Discrete Time Markov Processes

Adam Shwartz

Electrical Engineering
Technion—Israel Institute of Technology
Haifa 32000, Israel

0. Caveat

What follows is a quick survey of the main ingredients in the theory of discrete-time Markov processes. It is a birds' view, rather than the definitive "state of the art." To maximize accessibility, the nomenclature of mathematical probability is avoided, although rigor is not sacrificed. To compensate, examples (and counterexamples) abound and the bibliography is annotated. Relevance to control is discussed in §8.

1. Introduction

Discrete-time Markov processes, or Markov chains, are a powerful tool for modeling and analyzing discrete-time systems whose behavior is influenced by randomness. A Markov chains is probably the simplest object which incorporates both dynamics (that is, notions of "state" and time) and randomness. Let us illustrate the idea through a gambling example.

Example 1. A gambler bets one dollar on "red" at every turn of a (fair) game of roulette. Then, at every turn he gains either one dollar (win, with probability 1/2) or (-1) (lose). In an ideal game, the gains form a sequence of independent, identically distributed (i.i.d.) random variables. We cannot predict the outcome of each bet, although we do know the odds. Denote by X_t the total fortune the gambler has at time t. If we know X_s, then we can calculate the distribution of X_{s+1} (that is, the probability that $X_{s+1} = y$, for all possible values of y), and even of X_{s+k} for any $k > 0$. The variable X_t serves as a "state" in the following sense: given X_s, knowledge of X_t for $t < s$ is irrelevant to the calculation of the distribution of future values of the state.

This notion of a state is similar to classical "state space" descriptions. Consider the standard linear model of a dynamical system

$$x_{t+1} = Ax_t + v_t, \tag{1}$$

or the more general nonlinear model

$$x_{t+1} = f(x_t, v_t). \tag{2}$$

It is intuitively clear that, given the present state, the past does not provide additional information about the future, as long as v_t is not predictable (this is why deterministic state-space models require that v_t be allowed to change arbitrarily at each t). This may remain true even when v_t is random: for example, when they are i.i.d., for then the past does not provide additional information about future v_s's. As we shall see in Theorem 1, in this case (1) and (2) define Markov chains.

In the next section we give the basic definitions and describe the dynamics of Markov chains. We assume that the state space $\mathbf{S}$ is countable. We restrict our attention to time-homogeneous dynamics (see comment following Theorem 2), and discuss the limiting properties of the Markov chain. Finally, we shall discuss extensions to continuous time and to more general state spaces.

We conclude this section with a brief review of standard notation. All of our random variables and events are defined on a probability space Ω with a collection $\mathcal{F}$ of events (subsets of Ω) and probability $\mathbb{P}$. The "probability triple" $(\Omega, \mathcal{F}, \mathbb{P})$ is fixed, and we denote expectation by $\mathbb{E}$. For events A and B with $\mathbb{P}(B) > 0$, the basic definition of a conditional probability (the multiplication rule) is

$$\mathbb{P}\,(\,A \mid B\,) \triangleq \frac{\mathbb{P}\,(A \bigcap B)}{\mathbb{P}\,(B)}\ . \tag{3}$$

The abbreviation i.i.d. stands for *independent, identically distributed*, and random variables are denoted by capital letters. The identity matrix is denoted by I and $\mathbf{1}_A$ is the indicator function of A, that is, $\mathbf{1}_A(\omega) = 1$ if $\omega \in A$ and $= 0$ otherwise.

2. Definitions and Construction

Let $X_0, X_1, \ldots$ be a sequence of random variables, with values in a state space $\mathbf{S}$. We assume that $\mathbf{S}$ is finite or countable, and for convenience we usually set $\mathbf{S} = \{1, 2, \ldots\}$.

Definition 1. *A sequence $X_0, X_1, \ldots$ on $\mathbf{S}$ is a Markov chain if it possesses the Markov property, that is, if for all $t > 0$,*

$$\mathbb{P}\,(\,X_t = j \mid X_{t-1} = i_{t-1}, X_{t-2} = i_{t-2}, \ldots, X_0 = i_0\,) = \mathbb{P}\,(\,X_t = j \mid X_{t-1} = i_{t-1}\,)\,.$$

A Markov chain is called homogeneous if $\mathbb{P}\,(\,X_t = j \mid X_{t-1} = i\,)$ does not depend on t. In this case we denote the transition probability from i to j by

$$\mathrm{p}_{ij} \triangleq \mathbb{P}\,(\,X_t = j \mid X_{t-1} = i\,)\,.$$

The matrix $\mathrm{P} \triangleq \{\mathrm{p}_{ij},\ i, j = 1, \ldots\}$ is called the transition matrix.

Henceforth we restrict our attention to homogeneous chains.

From the Markov property, a little algebra with the definition of conditional probability gives the more general Markov property: if $t_1 \le t_2 \le \cdots \le t_k \le \cdots \le t_\ell$ then

$$\begin{aligned}\mathbb{P}\,\big(\,&X_{t_\ell} = j_\ell,\ X_{t_{\ell-1}} = j_{\ell-1}, \ldots, X_{t_k} = j_k \,\big|\, X_{t_{k-1}} = j_{k-1}, \ldots, X_{t_1} = j_1\big) \\ &= \mathbb{P}\,\big(\,X_{t_\ell} = j_\ell,\ X_{t_{\ell-1}} = j_{\ell-1}, \ldots, X_{t_k} = j_k \,\big|\, X_{t_{k-1}} = j_{k-1}\big)\,. \end{aligned} \tag{4}$$

This is a precise statement of the intuitive idea given in the introduction: given the present state $X_{t_{k-1}} = j_{k-1}$, the past does not provide additional information.

A chain is called finite if $\mathbf{S}$ is a finite set. An alternative name for "homogeneous Markov chain" is "Markov chain with stationary transition probabilities." There are those who call a homogeneous Markov chain a "stationary Markov chain." However, since "stationary process" means something entirely different, we shall avoid such usage: see Example 6.

Suppose that a process is a Markov chain according to Definition 1, and that its initial probability distribution is given by the row vector $\mu(0)$, that is

$$\mathbb{P}\,(X_0 = i) = \mu_i(0), \quad i = 1, 2, \ldots.$$

Then we can calculate the joint probability distribution at times 0 and 1 from the definition of conditional probability,

$$\mathbb{P}\,(X_0 = i, X_1 = j) = \mathbb{P}\,(X_1 = j \mid X_0 = i) \cdot \mathbb{P}\,(X_0 = i) = \mu_i(0) \cdot \mathrm{p}_{ij}.$$

More generally, using the Markov property:

$$\mathbb{P}\,(X_0 = j_0, X_1 = j_1, \ldots, X_t = j_t) = \mu_{j_0}(0) \cdot \prod_{s=1}^{t} \mathrm{p}_{j_{s-1}j_s}.$$

So, the probability distribution of the whole process can be calculated from the two quantities: the initial probability distribution, and the transition probabilities. In particular, we can calculate the probability distribution at any time t

$$\mu_j(t) = \mathbb{P}\,(X_t = j) = \sum_{j_0, j_1, \ldots, j_{t-1}} \mu_{j_0}(0) \prod_{s=1}^{t} \mathrm{p}_{j_{s-1}j_s} \tag{5}$$

where the sum is over all states in $\mathbf{S}$, that is, each index is summed over the values $1, 2, \ldots$. In vector notation,

$$\mu(t) = \mu(t-1)\mathrm{P} = \ldots = \mu(0)\mathrm{P}^t. \tag{6}$$

(If $\mathbf{S}$ is countable then, of course, the matrix is infinite, but it is manipulated in exactly the same way as a finite matrix.) Thus the *probability distribution* of a Markov chain evolves as a linear dynamical system, even when its evolution equation (2) is nonlinear. The one-dimensional probability distribution and the transition probabilities clearly satisfy

$$\begin{cases} \mu_j(t) \geq 0 \\ \sum_{j \in \mathbf{S}} \mu_j(t) = 1 \end{cases} \qquad \begin{cases} \mathrm{p}_{ij} \geq 0 \\ \sum_{j \in \mathbf{S}} \mathrm{p}_{ij} = 1. \end{cases} \tag{7}$$

That is, the rows of the transition matrix P sum to one. Thus P is a *stochastic matrix*: its elements are non negative and its rows sum to one.

If we denote by $\mathrm{p}_{ij}^{(n)} \triangleq \mathbb{P}\left(X_{m+n} = j \mid X_m = i\right)$ the *n step transition probability* from i to j then we obtain from the definition of conditional probability and the Markov property (4) the Chapman-Kolmogorov equations

$$\mathrm{p}_{ij}^{(n+m)} = \sum_{k \in \mathbf{S}} \mathrm{p}_{ik}^{(n)} \mathrm{p}_{kj}^{(m)} \quad \text{or} \quad \mathrm{P}^{n+m} = \mathrm{P}^n \mathrm{P}^m. \tag{8}$$

Therefore the $\mathrm{p}_{ij}^{(n)}$ are the elements of the matrix P^n. This matrix notation yields a compact expression for expectations of functions of the state. To compute $\mathbb{E}\, g(X_t)$ for some function g on $\mathbf{S} = \{1, 2, \ldots\}$, we represent g by a column vector $\underline{g} \triangleq \{g(1), g(2), \ldots\}^T$ (where T denotes transpose). Then

$$\mathbb{E}\, g(X_t) = \mu(0) \cdot \mathrm{P}^t \cdot \underline{g}\ .$$

Note that this expression does not depend on the particular $\mathbf{S}$: because the state space is countable we can, by definition, relabel the states so that the state space becomes $\{1, 2, \ldots\}$. Let us now summarize the connection between the representations (1)–(2) and Markov chains.

Theorem 1. *Let $V_0, V_1, \ldots$ be a sequence of i.i.d. random variables, independent of X_0. Then for any (measurable) function f, the sequence $X_0, X_1, \ldots$ defined through (2) is a Markov chain. Conversely, let $\tilde{X}_0, \tilde{X}_1, \ldots$ be a Markov chain with values in $\mathbf{S}$. Then there is a (probability triple and a measurable) function f and a sequence $V_0, V_1, \ldots$ of i.i.d. random variables so that the process $X_0, X_1, \ldots$ defined by (2) with $X_0 = \tilde{X}_0$ has the same probability distribution as $\tilde{X}_0, \tilde{X}_1, \ldots$, that is, for all t and $j_0, j_1, \ldots, j_t$,*

$$\mathbb{P}\left(X_0 = j_0, X_1 = j_1, \ldots, X_t = j_t\right) = \mathbb{P}\left(\tilde{X}_0 = j_0, \tilde{X}_1 = j_1, \ldots, \tilde{X}_t = j_t\right).$$

Note that, whether the system (2) is linear or not, the evolution of the probability distribution (5)–(6) is always linear.

We have seen that a Markov chain defines a set of transition probabilities. The converse is also true: given a set of transition probabilities and an initial probability distribution, it is possible to construct a stochastic process which is a Markov chain with the specified transitions and probability distribution.

Theorem 2. *If $X_0, X_1, \ldots$ is a homogeneous Markov chain then its probability distribution and transition probabilities satisfy (7) and (8). Conversely, given $\mu(0)$ and a matrix P that satisfy (7), there exists a (probability triple and a) Markov chain with initial distribution $\mu(0)$ and transition matrix P.*

The restriction to homogeneous chains is not too bad: if we define a new state $\tilde{x} \triangleq \{t, x\}$, then it is not hard to see that we can incorporate explicit time dependence, and the new state space is still countable.

Example 2. Let $V_0, V_1, \ldots$ be i.i.d. and independent of X_0. Assume both V_t and X_0 have integer values. Then

$$X_{t+1} \triangleq X_t + V_t = X_0 + \sum_{s=0}^{t} V_s \tag{9}$$

defines a Markov chain called a chain with *stationary independent increments*, with state space $\ldots, -1, 0, 1, \ldots$. The transition probability p_{ij} depends only on the difference $j - i$. It turns out [Ch] that the converse is also true: if the transition probabilities of a Markov chain depend only on the difference $j - i$ then the process can be obtained via (9) with i.i.d. V_t.

A *random walk* is a process defined through Equation (9), but where the V_t are not necessarily integers—they are real valued.

3. Properties and Classification

Given two states i and j, it may or may not be possible to reach j from i. This leads to the notion of classes.

Definition 2. *We say a state i leads to j if*

$$\mathbb{P}\,(\, x_t = j \text{ for some } t \mid x_0 = i) > 0.$$

This holds if and only if $\mathrm{p}_{ij}^{(t)} > 0$ for some t. We say states i and j communicate, denoted by $i \leftrightarrow j$, if i leads to j and j leads to i.

Communication is a property of pairs: it is obviously symmetric ($i \leftrightarrow j$ if and only if $j \leftrightarrow i$) and is transitive ($i \leftrightarrow j$ and $j \leftrightarrow k$ implies $i \leftrightarrow k$) by the Chapman-Kolmogorov equations (8). By convention, $i \leftrightarrow i$. By these three properties, $\leftrightarrow$ defines an equivalence relation. We can therefore partition $\mathbf{S}$ into non-empty *communicating classes* $\mathbf{S}_0, \mathbf{S}_1, \ldots$ with the properties

(i) every state i belongs to exactly one class,

(ii) if i and j belong to the same class then $i \leftrightarrow j$, and

(iii) if i and j belong to different classes then i and j do not communicate.

We denote the class containing state i by $\mathbf{S}(i)$. Notice that if i leads to j but j does not lead to i then i and j *do not* communicate.

Definition 3. *A set of states C is closed, or absorbing, if $\mathrm{p}_{ij} = 0$ whenever $i \in C$ and $j \notin C$. Equivalently,*

$$\sum_{j \in C} \mathrm{p}_{ij} = 1 \quad \text{for all } i \in C.$$

If a set C is not closed, then it is called open. A Markov chain is irreducible if all states communicate. In this case its partition contains exactly one class. A Markov chain is indecomposable if its partition contains at most one closed set.

Define the *incidence matrix* $\mathcal{I}$ as follows:

$$\mathcal{I}_{ij} = \begin{cases} 1 & \text{if } \mathrm{p}_{ij} > 0 \\ 0 & \text{otherwise}\,. \end{cases}$$

We can also define a directed graph whose nodes are the states, with a directed arc between any two states for which $\mathcal{I}_{ij} = 1$. The communication properties can obviously be extracted from $\mathcal{I}$ or from the directed graph. The chain is irreducible if and only if the directed graph is connected in the sense that, going in the direction of the arcs, we can reach any node from any other node. Closed classes can also be defined in terms of the incidence matrix or the graph. The classification leads to the following maximal decomposition.

Theorem 4. *By renumbering the states if necessary, we can put the matrix* P *into the block form*

$$\mathrm{P} = \begin{bmatrix} \mathrm{P}_1 & 0 & \dots & 0 & 0 \\ 0 & \mathrm{P}_2 & \dots & 0 & 0 \\ \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & \dots & \mathrm{P}_m & 0 \\ R_1 & R_2 & \dots & R_m & Q \end{bmatrix} \tag{10}$$

where the blocks P_i *correspond to closed irreducible classes. It is maximal in the sense that smaller classes will not be closed, and no subset of states corresponding to* Q *is a closed irreducible class.*

If $\mathbf{S}$ is countable, then the number of classes may be infinite. Note that all the definitions in this section apply when we replace $\mathbf{S}$ with any closed class.

Much like other dynamical systems, Markov chains can have cyclic behavior, and can be unstable. The relevant definitions are

Definition 4. *A state* i *has period* d *if* d *is the greatest common divisor of the set* $\{t : \mathrm{p}_{ii}^{(t)} > 0\}$. *If* d *is finite and* $d > 1$ *then the state is called periodic; otherwise it is aperiodic.*

Definition 5. *A state* i *is called recurrent if the probability of starting at* i *and returning to* i *in finite time is* 1. *Formally, if*

$$\mathbb{P}\,(\,X_t = i \text{ for some } t > 1 \mid x_1 = i) = 1.$$

Otherwise it is called transient.

Example 3. In the chain on $\mathbf{S} = \{1, 2\}$ with $\mathrm{p}_{ij} = 1$ if and only if $i \neq j$, both states are periodic with period $d = 2$ and both states are recurrent. The states communicate, and so $\mathbf{S}$ contains exactly one class, which is therefore closed. Consequently the chain is irreducible and indecomposable. However, if $\mathrm{p}_{12} = \mathrm{p}_{22} = 1$ then state 1 does not lead to itself, the states are not periodic, state 2 is recurrent and state 1 is transient. In this case the partition contains two sets: the closed set $\{2\}$, and the open set $\{1\}$. Consequently, the chain is not irreducible, but it is indecomposable.

When $\mathbf{S}$ is finite, then either it is irreducible or it contains a closed proper subset.

Example 4. Let $\mathbf{S} = \{1, 2, \ldots\}$. Suppose $\mathrm{p}_{ij} = 1$ if and only if $j = i + 1$. Then all states are transient, and $\mathbf{S}$ is indecomposable but not irreducible. Every set of the form $\{i : i \geq k\}$ is closed, but in the partition of the state space each state is the only member in its class. Suppose now $p_{11} = 1$ and for $i > 1$, $p_{ij} = 1$ if and only if $j = i - 1$. Then state 1 is the only recurrent state, and again each state is alone in its class.

Theorem 5. *Let $\mathbf{S}_k$ be a class. Then either all states in $\mathbf{S}_k$ are recurrent, or all are transient. Moreover, all states in $\mathbf{S}_k$ have the same period d.*

4. Algebraic View and Stationarity

The matrix P is positive, in the sense that each of its entries are positive. When $\mathbf{S}$ is finite, the Perron-Frobenius theorem [Se] implies

Theorem 6. *Let $\mathbf{S}$ be finite. Then P has a non zero left eigenvector π whose entries are non negative, and $\pi \cdot \mathrm{P} = \pi$, that is, the corresponding eigenvalue is 1. Moreover, $|\lambda| \leq 1$ for all other eigenvalues λ. The multiplicity of the eigenvalue 1 is equal to the number of irreducible closed subsets of the chain. In particular, if the entries of P^n are all positive for some n, then the eigenvalue 1 has multiplicity 1. In this case the entries of π are all positive and $|\lambda| < 1$ for all other eigenvalues λ.*

If the entries of P^n are all positive for some n then the chain is irreducible and aperiodic, hence the second part of the theorem. If the chain is irreducible and periodic with period d, then the dth roots of unity are left eigenvalues of P, each is of multiplicity 1 and all other eigenvalues have strictly smaller modulus. The results for a general finite chain can be obtained by writing the chain in the block form (10).

Definition 6. *Let $\mathbf{S}$ be finite or countable. A probability distribution μ satisfying $\mu \cdot \mathrm{P} = \mu$ is called invariant (under P) or stationary.*

Theorem 6 thus implies that every finite Markov chain possesses at least one invariant probability distribution. For countable chains, Example 4 shows that this is not true.

Example 5. Returning to Example 3, in the first case $(1/2, 1/2)$ is the only invariant probability distribution, while in the second case $(0, 1)$ is the only invariant probability distribution. In Example 4, in the first case there is no invariant probability distribution, while in the second case $(1, 0, 0, \ldots)$ is the only invariant probability distribution. Finally, if $\mathrm{P} = I$, the 2×2 identity matrix, then $\pi \triangleq (p, 1 - p)$ is invariant for any $0 \leq p \leq 1$.

Example 6. Recall that a process $X_0, X_1, \ldots$ is called stationary if, for all positive t and s, the distribution of $\{X_0, X_1, \ldots, X_t\}$ is the same as the distribution of $\{X_s, X_{1+s}, \ldots, X_{t+s}\}$. From the definitions it follows that a (homogeneous) Markov chain (finite or not) is stationary if and only if $\mu(0)$ is invariant.

A very useful tool in the calculation of invariant probability distributions is the "balance equations:"

$$\pi_i = \sum_{j : j \to i} \pi_j \mathrm{p}_{ji} = \pi_i \sum_{j : i \to j} \mathrm{p}_{ij} \qquad (11)$$

where the first equality is just a restatement of the definition of invariant probability, and the second follows since by (7), the last sum equals 1. The intuition behind these equations is very useful: in steady state, the rate at which "probability mass enters" must be equal to the rate at which it "leaves." This is particularly useful for continuous-time chains. More generally, given any set S, the rate at which "probability mass enters" the set (under the stationary distribution) equals the rate at which it "leaves:"

Theorem 7. *Let S be a set of states. Then*

$$\sum_{i \in S} \sum_{j : j \to i} \pi_j \mathrm{p}_{ji} = \sum_{i \in S} \pi_i \sum_{j : i \to j} \mathrm{p}_{ij}.$$

Example 7. Random walk with a reflecting barrier. This example models a discrete-time queue where, at each instance, either arrival or departure occurs. The state space $\mathbf{S}$ is the set of non negative integers (including 0), and

$$\mathrm{p}_{00} = 1 - p, \quad \mathrm{p}_{i(i+1)} = p, \quad \mathrm{p}_{i(i-1)} = 1 - p \quad \text{for } i \geq 1.$$

Then all states communicate so that the chain is irreducible, the chain is aperiodic and recurrent. From (11) we obtain

$$\begin{aligned} \pi_0 &= \pi_0 \mathrm{p}_{00} + \pi_1 \mathrm{p}_{10} \\ \pi_i &= \pi_{i-1} \mathrm{p}_{(i-1)i} + \pi_{i+1} \mathrm{p}_{(i+1)i}, \qquad i \geq 1. \end{aligned}$$

When $p < 1/2$, this and (7) imply that $\pi_i = \frac{p^i}{(1-p)^i}$ for $i \geq 0$.

Example 8. Birth-death process. A Markov chain on $\mathbf{S} = \{0, 1, \ldots\}$ is a birth-death process if $\mathrm{p}_{ij} = 0$ whenever $|i - j| \geq 2$. If X_t is the number of individuals alive at time t then, at any point in time, this number can increase by one (birth), decrease by one (death) or remain constant (simultaneous birth and death). Unlike Example 7, here the probability of a change in size may depend on the state.

5. Random Variables

In this section we shift our emphasis back from algebra to the stochastic process. We define some useful random variables associated with the Markov chain. It will be convenient to use $\mathbb{P}_j$ for the probability conditioned on the process starting at state j. That is, for an event A,

$$\mathbb{P}_j(A) \stackrel{\triangle}{=} \mathbb{P}\left(A \mid X_0 = j\right),$$

with a similar convention for expectation $\mathbb{E}_j$. The Markov property says that the past of a Markov chain is immaterial given the present. But suppose we observe a process until a random time, say the time a certain event occurs. Is this property preserved? The answer is positive, but only for non-anticipative times:

Definition 7. *Let S be a collection of states, that is, a subset of* $\mathbf{S}$. *The hitting time τ_S of S is the first time the Markov chain visits a state in S. Formally,*

$$\tau_S = \inf\{t > 0 : X_t \in S\}.$$

Note that by convention, if X_t never visits S then $\tau_S = \infty$. The initial time, here $t = 0$, does not qualify in testing whether the process did or did not visit S. By definition, hitting times have the following property. In order to decide whether or not $\tau_S = t$, it suffices to know the values of $X_0, \ldots, X_t$. This gives rise to the notion of Markov time or stopping time.

Definition 8. *A random variable τ with positive integer values is called a stopping time, or a Markov time (with respect to the process $X_0, X_1, \ldots$) if one of the following equivalent conditions hold. For each $t \geq 0$*

(i) it suffices to know the values of $X_0, X_1, \ldots, X_t$ in order to determine whether the event $\{\tau = t\}$ occurred or not,

(ii) there exists a function f_t so that

$$\mathbf{1}_{\tau=t}(\omega) = f_t\left(X_0(\omega), \ldots, X_t(\omega)\right).$$

An equivalent, and more standard definition is obtained by replacing $\tau = t$ by $\tau \leq t$. With respect to such times, the Markov property holds in a stronger sense.

Theorem 8 Strong Markov property. *If τ is a stopping time for $X_0, X_1, \ldots$, then*

$$\begin{aligned}\mathbb{P}\left(X_{\tau+1} = j_1, X_{\tau+2} = j_2, \ldots, X_{\tau+m} = j_m \mid X_t = i_t,\ t < \tau,\ X_\tau = i^*\right) \\ = \mathbb{P}\left(X_1 = j_1, X_2 = j_2, \ldots, X_m = j_m \mid X_0 = i^*\right).\end{aligned}$$

We can now rephrase and complement the definition of recurrence. We write τ_j when we really mean $\tau_{\{j\}}$.

Definition 9. *The state j is recurrent if $\mathbb{P}_j(\tau_j < \infty) = 1$. It is called positive recurrent if $\mathbb{E}_j\, \tau_j < \infty$, and null-recurrent if $\mathbb{E}_j\, \tau_j = \infty$.*

If state j is recurrent, then the hitting time of j is finite. By the strong Markov property, when the processes hits j for the first time, it "restarts:" therefore, it will hit j again! and again! So, let N_j be the number of times the process hits state j:

$$N_j \triangleq \sum_{t=1}^{\infty} \mathbf{1}_{X_t = j}.$$

Theorem 9. *(i) If a state is positive recurrent, then all states in its class are positive recurrent. The same holds for null recurrence.*
(ii) Suppose j is recurrent. Then $\mathbb{P}_j(N_j = \infty) = 1$, and consequently $\mathbb{E}_j\, N_j = \infty$. Moreover, for every state i,

$$\mathbb{P}_i\,(N_j = \infty) = \mathbb{P}_i\,(\tau_j < \infty) \cdot \mathbb{P}_j(N_j = \infty) = \mathbb{P}_i\,(\tau_j < \infty)\,,$$

and if $\mathbb{P}_i\,(\tau_j < \infty) > 0$ then $\mathbb{E}_i\, N_j = \infty$.
(iii) Suppose j is transient. Then $\mathbb{P}_j(N_j < \infty) = 1$, and for all i,

$$\mathbb{E}_i\, N_j = \frac{\mathbb{P}_i\,(\tau_j < \infty)}{1 - \mathbb{P}_j\,(\tau_j < \infty)}\;.$$

To see why the last relation should hold, note that by the strong Markov property,

$$\mathbb{P}_i\,(\tau_j < \infty \;\text{ and a second visit occurs}) = \mathbb{P}_i\,(\tau_j < \infty) \cdot \mathbb{P}_j\,(\tau_j < \infty)\,,$$

and similarly for later visits. This means that the distribution of the number of visits is geometric: with every visit we get another chance, with equal probability, to revisit. Therefore

$$\begin{aligned}
\mathbb{E}_i\, N_j &= \mathbb{P}_i\,(\tau_j < \infty) + \mathbb{P}_i\,(\tau_j < \infty \;\text{ and a second visit occurs}) + \ldots \\
&= \mathbb{P}_i\,(\tau_j < \infty)\,(1 + \mathbb{P}\,(\text{a second visit occurs} \mid \tau_j < \infty)) + \ldots. \\
&= \mathbb{P}_i\,(\tau_j < \infty)\,(1 + \mathbb{P}_j\,(\tau_j < \infty) + \ldots)
\end{aligned}$$

which is what we obtain if we expand the denominator. A similar interpretation gives rise to

$$\mathbb{E}_i\, \tau_j = 1 + \sum_{k : k \neq j \in \mathbf{S}} \mathrm{p}_{ik}\, \mathbb{E}_k\, \tau_j.$$

We have a simple criterion for recurrence in terms of transition probabilities, since

$$\mathbb{E}_i\, N_j = \mathbb{E}_i \sum_{t=1}^{\infty} \mathbf{1}_{X_t = j} = \sum_{t=1}^{\infty} \mathbb{P}_i\,(X_t = j) = \sum_{t=1}^{\infty} \mathrm{p}_{ij}^{(t)}.$$

6. Limit Theorems: Transitions

Classical limit theorems concern the behavior of t-step transition probabilities, for large t. Limits for the random variables are discussed in §7.

Theorem 10. *For every Markov chain, the limit*

$$\mathrm{P}^* \triangleq \lim_{t\to\infty} \frac{1}{t} \sum_{s=0}^{t-1} \mathrm{P}^s \tag{12}$$

exists and satisfies

$$\mathrm{P}^* \cdot \mathrm{P} = \mathrm{P} \cdot \mathrm{P}^* = \mathrm{P}^* \cdot \mathrm{P}^* = \mathrm{P}^*.$$

If $\mathbf{S}$ *is finite then* P^* *is a stochastic matrix.*
(i) Suppose the Markov chain is indecomposable, recurrent and non periodic. Then for all states i, j, k,

$$\lim_{t\to\infty} \sum_{j\in\mathbf{S}} |\mathrm{p}_{ij}^{(t)} - \mathrm{p}_{kj}^{(t)}| = 0.$$

(ii) An irreducible chain is positive recurrent if and only if it has an invariant probability distribution π*, and in this case* $\lim_{t\to\infty} \mathrm{p}_{ij}^{(t)} = \pi(j)$ *for all* i, j*. If it is null recurrent then for all* i, j, $\lim_{t\to\infty} \mathrm{p}_{ij}^{(t)} = 0$*. If state* j *is transient then* $\sum_t \mathrm{p}_{ij}^{(t)} < \infty$.

Since a finite Markov chain always contains a finite closed set of states, there always exists an invariant probability distribution. Moreover, if a set is recurrent than it is positive recurrent.

Example 3 Continued. For the periodic chain, $\mathrm{p}_{ij}^{(t)}$ clearly does not converge. However, $\mathrm{P}_{ij}^* = 1/2$ for all i, j, and the rows define an invariant measure.

Example 8 Continued. Assume that for the birth death process $\mathrm{p}_{i(i+1)} > 0$ and $\mathrm{p}_{(i+1)i} > 0$ for all $i \geq 0$ and $\mathrm{p}_{ii} > 0$ for some i. Then the chain is obviously irreducible, and aperiodic (if $\mathrm{p}_{ii} = 0$ for all i then $d = 2$). Using (11) we obtain that an invariant probability distribution, if it exists, must satisfy

$$\pi_i = \frac{\mathrm{p}_{01} \cdots \mathrm{p}_{(i-1)i}}{\mathrm{p}_{10} \cdots \mathrm{p}_{i(i-1)}} \cdot \pi_0. \tag{13}$$

Therefore, any invariant probability must satisfy $\pi_i > 0$ for all i, and in particular $\pi_0 > 0$. So, we can invoke (7) to obtain the following dichotomy. Either

$$Z \triangleq \sum_{i\in\mathbf{S}} \frac{\mathrm{p}_{01} \cdots \mathrm{p}_{i-1i}}{\mathrm{p}_{10} \cdots \mathrm{p}_{ii-1}} < \infty, \tag{14}$$

in which case (13)–(14) determine the unique invariant probability, and we conclude that the Markov chain is positive recurrent. Or $Z = \infty$, in which case there is no invariant probability and the chain is not positive recurrent.

In terms of the transition matrix P, if a chain is non periodic, indecomposable and recurrent then the matrix converges (uniformly over rows) to a matrix having identical rows, which are either all zeroes (null-recurrent case), or equal to the invariant probability distribution. Here are the missing cases from Theorem 10. Denote the mean hitting time of state j starting at i by $m_{ij} = \mathbb{E}_i \tau_j$. Clearly m_{jj} is infinite if j is not positive recurrent, and we shall use the convention that $a/\infty = 0$ whenever a is finite.

Theorem 11. *If a state j is transient then $\lim_{t\to\infty} \mathrm{p}_{ij}^{(t)} = 0$. If j is recurrent with period d then $\lim_{t\to\infty} \mathrm{p}_{jj}^{(nd)} = \frac{d}{m_{jj}}$. If j is non periodic, this remains true with $d = 1$, so that (by Theorem 10) $\pi(j) \cdot m_{jj} = 1$.*

The last statement should be intuitive: the steady state probability of visiting state j is a measure of how often this state is "visited," and this is inversely proportional to the mean time between visits. The rate at which convergence takes place depends on the second largest eigenvalue of P. Therefore, if the Markov chain is finite, indecomposable and aperiodic with invariant probability distribution π, then

$$\left|\mathrm{p}_{ij}^{(t)} - \pi_j\right| \leq R\rho^t \quad \text{for all } i, j$$

with $\rho < 1$. This of course implies that the one dimensional distributions converge geometrically fast. On the other hand, the Markov structure implies that if indeed the one dimensional distributions converge, then the distribution of the whole process converges:

Theorem 12. *Suppose that for all i and j we have $\lim_{t\to\infty} \mathrm{p}_{ij}^{(t)} = \pi_j$, for some probability distribution π. Then π is an invariant probability distribution, and for any i,*

$$\lim_{t\to\infty} \mathbb{P}_i\left(X_{t+1} = j_1, X_{t+2} = j_2 \ldots\right) = \mathbb{P}_\pi\left(X_1 = j_1, X_2 = j_2 \ldots\right)$$

where $\mathbb{P}_\pi$ is obtained by starting the process with the distribution π (in fact, the distribution of the process converges).

7. Ergodic Theorems

We do not expect the Markov chain to converge: since transition probabilities are homogeneous, the probability of leaving a given state does not change in time. However, in analogy with i.i.d. random variables, there are limit theorems under the right scaling. The connection to the i.i.d. case comes from the following construction. Fix an arbitrary state j and define

$$\begin{aligned} R_1 &= T_1 = \tau_j \\ R_k &= \inf\{t > R_{k-1} : X_t = j, \quad k > 1\} \\ T_k &= R_k - R_{k-1}, \quad k > 1. \end{aligned}$$

Theorem 13. *If j is recurrent and the Markov chain starts at j (with probability one), then $T_1, T_2 \ldots$ is a sequence of i.i.d. random variables. Moreover, the random vectors*

$$Z_k \triangleq \{X_{R_k}, X_{R_k+1}, \ldots, X_{R_{k+1}-1}\}$$

are independent and identically distributed (in the space of sequences of variable length!).

This is another manifestation of the fact that, once we know the Markov chain hits some state, future behavior is (probabilistically) determined. Fix a recurrent state i and a time t, and define

$$T^t \triangleq \max_k \{T_k : T_k \le t\}.$$

Denote by $N_s(j)$ the number of times in $1, 2, \ldots, s$ that $X_u = j$. By Theorem 13, the random variables $\{N_{R_k+1}(j) - N_{R_k}(j),\ k = 1, 2, \ldots\}$ are independent, and (except possibly for $k = 1$) are identically distributed for each j. By the law of large numbers this implies the following.

Theorem 14. *Let i be recurrent. Then starting at i ($\mathbb{P}_i$ a.s.)*

$$\lim_{t\to\infty} \frac{N_t(j)}{N_t(\ell)} = \frac{\mathbb{E}_i \sum_{s=1}^{\tau_i} \mathbf{1}_{X_s=j}}{\mathbb{E}_i \sum_{s=1}^{\tau_i} \mathbf{1}_{X_s=\ell}} = \frac{\pi_j}{\pi_\ell} \ ,$$

and π is an invariant probability distribution, concentrated on the closed class containing i, that is,

$$\pi_k = \sum_{j\in\mathbf{S}} \pi_j \mathrm{p}_{jk}, \qquad \sum_{k:i\to k\in\mathbf{S}} \pi_k = 1$$

so that $\pi_k = 0$ if $i \not\to k$. Moreover, if we start in some state j then

$$\lim_{t\to\infty} \frac{N_t(i)}{t} = \frac{\mathbf{1}_{\tau_i<\infty}}{m_{ii}} \quad \text{with } \mathbb{P}_j\text{-probability 1.}$$

The last relation implies (12): taking $\mathbb{E}_j$ expectations, we obtain that if i is recurrent then $\mathrm{P}^* = \mathbb{P}_j(\tau_i < \infty)/m_{ii}$. From here follows a limit theorem for functions of a Markov chain.

Theorem 15 Ergodic Theorem. *Let* $\mathbf{S}$ *be a single recurrent class, and assume* π *is an invariant probability distribution. Let* f *and* g *be functions such that* $\mathbb{E}_\pi |f(X_0)| < \infty$ *and* $\mathbb{E}_\pi |g(X_0)| < \infty$. *Then for an arbitrary starting state* i, *with probability one,*

$$\lim_{t\to\infty} \frac{\sum_{s=0}^{t} f(X_s)}{\sum_{s=0}^{t} g(X_s)} = \frac{\mathbb{E}_i \sum_{s=1}^{\tau_i} f(X_s)}{\mathbb{E}_i \sum_{s=1}^{\tau_i} g(X_s)} = \frac{\mathbb{E}_\pi f(X_0)}{\mathbb{E}_\pi g(X_0)}$$

provided not both numerator and denominator of the last terms are zero.

Setting $g \equiv 1$ we obtain a law of large numbers for a function of the Markov chain. Here is a statement of a Central Limit theorem and a Law of Iterated Logarithm.

Theorem 16. *Let* $\mathbf{S}$ *be a single positive recurrent class with an invariant probability distribution* π. *Fix a function* f *and suppose that for some* i, $\mathbb{E}_i\left[\tau_i^2\right] < \infty$ *and*

$$\mathbb{E}_\pi f(X_0) = 0 \quad \textit{and} \quad \mathbb{E}_i\left[\left(\sum_{t=1}^{\tau_i} |f(x_t)|\right)^2\right] < \infty.$$

Define $\gamma^2 \triangleq \pi_i \cdot \mathbb{E}_i\left[\left(\sum_{t=1}^{\tau_i} f(x_t)\right)^2\right]$. *Then* $\gamma^2 < \infty$, *and if* $\gamma^2 > 0$ *then*

CLT
$$\frac{\sum_{s=0}^{t} f(X_s)}{\sqrt{t\cdot\gamma^2}}$$

converges in distribution (as $t \to \infty$*) to a standard Gaussian random variable. Moreover, the sum satisfies the law of iterated logarithm, that is, the lim sup of*

LIL
$$\frac{\sum_{s=0}^{t} f(X_s)}{\sqrt{2\gamma^2 t \log\log t}},$$

as $t \to \infty$, *is* 1 *and the lim inf is* (-1).

8. Extensions and Comments

Markov processes are very general objects, of which our treatment covered just a fraction. Many of the basic ideas extend, but definitely not all, and usually some effort is required. In addition, the mathematical difficulties rise exponentially fast. There are two obvious directions to extend: more general state spaces, and continuous time.

When the state space is not countable, the probability that an arbitrary point in the state space is "visited" by the process is usually zero. Therefore the notions of communication, hitting, recurrence and periodicity have to be modified. The most extensive reference here is [MT].

In the discrete-space continuous time setting, the Markov property implies that if $x_t = i$ then the values of $x_s,\ s < t$ are not relevant. In particular, the length of time from the last change in state until t should be irrelevant. This implies that the distribution of the time between jumps (= change in state) should be exponential. If the only possible transition is then from i to $i+1$ and if all transition times have the same distribution (that is, they are all exponential with the same parameter), then we obtain the Poisson process. More generally we can describe most discrete-state, continuous time Markov chains as follows. The process stays at state i an exponential amount of time with parameter $\lambda(i)$. It then jumps to the next state according to a transition probability p_{ij}, and the procedure is repeated (this is correct if, for example, $\lambda(i) > \lambda > 0$ for all i). This subject is covered, for example in [Ch]. If we observe such a process at jump times, then we recover a Markov chain. This is one of the major tools in the analysis of continuous time chains.

Semi-Markov processes are a further generalization, where the time between events is drawn from a general distribution, which depends on the state, and possibly on the next state. This is no longer a Markov process; however, if we observe the process only at jump times, then we recover a Markov chain.

Finally, in applications, the information structure, and consequently the set of events, is richer: we can measure more than the values of the Markov chain. This is often manifested in a recursion of the type (2), but where the V_t are not independent. Do we still get a Markov chain? and in what sense? The rough answer is that, if the Markov property (4) holds, but where we condition on all the available information, then we are back on track: all of our results continue to hold. For this to happen we need the "noise sequence" $V_0, V_1, \ldots$ to be non anticipative in a probabilistic sense.

Criteria for Stability. As in the case of dynamical systems, there are criteria for stability and for recurrence, based on Lyapunov functions. This is one of the main tools in [MT]. These techniques are often the easiest and the most powerful.

Relevance to Control. Many models of control systems subject to noise can be modelled as Markov processes, and the discrete-time, discrete-space models are controlled Markov chains. Here the transition probabilities are parameterized by the control: see the section on Dynamic Programming. In addition, many filtering and identification algorithms give rise to Markov chains (usually with values in $\mathbb{R}^d$). Limit theorems for Markov chains can then be used to analyze the limiting properties of these algorithms. See, for example [MT, Ti].

Example 9. Extending Example 2, consider the recursion

$$X_{t+1} \triangleq X_t + V_t + U_t$$

where U_t is a control variable. This is a simple instance of a controlled recursion of the ARMA type. Suppose the U_t can only take the values ± 1. Of course, we require that the control depends only on past information. If the control values U_t depend on the past states, then $X_0, X_1, \ldots$ may not be a Markov chain. For example, if we choose $U_t = \operatorname{sign}(X_0)$, then the sequence $X_0, X_1, \ldots$ violates the Markov property (Definition 1). However, we do have a *controlled Markov chain.* This means that Definition 1 is replaced with the relation

$$\begin{aligned} &\mathbb{P}\left(X_t = j \mid X_{t-1} = i_{t-1}, \ldots, X_0 = i_0,\, U_{t-1} = u_{t-1}, \ldots, U_0 = u_0\right) \\ &\qquad\qquad = \mathbb{P}\left(X_t = j \mid X_{t-1} = i_{t-1},\, U_{t-1} = u_{t-1}\right). \end{aligned}$$

This in fact is the general definition of a controlled Markov chain. If we choose a feedback control, that is, $U_t = f(X_t)$ for some function f, then $X_0, X_1, \ldots$ is again a Markov chain; but the transitions and the limit behavior now depend on the choice of f. For more information on controlled Markov chains, see the section on Dynamic Programming.

References

Markov chains are covered by most introductory texts on stochastic processes. Here are some more specific references.

[Ch] K.L. Chung, *Markov chains with stationary transition probabilities*, Springer Verlag, New York, 1967.

A classic on discrete-space Markov chains, both discrete and continuous time. Very thorough and detailed, but mathematically not elementary.

[Çi] E. Çinlar, *Introduction to Stochastic Processes*, Prentice Hall, 1975.

A precise, elegant and accessible book, covering the basics.

[KS] J.G. Kemeny, J. Snell and A.W. Knapp, *Denumerable Markov Chains*, Van Nostrand, Princeton, NJ, 1966.

The most elementary in this list, but fairly thorough, not only in coverage of Markov chains, but also as introduction to Markov processes.

[MT] S.P. Meyn and R.L. Tweedie, *Markov Chains and Stochastic Stability*, Springer Verlag, London 1993.

Deals with general discrete-time Markov chains, and covers the state of the art. It is therefore demanding mathematically, although not much measure theory is required. The most comprehensive book if you can handle it.

[Nu] E. Nummelin, *General Irreducible Markov Chains and Non Negative Operators*, Cambridge University Press, 1984.

More general than [Se], treats general state spaces. Introduced many new techniques; perhaps less encyclopedic than [MT], but fairly mathematical.

[Or] S. Orey, *Limit Theorems for Markov Chain Transition Probabilities*, Van Nostrand Reinhold, London, 1971.

A thin gem on limit theorems, fairly accessible.

[Re] D. Revuz, *Markov Chains*, North Holland, Amsterdam, 1984.

A mathematical, thorough treatment.

[Se] E. Seneta, *Non Negative Matrices and Markov Chains*, Springer Verlag, New York, 1981.

Gives the algebraic point of view on Markov chains.

[Ti] H.C. Tijms, *Stochastic Modelling and Analysis: a Computational Approach*, John Wiley, New York, 1986.

Contains a basic introduction to the subject of Markov chains, both discrete and continuous time, with a wealth of examples for applications and computations.