

Markov Decision Processes With Applications To Finance Universitext

As recognized, adventure as with ease as experience very nearly lesson, amusement, as competently as arrangement can be gotten by just checking out a book **Markov Decision Processes With Applications To Finance Universitext** with it is not directly done, you could understand even more in the region of this life, something like the world.

We offer you this proper as with ease as easy habit to get those all. We manage to pay for Markov Decision Processes With Applications To Finance Universitext and numerous book collections from fictions to scientific research in any way. in the course of them is this Markov Decision Processes With Applications To Finance Universitext that can be your partner.

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.

Stochastic Modeling and Optimization David D. Yao 2012-12-06 This books covers the broad range of research in stochastic models and optimization.

Applications presented include networks, financial engineering, production planning, and supply chain management. Each contribution is aimed at graduate students working in operations research, probability, and statistics.

Constructive Computation in Stochastic Models with Applications Quan-Lin Li 2011-02-02 "Constructive Computation in Stochastic Models with Applications: The RG-Factorizations" provides a unified, constructive and algorithmic framework for numerical computation of many practical stochastic systems. It summarizes recent important advances in computational study of stochastic models from several crucial directions, such as stationary computation, transient solution, asymptotic analysis, reward processes, decision processes, sensitivity analysis as well as game theory. Graduate students, researchers and practicing engineers in the field of operations research, management sciences, applied probability, computer networks, manufacturing systems, transportation systems, insurance and finance, risk management and biological sciences will find this book valuable. Dr. Quan-Lin Li is an Associate Professor at the Department of Industrial Engineering of Tsinghua University, China.

Markov Decision Processes in Artificial Intelligence Olivier Sigaud 2013-03-04 Markov Decision Processes (MDPs) are a mathematical framework for modeling sequential decision problems under uncertainty as well as Reinforcement Learning problems. Written by experts in the field, this book provides a global view of current research using MDPs in Artificial Intelligence. It starts with an introductory presentation of the fundamental aspects of MDPs (planning in MDPs, Reinforcement Learning, Partially Observable MDPs, Markov games and the use of non-classical criteria). Then it presents more advanced research trends in the domain and gives some concrete examples using illustrative applications.

Alternative Decision-Making Models for Financial Portfolio Management: Emerging Research and Opportunities Spaseski, Narela 2017-08-11 Economics is an integral aspect to every successful society, yet basic financial practices have gone unchanged for decades. Analyzing unconventional finance methods can provide new ways to ensure personal financial futures on an individual level, as well as boosting international economies. *Alternative Decision-Making Models for Financial Portfolio Management: Emerging Research and Opportunities* is an essential reference source that discusses methods and techniques that make financial administration more efficient for professionals in economic fields. Featuring relevant topics such as mean-variance portfolio theory, decision tree analysis, risk protection strategies, and asset-liability management, this publication is ideal for academicians, students, economists, and researchers that would like to stay current on new and innovative methods to transform the financial realm.

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.

Stochastic Processes in Science, Engineering and Finance Frank Beichelt 2006-02-22 This book presents a self-contained introduction to stochastic processes with emphasis on their applications in science, engineering, finance, computer science, and operations research. It provides theoretical foundations for modeling time-dependent random phenomena in these areas and illustrates their application by analyzing numerous practical examples. The treatment assumes few prerequisites, requiring only the standard mathematical maturity acquired by undergraduate applied science students. It includes an introductory chapter that summarizes the basic probability theory needed as background. Numerous exercises reinforce the concepts and techniques discussed and allow readers to assess their grasp of the subject. Solutions to most of the exercises are provided in an appendix. While focused primarily on practical aspects, the presentation includes some important proofs along with more challenging examples and exercises for those more theoretically inclined. Mastering the contents of this book prepares readers to apply stochastic modeling in their own fields and enables them to work more creatively with software designed for dealing with the data analysis aspects of stochastic processes.

Hidden Markov Models in Finance Rogemar S. Mamon 2007-04-26 A number of methodologies have been employed to provide decision making solutions globalized markets. Hidden Markov Models in Finance offers the first systematic application of these methods to specialized financial problems: option pricing, credit risk modeling, volatility estimation and more. The book provides tools for sorting through turbulence, volatility, emotion, chaotic events – the random "noise" of financial markets – to analyze core components.

Secondary Analysis of Electronic Health Records MIT Critical Data 2016-09-09 This book trains the next generation of scientists representing different disciplines to leverage the data generated during routine patient care. It formulates a more complete lexicon of evidence-based recommendations and support shared, ethical decision making by doctors with their patients. Diagnostic and therapeutic technologies continue to evolve rapidly, and both individual practitioners and clinical teams face increasingly complex ethical decisions. Unfortunately, the current state of medical knowledge does not provide the guidance to make the majority of clinical decisions on the basis of evidence. The present research infrastructure is inefficient and frequently produces unreliable results that cannot be replicated. Even randomized controlled trials (RCTs), the traditional gold standards of the research reliability hierarchy, are not without limitations. They can be costly, labor intensive, and slow, and can return results that are seldom generalizable to every patient population. Furthermore, many pertinent but unresolved clinical and medical systems issues do not seem to have attracted the interest of the research

enterprise, which has come to focus instead on cellular and molecular investigations and single-agent (e.g., a drug or device) effects. For clinicians, the end result is a bit of a “data desert” when it comes to making decisions. The new research infrastructure proposed in this book will help the medical profession to make ethically sound and well informed decisions for their patients.

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 ofSection 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-studiessequential optimization ofdiscrete 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 ofimpacts: (i) they cost orsavetime, 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 offuture events. MDPs model this paradigm and provide results on the structure and existence of good policies and on methods for their calculation.

Markov Chains and Decision Processes for Engineers and Managers Theodore J Sheskin 2019-08-30 Recognized as a powerful tool for dealing with uncertainty, Markov modeling can enhance your ability to analyze complex production and service systems. However, most books on Markov chains or decision processes are often either highly theoretical, with few examples, or highly prescriptive, with little justification for the steps of the algorithms used to solve Markov models. Providing a unified treatment of Markov chains and Markov decision processes in a single volume, Markov Chains and Decision Processes for Engineers and Managers supplies a highly detailed description of the construction and solution of Markov models that facilitates their application to diverse processes. Organized around Markov chain structure, the book begins with descriptions of Markov chain states, transitions, structure, and models, and then discusses steady state distributions and passage to a target state in a regular Markov chain. The author treats canonical forms and passage to target states or to classes of target states for reducible Markov chains. He adds an economic dimension by associating rewards with states, thereby linking a Markov chain to a Markov decision process, and then adds decisions to create a Markov decision process, enabling an analyst to choose among alternative Markov chains with rewards so as to maximize expected rewards. An introduction to state reduction and hidden Markov chains rounds out the coverage. In a presentation that balances algorithms and applications, the author provides explanations of the logical relationships that underpin the formulas or algorithms through informal derivations, and devotes considerable attention to the construction of Markov models. He constructs simplified Markov models for a wide assortment of processes such as the weather, gambling, diffusion of gases, a waiting line, inventory, component replacement, machine maintenance, selling a stock, a charge account, a career path, patient flow

Applied Stochastic Models and Control for Finance and Insurance Charles S. Tapiero 2012-12-06 Applied Stochastic Models and Control for Finance and Insurance presents at an introductory level some essential stochastic models applied in economics, finance and insurance. Markov chains, random walks, stochastic differential equations and other stochastic processes are used throughout the book and systematically applied to economic and financial applications. In addition, a dynamic programming framework is used to deal with some basic optimization problems. The book begins by introducing problems of economics, finance and insurance which involve time, uncertainty and risk. A number of cases are treated in detail, spanning risk management, volatility, memory, the time structure of preferences, interest rates and yields, etc. The second and third chapters provide an introduction to stochastic models and their application. Stochastic differential equations and stochastic calculus are presented in an intuitive manner, and numerous applications and exercises are used to facilitate their understanding and their use in Chapter 3. A number of other processes which are increasingly used in finance and insurance are introduced in Chapter 4. In the fifth chapter, ARCH and GARCH models are presented and their application to modeling volatility is emphasized. An outline of decision-making procedures is presented in Chapter 6. Furthermore, we also introduce the essentials of stochastic dynamic programming and control, and provide first steps for the student who seeks to apply these techniques. Finally, in Chapter 7, numerical techniques and approximations to stochastic processes are examined. This book can be used in business, economics, financial engineering and decision sciences schools for second year Master's students, as well as in a number of courses widely given in departments of statistics, systems and decision sciences.

Markov Chains Wai-Ki Ching 2015-04-05 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.

Stochastic Model Checking Anne Remke 2014-11-03 The use of stochastic models in computer science is wide spread, for instance in performance

modeling, analysis of randomized algorithms and communication protocols which form the structure of the Internet. Stochastic model checking is an important field in stochastic analysis. It has rapidly gained popularity, due to its powerful and systematic methods to model and analyze stochastic systems. This book presents 7 tutorial lectures given by leading scientists at the ROCKS Autumn School on Stochastic Model Checking, held in Vahrn, Italy, in October 2012. The 7 chapters of this tutorial went through two rounds of reviewing and improvement and are summarizing the state-of-the-art in the field, centered around the tree areas of stochastic models, abstraction techniques and stochastic model checking.

Spatial Statistics and Spatio-Temporal Data Michael Sherman 2011-01-06 In the spatial or spatio-temporal context, specifying the correct covariance function is fundamental to obtain efficient predictions, and to understand the underlying physical process of interest. This book focuses on covariance and variogram functions, their role in prediction, and appropriate choice of these functions in applications. Both recent and more established methods are illustrated to assess many common assumptions on these functions, such as, isotropy, separability, symmetry, and intrinsic correlation. After an extensive introduction to spatial methodology, the book details the effects of common covariance assumptions and addresses methods to assess the appropriateness of such assumptions for various data structures. Key features: An extensive introduction to spatial methodology including a survey of spatial covariance functions and their use in spatial prediction (kriging) is given. Explores methodology for assessing the appropriateness of assumptions on covariance functions in the spatial, spatio-temporal, multivariate spatial, and point pattern settings. Provides illustrations of all methods based on data and simulation experiments to demonstrate all methodology and guide to proper usage of all methods. Presents a brief survey of spatial and spatio-temporal models, highlighting the Gaussian case and the binary data setting, along with the different methodologies for estimation and model fitting for these two data structures. Discusses models that allow for anisotropic and nonseparable behaviour in covariance functions in the spatial, spatio-temporal and multivariate settings. Gives an introduction to point pattern models, including testing for randomness, and fitting regular and clustered point patterns. The importance and assessment of isotropy of point patterns is detailed. Statisticians, researchers, and data analysts working with spatial and space-time data will benefit from this book as well as will graduate students with a background in basic statistics following courses in engineering, quantitative ecology or atmospheric science.

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 rig orous treatment that combines two significant research topics: Stochastic Games and Markov Decision Processes, which have been studied exten sively, and at times quite independently, by mathematicians, operations researchers, engineers, and economists. Since Markov decision processes can be viewed as a special noncompeti tive case of stochastic games, we introduce the new terminology Competi tive 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 program ming, 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.

A Primer on Experiments with Mixtures John A. Cornell 2011-09-26 The concise yet authoritative presentation of key techniques for basic mixtures experiments Inspired by the author's bestselling advanced book on the topic, A Primer on Experiments with Mixtures provides an introductory presentation of the key principles behind experimenting with mixtures. Outlining useful techniques through an applied approach with examples from real research situations, the book supplies a comprehensive discussion of how to design and set up basic mixture experiments, then analyze the data and draw inferences from results. Drawing from his extensive experience teaching the topic at various levels, the author presents the mixture experiments in an easy-to-follow manner that is void of unnecessary formulas and theory. Succinct presentations explore key methods and techniques for carrying out basic mixture experiments, including: Designs and models for exploring the entire simplex factor space, with coverage of simplex-lattice and simplex-centroid designs, canonical polynomials, the plotting of individual residuals, and axial designs Multiple constraints on the component proportions in the form of lower and/or upper bounds, introducing L-Pseudocomponents, multicomponent constraints, and multiple lattice designs for major and minor component classifications Techniques for analyzing mixture data such as model reduction and screening components, as well as additional topics such as measuring the leverage of certain design points Models containing ratios of the components, Cox's mixture polynomials, and the fitting of a slack variable model A review of least squares and the analysis of variance for fitting data Each chapter concludes with a summary and appendices with details on the technical aspects of the material. Throughout the book, exercise sets with selected answers allow readers to test their comprehension of the material, and References and Recommended Reading sections outline further resources for study of the presented topics. A Primer on Experiments with Mixtures is an excellent book for one-semester courses on mixture designs and can also serve as a supplement for design of experiments courses at the upper-undergraduate and graduate levels. It is also a suitable reference for practitioners and researchers who have an interest in experiments with mixtures and would like to learn more about the related mixture designs and models.

Modelling Under Risk and Uncertainty Etienne de Rocquigny 2012-04-30 "This volume addresses a concern of very high relevance and growing interest for large industries or environmentalists: risk and uncertainty in complex systems. It gives new insight on the peculiar mathematical challenges generated by recent industrial safety or environmental control analysis, focusing on implementing decision theory choices related to risk and uncertainty analysis through statistical estimation and computation, in the presence of physical modeling and risk analysis. The result will lead statisticians and associated professionals to formulate and solve new challenges at the frontier between statistical modeling, physics, scientific computing, and risk analysis"--

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.

Applied Probability Models with Optimization Applications Sheldon M. Ross 1992-01-01 Includes bibliographical references and index.

Mathematical Modeling of Random and Deterministic Phenomena Solym Mawaki Manou-Abi 2020-02-19 This book highlights mathematical research interests that appear in real life, such as the study and modeling of random and deterministic phenomena. As such, it provides current research in mathematics, with applications in biological and environmental sciences, ecology, epidemiology and social perspectives. The chapters can be read independently of each other, with dedicated references specific to each chapter. The book is organized in two main parts. The first is devoted to some advanced mathematical problems regarding epidemic models; predictions of biomass; space-time modeling of extreme rainfall; modeling with the piecewise deterministic Markov process; optimal control problems; evolution equations in a periodic environment; and the analysis of the heat equation. The second is devoted to a modelization with interdisciplinarity in ecological, socio-economic, epistemological, demographic and social problems. Mathematical Modeling of Random and Deterministic Phenomena is aimed at expert readers, young researchers, plus graduate and advanced undergraduate students who are interested in probability, statistics, modeling and mathematical analysis.

Modern Trends in Controlled Stochastic Processes: Alexey Piunovskiy 2021-06-04 This book presents state-of-the-art solution methods and applications of stochastic optimal control. It is a collection of extended papers discussed at the traditional Liverpool workshop on controlled stochastic processes with participants from both the east and the west. New problems are formulated, and progresses of ongoing research are reported. Topics covered in this book include theoretical results and numerical methods for Markov and semi-Markov decision processes, optimal stopping of Markov processes, stochastic

games, problems with partial information, optimal filtering, robust control, Q-learning, and self-organizing algorithms. Real-life case studies and applications, e.g., queueing systems, forest management, control of water resources, marketing science, and healthcare, are presented. Scientific researchers and postgraduate students interested in stochastic optimal control,- as well as practitioners will find this book appealing and a valuable reference.

Examples in Markov Decision Processes A. B. Piunovskiy 2013 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.

Markov Processes for Stochastic Modeling Oliver Ibe 2013-05-22 Markov processes are processes that have limited memory. In particular, their dependence on the past is only through the previous state. They are used to model the behavior of many systems including communications systems, transportation networks, image segmentation and analysis, biological systems and DNA sequence analysis, random atomic motion and diffusion in physics, social mobility, population studies, epidemiology, animal and insect migration, queueing systems, resource management, dams, financial engineering, actuarial science, and decision systems. Covering a wide range of areas of application of Markov processes, this second edition is revised to highlight the most important aspects as well as the most recent trends and applications of Markov processes. The author spent over 16 years in the industry before returning to academia, and he has applied many of the principles covered in this book in multiple research projects. Therefore, this is an applications-oriented book that also includes enough theory to provide a solid ground in the subject for the reader. Presents both the theory and applications of the different aspects of Markov processes Includes numerous solved examples as well as detailed diagrams that make it easier to understand the principle being presented Discusses different applications of hidden Markov models, such as DNA sequence analysis and speech analysis.

Planning with Markov Decision Processes Mausam Natarajan 2022-06-01 Markov Decision Processes (MDPs) are widely popular in Artificial Intelligence for modeling sequential decision-making scenarios with probabilistic dynamics. They are the framework of choice when designing an intelligent agent that needs to act for long periods of time in an environment where its actions could have uncertain outcomes. MDPs are actively researched in two related subareas of AI, probabilistic planning and reinforcement learning. Probabilistic planning assumes known models for the agent's goals and domain dynamics, and focuses on determining how the agent should behave to achieve its objectives. On the other hand, reinforcement learning additionally learns these models based on the feedback the agent gets from the environment. This book provides a concise introduction to the use of MDPs for solving probabilistic planning problems, with an emphasis on the algorithmic perspective. It covers the whole spectrum of the field, from the basics to state-of-the-art optimal and approximation algorithms. We first describe the theoretical foundations of MDPs and the fundamental solution techniques for them. We then discuss modern optimal algorithms based on heuristic search and the use of structured representations. A major focus of the book is on the numerous approximation schemes for MDPs that have been developed in the AI literature. These include determinization-based approaches, sampling techniques, heuristic functions, dimensionality reduction, and hierarchical representations. Finally, we briefly introduce several extensions of the standard MDP classes that model and solve even more complex planning problems. Table of Contents: Introduction / MDPs / Fundamental Algorithms / Heuristic Search Algorithms / Symbolic Algorithms / Approximation Algorithms / Advanced Notes

Mathematical Finance Ernst Eberlein 2019-12-03 Taking continuous-time stochastic processes allowing for jumps as its starting and focal point, this book provides an accessible introduction to the stochastic calculus and control of semimartingales and explains the basic concepts of Mathematical Finance such as arbitrage theory, hedging, valuation principles, portfolio choice, and term structure modelling. It bridges thegap between introductory texts and the advanced literature in the field. Most textbooks on the subject are limited to diffusion-type models which cannot easily account for sudden price movements. Such abrupt changes, however, can often be observed in real markets. At the same time, purely discontinuous processes lead to a much wider variety of flexible and tractable models. This explains why processes with jumps have become an established tool in the statistics and mathematics of finance. Graduate students, researchers as well as practitioners will benefit from this monograph.

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.

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 Lagranian 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.

Optimization, Control, and Applications of Stochastic Systems Daniel Hernández-Hernández 2012-08-15 This volume provides a general overview of discrete- and continuous-time Markov control processes and stochastic games, along with a look at the range of applications of stochastic control and some of its recent theoretical developments. These topics include various aspects of dynamic programming, approximation algorithms, and infinite-dimensional linear programming. In all, the work comprises 18 carefully selected papers written by experts in their respective fields. Optimization, Control, and Applications of Stochastic Systems will be a valuable resource for all practitioners, researchers, and professionals in applied mathematics and operations research who work in the areas of stochastic control, mathematical finance, queueing theory, and inventory systems. It may also serve as a supplemental text for graduate courses in optimal control and dynamic games.

Controlled Markov Processes and Viscosity Solutions Wendell H. Fleming 2006-02-04 This book is an introduction to optimal stochastic control for continuous time Markov processes and the theory of viscosity solutions. It covers dynamic programming for deterministic optimal control problems, as well as to the corresponding theory of viscosity solutions. New chapters in this second edition introduce the role of stochastic optimal control in portfolio optimization and in pricing derivatives in incomplete markets and two-controller, zero-sum differential games.

Distributed Sensing and Intelligent Systems Mohamed Elhoseny

Reinforcement Learning and Stochastic Optimization Warren B. Powell 2022-03-15 REINFORCEMENT LEARNING AND STOCHASTIC OPTIMIZATION Clearing the jungle of stochastic optimization Sequential decision problems, which consist of “decision, information, decision, information,” are ubiquitous, spanning virtually every human activity ranging from business applications, health (personal and public health, and medical decision making), energy, the sciences, all fields of engineering, finance, and e-commerce. The diversity of applications attracted the attention of at least 15 distinct fields of research, using eight distinct notational systems which produced a vast array of analytical tools. A byproduct is that powerful tools developed in one community may be unknown to other communities. Reinforcement Learning and Stochastic Optimization offers a single canonical framework that can model any sequential decision problem using five core components: state variables, decision variables, exogenous information variables, transition function, and objective function. This book highlights twelve types of uncertainty that might enter any model and pulls together the diverse set of methods for making decisions, known as policies, into four fundamental classes that span every method suggested in the academic literature or used in practice. Reinforcement Learning and Stochastic Optimization is the first book to provide a balanced treatment of the different methods for modeling and solving sequential decision problems, following the style used by most books on machine learning, optimization, and simulation. The presentation is designed for readers with a course in probability and statistics, and an interest in modeling and applications. Linear programming is occasionally used for specific problem classes. The book is designed for readers who are new to the field, as well as those with some background in optimization under uncertainty. Throughout this book, readers will find references to over 100 different applications, spanning pure learning problems, dynamic resource allocation problems, general state-dependent problems, and hybrid learning/resource allocation problems such as those that arose in the COVID pandemic. There are 370 exercises, organized into seven groups, ranging from review questions, modeling, computation, problem solving, theory, programming exercises and a “diary problem” that a reader chooses at the beginning of the book, and which is used as a basis for questions throughout the rest of the book.

Stochastic Processes for Finance

Hidden Markov Models in Finance Rogemar S. Mamon 2014-05-14 Since the groundbreaking research of Harry Markowitz into the application of operations research to the optimization of investment portfolios, finance has been one of the most important areas of application of operations research. The use of hidden Markov models (HMMs) has become one of the hottest areas of research for such applications to finance. This handbook offers systemic applications of different methodologies that have been used for decision making solutions to the financial problems of global markets. As the follow-up to the authors’ Hidden Markov Models in Finance (2007), this offers the latest research developments and applications of HMMs to finance and other related fields. Amongst the fields of quantitative finance and actuarial science that will be covered are: interest rate theory, fixed-income instruments, currency market, annuity and insurance policies with option-embedded features, investment strategies, commodity markets, energy, high-frequency trading, credit risk, numerical algorithms, financial econometrics and operational risk. Hidden Markov Models in Finance: Further Developments and Applications, Volume II presents recent applications and case studies in finance and showcases the formulation of emerging potential applications of new research over the book’s 11 chapters. This will benefit not only researchers in financial modeling, but also others in fields such as engineering, the physical sciences and social sciences. Ultimately the handbook should prove to be a valuable resource to dynamic researchers interested in taking full advantage of the power and versatility of HMMs in accurately and efficiently capturing many of the processes in the financial market.

Selected Topics on Continuous-Time Controlled Markov Chains and Markov Games Tomás Prieto-Rumeau 2012-03-16 This book concerns continuous-time controlled Markov chains, also known as continuous-time Markov decision processes. They form a class of stochastic control problems in which a single decision-maker wishes to optimize a given objective function. This book is also concerned with Markov games, where two decision-makers (or players) try to optimize their own objective function. Both decision-making processes appear in a large number of applications in economics, operations research, engineering, and computer science, among other areas. An extensive, self-contained, up-to-date analysis of basic optimality criteria (such as discounted and average reward), and advanced optimality criteria (e.g., bias, overtaking, sensitive discount, and Blackwell optimality) is presented. A particular emphasis is made on the application of the results herein: algorithmic and computational issues are discussed, and applications to population models and epidemic processes are shown. This book is addressed to students and researchers in the fields of stochastic control and stochastic games. Moreover, it could be of interest also to undergraduate and beginning graduate students because the reader is not supposed to have a high mathematical background: a working knowledge of calculus, linear algebra, probability, and continuous-time Markov chains should suffice to understand the contents of the book.

Contents:IntroductionControlled Markov ChainsBasic Optimality CriteriaPolicy Iteration and Approximation TheoremsOvertaking, Bias, and Variance

OptimalitySensitive Discount OptimalityBlackwell OptimalityConstrained Controlled Markov ChainsApplicationsZero-Sum Markov GamesBias and Overtaking Equilibria for Markov Games Readership: Graduate students and researchers in the fields of stochastic control and stochastic analysis.

Keywords:Markov Decision Processes;Continuous-Time Controlled Markov Chains;Stochastic Dynamic Programming;Stochastic GamesKey Features:This book presents a reader-friendly, extensive, self-contained, and up-to-date analysis of advanced optimality criteria for continuous-time controlled Markov chains and Markov games. Most of the material herein is quite recent (it has been published in high-impact journals during the last five years) and it appears in book form for the first timeThis book introduces approximation theorems which, in particular, allow the reader to obtain numerical approximations of the solution to several control problems of practical interest. To the best of our knowledge, this is the first time that such computational issues are studied for denumerable state continuous-time controlled Markov chains. Hence, the book has an adequate balance between, on the one hand, theoretical results and, on the other hand, applications and computational issuesThe books that analyze continuous-time controlled Markov chains usually restrict themselves to the case of bounded transition and reward rates, which can be reduced to discrete-time models by using the uniformization technique. In our case, however, the transition and the reward rates might be unbounded, and so the uniformization technique cannot be used. By the way, let us mention that in models of practical interest the transition and the reward rates are, typically, unboundedReviews:“The book contains a large number of recent research results on CMCs and Markov games and puts them in perspective. It is written in a very conscious manner, contains detailed proofs of all main results, as well as extensive bibliographic remarks. The book is a very valuable piece of work for researchers on continuous-time CMCs and Markov games.”Zentralblatt MATH

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.

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).

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.

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

Applications of Stochastic Models in Finance Reza Habibi 2014-08-14 This note contains some applications of stochastic models in finance. For example, we survey Markov Decision Processes, Bayesian Networks, Adaptive Control, Black-Scholes Pricing methods. This note involves the change point analysis in some financial models, risk management, portfolio selection and credit scoring in financial institutions. Some papers are too short, however, we have studied an observation.