OMB No. 5085493426218

Optimization Over Time Dynamic Programming And Stochastic Control Wiley Series In Probability And Statistics Applied Probability And Statistics Section Volume 1

Mastering Dynamic Programming - How to solve any interview problem (Part 1) Finding Optimal Paths - Dynamic Programming Principle of Optimality - Dynamic Programming Economic Applications of Continuous Time Dynamic Programming (3/3): A Planning Problem Dynamic Programming Tutorial - Basics, Backward Recursion, and Principle of Optimality IQ TEST How DTW (Dynamic Time Warping) algorithm works The HARDEST part about programming 🤷‍♂️ #code #programming #technology #tech #software #developer Elon Musk Laughs at the Idea of Getting a PhD and Explains How to Actually Be Useful! I gave 127 interviews. Top 5 Algorithms they asked me. CSES Problemset | Goldmine in Competitive Programming | #shorts Economic Applications of Continuous Time Dynamic Programming (1/3): A Cake Eating Problem 4.7 Traveling Salesperson Problem - Dynamic Programming That's Why Mohit Sir Called \"God Of Mathematics\"| Puzzle Brain teaser | #competishun #shorts #tricks IIT Bombay CSE 🔥 #shorts #iit #iitbombay Best Programming Languages #programming #coding #javascript Dynamic Optimization Part 3: Continuous Time

Elements of Dynamic Optimization

Dynamic Programming and Optimal Control

Optimization Over Time. Dynamic Programming and Stochastic Control Vol. 2

Iterative Dynamic Programming

Abstract Dynamic Programming

Lectures on Stochastic Programming

Reinforcement Learning and Optimal Control

Approximate Dynamic Programming

sgfrgds

Dynamic Optimization for Beginners

Dynamic Optimization, Second Edition
Convex and Stochastic Optimization
Dynamic Optimization
Dynamic Optimization
Extensions of Dynamic Programming for Combinatorial Optimization and Data Mining

Optimization Over Time Dynamic Programming And Stochastic Control Wiley Series In Probability And Statistics Applied Probability And Statistics Section Volume 1

OMB No. 5085493426218 edited by

LI CUNNINGHAM

Elements of Dynamic Optimization Springer
Optimization Over

TimeOptimization Over TimeOptimization Over Time. Dynamic Programming and Stochastic Control Vol. 2Optimization Over TimeApproximate Dynamic ProgrammingJohn Wiley & Sons

DYNAMIC PROGRAMMING AND

OPTIMAL CONTROL

Athena Scientific
This monograph applies the relative optimization approach to time nonhomogeneous continuous-time and continuous-state dynamic systems. The approach is intuitively clear and does not require deep knowledge of the mathematics of partial

differential equations. The topics covered have the following distinguishing features: long-run average with no under-selectivity, non-smooth value functions with no viscosity solutions, diffusion processes with degenerate points, multi-class optimization with state classification, and optimization with no dynamic programming. The book begins with an introduction to relative optimization, including a comparison with the traditional approach of dynamic programming. The text then studies the Markov process, focusing on infinite-horizon optimization problems, and moves on to discuss optimal control of diffusion processes with semi-smooth value functions and degenerate points, and optimization of multi-dimensional diffusion processes. The book concludes with a brief overview of performance derivative-based optimization. Among the more important novel considerations presented are: the extension of the Hamilton–Jacobi–Bellman optimality condition from smooth to semi-smooth value functions by derivation of explicit optimality conditions at semi-smooth points and application of this result to degenerate and reflected processes; proof of semi-smoothness of the value function at degenerate points; attention to the under-selectivity issue for the long-run average and bias optimality; discussion of state classification for time nonhomogeneous continuous processes and multi-class optimization; and development of the multi-dimensional Tanaka formula for semi-smooth functions and application of this formula to stochastic control of multi-dimensional systems with degenerate points. The book will be of interest to researchers and students in the field of stochastic control and performance optimization alike.

Optimization Over Time. Dynamic Programming and Stochastic Control Vol. 2 Athena Scientific

Praise for the First Edition "Finally, a book devoted to dynamic programming and written using the language of operations research (OR)! This beautiful book fills a gap in the libraries of OR specialists and practitioners."
—Computing Reviews This new edition showcases a focus on modeling and computation for complex classes of approximate dynamic programming problems Understanding approximate dynamic programming (ADP) is vital in order to develop practical and high-quality solutions to complex industrial problems, particularly when those problems involve making decisions in the presence of uncertainty. Approximate Dynamic Programming, Second Edition uniquely integrates four distinct disciplines—Markov decision processes, mathematical programming, simulation, and statistics—to demonstrate how to successfully approach, model, and solve a wide range of real-life problems

using ADP. The book continues to bridge the gap between computer science, simulation, and operations research and now adopts the notation and vocabulary of reinforcement learning as well as stochastic search and simulation optimization. The author outlines the essential algorithms that serve as a starting point in the design of practical solutions for real problems. The three curses of dimensionality that impact complex problems are introduced and detailed coverage of implementation challenges is provided. The Second Edition also features: A new chapter describing four fundamental classes of policies for working with diverse stochastic optimization problems: myopic policies, look-ahead policies, policy function approximations, and policies based on value function approximations A new chapter on policy search that brings together stochastic search and simulation optimization concepts and introduces a new class of optimal learning strategies Updated coverage of the exploration exploitation problem in ADP, now including a recently developed method for doing active learning in the presence of a physical state, using the concept of the knowledge gradient A new sequence of chapters describing statistical methods for approximating value functions, estimating the value of a fixed policy, and value function approximation while searching for optimal policies The presented coverage of ADP emphasizes models and algorithms, focusing on related applications and computation while also discussing the theoretical side of the topic that explores proofs of convergence and rate of convergence. A related website features an ongoing discussion of the evolving fields of approximation dynamic programming and reinforcement learning, along with additional readings, software, and datasets. Requiring only a basic understanding of statistics and probability, Approximate Dynamic Programming, Second Edition is an excellent book for industrial engineering and operations research courses at the upper-undergraduate and graduate levels. It also serves as a valuable reference for researchers and professionals who utilize dynamic programming, stochastic programming, and control theory to solve problems in their everyday work.

Iterative Dynamic Programming

Academic Press Computing Methods in Optimization Problems deals with hybrid computing methods and optimization techniques using computers. One paper discusses different numerical approaches to optimizing trajectories, including the gradient method, the second variation method, and a generalized Newton-Raphson method. The paper cites the advantages and

disadvantages of each method, and compares the second variation method (a direct method) with the generalized Newton-Raphson method (an indirect method). An example problem illustrates the application of the three methods in minimizing the transfer time of a low-thrust ion rocket between the orbits of Earth and Mars. Another paper discusses an iterative process for steepest-ascent optimization of orbit transfer trajectories to minimize storage requirements such as in reduced memory space utilized in guidance computers. By eliminating state variable storage and control schedule storage, the investigator can achieve reduced memory requirements. Other papers discuss dynamic programming, invariant imbedding, quasilinearization, Hilbert space, and the computational aspects of a time-optimal control problem. The collection is suitable for computer programmers, engineers, designers of industrial processes, and researchers involved in aviation or control systems technology.
Abstract Dynamic Programming John Wiley & Sons
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.

Lectures on Stochastic Programming Courier Corporation

The availability of today's online information systems rapidly increases the relevance of dynamic decision making within a large number of operational contexts. Whenever a sequence of interdependent decisions occurs, making a single decision raises the need for anticipation of its future impact on the entire decision process. Anticipatory support is needed for a broad variety of dynamic and stochastic decision problems from different operational contexts such as finance, energy management, manufacturing and transportation. Example problems include asset allocation, feed-in of electricity produced by wind power as well as scheduling and routing. All these problems entail a sequence of decisions contributing to an overall goal and taking place in the course of a certain period of time. Each of the decisions is derived by solution of an optimization problem. As a consequence a stochastic and dynamic decision problem resolves into a series of optimization problems to be formulated and solved by anticipation of the remaining decision process. However, actually solving a dynamic decision problem by means of approximate dynamic programming still is a major scientific challenge. Most of the work done so far is devoted to problems allowing for formulation of the underlying optimization problems as linear programs. Problem domains like scheduling and routing, where linear programming typically does not produce a significant benefit for problem solving, have not been considered so far. Therefore, the industry demand for dynamic scheduling and routing is still predominantly satisfied by purely heuristic approaches to anticipatory decision making. Although this may work well for certain dynamic decision problems, these approaches lack transferability of findings to other, related problems. This book has serves two major purposes: - It provides a comprehensive and unique view of anticipatory optimization for dynamic decision making. It fully integrates Markov decision processes, dynamic programming, data mining and optimization and introduces a new perspective on approximate dynamic programming. Moreover, the book identifies different degrees of anticipation, enabling an assessment of specific approaches to dynamic decision making. - It shows for the first time how to successfully solve a dynamic vehicle routing problem by approximate dynamic programming. It elaborates on every building block required for this kind of approach to dynamic vehicle routing. Thereby the book has a pioneering character and is intended to provide a footing for the dynamic vehicle routing community.

Reinforcement Learning and Optimal Control CRC Press

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.

Approximate Dynamic Programming Courier Corporation

A path-breaking account of Markov decision processes-theory and computation This book's clear presentation of theory, numerous chapter-end problems, and development of a unified method for the computation of optimal policies in both discrete and continuous time make it an excellent course text for graduate students and advanced undergraduates. Its comprehensive coverage of important recent advances in stochastic dynamic programming makes it a valuable working resource for operations research professionals, management scientists, engineers, and others. Stochastic Dynamic Programming and the Control of Queueing Systems presents the theory of optimization under the finite horizon, infinite horizon discounted, and average cost criteria. It then shows how optimal rules of operation (policies) for each criterion may be numerically determined. A great wealth of examples from the application area of the control of queueing systems is presented. Nine numerical programs for the computation of optimal policies are fully explicated. The Pascal source code for the programs is available for viewing and downloading on the Wiley Web site at www.wiley.com/products/subject/mathematics. The site contains a link to the author's own Web site and is also a place where readers may discuss developments on the programs or other aspects of the material. The source files are also available via ftp at ftp://ftp.wiley.com/public/sci_tech_med/stochastic Stochastic Dynamic Programming and the Control of Queueing Systems features: * Path-breaking advances in Markov decision process techniques, brought together for the first time in book form * A theorem/proof format (proofs may be omitted without loss of continuity) * Development of a unified method for the computation of optimal rules of system operation * Numerous examples drawn mainly from the control of queueing systems * Detailed discussions of nine numerical programs * Helpful chapter-end problems * Appendices with complete treatment of background material *sgfrgds* Springer Nature The concept of a system as an entity in its own right has emerged with increasing force in the past few decades in, for example, the areas of electrical and control engineering, economics, ecology, urban structures, automaton theory, operational research and industry. The more definite concept of a large-scale system is implicit in these applications, but is particularly evident in fields such as the study of communication networks, computer networks and neural networks. The Wiley-Interscience Series in Systems and Optimization has been established to serve the needs of researchers in these rapidly developing fields. It is intended for works concerned with developments in quantitative systems theory, applications of such theory in areas of interest, or associated methodology. This is the first book-length treatment of risk-sensitive control, with many new results. The quadratic cost function of the standard LQG (linear/quadratic/Gaussian) treatment is replaced by the exponential of a quadratic, giving the so-called LEQG formulation allowing for a degree of optimism or pessimism on the part of the optimiser. The author is the first to achieve formulation and proof of risk-sensitive versions of the certainty-equivalence and separation principles. Further analysis allows one to formulate the optimization as the extremization of a path integral and to characterize the solution in terms of canonical factorization. It is thus possible to achieve the long-sought goal of an operational stochastic maximum principle, valid for a higher-order model, and in fact only evident when the models are extended to the risk-sensitive class. Additional results include deduction of compact relations between value functions and canonical factors, the exploitation of the equivalence between policy improvement and Newton Raphson methods and the direct relation of LEQG methods to the H??? and minimum-entropy methods. This book will prove essential reading for all graduate students, researchers and practitioners who have an interest in control theory including mathematicians, engineers, economists, physicists and psychologists. 1990 Stochastic Programming Peter Kall, University of Zurich, Switzerland and Stein W. Wallace, University of Trondheim, Norway Stochastic Programming is the first textbook to provide a thorough and self-contained introduction to the subject. Carefully written to cover all necessary background material from both linear and non-linear programming, as well as probability theory, the book draws together the methods and techniques previously described in disparate sources. After introducing the terms and modelling issues when randomness is introduced in a deterministic mathematical programming model, the authors cover decision trees and dynamic programming, recourse programming, recourse problems, probabilistic constraints, preprocessing and network problems. Exercises are provided at the end of each chapter. Throughout, the emphasis is on the appropriate use of the techniques, rather than on the underlying mathematical proofs and theories, making the book ideal for researchers and students in mathematical programming and operations research who wish to develop their skills in stochastic programming. 1994 **Dynamic Optimization for Beginners** Springer Nature Dynamic programming is an efficient technique for solving optimization problems. It is based on breaking the initial problem down into simpler ones and solving these sub-problems, beginning with the simplest ones. A conventional dynamic programming algorithm returns an optimal object from a given set of objects. This book develops extensions of dynamic programming, enabling us to (i) describe the set of objects under consideration; (ii) perform a multi-stage optimization of objects relative to different criteria; (iii) count the number of optimal objects; (iv) find the set of Pareto optimal points for bi-criteria optimization problems; and (v) to study relationships between two criteria. It considers various applications, including optimization of decision trees and decision rule systems as algorithms for problem solving, as ways for knowledge representation, and as classifiers; optimization of element partition trees for rectangular meshes, which are used in finite element methods for solving PDEs; and multi-stage optimization for such classic combinatorial optimization problems as matrix chain multiplication, binary search trees, global sequence alignment, and shortest paths. The results presented are useful for researchers in combinatorial optimization, data mining, knowledge discovery, machine learning, and finite element methods, especially those working in rough set theory, test theory, logical analysis of data, and PDE solvers. This book can be used as the basis for graduate courses.

Elsevier

This book explores discrete-time dynamic optimization and provides a detailed introduction to both deterministic and stochastic models. Covering problems with finite and infinite horizon, as well as Markov renewal programs, Bayesian control models and partially observable processes, the book focuses on the precise modelling of applications in a variety of areas, including operations research, computer science, mathematics, statistics, engineering, economics and finance. Dynamic Optimization is a carefully presented textbook which starts with discrete-time deterministic dynamic optimization problems, providing readers with the tools for sequential decision-making, before proceeding to the more complicated stochastic models. The authors present complete and simple proofs and illustrate the main results with numerous examples and exercises (without solutions). With relevant material covered in four appendices, this book is completely self-contained.

Dynamic Optimization, Second Edition

John Wiley & Sons
The main purpose of the book is to show how a viscosity approach can be used to tackle control problems in insurance. The problems covered are the maximization of survival probability as well as the maximization of dividends in the classical collective risk model. The authors consider the possibility of controlling the risk process by reinsurance as well as by investments. They show that optimal value functions are characterized as either the unique or the smallest viscosity solution of the associated Hamilton-Jacobi-Bellman equation; they also study the structure of the optimal strategies and show how to find them. The viscosity approach was widely used in control problems related to mathematical finance but until quite recently it was not used to solve control problems related to actuarial mathematical science. This book is designed to familiarize the reader on how to use this approach. The intended audience is graduate students as well as researchers in this area.

Convex and

Stochastic Optimization

John Wiley & Sons
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 Optimization

Academic Press
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).
Dynamic Optimization
John Wiley & Sons

Richard Bellmann developed a theory of dynamic programming which is for many reasons still in the center of great interest. The authors present a new approach in the ?eld of the optimization and multi-objective control of time-discrete systems which is closely related to the work of Richard Bellmann. They develop their own concept and their extension to the optimization and multi-objective control of time-discrete systems as well as to dynamic networks and multilayered structures are very stimulating for further research. Di?erent perspectives of discrete control and optimal dynamic ?ow problems on networks are treated and characterized. Together with the algorithmic solutions a framework of multi-objective control problems is - rived. The conclusion with a real world example underlines the necessity and - portance of their theoretic framework. As they come back to the classical Bellmann concept of dynamic programming they stress and honor his basic concept without debase their own work. Multilayereddecisionprocessesaspartofthedesignandanalysisofcomplexsystems and networks will be essential in many ways and ?elds in the future.
Extensions of Dynamic Programming for Combinatorial Optimization and Data Mining Springer Science & Business Media
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.
Computing Methods in Optimization Problems
Athena Scientific
A path-breaking account of Markov decision processes-theory and computation This book's clear presentation of theory, numerous chapter-end problems, and development of a unified method for the computation of optimal policies in both discrete and continuous time make it an excellent course text for graduate students and advanced undergraduates. Its comprehensive coverage of important recent advances in stochastic dynamic programming makes it a valuable working resource for operations research professionals, management scientists, engineers, and others. Stochastic Dynamic Programming and the Control of Queueing Systems presents the theory of optimization under the finite horizon, infinite horizon discounted, and average cost criteria. It then shows how optimal rules of operation (policies) for each criterion may be numerically determined. A great wealth of examples from the application area of the control of queueing systems is presented. Nine numerical programs for the computation of optimal policies are fully explicated. The Pascal source code for the programs is available for viewing and downloading

on the Wiley Web site at www.wiley.com/products/subject/mathematics. The site contains a link to the author's own Web site and is also a place where readers may discuss developments on the programs or other aspects of the material. The source files are also available via ftp at ftp://ftp.wiley.com/public/sci_tech_med/stochastic Stochastic Dynamic Programming and the Control of Queueing Systems features: * Path-breaking advances in Markov decision process techniques, brought together for the first time in book form * A theorem/proof format (proofs may be omitted without loss of continuity) * Development of a unified method for the computation of optimal rules of system operation * Numerous examples drawn mainly from the control of queueing systems * Detailed discussions of nine numerical programs * Helpful chapter-end problems * Appendices with complete treatment of background material

Dynamic Optimization
John Wiley & Sons
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.

Relative Optimization of Continuous-Time and Continuous-State Stochastic Systems SIAM This book introduces a fairly universal approach to the design and analysis of exact optimization algorithms for multi-objective combinatorial optimization problems. It proposes the circuits without repetitions representing the sets of feasible solutions along with the increasing and strictly increasing cost functions as a model for such problems. The book designs the algorithms for multi-stage and bi-criteria optimization and for counting the solutions in the framework of this model. As applications, this book studies eleven known combinatorial optimization problems: matrix chain multiplication, global sequence alignment, optimal paths in directed graphs, binary search trees, convex polygon triangulation, line breaking (text justification), one-dimensional clustering, optimal bitonic tour, segmented least squares, optimization of matchings in trees, and 0/1 knapsack problem. The results presented are useful for researchers in combinatorial optimization. This book is also useful as the basis for graduate courses.

Lessons from AlphaZero for Optimal, Model Predictive, and Adaptive Control Springer Optimization problems involving stochastic models occur in almost all areas of science and engineering, such as telecommunications, medicine, and finance. Their existence compels a need for rigorous ways of formulating, analyzing, and solving such problems. This book focuses on optimization problems involving uncertain parameters and covers the theoretical foundations and recent advances in areas where stochastic models are available. In Lectures on Stochastic Programming: Modeling and Theory, Second Edition, the authors introduce new material to reflect recent developments in stochastic programming, including: an analytical description of the tangent and normal cones of chance constrained sets; analysis of optimality conditions applied to nonconvex problems; a discussion of the stochastic dual dynamic programming method; an extended discussion of law invariant coherent risk measures and their Kusuoka representations; and in-depth analysis of dynamic risk measures and concepts of time consistency, including several new results.

Related with Optimization Over Time Dynamic Programming And Stochastic Control Wiley Series In Probability And Statistics Applied Probability And Statistics Section Volume 1: