.g. $A = \{4\}$.)

The **hitting probability** from state i to set A is the probability of ever reaching the set A, starting from initial state i. We write this probability as
Thus

Example: Let set $A = \{1, 3\}$ as shown.

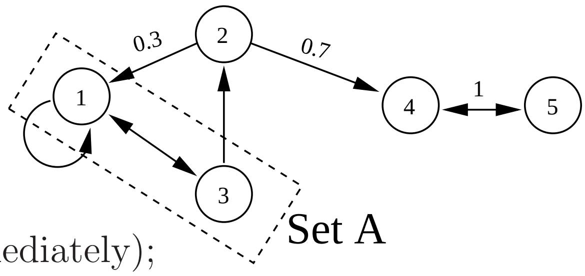

The hitting probability for set A is:

-
  (We are starting in set A, so we hit it immediately);

-
  (The set $\{4, 5\}$ is a closed class, so we can never escape out to set A);

-
  (We could hit A at the first step (probability 0.3), but otherwise we move to state 4 and get stuck in the closed class $\{4, 5\}$ (probability 0.7).)

We can summarize all the information from the example above in a

Note: When A is a closed class, the hitting probability h_{iA} is called the

In general, if there are N possible states, the vector of hitting probabilities is

Example: finding the hitting probability vector using First-Step Analysis

Suppose $\{X_t : t \geq 0\}$ has the following transition diagram:

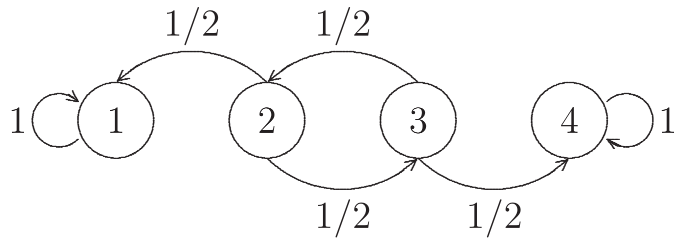

Find the vector of hitting probabilities for state 4.

Solution:

Formula for hitting probabilities

In the previous example, we used our common sense to state that $h_{14} = 0$. While this is easy for a human brain, it is harder to explain a general rule that would describe this 'common sense' mathematically, or that could be used to write computer code that will work for all problems.

Although it is usually best to continue to use common sense when solving problems, this section provides a general formula that will *always* work to find a vector of hitting probabilities $\boldsymbol{h_A}$.

Theorem 5.11: The vector of hitting probabilities $\boldsymbol{h_A} = (h_{iA} \; : \; i \in S)$ is the minimal non-negative solution to the following equations:

$$h_{iA} = \begin{cases} 1 & \text{for} \quad i \in A, \\ \sum_{j \in S} p_{ij} h_{jA} & \text{for} \quad i \notin A. \end{cases}$$

The 'minimal non-negative solution' means that:

1. the values $\{h_{iA}\}$ collectively satisfy the equations above;
2. each value h_{iA} is ≥ 0 (non-negative);
3. given any other non-negative solution to the equations above, say $\{g_{iA}\}$ where $g_{iA} \geq 0$ for all i, then $h_{iA} \leq g_{iA}$ for all i (minimal solution).

Example: How would this formula be used to substitute for 'common sense' in the previous example?

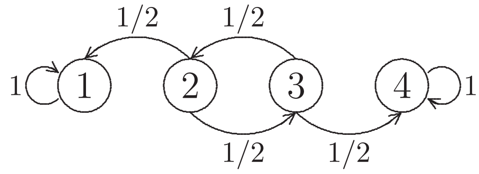

The equations give:

$$h_{i4} = \begin{cases} 1 & \textit{if} \quad i = 4, \\ \sum_{j \in S} p_{ij} h_{j4} & \textit{if} \quad i \neq 4. \end{cases}$$

Thus,

$$\begin{aligned} h_{44} &= 1 \\ h_{14} &= h_{14} \quad \textit{unspecified! Could be anything!} \\ h_{24} &= \tfrac{1}{2} h_{14} + \tfrac{1}{2} h_{34} \\ h_{34} &= \tfrac{1}{2} h_{24} + \tfrac{1}{2} h_{44} \;=\; \tfrac{1}{2} h_{24} + \tfrac{1}{2} \end{aligned}$$

Because h_{14} could be anything, we have to use the minimal non-negative value, which is $h_{14} = 0$.
(Need to check $h_{14} = 0$ does not force $h_{i4<0}$ for any other i: OK.)

The other equations can then be solved to give the same answers as before. $\square$

Proof of Theorem 5.11 (non-examinable):

Consider the equations $h_{iA} = \begin{cases} 1 & \text{for } i \in A, \\ \sum_{j \in S} p_{ij} h_{jA} & \text{for } i \notin A. \end{cases}$ $(\star)$

We need to show that:

(i) the hitting probabilities $\{h_{iA}\}$ collectively satisfy the equations $(\star)$;

(ii) if $\{g_{iA}\}$ is any other non-negative solution to $(\star)$, then the hitting probabilities $\{h_{iA}\}$ satisfy $h_{iA} \le g_{iA}$ for all i (minimal solution).

Proof of (i): Clearly, $h_{iA} = 1$ if $i \in A$ (as the chain hits A immediately).

Suppose that $i \notin A$. Then

$$\begin{aligned} h_{iA} &= \mathbb{P}(X_t \in A \text{ for some } t \ge 1 \,|\, X_0 = i) \\ &= \sum_{j \in S} \mathbb{P}(X_t \in A \text{ for some } t \ge 1 \,|\, X_1 = j)\mathbb{P}(X_1 = j \,|\, X_0 = i) \\ & \qquad \text{(Partition Rule)} \\ &= \sum_{j \in S} h_{jA}\, p_{ij} \qquad \text{(by definitions).} \end{aligned}$$

Thus the hitting probabilities $\{h_{iA}\}$ must satisfy the equations $(\star)$.

Proof of (ii): Let $h_{iA}^{(t)} = \mathbb{P}(\text{hit } A \text{ at or before time } t \,|\, X_0 = i)$.

We use mathematical induction to show that $h_{iA}^{(t)} \le g_{iA}$ for all t, and therefore $h_{iA} = \lim_{t\to\infty} h_{iA}^{(t)}$ must also be $\le g_{iA}$.

Time $t = 0$: $\quad h_{iA}^{(0)} = \begin{cases} 1 & \text{if } i \in A, \\ 0 & \text{if } i \notin A. \end{cases}$

But because g_{iA} is non-negative and satisfies $(\star)$, $\begin{cases} g_{iA} = 1 & \text{if } i \in A, \\ g_{iA} \geq 0 & \text{for all } i. \end{cases}$

So $g_{iA} \geq h_{iA}^{(0)}$ for all i.

The inductive hypothesis is true for time $t = 0$.

Time t: Suppose the inductive hypothesis holds for time t, i.e.

$$h_{jA}^{(t)} \leq g_{jA} \quad \text{for all } j.$$

Consider

$$\begin{aligned} h_{iA}^{(t+1)} &= \mathbb{P}(\text{hit } A \text{ by time } t+1 \,|\, X_0 = i) \\ &= \sum_{j \in S} \mathbb{P}(\text{hit } A \text{ by time } t+1 \,|\, X_1 = j)\mathbb{P}(X_1 = j \,|\, X_0 = i) \\ &\qquad \text{(Partition Rule)} \\ &= \sum_{j \in S} h_{jA}^{(t)} p_{ij} \qquad \text{by definitions} \\ &\leq \sum_{j \in S} g_{jA} p_{ij} \qquad \text{by inductive hypothesis} \\ &= g_{iA} \qquad \text{because } \{g_{iA}\} \text{ satisfies } (\star). \end{aligned}$$

Thus $h_{iA}^{(t+1)} \leq g_{iA}$ for all i, so the inductive hypothesis is proved.

By the Continuity Theorem (Chapter 2), $h_{iA} = \lim_{t \to \infty} h_{iA}^{(t)}$.

So $h_{iA} \leq g_{iA}$ as required. □

5.12 Expected hitting times

In the previous section we found the **probability** of hitting set A, starting at state i. Now we study **how long** it takes to get from i to A. As before, it is best to solve problems using first-step analysis and common sense. However, a general formula is also available.

Definition: Let A be a subset of the state space S. The **hitting time** of A is the random variable T_A, where

T_A is the time taken before hitting set A

The hitting time T_A can take values
If the chain *never* hits set A, then

Note: The hitting time is also called the **reaching time**. If A is a closed class, it is also called the

Definition: The **mean hitting time** for A, starting from state i, is

Note: If there is any possibility that the chain *never* reaches A, starting from i,

Calculating the mean hitting times

Theorem 5.12: The vector of expected hitting times $\boldsymbol{m_A} = (m_{iA} : i \in S)$ is

Proof (sketch):

Consider the equations $m_{iA} = \begin{cases} 0 & \text{for } \; i \in A, \\ 1 + \sum_{j \notin A} p_{ij} m_{jA} & \text{for } \; i \notin A. \end{cases}$ $(\star)$.

We need to show that:

(i) the mean hitting times $\{m_{iA}\}$ collectively satisfy the equations $(\star)$;

(ii) if $\{u_{iA}\}$ is any other non-negative solution to $(\star)$, then the mean hitting times $\{m_{iA}\}$ satisfy $m_{iA} \leq u_{iA}$ for all i (minimal solution).

We will prove point (i) only. A proof of (ii) can be found online at: `http://www.statslab.cam.ac.uk/~james/Markov/` , Section 1.3.

Proof of (i): Clearly, $m_{iA} = 0$ if $i \in A$ (as the chain hits A immediately).

Suppose that $i \notin A$. Then

$$\begin{aligned} m_{iA} &= \mathbb{E}(T_A \,|\, X_0 = i) \\ &= 1 + \sum_{j \in S} \mathbb{E}(T_A \,|\, X_1 = j)\mathbb{P}(X_1 = j \,|\, X_0 = i) \\ &\quad \text{(conditional expectation: take 1 step to get to state } j \\ &\quad \text{at time 1, then find } \mathbb{E}(T_A) \text{ from there)} \\ &= 1 + \sum_{j \in S} m_{jA}\, p_{ij} \qquad \text{(by definitions)} \\ &= 1 + \sum_{j \notin A} p_{ij}\, m_{jA}\,, \qquad \text{because } m_{jA} = 0 \text{ for } j \in A. \end{aligned}$$

Thus the mean hitting times $\{m_{iA}\}$ must satisfy the equations $(\star)$.

□

Example: Let $\{X_t : t \geq 0\}$ have the same transition diagram as before:

Starting from state 2, find the expected time to absorption.

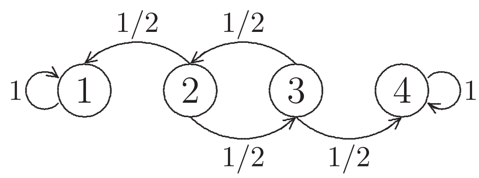

Solution:

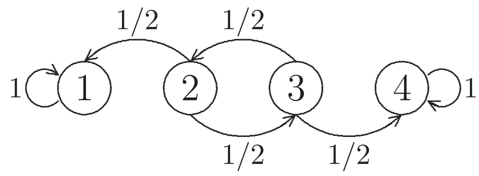

Example: Glee-flea hops around on a triangle. At each step he moves to one of the other two vertices at random. What is the expected time taken for Glee-flea to get from vertex 1 to vertex 2?

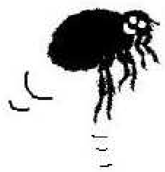

Solution:

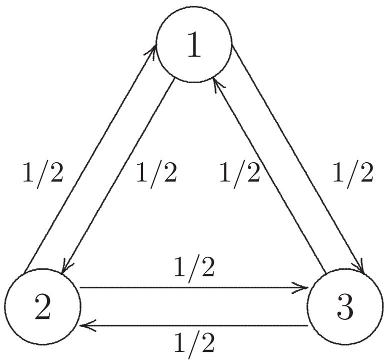

transition matrix, $P = \begin{pmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & 0 \end{pmatrix}$.

We wish to find m_{12}.

Now $$m_{i2} = \begin{cases} 0 & \textit{if} \quad i = 2, \\ 1 + \sum_{j \neq 2} p_{ij} m_{j2} & \textit{if} \quad i \neq 2. \end{cases}$$

Thus

$$\begin{aligned} m_{22} &= 0 \\ m_{12} &= 1 + \tfrac{1}{2} m_{22} + \tfrac{1}{2} m_{32} \;=\; 1 + \tfrac{1}{2} m_{32}. \\ m_{32} &= 1 + \tfrac{1}{2} m_{22} + \tfrac{1}{2} m_{12} \\ &= 1 + \tfrac{1}{2} m_{12} \\ &= 1 + \tfrac{1}{2} \left(1 + \tfrac{1}{2} m_{32}\right) \\ \Rightarrow \quad m_{32} &= 2. \end{aligned}$$

Thus $m_{12} = 1 + \frac{1}{2} m_{32} = 2$ *steps.*