

MAS275 Probability Modelling

3 Limiting behaviour of Markov chains

3.1 Renewal processes and Markov chains

Suppose a Markov chain (X_n) is started in a particular fixed state i . If it returns to i at some later (random) time, then, because of the Markov property and time homogeneity, the future behaviour will be as if the process had been started off in state i at this later time. Hence, returns to the starting state form a renewal process. In this renewal process,

$$u_n = p_{ii}^{(n)},$$

where $p_{ii}^{(n)}$ denotes the (i, i) element of the n -step transition matrix $P^{(n)} = P^n$. (Note that although $P^{(n)} = P^n$ it is not true in general that $p_{ii}^{(n)} = p_{ii}^n$.)

The “ f_n ” of the renewal process will then be the probability, starting in i , of the first return to state i being at time n , namely

$$P(X_1 \neq i, X_2 \neq i, \dots, X_{n-1} \neq i, X_n = i | X_0 = i).$$

We will write $f_{ii}^{(n)}$ for this.

Also define

$$f_{ii} = \sum_{n=1}^{\infty} f_{ii}^{(n)},$$

i.e. the probability, starting in state i , of ever returning to state i . If the renewal process is recurrent this will be 1, and the process will keep returning to i , and if it is transient then it will be strictly less than 1, and eventually, with probability 1, the process will visit state i for the last time and never return.

We can refer to a any state i of a Markov chain as being **transient** or **recurrent** if the corresponding renewal process is. Similarly, if state i is recurrent,

we can describe it as **positive recurrent** or **null recurrent** if the corresponding renewal process is, and we can define the **period** of state i as the same as the period of the corresponding renewal process. Again, we describe a state with period 1 as **aperiodic**.

These are called the **recurrence properties** of the states.

Example 16. *Simple random walk*

Example 17. *Mean recurrence time*

If we now take two different states, i and j say, and consider visits to state j having started in state i , then, by a similar argument to the above, we have a delayed renewal process, the delay being the time until the first visit to j .

We can identify the “ b_n ” and “ v_n ” of this delayed renewal process as

$$b_n = f_{ij}^{(n)} = P(X_1 \neq j, X_2 \neq j, \dots, X_{n-1} \neq j, X_n = j | X_0 = i)$$

(that is, the probability that the first visit to state j happens at time n , starting from state i) and

$$v_n = p_{ij}^{(n)}$$

where $p_{ij}^{(n)}$ is the (i, j) -th element of the n -step transition matrix.

Note that it may be possible, starting in state i , never to visit state j , in which case the delay distribution is defective: in fact

$$f_{ij} = \sum_{n=1}^{\infty} f_{ij}^{(n)}$$

is the probability, starting in state i , of ever visiting state j .

The probabilities $f_{ij}^{(n)}$ are called **first passage** probabilities.

Finite state spaces tend to be slightly easier to handle, as for example in the following result.

Theorem 8. *If the state space is finite then some of the states must be positive recurrent, and the rest, if there are any, must be transient.*

Proof. First, in a finite state space the states cannot all be transient, because, if they were, after the process had left each of the states for the last time there would be nowhere left for it to go.

If a state i is recurrent, then, starting in i , then consider V_j , the number of visits to another state, j say, before the first return to i . If q_{ij} is the probability of visiting j before returning to i , starting in i , and q_{ji} is the probability of visiting i before returning to j , then we can write

$$P(V_j = 0) = 1 - q_{ij}$$

and, for $n \geq 1$,

$$P(V_j = n) = q_{ij}(1 - q_{ji})^{n-1}q_{ji}.$$

The distribution of V_j has finite mean (by essentially the same calculations as for the geometric distribution) and so, as we are in a finite state space, the total time spent away from i before the first return has finite mean, i.e. the state i is positive recurrent. Hence there are no null recurrent states in a chain with a finite state space. \square

3.2 Communication

Two states i and j of a Markov chain are said to **communicate** with each other if it is possible, starting in i , to get to j (not necessarily at the next step) and similarly possible, starting in j , to get to i . Formally, there exist positive integers m and n such that $p_{ij}^{(m)} > 0$ and $p_{ji}^{(n)} > 0$.

This relationship is obviously **symmetric**, in the sense that if i and j communicate then so do j and i (which is implicit in the way we have defined communication). It is also not hard to see that if i communicates with j and also j communicates with k , then i communicates with k ; formally, this

follows most easily from the inequality

$$p_{ik}^{(m+r)} \geq p_{ij}^{(m)} p_{jk}^{(r)}$$

which is a straightforward consequence of the Chapman-Kolmogorov equations; if we can find m and r such that the right hand side is positive, then the left hand side must be positive. This property of the communication relationship is called **transitivity**.

If we also, by convention, define a state to communicate with itself, then we have a relationship which is **reflexive**.

A binary relationship between members of a set is called an **equivalence relation** if it is reflexive, symmetric and transitive. The most important consequence of having all three of these properties is that the set in question, in this case the state space of the Markov chain, is partitioned into **equivalence classes**, such that any two members of the same class communicate with each other, and any two members of two different classes do not communicate with each other. These are called **communicating classes** or simply classes.

A Markov chain is called **irreducible** if all its states communicate with each other, i.e. they form a single large class.

Example 18. *Irreducible Markov chains*

Example 19. *Gambler's ruin*

3.3 Solidarity of recurrence properties within classes

The solidarity property can be very simply stated as follows.

Theorem 9. *All states within a communicating class share the same recurrence properties.*

This theorem means, in particular, that if the Markov chain is irreducible then all states share the same recurrence properties, and it is meaningful to

refer to the recurrence properties of the whole Markov chain. In any event, we can refer to the recurrence properties of classes as well as those of individual states.

Proof. We will prove this for recurrence and transience and for period.

It is clearly sufficient to show that any two members of a class share the same recurrence properties. So let i and j be two members of a class, and let m and n be such that $p_{ij}^{(m)} > 0$ and $p_{ji}^{(n)} > 0$. The proof of the theorem depends on the following inequality: for any $r = 0, 1, 2, \dots$,

$$p_{ii}^{(m+r+n)} \geq p_{ij}^{(m)} p_{jj}^{(r)} p_{ji}^{(n)}. \quad (1)$$

To see why this holds, note that to return to state i after $m + r + n$ steps, starting in state i , one possibility is to visit state j after m steps, then to visit state j again after a further r steps, and then to go back to state i after a further n steps.

The case $r = 0$ can be included in (1) if we adopt the natural convention that $p_{jj}^{(0)} = 1$ for all j and $p_{ij}^{(0)} = 0$ if $i \neq j$.

The important thing about the inequality (1) is that the first and third factors on the right hand side are positive and fixed, by definition of m and n , and this turns out to be enough to tell us all we need to know about the relationship between $p_{ii}^{(m+r+n)}$ and $p_{jj}^{(r)}$.

Suppose first that state i is transient. Then, from renewal theory,

$$\sum_{r=0}^{\infty} p_{ii}^{(r)} < \infty.$$

Missing out the first $m + n$ terms,

$$\sum_{r=0}^{\infty} p_{ii}^{(m+r+n)} < \infty.$$

Then using our key inequality above,

$$\sum_{r=0}^{\infty} p_{jj}^{(r)} \leq (p_{ij}^{(m)} p_{ji}^{(n)})^{-1} \sum_{r=0}^{\infty} p_{ii}^{(m+r+n)} < \infty.$$

Hence state j is transient. Similarly, if state j is transient then state i is transient.

Now consider periodicity. Suppose state i has period d . Then the case $r = 0$ of our key inequality gives

$$p_{ii}^{(m+n)} \geq p_{ij}^{(m)} p_{ji}^{(n)} > 0$$

and so $m + n$ must be divisible by d . Then for all r which are not divisible by d , $p_{ii}^{(m+r+n)} = 0$, and so by our key inequality, $p_{jj}^{(r)} = 0$. It follows that the period of state j is divisible by d , i.e. the period of state i . Similarly, the period of state i is divisible by the period of state j . Hence both states have the same period.

The theorem also applies to positive versus null recurrence; this is a bit harder to prove based on what we have seen so far. A proof is given in supplementary notes. \square

We can immediately conclude the following:

Corollary 10. *An irreducible chain with finite state space has all its states positive recurrent.*

Proof. By Theorem 8, some of the states must be positive recurrent. Hence by solidarity (Theorem 9) they all are. \square

A class C is called **closed** if, once entered, the probability of leaving is zero. It is also sometimes known as an **absorbing class**. Formally, C is closed if $p_{ij} = 0$ for all $i \in C$ and $j \notin C$. In the irreducible case there is a single class $C = S$, which is automatically closed.

If a class is not closed, then, once it has been left, there is no possibility of returning to it; if there were such a possibility then there would be at least one state outside the class which communicates with it, which is a contradiction.

If the state space S is finite, then the concept of a closed class is equivalent to that of a recurrent class, and that of a class which is not closed is equivalent to that of a transient class. This is because, if a class is recurrent, then it is impossible to leave it, because if it were possible and it happened, then return to the class could not take place, and this contradicts recurrence; conversely if a class is transient then each state of the class is only visited finitely many times and so the class must eventually be left.

If the state space is infinite, there are more possibilities. For example, in the simple random walk with $p \neq \frac{1}{2}$, all states are transient but they form a single closed class.

Example 20. *Finding classes and recurrence properties*

3.4 Limiting/equilibrium behaviour

The aim of this section is to investigate how the distribution of the state a Markov chain evolves with time. In particular we will see that many chains evolve towards an equilibrium behaviour described by a stationary distribution.

We will first of all consider one case where this does not happen: when the chain is transient.

Theorem 11. *Assume we have a Markov chain (X_n) which is irreducible and transient (which necessarily means that its state space S is infinite). Then, for each state $i \in S$, $P(X_n = i) \rightarrow 0$ as $n \rightarrow \infty$. Furthermore, there is no stationary distribution.*

Proof. The transience of each state i tells us that there is some random time T_i , finite with probability 1, when we visit i for the last time. (Set $T_i = 0$ if

we never visit i .) We can then write

$$P(X_n = i) = P(X_n = i | T_i \geq n)P(T_i \geq n) + P(X_n = i | T_i < n)P(T_i < n).$$

But $P(X_n = i | T_i < n) = 0$ by the definition of T_i , and $P(T_i \geq n) \rightarrow 0$ as $n \rightarrow \infty$ because of T_i being finite. Hence $P(X_n = i) \rightarrow 0$ as $n \rightarrow \infty$.

If a stationary distribution $\boldsymbol{\pi}$ existed, then we could start the chain in its stationary distribution. Then $P(X_n = i) = \pi_i$ for all n , by stationarity, so, as $P(X_n = i) \rightarrow 0$ as $n \rightarrow \infty$, we must have $\pi_i = 0$ for all i , but this cannot work as stationary distributions must sum to 1. \square

We now consider the case where the Markov chain is irreducible and aperiodic, and we know that a stationary distribution exists. By Theorem 11 we know that the chain must be recurrent.

Theorem 12. *If a Markov chain (X_n) is irreducible and aperiodic and has a stationary distribution $\boldsymbol{\pi}$ then the distribution of the state of the Markov chain at time n converges to $\boldsymbol{\pi}$, i.e.*

$$\boldsymbol{\pi}^{(n)} \rightarrow \boldsymbol{\pi}$$

as $n \rightarrow \infty$.

Proof. We can prove this theorem using a “coupling” argument. The idea of this argument is to consider a second Markov chain, (Y_n) , which has the same transition matrix as the original chain (X_n) but has initial distribution given by the stationary distribution $\boldsymbol{\pi}$. The two Markov chains run independently of each other. Then $P(X_n = i) = \pi_i^{(n)}$, while $P(Y_n = i) = \pi_i$.

Consider what happens if at some time m , the two chains are in the same state: $X_m = Y_m$. Then, by the Markov property, what happened before time m does not affect the probabilities of what happens after time n , so for any $n > m$ the conditional probabilities that $X_n = i$ and that $Y_n = i$ will be the same.

More formally, let T be the first time at which the two chains are in the same state (note that, a priori, T could be infinite) and note that the Markov property shows that

$$P(X_n = i|T \leq n) = P(Y_n = i|T \leq n).$$

Now,

$$\begin{aligned} P(X_n = i) - P(Y_n = i) &= (P(X_n = i|T \leq n) - P(Y_n = i|T \leq n))P(T \leq n) \\ &\quad + (P(X_n = i|T > n) - P(Y_n = i|T > n))P(T > n) \\ &= (P(X_n = i|T > n) - P(Y_n = i|T > n))P(T > n), \end{aligned}$$

so if we can show that $P(T > n)$ converges to zero as $n \rightarrow \infty$ (which is the same thing as saying that the chains will meet eventually with probability 1) we will have that $P(X_n = i) - P(Y_n = i) = \pi_i^{(n)} - \pi_i \rightarrow 0$ as $n \rightarrow \infty$, which will give the result.

To show that the chains will meet eventually with probability 1, consider the bivariate Markov chain on $S \times S$ obtained by considering the two chains together, so that its state at time n is (X_n, Y_n) .

By the irreducibility and aperiodicity of our original chain, we know that for any choice of i and j it is the case that for n large enough $p_{ij}^{(n)} > 0$. So for any pair of states (i_1, i_2) and (j_1, j_2) of the bivariate Markov chain, if n is large enough both $p_{i_1, j_1}^{(n)}$ and $p_{i_2, j_2}^{(n)}$ will be positive, showing that it is possible for the bivariate Markov chain to get from (i_1, i_2) to (j_1, j_2) . Hence the bivariate Markov chain is irreducible. (Note that aperiodicity is crucial to this argument.)

It also has a stationary distribution, given by

$$P(X_n = i, Y_n = j) = \pi_i \pi_j$$

by independence. An irreducible Markov chain with a stationary distribution cannot be transient, by Theorem 11. Hence it is recurrent, and so all states are visited, with probability 1. This includes states where the two co-ordinates are the same, so this shows that the chains (X_n) and (Y_n) will meet, with probability 1, and thus completes the proof of the theorem. \square

Corollary 13. *If a Markov chain (X_n) is irreducible and aperiodic, and has a stationary distribution $\boldsymbol{\pi}$, then that stationary distribution is unique.*

Proof. If another stationary distribution $\boldsymbol{\psi}$ exists, then we can start a chain with the same transition matrix with initial distribution $\boldsymbol{\pi}^{(0)} = \boldsymbol{\psi}$. By stationarity we have $\boldsymbol{\pi}^{(n)} = \boldsymbol{\psi}$ for all n and hence $\lim_{n \rightarrow \infty} \pi_i^{(n)} = \psi_i$, but we have shown above that this must be π_i . \square

Now, we show that if we have an irreducible and positive recurrent Markov chain we can construct a stationary distribution. For each state j , let μ_j be the expected time until the first return to state j , given that the chain starts there. Write $\boldsymbol{\pi}$ for the row vector with entry j being $\frac{1}{\mu_j}$, for $j \in S$. The following theorem gives the main results of this section.

Theorem 14. *Assume the Markov chain is irreducible, aperiodic and positive recurrent. Then*

- (a) *The distribution $\boldsymbol{\pi}$ is a stationary distribution, and is unique.*
- (b) *The distribution of the state of the Markov chain at time n converges to $\boldsymbol{\pi}$, i.e.*

$$\boldsymbol{\pi}^{(n)} \rightarrow \boldsymbol{\pi}$$

as $n \rightarrow \infty$.

Proof. Select a state i of the Markov chain, and assume that the chain starts in i (or start a chain in i which has the same transition matrix). Because the chain is positive recurrent, it will return to i , with probability 1, at a random time T with finite mean μ_i .

For any state j , let V_j be the number of visits to j in times $1, 2, \dots, T$. By the definition of T , $V_i = 1$, and it is also the case that $E\left(\sum_{j \in S} V_j\right) = E(T) = \mu_i$. Also, for $n \geq 1$, let E_n be the event that $T \geq n$, i.e. that the chain does not visit i between times 1 and $n - 1$ inclusive. Note that $E(V_j)$ is the sum of the probabilities $P(X_n = j, E_n)$ over all n .

Define a row vector $\boldsymbol{\psi}$ by $\psi_j = E(V_j)$; now

$$\begin{aligned}
(\boldsymbol{\psi}P)_k &= \sum_{j \in S} \psi_j p_{jk} \\
&= \sum_{j \in S} E(V_j) p_{jk} \\
&= \sum_{j \in S, j \neq i} \sum_{n=1}^{\infty} P(X_n = j, E_n) p_{jk} + p_{ik} \\
&= \sum_{n=1}^{\infty} \sum_{j \in S, j \neq i} P(X_n = j, E_n) p_{jk} + p_{ik} \\
&= \sum_{n=1}^{\infty} P(X_{n+1} = k, E_{n+1}) + p_{ik} \\
&= \sum_{n=2}^{\infty} P(X_n = k, E_n) + p_{ik} \\
&= \sum_{n=1}^{\infty} P(X_n = k, E_n) \\
&= E(T_k) = \psi_k.
\end{aligned}$$

(Note that it is possible to reverse the order of summation when terms in the sum are non-negative.)

Hence $\boldsymbol{\psi}P = \boldsymbol{\psi}$, so $\boldsymbol{\psi}$ is a left eigenvector of the transition matrix with eigenvalue 1. All its elements are non-negative, and they sum to μ_i , so we can obtain a stationary distribution $\boldsymbol{\pi}$ by defining

$$\pi_j = \frac{\psi_j}{\mu_i} = \frac{E(V_j)}{\mu_i}.$$

Because $V_i = 1$, we must have $\pi_i = 1/\mu_i$. However, our choice of i was arbitrary, and we know from Corollary 13 that the stationary distribution is unique. Hence π_i must be $1/\mu_i$ for all i , and this is a stationary distribution. So applying Theorem 12 gives the result. \square

Note that because we know that π is a unique stationary distribution, it is usually easiest to find it by solving the equilibrium equations of section 1.7, and then to apply the above results.

Example 21. *Modelling the game of Monopoly*

Finally we note that the results of Theorem 11 (that for any state i $P(X_n = i) \rightarrow 0$ as $n \rightarrow \infty$, and that there are no stationary distributions) apply to irreducible null recurrent chains as well as irreducible transient chains. This is a bit harder to prove, and we omit the proof; however a proof that there are no stationary distributions is given in supplementary notes.

To summarise:

- If a chain has finite state space and is irreducible and all states are aperiodic, then it is positive recurrent by Corollary 10, and by Theorem 14 it has a unique stationary distribution, to which the distribution of the chain at time n converges as $n \rightarrow \infty$.
- If a chain with infinite state space is irreducible, aperiodic and positive recurrent, then again by Theorem 14 it has a unique stationary distribution, to which the distribution of the chain at time n converges as $n \rightarrow \infty$.
- If a chain with infinite state space is irreducible and is either transient or null recurrent, then it does not have a stationary distribution and for any state i $P(X_n = i) \rightarrow 0$ as $n \rightarrow \infty$.

3.5 Non-irreducible and periodic chains

Consider a Markov chain which has a closed class C which is positive recurrent and not the whole state space. Then the behaviour of the process if it starts inside C is that of a Markov chain in which the states outside C have been removed from the state space, and the corresponding rows and columns of

the transition matrix have been deleted. This “reduced” Markov chain will be irreducible and will have a stationary distribution, and if this distribution is extended to the original state space by assigning probability zero to each of the states outside C , we have a stationary distribution for the original Markov chain.

It follows that if there are two or more positive recurrent closed classes, then the stationary distribution is not unique: there will be, at least, one corresponding to each of these classes. Furthermore, any “mixture” of the form

$$\alpha\pi + (1 - \alpha)\psi$$

where π and ψ are stationary distributions and $0 < \alpha < 1$, will also be a stationary distribution, and so there will be a continuum of them. In the case of a finite state space, this situation corresponds to the space of left eigenvectors of the transition matrix corresponding to eigenvalue 1 being more than one-dimensional, and the eigenvalue 1 having multiplicity greater than 1.

Example 22. *A non-irreducible chain*

Periodicity is a nuisance when it comes to looking at limiting behaviour; the following example illustrates how things can differ from the aperiodic case.

Example 23. *A periodic chain*

3.6 The renewal theorem

It is possible to apply the ideas of long-run behaviour of Markov chains to renewal theory. Given a renewal process, recall the following notation:

- f_k is the probability that a given inter-renewal time takes the value k
- u_k is the probability that there is a renewal at time k
- d is the period of the renewal process
- μ is the mean inter-recurrence time, if the process is positive recurrent

Also define K to be the largest value of k such that f_k is non-zero (i.e. the largest possible inter-renewal time) if this is finite.

Given a renewal process, let X_n be the number of time steps from time n until the next renewal; this is called the **forward recurrence time**. (Set $X_n = 0$ if there is a renewal at time n .) Then $P(X_{n+1} = i - 1 | X_n = i) = 1$ unless $i = 0$, and for $i \geq 0$ we have $P(X_{n+1} = i | X_n = 0) = f_{i+1}$. (To see the latter, note that $X_n = 0$ means there is a renewal at time n and that $X_{n+1} = i$ means that the next renewal after the one at time n occurs at time $n + i + 1$, which has probability f_{i+1} .) By the construction of the renewal process, the process (X_n) satisfies the Markov property, so is a Markov chain with transition matrix

$$\begin{pmatrix} f_1 & f_2 & f_3 & f_4 & \dots \\ 1 & 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & 0 & \dots \\ 0 & 0 & 1 & 0 & \dots \\ 0 & 0 & 0 & 1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

and state space $\{0, 1, 2, \dots, K-1\}$ if K is finite, and the non-negative integers \mathbb{N}_0 otherwise.

This chain is irreducible, and if the renewal process is positive recurrent, then the Markov chain is too (we can see this by considering returns to state zero) and it has a stationary distribution given by

$$\pi_i = \frac{1}{\mu}(1 - f_1 - \dots - f_i)$$

for $i = 0, 1, 2, \dots$. Similarly, if the renewal process is transient or null recurrent, then the Markov chain will be too. We can use the results on Markov chains to deduce the following theorem.

Theorem 15. (*Erdős, Feller, Pollard*) (*The Discrete Renewal Theorem*)

If a renewal process is recurrent and has period d and finite mean inter-

renewal time μ , then

$$u_{nd} \rightarrow \frac{d}{\mu} \text{ as } n \rightarrow \infty.$$

In all other cases,

$$u_n \rightarrow 0 \text{ as } n \rightarrow \infty.$$

Proof. We observe that $u_n = P(X_n = 0)$, as both are the probability that we have a renewal at time n . So, if the renewal process is aperiodic and positive recurrent, we can immediately deduce $u_n \rightarrow \pi_0 = \frac{1}{\mu}$ from the results on Markov chains. If the renewal process is transient or null recurrent, then we can similarly deduce $u_n \rightarrow 0$ as $n \rightarrow \infty$.

We can handle the periodic case by defining a new aperiodic renewal process which has renewals at time n if our periodic renewal process has a renewal at time nd , and applying the theorem to that. \square

Example 24. *Bernoulli trials with blocking*