Dynamic Programming Optimal Control Vol

Differential Dynamic Programming
Reinforcement Learning and Optimal Control
Rollout, Policy Iteration, and Distributed Reinforcement Learning
Dynamic Programming and Optimal Control
Convex Optimization Algorithms
Optimal Control
Stochastic Controls
Dynamic Optimization, Second Edition
Control and Dynamic Systems
Automotive Model Predictive Control
Semiconcave Functions, Hamilton-Jacobi Equations, and Optimal Control
Adaptive Dynamic Programming with Applications in Optimal Control
Theory of Optimal Control and Mathematical Programming
Numerical Methods for Optimal Control Problems
Optimal Control Theory
Approximate Dynamic Programming
Dynamic Programming and Stochastic Control
Stochastic Analysis, Control, Optimization and Applications
Computational Methods in Optimal Control Problems
Abstract Dynamic Programming
Lessons from AlphaZero for Optimal, Model Predictive, and Adaptive Control
A Survey of Optimal Control of Distributed-parameter Systems
Stochastic Optimal Control in Infinite Dimension
Continuous-time Stochastic Control and Optimization with Financial Applications
Linear Systems and Optimal Control
Controlled Markov Processes and Viscosity Solutions
Optimal Control
Self-Learning Optimal Control of Nonlinear Systems
Deterministic and Stochastic Optimal Control
Applied Optimal Control
Stochastic Control Theory
Convex Optimization Theory
Dynamic Programming and Optimal Control
Optimal Control Methods for Linear Discrete-Time Economic Systems
Optimal Control of Dynamic Operations Research Models
Adaptive Dynamic Programming: Single and Multiple Controllers
Dynamic Programming and Its Application to Optimal Control
Dynamic Programming and Optimal Control

Optimal Control Theory
Optimal Stochastic Control, Stochastic Target Problems, and Backward SDE

AINSLEY DALTON

Differential Dynamic Programming Springer
The purpose of this modest report is to present in a simplified manner some of the computational methods that have been developed in the last ten years for the solution of optimal control problems. Only those methods that are based on the minimum (maximum) principle of Pontriagin are discussed here. The autline of the report is as follows: In the first two sections a control problem of Bolza is formulated and the necessary conditions in the form of the minimum principle are given. The method of steepest descent and a conjugate gradient-method are dis cussed in Section 3. In the remaining sections, the successive sweep method, the Newton-Raphson method and the generalized Newton-Raphson method (also called quasilinearization method) ar~ presented from a unified approach which is based on the application of Newton Raphson approximation to the necessary conditions of optimality. The second-variation method and other shooting methods based on minimizing an error function are also considered. TABLE OF CONTENTS 1. 0 INTRODUCTION 1 2. 0 NECESSARY CONDITIONS FOR OPTIMALITY •••••••• 2 3. 0 THE GRADIENT METHOD 4 3. 1 Min H Method and Conjugate Gradient Method •. •••••••••. . . . ••••••. ••••••••. • 8 3. 2 Boundary Constraints •••••••••••. ••••. • 9 3. 3 Problems with Control Constraints ••. •• 15 4. 0 SUCCESSIVE SWEEP METHOD ••••••••••••••••••• 18 4. 1 Final Time Given Implicitly ••••. •••••• 22 5. 0 SECOND-VARIATION METHOD ••••••••••••••••••• 23 6. 0 SHOOTING METHODS ••••••••••••••••••••••••• 27 6. 1 Newton-Raphson Method •••••••••••••••• 27 6.
Reinforcement Learning and Optimal Control Academic Press
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.
Rollout, Policy Iteration, and Distributed Reinforcement Learning Athena Scientific
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.
Dynamic Programming and Optimal Control Springer Nature
A knowledge of linear systems provides a firm foundation for the study of optimal control theory and many areas of system theory and signal processing. State-space techniques developed since the early sixties have been proved to be very effective. The main objective of this book is to present a brief and somewhat complete investigation on the theory of linear systems, with emphasis on these techniques, in both continuous-time and discrete-time settings, and to demonstrate an application to the study of elementary (linear and nonlinear) optimal control theory. An essential feature of the state-space approach is that both time-varying and time-invariant systems are treated systematically. When time-varying systems are considered, another important subject that depends very much on the state-space formulation is perhaps real-time filtering, prediction, and smoothing via the Kalman filter. This subject is treated in our monograph entitled "Kalman Filtering with Real-Time Applications" published in this Springer Series in Information Sciences (Volume 17). For time-invariant systems, the recent frequency domain approaches using the techniques of Adamjan, Arov, and Krein (also known as AAK), balanced realization, and oo H theory via Nevanlinna-Pick interpolation seem very promising, and this will be studied in our forthcoming monograph entitled "Mathematical Ap proach to Signal Processing and System Theory". The present elementary treatise on linear system theory should provide enough engineering and mathe of these two subjects.
Convex Optimization Algorithms Athena Scientific
An insightful, concise, and rigorous treatment of the basic theory of convex sets and functions in finite dimensions, and the analytical/geometrical foundations of convex optimization and duality theory. Convexity theory is first developed in a simple accessible manner, using easily visualized proofs. Then the focus shifts to a transparent geometrical line of analysis to develop the fundamental duality between descriptions of convex functions in terms of points, and in terms of

hyperplanes. Finally, convexity theory and abstract duality are applied to problems of constrained optimization, Fenchel and conic duality, and game theory to develop the sharpest possible duality results within a highly visual geometric framework. This on-line version of the book, includes an extensive set of theoretical problems with detailed high-quality solutions, which significantly extend the range and value of the book. The book may be used as a text for a theoretical convex optimization course; the author has taught several variants of such a course at MIT and elsewhere over the last ten years. It may also be used as a supplementary source for nonlinear programming classes, and as a theoretical foundation for classes focused on convex optimization models (rather than theory). It is an excellent supplement to several of our books: Convex Optimization Algorithms (Athena Scientific, 2015), Nonlinear Programming (Athena Scientific, 2017), Network Optimization(Athena Scientific, 1998), Introduction to Linear Optimization (Athena Scientific, 1997), and Network Flows and Monotropic Optimization (Athena Scientific, 1998).

Optimal Control John Wiley & Sons

This book focuses on how to implement optimal control problems via the variational method. It studies how to implement the extrema of functional by applying the variational method and covers the extrema of functional with different boundary conditions, involving multiple functions and with certain constraints etc. It gives the necessary and sufficient condition for the (continuous-time) optimal control solution via the variational method, solves the optimal control problems with different boundary conditions, analyzes the linear quadratic regulator & tracking problems respectively in detail, and provides the solution of optimal control problems with state constraints by applying the Pontryagin's minimum principle which is developed based upon the calculus of variations. And the developed results are applied to implement several classes of popular optimal control problems and say minimum-time, minimum-fuel and minimum-energy problems and so on. As another key branch of optimal control methods, it also presents how to solve the optimal control problems via dynamic programming and discusses the relationship between the variational method and dynamic programming for comparison. Concerning the system involving individual agents, it is also worth to study how to implement the decentralized solution for the underlying optimal control problems in the framework of differential games. The equilibrium is implemented by applying both Pontryagin's minimum principle and dynamic programming. The book also analyzes the discrete-time version for all the above materials as well since the discrete-time optimal control problems are very popular in many fields.

Stochastic Controls Springer Science & Business Media

February 27 - March 1, 1997, the conference Optimal Control: The ory, Algorithms, and Applications took place at the University of Florida, hosted by the Center for Applied Optimization. The conference brought together researchers from universities, industry, and government laborato ries in the United States, Germany, Italy, France, Canada, and Sweden. There were forty-five invited talks, including seven talks by students. The conference was sponsored by the National Science Foundation and endorsed by the SIAM Activity Group on Control and Systems Theory, the Mathe matical Programming Society, the International Federation for Information Processing (IFIP), and the International Association for Mathematics and Computers in Simulation (IMACS). Since its inception in the 1940s and 1950s, Optimal Control has been closely connected to industrial applications, starting with aerospace. The program for the Gainesville conference, which reflected the rich cross-disci plinary flavor of the field, included aerospace applications as well as both novel and emerging applications to superconductors, diffractive optics, non linear optics, structural analysis, bioreactors, corrosion detection, acoustic flow, process design in chemical engineering, hydroelectric power plants, sterilization of canned foods, robotics, and thermoelastic plates and shells. The three days of the conference were organized around the three confer ence themes, theory, algorithms, and applications. This book is a collection of the papers presented at the Gainesville conference. We would like to take this opportunity to thank the sponsors and participants of the conference, the authors, the referees, and the publisher for making this volume possible.

Dynamic Optimization, Second Edition Springer

This book presents a class of novel, self-learning, optimal control schemes based on adaptive dynamic programming techniques, which quantitatively obtain the optimal control schemes of the systems. It analyzes the properties identified by the programming methods, including the convergence of the iterative value functions and the stability of the system under iterative control laws, helping to guarantee the effectiveness of the methods developed. When the system model is known, self-learning optimal control is designed on the basis of the system model; when the system model is not known, adaptive dynamic programming is implemented according to the system data, effectively making the performance of the system converge to the optimum. With various real-world examples to complement and substantiate the mathematical analysis, the book is a valuable guide for engineers, researchers, and students in control science and engineering.

Control and Dynamic Systems Springer

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.

Automotive Model Predictive Control Courier Corporation

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.

Semiconcave Functions, Hamilton-Jacobi Equations, and Optimal Control Athena Scientific

This book provides a comprehensive and accessible presentation of algorithms for solving convex optimization problems. It relies on rigorous mathematical analysis, but also aims at an intuitive exposition that makes use of visualization where possible. This is facilitated by the extensive use of analytical and algorithmic concepts of duality, which by nature lend themselves to geometrical interpretation. The book places particular emphasis on modern developments, and their widespread applications in fields such as large-scale resource allocation problems, signal processing, and machine learning. The book is aimed at students, researchers, and practitioners, roughly at the first year graduate level. It is similar in style to the author's 2009"Convex Optimization Theory" book, but can be read independently. The latter book focuses on convexity theory and optimization duality, while the present book focuses on algorithmic issues. The two books share notation, and together cover the entire finite-dimensional convex optimization methodology. To facilitate readability, the statements of definitions and results of the "theory book" are reproduced without proofs in Appendix B.

Adaptive Dynamic Programming with Applications in Optimal Control Athena Scientific

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

Theory of Optimal Control and Mathematical Programming Springer

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.

Numerical Methods for Optimal Control Problems Springer Science & Business Media

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.

Optimal Control Theory Springer Science & Business Media

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.

Approximate Dynamic Programming Elsevier Publishing Company

* A comprehensive and systematic exposition of the properties of semiconcave functions and their various applications, particularly to optimal control problems, by leading experts in the field * A central role in the present work is reserved for the study of singularities * Graduate students and researchers in optimal control, the calculus of variations, and PDEs will find this book useful as a reference work on modern dynamic programming for nonlinear control systems

Dynamic Programming and Stochastic Control Springer Science & Business Media

Dynamic Programming and Stochastic Control

Stochastic Analysis, Control, Optimization and Applications 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.

Computational Methods in Optimal Control Problems Springer Science & Business Media
Control and Dynamic Systems: Advances in Theory and Applications, Volume 10 brings together diverse information on important progress in the field of control and systems theory and applications. This volume is comprised of contributions from leading researchers in the field. Topics discussed include the evaluation of suboptimal strategies using quasilinearization; aircraft symmetric flight optimization; aircraft maneuver optimization by reduced-order approximation; and differential dynamic programming. Estimation of uncertain systems; application of modern control and optimization techniques to transportation systems; and integrated system identification and optimization are also elucidated. Aerospace engineers and scientists and researchers in applied sciences will find the book interesting.

Abstract Dynamic Programming Athena Scientific
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.