

Applications to Markov chains

From Lay, §4.9

(This section is not examinable on the mid-semester exam.)

Theory and definitions

Markov chains are useful tools in certain kinds of probabilistic models. They make use of matrix algebra in a powerful way. The basic idea is the following: suppose that you are watching some collection of objects that are changing through time.

- Assume that the total number of objects is not changing, but rather their “states” (position, colour, disposition, etc) are changing.
- Further, assume that the proportion of state A objects changing to state B is constant and these changes occur at discrete stages, one after the next.

Then we are in a good position to model changes by a Markov chain.

As an example, consider the three storey aviary at a local zoo which houses 300 small birds. The aviary has three levels, and the birds spend their day flying around from one favourite perch to the next. Thus at any given time the birds seem to be randomly distributed throughout the three levels, except at feeding time when they all fly to the bottom level.

Our problem is to determine what the probability is of a given bird being at a given level of the aviary at a given time. Of course, the birds are always flying from one level to another, so the bird population on each level is constantly fluctuating. We shall use a Markov chain to model this situation.

Consider a 3×1 matrix

$$\mathbf{p} = \begin{bmatrix} p_1 \\ p_2 \\ p_3 \end{bmatrix}$$

where p_1 is the percentage of total birds on the first level, p_2 is the percentage on the second level, and p_3 is the percentage on the third level. Note that $p_1 + p_2 + p_3 = 1 = 100\%$.

After 5 min we have a new matrix

$$\mathbf{p}' = \begin{bmatrix} p'_1 \\ p'_2 \\ p'_3 \end{bmatrix}$$

giving a new distribution of the birds.

- We shall assume that the change from the \mathbf{p} matrix to the \mathbf{p}' matrix is given by a linear operator on \mathbb{R}^3 .
- In other words there is a 3×3 matrix T , known as the **transition matrix** for the Markov chain, for which $T\mathbf{p} = \mathbf{p}'$.
- After another 5 minutes we have another distribution $\mathbf{p}'' = T\mathbf{p}'$ (using the same matrix T), and so forth.

The same matrix T is used since we are assuming that the probability of a bird moving to another level is independent of time.

In other words, the probability of a bird moving to a particular level depends only on the present state of the bird, and not on any past states—it's as if the birds had no memory of their past states.

This type of model is known as a finite Markov Chain.

A sequence of trials of an experiment is a finite Markov Chain if it has the following features:

- the outcome of each trial is one of a finite set of outcomes (such as {level 1, level 2, level 3} in the aviary example);
- the outcome of one trial depends only on the immediately preceding trial.

In order to give a more formal definition we need to introduce the appropriate terminology.

Definition

A vector $\mathbf{p} = \begin{bmatrix} p_1 \\ \vdots \\ p_n \end{bmatrix}$ with nonnegative entries that add up to 1 is called a **probability vector**.

Definition

A **stochastic matrix** is a square matrix whose columns are probability vectors.

The transition matrix T described above that takes the system from one distribution to another is a stochastic matrix.

Definition

In general, a finite **Markov chain** is a sequence of probability vectors $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots$ together with a stochastic matrix T , such that

$$\mathbf{x}_1 = T\mathbf{x}_0, \mathbf{x}_2 = T\mathbf{x}_1, \mathbf{x}_3 = T\mathbf{x}_2, \dots$$

We can rewrite the above conditions as a recurrence relation

$$\mathbf{x}_{k+1} = T\mathbf{x}_k, \quad \text{for } k = 0, 1, 2, \dots$$

The vector \mathbf{x}_k is often called a **state vector**.

More generally, a recurrence relation of the form

$$\mathbf{x}_{k+1} = A\mathbf{x}_k \quad \text{for } k = 0, 1, 2, \dots$$

where A is an $n \times n$ matrix (not necessarily a stochastic matrix), and the \mathbf{x}_k s are vectors in \mathbb{R}^n (not necessarily probability vector) is called a *first order difference equation*.

Examples

Example 1

We return to the aviary example. Assume that whenever a bird is on any level of the aviary, the probability of that bird being on the same level 5 min later is $1/2$. If the bird is on the first level, the probability of moving to the second level in 5 min is $1/3$ and of moving to the third level in 5 min is $1/6$. For a bird on the second level, the probability of moving to either the first or third level is $1/4$. Finally for a bird on the third level, the probability of moving to the second level is $1/3$ and of moving to the first is $1/6$.

We want to find the transition matrix for this example and use it to determine the distribution after certain periods of time.

From the information given, we derive the following matrix as the transition matrix:

$$\begin{array}{c}
 \text{From:} \\
 \text{lev 1} \quad \text{lev 2} \quad \text{lev 3} \quad \text{To:} \\
 T = \begin{array}{ccc|c}
 1/2 & 1/4 & 1/6 & \text{lev 1} \\
 1/3 & 1/2 & 1/3 & \text{lev 2} \\
 1/6 & 1/4 & 1/2 & \text{lev 3}
 \end{array}
 \end{array}$$

Note that in each column, the sum of the probabilities is 1.

Using T we can now compute what happens to the bird distribution at 5-min intervals.

Suppose that immediately after breakfast all the birds are in the dining area on the first level. Where are they in 5 min? The probability matrix at time 0 is

$$\mathbf{p} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$$

According to the Markov chain model the bird distribution after 5 min is

$$T\mathbf{p} = \begin{bmatrix} 1/2 & 1/4 & 1/6 \\ 1/3 & 1/2 & 1/3 \\ 1/6 & 1/4 & 1/2 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 1/2 \\ 1/3 \\ 1/6 \end{bmatrix}$$

After another 5 min the bird distribution becomes

$$T \begin{bmatrix} 1/2 \\ 1/3 \\ 1/6 \end{bmatrix} = \begin{bmatrix} 13/36 \\ 7/18 \\ 1/4 \end{bmatrix}$$

Example 2

We investigate the weather in the Land of Oz. to illustrate the principles without too much heavy calculation.) The weather here is not ver good: there are never two fine days in a row.

If the weather on a particular day is known, we cannot predict exactly what the weather will be the next day, but we can predict the probabilities of various kinds of weather. We will say that there are only three kinds: fine, cloudy and rain.

Here is the behaviour:

- After a fine day, the weather is equally likely to be cloudy or rain.
- After a cloudy day, the probabilities are 1/4 fine, 1/4 cloudy and 1/2 rain.
- After rain, the probabilities are 1/4 fine, 1/2 cloudy and 1/4 rain.

We aim to find the transition matrix and use it to investigate some of the weather patterns in the Land of Oz.

The information gives a transition matrix:

From: To:

$$T = \begin{array}{ccc|c} \textit{fine} & \textit{cloudy} & \textit{rain} & \\ \hline 0 & 1/4 & 1/4 & \textit{fine} \\ 1/2 & 1/4 & 1/2 & \textit{cloudy} \\ 1/2 & 1/2 & 1/4 & \textit{rain} \end{array}$$

Suppose on day 0 that the weather is rainy. That is

$$\mathbf{x}_0 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}.$$

Then the probabilities for the weather the next day are

$$\mathbf{x}_1 = T\mathbf{x}_0 = \begin{bmatrix} 0 & 1/4 & 1/4 \\ 1/2 & 1/4 & 1/2 \\ 1/2 & 1/2 & 1/4 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 1/4 \\ 1/2 \\ 1/4 \end{bmatrix},$$

and for the next day

$$\mathbf{x}_2 = T\mathbf{x}_1 = \begin{bmatrix} 0 & 1/4 & 1/4 \\ 1/2 & 1/4 & 1/2 \\ 1/2 & 1/2 & 1/4 \end{bmatrix} \begin{bmatrix} 1/4 \\ 1/2 \\ 1/4 \end{bmatrix} = \begin{bmatrix} 3/16 \\ 3/8 \\ 7/16 \end{bmatrix}$$

If we want to find the probabilities for the weather for a week after the initial rainy day, we can calculate like this

$$\mathbf{x}_7 = T\mathbf{x}_6 = T^2\mathbf{x}_5 = T^3\mathbf{x}_4 = \dots = T^7\mathbf{x}_0.$$

Predicting the distant future

The most interesting aspect of Markov chains is the study of the chain's long term behaviour.

Example 3

Consider a system whose state is described by the Markov chain $\mathbf{x}_{k+1} = T\mathbf{x}_k$, for $k = 0, 1, 2, \dots$, where T is the matrix

$$T = \begin{bmatrix} .7 & .2 & .2 \\ 0 & .2 & .4 \\ .3 & .6 & .4 \end{bmatrix} \quad \text{and} \quad \mathbf{x}_0 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}.$$

We want to investigate what happens to the system as time passes.

To do this we compute the state vector for several different times. We find

$$\mathbf{x}_1 = T\mathbf{x}_0 = \begin{bmatrix} .7 & .2 & .2 \\ 0 & .2 & .4 \\ .3 & .6 & .4 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 0.2 \\ 0.4 \\ 0.4 \end{bmatrix}$$

$$\mathbf{x}_2 = T\mathbf{x}_1 = \begin{bmatrix} .7 & .2 & .2 \\ 0 & .2 & .4 \\ .3 & .6 & .4 \end{bmatrix} \begin{bmatrix} 0.2 \\ 0.4 \\ 0.4 \end{bmatrix} = \begin{bmatrix} 0.3 \\ 0.24 \\ 0.46 \end{bmatrix}$$

$$\mathbf{x}_3 = T\mathbf{x}_2 = \begin{bmatrix} .7 & .2 & .2 \\ 0 & .2 & .4 \\ .3 & .6 & .4 \end{bmatrix} \begin{bmatrix} 0.3 \\ 0.24 \\ 0.46 \end{bmatrix} = \begin{bmatrix} 0.350 \\ 0.232 \\ 0.416 \end{bmatrix}$$

Subsequent calculations give

$$\mathbf{x}_4 = \begin{bmatrix} 0.3750 \\ 0.2136 \\ 0.4114 \end{bmatrix}, \quad \mathbf{x}_5 = \begin{bmatrix} 0.38750 \\ 0.20728 \\ 0.40522 \end{bmatrix},$$

$$\mathbf{x}_6 = \begin{bmatrix} 0.393750 \\ 0.203544 \\ 0.4027912 \end{bmatrix}, \quad \mathbf{x}_7 = \begin{bmatrix} 0.3968750 \\ 0.2017912 \\ 0.4013338 \end{bmatrix},$$

$$\mathbf{x}_8 = \begin{bmatrix} 0.39843750 \\ 0.20089176 \\ 0.4006704 \end{bmatrix}, \quad \mathbf{x}_9 = \begin{bmatrix} 0.399218750 \\ 0.200448848 \\ 0.400034602 \end{bmatrix},$$

$$\dots, \mathbf{x}_{20} = \begin{bmatrix} 0.3999996185 \\ 0.2000002179 \\ 0.4000001634 \end{bmatrix}.$$

These vectors seem to be approaching

$$\mathbf{q} = \begin{bmatrix} 0.4 \\ 0.2 \\ 0.4 \end{bmatrix}.$$

Observe the following calculation:

$$T\mathbf{q} = \begin{bmatrix} .7 & .2 & .2 \\ 0 & .2 & .4 \\ .3 & .6 & .4 \end{bmatrix} \begin{bmatrix} 0.4 \\ 0.2 \\ 0.4 \end{bmatrix} = \begin{bmatrix} 0.4 \\ 0.2 \\ 0.4 \end{bmatrix}.$$

This calculation is exact, with no rounding error. When the system is in state \mathbf{q} there is no change in the system from one measurement to the next.

We might also note that T^{20} is given by

$$\begin{bmatrix} 0.4000005722 & 0.3999996185 & 0.3999996185 \\ 0.1999996730 & 0.2000002180 & 0.2000002179 \\ 0.3999997548 & 0.4000001635 & 0.4000001634 \end{bmatrix}.$$

Example 4

For the weather in the Land of Oz, where

$$T = \begin{bmatrix} 0 & 0.25 & 0.25 \\ 0.5 & 0.25 & 0.5 \\ 0.5 & 0.5 & 0.25 \end{bmatrix}, \quad \mathbf{x}_0 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

we have already calculated

$$\mathbf{x}_7 = \begin{bmatrix} 0.2000122070 \\ 0.4000244140 \\ 0.3999633789 \end{bmatrix}.$$

We want to look further ahead.

A further calculation gives

$$\mathbf{x}_{15} = \begin{bmatrix} 0.2000000002 \\ 0.4000000003 \\ 0.3999999994 \end{bmatrix}.$$

This suggests that

$$\mathbf{q} = \begin{bmatrix} 0.2 \\ 0.4 \\ 0.4 \end{bmatrix}.$$

An easy calculation shows that $T\mathbf{q} = \mathbf{q}$.

Steady-state vectors

Definition

If T is a stochastic matrix, then a **steady state vector** for T is a probability vector \mathbf{q} such that

$$T\mathbf{q} = \mathbf{q}.$$

A steady state vector \mathbf{q} for T represents an *equilibrium* of the system modeled by the Markov Chain with transition matrix T . If at time 0 the system is in state \mathbf{q} (that is if we have $\mathbf{x}_0 = \mathbf{q}$) then the system will remain in state \mathbf{q} at all times (that is we will have $\mathbf{x}_n = \mathbf{q}$ for every $n \geq 0$).

It can be shown that every stochastic matrix has a steady state vector. In the examples in Section 2, the vector \mathbf{q} is the steady state vector.

To find a suitable vector \mathbf{q} , we want to solve the equation $T\mathbf{x} = \mathbf{x}$.

$$\begin{aligned} T\mathbf{x} - \mathbf{x} &= \mathbf{0} \\ T\mathbf{x} - I\mathbf{x} &= \mathbf{0} \\ (T - I)\mathbf{x} &= \mathbf{0} \end{aligned}$$

In the case $n = 2$, the problem is easily solved directly. Suppose first that all the entries of the transition matrix T are non-zero. Then T must be of the form

$$T = \begin{bmatrix} 1-p & q \\ p & 1-q \end{bmatrix} \text{ for } 0 < p, q < 1.$$

Then

$$T - I = \begin{bmatrix} -p & q \\ p & -q \end{bmatrix} \xrightarrow{\text{rref}} \begin{bmatrix} -p & q \\ 0 & 0 \end{bmatrix}.$$

So when solving $(T - I)\mathbf{x} = \mathbf{0}$, x_2 is free and $px_1 = qx_2$, so that

$$\mathbf{q} = \frac{1}{p+q} \begin{bmatrix} q \\ p \end{bmatrix}$$

is a steady state probability vector. Note that in this particular case the steady state vector is unique.

The case when one or more of the entries of T are zero is handled in a similar way. Note that if $p = q = 0$ then T is the identity matrix for which every probability vector is clearly a steady state vector.

A stochastic matrix does not necessarily have a *unique* steady state vector. In other words, a system modeled by a Markov Chain can have more than one equilibrium.

For example the probability vectors

$$\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1/2 \\ 1/2 \end{bmatrix}, \begin{bmatrix} 1/3 \\ 1/3 \\ 1/3 \end{bmatrix}$$

are all steady state vectors for the stochastic matrix

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

Indeed all the probability vectors

$$\begin{bmatrix} a \\ b \\ b \end{bmatrix} \text{ with } a, b \geq 0 \text{ and } a + 2b = 1$$

are steady state vectors for the above matrix T .

We would like to have some conditions on P that ensure that T has a unique steady state vector \mathbf{q} and that the Markov Chain \mathbf{x}_n associated to T converges to the steady state \mathbf{q} , independently of the initial state \mathbf{x}_0 . For this kind of Markov chains, we can easily predict the long term behaviour.

It turns out that there is a large set of stochastic matrices for which long range predictions *are* possible. Before stating the main theorem we have to give a definition.

Definition

A stochastic matrix T is **regular** if some matrix power T^k contains only strictly positive entries.

In other words, if the transition matrix of a Markov chain is regular then, for some k , it is possible to go from any state to any state (including remaining in the current state) in exactly k steps.

For the transition matrix showing the probabilities for change in the weather in the Land of Oz, we have

$$T = \begin{bmatrix} 0 & 1/4 & 1/4 \\ 1/2 & 1/4 & 1/2 \\ 1/2 & 1/2 & 1/4 \end{bmatrix}$$

However,

$$T^2 = \begin{bmatrix} 1/4 & 3/16 & 3/16 \\ 3/8 & 7/16 & 3/8 \\ 3/8 & 3/8 & 7/16 \end{bmatrix}$$

which shows that T is a regular stochastic matrix.

Here's an example of a stochastic matrix that is not regular:

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

Not only does T have some zero entries, but also

$$T^2 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = I_2$$

$$T^3 = TT^2 = TI_2 = T$$

so that

$$T^k = T \quad \text{if } k \text{ is odd,} \quad T^k = I_2 \quad \text{if } k \text{ is even.}$$

Thus any matrix power T^k has some entries equal to zero.

Theorem

If T is an $n \times n$ regular stochastic matrix, then T has a unique steady state vector \mathbf{q} . The entries of \mathbf{q} are strictly positive

Moreover, if \mathbf{x}_0 is any initial probability vector and $\mathbf{x}_{k+1} = T\mathbf{x}_k$ for $k = 0, 1, 2, \dots$ then the Markov chain $\{\mathbf{x}_k\}$ converges to \mathbf{q} as $k \rightarrow \infty$.

Equivalently, the steady state vector \mathbf{q} is the limit of $T^k\mathbf{x}_0$ when $k \rightarrow \infty$ for any probability vector \mathbf{x}_0 .

Notice that if $T = [\mathbf{p}_1 \dots \mathbf{p}_n]$, where $\mathbf{p}_1, \dots, \mathbf{p}_n$ are the columns of T , then taking $\mathbf{x}_0 = \mathbf{e}_i$, where \mathbf{e}_i is the i th vector of the standard basis we have that

$$\mathbf{x}_1 = T\mathbf{x}_0 = T\mathbf{e}_i = \mathbf{p}_i$$

so \mathbf{x}_1 is the i th column of T .

Similarly $\mathbf{x}_k = T^k\mathbf{x}_0 = T^k\mathbf{e}_i$ is the i th column of T^k .

The previous theorem implies that $T^k\mathbf{e}_i \rightarrow \mathbf{q}$ for every $i = 1, \dots, n$ when $k \rightarrow \infty$, that is every column of T^k approaches the limiting vector \mathbf{q} when $k \rightarrow \infty$.

Examples

Example 5

Let $T = \begin{bmatrix} 0.8 & 0.5 \\ 0.2 & 0.5 \end{bmatrix}$. We want to find the steady state vector associated with T .

We want to solve $(T - I)\mathbf{x} = \mathbf{0}$:

$$T - I = \begin{bmatrix} -0.2 & 0.5 \\ 0.2 & -0.5 \end{bmatrix} \rightarrow R = \begin{bmatrix} 1 & -5/2 \\ 0 & 0 \end{bmatrix}$$

The homogeneous system having the reduced row echelon matrix R as coefficient matrix is $x_1 - (5/2)x_2 = 0$. Taking x_2 as a free variable, the general solution is $x_1 = (5/2)t$, $x_2 = t$.

For \mathbf{x} to be a probability vector we also require $x_1 + x_2 = 1$.

Put $x_1 = (5/2)t$, $x_2 = t$, then $x_1 + x_2 = 1$ becomes $(5/2)t + t = 1$.

This gives $t = 2/7 = x_2$ and $x_1 = 5/7$, so $\mathbf{x} = \begin{bmatrix} 5/7 \\ 2/7 \end{bmatrix}$.

An alternative Solution

If we consider $T = \begin{bmatrix} 0.8 & 0.5 \\ 0.2 & 0.5 \end{bmatrix}$ as a matrix of the form


$$\begin{bmatrix} 1-p & q \\ p & 1-q \end{bmatrix}$$

we can identify $p = 0.2$ and $q = 0.5$. The solution is then given by

$$\mathbf{p} = \frac{1}{p+q} \begin{bmatrix} q \\ p \end{bmatrix} = \frac{1}{0.7} \begin{bmatrix} 0.5 \\ 0.2 \end{bmatrix} = \begin{bmatrix} 5/7 \\ 2/7 \end{bmatrix}.$$

Example 6

A psychologist places a rat in a cage with three compartments, as shown in the diagram.



The rat has been trained to select a door at random whenever a bell is rung and to move through it into the next compartment.

Example (continued)

From the diagram, if the rat is in space 1, there are equal probabilities that it will go to either space 2 or 3 (because there is just one opening to each of these spaces).

On the other hand, if the rat is in space 2, there is one door to space 1, and 2 to space 3, so the probability that it will go to space 1 is $1/3$, and to space 3 is $2/3$.

The situation is similar if the rat is in space 3. Wherever the rat is there is 0 probability that the rat will stay in that space.

The transition matrix is

$$P = \begin{bmatrix} 0 & 1/3 & 1/3 \\ 1/2 & 0 & 2/3 \\ 1/2 & 2/3 & 0 \end{bmatrix}.$$

It is easy to check that P^2 has entries which are strictly positive, so P is a regular stochastic matrix.

It is also easy to see that a rat can get from any room to any other room (including the one it starts from) through one or more moves.

To find the steady state vector we need to solve $(P - I)\mathbf{x} = \mathbf{0}$, that is we need to find the null space of $P - I$.

$$P - I = \begin{bmatrix} -1 & 1/3 & 1/3 \\ 1/2 & -1 & 2/3 \\ 1/2 & 2/3 & -1 \end{bmatrix}$$
$$\xrightarrow{\text{rref}} \begin{bmatrix} 1 & 0 & -2/3 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \end{bmatrix}$$

Hence if $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$, then $x_3 = t$ is free, $x_1 = \frac{2}{3}t$, $x_2 = t$. Since \mathbf{x} must be a probability vector, we need $1 = x_1 + x_2 + x_3 = \frac{8}{3}t$. Thus, $t = \frac{3}{8}$ and

$$\mathbf{x} = \begin{bmatrix} 1/4 \\ 3/8 \\ 3/8 \end{bmatrix}.$$

In the long run, the rat spends $\frac{1}{4}$ of its time in space 1, and $\frac{3}{8}$ of its time in each of the other two spaces.