Markov Decision Processes

Markov Decision Processes
Reinforcement Learning, second edition
Markov Processes for Stochastic Modeling
Competitive Markov Decision Processes
Continuous Average Control of Piecewise Deterministic Markov Processes
Constrained Markov Decision Processes
Markov Decision Processes with Continuous Time Parameter
Markov Decision Processes in Practice
Planning with Markov Decision Processes
Markov Decision Process
Markov Decision Processes
Recent developments in Markov decision processes : based on the proceedings of an Internat. Conference on Markov Decision Processes, org. jointly by the Dep. of Decision Theory of the Univ. of Manchester ..., held at the Univ. of Manchester on July 17-19, 1978
Operations Research and Health Care
Constrained Markov Decision Processes
Reinforcement Learning
Markovian Decision Processes
Continuous-Time Markov Decision Processes
Markov Decision Processes with Their Applications
Examples in Markov Decision Processes
Decision Processes in Dynamic Probabilistic Systems
Markov Decision Processes and the Belief-Desire-Intention Model
Markov Decision Processes
Applied Probability Models with Optimization Applications
Markov Decision Processes with Policy Constraints
Continuous-Time Markov Decision Processes
Markov Chains and Decision Processes for Engineers and Managers
Markov Decision Processes in Artificial Intelligence
Partially Observed Markov Decision Processes
Continuous-Time Markov Decision Processes
Recent Developments in Markov Decision Processes Based on the Proceedings of an International Conference on Markov Decision Pr an
Simulation-based Algorithms for Markov Decision Processes
Markov Decision Processes with Applications to Finance
Partially Observed Markov Decision Processes
Handbook of Markov Decision Processes
Learning Representation and Control in Markov Decision Processes
Decision Making Under Uncertainty
Markov Decision Processes
Recent Developments in Markov Decision Processes
Markov Processes and Controlled Markov Chains

JAX HOLLAND

Markov Decision Processes Springer
'Et moi - ... - si j'avait su comment en revenir. One service mathematics has rendered the je n'y serais point aile: human race. It has put common sense back where it belongs. on the topmost shelf next Jules Verne (0 the dusty canister labelled 'discarded non sense'. The series is divergent; therefore we may be able to do something with it. Eric T. Bell O. Heaviside Mathematics is a tool for thought. A highly necessary tool in a world where both feedback and non linearities abound. Similarly, all kinds of parts of mathematics serve as tools for other parts and for other sciences. Applying a simple rewriting rule to the quote on the right above one finds such statements as: 'One service topology has rendered mathematical physics .. .'; 'One service logic has rendered com puter science .. .'; 'One service category theory has rendered mathematics .. .'. All arguably true. And all statements obtainable this way form part of the raison d'etre of this series.
Reinforcement Learning, second edition MIT Press
The intent of this book is to present recent results in the control theory for the long run average continuous control problem of piecewise deterministic Markov processes (PDMPs). The book focuses mainly on the long run average cost criteria and extends to the PDMPs some well-known techniques related to discrete-time and continuous-time Markov decision processes, including the so-called ``average inequality approach'', ``vanishing discount technique'' and ``policy iteration algorithm''. We believe that what is unique about our approach is that, by using the special features of the PDMPs, we trace a parallel with the general theory for discrete-time Markov Decision Processes rather than the continuous-time case. The two main reasons for doing that is to use the powerful tools developed in the discrete-time framework and to avoid working with the infinitesimal generator associated to a PDMP, which in most cases has its domain of definition difficult to be characterized. Although the book is mainly intended to be a theoretically oriented text, it also contains some motivational examples. The book is targeted primarily for advanced students and practitioners of control theory. The book will be a valuable source for experts in the field of Markov decision processes. Moreover, the book should be suitable for certain advanced courses or seminars. As background, one needs an acquaintance with the theory of Markov decision processes and some knowledge of stochastic processes and modern analysis.
Markov Processes for Stochastic Modeling Newnes
What Is Markov Decision Process A discrete-time stochastic control process is referred to as a Markov decision process (MDP) in the field of mathematics. It offers a mathematical framework for modeling decision making in scenarios in which the outcomes are partially controlled by a decision maker and partly determined by random chance. The study of optimization issues that can be handled by dynamic programming lends itself well to the use of MDPs. At the very least, MDPs were

recognized to exist in the 1950s. Ronald Howard's book, published in 1960 and titled Dynamic Programming and Markov Processes, is credited for initiating a core body of study on Markov decision processes. They have applications in a wide variety of fields, including as robotics, automatic control, economics, and manufacturing, among others. Because Markov decision processes are an extension of Markov chains, the Russian mathematician Andrey Markov is where the term "Markov decision processes" (MDPs) originated. How You Will Benefit (I) Insights, and validations about the following topics: Chapter 1: Markov decision process Chapter 2: Markov chain Chapter 3: Reinforcement learning Chapter 4: Bellman equation Chapter 5: Admissible decision rule Chapter 6: Partially observable Markov decision process Chapter 7: Temporal difference learning Chapter 8: Multi-armed bandit Chapter 9: Optimal stopping Chapter 10: Metropolis-Hastings algorithm (II) Answering the public top questions about markov decision process. (III) Real world examples for the usage of markov decision process in many fields. (IV) 17 appendices to explain, briefly, 266 emerging technologies in each industry to have 360-degree full understanding of markov decision process' technologies. Who This Book Is For Professionals, undergraduate and graduate students, enthusiasts, hobbyists, and those who want to go beyond basic knowledge or information for any kind of markov decision process. What is Artificial Intelligence Series The artificial intelligence book series provides comprehensive coverage in over 200 topics. Each ebook covers a specific Artificial Intelligence topic in depth, written by experts in the field. The series aims to give readers a thorough understanding of the concepts, techniques, history and applications of artificial intelligence. Topics covered include machine learning, deep learning, neural networks, computer vision, natural language processing, robotics, ethics and more. The ebooks are written for professionals, students, and anyone interested in learning about the latest developments in this rapidly advancing field. The artificial intelligence book series provides an in-depth yet accessible exploration, from the fundamental concepts to the state-of-the-art research. With over 200 volumes, readers gain a thorough grounding in all aspects of Artificial Intelligence. The ebooks are designed to build knowledge systematically, with later volumes building on the foundations laid by earlier ones. This comprehensive series is an indispensable resource for anyone seeking to develop expertise in artificial intelligence.

Competitive Markov Decision Processes World Scientific

Markov Decision ProcessesJohn Wiley & Sons

Continuous Average Control of Piecewise Deterministic Markov Processes Morgan & Claypool Publishers

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.

Constrained Markov Decision Processes Springer Science & Business Media

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.

Markov Decision Processes with Continuous Time Parameter CRC Press

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 in Practice CRC Press

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 re spective area. The papers cover major research areas and methodologies, and discuss open questions and future research directions. The papers can be read independently, with the basic notation and concepts ofSection 1.2. Most chap ters should be accessible by graduate or advanced undergraduate students in fields of operations research, electrical engineering, and computer science. 1.1 AN OVERVIEW OF MARKOV DECISION PROCESSES The theory of Markov Decision Processes-also known under several other names including sequential stochastic optimization, discrete-time stochastic control, and stochastic dynamic programming-studiessequential optimization ofdiscrete time stochastic systems. The basic object is a discrete-

time stochas tic system whose transition mechanism can be controlled over time. Each control policy defines the stochastic process and values of objective functions associated with this process. The goal is to select a "good" control policy. In real life, decisions that humans and computers make on all levels usually have two types ofimpacts: (i) they cost orsavetime, money, or other resources, or they bring revenues, as well as (ii) they have an impact on the future, by influencing the dynamics. In many situations, decisions with the largest immediate profit may not be good in view offuture events. MDPs model this paradigm and provide results on the structure and existence of good policies and on methods for their calculation.

Planning with Markov Decision Processes Springer

Which customers cant participate in our Partially observable Markov decision process domain because they lack skills, wealth, or convenient access to existing solutions? Can we add value to the current Partially observable Markov decision process decision-making process (largely qualitative) by incorporating uncertainty modeling (more quantitative)? Who are the people involved in developing and implementing Partially observable Markov decision process? How does Partially observable Markov decision process integrate with other business initiatives? Does the Partially observable Markov decision process performance meet the customer's requirements? This premium Partially observable Markov decision process self-assessment will make you the assured Partially observable Markov decision process domain master by revealing just what you need to know to be fluent and ready for any Partially observable Markov decision process challenge. How do I reduce the effort in the Partially observable Markov decision process work to be done to get problems solved? How can I ensure that plans of action include every Partially observable Markov decision process task and that every Partially observable Markov decision process outcome is in place? How will I save time investigating strategic and tactical options and ensuring Partially observable Markov decision process costs are low? How can I deliver tailored Partially observable Markov decision process advice instantly with structured going-forward plans? There's no better guide through these mind-expanding questions than acclaimed best-selling author Gerard Blokdyk. Blokdyk ensures all Partially observable Markov decision process essentials are covered, from every angle: the Partially observable Markov decision process self-assessment shows succinctly and clearly that what needs to be clarified to organize the required activities and processes so that Partially observable Markov decision process outcomes are achieved. Contains extensive criteria grounded in past and current successful projects and activities by experienced Partially observable Markov decision process practitioners. Their mastery, combined with the easy elegance of the self-assessment, provides its superior value to you in knowing how to ensure the outcome of any efforts in Partially observable Markov decision process are maximized with professional results. Your purchase includes access details to the Partially observable Markov decision process self-assessment dashboard download which gives you your dynamically prioritized projects-ready tool and shows you exactly what to do next. Your exclusive

instant access details can be found in your book.

Markov Decision Process Springer Science & Business Media

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

Markov Decision Processes C O M A P, Incorporated

This book provides a unified approach for the study of constrained Markov decision processes with a finite state space and unbounded costs. Unlike the single controller case considered in many other books, the author considers a single controller with several objectives, such as minimizing delays and loss, probabilities, and maximization of throughputs. It is desirable to design a controller that minimizes one cost objective, subject to inequality constraints on other cost objectives. This framework describes dynamic decision problems arising frequently in many engineering fields. A thorough overview of these applications is presented in the introduction. The book is then divided into three sections that build upon each other. The first part explains the theory for the finite state space. The author characterizes the set of achievable expected occupation measures as well as performance vectors, and identifies simple classes of policies among which optimal policies exist. This allows the reduction of the original dynamic into a linear program. A Lagranian approach is then used to derive the dual linear program using dynamic programming techniques. In the second part, these results are extended to the infinite state space and action spaces. The author provides two frameworks: the case where costs are bounded below and the contracting framework. The third part

builds upon the results of the first two parts and examines asymptotical results of the convergence of both the value and the policies in the time horizon and in the discount factor. Finally, several state truncation algorithms that enable the approximation of the solution of the original control problem via finite linear programs are given.

Recent developments in Markov decision processes : based on the proceedings of an Internat. Conference on Markov Decision Processes, org. jointly by the Dep. of Decision Theory of the Univ. of Manchester ..., held at the Univ. of Manchester on July 17-19, 1978 Springer Science & Business Media

Provides a comprehensive survey of techniques to automatically construct basis functions or features for value function approximation in Markov decision processes and reinforcement learning.

Operations Research and Health Care Springer Nature

In this work, we provide a treatment of the relationship between two models that have been widely used in the implementation of autonomous agents: the Belief DesireIntention (BDI) model and Markov Decision Processes (MDPs). We start with an informal description of the relationship, identifying the common features of the two approaches and the differences between them. Then we hone our understanding of these differences through an empirical analysis of the performance of both models on the TileWorld testbed. This allows us to show that even though the MDP model displays consistently better behavior than the BDI model for small worlds, this is not the case when the world becomes large and the MDP model cannot be solved exactly. Finally we present a theoretical analysis of the relationship between the two approaches, identifying mappings that allow us to extract a set of intentions from a policy (a solution to an MDP), and to extract a policy from a set of intentions.

Constrained Markov Decision Processes MIT Press

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.

Reinforcement Learning Courier Corporation

The general theory of stochastic processes and the more specialized theory of Markov processes evolved enormously in the second half of the last century. In parallel, the theory of controlled Markov chains (or Markov decision processes) was being pioneered by control engineers and operations researchers. Researchers in Markov processes and controlled Markov chains have been, for a long time, aware of the synergies between these two subject areas. However, this may be the first volume dedicated to highlighting these synergies and, almost certainly, it is the

first volume that emphasizes the contributions of the vibrant and growing Chinese school of probability. The chapters that appear in this book reflect both the maturity and the vitality of modern day Markov processes and controlled Markov chains. They also will provide an opportunity to trace the connections that have emerged between the work done by members of the Chinese school of probability and the work done by the European, US, Central and South American and Asian scholars.

Markovian Decision Processes Elsevier Publishing Company

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 Springer Science & Business Media

Examines several fundamentals concerning the manner in which Markov decision problems may be properly formulated and the determination of solutions or their properties. Coverage includes optimal equations, algorithms and their characteristics, probability distributions, modern development in the Markov decision process area, namely structural policy analysis, approximation modeling, multiple objectives and Markov games. Copiously illustrated with examples.

Markov Decision Processes with Their Applications Createspace Independent Publishing Platform

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.

Examples in Markov Decision Processes Springer Science & Business Media

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

Decision Processes in Dynamic Probabilistic Systems Springer Science & Business Media

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.