

Markov chains

©University of Bristol, 2012

1 Discrete time Markov chains

Example: A drunk is walking home from the pub. There are n lampposts between the pub and his home, at each of which he stops to steady himself. After every such stop, he may change his mind about whether to walk home or turn back towards the pub, independent of all his previous decisions. He moves homeward with probability p and pubward with probability $1 - p$, stopping when he reaches either. How do you describe his trajectory?

Let us look at the times that he reaches either a lamppost or the pub or home. The lampposts are numbered 1 through n , the pub is at location 0 and his home is denoted $n + 1$. At each time $t = 0, 1, 2, \dots$, he may be at any of these locations, and we'll let X_t denote his location at time t . We want to know $P(X_{t+1} = x | X_0, \dots, X_t)$ for all x and t .

From the description above, it should be clear that, conditional on the drunk's trajectory up to time t , his position at the next time only depends on X_t , i.e.,

$$P(X_{t+1} = x | X_0, \dots, X_t) = P(X_{t+1} = x | X_t). \quad (1)$$

This is called the Markov property, and a process having this property is called a Markov process or Markov chain. By repeatedly using (1), we get

$$\begin{aligned} P(X_{t+1} = x_1, \dots, X_{t+n} = x_n | X_0, \dots, X_t) \\ = P(X_{t+1} = x_1, \dots, X_{t+n} = x_n | X_t), \end{aligned} \quad (2)$$

for all t and n , and all possible states X_0, \dots, X_{t+n} . In words, it says that the future evolution of the process is *conditionally* independent of the past given the present.

Formally, a discrete-time Markov chain on a state space S is a process X_t , $t = 0, 1, 2, \dots$ taking values in S which satisfies (1) or, equivalently, (2).

In all the examples we see in this course the state space S will be discrete (usually finite, occasionally countable). There is a theory of Markov chains for general state spaces as well, but it is outside our scope.

Other examples

1. A toy model of the weather with, e.g., 3 states, Sunny, Cloudy, Rainy, and transition probabilities between them.
2. A model of language with transition probabilities for, say, successive letters in a word, or successive words in a sentence. A 1-step Markov model may be simplistic for this. Suppose that, instead of having memory 1, a process has some fixed, finite memory. Can it still be modelled as a Markov chain?
3. The Ethernet protocol: A computer that has a packet to transmit over the local area network starts in back-off stage 0 and attempts to transmit it. Every time it fails (because the packet collides with another computer that is trying to transmit at the same time), it increments its back-off counter. The back-off counter tells it how long to wait before attempting again. Once the packet is successfully transmitted, the back-off counter is reset to 0. In this example, the choice of state is important. If you choose the state as the value of the back-off counter in each time step, then it is not Markovian (because its future evolution depends on how long that back-off counter has had its present value, not just what this value is). However, if you consider the “embedded chain” of the counter values just after the back-off counter changes, then this chain is Markovian.

What does it take to fully describe a Markov chain? Clearly, it suffices to describe all the conditional probabilities in (2), for all t and n and all possible combinations of states. In fact, it suffices to just specify the *one step* transition probabilities $P(X_{t+1} = y | X_t = x)$ for all t , and all $x, y \in S$. Why is this sufficient?

First note that we can represent all the one step transition probabilities in the form of a matrix $P(t)$ with entries $P_{xy}(t) = P(X_{t+1} = y | X_t = x)$. From this, we can compute

$$\begin{aligned}
 & P(X_{t+2} = z, X_{t+1} = y | X_t = x) \\
 &= P(X_{t+1} = y | X_t = x) P(X_{t+2} = z | X_{t+1} = y, X_t = x) \\
 &= P(X_{t+1} = y | X_t = x) P(X_{t+2} = z | X_{t+1} = y) = P_{xy}(t) P_{yz}(t+1),
 \end{aligned}$$

and so on.

Thus, to describe a Markov process, it suffices to specify its initial distribution μ on S (which may be unit mass on a single state on S), and all the one step transition probability matrices $P(t)$, $t = 0, 1, 2, \dots$. We will typically be interested in the case in which $P(t) = P$ for all t . In this case, $P(X_{t+s} = y | X_s = x)$ is the same as $P(X_t = y | X_0 = x)$ for any s, t, x and y . Such a Markov chain is called time homogeneous.

It is easy to see that, for time homogeneous Markov chains, $P(X_t = y | X_0 = x) = (P^t)_{xy}$. (Verify this explicitly for $t = 2$). In other words, the t -step transition probability matrix is given by the t^{th} power of the one step transition matrix.

The transition probability matrix P has the property that all its entries are non-negative and all its row sums are 1, i.e., $P_{xy} \geq 0$ for all $x, y \in S$, and $\sum_{y \in S} P_{xy} = 1$ for all $x \in S$. A matrix having these properties is called a stochastic matrix.

Classification of states

Consider a Markov chain $\{X_t, t \geq 0\}$ on a discrete state space S . A state $x \in S$ is said to be recurrent if, starting from this state, the chain returns to it with probability 1, i.e.,

$$P(\exists t \geq 1 : X_t = x | X_0 = x) = 1.$$

It is called transient if this return probability is strictly smaller than 1. If the state space S is finite, then the mean time to return to a recurrent state is also finite, but this need not be the case if $|S|$ is infinite. For example, consider the simple symmetric random walk on the integers, specified by the transition probabilities,

$$P(X_{t+1} = x + 1 | X(t) = x) = P(X_{t+1} = x - 1 | X(t) = x) = \frac{1}{2}, \quad x \in \mathbb{Z}.$$

It can be shown that the probability of returning to 0 (say) is 1, but that the mean return time is infinite.

State j is said to be accessible from state i if it is possible to go from i to j , i.e., $P_{ij}^t > 0$ for some $t \geq 0$. In particular, each state is accessible from itself since $P_{ii}^0 = 1$. States i and j are said to communicate if i is accessible from j and j from i . Note that this defines an equivalence relation. (A relation R is said to be an equivalence relation if it is (a) symmetric: xRy implies

yRx for all x and y , (b) reflexive: xRx for all x , and (c) transitive: xRy and yRz together imply xRz for all x, y and z .) Hence, it partitions the state space into equivalence classes, which are called communicating classes. (Sets A_1, A_2, \dots are said to form a partition of A if $\cup_{i=1}^{\infty} A_i = A$ and $A_i \cap A_j = \emptyset$ for all $i \neq j$.) Note that states in the same communicating class have to all be transient or all be recurrent.

In our example of the drunkard's walk, his home and the pub each form a recurrent communicating class (a trivial one, consisting of a single state), while all intermediate states (lampposts) are transient and form a single communicating class, since there is a positive probability of going from any lamppost to any other.

A finite state Markov chain eventually has to leave the set of transient states and end up in one or other recurrent communicating class. A communicating class is said to be closed if the probability of leaving it is zero. Every recurrent communicating class is closed. The converse need not be true in infinite-state chains. Finally, the Markov chain is said to be irreducible if it consists of a single communicating class. In that case, we can talk of the chain itself being transient or recurrent. If it is a finite-state chain, it necessarily has to be recurrent.

Invariant distributions

Suppose we observe a finite-state Markov chain over a long period of time. What can we say about its behaviour? If it starts in a transient state, it will spend some time in a transient communicating class before entering one or another recurrent class, where it will remain from then on. Thus, as far as long-term behaviour is concerned, it is enough to look at the recurrent classes. Moreover, each such communicating class can be studied in isolation since the Markov chain never leaves such a class upon entering it. Finally, the Markov chain within a single class is irreducible.

Hence, we restrict our attention in what follows to irreducible Markov chains. Consider such a chain on a finite state space S , and let P denote its transition probability matrix. We have:

Theorem 1 *There is a unique probability distribution π on the state space S such that $\pi P = \pi$.*

It is easy to see that every stochastic matrix has the all-1 vector as a right eigenvector corresponding to the eigenvalue 1. The above theorem says

that the corresponding left eigenvector is also non-negative, and that there is only such eigenvector corresponding to an eigenvalue of 1 if the matrix corresponds to an irreducible chain. (The assumption of irreducibility is necessary for uniqueness. The identity matrix is a stochastic matrix, but all its eigenvalues are 1, and there are multiple non-negative eigenvectors.) We won't give a proof of this theorem, but one way is to use the Perron-Frobenius theorem for non-negative matrices (of which stochastic matrices are a special case).

The probability distribution π solving $\pi P = \pi$ is called the *invariant distribution* or *stationary distribution* of the Markov chain. The reason is that, if the chain is started in this distribution, then it remains in this distribution forever. More precisely, if $P(X_t = x) = \pi_x$ for all $x \in S$, then, for all $y \in S$,

$$\begin{aligned} P(X_{t+1} = y) &= \sum_{x \in S} P(X_t = x, X_{t+1} = y) \\ &= \sum_{x \in S} P(X_t = x) P(X_{t+1} = y | X(t) = x) \\ &= \sum_{x \in S} \pi_x P_{xy} = (\pi P)_y = \pi(y). \end{aligned}$$

The invariant distribution describes the long-run behaviour of the Markov chain in the following sense.

Theorem 2 (Ergodic theorem for Markov chains) *If $\{X_t, t \geq 0\}$ is a Markov chain on the state space S with unique invariant distribution π , then*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=0}^{n-1} 1(X_t = x) = \pi(x) \quad \forall x \in S,$$

irrespective of the initial condition. The convergence holds almost surely.

The above theorem holds for both finite and countable state spaces, assuming the invariant distribution exists, which it may fail to do in the countable case, even if the Markov chain is irreducible and recurrent. An example is the simple symmetric random walk, which has no invariant distribution.

The theorem says that $\pi(x)$ specifies the fraction of time that the Markov chain spends in state x in the long run. It states a law of large numbers for the mean of the random variables $1(X_t = x)$, for any given $x \in S$. Note

that this is a sequence of random variables which aren't independent or identically distributed, so this theorem is not just a corollary of the law of large numbers in the iid case.

The two theorems above tell us that irreducible finite-state Markov chains have unique invariant distributions, and that these can be related to the average time spent by the Markov chain in each state. But how do we compute the invariant distribution?

In general, we can do this by solving a system of linear equations. Denote by n the total number of states. Then, the matrix equation $\pi = \pi P$ or, equivalently, $\pi(I - P) = 0$ corresponds to a system of n equations in n variables. However, only $n - 1$ of these equations are linearly independent, which is reflected in the fact that there isn't a single solution, but a one-dimensional subspace of solutions: if π is a solution of $\pi(I - P) = 0$, then so is any multiple of π . But the particular solution we are interested in is also a probability distribution, which means that it must satisfy the additional equation $\pi \mathbf{1} = 1$, where $\mathbf{1}$ denotes the all-ones vector of size n . The system of equations $\pi P = \pi$ can be written out as

$$\pi_j = \sum_{i \in S} \pi_i p_{ij}. \quad (3)$$

These are called the global balance equations for the Markov chain.

Reversibility

An irreducible Markov chain is said to be reversible if the probability law of the chain is the same observed either forwards or backwards in time. More precisely, all finite-dimensional distributions are the same in forward and reverse time:

$$\begin{aligned} &P(X_{t_1} = x_1, X_{t_2} = x_2, \dots, X_{t_k} = x_k) \\ &= P(X_{s-t_1} = x_1, X_{s-t_2} = x_2, \dots, X_{s-t_k} = x_k), \end{aligned} \quad (4)$$

for all $k \in \mathbb{N}$, $s, t_1, t_2, \dots, t_k \in \mathbb{Z}$ and $x_1, \dots, x_k \in S$.

We first note that a reversible Markov chain must be stationary. To see this, observe from (4) that

$$\begin{aligned} &P(X_{t_1+s} = x_1, X_{t_2+s} = x_2, \dots, X_{t_k+s} = x_k) \\ &= P(X_{-t_1} = x_1, X_{-t_2} = x_2, \dots, X_{-t_k} = x_k). \end{aligned} \quad (5)$$

But (4) holds for all $s \in \mathbb{Z}$. In particular, it holds for $s = 0$, which means that

$$\begin{aligned} &P(X_{t_1} = x_1, X_{t_2} = x_2, \dots, X_{t_k} = x_k) \\ &= P(X_{-t_1} = x_1, X_{-t_2} = x_2, \dots, X_{-t_k} = x_k). \end{aligned} \quad (6)$$

By (5) and (6), we have shown that

$$\begin{aligned} &P(X_{t_1+s} = x_1, X_{t_2+s} = x_2, \dots, X_{t_k+s} = x_k) \\ &= P(X_{t_1} = x_1, X_{t_2} = x_2, \dots, X_{t_k} = x_k), \end{aligned}$$

for arbitrary $s \in \mathbb{Z}$, which is the definition of stationarity.

Theorem 3 *A stationary Markov chain is reversible if and only if there exists a collection of non-negative numbers π_j , $j \in S$, summing to unity which solve the detailed balance equations*

$$\pi_j p_{jk} = \pi_k p_{kj}, \quad \forall j, k \in S. \quad (7)$$

When there exists such a collection (π_j) , it is the invariant distribution of the Markov chain.

Proof. Suppose first that the process is reversible. Then, as seen above, it must be stationary, so $P(X_t = j)$ doesn't depend on t . Define $\pi_j = P(X_0 = j)$. Clearly $\pi_j \geq 0$ and $\sum_{j \in S} \pi_j = 1$. Since the process is reversible,

$$P(X_t = j, X_{t+1} = k) = P(X_{t+1} = j, X_t = k)$$

for all $j, k \in S$, i.e.,

$$\pi_j p_{jk} = \pi_k p_{kj}, \quad \forall j, k \in S.$$

Conversely, suppose there is a collection of non-negative π_j summing to unity and satisfying (7). Then, summing (7) over $k \in S$, we get

$$\pi_j = \sum_{k \in S} \pi_k p_{kj} \quad \forall j \in S.$$

Since these are the global balance equations and π is a probability vector solving them, π is the invariant distribution of the chain. Thus, if the Markov chain is stationary, then

$$P(X_t = j_0, X_{t+1} = j_1, \dots, X_{t+m} = j_m) = \pi(j_0) p_{j_0, j_1} \cdots p_{j_{m-1}, j_m},$$

and

$$P(X_\tau = j_m, X_{\tau+1} = j_{m-1}, \dots, X_{\tau+m} = j_0) = \pi_{j_m} p_{j_m, j_{m-1}} \cdots p_{j_1, j_0},$$

for arbitrary t, τ and m , and arbitrary $j_0, \dots, j_m \in S$. Now, by (7), the two right hand sides above are equal. Thus, taking $s = t + \tau + m$, we get

$$\begin{aligned} & P(X_t = j_0, X_{t+1} = j_1, \dots, X_{t+m} = j_m) \\ &= P(X_{s-t} = j_0, X_{s-(t+1)} = j_1, \dots, X_{s-(t+m)} = j_m), \end{aligned}$$

for all j_0, \dots, j_m . In other words, $X_t, X_{t+1}, \dots, X_{t+m}$ have the same joint distribution as $X_{s-t}, X_{s-(t+1)}, \dots, X_{s-(t+m)}$. This completes the proof. \square

If we know that a certain Markov chain is reversible, then it is much easier to compute the invariant distribution; we only need to solve the system of equation $\pi_x P_{xy} = \pi_y P_{yx}$ for all $x, y \in S$ along with the normalisation $\sum_{x \in S} \pi_x = 1$.

But this seems to bring up a circularity. How do we know that a Markov chain is reversible before computing its invariant distribution, since reversibility is defined in terms of the invariant distribution? The answer is that we don't, but if we optimistically assume it to be the case and try to solve for π , then we will find a solution if and only if the Markov chain is reversible. (Otherwise, we will come across an inconsistency.) We now illustrate this with an example, which is in fact on a countable state space.

Birth and death Markov chains

Let $\{X_t, t \geq 0\}$ be a Markov chain on the state space $\{0, 1, 2, \dots\}$ with transition probabilities given by $P_{k, k+1} = p_k$ and $P_{k, k-1} = q_k = 1 - p_k$ for all $k \geq 1$, while $P_{01} = 1$.

Let us suppose that this Markov chain is reversible and try to compute its invariant distribution. Defining $p_0 = 1$, we obtain the equations

$$\pi_k q_k = \pi_{k-1} p_{k-1}, \quad k \geq 1$$

These equations have the non-negative solution

$$\pi_k = \pi_0 \frac{\prod_{j=0}^{k-1} p_j}{\prod_{j=1}^k q_j}.$$

The solution can be normalised to a probability distribution by a suitable choice of π_0 provided that

$$\sum_{k=1}^{\infty} \frac{\prod_{j=0}^{k-1} p_j}{\prod_{j=1}^k q_j} < \infty. \quad (8)$$

For an arbitrary sequence p_n , $n = 0, 1, 2, \dots$, it is hard to tell whether the condition in (8) is satisfied. One special case is when $p_n = p$ and $q_n = q = 1 - p$ for all $n \geq 1$. In that case, we can rewrite (8) as

$$\sum_{k=1}^{\infty} \left(\frac{p}{1-p} \right)^k < \infty,$$

which is satisfied if and only if $p < 1 - p$, i.e., $p < 1/2$. Under this condition, the Markov chain is reversible with unique invariant distribution

$$\pi_k = (1 - \rho)\rho^k \text{ where } \rho = \frac{p}{1-p} < 1.$$

If $p > 1 - p$, then the Markov chain is transient and escapes to infinity. If $p = 1 - p$, we have the simple symmetric random walk, but reflected at the origin. In this case, the Markov chain is recurrent but fails to have an invariant distribution. (The mean return time to any state is infinite. This situation is called null-recurrence.)

To summarise the above discussion, we have shown that birth-death chains are always reversible provided they have an invariant distribution, i.e., if they are not transient or null-recurrent. If we were to restrict the birth-death chain to a finite state space by having a reflecting boundary at some level n (so that $P_{n,n-1} = 1$), then the resulting finite-state chain would always be reversible.

2 Continuous time Markov chains

Consider a continuous time process X_t , $t \geq 0$ on a discrete state space S . The process is called a continuous time Markov chain (CTMC) or a Markov process on this state space if

$$P(X_{t+s} = y | \{X_u, u \leq s\}) = P(X_{t+s} = y | X_s), \quad (9)$$

for all states y and all $s, t \geq 0$. In words, this again says that the future is *conditionally* independent of the past given the present.

If the above probability only depends on t and X_s but not on s , then we say that the Markov chain is time homogeneous. In that case, we can represent the conditional probabilities in (9) in the form of a matrix $P(t)$ with entries

$$P_{xy}(t) = P(X_{t+s} = y | X(s) = x) = P(X_t = y | X_0 = x).$$

(Warning: I have used notation somewhat inconsistently. $P(t)$ doesn't correspond to the same thing as what I called $P(t)$ in the discrete time setting. There, it referred to the one-step transition probabilities *at time step* t in a time-inhomogeneous chain. Here, I am using it for transition probabilities over a *period of length* t (or t time steps in the discrete analogue) for a time-homogeneous chain.)

Whereas in the discrete time case, if the Markov chain was time homogeneous, its transition probabilities could be described by a single matrix P , here we seem to need a separate matrix $P(t)$ for each time interval t . The matrices $\{P(t), t \geq 0\}$ satisfy the following properties:

$$P(0) = I, \quad P(t+s) = P(t)P(s). \quad (10)$$

(Why?)

Poisson process: This is a counting process $N_t, t \geq 0$, (an integer-valued process which counts the number of 'points' or 'events' up to time t). We say that a counting process is a Poisson process of rate (or intensity) λ if the following hold:

- [1] The random variable $N_{t+s} - N_t$ is independent of $\{N_u, 0 \leq u \leq t\}$, for all $s, t \geq 0$.
- [2] The random variable $N_{t+s} - N_t$ has a Poisson distribution with mean λs , i.e.,

$$P(N_{t+s} - N_t = k) = \frac{(\lambda s)^k}{k!} e^{-\lambda s}, \quad k = 0, 1, 2, \dots \quad (11)$$

Property 2 above can be equivalently restated in either of the following two ways:

- [2a] For all $t \geq 0$,

$$\begin{aligned} P(N_{t+h} - N_t = 1) &= \lambda h + o(h) \\ P(N_{t+h} - N_t = 0) &= 1 - \lambda h + o(h) \\ P(N_{t+h} - N_t \geq 2) &= o(h), \end{aligned}$$

where, for a function f , we write $f(h) = o(h)$ if $f(h)/h \rightarrow 0$ as $h \rightarrow 0$. Loosely speaking, we say a function is $o(h)$ if it tends to zero faster than h .

- [2b] Let T_1, T_2, T_3, \dots be the increment times of the counting process $(N_t, t \geq 0)$, i.e.,

$$T_n = \inf\{t \geq 0 : N_t \geq n\}.$$

The process $(N_t, t \geq 0)$ is a Poisson process of rate λ if the random variables $T_{n+1} - T_n$ are independent and identically distributed (iid) with the $\text{Exp}(\lambda)$ distribution, i.e., $P(T_{n+1} - T_n \geq t) = \exp(-\lambda t)$ for all $t \geq 0$.

Recall that a random variable X has the $\text{Exp}(\lambda)$ distribution if $P(X > t) = e^{-\lambda t}$. Now, by Bayes' theorem,

$$P(X > t+s | X > t) = \frac{P(X > t+s \cap X > t)}{P(X > t)} = \frac{\exp(-\lambda(t+s))}{\exp(-\lambda t)} = P(X > s).$$

If we think of X as the time to occurrence of an event, then knowing that the event hasn't occurred up to time t tells us nothing about how much longer we need to wait for it to occur; the distribution of the *residual time* until it occurs doesn't depend on how long we have already waited. This is referred to as the *memoryless property* of the exponential distribution.

We shall establish the equivalence of [2], [2a] and [2b] by showing that [2] \Rightarrow [2a] \Rightarrow [2b] \Rightarrow [2]. The first implication is obvious by letting s tend to zero.

For the second implication, it suffices to show that T_1 has an $\text{Exp}(\lambda)$ distribution since, by [1], $T_1, T_2 - T_1, T_3 - T_2, \dots$ are iid. Let F denote the cdf of T_1 . Observe that

$$\begin{aligned} P(T_1 \in (t, t+h]) &= P(N_t = 0, N_{t+h} - N_t \geq 1) \\ &= P(N_t = 0)P(N_{t+h} - N_t \geq 1 | N_u = 0, 0 \leq u \leq t) \\ &= P(N_t = 0)P(N_{t+h} - N_t \geq 1) \\ &= (1 - \lambda h + o(h))P(T_1 > t), \end{aligned}$$

where the third equality follows from [1] and the last equality from [2a]. The above equation implies that

$$P(T_1 \in (t, t + h] | T_1 > t) = 1 - \lambda h + o(h),$$

i.e., that $1 - F(t + h) = (1 - \lambda h + o(h))(1 - F(t))$. Letting h tend to zero, we obtain $F'(t) = -\lambda(1 - F(t))$. Solving this differential equation with the boundary condition $F(0) = 0$, we get $F(t) = 1 - \exp(-\lambda t)$, which is the cdf of an $\text{Exp}(\lambda)$ random variable. Thus, we have shown that T_1 has an $\text{Exp}(\lambda)$ distribution, as required.

To show the last implication, we show (11) by induction on k . Without loss of generality, let $t = 0$. For $k = 0$, (11) reads $P(N_s = 0) = e^{-\lambda s}$. Since the event $N_s = 0$ is the same as $T_1 > s$, this follows from the exponential distribution of T_1 . Next, suppose that (11) holds for all $j \leq k$. By conditioning on the time to the first increment, we have

$$\begin{aligned} P(N_s = k + 1) &= \int_0^s f(u)P(N_s = k + 1 | N_u = 1)du \\ &= \int_0^s f(u)P(N_{s-u} = k)du \\ &= \int_0^s \lambda e^{-\lambda u} \frac{(\lambda(s-u))^k}{k!} e^{-\lambda(s-u)} du \\ &= e^{-\lambda s} \frac{(\lambda s)^{k+1}}{(k+1)!}. \end{aligned}$$

Here, f is the density function of T_1 , the time to the first increment. The second equality follows from [1], and the third from [2b] and the induction hypothesis. Thus, we have shown that [2b] implies [2], completing the proof of the equivalence of [2], [2a] and [2b].

The reason that this process is called a Poisson process is that the number of points within an interval of duration t has a Poisson distribution, with mean λt . Intuitively, you can think of each infinitesimal interval of length dt as having a point with probability λdt , and no points with probability $1 - (\lambda dt)$. The probability of having two or more points is of order $(dt)^2$, which is negligible. (This is a restatement of [2a]).

The Poisson process also has the property that the numbers of points in disjoint intervals are mutually independent random variables. Key to this independence is the memoryless nature of the exponential distribution, which

describes the time between successive points or events. This is what makes the Poisson process a Markov process.

Chapman-Kolmogorov equations

It turns out that the Poisson process contains most of the features of general continuous-time Markov processes. Let us return to such a process on a countable state space S , with transition probability matrices $\{P(t), t \geq 0\}$. We saw earlier that these matrices obeyed equation (10). Thus,

$$P(t + \delta) - P(t) = P(t)(P(\delta) - I) = (P(\delta) - I)P(t).$$

This suggests that, if $\frac{1}{\delta}(P(\delta) - I)$ converges to some matrix Q as δ decreases to zero, then

$$P'(t) = QP(t) = P(t)Q. \quad (12)$$

This is correct if the Markov process is finite-state, so the matrices $P(t)$ are finite-dimensional. For countable state chains, the first equality still holds but the second may not. The two equalities are known, respectively, as the Chapman-Kolmogorov backward and forward equations. The equations have the solution

$$P(t) = e^{Qt} = I + Qt + \frac{(Qt)^2}{2!} + \dots$$

What does the matrix Q look like? Since $P(\delta)$ and I both have all their row sums equal to 1, it follows that the row sums of $P(\delta) - I$, and hence of Q must all be zero. Moreover, all off-diagonal terms in $P(\delta)$ are non-negative and those in I are zero, so Q must have non-negative off-diagonal terms. Thus, in general, we must have

$$q_{ij} \geq 0 \forall i \neq j; \quad q_{ii} = -\sum_{j \neq i} q_{ij},$$

where q_{ij} denotes the ij^{th} term of Q . Define $q_i = -q_{ii}$, so $q_i \geq 0$. In fact, $q_i > 0$ unless i is an absorbing state. (Why?)

The matrix Q is called the generator (or infinitesimal generator) of the Markov chain. The evolution of a Markov chain can be described in terms of its Q matrix as follows. Suppose the Markov chain is currently in state i . Then it remains in state i for a random time which is exponentially distributed with parameter q_i . In particular, knowing how long the Markov chain has been in this state gives us no information about how much longer

it will do so (memoryless property of the exponential distribution). At the end of this time, the chain jumps to one of the other states; the probability of jumping to state j is given by $q_{ij}/\sum_{k \neq i} q_{ik} = q_{ij}/q_i$, and is independent of how long the Markov chain spent in this state, and of everything else in the past.

If we observed the CTMC only at the jump times, we'd get a DTMC on the same state space, with transition probabilities $p_{ij} = q_{ij}/q_i$. Conversely, we can think of a CTMC as a DTMC which spends random, exponentially distributed times (instead of unit time) in each state before transiting to another state.

In particular, everything we said about recurrence, transience, communicating classes and irreducibility translates to CTMCs.

Theorem 4 *Suppose $\{X_t, t \geq 0\}$ is an irreducible Markov process on a finite state space S . Then there is a unique probability distribution π on S , called the invariant distribution, such that $\pi Q = 0$. Moreover, for all $j \in S$,*

$$\pi(j) = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t 1(X(t) = j) dt.$$

It is left to you as an exercise to show that, if $\pi Q = 0$, then $\pi P(t) = \pi$ for all $t \geq 0$, where $P(t) = e^{Qt}$.

Reversibility

A CTMC on S is said to be reversible if there is a probability distribution π on S such that $\pi_i q_{ij} = \pi_j q_{ji}$ for all $i, j \in S$. Then π is also the invariant distribution of the Markov chain. To see this, observe that for all $j \in S$,

$$(\pi Q)_j = \sum_{i \in S} \pi_i q_{ij} = \sum_{i \in S} \pi_j q_{ji} = \pi_j \sum_{i \in S} q_{ji} = 0.$$

To obtain the second equality above, we have used the definition of reversibility. Then, since π_j is a constant that does not depend on the index i in the sum, we can pull it out of the sum. Finally, we use the fact that the row sums of Q are zero.

Example: The $M/M/1$ queue

This is just a birth-death process in continuous time but can be interpreted in terms of a queue. This and other models of queues find a variety of appli-

cations in describing communication networks, insurance, water reservoirs etc.

Consider a queue into which customers arrive according to a Poisson process of intensity λ . In other words, customers arrive at random times T_1, T_2, \dots , and $T_k - T_{k-1}$ form an iid sequence of $\text{Exponential}(\lambda)$ random variables. The k^{th} customer requires service for a random length of time, which is $\text{Exponential}(\mu)$, and service times of different customers are iid. There is a single server, which serves customers in the order of their arrival. We'll assume that $\lambda < \mu$, i.e., it takes less time to serve a customer on average ($1/\mu$) than it takes until the next customer arrives ($1/\lambda$ on average). If this were not the case, then the queue would become infinitely long.

In the notation for this type of queue, the first M refers to the arrival process being Markovian, the second M to the service process being Markovian, and the 1 to the number of servers. An $M/G/2$ queue would have Markovian arrivals, general service time distribution, and two servers.

Let X_t denote the number of customers in the queue at time t . We show that X_t evolves as a Markov process. If $X_t = 0$, then the next event has to be the arrival of a new customer, and the time to this arrival is $\text{Exponential}(\lambda)$, independent of how long the chain has been in state 0, or anything else about its past. Thus, $q_0 = \lambda$, i.e., $q_{00} = -\lambda$. Since the chain can only move to state 1 from state 0, and the Q matrix has rows summing to zero, it must be that $q_{01} = \lambda$. Next, suppose $X_t = n$ for some $n > 0$. The next event could be either the arrival of a new customer or the departure of the customer currently being served. Think of these as two clocks running in parallel, one of which (new arrival) will go off after an $\text{Exponential}(\lambda)$ time, and the other (departure) after an $\text{Exponential}(\mu)$ time. How long is it till one of these goes off, and what are the chances of each of them being first?

To answer that, let X and Y be independent exponential random variables with parameters λ and μ , describing the random times that clocks A (arrival) and S (service) go off, respectively. Then,

$$F_X(t) = 1 - e^{-\lambda t}, \quad F_Y(t) = 1 - e^{-\mu t}.$$

Let $Z = \min\{X, Y\}$ denote the time until the first of these two clocks goes off. We want to compute the cdf of Z and find out what the probability is that $Z = X$ (i.e., a new arrival happens before the service of the customer

at the head of the line is completed). Observe that

$$\begin{aligned} P(Z > t) &= P(\min\{X, Y\} > t) = P(X > t, Y > t) \\ &= P(X > t)P(Y > t) = e^{-(\lambda+\mu)t}, \end{aligned}$$

so that $F_Z(t) = 1 - e^{-(\lambda+\mu)t}$. Thus, the time Z to the next jump of the Markov chain is exponentially distributed with parameter $\lambda + \mu$. Hence, $q_n = \lambda + \mu$, where $q_{nn} = -q_n$ denotes the n^{th} diagonal element in the Q matrix.

What are the probabilities that the next jump corresponds to an arrival or departure respectively? We can compute $P(Y > X)$ by conditioning on all possible values of X . We have,

$$\begin{aligned} P(Y > X) &= \int_0^\infty f_X(x)P(Y > X|X = x)dx \\ &= \int_0^\infty f_X(x)P(Y > x)dx \text{ (by the independence of } X \text{ and } Y) \\ &= \int_0^\infty \lambda e^{-\lambda x} e^{-\mu x} dx = \frac{\lambda}{\lambda + \mu}. \end{aligned}$$

This is the probability that the next jump corresponds to an arrival (clock A goes off before clock B). The complementary probability $\frac{\mu}{\lambda+\mu}$ is the probability that the customer in service departs before the next arrival. Letting $q_{n,n+1}$ and $q_{n,n-1}$ denote the elements of the Q matrix corresponding to the transitions from n to $n + 1$ and $n - 1$ customers respectively, it must then be the case that

$$\frac{q_{n,n+1}}{q_{n,n+1} + q_{n,n-1}} = \frac{\lambda}{\lambda + \mu}, \quad \frac{q_{n,n-1}}{q_{n,n+1} + q_{n,n-1}} = \frac{\mu}{\lambda + \mu}.$$

Combining this with the fact that $q_{n,n+1} + q_{n,n-1} = q_n = \lambda + \mu$, we get that $q_{n,n+1} = \lambda$ and $q_{n,n-1} = \mu$.

Thus, we have fully specified the Q matrix, which means we have fully described the Markov process. It is left as an exercise to you to check that this Markov process is reversible with the invariant distribution π given by

$$\pi_n = (1 - \rho)\rho^n, \quad n \geq 0, \quad \text{where } \rho = \frac{\lambda}{\mu} < 1.$$