

Markov Decision Processes With Applications To Finance Universitext

If you ally need such a referred **Markov Decision Processes With Applications To Finance Universitext** book that will present you worth, get the no question best seller from us currently from several preferred authors. If you want to humorous books, lots of novels, tale, jokes, and more fictions collections are plus launched, from best seller to one of the most current released.

You may not be perplexed to enjoy every ebook collections Markov Decision Processes With Applications To Finance Universitext that we will totally offer. It is not a propos the costs. Its about what you habit currently. This Markov Decision Processes With Applications To Finance Universitext , as one of the most operational sellers here will entirely be along with the best options to review.

*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. *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.

Stochastic Processes - Robert G. Gallager 2013-12-12

This definitive textbook provides a solid introduction to discrete and continuous stochastic processes, tackling a complex field in a way that

instils a deep understanding of the relevant mathematical principles, and develops an intuitive grasp of the way these principles can be applied to modelling real-world systems. It includes a careful review of elementary probability and detailed coverage of Poisson, Gaussian and Markov processes with richly varied queuing applications. The theory and applications of inference, hypothesis testing, estimation, random walks, large deviations, martingales and investments are developed. Written by one of the world's leading information theorists, evolving over twenty years of graduate classroom teaching and enriched by over 300 exercises, this is an exceptional resource for anyone looking to develop their understanding of stochastic processes.

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.

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.

Modeling, Stochastic Control, Optimization, and Applications

- George Yin 2019-07-16
This volume collects papers, based on invited talks given at the IMA workshop in Modeling, Stochastic Control, Optimization, and Related Applications, held at the Institute for Mathematics and Its Applications, University of Minnesota, during May and June, 2018. There were four week-long workshops during

the conference. They are (1) stochastic control, computation methods, and applications, (2) queueing theory and networked systems, (3) ecological and biological applications, and (4) finance and economics applications. For broader impacts, researchers from different fields covering both theoretically oriented and application intensive areas were invited to participate in the conference. It brought together researchers from multi-disciplinary communities in applied mathematics, applied probability, engineering, biology, ecology, and networked science, to review, and substantially update most recent progress. As an archive, this volume presents some of the highlights of the workshops, and collect papers covering a broad range of topics.

Stochastic Processes for Finance -

Limit Order Books - Frédéric Abergel 2016-05-09

A limit order book is essentially

a file on a computer that contains all orders sent to the market, along with their characteristics such as the sign of the order, price, quantity and a timestamp. The majority of organized electronic markets rely on limit order books to store the list of interests of market participants on their central computer. A limit order book contains all the information available on a specific market and it reflects the way the market moves under the influence of its participants. This book discusses several models of limit order books. It begins by discussing the data to assess their empirical properties, and then moves on to mathematical models in order to reproduce the observed properties. Finally, the book presents a framework for numerical simulations. It also covers important modelling techniques including agent-based modelling, and advanced modelling of limit order books based on Hawkes processes. The book also provides in-depth coverage of simulation

techniques and introduces general, flexible, open source library concepts useful to readers studying trading strategies in order-driven markets.

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.

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.

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 three areas of stochastic models, abstraction techniques and stochastic model checking.

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.

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.

Statistical Machine

Learning with Applications -

Gordon Ritter 2021-07-30

This unique compendium develops a general approach to building models of economic and financial processes, with a focus on statistical learning techniques that scale to large data sets. It introduces the key elements of a parametric statistical model: likelihood, prior, and posterior, and show how to use them to make predictions. The book covers classical techniques such as multiple regression and the Kalman filter in a clear, accessible style that has been popular with students, but also includes detailed treatments of state-of-the-art models, highlighting tree-based methods, support vector machines and kernel methods, deep learning, and reinforcement learning. Theories are supplemented by real-world examples. This reference text is useful for undergraduate, graduate and even PhD students in quantitative finance, and also to practitioners who are facing the reality that data science

and machine learning are disrupting the industry.

Decision Making Under Uncertainty - Mykel J. Kochenderfer 2015-07-24

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.

Markov Decision Processes with Applications to Finance - Nicole Bäuerle 2011-06-08

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).

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).

Analytic Perturbation Theory and Its Applications -

Konstantin E. Avrachenkov
2013-12-11

Mathematical models are often used to describe complex phenomena such as climate change dynamics, stock market fluctuations, and the Internet. These models typically depend on estimated values of key parameters that determine system behavior. Hence it is important to know what happens when these values are changed. The study of single-parameter deviations provides a natural starting point for this analysis in many special settings in the sciences, engineering, and economics. The difference between the actual and nominal values of the perturbation parameter is

small but unknown, and it is important to understand the asymptotic behavior of the system as the perturbation tends to zero. This is particularly true in applications with an apparent discontinuity in the limiting behavior?the so-called singularly perturbed problems. Analytic Perturbation Theory and Its Applications includes a comprehensive treatment of analytic perturbations of matrices, linear operators, and polynomial systems, particularly the singular perturbation of inverses and generalized inverses. It also offers original applications in Markov chains, Markov decision processes, optimization, and applications to Google PageRank? and the Hamiltonian cycle problem as well as input retrieval in linear control systems and a problem section in every chapter to aid in course preparation.

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.

Applied Probability Models with Optimization Applications

- Sheldon M. Ross 1992-01-01
Includes bibliographical references and index.

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

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.

Discrete-Time Markov Chains - George Yin 2005

Focusing on discrete-time-scale Markov chains, the contents of this book are an outgrowth of some of the authors' recent research. The motivation stems from existing and emerging applications in optimization and control of complex hybrid Markovian systems in manufacturing, wireless communication, and financial engineering. Much effort in this book is devoted to designing system models arising from these applications, analyzing them via analytic and probabilistic techniques, and developing feasible computational algorithms so as to reduce the inherent complexity. This book presents results including asymptotic expansions of probability vectors, structural properties of occupation measures, exponential bounds, aggregation and decomposition and associated limit processes, and interface of discrete-time and continuous-time systems. One of the salient features is that it contains a diverse range of applications on filtering, estimation, control,

optimization, and Markov decision processes, and financial engineering. This book will be an important reference for researchers in the areas of applied probability, control theory, operations research, as well as for practitioners who use optimization techniques. Part of the book can also be used in a graduate course of applied probability, stochastic processes, and applications.

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 in Practice - Richard J. Boucherie
2018-07-20

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.

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

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.

Partially Observed Markov Decision Processes - Vikram Krishnamurthy 2016-03-21

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?

Hidden Markov Models - Ramaprasad Bhar 2006-04-18

Markov chains have increasingly become useful way of capturing stochastic nature of many economic and financial variables. Although the hidden Markov processes have been widely employed for some time in many engineering applications e.g. speech recognition, its effectiveness has now been recognized in areas of social science research as well. The main aim of Hidden Markov Models: Applications to Financial Economics is to make such techniques available to more researchers in financial economics. As such we only

cover the necessary theoretical aspects in each chapter while focusing on real life applications using contemporary data mainly from OECD group of countries. The underlying assumption here is that the researchers in financial economics would be familiar with such application although empirical techniques would be more traditional econometrics. Keeping the application level in a more familiar level, we focus on the methodology based on hidden Markov processes. This will, we believe, help the reader to develop more in-depth understanding of the modeling issues thereby benefiting their future research.

Approximate Dynamic Programming - Warren B.

Powell 2007-10-05

A complete and accessible introduction to the real-world applications of approximate dynamic programming With the growing levels of sophistication in modern-day operations, it is vital for practitioners to understand how to approach, model, and

solve complex industrial problems. Approximate Dynamic Programming is a result of the author's decades of experience working in large industrial settings to develop practical and high-quality solutions to problems that involve making decisions in the presence of uncertainty. This groundbreaking book uniquely integrates four distinct disciplines—Markov design processes, mathematical programming, simulation, and statistics—to demonstrate how to successfully model and solve a wide range of real-life problems using the techniques of approximate dynamic programming (ADP). The reader is introduced to the three curses of dimensionality that impact complex problems and is also shown how the post-decision state variable allows for the use of classical algorithmic strategies from operations research to treat complex stochastic optimization problems. Designed as an introduction and assuming no prior training in dynamic programming of

any form, Approximate Dynamic Programming contains dozens of algorithms that are intended to serve as a starting point in the design of practical solutions for real problems. The book provides detailed coverage of implementation challenges including: modeling complex sequential decision processes under uncertainty, identifying robust policies, designing and estimating value function approximations, choosing effective stepsize rules, and resolving convergence issues. With a focus on modeling and algorithms in conjunction with the language of mainstream operations research, artificial intelligence, and control theory, Approximate Dynamic Programming: Models complex, high-dimensional problems in a natural and practical way, which draws on years of industrial projects Introduces and emphasizes the power of estimating a value function around the post-decision state, allowing solution algorithms to be broken down into three fundamental steps: classical

simulation, classical optimization, and classical statistics Presents a thorough discussion of recursive estimation, including fundamental theory and a number of issues that arise in the development of practical algorithms Offers a variety of methods for approximating dynamic programs that have appeared in previous literature, but that have never been presented in the coherent format of a book Motivated by examples from modern-day operations research, Approximate Dynamic Programming is an accessible introduction to dynamic modeling and is also a valuable guide for the development of high-quality solutions to problems that exist in operations research and engineering. The clear and precise presentation of the material makes this an appropriate text for advanced undergraduate and beginning graduate courses, while also serving as a reference for researchers and practitioners. A companion Web site is

available for readers, which includes additional exercises, solutions to exercises, and data sets to reinforce the book's main concepts.

On the process of maintenance decision-making in the offshore operational environment of the oil and gas industry -

Mario Marcondes Machado
2022-07-08

The overall objective of this research is to develop a decision support framework for preventive maintenance program implementations in the offshore operational environment of the oil and gas industry. The author investigates the problems surrounding the maintenance decision-making process, which elements are systematically identified and categorized, including organizational aspects. The interface with operations is considered in order to promote integration between maintenance and production schedules. The research subjacent objectives have been to identify, among major oil and gas offshore

operators, the state-of-practices regarding the maintenance decision-making process and identify, in the literature, the main techniques used for maintenance decisions and optimization in order to propose alternatives for supporting future implementations. Through a Systems Engineering approach, assisted by a literature review, interviews with experts (in Norway and in Brazil), an on-line survey and case studies, the results of this research complement a toolkit for maintenance engineers and managers aiming to facilitate information sharing and interdisciplinary cooperation in the operational environment. Among the results: A concept map for the maintenance decision-making process ontology; A plan for PM program implementations; The suggested cross-sector solution: the minimum equipment list; A Markovian dependability nomogram; and A Markov decision model application.

[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.

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.

**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: Introduction Controlled Markov Chains Basic Optimality Criteria Policy Iteration and Approximation Theorems Overtaking, Bias, and Variance Optimality Sensitive Discount Optimality Blackwell

Optimality Constrained Controlled Markov Chains Applications Zero-Sum Markov Games Bias 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 Games

Key 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 time. This 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 issues. The 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, unbounded. Reviews: "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 Distributed Sensing and Intelligent Systems - Mohamed Elhoseny 2022

This book is the proceeding of the 1st International Conference on Distributed Sensing and Intelligent Systems (ICDSIS2020) which will be held in The National School of Applied Sciences of Agadir, Ibn Zohr University, Agadir, Morocco on February 01-03, 2020. ICDSIS2020 is co-organized by Computer Vision and Intelligent Systems Lab, University of North Texas, USA as a scientific collaboration event with The National School of Applied Sciences of Agadir, Ibn Zohr University. ICDSIS2020 aims to foster students, researchers, academicians and industry persons in the field of Computer and Information Science, Intelligent Systems,

and Electronics and Communication Engineering in general. The volume collects contributions from leading experts around the globe with the latest insights on emerging topics, and includes reviews, surveys, and research chapters covering all aspects of distributed sensing and intelligent systems. The volume is divided into 5 key sections: Distributed Sensing Applications; Intelligent Systems; Advanced theories and algorithms in machine learning and data mining; Artificial intelligence and optimization, and application to Internet of Things (IoT); and Cybersecurity and Secure Distributed Systems. This conference proceeding is an academic book which can be read by students, analysts, policymakers, and regulators interested in Distributed Sensing, Smart Network approaches, Smart Cities, IoT Applications, and Intelligent Applications. It is written in plain and easy language, and describes new concepts when they appear first so that a

reader without prior background of the field finds it readable. The book is primarily intended for research students in sensor networks and IoT applications (including intelligent information systems, and smart sensors applications), academics in higher education institutions including universities and vocational colleges, policy makers and legislators.

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.

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

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.

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.

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.