OMB No. 8974702613615

Dynamic Programming Optimal Control Vol I

Nonlinear Control: Hamilton Jacobi Bellman (HJB) and Dynamic Programming Dimitri Bertsekas: Stable Optimal Control and Semicontractive Dynamic Programming Stochastic dynamic programming, optimization in policy space \u0026 DFO's "precautionary" harvest control Continuous Time Dynamic Programming -- The Hamilton-Jacobi-Bellman Equation Optimal Control (CMU 16-745) - Lecture 8: Controllability and Dynamic Programming Dynamic programing and LQ optimal control Dynamic Programming in Discrete Time [Optimal Control] Cart-pole Control using Differential Dynamic Programming Gear I Like: SOMA SimplexFM IT'S SO COMPLICATED! - The Game Changer Auto Reverb has so many features I can't cover them all 4 Books that will make you a smarter person in 2025 BEST BOOKS for Software Engineers by FAANG Senior Books every software engineer must read in 2023. Benjamin Recht: Optimization Perspectives on Learning to Control (ICML 2018 tutorial) Transforming an infinite horizon problem into a Dynamic Programming one Source Audio Artifakt Lo-Fi Elements Demo (IN STEREO) 5 Favorite ML Books for learning Machine Learning 15. Dynamic Programming, Part 1: SRTBOT, Fib, DAGs, Bowling Bryson Denham Optimal Control L5.1 - Introduction to dynamic programming and its application to discrete-time optimal control Optimal Control Solved with Excel and Python GEKKO Luus Optimal Control Problem Stable Optimal Control and Semicontractive Dynamic Programming ECE 493 Dynamic Programming 2 Stable Optimal Control and Semicontractive Dynamic Programming
Sparsity-Inducing Optimal Control via Differential Dynamic Programming
Reinforcement Learning and Optimal Control
Numerical Methods for Optimal Control Problems
Stochastic Controls
Handbook of Model Predictive Control
Reinforcement Learning and Approximate Dynamic Programming for Feedback Control
Adaptive Dynamic Programming: Single and Multiple Controllers
Dynamic Programming and Optimal Control
Dynamic Programming and Optimal Control
Optimal Control Theory and Static Optimization in Economics
Iterative Dynamic Programming
Optimal Stochastic Control, Stochastic Target Problems, and Backward SDE
Dynamic Programming and Optimal Control
Calculus of Variations and Optimal Control Theory
Rollout, Policy Iteration, and Distributed Reinforcement Learning
Reinforcement Learning and Stochastic Optimization
Dynamic Programming and the Calculus of Variations

Dynamic Programming Optimal Control Vol I *OMB No. 8974702613615 edited by*

MALONE BAKER

Reinforcement Learning and Optimal Control Princeton University Press
Dynamic Programming and Optimal ControlAthena Scientific

Numerical Methods for Optimal Control Problems

Springer Science & Business Media
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.

Stochastic Controls Springer

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

Handbook of Model Predictive Control

John Wiley & Sons

Optimal control theory is a technique being used increasingly by academic economists to study problems involving optimal decisions in a multi-period framework. This textbook is designed to make the difficult subject of optimal control theory easily accessible to economists while at the same time maintaining rigour. Economic intuitions are emphasized, and examples and problem sets covering a wide range of applications in economics are provided to assist in the learning process. Theorems are clearly stated and their proofs are carefully explained. The development of the text is gradual and fully integrated, beginning with simple formulations and progressing to advanced topics such as control parameters, jumps in state variables, and bounded state space. For greater economy and elegance, optimal control theory is introduced directly, without recourse to the calculus of variations. The connection with the latter and with dynamic programming is explained in a separate chapter. A second purpose of the book is to draw the parallel between optimal control theory and static optimization. Chapter 1 provides an extensive treatment of constrained and unconstrained maximization, with emphasis on economic insight and applications. Starting from basic concepts, it derives and explains important results, including the envelope theorem and the method of comparative statics. This chapter may be used for a course in static optimization. The book is largely self-contained. No previous knowledge of differential equations is required.

Reinforcement Learning and Approximate Dynamic Programming for Feedback Control

Athena Scientific

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).

Adaptive Dynamic Programming: Single and Multiple Controllers John Wiley & Sons

Automotive control has developed over the decades from an auxiliary te- nology 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 di?cult tuning. The progressive evolution has been mainly ledby speci?capplicationsandshorttermtargets,withthe consequencethat 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 ba- ground 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, severaldevelopments haveallowedusing these methods also for “fast”systemsandthis hassupporteda growinginterestinitsusealsofor automotive applications, with several promising results reported. Still there is no consensus on whether model predictive control with its high requi- ments on model quality and on computational power is a sensible choice for automotive control.

Dynamic Programming and Optimal Control Courier Corporation

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 fol lowing: (Q) What is the relationship betwccn the maximum principlc and dy namic programming in stochastic optimal controls? There did exist some researches (prior to the 1980s) on the relationship between these two. Nevertheless, the results usually werestated 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 equa tion, 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 or der in the stochastic case. This is known as a Hamilton-Jacobi-Bellman (HJB) equation.

Dynamic Programming and Optimal Control Springer

This book may be regarded as consisting of two parts. In Chapters I-IV we pre sent 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 opti mum, 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 pro gramming 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 inde pendently 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.

Optimal Control Theory and Static Optimization in Economics John Wiley & Sons

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 ant 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

Iterative Dynamic Programming Springer Science & Business Media

This monograph is concerned with the basic results on Cauchy problems associated with nonlinear monotone operators in Banach spaces with applications to partial differential equations of evolutive type. It focuses on major results in recent decades.

Optimal Stochastic Control, Stochastic Target Problems, and Backward SDE

Dynamic Programming and Optimal Control

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.dcsc.tudelft.nl/rlbook/ for additional material, including computer code used in the studies and information concerning new developments.

Dynamic Programming and Optimal Control Athena Scientific

Introduction to mathematical theory of multistage decision processes takes a "functional equation" approach. Topics include existence and uniqueness theorems, optimal inventory equation, bottleneck problems, multistage games, Markovian decision processes, and more. 1957 edition.

Calculus of Variations and Optimal Control Theory Springer Science & Business Media

Recent developments in model-predictive control promise remarkable opportunities for designing multi-input, multi-output control systems and improving the control of single-input, single-output systems. This volume provides a definitive survey of the latest model-predictive control methods available to engineers and scientists today. The initial set of chapters present various methods for managing uncertainty in systems, including stochastic model-predictive control. With the advent of affordable and fast computation, control engineers now need to think about using “computationally intensive controls,” so the second part of this book addresses the solution of optimization problems in “real” time for model-predictive control. The theory and applications of control theory often influence each other, so the last section of Handbook of Model Predictive Control rounds out the book with representative applications to automobiles, healthcare, robotics, and finance. The chapters in this volume will be useful to working engineers, scientists, and mathematicians, as well as students and faculty interested in the progression of control theory. Future developments in MPC will no doubt build from concepts demonstrated in this book and anyone with an interest in MPC will find fruitful information and suggestions for additional reading.

Rollout, Policy Iteration, and Distributed Reinforcement Learning Athena Scientific

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 schem

Reinforcement Learning and Stochastic Optimization Waveland Press

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.

Dynamic Programming and the Calculus of Variations

Springer Science & Business Media

The purpose of this book is to propose and develop a new conceptual framework for approximate Dynamic Programming (DP) and Reinforcement Learning (RL). This framework centers around two algorithms, which are designed largely independently of each other and operate in synergy through the powerful mechanism of Newton's method. We call these the off-line training and the on-line play algorithms; the names are borrowed from some of the major successes of RL involving games. Primary examples are the recent (2017) AlphaZero program (which plays chess), and the similarly structured and earlier (1990s) TD-Gammon program (which plays backgammon). In these game contexts, the off-line training algorithm is the method used to teach the program how to evaluate positions and to generate good moves at any given position, while the on-line play algorithm is the method used to play in real time against human or computer opponents. Both AlphaZero and TD-Gammon were trained off-line extensively using neural networks and an approximate version of the fundamental DP algorithm of policy iteration. Yet the AlphaZero player that was obtained off-line is not used directly during on-line play (it is too inaccurate due to approximation errors that are inherent in off-line neural network training). Instead a separate on-line player is used to select moves, based on multistep lookahead minimization and a terminal position evaluator that was trained using experience with the off-line player. The on-line player performs a form of policy improvement, which is not degraded by neural network approximations. As a result, it greatly improves the performance of the off-line player. Similarly, TD-Gammon performs on-line a policy improvement step using one-step or two-step lookahead minimization, which is not degraded by neural network approximations. To this end it uses an off-line neural network-trained terminal position evaluator, and importantly it also extends its on-line lookahead by rollout (simulation with the one-step lookahead player that is based on the position evaluator). Significantly, the synergy between off-line training and on-line play also underlies Model Predictive Control (MPC), a major control system design methodology that has been extensively developed since the 1980s. This synergy can be understood in terms of abstract models of infinite horizon DP and simple geometrical constructions, and helps to explain the all-important stability issues within the MPC context. An additional benefit of policy improvement by approximation in value space, not observed in the context of games (which have stable rules and environment), is that it works well with changing problem parameters and on-line replanning, similar to indirect adaptive control. Here the Bellman equation is perturbed due to the parameter changes, but approximation in value space still operates as a Newton step. An essential requirement here is that a system model is estimated on-line through some identification method, and is used during the one-step or multistep lookahead minimization process. In this monograph we aim to provide insights (often based on visualization), which explain the beneficial effects of on-line decision making on top of off-line training. In the process, we will bring out the strong connections between the artificial intelligence view of RL, and the control theory views of MPC and adaptive control. Moreover, we will show that in addition to MPC and adaptive control, our conceptual framework can be effectively integrated with other important methodologies such as multiagent systems and decentralized control, discrete and Bayesian optimization, and heuristic algorithms for discrete optimization. One of our principal aims is to show, through the algorithmic ideas of Newton's method and the unifying principles of abstract DP, that the AlphaZero/TD-Gammon methodology of approximation in value space and rollout applies very broadly to deterministic and stochastic optimal control problems. Newton's method here is used for the solution of Bellman's equation, an operator equation that applies universally within DP with both discrete and continuous state and control spaces, as well as finite and infinite horizon.

Reinforcement Learning and Optimal Control

John Wiley & Sons
This is the leading and most up-to-date textbook on the far-ranging algorithmic methododogy 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 and Optimal Control

Athena Scientific
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.
Optimal Control Athena Scientific
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.

Reinforcement Learning and Dynamic Programming Using Function Approximators

Springer

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.