

Dynamic Programming Optimal Control Vol

As is well known, Pontryagin's maximum principle and Bellman's dynamic programming are the two principal and most commonly used approaches in solving stochastic optimal control problems. * An interesting phenomenon one can observe from the literature is that these two approaches have been developed separately and independently. Since both methods are used to investigate the same problems, a natural question one will ask is the following: (Q) What is the relationship between the maximum principle and dynamic programming in stochastic optimal controls? There did exist some researches (prior to the 1980s) on the relationship between these two. Nevertheless, the results usually were stated in heuristic terms and proved under rather restrictive assumptions, which were not satisfied in most cases. In the statement of a Pontryagin-type maximum principle there is an adjoint equation, which is an ordinary differential equation (ODE) in the (finite-dimensional) deterministic case and a stochastic differential equation (SDE) in the stochastic case. The system consisting of the adjoint equation, the original state equation, and the maximum condition is referred to as an (extended) Hamiltonian system. On the other hand, in Bellman's dynamic programming, there is a partial differential equation (PDE), of first order in the (finite-dimensional) deterministic case and of second order in the stochastic case. This is known as a Hamilton-Jacobi-Bellman (HJB) equation.

Much debate has centered around the decreasing mathematical ability of students entering higher education, as well as the discrepancy between skills found in the UK and Europe in mathematics. This collection of articles from leading researchers and teachers considers solutions to this problem, with suggestions outlined for new methods of teaching the subject. Topics include the application of mathematics to engineering careers; the problems of wider access to higher education and current practices that are helping to tackle them; teaching experience from varying educational establishments; and computer-based teaching and assessment. The discussions presented here should be read by anyone involved in mathematics, education, and engineering

This book considers large and challenging multistage decision problems, which can be solved in principle by dynamic programming (DP), but their exact solution is computationally intractable. We discuss solution methods that rely on approximations to produce suboptimal policies with adequate performance. These methods are collectively known by several essentially equivalent names: reinforcement learning, approximate dynamic programming, neuro-dynamic programming. They have been at the forefront of research for the last 25 years, and they underlie, among others, the recent impressive successes of self-learning in the context of games

such as chess and Go. Our subject has benefited greatly from the interplay of ideas from optimal control and from artificial intelligence, as it relates to reinforcement learning and simulation-based neural network methods. One of the aims of the book is to explore the common boundary between these two fields and to form a bridge that is accessible by workers with background in either field. Another aim is to organize coherently the broad mosaic of methods that have proved successful in practice while having a solid theoretical and/or logical foundation. This may help researchers and practitioners to find their way through the maze of competing ideas that constitute the current state of the art. This book relates to several of our other books: *Neuro-Dynamic Programming* (Athena Scientific, 1996), *Dynamic Programming and Optimal Control* (4th edition, Athena Scientific, 2017), *Abstract Dynamic Programming* (2nd edition, Athena Scientific, 2018), and *Nonlinear Programming* (Athena Scientific, 2016). However, the mathematical style of this book is somewhat different. While we provide a rigorous, albeit short, mathematical account of the theory of finite and infinite horizon dynamic programming, and some fundamental approximation methods, we rely more on intuitive explanations and less on proof-based insights. Moreover, our mathematical requirements are quite modest: calculus, a minimal use of matrix-vector algebra, and elementary probability (mathematically complicated arguments involving laws of large numbers and stochastic convergence are bypassed in favor of intuitive explanations). The book illustrates the methodology with many examples and illustrations, and uses a gradual expository approach, which proceeds along four directions: (a) From exact DP to approximate DP: We first discuss exact DP algorithms, explain why they may be difficult to implement, and then use them as the basis for approximations. (b) From finite horizon to infinite horizon problems: We first discuss finite horizon exact and approximate DP methodologies, which are intuitive and mathematically simple, and then progress to infinite horizon problems. (c) From deterministic to stochastic models: We often discuss separately deterministic and stochastic problems, since deterministic problems are simpler and offer special advantages for some of our methods. (d) From model-based to model-free implementations: We first discuss model-based implementations, and then we identify schemes that can be appropriately modified to work with a simulator. The book is related and supplemented by the companion research monograph *Rollout, Policy Iteration, and Distributed Reinforcement Learning* (Athena Scientific, 2020), which focuses more closely on several topics related to rollout, approximate policy iteration, multiagent problems, discrete and Bayesian optimization, and distributed computation, which are either discussed in less detail or not covered at all in the present book. The author's website contains class notes, and a series of videolectures and slides from a 2021 course at ASU,

which address a selection of topics from both books.

Providing an introduction to stochastic optimal control in infinite dimension, this book gives a complete account of the theory of second-order HJB equations in infinite-dimensional Hilbert spaces, focusing on its applicability to associated stochastic optimal control problems. It features a general introduction to optimal stochastic control, including basic results (e.g. the dynamic programming principle) with proofs, and provides examples of applications. A complete and up-to-date exposition of the existing theory of viscosity solutions and regular solutions of second-order HJB equations in Hilbert spaces is given, together with an extensive survey of other methods, with a full bibliography. In particular, Chapter 6, written by M. Fuhrman and G. Tessitore, surveys the theory of regular solutions of HJB equations arising in infinite-dimensional stochastic control, via BSDEs. The book is of interest to both pure and applied researchers working in the control theory of stochastic PDEs, and in PDEs in infinite dimension. Readers from other fields who want to learn the basic theory will also find it useful. The prerequisites are: standard functional analysis, the theory of semigroups of operators and its use in the study of PDEs, some knowledge of the dynamic programming approach to stochastic optimal control problems in finite dimension, and the basics of stochastic analysis and stochastic equations in infinite-dimensional spaces.

Rollout, Policy Iteration, and Distributed Reinforcement Learning
sgfrgds

Dynamic programming and optimal control, vol. 2

Optimization, Estimation and Control

Dynamic Programming and Its Applications

Reinforcement Learning, second edition

Stochastic optimization problems arise in decision-making problems under uncertainty, and find various applications in economics and finance. On the other hand, problems in finance have recently led to new developments in the theory of stochastic control. This volume provides a systematic treatment of stochastic optimization problems applied to finance by presenting the different existing methods: dynamic programming, viscosity solutions, backward stochastic differential equations, and martingale duality methods. The theory is discussed in the context of recent developments in this field, with complete and detailed proofs, and is illustrated by means of concrete examples from the world of finance: portfolio allocation, option hedging, real options, optimal investment, etc. This book is directed towards graduate students and researchers in mathematical finance, and will also benefit applied mathematicians interested in financial applications and practitioners wishing to know more about the use of stochastic optimization methods in finance.

Reinforcement learning (RL) and adaptive dynamic programming (ADP) has been one of the most critical research fields in science and engineering for modern complex systems. This book describes the latest RL and ADP techniques for decision and control in human engineered systems, covering both single player decision and control and multi-player games. Edited by the pioneers of RL and ADP research, the book brings together ideas and methods from many fields and provides an important and timely guidance on controlling a wide variety of systems, such as robots, industrial processes, and economic decision-making.

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.

Linear Network Optimization presents a thorough treatment of classical approaches to network problems such as shortest path, max-flow, assignment, transportation, and minimum cost flow problems.

The Calculus of Variations and Optimal Control in Economics and Management
Volume II; Approximate Dynamic Programming

Solving the Curses of Dimensionality
Continuous-time Stochastic Control and Optimization with Financial Applications
An Introduction
Reinforcement Learning and Stochastic Optimization

Automotive control has developed over the decades from an auxiliary technology to a key element without which the actual performances, emission, safety and consumption targets could not be met. Accordingly, automotive control has been increasing its authority and responsibility – at the price of complexity and difficult tuning. The progressive evolution has been mainly led by specific applications and short term targets, with the consequence that automotive control is to a very large extent more heuristic than systematic. Product requirements are still increasing and new challenges are coming from potentially huge markets like India and China, and against this background there is wide consensus both in the industry and academia that the current state is not satisfactory. Model-based control could be an approach to improve performance while reducing development and tuning times and possibly costs. Model predictive control is a kind of model-based control design approach which has experienced a growing success since the middle of the 1980s for “slow” complex plants, in particular of the chemical and process industry. In the last decades, several developments have allowed using these methods also for

“fast” systems and this has supported a growing interest in its use also for automotive applications, with several promising results reported. Still there is no consensus on whether model predictive control with its high requirements on model quality and on computational power is a sensible choice for automotive control.

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.

The Art and Theory of Dynamic Programming

This handbook presents state-of-the-art research in reinforcement learning, focusing on its applications in the control and game theory of dynamic systems and future directions for related research and technology. The contributions gathered in this book deal with challenges faced when using learning and adaptation methods to solve academic and industrial problems, such as optimization in dynamic environments with single and multiple agents, convergence and performance analysis, and online implementation. They explore means by which these difficulties can be solved, and cover a wide range of related topics including: deep learning; artificial intelligence; applications of game theory; mixed modality learning; and multi-agent reinforcement learning.

Practicing engineers and scholars in the field of machine learning, game theory, and autonomous control will find the Handbook of Reinforcement Learning and Control to be thought-provoking, instructive and informative.

Optimal Control: Novel Directions and Applications

Mathematical Education of Engineers

Automotive Model Predictive Control

Dynamic Programming and HJB Equations

Reinforcement Learning and Dynamic Programming Using Function Approximators

Dynamic Programming and Optimal Control

This is the leading and most up-to-date textbook on the far-ranging algorithmic methodology of Dynamic Programming, which can be used for optimal control, Markovian decision problems, planning and sequential decision making under uncertainty, and discrete/combinatorial optimization. The treatment focuses on basic unifying themes, and conceptual foundations. It illustrates the versatility, power, and generality of the method with many examples and applications from engineering, operations research, and other fields. It also addresses extensively the practical application of the

methodology, possibly through the use of approximations, and provides an extensive treatment of the far-reaching methodology of Neuro-Dynamic Programming/Reinforcement Learning. Among its special features, the book 1) provides a unifying framework for sequential decision making, 2) treats simultaneously deterministic and stochastic control problems popular in modern control theory and Markovian decision popular in operations research, 3) develops the theory of deterministic optimal control problems including the Pontryagin Minimum Principle, 4) introduces recent suboptimal control and simulation-based approximation techniques (neuro-dynamic programming), which allow the practical application of dynamic programming to complex problems that involve the dual curse of large dimension and lack of an accurate mathematical model, 5) provides a comprehensive treatment of infinite horizon problems in the second volume, and an introductory treatment in the first volume. The electronic version of the book includes 29 theoretical problems, with high-quality solutions, which enhance the range of coverage of the book.

Dynamic programming is a powerful method for solving optimization problems, but has a number of drawbacks that limit its use to solving problems of very low dimension. To overcome these limitations, author Rein Luus suggested using it in an iterative fashion. Although this method required vast computer resources, modifications to his original scheme. This book offers a systematic introduction to the optimal stochastic control theory via the dynamic programming principle, which is a powerful tool to analyze control problems. First we consider completely observable control problems with finite horizons. Using a time discretization we construct a nonlinear semigroup related to the dynamic programming principle (DPP), whose generator provides the Hamilton-Jacobi-Bellman (HJB) equation, and we characterize the value function via the nonlinear semigroup, besides the viscosity solution theory. When we control not only the dynamics of a system but also the terminal time of its evolution, control-stopping problems arise. This problem is treated in the same frameworks, via the nonlinear semigroup. Its results are applicable to the American option price problem. Zero-sum two-player time-homogeneous stochastic differential games and viscosity solutions of the Isaacs equations arising from such games are studied via a

nonlinear semigroup related to DPP (the min-max principle, to be precise). Using semi-discretization arguments, we construct the nonlinear semigroups whose generators provide lower and upper Isaacs equations. Concerning partially observable control problems, we refer to stochastic parabolic equations driven by colored Wiener noises, in particular, the Zakai equation. The existence and uniqueness of solutions and regularities as well as Itô's formula are stated. A control problem for the Zakai equations has a nonlinear semigroup whose generator provides the HJB equation on a Banach space. The value function turns out to be a unique viscosity solution for the HJB equation under mild conditions. This edition provides a more generalized treatment of the topic than does the earlier book Lectures on Stochastic Control Theory (ISI Lecture Notes 9), where time-homogeneous cases are dealt with. Here, for finite time-horizon control problems, DPP was formulated as a one-parameter nonlinear semigroup, whose generator provides the HJB equation, by using a time-discretization method. The semigroup corresponds to the value function and is characterized as the envelope of Markovian transition semigroups of responses for constant control processes. Besides finite time-horizon controls, the book discusses control-stopping problems in the same frameworks.

asfbgsdfg

Reinforcement Learning and Approximate Dynamic Programming for Feedback Control

Dynamic Programming Principle

Proceedings of the International Conference on Dynamic Programming and Its Applications, University of British Columbia, Vancouver, British Columbia, Canada, April 14-16, 1977

Linear Network Optimization

Reinforcement Learning and Optimal Control

Abstract Dynamic Programming

The purpose of this book is to develop in greater depth some of the methods from the author's Reinforcement Learning and Optimal Control recently published textbook (Athena Scientific, 2019). In particular, we present new research, relating to systems involving multiple agents, partitioned architectures, and distributed asynchronous computation. We pay special attention to the contexts of dynamic programming/policy iteration and control theory/model predictive control. We also discuss in some

detail the application of the methodology to challenging discrete/combinatorial optimization problems, such as routing, scheduling, assignment, and mixed integer programming, including the use of neural network approximations within these contexts. The book focuses on the fundamental idea of policy iteration, i.e., start from some policy, and successively generate one or more improved policies. If just one improved policy is generated, this is called rollout, which, based on broad and consistent computational experience, appears to be one of the most versatile and reliable of all reinforcement learning methods. In this book, rollout algorithms are developed for both discrete deterministic and stochastic DP problems, and the development of distributed implementations in both multiagent and multiprocessor settings, aiming to take advantage of parallelism. Approximate policy iteration is more ambitious than rollout, but it is a strictly off-line method, and it is generally far more computationally intensive. This motivates the use of parallel and distributed computation. One of the purposes of the monograph is to discuss distributed (possibly asynchronous) methods that relate to rollout and policy iteration, both in the context of an exact and an approximate implementation involving neural networks or other approximation architectures. Much of the new research is inspired by the remarkable AlphaZero chess program, where policy iteration, value and policy networks, approximate lookahead minimization, and parallel computation all play an important role.

A complete resource to Approximate Dynamic Programming (ADP), including on-line simulation code Provides a tutorial that readers can use to start implementing the learning algorithms provided in the book Includes ideas, directions, and recent results on current research issues and addresses applications where ADP has been successfully implemented The contributors are leading researchers in the field

This is the 3rd edition of a research monograph providing a synthesis of old research on the foundations of dynamic programming (DP), with the modern theory of approximate DP and new research on semicontractive models. It aims at a unified and economical development of the core theory and algorithms of total cost sequential decision problems, based on the strong connections of the subject with fixed point theory. The analysis focuses on the abstract mapping that underlies DP and defines the mathematical character of the associated problem. The discussion centers on two fundamental properties that this mapping may have: monotonicity and (weighted sup-norm) contraction. It turns out that the nature of the analytical and algorithmic DP theory is determined primarily by the presence or

absence of these two properties, and the rest of the problem's structure is largely inconsequential. New research is focused on two areas: 1) The ramifications of these properties in the context of algorithms for approximate DP, and 2) The new class of semicontractive models, exemplified by stochastic shortest path problems, where some but not all policies are contractive. The 3rd edition is very similar to the 2nd edition, except for the addition of a new chapter (Chapter 5), which deals with abstract DP models for sequential minimax problems and zero-sum games. The book is an excellent supplement to several of our books: *Neuro-Dynamic Programming* (Athena Scientific, 1996), *Dynamic Programming and Optimal Control* (Athena Scientific, 2017), *Reinforcement Learning and Optimal Control* (Athena Scientific, 2019), and *Rollout, Policy Iteration, and Distributed Reinforcement Learning* (Athena Scientific, 2020). This book may be regarded as consisting of two parts. In Chapters I-IV we present what we regard as essential topics in an introduction to deterministic optimal control theory. This material has been used by the authors for one semester graduate-level courses at Brown University and the University of Kentucky. The simplest problem in calculus of variations is taken as the point of departure, in Chapter I. Chapters II, III, and IV deal with necessary conditions for an optimum, existence and regularity theorems for optimal controls, and the method of dynamic programming. The beginning reader may find it useful first to learn the main results, corollaries, and examples. These tend to be found in the earlier parts of each chapter. We have deliberately postponed some difficult technical proofs to later parts of these chapters. In the second part of the book we give an introduction to stochastic optimal control for Markov diffusion processes. Our treatment follows the dynamic programming method, and depends on the intimate relationship between second order partial differential equations of parabolic type and stochastic differential equations. This relationship is reviewed in Chapter V, which may be read independently of Chapters I-IV. Chapter VI is based to a considerable extent on the authors' work in stochastic control since 1961. It also includes two other topics important for applications, namely, the solution to the stochastic linear regulator and the separation principle.

Stochastic Optimal Control in Infinite Dimension

Handbook of Reinforcement Learning and Control

Approximate Dynamic Programming

optimization and computation series

Optimal Control

The Art and Theory of Dynamic Programming

This comprehensive study of dynamic programming applied to numerical solution of optimization problems. It will interest aerodynamic, control, and industrial engineers, numerical analysts, and computer specialists, applied mathematicians, economists, and operations and systems analysts. Originally published in 1962. The Princeton Legacy Library uses the latest print-on-demand technology to again make available previously out-of-print books from the distinguished backlist of Princeton University Press. These editions preserve the original texts of these important books while presenting them in durable paperback and hardcover editions. The goal of the Princeton Legacy Library is to vastly increase access to the rich scholarly heritage found in the thousands of books published by Princeton University Press since its founding in 1905.

Since its initial publication, this text has defined courses in dynamic optimization taught to economics and management science students. The two-part treatment covers the calculus of variations and optimal control. 1998 edition.

This book covers the most recent developments in adaptive dynamic programming (ADP). The text begins with a thorough background review of ADP making sure that readers are sufficiently familiar with the fundamentals. In the core of the book, the authors address first discrete- and then continuous-time systems. Coverage of discrete-time systems starts with a more general form of value iteration to demonstrate its convergence, optimality, and stability with complete and thorough theoretical analysis. A more realistic form of value iteration is studied where value function approximations are assumed to have finite errors. Adaptive Dynamic Programming also details another avenue of the ADP approach: policy iteration. Both basic and generalized forms of policy-iteration-based ADP are studied with complete and thorough theoretical analysis in terms of convergence, optimality, stability, and error bounds. Among continuous-time systems, the control of affine and nonaffine nonlinear systems is studied using the ADP approach which is then extended to other branches of control theory including decentralized control, robust and guaranteed cost control, and game theory. In the last part of the book the real-world significance of ADP theory is presented, focusing on three application examples developed from the authors' work: • renewable energy scheduling for smart power grids; • coal gasification processes; and • water–gas shift reactions. Researchers studying intelligent control methods and practitioners looking to apply them in the chemical-process and power-supply industries will find much to interest them in this thorough treatment of an advanced approach to control.

Focusing on applications to science and engineering, this book presents the results of the ITN-FP7 SADCO network's innovative research in optimization and control in the following interconnected topics: optimality conditions in optimal control, dynamic programming approaches to optimal feedback synthesis and reachability analysis, and computational developments in model predictive control. The novelty of the book resides in the fact that it has been developed by early career researchers, providing a good balance between clarity and scientific rigor. Each chapter features an introduction addressed to PhD students and some original contributions aimed at specialist researchers. Requiring only a graduate mathematical background, the book is self-contained. It will be of particular interest to graduate and advanced undergraduate students, industrial practitioners and to senior scientists wishing to update their knowledge.

Stochastic Optimal Control

Models, Methods and Applications

Hamiltonian Systems and HJB Equations

Based on the Proceedings of a Conference on the Mathematical Education of Engineers, Organized by the Institute of Mathematics and Its Applications and Held at Loughborough University of Technology in March 1994

Calculus of Variations and Optimal Control Theory

Algorithms and Codes

A NEW EDITION OF THE CLASSIC TEXT ON OPTIMAL CONTROL THEORY As a superb introductory text and an indispensable reference, this new edition of Optimal Control will serve the needs of both the professional engineer and the advanced student in mechanical, electrical, and aerospace

engineering. Its coverage encompasses all the fundamental topics as well as the major changes that have occurred in recent years. An abundance of computer simulations using MATLAB and relevant Toolboxes is included to give the reader the actual experience of applying the theory to real-world situations. Major topics covered include: Static Optimization Optimal Control of Discrete-Time Systems Optimal Control of Continuous-Time Systems The Tracking Problem and Other LQR Extensions Final-Time-Free and Constrained Input Control Dynamic Programming Optimal Control for Polynomial Systems Output Feedback and Structured Control Robustness and Multivariable Frequency-Domain Techniques Differential Games Reinforcement Learning and Optimal Adaptive Control

This book presents a class of novel optimal control methods and games schemes based on adaptive dynamic programming techniques. For systems with one control input, the ADP-based optimal control is designed for different objectives, while for systems with multi-players, the optimal control inputs are proposed based on games. In order to verify the effectiveness of the proposed methods, the book analyzes the properties of the adaptive dynamic programming methods, including convergence of the iterative value functions and the stability of the system under the iterative control laws. Further, to substantiate the mathematical analysis, it presents various application examples, which provide reference to real-world practices.

Dynamic Programming and Optimal Control Volume II; Approximate Dynamic Programming Athena Scientific

"The leading and most up-to-date textbook on the far-ranging algorithmic methodology of Dynamic Programming, which can be used for optimal control, Markovian decision problems, planning and sequential decision making under uncertainty, and discrete/combinatorial optimization. The treatment focuses on basic unifying themes, and conceptual foundations. It illustrates the versatility, power, and generality of the method with many examples and applications from engineering, operations research, and other fields. It also addresses extensively the practical application of the methodology, possibly through the use of approximations, and provides an extensive treatment of the far-reaching methodology of Neuro-Dynamic Programming/Reinforcement Learning. The first volume is oriented towards modeling, conceptualization, and finite-horizon problems, but also includes a substantive introduction to infinite horizon problems that is suitable for classroom use. The second volume is oriented towards mathematical analysis and computation, treats infinite horizon problems extensively, and provides an up-to-date account of approximate large-scale dynamic programming and reinforcement learning. The text contains many illustrations, worked-out examples, and exercises."--Publisher's website.

Volume I

Dynamic Programming and Its Application to Optimal Control

3rd Edition

Stochastic Control of Partially Observable Systems

A Unified Framework for Sequential Decisions

This book collects some recent developments in stochastic control theory with applications to financial mathematics. We first address standard stochastic control problems from the viewpoint of the recently developed weak dynamic programming principle. A special emphasis is put on the regularity issues and, in particular, on the behavior of the value function near the boundary. We then provide a quick review of the main tools from viscosity solutions which allow to overcome all regularity problems. We next address the class of stochastic target problems which extends in a nontrivial way the standard stochastic control problems. Here the theory of viscosity solutions plays a crucial role in the derivation of the dynamic programming equation as the infinitesimal counterpart of the corresponding geometric dynamic programming equation. The various developments of this theory have been stimulated by applications in finance and by relevant connections with geometric flows. Namely, the second order extension was motivated by illiquidity

modeling, and the controlled loss version was introduced following the problem of quantile hedging. The third part specializes to an overview of Backward stochastic differential equations, and their extensions to the quadratic case.

From household appliances to applications in robotics, engineered systems involving complex dynamics can only be as effective as the algorithms that control them. While Dynamic Programming (DP) has provided researchers with a way to optimally solve decision and control problems involving complex dynamic systems, its practical value was limited by algorithms that lacked the capacity to scale up to realistic problems. However, in recent years, dramatic developments in Reinforcement Learning (RL), the model-free counterpart of DP, changed our understanding of what is possible. Those developments led to the creation of reliable methods that can be applied even when a mathematical model of the system is unavailable, allowing researchers to solve challenging control problems in engineering, as well as in a variety of other disciplines, including economics, medicine, and artificial intelligence. Reinforcement Learning and Dynamic Programming Using Function Approximators provides a comprehensive and unparalleled exploration of the field of RL and DP. With a focus on continuous-variable problems, this seminal text details essential developments that have substantially altered the field over the past decade. In its pages, pioneering experts provide a concise introduction to classical RL and DP, followed by an extensive presentation of the state-of-the-art and novel methods in RL and DP with approximation. Combining algorithm development with theoretical guarantees, they elaborate on their work with illustrative examples and insightful comparisons. Three individual chapters are dedicated to representative algorithms from each of the major classes of techniques: value iteration, policy iteration, and policy search. The features and performance of these algorithms are highlighted in extensive experimental studies on a range of control applications. The recent development of applications involving complex systems has led to a surge of interest in RL and DP methods and the subsequent need for a quality resource on the subject. For graduate students and others new to the field, this book offers a thorough introduction to both the basics and emerging methods. And for those researchers and practitioners working in the fields of optimal and adaptive control, machine learning, artificial intelligence, and operations research, this resource offers a combination of practical algorithms, theoretical analysis, and comprehensive examples that they will be able to adapt and apply to their own work. Access the authors' website at www.dsc.tudelft.nl/rlbook/ for additional material, including computer code used in the studies and information concerning new developments.

This best-selling text focuses on the analysis and design of complicated dynamics systems. CHOICE called it "a high-level, concise book that could well be used as a reference by engineers, applied mathematicians, and undergraduates. The format is good, the presentation clear, the diagrams instructive, the examples and problems helpful...References and a multiple-choice examination are included.

The problem of stochastic control of partially observable systems plays an important role in many applications. All real problems are in fact of this type, and deterministic control as well as stochastic control with full observation can only be approximations to the real world. This justifies the importance of having a theory as complete as possible, which can

be used for numerical implementation. This book first presents those problems under the linear theory that may be dealt with algebraically. Later chapters discuss the nonlinear filtering theory, in which the statistics are infinite dimensional and thus, approximations and perturbation methods are developed.

Optimal Stochastic Control, Stochastic Target Problems, and Backward SDE

Stochastic Control Theory

A Concise Introduction

Deterministic and Stochastic Optimal Control

Iterative Dynamic Programming

Handbook of Learning and Approximate Dynamic Programming

“Each chapter contains a well-written introduction and notes. They include the author's deep insights on the subject matter and provide historical comments and guidance to related literature. This book may well become an important milestone in the literature of optimal control.” —Mathematical Reviews “Thanks to a great effort to be self-contained, [this book] renders accessibly the subject to a wide audience. Therefore, it is recommended to all researchers and professionals interested in Optimal Control and its engineering and economic applications. It can serve as an excellent textbook for graduate courses in Optimal Control (with special emphasis on Nonsmooth Analysis).” —Automatica

This textbook offers a concise yet rigorous introduction to calculus of variations and optimal control theory, and is a self-contained resource for graduate students in engineering, applied mathematics, and related subjects. Designed specifically for a one-semester course, the book begins with calculus of variations, preparing the ground for optimal control. It then gives a complete proof of the maximum principle and covers key topics such as the Hamilton-Jacobi-Bellman theory of dynamic programming and linear-quadratic optimal control. Calculus of Variations and Optimal Control Theory also traces the historical development of the subject and features numerous exercises, notes and references at the end of each chapter, and suggestions for further study. Offers a concise yet rigorous introduction Requires limited background in control theory or advanced mathematics Provides a complete proof of the maximum principle Uses consistent notation in the exposition of classical and modern topics Traces the historical development of the subject Solutions manual (available only to teachers) Leading universities that have adopted this book include: University of Illinois at Urbana-Champaign ECE 553: Optimum Control Systems Georgia Institute of Technology ECE 6553: Optimal Control and Optimization University of Pennsylvania ESE 680: Optimal Control Theory University of Notre Dame EE 60565: Optimal Control Dynamic Programming and Its Applications provides information pertinent to the theory and application of dynamic programming. This book presents the development and future directions for dynamic programming. Organized into four parts encompassing 23 chapters, this book begins with an overview of recurrence conditions for countable state Markov decision problems, which ensure that the optimal average reward exists and satisfies the functional equation of dynamic programming. This text then provides an extensive analysis

of the theory of successive approximation for Markov decision problems. Other chapters consider the computational methods for deterministic, finite horizon problems, and present a unified and insightful presentation of several foundational questions. This book discusses as well the relationship between policy iteration and Newton's method. The final chapter deals with the main factors severely limiting the application of dynamic programming in practice. This book is a valuable resource for growth theorists, economists, biologists, mathematicians, and applied management scientists.

In this book, we study theoretical and practical aspects of computing methods for mathematical modelling of nonlinear systems. A number of computing techniques are considered, such as methods of operator approximation with any given accuracy; operator interpolation techniques including a non-Lagrange interpolation; methods of system representation subject to constraints associated with concepts of causality, memory and stationarity; methods of system representation with an accuracy that is the best within a given class of models; methods of covariance matrix estimation; methods for low-rank matrix approximations; hybrid methods based on a combination of iterative procedures and best operator approximation; and methods for information compression and filtering under condition that a filter model should satisfy restrictions associated with causality and different types of memory. As a result, the book represents a blend of new methods in general computational analysis, and specific, but also generic, techniques for study of systems theory and its particular branches, such as optimal filtering and information compression. - Best operator approximation, - Non-Lagrange interpolation, - Generic Karhunen-Loeve transform - Generalised low-rank matrix approximation - Optimal data compression - Optimal nonlinear filtering

Stochastic Controls

Adaptive Dynamic Programming: Single and Multiple Controllers

Applied Optimal Control

Adaptive Dynamic Programming with Applications in Optimal Control

The Discrete Time Case

Dynamic Optimization, Second Edition

This is the leading and most up-to-date textbook on the far-ranging algorithmic methodology of Dynamic Programming, which can be used for optimal control, Markovian decision problems, planning and sequential decision making under uncertainty, and discrete/combinatorial optimization. The treatment focuses on basic unifying themes, and conceptual foundations. It illustrates the versatility, power, and generality of the method with many examples and applications from engineering, operations research, and other fields. It also addresses extensively the practical application of the methodology, possibly through the use of approximations, and provides an extensive treatment of the far-reaching methodology of Neuro-Dynamic Programming/Reinforcement Learning. Among its special features, the book 1) provides a unifying framework for sequential decision making, 2) treats simultaneously deterministic and stochastic control problems popular in modern control theory and Markovian decision popular in

operations research, 3) develops the theory of deterministic optimal control problems including the Pontryagin Minimum Principle, 4) introduces recent suboptimal control and simulation-based approximation techniques (neuro-dynamic programming), which allow the practical application of dynamic programming to complex problems that involve the dual curse of large dimension and lack of an accurate mathematical model, 5) provides a comprehensive treatment of infinite horizon problems in the second volume, and an introductory treatment in the first volume.

Applied Dynamic Programming