

Markov Decision Processes With Applications To Finance Universitext

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 Glittis 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 practicioning, academic research and educational purposes, with a background in, among others, operations research, mathematics, computer science, and industrial engineering.

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).

The study examines a class of partially observable sequential decision models motivated by the process of machine maintenance and corrective action or medical diagnosis and treatment. Emphasis is placed on the dynamics of the state, i.e., the possibility that the machine (disease) state changes during the decision process. This is incorporated in the form of a Markov chain. It is also assumed that the state is only indirectly observable via outputs probabilistically related to the state. The end result is a model which is a discrete time Markov decision process with a continuous state space, a finite action space, and a special transition structure. (Modified author abstract).

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.

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.

The Wiley-Interscience Paperback Series consists of selected booksthat have been made more accessible to consumers in an effort toincrease global appeal and general circulation. With these newunabridged softcover volumes, Wiley hopes to extend the lives ofthese works by making them available to future generations ofstatisticians, mathematicians, and scientists. "This text is unique in bringingtogether so many resultshitherto found only in part in other texts and papers. . . . Thetext is fairly self-contained, inclusive of some basic mathematicalresults needed, and provides a rich diet of examples, applications, and exercises. The bibliographical material at the end ofeachchapter is excellent, not only from a historical perspective, butbecause it is valuable for researchers in acquiring a goodperspective of the MDP research potential." —Zentralblatt fur Mathematik ". . . it is of great value to advanced-level students,researchers, and professional practitioners of this field to havenow a complete volume (with more than 600 pages) devoted to thistopic. . . . Markov Decision Processes: Discrete Stochastic DynamicProgramming represents an up-to-date, unified, and rigorousstatement of theoretical and computational aspects of discrete-timeMarkov decision processes." —Journal of the American Statistical Association

[Reinforcement Learning, second edition](#)

[Technical Report, No.15](#)

[Markov Chains: Models, Algorithms and Applications](#)

[Theoretical Minimum for Conducting 21st Century Research on Healthcare Operations](#)

[An AI Perspective](#)

[Methods of Machine Learning and their Application: The Basics of Markov Decision Processes and Reinforcement Learning](#)

[Simulation-based Algorithms for Markov Decision Processes](#)

[Theory and Applications](#)

[An Introduction](#)

[Recent Developments in Markov Decision Processes](#)

Semi-Markov Processes: Applications in System Reliability and Maintenance is a modern view of discrete state space and continuous time semi-Markov processes and their applications in reliability and maintenance. The book explains how to construct semi-Markov models and discusses the different reliability parameters and characteristics that can be obtained from those models. The book is a useful resource for mathematicians, engineering practitioners, and PhD and MSc students who want to understand the basic concepts and results of semi-Markov process theory. Clearly defines the properties and theorems from discrete state Semi-Markov Process (SMP) theory. Describes the method behind constructing Semi-Markov (SM) models and SM decision models in the field of reliability and maintenance. Provides numerous individual versions of SM models, including the most recent and their impact on system reliability and maintenance.

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.

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 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 determination-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

One of the central problems in operations research and management science is how to quantify the effects of uncertainty about the future. This, the second volume in a series of handbooks, is devoted to models where chance events play a major role. The thirteen chapters survey topics in applied probability that have been particularly useful in operations research and management science. Each chapter was written by an expert, both in subject matter and in its exposition. The chapters fall into four groups. The first four cover the fundamentals of stochastic processes, and lay the foundation for the following chapters. The next three chapters are concerned with methods of getting numbers. This includes numerical solution of models, parameter estimation for models, and simulation of models. Chapters 8 and 9 describe the fundamentals of dynamic optimization. The last four chapters are concerned with the most important structured models in operations research and management science: queues, queueing networks, inventories, and reliability.

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.

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 asymptotic 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.

[Examples in Markov Decision Processes](#)

[Hierarchical Markov Decision Processes](#)

[Markov Decision Processes with Information Costs](#)

[Theory, Algorithms, and Applications](#)

[Methods and Applications](#)

[Theory and Application](#)

[Markov Decision Processes in Artificial Intelligence](#)

[Robust and Risk-sensitive Markov Decision Processes with Applications to Dynamic Optimal Reinsurance](#)

[Markov Processes for Stochastic Modeling](#)

[Markov Chains and Decision Processes for Engineers and Managers](#)

Bachelor Thesis from the year 2021, in the Subject Engineering - Computer Engineering, grade: 1,3, Hamburg University of Technology (Embedded Systems), language: English, abstract: This bachelor thesis aims to illustrate the idea behind Markov Decision Processes (MDP) and to present a few basic methods of Reinforcement Learning (RL) namely Monte Carlo Learning and Q-Learning, which are the solutions for decision problems modelled by MDPs. For the last section we apply these methods on an application and in the end discuss the results. Let us imagine the scenario where we put a hamster inside a maze, we expect the hamster to go through the maze till it reaches some point we considered as the goal. Well, it may randomly work but most of the time it won't. At this place, the hamster does not know how important this particular point remains namely the goal. But how will it be, when we remunerate the hamster once the goal is reached, he receives a reward for example a piece of cheese. The hamster will start to remember the route, which leads to the cheese and he maybe will learn to go the easy and quick way to achieve this goal. What we did, is that we reinforce the good behavior of the hamster by giving it some reward.

An introduction to decision making under uncertainty from a computational perspective, covering both theory and applications ranging from speech recognition to airborne collision avoidance. Many important problems involve decision making under uncertainty—that is, choosing actions based on often imperfect observations, with unknown outcomes. Designers of automated decision support systems must take into account the various sources of uncertainty while balancing the multiple objectives of the system. This book provides an introduction to the challenges of decision making under uncertainty from a computational perspective. It presents both the theory behind decision making models and algorithms and a collection of example applications that range from speech recognition to aircraft collision avoidance. Focusing on two methods for designing decision agents, planning and reinforcement learning, the book covers probabilistic models, introducing Bayesian networks as a graphical model that captures probabilistic relationships between variables; utility theory as a framework for understanding optimal decision making under uncertainty; Markov decision processes as a method for modeling sequential problems; model uncertainty; state uncertainty; and cooperative decision making involving multiple interacting agents. A series of applications shows how the theoretical concepts can be applied to systems for attribute-based person search, speech applications, collision avoidance, and unmanned aircraft persistent surveillance. Decision Making Under Uncertainty unifies research from different communities using consistent notation, and is accessible to students and researchers across engineering disciplines who have some prior exposure to probability theory and calculus. It can be used as a text for advanced undergraduate and graduate students in fields including computer science, aerospace and electrical engineering, and management science. It will also be a valuable professional reference for researchers in a variety of disciplines.

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-structured 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 resources, 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.

neither on the past history nor the present state of the line-lengths. Covering formulation, algorithms, and structural results, and linking theory to real-world applications in controlled sensing (including social learning, adaptive radars and sequential detection), this book focuses on the conceptual foundations of partially observed Markov decision processes (POMDPs). It emphasizes structural results in stochastic dynamic programming, enabling graduate students and researchers in engineering, operations research, and economics to understand the underlying unifying themes without getting weighed down by mathematical technicalities. Bringing together research from across the literature, the book provides an introduction to nonlinear filtering followed by a systematic development of stochastic dynamic programming, lattice programming and reinforcement learning for POMDPs. Questions addressed in the book include: when does a POMDP have a threshold optimal policy? When are myopic policies optimal? How do local and global decision makers interact in adaptive decision making in multi-agent social learning where there is herding and data incest? And how can sophisticated radars and sensors adapt their sensing in real time?

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.

[With Applications to a Class of Optimization Problems in Queueing Theory](#)

[Finite State Continuous-time Markov Decision Processes, with Applications to a Class of Optimization Problems in Queueing Theory](#)

[From Filtering to Controlled Sensing](#)

[Continuous-Time Markov Decision Processes](#)

[Handbook of Markov Decision Processes](#)

[Decision Making Under Uncertainty](#)

[Application of Markov Decision Processes to Search Problems](#)

[NON-HOMOGENEOUS MARKOV DECISION PROCESSES WITH APPLICATIONS TO R&D PLANNING \(FORECAST HORIZONS, INFINITE OPTIMIZATION, RESEARCH DEVELOPMENT\)](#)

[Stochastic Models](#)

[Partially Observable Markov Decision Processes with Applications](#)

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 (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.

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.

structure is exploited to develop practical solution techniques. Economic and policy implications of the model are examined and comparisons with existing models are made.

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 problems. 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 in a hospital, marketing, and a production line. This treatment helps you harness the power of Markov modeling and apply it to your organization's processes.

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.

[Markov Decision Processes with Their Applications](#)

[Markov Decision Processes in Practice](#)

[Finite State Continuous-time Markov Decision Processes](#)

[Planning with Markov Decision Processes](#)

[Discrete Stochastic Dynamic Programming](#)

[Competitive Markov Decision Processes](#)

[Partially Observed Markov Decision Processes](#)

[Semi-Markov Processes: Applications in System Reliability and Maintenance](#)

[Selected Topics on Continuous-time Controlled Markov Chains and Markov Games](#)

[CONSTRAINED MARKOV DECISION PROCESSES WITH QUEUEING APPLICATIONS](#)

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.

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 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.

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.

[Markov Decision Processes](#)

[Borel Space Models and General Control Strategies](#)

[Markov Decision Processes with Applications to Finance](#)

[Theory and Applications of Markov Decision Processes and Their Perturbations](#)

[The Theory and Applications](#)

[Constrained Markov Decision Processes](#)

[Handbook of Healthcare Analytics](#)