

# Operations Research Markov Decision Theory

Right here, we have countless ebook **operations research markov decision theory** and collections to check out. We additionally pay for variant types and also type of the books to browse. The up to standard book, fiction, history, novel, scientific research, as with ease as various supplementary sorts of books are readily approachable here.

As this operations research markov decision theory, it ends taking place swine one of the favored book operations research markov decision theory collections that we have. This is why you remain in the best website to see the unbelievable ebook to have.

*Continuous-Time Markov Decision Processes* Xianping Guo 2010-04-29 Continuous-time Markov decision processes (MDPs), also known as controlled Markov chains, are used for modeling decision-making problems that arise in operations research (for instance, inventory, manufacturing, and queueing systems), computer science, communications engineering, control of populations (such as fisheries and epidemics), and management science, among many other fields. This volume provides a unified, systematic, self-contained presentation of recent developments on the theory and applications of continuous-time MDPs. The MDPs in this volume include most of the cases that arise in applications, because they allow unbounded transition and reward/cost rates. Much of the material appears for the first time in book form.

**Decision Theory Models for Applications in Artificial Intelligence: Concepts and Solutions** Sucar, L. Enrique 2011-10-31 One of the goals of artificial intelligence (AI) is creating autonomous agents that must make decisions based on uncertain and incomplete information. The goal is to design rational agents that must take the best action given the information available and their goals. *Decision Theory Models for Applications in Artificial Intelligence: Concepts and Solutions* provides an introduction to different types of decision theory techniques, including MDPs, POMDPs, Influence Diagrams, and Reinforcement Learning, and illustrates their application in artificial intelligence. This book provides insights into the advantages and challenges of using decision theory models for developing intelligent systems.

Queues Moshe Haviv 2013-05-20 Queueing theory (the mathematical theory of waiting lines in all its configurations) continues to be a standard major area of operations research on the stochastic side. Therefore, universities with an active program in operations research sometimes will have an entire course devoted mainly or entirely to queueing theory, and the course is also taught in computer science, electrical engineering, mathematics, and industrial engineering programs. The basic course in queueing theory is often taught at first year graduate level, though can be taught at senior level undergraduate as well. This text evolved from the author's preferred syllabus for teaching the course, presenting the material in a more logical order than other texts and so being more effective in teaching the basics of queueing theory. The first three chapters focus on the needed preliminaries, including exposition distributions, Poisson processes and generating functions, renewal theory, and Markov chains, Then, rather than switching to first-come first-served memoryless queues

here as most texts do, Haviv discusses the M/G/1 model instead of the M/M/1, and then covers priority queues. Later chapters cover the G/M/1 model, thirteen examples of continuous-time Markov processes, open networks of memoryless queues and closed networks, queueing regimes with insensitive parameters, and then concludes with two-dimensional queueing models which are quasi birth and death processes. Each chapter ends with exercises.

## **U.S. Government Research & Development Reports 1970**

**Markov Decision Processes with Applications to Finance** Nicole Bäuerle 2011-06-06 The theory of Markov decision processes focuses on controlled Markov chains in discrete time. The authors establish the theory for general state and action spaces and at the same time show its application by means of numerous examples, mostly taken from the fields of finance and operations research. By using a structural approach many technicalities (concerning measure theory) are avoided. They cover problems with finite and infinite horizons, as well as partially observable Markov decision processes, piecewise deterministic Markov decision processes and stopping problems. The book presents Markov decision processes in action and includes various state-of-the-art applications with a particular view towards finance. It is useful for upper-level undergraduates, Master's students and researchers in both applied probability and finance, and provides exercises (without solutions).

Decision Theory and Multi-Agent Planning Giacomo Della Riccia 2007-05-03 The work presents a modern, unified view on decision support and planning by considering its basics like preferences, belief, possibility and probability as well as utilities. These features together are immanent for software agents to believe the user that the agents are "intelligent".

Simulation-based Algorithms for Markov Decision Processes Hyeong Soo Chang 2007-05-01 Markov decision process (MDP) models are widely used for modeling sequential decision-making problems that arise in engineering, economics, computer science, and the social sciences. This book brings the state-of-the-art research together for the first time. It provides practical modeling methods for many real-world problems with high dimensionality or complexity which have not hitherto been treatable with Markov decision processes.

Markov Chains Wai-Ki Ching 2013-03-27 This new edition of Markov Chains: Models, Algorithms and Applications has been completely reformatted as a text, complete with end-of-chapter exercises, a new focus on management science, new applications of the models, and new examples with applications in financial risk management and modeling of financial data. This book consists of eight chapters. Chapter 1 gives a brief introduction to the classical theory on both discrete and continuous time Markov chains. The relationship between Markov chains of finite states and matrix theory will also be highlighted. Some classical iterative methods for solving linear systems will be introduced for finding the stationary distribution of a Markov chain. The chapter then covers the basic theories and algorithms for hidden Markov models (HMMs) and Markov decision processes (MDPs). Chapter 2 discusses the applications of continuous time Markov chains to model queueing systems and discrete time Markov chain for computing the PageRank, the ranking of websites on the Internet. Chapter 3 studies Markovian models for manufacturing and re-manufacturing systems and presents closed form solutions and fast numerical algorithms for solving the captured systems. In Chapter 4, the authors present a simple hidden Markov model (HMM) with fast

numerical algorithms for estimating the model parameters. An application of the HMM for customer classification is also presented. Chapter 5 discusses Markov decision processes for customer lifetime values. Customer Lifetime Values (CLV) is an important concept and quantity in marketing management. The authors present an approach based on Markov decision processes for the calculation of CLV using real data. Chapter 6 considers higher-order Markov chain models, particularly a class of parsimonious higher-order Markov chain models. Efficient estimation methods for model parameters based on linear programming are presented. Contemporary research results on applications to demand predictions, inventory control and financial risk measurement are also presented. In Chapter 7, a class of parsimonious multivariate Markov models is introduced. Again, efficient estimation methods based on linear programming are presented. Applications to demand predictions, inventory control policy and modeling credit ratings data are discussed. Finally, Chapter 8 re-visits hidden Markov models, and the authors present a new class of hidden Markov models with efficient algorithms for estimating the model parameters. Applications to modeling interest rates, credit ratings and default data are discussed. This book is aimed at senior undergraduate students, postgraduate students, professionals, practitioners, and researchers in applied mathematics, computational science, operational research, management science and finance, who are interested in the formulation and computation of queueing networks, Markov chain models and related topics. Readers are expected to have some basic knowledge of probability theory, Markov processes and matrix theory.

**Markov Decision Processes** Martin L. Puterman 2014-08-28 The Wiley-Interscience Paperback Series consists of selected books that have been made more accessible to consumers in an effort to increase global appeal and general circulation. With these new unabridged softcover volumes, Wiley hopes to extend the lives of these works by making them available to future generations of statisticians, mathematicians, and scientists. "This text is unique in bringing together so many results hitherto found only in part in other texts and papers. . . . The text is fairly self-contained, inclusive of some basic mathematical results needed, and provides a rich diet of examples, applications, and exercises. The bibliographical material at the end of each chapter is excellent, not only from a historical perspective, but because it is valuable for researchers in acquiring a good perspective of the MDP research potential." —Zentralblatt für Mathematik ". . . it is of great value to advanced-level students, researchers, and professional practitioners of this field to have now a complete volume (with more than 600 pages) devoted to this topic. . . . Markov Decision Processes: Discrete Stochastic Dynamic Programming represents an up-to-date, unified, and rigorous treatment of theoretical and computational aspects of discrete-time Markov decision processes." —Journal of the American Statistical Association

*Planning Based on Decision Theory* Giacomo Della Riccia 2014-05-04 Planning of actions based on decision theory is a hot topic for many disciplines. Seemingly unlimited computing power, networking, integration and collaboration have meanwhile attracted the attention of fields like Machine Learning, Operations Research, Management Science and Computer Science. Software agents of e-commerce, mediators of Information Retrieval Systems and Database based Information Systems are typical new application areas. Until now, planning methods were successfully applied in production, logistics, marketing, finance, management, and used in robots, software agents etc. It is the special feature of the book that planning is embedded into decision theory, and this will give the interested reader new perspectives to follow-up.

**Examples In Markov Decision Processes** Alexey B Piunovskiy 2012-09-21 This invaluable book provides approximately eighty examples illustrating the theory of controlled discrete-time Markov processes. Except for applications of the theory to real-life problems like stock exchange, queues, gambling, optimal search etc, the main attention is paid to counter-intuitive, unexpected properties of optimization problems. Such examples illustrate the importance of conditions imposed in the theorems on Markov Decision Processes. Many of the examples are based upon examples published earlier in journal articles or textbooks while several other examples are new. The aim was to collect them together in one reference book which should be considered as a complement to existing monographs on Markov decision processes. The book is self-contained and unified in presentation. The main theoretical statements and constructions are provided, and particular examples can be read independently of others. Examples in Markov Decision Processes is an essential source of reference for mathematicians and all those who apply the optimal control theory to practical purposes. When studying or using mathematical methods, the researcher must understand what can happen if some of the conditions imposed in rigorous theorems are not satisfied. Many examples confirming the importance of such conditions were published in different journal articles which are often difficult to find. This book brings together examples based upon such sources, along with several new ones. In addition, it indicates the areas where Markov decision processes can be used. Active researchers can refer to this book on applicability of mathematical methods and theorems. It is also suitable reading for graduate and research students where they will better understand the theory.

**Handbook of Markov Decision Processes** Eugene A. Feinberg 2012-12-06 Eugene A. Feinberg Adam Shwartz This volume deals with the theory of Markov Decision Processes (MDPs) and their applications. Each chapter was written by a leading expert in the respective area. The papers cover major research areas and methodologies, and discuss open questions and future research directions. The papers can be read independently, with the basic notation and concepts of Section 1.2. Most chapters should be accessible by graduate or advanced undergraduate students in fields of operations research, electrical engineering, and computer science.

1.1 AN OVERVIEW OF MARKOV DECISION PROCESSES The theory of Markov Decision Processes—also known under several other names including sequential stochastic optimization, discrete-time stochastic control, and stochastic dynamic programming—studies sequential optimization of discrete time stochastic systems. The basic object is a discrete-time stochastic system whose transition mechanism can be controlled over time. Each control policy defines the stochastic process and values of objective functions associated with this process. The goal is to select a "good" control policy. In real life, decisions that humans and computers make on all levels usually have two types of impacts: (i) they cost or save time, money, or other resources, or they bring revenues, as well as (ii) they have an impact on the future, by influencing the dynamics. In many situations, decisions with the largest immediate profit may not be good in view of future events. MDPs model this paradigm and provide results on the structure and existence of good policies and on methods for their calculation.

**Handbook of Healthcare Analytics** Tinglong Dai 2018-07-30 How can analytics scholars and healthcare professionals access the most exciting and important healthcare topics and tools for the 21st century? Editors Tinglong Dai and Sridhar Tayur, aided by a team of internationally acclaimed experts, have curated this timely volume to help newcomers and seasoned researchers alike to rapidly comprehend a diverse set of thrusts and tools in this

rapidly growing cross-disciplinary field. The Handbook covers a wide range of macro-, meso- and micro-level thrusts—such as market design, competing interests, global health, personalized medicine, residential care and concierge medicine, among others—and structures what has been a highly fragmented research area into a coherent scientific discipline. The handbook also provides an easy-to-comprehend introduction to five essential research tools—Markov decision process, game theory and information economics, queueing games, econometric methods, and data science—by illustrating their uses and applicability on examples from diverse healthcare settings, thus connecting tools with thrusts. The primary audience of the Handbook includes analytics scholars interested in healthcare and healthcare practitioners interested in analytics. This Handbook: Instills analytics scholars with a way of thinking that incorporates behavioral, incentive, and policy considerations in various healthcare settings. This change in perspective—a shift in gaze away from narrow, local and one-off operational improvement efforts that do not replicate, scale or remain sustainable—can lead to new knowledge and innovative solutions that healthcare has been seeking so desperately. Facilitates collaboration between healthcare experts and analytics scholar to frame and tackle their pressing concerns through appropriate modern mathematical tools designed for this very purpose. The handbook is designed to be accessible to the independent reader, and it may be used in a variety of settings, from a short lecture series on specific topics to a semester-long course.

### **Anticipatory Optimization for Dynamic Decision Making** Stephan Meisel 2011-06-23

The availability of today's online information systems rapidly increases the relevance of dynamic decision making within a large number of operational contexts. Whenever a sequence of interdependent decisions occurs, making a single decision raises the need for anticipation of its future impact on the entire decision process. Anticipatory support is needed for a broad variety of dynamic and stochastic decision problems from different operational contexts such as finance, energy management, manufacturing and transportation. Example problems include asset allocation, feed-in of electricity produced by wind power as well as scheduling and routing. All these problems entail a sequence of decisions contributing to an overall goal and taking place in the course of a certain period of time. Each of the decisions is derived by solution of an optimization problem. As a consequence a stochastic and dynamic decision problem resolves into a series of optimization problems to be formulated and solved by anticipation of the remaining decision process. However, actually solving a dynamic decision problem by means of approximate dynamic programming still is a major scientific challenge. Most of the work done so far is devoted to problems allowing for formulation of the underlying optimization problems as linear programs. Problem domains like scheduling and routing, where linear programming typically does not produce a significant benefit for problem solving, have not been considered so far. Therefore, the industry demand for dynamic scheduling and routing is still predominantly satisfied by purely heuristic approaches to anticipatory decision making. Although this may work well for certain dynamic decision problems, these approaches lack transferability of findings to other, related problems. This book has serves two major purposes: - It provides a comprehensive and unique view of anticipatory optimization for dynamic decision making. It fully integrates Markov decision processes, dynamic programming, data mining and optimization and introduces a new perspective on approximate dynamic programming. Moreover, the book identifies different degrees of anticipation, enabling an assessment of specific approaches to dynamic decision making. - It shows for the first time how to successfully solve a dynamic vehicle routing problem by approximate dynamic programming. It elaborates on every

building block required for this kind of approach to dynamic vehicle routing. Thereby the book has a pioneering character and is intended to provide a footing for the dynamic vehicle routing community.

Partially Observed Markov Decision Processes Vikram Krishnamurthy 2016-03-21 This book covers formulation, algorithms, and structural results of partially observed Markov decision processes, whilst linking theory to real-world applications in controlled sensing. Computations are kept to a minimum, enabling students and researchers in engineering, operations research, and economics to understand the methods and determine the structure of their optimal solution.

Operations Research Kuodi Jian 2016-12-28 This book is dedicated to operations research of broad applications, such as improving informational bases of performance measurement with grey relational analysis, application of lean methodologies in a neurosurgery high dependency unit, iteration algorithms in Markov decision processes with state-action-dependent discount factors and unbounded costs, financial feasibility analysis of Natura Rab business case study, and mathematical modeling of isothermal drying and its potential application in the design of the industrial drying regimes of clay products. Operations research is an important topic. In addition to its obvious benefits of winning a war, making most profit in a business endeavor, and constructing a correct mathematical model, it also provides a tool for efficient use of natural resources. Furthermore, both theory and practice of operations research and its related concepts are covered in the book, and a reader can benefit from this balanced coverage.

Markov Chains: Models, Algorithms and Applications Wai-Ki Ching 2006-06-05 Markov chains are a particularly powerful and widely used tool for analyzing a variety of stochastic (probabilistic) systems over time. This monograph will present a series of Markov models, starting from the basic models and then building up to higher-order models. Included in the higher-order discussions are multivariate models, higher-order multivariate models, and higher-order hidden models. In each case, the focus is on the important kinds of applications that can be made with the class of models being considered in the current chapter. Special attention is given to numerical algorithms that can efficiently solve the models. Therefore, Markov Chains: Models, Algorithms and Applications outlines recent developments of Markov chain models for modeling queueing sequences, Internet, re-manufacturing systems, reverse logistics, inventory systems, bio-informatics, DNA sequences, genetic networks, data mining, and many other practical systems.

Operations Research: Introduction To Models And Methods Richard Johannes Boucherie 2021-10-26 This attractive textbook with its easy-to-follow presentation provides a down-to-earth introduction to operations research for students in a wide range of fields such as engineering, business analytics, mathematics and statistics, computer science, and econometrics. It is the result of many years of teaching and collective feedback from students. The book covers the basic models in both deterministic and stochastic operations research and is a springboard to more specialized texts, either practical or theoretical. The emphasis is on useful models and interpreting the solutions in the context of concrete applications. The text is divided into several parts. The first three chapters deal exclusively with deterministic models, including linear programming with sensitivity analysis, integer programming and heuristics, and network analysis. The next three chapters primarily cover

basic stochastic models and techniques, including decision trees, dynamic programming, optimal stopping, production planning, and inventory control. The final five chapters contain more advanced material, such as discrete-time and continuous-time Markov chains, Markov decision processes, queueing models, and discrete-event simulation. Each chapter contains numerous exercises, and a large selection of exercises includes solutions.

**Operations Research** Richard J. Boucherie 2021-10-26 Linear programming -- Integer programming -- Network analysis -- Decision trees -- Dynamic programming -- Inventory management -- Discrete-time Markov chains -- Continuous-time Markov chains -- Queueing theory -- Markov decision processes.

Markov Decision Processes with Their Applications Qiyang Hu 2007-09-14 Put together by two top researchers in the Far East, this text examines Markov Decision Processes - also called stochastic dynamic programming - and their applications in the optimal control of discrete event systems, optimal replacement, and optimal allocations in sequential online auctions. This dynamic new book offers fresh applications of MDPs in areas such as the control of discrete event systems and the optimal allocations in sequential online auctions.

The Operations Research Center at MIT J. Cole Smith 2005

**Introduction to Operations Research: Probabilistic models: 15. Markov chains. 16. Queueing theory. 17. The application of queueing theory. 18. Inventory theory. 19. Forecasting. 20. Markovian decision processes and applications. 21. Reliability. 22. Decision analysis. 23. Simulation** Frederick S. Hillier 1990

Reinforcement Learning, second edition Richard S. Sutton 2018-11-13 The significantly expanded and updated new edition of a widely used text on reinforcement learning, one of the most active research areas in artificial intelligence. Reinforcement learning, one of the most active research areas in artificial intelligence, is a computational approach to learning whereby an agent tries to maximize the total amount of reward it receives while interacting with a complex, uncertain environment. In Reinforcement Learning, Richard Sutton and Andrew Barto provide a clear and simple account of the field's key ideas and algorithms. This second edition has been significantly expanded and updated, presenting new topics and updating coverage of other topics. Like the first edition, this second edition focuses on core online learning algorithms, with the more mathematical material set off in shaded boxes. Part I covers as much of reinforcement learning as possible without going beyond the tabular case for which exact solutions can be found. Many algorithms presented in this part are new to the second edition, including UCB, Expected Sarsa, and Double Learning. Part II extends these ideas to function approximation, with new sections on such topics as artificial neural networks and the Fourier basis, and offers expanded treatment of off-policy learning and policy-gradient methods. Part III has new chapters on reinforcement learning's relationships to psychology and neuroscience, as well as an updated case-studies chapter including AlphaGo and AlphaGo Zero, Atari game playing, and IBM Watson's wagering strategy. The final chapter discusses the future societal impacts of reinforcement learning.

**Algorithms for Reinforcement Learning** Csaba Szepesvari 2010 Reinforcement learning is a learning paradigm concerned with learning to control a system so as to maximize a

numerical performance measure that expresses a long-term objective. What distinguishes reinforcement learning from supervised learning is that only partial feedback is given to the learner about the learner's predictions. Further, the predictions may have long term effects through influencing the future state of the controlled system. Thus, time plays a special role. The goal in reinforcement learning is to develop efficient learning algorithms, as well as to understand the algorithms' merits and limitations. Reinforcement learning is of great interest because of the large number of practical applications that it can be used to address, ranging from problems in artificial intelligence to operations research or control engineering. In this book, we focus on those algorithms of reinforcement learning that build on the powerful theory of dynamic programming. We give a fairly comprehensive catalog of learning problems, describe the core ideas, note a large number of state of the art algorithms, followed by the discussion of their theoretical properties and limitations.

**Dynamic Programming and Inventory Control** Alain Bensoussan 2011 Presents a unified theory of dynamic programming and Markov decision processes and its application to a major field of operations research and operations management: inventory control. For continuous time, this book concentrates only on models of interest to inventory control. For discrete time, the focus is mainly on infinite horizon models.

*Algorithmic Decision Theory* RONEN BRAFMAN 2011-10-07 This book constitutes the refereed proceedings of the Second International Conference on Algorithmic Decision Theory, ADT 2011, held in Piscataway, NJ, USA, in October 2011. The 24 revised full papers presented were carefully reviewed and selected from 50 submissions.

*Operations Research and Health Care* Margaret L. Brandeau 2006-04-05 In both rich and poor nations, public resources for health care are inadequate to meet demand. Policy makers and health care providers must determine how to provide the most effective health care to citizens using the limited resources that are available. This chapter describes current and future challenges in the delivery of health care, and outlines the role that operations research (OR) models can play in helping to solve those problems. The chapter concludes with an overview of this book - its intended audience, the areas covered, and a description of the subsequent chapters. KEY WORDS Health care delivery, Health care planning HEALTH CARE DELIVERY: PROBLEMS AND CHALLENGES 3 1.1 WORLDWIDE HEALTH: THE PAST 50 YEARS Human health has improved significantly in the last 50 years. In 1950, global life expectancy was 46 years [1]. That figure rose to 61 years by 1980 and to 67 years by 1998 [2]. Much of these gains occurred in low- and middle-income countries, and were due in large part to improved nutrition and sanitation, medical innovations, and improvements in public health infrastructure.

Competitive Markov Decision Processes Jerzy Filar 2012-12-06 This book is intended as a text covering the central concepts and techniques of Competitive Markov Decision Processes. It is an attempt to present a rigorous treatment that combines two significant research topics: Stochastic Games and Markov Decision Processes, which have been studied extensively, and at times quite independently, by mathematicians, operations researchers, engineers, and economists. Since Markov decision processes can be viewed as a special noncompetitive case of stochastic games, we introduce the new terminology Competitive Markov Decision Processes that emphasizes the importance of the link between these two topics and of the properties of the underlying Markov processes. The book is designed to be used either in a

classroom or for self-study by a mathematically mature reader. In the Introduction (Chapter 1) we outline a number of advanced undergraduate and graduate courses for which this book could usefully serve as a text. A characteristic feature of competitive Markov decision processes - and one that inspired our long-standing interest - is that they can serve as an "orchestra" containing the "instruments" of much of modern applied (and at times even pure) mathematics. They constitute a topic where the instruments of linear algebra, applied probability, mathematical programming, analysis, and even algebraic geometry can be "played" sometimes solo and sometimes in harmony to produce either beautifully simple or equally beautiful, but baroque, melodies, that is, theorems.

*Continuous-Time Markov Decision Processes* Alexey Piunovskiy 2020-12-11 This book offers a systematic and rigorous treatment of continuous-time Markov decision processes, covering both theory and possible applications to queueing systems, epidemiology, finance, and other fields. Unlike most books on the subject, much attention is paid to problems with functional constraints and the realizability of strategies. Three major methods of investigations are presented, based on dynamic programming, linear programming, and reduction to discrete-time problems. Although the main focus is on models with total (discounted or undiscounted) cost criteria, models with average cost criteria and with impulsive controls are also discussed in depth. The book is self-contained. A separate chapter is devoted to Markov pure jump processes and the appendices collect the requisite background on real analysis and applied probability. All the statements in the main text are proved in detail. Researchers and graduate students in applied probability, operational research, statistics and engineering will find this monograph interesting, useful and valuable.

## **Markov Decision Processes 1991**

OPTIMIZATION AND OPERATIONS RESEARCH - Volume IV Ulrich Derigs 2009-04-15 Optimization and Operations Research is a component of Encyclopedia of Mathematical Sciences in the global Encyclopedia of Life Support Systems (EOLSS), which is an integrated compendium of twenty one Encyclopedias. The Theme on Optimization and Operations Research is organized into six different topics which represent the main scientific areas of the theme: 1. Fundamentals of Operations Research; 2. Advanced Deterministic Operations Research; 3. Optimization in Infinite Dimensions; 4. Game Theory; 5. Stochastic Operations Research; 6. Decision Analysis, which are then expanded into multiple subtopics, each as a chapter. These four volumes are aimed at the following five major target audiences: University and College students Educators, Professional Practitioners, Research Personnel and Policy Analysts, Managers, and Decision Makers and NGOs.

**Issues in Industrial Relations and Management: 2011 Edition** 2012-01-09 Issues in Industrial Relations and Management: 2011 Edition is a ScholarlyEditions™ eBook that delivers timely, authoritative, and comprehensive information about Industrial Relations and Management. The editors have built Issues in Industrial Relations and Management: 2011 Edition on the vast information databases of ScholarlyNews.™ You can expect the information about Industrial Relations and Management in this eBook to be deeper than what you can access anywhere else, as well as consistently reliable, authoritative, informed, and relevant. The content of Issues in Industrial Relations and Management: 2011 Edition has been produced by the world's leading scientists, engineers, analysts, research institutions, and companies. All of the content is from peer-reviewed sources, and all of it is written,

assembled, and edited by the editors at ScholarlyEditions™ and available exclusively from us. You now have a source you can cite with authority, confidence, and credibility. More information is available at <http://www.ScholarlyEditions.com/>.

*Stochastic Models in Operations Research: Stochastic optimization* Daniel P. Heyman  
2004-01-01 This two-volume set of texts explores the central facts and ideas of stochastic processes, illustrating their use in models based on applied and theoretical investigations. They demonstrate the interdependence of three areas of study that usually receive separate treatments: stochastic processes, operating characteristics of stochastic systems, and stochastic optimization. Comprehensive in its scope, they emphasize the practical importance, intellectual stimulation, and mathematical elegance of stochastic models and are intended primarily as graduate-level texts.

**Finite State Continuous-time Markov Decision Processes, with Applications to a Class of Optimization Problems in Queueing Theory** Bruce L. Miller 1967

**Markov Decision Processes in Practice** Richard J. Boucherie 2017-03-10 This book presents classical Markov Decision Processes (MDP) for real-life applications and optimization. MDP allows users to develop and formally support approximate and simple decision rules, and this book showcases state-of-the-art applications in which MDP was key to the solution approach. The book is divided into six parts. Part 1 is devoted to the state-of-the-art theoretical foundation of MDP, including approximate methods such as policy improvement, successive approximation and infinite state spaces as well as an instructive chapter on Approximate Dynamic Programming. It then continues with five parts of specific and non-exhaustive application areas. Part 2 covers MDP healthcare applications, which includes different screening procedures, appointment scheduling, ambulance scheduling and blood management. Part 3 explores MDP modeling within transportation. This ranges from public to private transportation, from airports and traffic lights to car parking or charging your electric car. Part 4 contains three chapters that illustrates the structure of approximate policies for production or manufacturing structures. In Part 5, communications is highlighted as an important application area for MDP. It includes Gittins indices, down-to-earth call centers and wireless sensor networks. Finally Part 6 is dedicated to financial modeling, offering an instructive review to account for financial portfolios and derivatives under proportional transactional costs. The MDP applications in this book illustrate a variety of both standard and non-standard aspects of MDP modeling and its practical use. This book should appeal to readers for practicing, academic research and educational purposes, with a background in, among others, operations research, mathematics, computer science, and industrial engineering.

*A Theory for Semi-Markov Decision Processes with Unbounded Costs and Its Application to the Optimal Control of Queueing Systems* 1976 Semi-Markov decision processes with countable state and action spaces are investigated. The optimality criteria considered are the average cost criterion, the undiscounted cost criterion, and the discounted cost criterion. The common assumption of bounded costs has been replaced by some considerably weaker conditions. In particular, our assumptions are weaker than those made by Harrison, Hordijk, Lippman and Reed when they considered the same problem. The existence of optimal, stationary optimal and stationary E-optimal policies is investigated. Policy improvement is considered. Necessary and sufficient conditions for the optimality of a policy are given. Then

Downloaded from [avenza-dev.avenza.com](https://www.avenza-dev.avenza.com)  
on October 3, 2022 by guest

the optimal control of queueing systems is considered by formulating this general problem as a semi- Markov decision process. Finally, four different ways of proving the optimality of an unimprovable policy are developed in the context of queueing systems.

### **Optimization Theory, Decision Making, and Operations Research Applications**

Athanasios Migdalas 2012-11-28 These proceedings consist of 30 selected research papers based on results presented at the 10th Balkan Conference & 1st International Symposium on Operational Research (BALCOR 2011) held in Thessaloniki, Greece, September 22-24, 2011. BALCOR is an established biennial conference attended by a large number of faculty, researchers and students from the Balkan countries but also from other European and Mediterranean countries as well. Over the past decade, the BALCOR conference has facilitated the exchange of scientific and technical information on the subject of Operations Research and related fields such as Mathematical Programming, Game Theory, Multiple Criteria Decision Analysis, Information Systems, Data Mining and more, in order to promote international scientific cooperation. The carefully selected and refereed papers present important recent developments and modern applications and will serve as excellent reference for students, researchers and practitioners in these disciplines.

*Continuous-Time Markov Decision Processes* Xianping Guo 2009-09-18 Continuous-time Markov decision processes (MDPs), also known as controlled Markov chains, are used for modeling decision-making problems that arise in operations research (for instance, inventory, manufacturing, and queueing systems), computer science, communications engineering, control of populations (such as fisheries and epidemics), and management science, among many other fields. This volume provides a unified, systematic, self-contained presentation of recent developments on the theory and applications of continuous-time MDPs. The MDPs in this volume include most of the cases that arise in applications, because they allow unbounded transition and reward/cost rates. Much of the material appears for the first time in book form.

**Computational Probability** Winfried K. Grassmann 2000 Great advances have been made in recent years in the field of computational probability. In particular, the state of the art - as it relates to queueing systems, stochastic Petri-nets and systems dealing with reliability - has benefited significantly from these advances. The objective of this book is to make these topics accessible to researchers, graduate students, and practitioners. Great care was taken to make the exposition as clear as possible. Every line in the book has been evaluated, and changes have been made whenever it was felt that the initial exposition was not clear enough for the intended readership. The work of major research scholars in this field comprises the individual chapters of Computational Probability. The first chapter describes, in nonmathematical terms, the challenges in computational probability. Chapter 2 describes the methodologies available for obtaining the transition matrices for Markov chains, with particular emphasis on stochastic Petri-nets. Chapter 3 discusses how to find transient probabilities and transient rewards for these Markov chains. The next two chapters indicate how to find steady-state probabilities for Markov chains with a finite number of states. Both direct and iterative methods are described in Chapter 4. Details of these methods are given in Chapter 5. Chapters 6 and 7 deal with infinite-state Markov chains, which occur frequently in queueing, because there are times one does not want to set a bound for all queues. Chapter 8 deals with transforms, in particular Laplace transforms. The work of Ward Whitt and his collaborators, who have recently developed a number of numerical methods for

Laplace transform inversions, is emphasized in this chapter. Finally, if one wants to optimize a system, one way to do the optimization is through Markov decision making, described in Chapter 9. Markov modeling has found applications in many areas, three of which are described in detail: Chapter 10 analyzes discrete-time queues, Chapter 11 describes networks of queues, and Chapter 12 deals with reliability theory.

**Constrained Markov Decision Processes** Eitan Altman 1999-03-30 This book provides a unified approach for the study of constrained Markov decision processes with a finite state space and unbounded costs. Unlike the single controller case considered in many other books, the author considers a single controller with several objectives, such as minimizing delays and loss, probabilities, and maximization of throughputs. It is desirable to design a controller that minimizes one cost objective, subject to inequality constraints on other cost objectives. This framework describes dynamic decision problems arising frequently in many engineering fields. A thorough overview of these applications is presented in the introduction. The book is then divided into three sections that build upon each other. The first part explains the theory for the finite state space. The author characterizes the set of achievable expected occupation measures as well as performance vectors, and identifies simple classes of policies among which optimal policies exist. This allows the reduction of the original dynamic into a linear program. A Lagrangian approach is then used to derive the dual linear program using dynamic programming techniques. In the second part, these results are extended to the infinite state space and action spaces. The author provides two frameworks: the case where costs are bounded below and the contracting framework. The third part builds upon the results of the first two parts and examines asymptotical results of the convergence of both the value and the policies in the time horizon and in the discount factor. Finally, several state truncation algorithms that enable the approximation of the solution of the original control problem via finite linear programs are given.