Section 2: Discrete Time Markov Chains

Contents

2.1 Examples of Discrete State Space Markov Chains

The theory of Markov chains is, not surprisingly, most straightforward when the state space S is discrete (i.e. finite or countably infinite). Henceforth, we shall focus exclusively here on such discrete state space discrete-time Markov chains (DTMC's). We devote this section to introducing some examples.

Example 2.1.1 (Gambler's ruin) Suppose that a gambler starts with initial wealth x, and wagers one dollar on each bet. The gambler wins a bet with probability p and loses with probability $q \triangleq 1 - p$. The game terminates either when the gambler ruins (i.e. wealth = 0) or when the gambler wins the house limit (i.e. wealth = N). The transition graph for this (time-homogeneous) Markov chain corresponding to the number of wagers is given by:

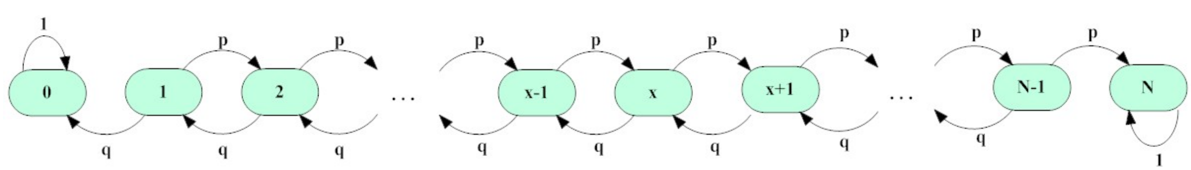

Example 2.1.2 (Ehrenfest chain) This model, introduced by the physicist Ehrenfest, is a caricature of molecular dynamics. Imagine a room divided into two, with N molecules in total within the room. In every slot of time, a molecule is chosen uniformly and at random and moved to the opposite side. Let X_n be the number of molecules on Side 1 at time n. Then, $X = (X_n : n \geq 0)$ is a DTMC on state space $S = \{0, 1, \ldots, N\}$, with transition graph:

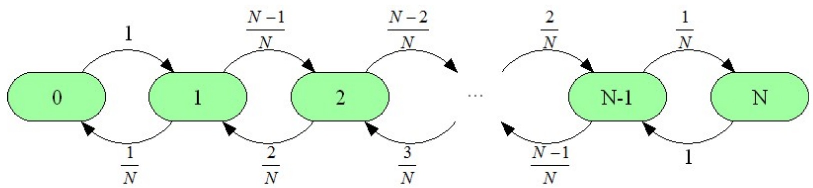

Example 2.1.3 The above two examples are so-called "birth-death" Markov chains. A *birth-death* chain is a chain taking values in a subset of $\mathbb{Z}$ (often $\mathbb{Z}_+$), permitting only one-step transition to nearest neighbors. The transition graph of the general birth-death chain on $\mathbb{Z}_+$ is:

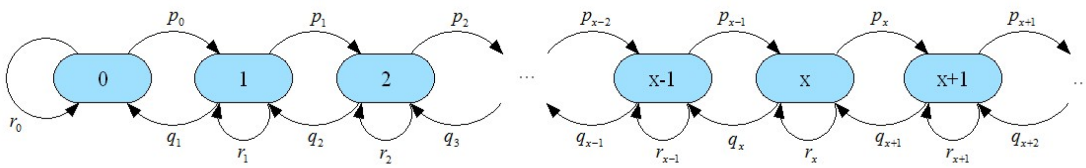

Example 2.1.4 Let $Z = (Z_n : n \geq 1)$ be a sequence of iid $\mathbb{R}^d$-valued random variables. Consider the sequence $S = (S_n : n \geq 0)$ defined through the recursion

$$S_{n+1} = S_n + Z_{n+1},$$

where S_0 is independent of Z. Such a Markov chain is called a *random walk*. When $S_0 \in \mathbb{Z}^d$ and the Z_n's are $\mathbb{Z}^d$-valued, then S is a Markov chain on $\mathbb{Z}^d$. Furthermore, if $d = 1$ and $\mathrm{P}(Z_n = 1) = p = 1 - \mathrm{P}(Z_n = -1)$, S is a birth-death chain on $\mathbb{Z}$ with transition graph:

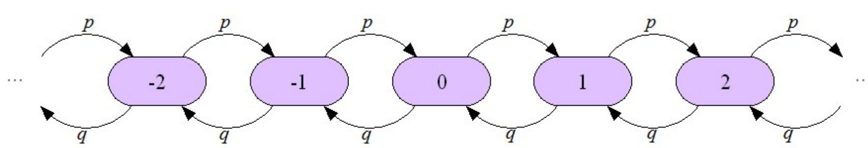

This birth-death chain is called *a nearest-neighbor random walk.*

Example 2.1.5 In most real-world retail settings, inventory can be carried over to future time periods to satisfy future demand. In the presence of reasonable assumptions on building costs, penalty costs, and fixed ordering costs, the optimal ordering policy can be characterized, and the resulting inventory position process $X = (X_n : n \geq 0)$ is of (s, S) form. (In other words, the associated stochastic control problem has an optimal control that can be explicitly characterized, and is of (s, S) type.) To describe the (s, S) policy, let D_i be the demand for the product in $(i-1, i]$. Assume the demand D_i is realized immediately after time $i - 1$. We first attempt to satisfies the demand from on-hand inventory. If the resulting inventory position $X_{i-1} - D_i$ is less than the *re-order level* s, we immediately place an order to bring the resulting inventory passing up to level S, and we assume that the order is satisfied immediately. It follows that X satisfies the stochastic recursion

$$X_{n+1} = \begin{cases} X_n - D_{n+1}, \ X_n - D_{n+1} \geq s \\ S, \ X_n - D_{n+1} < s \end{cases}$$

It follows that if the D_n's are iid, X is a Markov chain with one-step transition matrix P given by

$$P = \begin{array}{c} \\ s \\ s+1 \\ \vdots \\ S \end{array} \begin{array}{c} \begin{array}{ccccc} s & s+1 & s+2 & \dots & S \end{array} \\ \begin{bmatrix} p_0 & 0 & 0 & \dots & 1-p_0 \\ p_1 & p_0 & 0 & \dots & 1-p_0-p_1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ p_{S-s} & p_{S-s-1} & p_{S-s-2} & \dots & p_0 + \sum_{j>S-s} p_j \end{bmatrix} \end{array}$$

where $p_j = P(D_i = j)$.

Exercise 2.1.1 A common assumption in the literature is to presume that the D_n's are Poisson distributed. This question explores the underlying validity of this assumption.

a.) Suppose that Z_i is Bernoulli(p_i). Prove that one can find a (jointly distributed) Poisson rv $\widetilde{Z}_i$ such that
$$\mathrm{P}(Z_i \neq \widetilde{Z}_i) = O(p_i^2).$$

b.) Suppose that $(Z_i : 1 \leq i \leq n)$ is a collection of independent Bernoulli(p_i) rv's. Prove that one can find a (jointly distributed) Poisson rv $\widetilde{D}$ such that
$$\mathrm{P}\left(\widetilde{D} \neq \sum_{i=1}^{n} Z_i\right) = O\left(\sum_{i=1}^{n} p_i^2\right)$$

c.) For each $n \geq 1$, let $Z_{n,i}$, $1 \leq i \leq n$ be independent Bernoulli($p_{n,i}$) r.v.'s. Suppose $\sum_{i=1}^{n} p_{n,i} \to \lambda \in (0, \infty)$ and $\max_{1 \leq i \leq n} p_{n,i} \to 0$ as $n \to \infty$. Prove that
$$\sum_{i=1}^{n} Z_{n,i} \Rightarrow \text{Poisson}(\lambda).$$

d.) How does c.) support the use of the Poisson distribution in the inventory setting?

Example 2.1.6 Consider a finite connected graph with vertex set V and (undirected) edges E. A "random walk" moves from vertex to vertex according to the following rule: Choose uniformly and at random from the edges that connect to the currently occupied vertex, and move to the vertex associated with that edge. Let X_n be the vertex occupied at time n by the random walk. Then, $X = (X_n : n \geq 0)$ is a Markov chain on $S = V$.

Example 2.1.7 The following Markov chain arises in connection with the "page-rank" algorithm used by Google to rank webpages as part of its search engine. Think of each website as a vertex in a graph. Vertex v has a directed edge to vertex w if there is a link to website w from website v. Consider a "random walk" that moves from vertex to vertex by choosing uniformly and at random from the directed edges available to the walker in the currently occupied vertex. The Markov chain $X = (X_n : n \geq 0)$ then describes the position of the random walk as a function of time.

Example 2.1.8 (Queueing chain) Consider a single server queue that operates in "slotted time". Let X_n be the number of customers in the system at the start of slot n. Immediately after the start of slot n, Z_{n+1} customers arrive. The server serves precisely one customer per slot of time, assuming that one is available to serve. Assuming that the Z_i's are iid (and independent of X_0), it follows that $X = (X_n : n \geq 0)$ is a Markov chain on $\mathbb{Z}_+$ that obeys the recursion

$$X_{n+1} = [X_n + Z_{n+1} - 1]^+,$$

where $[y]^+ \triangleq \max(y, 0)$.

Example 2.1.9 (Branching chain) Suppose that X_n represents the number of individuals in a population at generation n. Each individual i in generation n independently produces $Z_{n+1,i}$ progeny in generation $n+1$. It follows that

$$X_{n+1} = \sum_{i=1}^{X_n} Z_{n+1,i}.$$

If the $(Z_{n,i} : n \geq 1, i \geq 1)$ is a family of iid rv's, then $X = (X_n : n \geq 0)$ is a Markov chain taking values in $\mathbb{Z}_+$.

Example 2.1.10 Consider a population of $N/2$ diploid individuals who have two copies of each of their chromosomes, or of N haploid individuals who have one copy. In either case, we are led to consideration of a population of N genes that can be one of two types: A or a. In the *Wright-Fisher model*, the number of genes X_{n+1} of type A at time $n+1$ is obtained by drawing (with replacement) from the population at time n. In other words, conditional on X_n,

$$X_{n+1} \stackrel{\mathcal{D}}{=} Bin(N, X_n/N).$$

Exercise 2.1.2 a.) Prove that if X is the solution of a time-homogeneous stochastic recursion, then X is a time-homogeneous Markov chain.

b.) Prove that any discrete state space time-homogeneous Markov chain can be represented as the solution of a time-homogeneous stochastic recursion.

2.2 Stochastic Processes

Recall that a stochastic process is a family of random variables that is intended to model a time dependent stochastically evolving dynamical system. More precisely, given a sample space Ω, a *stochastic process* is a mapping $X : T \times \Omega \to S$, where T is the time parameter set (usually, one of $\mathbb{Z}_+$, $\mathbb{Z}$, $\mathbb{R}_+$, or $\mathbb{R}$) and S is the "state space" of X. Following probability convention, we shall often suppress the ω-dependence and write $X = (X(t) : t \in T)$ (in order to reduce the notational burden). When the system evolves in discrete time, we often write X_n rather than $X(n)$.

Note that $X : (t, \omega) \to X(t, \omega)$ is properly viewed as a function of both t and ω. In elementary probability, one frequently focuses one's attention on statements that involve fixing t and looking at properties of $X(t, \omega)$ as a function of ω (e.g. $\mathrm{P}(X(t) \in A) = \mathrm{P}(\{\omega : X(t, \omega) \in A\})$). But it is also meaningful (and interesting) to study properties of X for fixed ω (e.g. $\mathrm{P}(X(\cdot, \omega))$ is a continuous function on $[0, \infty)$).

When the index set for X is no longer time but instead has a spatial interpretation, it is standard to call the associated random object a random field. More precisely, a *random field* is a mapping $X : \Lambda \times \Omega \to S$, where $\Lambda \subseteq \mathbb{R}^d$ with $d \geq 2$.

2.3 Stochastic Recursions

A very useful class of discrete-time deterministic dynamical systems is the class defined as solutions of recursive equations of the form

$$x_{n+1} = f_{n+1}(x_n) \tag{2.3.1}$$

for $n \geq 0$, where $(f_n : n \geq 1)$ is a sequence of functions $f_n : S \to S$ provided as a "model primitive" by the modeler. The natural stochastic analog of (2.3.1) is a sequence $(X_n : n \geq 0)$ defined as the solution of a *stochastic recursion* in which

$$X_{n+1} = f_{n+1}(X_n, Z_{n+1}) \tag{2.3.2}$$

for $n \geq 0$, where $(f_n : n \geq 1)$ is a sequence of (deterministic) functions $f_n : S \times S' \to S$, $Z = (Z_n : n \geq 1)$ is a sequence of independent S'-valued random variables, and Z is independent of X_0. We say that $X = (X_n : n \geq 0)$ is *time-homogeneous* if $f_n \equiv f$ for $n \geq 1$ and the Z_n's are identically distributed; otherwise, X is *time-inhomogeneous.* Time-inhomogeneous models arise naturally in settings where one wishes to explicitly model "time-of-day effects" or "seasonality effects" or economic growth over time as exogenous variables.

Exercise 2.3.1 Suppose $S = \mathbb{R}$. Prove that any stochastic recursion of the form (2.3.2) can also be written as the solution of

$$X_{n+1} = \widetilde{f}_{n+1}(X_n, \widetilde{Z}_{n+1})$$

for suitably defined $(\widetilde{f}_n : n \geq 1)$ and $\widetilde{Z} = (\widetilde{Z}_n : n \geq 1)$, where the $\widetilde{Z}_n$'s are iid.

Solutions of stochastic recursions of the form (2.3.2) arise naturally in many different modeling contexts and are the prevalent models used within the physical sciences to describe physical phenomena. However, they also arise naturally in the setting of the social sciences, economics, and management science. Furthermore, with increasing interest in randomized distributed algorithms that often take the form (2.3.2), such models are often ubiquitous within performance engineering applications arising in computer science and electrical engineering. In short, stochastic recursions are fundamental to the study of many different applications domains.

2.4 The Markov Property

A key mathematical property of a stochastic recursion is that such processes X enjoy the Markov property.

Definition 2.4.1 An S-valued stochastic sequence $X = (X_n : n \geq 0)$ is said to have the *Markov property* if for each $n \geq 0$ and (measurable) subset $A \subseteq S$,

$$\mathrm{P}(X_{n+1} \in A | X_0 = x_0, \ldots, X_{n-1} = x_{n-1}, X_n = x) = \mathrm{P}(X_{n+1} \in A | X_n = x) \tag{2.4.1}$$

for all $(x_0, x_1, \ldots, x_{n-1}, x) \in S^{n+1}$.

The probability $\mathrm{P}(X_{n+1} \in A | X_n = x)$ is often called the one-step transition probability (associated with time $n+1$) and is usually denoted as $P(n+1, x, A)$.

Remark 2.4.1 When S is discrete, the one-step transition structure of X at time $n+1$ can be summarized through a square matrix $P(n+1) = (P(n+1,x,y) : x, y \in S)$, where $P(n+1,x,y) = \mathrm{P}(X_{n+1} = y | X_n = x)$.

Remark 2.4.2 To reduce the notational burden associated with writing conditional probabilities and expectations, we will write

$$\mathrm{P}(X_{n+1} = y | X_0, \ldots, X_n)$$

and

$$\mathrm{E}[g(X_{n+1}) | X_0, \ldots, X_n]$$

instead of

$$\mathrm{P}(X_{n+1} = y | X_0 = x_0, \ldots, X_n = x_n)$$

and

$$\mathrm{E}[g(X_{n+1}) | X_0 = x_0, \ldots, X_n = x_n].$$

With this notation in hand, the Markov property reads as

$$\mathrm{P}(X_{n+1} \in A | X_0, \ldots, X_n) = \mathrm{P}(X_{n+1} \in A | X_n).$$

As we shall see later in the course, use of this notational convention is actually a more accurate reflection of the general definition of conditional expectation (to cover settings in which we wish to condition on an infinite number of random variables).

The Markov property can be reformulated in several mathematically equivalent ways.

Proposition 2.4.1 An S-valued stochastic sequence $X = (X_n : n \geq 0)$ has the Markov property (or, simply, is *Markov*) if for each non-negative function g,

$$\mathrm{E}[g(X_{n+1}, X_{n+2}, \ldots) | X_0, \ldots, X_n] = \mathrm{E}[g(X_{n+1}, X_{n+2}, \ldots) | X_n]. \tag{2.4.2}$$

Such a Markov sequence X is called a (discrete-time) *Markov chain*.

Exercise 2.4.1 Prove Proposition 1.3.1.

Proposition 2.4.2 An S-valued stochastic sequence $X = (X_n : n \geq 0)$ is Markov if for each n,

$$\begin{aligned} &\mathrm{P}(X_0 = x_0, \ldots, X_{n-1} = x_{n-1}, X_{n+1} = x_{n+1}, \ldots, X_{n+m} = x_{n+m} | X_n) \\ &= \mathrm{P}(X_0 = x_0, \ldots, X_{n-1} = x_{n-1} | X_n) \mathrm{P}(X_{n+1} = x_{n+1}, \ldots, X_{n+m} = x_{n+m} | X_n) \end{aligned} \tag{2.4.3}$$

for all $(x_0, \ldots, x_{n-1}, x_{n+1}, \ldots, x_{n+m}) \in S^{n+m}$.

Exercise 2.4.2 Prove Proposition 1.3.2.

Property (2.4.3) is often stated as: X is Markov if and only if the "past" and the "future" are conditionally independent, given the "present".

The relation (2.4.2) can be generalized to certain random times. To get a sense of the generalization, write (2.4.2) as

$$\mathrm{E}[g(X_{n+1}, X_{n+2}, \ldots) | X_0, \ldots, X_n] = u_n(X_n),$$

where

$$u_n(x) = \mathrm{E}[g(X_{n+1}, X_{n+2}, \ldots)|X_n = x].$$

The generalization reads as:

$$\mathrm{E}[g(X_{T+1}, X_{T+2}, \ldots)|X_0, \ldots, X_T] = u_T(X_T) \tag{2.4.4}$$

for a given random time T. Unfortunately, (2.4.4) is, in general, not correct. Consider, for example, the setting in which

$$g(x_1, x_2, \ldots) = \begin{cases} 1 & \text{if } x_1 \in A \\ 0 & \text{else,} \end{cases}$$

and in which $T = \inf\{n \geq 0 : X_{n+1} \in A\}$. Note that by definition of T, $X_{T+1} \in A$ so the left-hand side of (2.4.4) equals one. On the other hand, $u_n(x)$ will usually be less than one for $n \geq 0$ and $x \in S$. So, $u_T(X_T)$ will typically be less than one. The difficulty here turns out to be related to the fact that the particular T we have chosen is "clairvoyant", in the sense that T is peeking one time unit into the future. If we restrict ourselves to random times that are not clairvoyant (i.e. "non-anticipating"), then it turns out (2.4.4) holds.

Definition 2.4.2 A rv T is said to be a *random time* if $T : \Omega \to \mathbb{Z}_+ \bigcup \{\infty\}$. A random time T is finite-valued if $\mathrm{P}(T < \infty) = 1$. Furthermore, a random time T is said to be a *stopping time* (adapted to a sequence $W = (W_n : n \geq 0)$) if for each $n \geq 0$, there exists a (deterministic) function k_n such that

$$\mathrm{I}(T = n) = k_n(W_0, W_1, \ldots, W_n).$$

The notion of "stopping time" encapsulates the concept of non-clairvoyance.

Proposition 2.4.3 Let $X = (X_n : n \geq 0)$ be an S-valued Markov chain, and let T be a stopping time adapted to X. Then, for each non-negative g,

$$\mathrm{E}[g(X_{T+1}, X_{T+2}, \ldots)|X_0, \ldots, X_T] = u_T(X_T)$$

on $\{T < \infty\}$.

Exercise 2.4.3 Prove Proposition 1.3.3.

The fact that property (2.4.4) holds for stopping times is called the *strong Markov property.*

Remark 2.4.3 The Markov property allows us to rigorously define the notion of "state". Given a stochastic dynamical system, the appropriate definition of state for the system is that under which the dynamics become Markovian.

2.5 Computing the Transient Distribution

As we shall see, a key property of a Markov chain is the fact that many probabilities and expectations of interest can be computed as solutions of associated linear equations. We start with the computation of the n-step transition probabilities. Because many applications naturally lead to time-dependent formulations, we develop our methodology in the general setting of a time-inhomogeneous Markov chain living on discrete state space S. In such a setting, the one-step transition structure of $X = (X_n : n \geq 0)$ is characterized by a sequence of matrices $P(1), P(2), \ldots$. Then,

$$\begin{aligned} & \mathrm{P}(X_{n+i} = y | X_i = x) \\ = & \sum_{z_1, z_2, \ldots, z_{n-1}} P(i+1, x, z_1) P(i+2, z_1, z_2) \cdots P(i+n, z_{n-1}, y) \\ = & (P(i+1)P(i+2) \cdots P(i+n))(x, y), \end{aligned} \tag{2.5.1}$$

so that $\mathrm{P}(X_{n+i} = y | X_i = x)$ can be computed by calculating the (x, y)'th entry of the matrix product $P(i+1)P(i+2) \cdots P(i+n)$.

Suppose that we generalize the notion of time-inhomogeneity so that the state space S_i at time i can depend on i. In this case, the transition matrix $P(i) = (P(i, x, y) : x \in S_{i-1}, y \in S_i)$ is rectangular (rather than square).

Exercise 2.5.1 a.) Discuss the computation of $\mathrm{P}(X_{n+i} = y | X_i = x)$ in this setting.

b.) Where might this generalization be useful?

Remark 2.5.1 In principle, any time-inhomogeneous Markov chain can be viewed as a time-homogeneous Markov chain. If $(X_n : n \geq 0)$ is time-inhomogeneous, then $(Y_n : n \geq 0)$ is time-homogeneous, where $Y_n = (X_n, n)$; the Markov chain $Y = (Y_n :\geq 0)$ is called the *space-time* Markov chain. While the space-time is conceptually useful, it would be computationally disastrous to implement naively, since the state space of Y is $T \times S$ (so that the matrices involved would have enormous dimensions). Instead, the formula (2.5.1) involves a matrix product in which the individual matrices that appear are $|S| \times |S|$ (rather than $(|S| \times |T|) \times (|S| \times |T|)$); this is much more efficient than naively implementing the space-time computations.

Remark 2.5.2 When $P(i) \equiv P$ for $i \geq 1$, the transition probabilities are stationary, in the sense that

$$\mathrm{P}(X_{n+i} = y | X_i = x) = \mathrm{P}(X_n = y | X_0 = x)$$

for $i \geq 0$. For this reason, time-homogeneous Markov chains are often called *Markov chains with stationary transition probabilities.*

Remark 2.5.3 For time-homogeneous Markov chains,

$$\mathrm{P}(X_n = y | X_0 = x) = P^n(x, y),$$

where P^n is the n'th power of P. If one can compute the spectrum of P (i.e. the set of eigenvalues of P), as well as the corresponding eigenvectors, then P can be represented as

$$P = R^{-1} J R,$$

where J is the so-called Jordan form associated with P. In the special case that P is diagonalizable, J is a diagonal matrix. Since J^n can be easily computed, this provides a vehicle for efficiently computing P^n:

$$P^n = R^{-1} J^n R.$$

Unfortunately, for most Markov chain models, computing the spectrum is extremely difficult (both analytically and numerically).

To compute the unconditional distribution $P(X_n \in \cdot)$, assume that there exists an initial distribution $\mu = (\mu(x) : x \in S)$, where $\mu(x) = \mathrm{P}(X_0 = x)$. We shall choose to encode all probability mass functions in this course as row vectors. Given this convention, μ can then be viewed as a row vector. Set

$$\mu_n(x) = \mathrm{P}(X_n = x)$$

for $n \geq 0$ and $x \in S$; again, we represent $\mu_n = (\mu_n(x) : x \in S)$ as a row vector. Note that

$$\begin{aligned}
\mathrm{P}(X_n = y) &= \sum_x \mathrm{P}(X_n = y | X_{n-1} = x) P(X_{n-1} = x) \\
&= \sum_x P(n, x, y) \mu_{n-1}(x) \\
&= (\mu_{n-1} P(n))(y),
\end{aligned}$$

and consequently

$$\mu_n = \mu_{n-1} P(n).$$

The following proposition follows immediately.

Proposition 2.5.1 The row vector μ_n can be recursively solved via the recursion

$$\mu_i = \mu_{i-1} P(i) \tag{2.5.2}$$

for $1 \leq i \leq n$, subject to the initial condition $\mu_0 = \mu$.

Because (2.5.2) is solved recursively by going forward in time, (2.5.2) is called the *forward equations* for the Markov chain.

An alternative (but related) computation is the computation of $\mathrm{E}[r(X_{i+1}) | X_i = x]$ for a given function $r : S \to \mathbb{R}$. Such calculations arise in many settings and it is frequently conceptually useful to think of r as a "reward function". Our convention, in this course, will be to encode all reward functions as column vectors. Set

$$u(i, i+n, x) = \mathrm{E}[r(X_{i+1}) | X_i = x];$$

we view the function $u(i, i+n) = (u(i, i+n, x) : x \in S)$ as a column vector. Note that

$$\begin{aligned}
\mathrm{E}[r(X_{i+1}) | X_i = x] &= \sum_y \mathrm{E}[r(X_{i+1}) | X_{i+1} = y, X_i = x] \cdot \mathrm{P}(X_{i+1} = y | X_i = x) \\
&= u(i+1, i+n, y) P(i+1, x, y) \\
&= (P(i+1) u(i+1, i+n))(x),
\end{aligned}$$

and consequently,

$$u(i+1, i+n) = P(i+1) u(i, i+n)$$

We therefore have established the following proposition.

Proposition 2.5.2 The sequence of column vectors $(u(i, n) : 0 \leq i \leq n)$ can be computed via the recursion

$$u(i, n) = P(i+1) u(i+1, n) \tag{2.5.3}$$

for $0 \leq i \leq n-1$, subject to the terminal condition $u(n, n) = r$.

Because (2.5.3) is solved recursively by going backwards in time, (2.5.3) is called the *backwards equations* for the Markov chain.

Remark 2.5.4 The usual numerical procedure for computing μ_n or $u(0, n)$ involves using the above recursions (even when the Markov chain is time-homogenous). Note that the general complexity of either of these recursions is of order $O(n|S|^2)$ in the non-sparse setting, and can be even much more efficient when the matrices introduced ares sparse (as is typical of most models.)

2.6 First Step Analysis

One of the most powerful ideas in the theory of Markov chains and Markov processes is that many of the key probabilities and expectations can be computed as the solutions of linear equations: linear (matrix) equations for discrete state space Markov chains, linear integral equations for continuous state space Markov chains, and linear partial differential equations for Markov diffusion processes.

We will now illustrate this idea by heuristically deriving appropriate linear systems for a number of key probabilities and expectations that arise in the Markov chain settings. We will focus here on time-homogeneous Markov chains. In this time-homogeneous settings, the notation

$$\mathrm{P}_x(\cdot) \triangleq \mathrm{P}(\cdot|X_0 = x)$$

and

$$\mathrm{E}_x(\cdot) \triangleq \mathrm{E}(\cdot|X_0 = x)$$

is commonly used. More generally, for a given initial distribution μ, we use the notation

$$\mathrm{P}_\mu(\cdot) = \sum_x \mu(x)\mathrm{P}_x(\cdot)$$

and

$$\mathrm{E}_\mu(\cdot) = \sum_x \mu(x)\mathrm{E}_x(\cdot).$$

With this notation in hand, consider the following computations:

Problem 1 (Computing Absorption Probabilities) Given a non-empty subset $C^c \subseteq S$, let $T = \inf\{n \geq 0 : X_n \in C^c\}$ be the "first hitting time" of C^c. For $A \subseteq C^c$, let

$$u^*(x) = \mathrm{P}_x(X_T \in A, T < \infty)$$

be the probability that the chain is "absorbed" into A at the hitting time T, having started the chain at x. We can derive a linear system of equations for the $u^*(x)$'s by conditioning on the first step X_1 taken by the chain:

$$u^*(x) = \sum_y P(x, y)u^*(y), \quad x \in C.$$

Of course, there are obvious "boundary conditions" associated with this problem, namely

$$u^*(z) = \begin{cases} 1 & \text{for } z \in A, \\ 0 & \text{for } z \in C^c - A. \end{cases}$$

We can write the above linear system in matrix-vector notation as: u^* is a solution of

$$u = Pu \qquad \text{on } C \tag{2.6.1}$$

subject to $u = 1$ on A and $u = 0$ on $C^c - A$.

Remark 2.6.1 A function u satisfying $u = Pu$ is said to be a *P-harmonic* function (or, simply, a harmonic function).

The linear system (2.6.1) can also be expressed in a mathematically identical way that is free of boundary conditions:

$$u = f + Bu \tag{2.6.2}$$

on C, where $f(x) \triangleq \mathrm{P}_x(X_1 \in A)$ for $x \in C$ and $B = (P(x,y) : x, y \in C)$ is the restriction of P to "C to C transitions".

Example 2.1.1 (continued). For the gambler's ruin problem, perhaps the most interesting probability to compute is the probability of ruin, namely $\mathrm{P}_x(X_T, T < \infty)$, where $T = \inf\{n \geq 0 : X_n \in \{0, N\}\}$. For this example, (2.6.1) translates to

$$u(x) = pu(x+1) + (1-p)u(x-1)$$

for $1 \leq x \leq N-1$, subject to $u(0) = 1$ and $u(N) = 0$.

Example 2.1.8 (continued). For this example, the key probability of interest is the probability of eventual extinction. Here, we take $C^c = \{0\}$ and $A = C^c$, so that $u^*(x) = \mathrm{P}_x(T < \infty)$. Equation (2.6.2) reduces to

$$u^*(x) = \mathrm{P}(Z_{11} = 0)^x + \sum_{y=1}^{\infty} \mathrm{P}\left(\sum_{i=1}^{x} Z_{1i} = y\right) u^*(y).$$

Problem 2 (Computing Expected Absorption Time) With the same set-up as in Problem 1, let

$$u^*(x) = \mathrm{E}_x T$$

be the expected time to "absorption" into C^c. In this setting, first step analysis establishes that u^* satisfies the linear system

$$u = e + Pu \qquad \text{on } C \tag{2.6.3}$$

subject to $u = 0$ on C^c, where $e = (e(x) : x \in C)$ is the constant function for which $e(x) \equiv 1$.

Example 2.1.1 (continued). Let T be the number of wagers made prior to termination of the game (i.e. either ruin or hitting the house limit N). Here, (2.6.3) takes the form

$$u(x) = 1 + pu(x+1) + (1-p)u(x-1)$$

for $1 \leq x \leq N-1$, subject to $u(0) = u(N) = 0$.

Problem 3 (Computing Expected Reward to Absorption) Again, the set-up is the same as in Problem 1. Given a "reward function" $r : S \to \mathbb{R}_+$, let

$$u^*(x) = \mathrm{E}_x \sum_{i=0}^{T-1} r(X_i)$$

be the expected total reward up to the hitting time T. (Consider, for example, a satellite having an operational lifetime T; the above expectation might then correspond to total revenue generated by the satellite over its lifetime.)

First step analysis leads to the conclusion that u^* satisfies

$$u = r + Pu \qquad \text{on } C \tag{2.6.4}$$

subject to $u = 0$ on C^c. Alternatively, we can represent the above linear system as

$$u = r + Bu \tag{2.6.5}$$

on C, where $B = (P(x,y) : x, y \in C)$ is again the restriction of P to C.

Exercise 2.6.1 Suppose that we wish to compute the k'th moment of the total reward to absorption, so that

$$u_k^*(x) = \mathrm{E}_x \left(\sum_{i=0}^{T-1} r(X_i) \right)^k$$

is the quantity of interest. Use first step analysis to derive a linear system for u_k^* that involves $u_1^*, \ldots, u_{k-1}^*$ (so that $u_k^* = (u_k^*(x) : x \in C)$ can be recursively computed from $u_1^*, \ldots, u_{k-1}^*$).

Problem 4 (Computing the Moment Generating Function of Total Reward to Absorption) For $\theta \in \mathbb{R}$, our goal here is to derive an appropriate linear system for

$$u^*(\theta, x) = \mathrm{E}_x \exp\left(\theta \sum_{i=0}^{T-1} r(X_i) \right),$$

where T is as in Problem 1. First step analysis yields

$$u^*(\theta, x) = \sum_y e^{\theta r(x)} P(x,y) u^*(\theta, y) \tag{2.6.6}$$

subject to $u^*(\theta, \cdot) = 1$ on C^c. A mathematically equivalent linear system is

$$u(\theta) = f(\theta) + G(\theta) u(\theta) \tag{2.6.7}$$

on C, where $f(\theta, x) = \sum_{y \in C^c} e^{\theta r(x)} P(x,y)$, $x \in C$, and $G(\theta) = (G(\theta, x, y) : x, y \in C)$ has entries given by

$$G(\theta, x, y) = e^{\theta r(x)} P(x,y).$$

Exercise 2.6.2 Using the fact that successive moments of $\sum_{i=0}^{T-1} r(X_i)$ can be obtained by successively differentiating its moment generating function, formally differentiate (2.6.7) k times and rederive the recursion of Exercise 2.6.1.

Problem 5 (Computing Expected Infinite Horizon Discounted Reward with Constant Discounting) The presence of "discounting" is used by economists and finance specialists to reflect the "time value" of money, so that one dollar available today is more valuable than receiving one dollar in k time units. Specifically, if one earns a return of $c^{\alpha} - 1$ per unit time, a dollar invested today will be worth $e^{\alpha k}$ dollars in k time units. Hence, the discounted "present value" of a dollar earned in k time units is worth the equivalent of $e^{-\alpha k}$ current dollars.

In a stochastic environment in which the reward earned in time period k is given by $r(X_k)$ (with $X = (X_n : n \geq 0)$ a Markov chain), the present value of the reward is $e^{-\alpha k} r(X_k)$. Hence, the present value of the infinite horizon "cash stream" is

$$\sum_{k=0}^{\infty} e^{-\alpha k} r(X_k)$$

and its expected value is given by

$$u^*(x) = \mathrm{E}_x \sum_{k=0}^{\infty} e^{-\alpha k} r(X_k).$$

When $r \geq 0$, this can be rewritten as

$$u^*(x) = \sum_{k=0}^{\infty} e^{-\alpha k} \mathrm{E}_x r(X_k).$$

First step analysis leads to the conclusion that $u^* = (u^*(x) : x \in S)$ satisfies

$$u = r + e^{-\alpha} P u. \tag{2.6.8}$$

Problem 6 (Computing Expected Infinite Horizon Discounted Reward with Stochastic Discounting) Suppose now that the investment return in period k is $\exp(\alpha(X_k)) - 1$. Then, a dollar invested today will have a value of $\exp\left(\sum_{j=0}^{k-1} \alpha(X_j)\right)$ dollars in period k, so that the present value of a dollar received in period k is worth the equivalent of $\exp\left(-\sum_{j=0}^{k-1} \alpha(X_j)\right)$ current dollars. In this setting, the relevant expected infinite horizon present value of a cash stream $(r(X_n) : n \geq 0)$ is

$$u^*(x) = \mathrm{E}_x \sum_{k=0}^{\infty} \exp\left(-\sum_{j=0}^{k-1} \alpha(X_j)\right) r(X_k).$$

In this setting, first step analysis leads to the conclusion that u^* satisfies

$$u = r + Gu \tag{2.6.9}$$

where $G = (G(x, y) : x, y \in S)$ has entries given by

$$G(x, y) = \exp(-\alpha(x)) P(x, y).$$

Exercise 2.6.3 Suppose that we wish to compute

$$u^*(x) = \mathrm{E}_x \exp\left(-\sum_{j=0}^{T-1} \alpha(X_j)\right) r(X_T) I(T < \infty),$$

where $r \geq 0$ and T is the first hitting time of C^c. Use first step analysis to derive the appropriate linear system satisfied by u^*.

Exercise 2.6.4 Let $X = (X_n : n \geq 0)$ be a Markov chain with non-stationary transition probabilities and set $T = \inf\{n \geq 0 : X_n \in C\}$. Put

$$\alpha = \sum_x \mathrm{E}[T|X_0 = x]\mu(x),$$

where $\mu = (\mu(x) : x \in S)$ is the initial distribution of X

i.) Discuss the efficient numerical computation of α.

ii.) How might matrix norms be useful in this problem? (Think of their potential value in computing error estimates.)

We now illustrate the standard style of argument used to make first step analysis rigorous. Consider, for example, Problem 6, in which the goal is to compute $\mathrm{E}_x R$, where

$$R = \sum_{k=0}^{\infty} \exp\left(-\sum_{j=1}^{k-1} \alpha(X_j)\right) r(X_k).$$

Note that we can rewrite R as

$$R = g(X_0, X_1, \ldots)$$

where g is the (deterministic) map

$$g(x_0, x_1, \ldots) = \sum_{k=0}^{\infty} \exp\left(-\sum_{j=1}^{k-1} \alpha(x_j)\right) r(x_k).$$

With g at hand, observe that

$$R = r(X_0) + \exp(-\alpha(X_0))g(X_1, X_2, \ldots)$$

and hence

$$\begin{aligned}
\mathrm{E}_x R &= r(x) + e^{-\alpha(x)}\mathrm{E}_x g(X_1, X_2, \ldots) \\
&= r(x) + e^{-\alpha(x)} \sum_y P(x, y)\mathrm{E}_y g(X_0, X_1, \ldots) \qquad \text{(via Proposition 1.3.1)} \\
&= r(x) + e^{-\alpha(x)} \sum_y P(x, y)\mathrm{E}_y R.
\end{aligned}$$

2.7 Uniqueness/Existence for First Step Analysis

Consider the queueing chain of Example 1.4.7. Given the presence of x customers in the system currently, one may be interested in calculating the expected time to emptying of the queue, specifically $\mathrm{E}_x T$ (where $T = \inf\{n \geq 0 : X_n = 0\}$). Assume that $\mathrm{P}(Z_n \geq 3) = 0$, so that Z_n is supported on the three values 0, 1, and 2 with corresponding probabilities q, r, and p. The function $u^* = (u^*(x) : x \in S)$ then satisfies

$$u(x) = 1 + pu(x + 1) + ru(x) + qu(x - 1) \tag{2.7.1}$$

for $x \geq 1$, subject to the boundary condition $u(0) = 0$. For $q \neq p$, the general solution of (2.7.1) is

$$u(x) = \frac{x}{q-p} + a\left(\left(\frac{q}{p}\right)^x - 1\right) \tag{2.7.2}$$

where a is an undetermined constant. In other words, (2.7.1) does not have a unique solution. It follows that in applying first step analysis one must study the question of existence/uniqueness for such equations. More precisely, in settings where the associated equations have multiple solutions (as in (2.7.2)), we must know how to identify the probabilistically meaningful solution.

In the presence of non-negative rewards, all the linear systems associated with first step analysis take the form

$$u = f + Gu, \tag{2.7.3}$$

where f and G are non-negative (with G square). Here is what we can generally expect for such linear systems:

i.) the probabilistically meaningful solution u^* of (2.7.3) will be given by

$$u^* = \sum_{k=0}^{\infty} G^k f;$$

ii.) the probabilistically meaningful solution u^* will be the minimal non-negative solution of (2.7.3), in the sense that if $u \geq 0$ is any other solution of (2.7.3), then $u \geq u^*$.

We now utilize the criterion ii.) to identify the probabilistically meaningful solution $u(x) = \mathrm{E}_x T$ associated with (2.7.2). Note that (2.7.2) asserts that if $q > p$, any non-negative solution must have $a \geq 0$; minimality then requires setting $a = 0$. Hence,

$$u^*(x) = \frac{x}{q-p}$$

for $q > p$. On the other hand, if $q < p$, then $x/(q-p)$ tends to $-\infty$ as $x \to \infty$ and $(q/p)^x \to 0$ as $x \to \infty$. Hence, the only choice of a available that leads to a non-negative solution is $a = -\infty$. So, we conclude that

$$u^*(x) = \infty \tag{2.7.4}$$

for $q < p$.

We now proceed to prove ii.) above.

Exercise 2.7.1 In this problem, we compute the probability of extinction for the branching chain introduced earlier. Let $u^*(x) = \mathrm{P}_x(T < \infty)$, where $T = \inf\{n \geq 0 : X_n = 0\}$ is the extinction time.

a.) Write down a linear system satisfied by $(u^*(x) : x \geq 1)$.

b.) Argue that $u^*(x) = \mathrm{P}_1(T < \infty)^x$ for $x \geq 1$.

c.) Prove that $\rho \triangleq \mathrm{P}_1(T < \infty)$ satisfies the equation $\rho = \mathrm{E}\rho^{Z_{11}}$.

d.) Prove that when $\mathrm{E}Z_{11} \leq 1$, $\rho = 1$.

e.) Prove that when $\mathrm{E}Z_{11} > 1$, there exists a unique root $\rho_1 \in (0, 1)$ of the equation $\gamma = \mathrm{E}\gamma^{Z_{11}}$.

f.) Using the fact that u^* is the minimal non-negative solution, prove that $\rho = \rho_1$ when $\mathrm{E}Z_{11} > 1$.

Proposition 2.7.1 Consider the linear system

$$u = f + Gu, \tag{2.7.5}$$

where $f = (f(x) : x \in S)$ is non-negative and $G = (G(x, y) : x, y \in S)$ is non-negative. Then, if

$$v = \sum_{j=0}^{\infty} G^j f,$$

v is a non-negative solution of (2.7.5). Furthermore, v is the minimal non-negative solution of (2.7.5).

Remark 2.7.1 Note that $u \equiv \infty$ is one possible non-negative solution of (2.7.5). As seen in (2.7.4), the possibility that $u^* \equiv \infty$ can arise in many practical problems.

Proof of Proposition 2.7.1 Note first that

$$\begin{aligned} v &= f + G(f + Gf + G^2 f + \cdots) \\ &= f + Gv, \end{aligned}$$

so v is indeed a non-negative solution of (2.7.5). If $u \geq 0$ is any other non-negative solution, then

$$\begin{aligned} u &= f + Gu \\ &= f + G(f + Gu) \\ &= f + Gf + G^2 u \\ &= f + Gf + G^2(f + Gu) \\ &= f + Gf + G^2 f + G^3 u. \end{aligned}$$

Iterating this process k times, we obtain

$$u = f + Gf + \cdots + G^k f + G^{k+1} u.$$

Since $u \geq 0$ and G is non-negative, clearly $G^{k+1} u \geq 0$. Hence,

$$u \geq f + Gf + \cdots + G^k f.$$

Now, let $k \to \infty$, yielding

$$u \geq v. \quad \square$$

To argue part i.) is usually done on a case-by-case basis. Consider, for example, Problem 1. Note that for $x \in S$,

$$\begin{aligned} u^*(x) &= \mathrm{P}_x(X_T \in A, T < \infty) \\ &= \sum_{k=1}^{\infty} \mathrm{P}_x(X_T \in A, T = k) \\ &= \sum_{k=1}^{\infty} \mathrm{P}_x(X_1 \in C, \ldots, X_{k-1} \in C, X_k \in A). \end{aligned}$$

But

$$
\begin{aligned}
& \mathrm{P}_x(X_1 \in C, \ldots, X_{k-1} \in C, X_k \in A) \\
= & \sum_{\substack{z_1,\ldots,z_{k-1} \\ z_i \in C, 1 \le i \le k-1}} P(x,z_1)P(z_1,z_2)\cdots P(z_{k-2},z_{k-1})P_{z_{k-1}}(X_1 \in A) \\
= & \sum_{z_1,\ldots,z_{k-1}} B(x,z_1)B(z_1,z_2)\cdots B(z_{k-2},z_{k-1})f(z_{k-1}) \\
= & (B^{k-1}f)(x).
\end{aligned}
$$

We conclude that

$$u^* = \sum_{k=0}^{\infty} B^k f,$$

as required by i.). The other problems discussed earlier can be similarly handled.

2.8 Matrix and Vector Norms

In Section 2.7, we argued that the probabilistically meaningful u^* is a solution of an equation of the form

$$u = f + Gu$$

and that u^* equals $\sum_{k=0}^{\infty} G^k f$. This raises the question of when $I - G$ is non-singular and has an inverse that can be represented as $\sum_{k=0}^{\infty} G^k f$. While the notion of a matrix inverse is uniquely defined in the setting of finite matrices, there are subtleties that arise in the context of infinite dimensional matrices. To settle these questions carefully requires appealing to matrix and vector norms.

Let $\mathscr{L}$ be an (abstract) vector space. A function $\|\| : \mathscr{L} \to [0,\infty)$ is called a *norm* if:

i.) $\|v\| = 0$ iff $v = 0$;

ii.) $\|v + w\| \le \|v\| + \|w\|$; $v, w \in \mathscr{L}$

iii.) $\|cv\| = |c| \cdot \|v\|$; $c \in \mathbb{R}$, $v \in \mathscr{L}$.

Note that a norm on a vector space $\mathscr{L}$ induces a "distance" between elements of $\mathscr{L}$, namely the distance between v and w can be taken to be $\|v - w\|$.

Definition 2.8.1 We say that $(x_n : n \ge 1)$ *converges* to $x_\infty \in \mathscr{L}$ if $\|x_n - x_\infty\| \to 0$ as $n \to \infty$.

Definition 2.8.2 A sequence $(x_n : n \ge 1)$ of elements in $\mathscr{L}$ is said to be a *Cauchy sequence* if, for each $\epsilon > 0$. there exists $N = N(\epsilon)$ such that for any $m, n \ge N(\epsilon)$, $\|x_n - x_m\| < \epsilon$.

Definition 2.8.3 The normed linear space $\mathscr{L}$ is said to be *complete* if every Cauchy sequence $(x_n : n \ge 1)$ converges to a limit $x_\infty \in \mathscr{L}$.

A complete normed vector space $\mathscr{L}$ is called a *Banach space*. For our current purpose, the most important example of such a Banach space is the following.

Example 2.8.1 Let $w : C \to [1, \infty)$ be a so-called "weight function". Given w, let L_w^∞ be the set of all functions $g : C \to \mathbb{R}$ such that

$$\sup_{x \in C} \frac{|g(x)|}{w(x)} < \infty.$$

Note that L_w^∞ is a vector space. For $g \in L_w^\infty$, set

$$\|g\|_w \triangleq \sup_{x \in C} \frac{|g(x)|}{w(x)}.$$

It is easily verified that $\| \cdot \|_w$ is a norm on L_w^∞. Furthermore, L_w^∞ is complete with respect to the norm $\| \cdot \|_w$; see, for example, p.117 of *Real Analysis* by H. Royden (1968). Hence, L_w^∞ is a Banach space.

Remark 2.8.1 Because C is discrete, we may choose to encode the functions $g \in L_w^\infty$ as column vectors. Hence, we can alternatively view L_w^∞ as a linear space consisting of elements that are vectors (in the linear algebra sense). Furthermore, the norm $\| \cdot \|_w$ can be viewed as a norm on the space of vectors.

Another useful class of Banach spaces is offered by the next example. (We shall not have immediate need for these spaces, however.)

Example 2.8.2 For $1 \le p < \infty$, let L^p be the set of (real-valued) rv's Z such that $\mathrm{E}|Z|^p < \infty$. Note that L^p is a vector space. Furthermore, for $Z \in L^p$, set

$$\|Z\|_p = (\mathrm{E}|Z|^p)^{1/p}.$$

Then $\| \cdot \|_p$ is a norm on L^p under which L^p is complete; this is the Rieze-Fischer theorem (see p.117 of H. Royden's *Real Analysis*). Hence, L^p is a Banach space.

Suppose that $A = (A(x, y) : x, y \in C)$ is a matrix for which

$$\sup_x \frac{\sum_y |A(x, y)| w(y)}{w(x)} < \infty. \tag{2.8.1}$$

Note that if $g \in L_w^\infty$, then

$$\begin{aligned}
& |(Ag)(x)|/w(x) \\
\le \;& \sum_y |A(x, y)| \cdot |g(y)|/w(x) \\
= \;& \sum_y |A(x, y)| \frac{w(y)}{w(x)} \cdot \frac{|g(y)|}{w(y)} \\
\le \;& \sup_z \sum_y |A(z, y)| \frac{w(y)}{w(x)} \cdot \sup_y \frac{|g(y)|}{w(y)} \\
= \;& \sup_z \sum_y |A(z, y)| \frac{w(y)}{w(x)} \cdot \|g\|_w
\end{aligned}$$

and hence $Ag \in L_w^\infty$. So, if condition (2.8.1) is in force, then A maps L_w^∞ into L_w^∞.

Let $\mathscr{M}_w$ be the set of all matrices A such that (2.8.1) holds. Note that $\mathscr{M}_w$ is itself a vector space. Furthermore, the norm $\|\cdot\|_w$ on $\mathscr{L}_w$ induces an "operator norm" on $\mathscr{M}_w$ via

$$|||A|||_w = \sup_{\substack{g \in L_w^\infty \\ \|g\|_w \neq 0}} \frac{\|Ag\|_w}{\|g\|_w}.$$

Exercise 2.8.1 Prove that $|||\cdot|||_w$ is a norm on $\mathscr{M}_w$.

Remark 2.8.2 The norm $|||\cdot|||_w$ can be viewed as a matrix norm on the space of matrices $\mathscr{M}_w$.

The space $\mathscr{M}_w$ is complete under $|||\cdot|||_w$, and hence $\mathscr{M}_w$ is a Banach space; see p.189 of *Real Analysis* by H. Royden (1968).

Note that $\|Ag\|_w/\|g\|_w = \|A\widetilde{g}\|_w$, where $\widetilde{g} = g/\|g\|_w$ has unit norm. As a consequence, we can re-write $|||\cdot|||_w$ as

$$|||A|||_w = \sup_{\substack{g \in L_w^\infty \\ \|g\|_w = 1}} \|Ag\|_w.$$

For $x \in C$, $|(Ag)(x)|$ is maximized over functions g for which $\|g\|_w = 1$ by setting $g(y) = w(y) \cdot \text{sign}\, A(x,y)$, in which case we get

$$|(Ag)(x)| = \sum_y |A(x,y)| w(y),$$

proving that

$$|||A|||_w = \sup_x \sum_y |A(x,y)| \frac{w(y)}{w(x)}.$$

Exercise 2.8.2 i.) Prove that if $A_1, A_2 \in \mathscr{M}_w$, then

$$|||A_1 A_2|||_w \leq |||A_1\,|||_w \cdot |||\,A_2\,|||_w\,.$$

ii.) Prove that if $g \in L_w^\infty$, then

$$\|Ag\|_w \leq |||A\,|||_w \cdot \|g\|_w.$$

2.9 Using Matrix/Vector Norms to Bound the Solution to a Linear System

In all the instances of first transition analysis discussed earlier, the probabilistically meaningful solution u^* can be represented as

$$u^* = \sum_{n=0}^{\infty} G^n f, \tag{2.9.1}$$

for some matrix G and function f. We can now use the theory of matrix/vector norms to bound the solution u^*.

Suppose that $|||B|||_w < 1$ and $\|f\|_w < \infty$. Then, if we set

$$u_n^* = \sum_{j=0}^{n} G^j f,$$

it follows that $(u_n^* : n \geq 1)$ is a Cauchy sequence in L_w^∞. To see this, note that Exercise 1.8.2 implies that $|||G^j|||_w \leq |||G\,|||_w^j$. For $\epsilon > 0$, choose $N(\epsilon)$ so that

$$|||G\,|||_w^{N(\epsilon)}\,(1 - |||G|||_w)^{-1}\|f\|_w < \epsilon.$$

Then, for $n \geq m \geq N(\epsilon)$,

$$\|u_n^* - u_m^*\|_w \leq \sum_{j=m+1}^{n} |||G\,|||_w^j\,\|f\|_w \leq \sum_{j=N(\epsilon)}^{\infty} |||G\,|||_w^j\,\|f\|_w < \epsilon,$$

proving that $(u_n^* : n \geq 1)$ is Cauchy. It follows that there exists $u^* \in L_w^\infty$ such that $u_n^* \to u^*$ in L_w^∞. Hence, when $|||G|||_w < 1$ and $\|f\|_w < \infty$, the infinite sum (2.9.1) is well-defined as a limit.

Furthermore,

$$\|u^*\|_w \leq \sum_{n=0}^{\infty} |||G\,|||_w^j\,\|f\|_w = (1 - |||G|||_w)^{-1}\|f\|_w.$$

Recalling the definition of $\|\cdot\|_w$, we conclude that

$$|u^*(x)| \leq w(x) \cdot (1 - |||G|||_w)^{-1}\|f\|_w,$$

providing an upper bound on the function u^*. We have therefore proved the following result.

Theorem 2.9.1 Suppose that $G = (G(x, y) : x, y \in C)$ is a matrix for which there exists $c < 1$ such that

$$\sum_y |G(x, y)|w(y) \leq cw(x)$$

for $x \in C$. Assume also that $\sup\{|f(x)|/w(x) : x \in C\} < \infty$. Then,

$$u^* = \sum_{n=0}^{\infty} G^n f$$

is an element of L_w^∞ and

$$|u^*(x)| \leq w(x) \cdot (1 - c)^{-1}\|f\|_w$$

for $x \in C$.

Note that $(\sum_{j=0}^{n} G^j : n \geq 0)$ is a Cauchy sequence in $\mathscr{M}_w$, so that there exists $H \in \mathscr{M}_w$ such that

$$\sum_{j=0}^{n} G^j \to H$$

as $n \to \infty$ in $\mathscr{M}_w$, and

$$|||H|||_w \leq (1 - |||G|||_w)^{-1}.$$

Furthermore,

$$(I - G)(I + G + G^2 + \cdots + G^n) = I - G^{n+1} = (I + G + G^2 + \cdots + G^n)(I - G). \tag{2.9.2}$$

Taking limits in (2.9.2) (using the norm $|||\cdot|||_w$), we find that

$$(I - G)H = I = H(I - G) \tag{2.9.3}$$

In other words, we may conclude that $(I - G)$ has an inverse (namely H) when considered as an operator on $\mathscr{M}_w$. Thus, when $f \in L_w^\infty$ and $|||G|||_w < 1$, the linear system

$$u = f + Gu$$

has a unique solution $u^* \in L_w^\infty$ given by $u^* = Hf$.

Remark 2.9.1 When $|C| < \infty$, G is a finite dimensional matrix. The above discussion establishes that when $|||G|||_w < 1$, then $I - G$ is non-singular and $(I - G)^{-1} = H$.

Exercise 2.9.1 Suppose that $|||G|||_w < \infty$.

i.) If there exists $m < \infty$ such that $|||G^m|||_w < 1$, then $(\sum_{j=0}^n G^j : n \geq 0)$ is Cauchy in $\mathscr{M}_w$ and hence there exists $H \in \mathscr{M}_w$ such that

$$\sum_{j=0}^{n} G^j \to H$$

in $\mathscr{M}_w$ as $n \to \infty$. Furthermore, $(I - G)Hf = f = H(I - G)f$ for all $f \in L_w^\infty$.

ii.) If there exists $m < \infty$ such that $|||G^m|||_w < 1$ and $f \in L_w^\infty$, compute an upper bound on $u^* = \sum_{n=0}^\infty G^n f$.

iii.) If $|C| < \infty$, prove that there exists $m < \infty$ such that $|||G^m|||_w < 1$ if and only if $G^n \to 0$ as $n \to \infty$.

Exercise 2.9.2 Suppose that $u^*(x) = \mathrm{E}_x \sum_{j=0}^{T-1} f(X_j)$ for $x \in C$, where $T = \inf\{n \geq 0 : X_n \in C^c\}$ and $f : C \to [0, \infty)$.

i.) If there exists $c < 1$ and $w : C \to [1, \infty)$ such that

$$\mathrm{E}_x w(X_m) I(T > m) \leq c w(x)$$

for $x \in C$ and

$$\sup_{x \in C} \mathrm{E}_x w(X_1) I(T > 1)/w(x) < \infty,$$

compute an upper bound on $u^*(x)$.

ii.) Prove that if there exists $c < 1$ such that

$$\sup_{x \in C} P_x(T > m) \leq c,$$

then $|u^*(x)| \leq (1 - c)^{-1} \sup_{x \in C} |f(x)|$.

iii.) If $\sup_{x \in C} P_x(T > m) < 1$ for some $m \geq 1$, and $|C| < \infty$, then $u^* = (I - B)^{-1} f$.

Exercise 2.9.3 Suppose that $u^*(x) = \mathrm{E}_x \sum_{j=0}^\infty e^{-\alpha j} f(X_j)$ for $\alpha > 0$ and f bounded.

i.) Prove that if $H = \sum_{n=0}^\infty e^{-\alpha n} P^n$, then

$$(I - e^{-\alpha} P) H g = H(I - e^{-\alpha} P) g = g$$

for $g \in L_e^\infty$ (where $e(x) \equiv 1$ for $x \in C$).

ii.) If $|C| < \infty$, $I - e^{-\alpha}P$ is non-singular and $(I - e^{-\alpha}P)^{-1} = H$. Furthermore, $u^* = (I - e^{-\alpha}P)^{-1}f$.

In this section, we have studied the linear system

$$u = f + Gu, \tag{2.9.4}$$

using the concept of vector and matrix norms. Problem 1.9.2 establishes that if $|||G^m|||_w < 1$ for some $m \geq 1$ (which is equivalent to requiring that $|||G^n|||_w \to 0$ as $n \to \infty$), then for each $f \in L_w^\infty$, there exists a solution in L_w^∞ that is unique in the space L_w^∞. This, of course, is just asserting that the operator $I - G$ is invertible on L_w^∞. Note that in the infinite-dimensional setting (where $|C| = \infty$), questions of invertibility depend on the choice of the "function space" L_w^∞.

Consider, for example, the linear system $u = f + Bu$ that arises in the queueing chain example, in which

$$B = \begin{pmatrix} r & p & 0 & 0 & 0 & \cdots \\ q & r & p & 0 & 0 & \cdots \\ 0 & q & r & p & 0 & \cdots \\ 0 & 0 & q & r & p & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

Because of the fact that a solution can be recursively computed (by setting $u(1)$ arbitrarily and then computing $u(k)$ in terms of $u(k-1)$ and $u(k-2)$), it is evident that $u = f + Bu$ has a solution for every "right-hand side" f. However, if $I - B$ is invertible, there must exist a unique solution. In particular, there must exist a unique solution to the homogeneous linear system

$$u = Bu, \tag{2.9.5}$$

namely the trivial solution $u = 0$. But

$$\widetilde{u}(x) = \left(\frac{q}{p}\right)^x - 1$$

is a solution of (2.9.5) and thus (2.9.5) has a non-trivial solution.

But this non-trivial solution disappears if we use carefully the notions discussed earlier in this section. In particular, if $|||B^m|||_w < 1$ for some $m \geq 1$, then (2.9.4) has unique solutions in L_w^∞. It follows that for any weight function w satisfying $|||B^m|||_w < 1$, it must be that the non-trivial solution $\widetilde{u}$ satisfies $|||\widetilde{u}||| = \infty$, so that $\widetilde{u} \notin L_w^\infty$. On the other hand, our theory makes clear that if $|||B^n||| \to 0$ as $n \to \infty$, then $I - B$ is invertible on the function space L_w^∞. Hence, in the infinite-dimensional setting, the choice of function space (or, equivalently, the choice of w) plays a key role in the study of operator invertibility.

Exercise 2.9.4 This problem considers the effect of the choice of weight function on the linear system $u = f + Bu$, where B is the coefficient matrix that arises in consideration of the queueing chain with $q > p$.

i.) Prove that if $w = e$ (the constant function equal to one), then $|||B^m|||_e = 1$ for $m \geq 1$.

ii.) Prove that $\mathrm{E}_x \exp(\theta X_1) = \mathrm{E}\exp(\theta \widetilde{Z}_1)\exp(\theta x)$ for $x \geq 1$ for this chain, where $\mathrm{P}(\widetilde{Z}_1 = -1) = q$, $\mathrm{P}(\widetilde{Z}_1 = 0) = r$, and $\mathrm{P}(\widetilde{Z}_1 = 1) = p$.

iii.) Prove that if $\mathrm{E}\widetilde{Z}_1 < 0$, then $\mathrm{E}\exp(\theta\widetilde{Z}_1) < 1$ for $0 < \theta < \log(q/p)$. (Hint: Prove that $\mathrm{E}\exp(\theta\widetilde{Z}_1)$ is convex in θ.)

iv.) Prove that if $w(x) = \exp(\theta x)$ for $0 < \theta < \log(q/p)$, then $|||B|||_w < 1$.

v.) Suppose that $|f(x)| = O(\exp(rx))$ as $x \to \infty$, where $r < \log(q/p)$. Prove that there exists w for which $u = f + Bu$ has a unique solution in L^∞_w.

vi.) Suppose that $f(x) = x$. For each $\theta \in (0, \log(q/p))$, compute an explicit upper bound on $\mathrm{E}_x \sum_{j=0}^{T-1} f(X_j)$.

vii.) Optimize your bound on $\mathrm{E}_x \sum_{j=0}^{T-1} f(X_j)$ by judiciously choosing θ.