STOCHASTIC OPERATIONS RESEARCH

Ulrich Rieder
University of Ulm, Germany

Keywords: Markov chain, recurrence, limit distribution, Markov decision process, optimal policy, total and average reward criteria, backward induction algorithm, stochastic game, minimax theorem, Nash-equilibrium, queueing system, queueing discipline, Little's formula, multi-period inventory model, economic order quantity, Markowitz model, Black-Scholes model, two-fund separation, adaptive dynamic programming, estimation and control.

Contents

Summary

Stochastic Operations Research is concerned with phenomena that vary as time advances and where the variation has a significant chance component. This covers an enormous variety of applications in engineering systems, management science, economics, and computer science. Many applications focus on decision making. Stochastic models are then used to compare alternative decisions. In this paper, we discuss the most important stochastic models in Operations Research: Markov models, Markov decision processes, stochastic games, queueing systems, inventory models and investment models. Moreover, the section on adaptive dynamic programming includes statistical methods for analyzing sequential decision problems under uncertainty.

1. Introduction

One of the central problems in Operations Research is how to quantify the effects of uncertainty about the future. Uncertainties are important characteristics in a variety of applications in engineering systems, management science, computer science, economics and biological sciences. A few examples are queueing systems (the random occurrences being arrival of customers or completion of service), investment planning (changes of interest rate or market price), stochastic scheduling (completion of jobs or failure of machines), inventory systems (changes of demand, random lifetime of goods), insurance analysis (occurrence of claims, changes in premium rates) and optimal exploitation of resources such as fisheries, forestry or oil (amounts of resources found, random factors associated with harvesting or production). All these examples-and there are many more-are dynamic in that decisions are taken over time and decisions taken now have consequences for the future because they could affect the availability of resources or limit the options for subsequent decisions. *Stochastic Operations Research* is concerned with such dynamic and stochastic models. Methods of this paper have been applied successfully to these applications.

In this exposition, we present the most important stochastic models. Section 2 covers Markov chains and models and lays the foundation for the following sections. Section 3 describes the fundamentals of Markov decision processes. In section 4 we introduce concepts of game theory in a dynamic setting. The next three sections are concerned with important structured models: queueing systems, inventory and investment models. The last section on adaptive dynamic programming includes some statistical methods for analyzing sequential decision problems under uncertainty. Stochastic reliability models, statistical methods that arise in fitting stochastic models to real data, simulation experiments and also methods of stochastic programming are not treated in this paper.

2. Markov Models

Independence of random variables is a very restrictive assumption in stochastic modeling. The field of stochastic processes has focused on temporal relationships among random variables and on formulation of tractable forms of dependence, three of which have proved uncommonly fruitful in that they are simultaneously broad enough to be applicable and narrow enough to be interesting: *stationary processes*, *martingales* and *Markov processes*. Each has been studied extensively, and they carry a vast literature encompassing theory and application. Here we emphasize Markov processes and for simplicity, we shall restrict to Markov chains with finite or countable state spaces. The state space is always denoted by S.

We start with discrete-time Markov chains (X_n). The sequence (X_n) is called a *discrete-time Markov chain* if for all $n \in \mathbb{N}_0$ and for all states $i, j, i_k \in S$

$$\begin{aligned}&\mathbb{P}(X_{n+1} = j \mid X_0 = i_0, ..., X_{n-1} = i_{n-1}, X_n = i)\\ &= \mathbb{P}(X_{n+1} = j \mid X_n = i) = p_{ij}\end{aligned} \tag{1}$$

for a suitable transition probability p_{ij}. This condition (1), the Markov-property, was proposed by A.A. Markov (1856-1922) as part of work aimed at generalizing the classical limit theorems for independent random variables. It means that prediction of the future of the process, once the current state is known, cannot be improved by additional knowledge of the past. The matrix $P = (p_{ij})$ is called transition matrix. Two fundamental properties of P are that

$$p_{ij} \geq 0 \text{ and } \sum_{j\in S} p_{ij} = 1 \text{ for all } i, j \in S. \tag{2}$$

A matrix satisfying these properties is a *stochastic matrix*, and is a transition matrix of some Markov chain.

An important role in the analysis of Markov chains play the matrix powers of P. We denote $P^0 = I$, where I is the identity matrix, and the elements of the matrix P^n by $p_{ij}^{(n)}$. It is important to note that the matrix P^n is again a stochastic matrix and $p_{ij}^{(n)}$ exactly gives the probability of getting from state i to state j in n steps, i.e.

$$\mathbb{P}(X_{n+m} = j \mid X_m = i) = p_{ij}^{(n)} \text{ for all m} \geq 0. \tag{3}$$

The probabilities $p_{ij}^{(n)}$ are therefore called *n-step transition probabilities* and can be computed using the *Chapman-Kolmogorov equation*

$$p_{ij}^{(n+m)} = \sum_{k \in S} p_{ik}^{(n)} p_{kj}^{(m)} \text{ for all } i, j \in S. \tag{4}$$

We now introduce two examples, which will be used to illustrate concepts and results.

Example (*Random Walk*)

Let (Y_n) be a sequence of independent and identically distributed random variables with $\mathbb{P}(Y_n = 1) = p$, $\mathbb{P}(Y_n = -1) = 1 - p$ and $p \in (0, 1)$. We assume that $X_0 = 0$ and for $n = 1, 2, ...$,

$$X_n := \sum_{k=1}^{n} Y_k = X_{n-1} + Y_n. \tag{5}$$

Then (X_n) is obviously a Markov chain. The transition probabilities are given by $p_{ii+1} = p$ and $p_{ii-1} = 1 - p$.

Example (*Gambler's Ruin*)

A gambler makes repeated independent bets and wins \$1 with probability $p \in (0, 1)$ or loses \$1 with probability $q := 1 - p$. The gambler starts with an initial state i and will play repeatedly until all money is lost or until the fortune increases to \$$M$. Let X_n be the gambler's wealth after n plays. It is easily seen that (X_n) is Markov chain with state space $\{0, 1, ..., M\}$. The Markov property follows from the assumption that outcomes of successive bets are independent events. The transition probabilities are given by $p_{00} = p_{MM} = 1$ and for $i \neq 0$, $i \neq M$,

$$p_{ii+1} = p \quad \text{and} \quad p_{ii-1} = q. \tag{6}$$

State classification is a fundamental problem for Markov chains. The crucial issues are whether return to a state is certain or less than certain and, when return is certain, whether the mean time to return is finite. In the nicest case of finite mean time to return, the Markov chain exhibits stable long-run behavior, which we describe presently. Let $\tau_i :=$

$\inf\{n \in \mathbb{N} \mid X_n = i\}$ be the time of the first visit to state $i \in S$. A state i is called *recurrent* if $\mathbb{P}_i(\tau_i < \infty) = 1$ and *transient* if $\mathbb{P}_i(\tau_i < \infty) < 1$. Recurrent states are classified further. A recurrent state i is called *positive recurrent* if $m_i := \mathbb{E}_i[\tau_i] < \infty$ and *null-recurrent* if $m_i = \infty$. All these state properties are solidarity properties (see *Markov Models*). Moreover, a Markov chain is called *irreducible* if for each i and j there exit $n \in \mathbb{N}$ such that $p_{ij}^{(n)} > 0$. To classify the states of an irreducible Markov chain one employs the following criteria:

> (X_n) is positive recurrent if and only if there exists a probability distribution x satisfying $x_j = \sum_{i \in S} x_i p_{ij}$ for $j \in S$.
>
> Let i_0 be any fixed state. Then (X_n) is transient if and only if there exists a bounded solution x, not identical zero, such that $x_i = \sum_{j \neq i_0} p_{ij} x_j$ for $i \neq i_0$.

Example (*Random Walk*)

The random walk (X_n) is irreducible. All states are transient if $p \neq \frac{1}{2}$ and null recurrent if $p = \frac{1}{2}$.

Example (*Gambler's Ruin*)

The states $i = 0$ and $i = M$ are positive recurrent, all other states are transient. An interesting question is what is the probability that the gambler's fortune will reach M before all the resources are lost. Let us denote this probability by P_i. It is not too difficult to show that

$$P_i = \begin{cases} \dfrac{1-(q/p)^i}{1-(q/p)^M}, & \text{if } p \neq \dfrac{1}{2} \\ \dfrac{i}{M}, & \text{if } p = \dfrac{1}{2} \end{cases} \tag{7}$$

The probability $1 - P_i$ is then the ruin probability of the gambler. What is the expected number T_i of bets in this play? This expected number is surprisingly large and is given by

$$T_i = \begin{cases} \dfrac{1}{q-p}\left[i - M\,\dfrac{1-(q/p)^i}{1-(q/p)^M} \right], & \text{if } p \neq \dfrac{1}{2} \\ i(M-i), & \text{if } p = \dfrac{1}{2} \end{cases} \tag{8}$$

An important question concerns the long run behavior of a Markov chain. More precisely, we are interested in the probability π_j with which the Markov chain will be in state j after a large number of transitions. Before the main limit theorem can be given, one further concept is needed. A recurrent state i is periodic with period $d \geq 2$, if

$$\mathbb{P}_i\left(\tau_i \text{ an integer multiple of } d\right) = 1. \tag{9}$$

The aperiodic case is much nicer. Random walks are periodic with $d = 2$. The solidarity property extends: in an irreducible Markov chain, either all states are aperiodic or all are periodic with the same period. We come now to the main limit result:

> Let (X_n) be an irreducible, positive recurrent and aperiodic Markov chain. Then there exists the limiting probability $\pi_j = \lim_{n\to\infty} p_{ij}^{(n)}$ for all $i,\ j \in S$, which satisfies $\pi_j = \frac{1}{m_j}$. Moreover, $\pi = (\pi_j)$ is the unique stationary distribution.

A *stationary distribution* $\pi = (\pi_i)$ is a probability distribution on S satisfying the linear equation

$$\pi_i = \sum_{j\in S} \pi_j p_{ji} \text{ for } i \in S. \tag{10}$$

An irreducible, positive recurrent and aperiodic Markov chain has limiting probabilities which are independent of the initial state (a phenomenon termed *ergodic*) and which are the inverse of the mean time to return.

In the second part of this section, we consider continuous-time Markov chains (X_t). The stochastic process (X_t) is called a *continuous-time Markov chain* if for all states $i, j, i_k \in S$ and time points $0 \le t_0 < t_1 < \ldots < t_{n+1}$

$$\begin{aligned}&\mathbb{P}(X_{t_{n+1}} = j \mid X_{t_0} = i_0, \ldots, X_{t_{n-1}} = i_{n-1}, X_{t_n} = i) = \\ &= \mathbb{P}(X_{t_{n+1}} = j \mid X_{t_n} = i) = p_{ij}(t_{n+1} - t_n)\end{aligned} \tag{11}$$

for a suitable transition function $p_{ij}(t)$. Each matrix $P(t) = (p_{ij}(t))$ is a stochastic matrix and $P(0) = I$. Moreover, the Chapman-Kolmogoroff equation (or semigroup property)

$$P(s+t) = P(s)P(t) \quad \text{for all } s, t \ge 0 \tag{12}$$

is fulfilled. In what follows we will always assume that $P(t)$ is right-continuous in $t = 0$. Then the function $t \mapsto p_{ij}(t)$ is continuously differentiable and

$$q_{ij} := p'_{ij}(0) = \lim_{t\downarrow 0} \frac{p_{ij}(t) - \delta ij}{t} \tag{13}$$

exists. The derivatives satisfy

$$q_{ii} =: -q_i \le 0,\ q_{ij} \in [0,\infty) \text{ for } i \ne j \text{ and } \sum_{j\ne i} q_{ij} \le q_i. \tag{14}$$

The matrix $Q = (q_{ij})$ is called *intensity matrix* or *generator* and the interpretation is that for $i \neq j$ and $h > 0$

$$\mathbb{P}(X_{t+h} = j \mid X_t = i) = q_{ij} h + o(h) \tag{15}$$

$$\mathbb{P}(X_{t+h} = j \mid X_t = i) = 1 - q_i h + o(h). \tag{16}$$

The matrix Q is called *conservative* if $\sum_{j \neq i} q_{ij} = q_i < \infty$ for all $i \in S$. For finite state spaces Q is always conservative.

Example (*Poisson Process*)

Let (X_t) be an arrival or counting process, i.e. X_t is the random number of arrivals in $[0, t]$. (X_t) is a *Poisson process* if $X_0 = 0$ and (X_t) has the transition function

$$p_{ij}(t) = e^{-\lambda t} \frac{(\lambda t)^{j-i}}{(j-i)!} \text{ for } j \geq i. \tag{17}$$

for some constant $\lambda > 0$, known as the rate of (X_t). The state space is $\mathbb{N}_0$ and the intensities are

$$q_{ii} = -\lambda, \quad q_{ii+1} = \lambda \quad \text{for} \quad i \in S.$$

Example (*Birth and Death Process*)

The Markov chain (X_t) with state space $\mathbb{N}_0$ is a *birth and death process* if for each $i \in S$

$$q_{ii-1} = \mu_i, \quad q_{ii} = -(\mu_i + \lambda_i), \quad q_{ii+1} = \lambda_i. \tag{18}$$

λ_i are the birth rates and μ_i the death rates. One interprets X_t as the size at time t of a randomly varying population, which changes only by single births and deaths.

Under some mild conditions the transition functions can be reconstructed from the intensity matrix Q by using the *forward differential equation* $P'(t) = P(t)Q$ for $t \geq 0$ and the *backward equation* $P'(t) = QP(t)$ for $t \geq 0$. Limit theory parallels that for discrete-time Markov chains. To make this precise, we note first that (X_t) moves only by jumping from one state to another. Let $0 = T_0 < T_1 < T_2 \ldots$ be the time of these transitions, with $T_n = \infty$ if there are fewer that n jumps. If $T_n < \infty$, let $Y_n = X_{T_n}$ be the state entered at T_n, and otherwise let $Y_n = Y_{n-1}$. Then (Y_n) is a Markov chain, the so-called *embedded Markov chain*, with transition probabilities

$$p_{ij} = \begin{cases} \dfrac{q_{ij}}{q_i}, & \text{if } i \neq j \\ 0, & \text{if } i = j. \end{cases} \tag{19}$$

State classification for (X_t) is effected in part via (Y_n). State $i \in S$ is *recurrent* (*transient*) if i is recurrent (transient) for the embedded Markov chain (Y_n). (X_t) is called *irreducible* if (Y_n) is irreducible. Then we come to the main limit result:

> Let (X_t) be irreducible and recurrent. Then there exist the limits $\pi_j = \lim_{t\to\infty} p_{ij}(t)$ for all $i, j \in S$, which are independent of the initial state. Either $\pi \equiv 0$ or $\sum_{j\in S} \pi_j = 1$.

The dichotomy above can be used to complete the classification of recurrent states in the irreducible case. All states are *positive recurrent* if $\sum_{j\in S} \pi_j = 1$, and *null-recurrent* if $\pi \equiv 0$. Moreover, the limit distribution π is a *stationary distribution* if

$$\pi_i = \sum_{j\in S} \pi_j p_{ji}(t) \text{ for all } i \in S \text{ and } t \geq 0. \tag{20}$$

An effective technique for computing a stationary distribution is given in the next result:

> Let (X_t) be irreducible and positive recurrent with intensity matrix Q. Then the stationary distribution π is the unique probability distribution on S satisfying the linear equation

$$\sum_{i\in S} \pi_i q_{ij} = 0 \text{ for } j \in S \tag{21}$$

Example (*Birth and Death Process*)

The birth and death process is obviously irreducible and positive recurrent if and only if

$$\sum_{i=1}^{\infty} \frac{\lambda_0 \cdots \lambda_i}{\mu_1 \cdots \mu_{i+1}} < \infty. \tag{22}$$

If the birth and death rates are independent of i, i.e. $\lambda_i = \lambda$ and $\mu_i = \mu$, this continuous-time Markov chain is positive recurrent if and only if $\rho = \frac{\lambda}{\mu}$ is less than one. Then the limiting distribution has the form

$$\pi_i = (1-\rho)\rho^i \quad \text{for} \quad i \in S \tag{23}$$

and $\pi = (\pi_i)$ is the unique stationary distribution. Applications of this Markov chain can be found in *Queueing Systems*.

-
-
-

Bibliography

Arrow K.I., Karlin S., and Scarf H. (1958). *Studies in the Mathematical Theory of Inventory and Production*, 340 pp. Stanford, CA: Stanford University Press. [Classical textbook on stochastic inventory control.]

Asmussen S. (1987). *Applied Probability and Queues*, 318 pp. Chichester: John Wiley & Sons. [Presents a comprehensive treatment of queueing systems.]

Filar J.A. and Vrieze K. (1997). *Competitive Markov Decision Processes*, 393 pp. New York: Springer-Verlag. [This textbook presents the main topics for stochastic games.]

Gross D. and Harris C.M. (1985) *Fundamentals of Queueing Theory*, 587 pp. New York: Wiley. [Classical textbook on queueing systems.]

Hernández-Lerma O. (1989). *Adaptive Markov Control Processes*, 148 pp. New York: Springer-Verlag. [Graduate textbook on Markov decision processes under uncertainty.]

Heyman D.P. and Sobel M. (1984). *Stochastic Models in Operations Research*. New York: McGraw-Hill. [Textbook on stochastic models.]

Korn R. (1997). *Optimal Portfolios: Stochastic Models For Optimal Investment and Risk Management In Continuous Time*. Singapore: World Scientific. [Recent book on continuous-time portfolio theory.]

Luenberger D.G. (1998). *Investment Science*. New York: Oxford University Press. [Excellent textbook on investment models.]

Puterman M. (1994). *Markov Decision Processes*. New York: Wiley. [Textbook on Markov decision processes with countable state spaces.]

Resnick S.I. (1992). *Adventures in Stochastic Processes*. Boston: Birkhäuser. [Excellent text book with many examples.]

Ross S.M. (1997). *An Introduction to Probability Models*. San Diego: Academic Press,. [Elementary textbook on Markov models, including chapters on reliability and simulation.]

Sennott L.I. (1998). *Stochastic Dynamic Programming and the Control of Queues*. New York: John Wiley & Sons. [This book considers Markov decision processes with queueing applications and has a strong computational emphasis.]

Zipkin P.H. (2000). *Foundations of Inventory Management*. Boston: McGraw-Hill. [Comprehensive summary documenting the state of the art in book form.]

Biographical Sketch

Ulrich Rieder, born in 1945, received the Ph.D. degree in mathematics in 1972 (University of Hamburg) and the Habilitation in 1979 (University of Karlsruhe). Since 1980, he has been Full Professor of Mathematics and head of the Department of Optimization and Operations Research at the University of Ulm. His research interests include the analysis and control of stochastic processes with applications in telecommunication, logistics, applied probability, finance and insurance. Dr. Rieder is Editor-in-Chief of the journal Mathematical Methods of Operations Research.