

## Dynamic Programming Models And Applications

Incorporating a number of the author's recent ideas and examples, *Dynamic Programming: Foundations and Principles, Second Edition* presents a comprehensive and rigorous treatment of dynamic programming. The author emphasizes the crucial role that modeling plays in understanding this area. He also shows how Dijkstra's algorithm is an excellent example of a dynamic programming algorithm, despite the impression given by the computer science literature. New to the Second Edition

Expanded discussions of sequential decision models and the role of the state variable in modeling  
A new chapter on forward dynamic programming models  
A new chapter on the Push method that gives a dynamic programming perspective on Dijkstra's algorithm for the shortest path problem  
A new appendix on the Corridor method

Taking into account recent developments in dynamic programming, this edition continues to provide a systematic, formal outline of Bellman's approach to dynamic programming. It looks at dynamic programming as a problem-solving methodology, identifying its constituent components and explaining its theoretical basis for tackling problems.

*Operations Research: A Practical Introduction* is just that: a hands-on approach to the field of operations research (OR) and a useful guide for using OR techniques in scientific decision making, design, analysis and management. The text accomplishes two goals. First, it provides readers with an introduction to standard mathematical models and algorithms. Second, it is a thorough examination of practical issues relevant to the development and use of computational methods for problem solving.

**Highlights:** All chapters contain up-to-date topics and summaries  
A succinct presentation to fit a one-term course  
Each chapter has references, readings, and list of key terms  
Includes illustrative and current applications  
New exercises are added throughout the text  
Software tools have been updated with the newest and most popular software  
Many students of various disciplines such as mathematics, economics, industrial engineering and computer science often take one course in operations research. This book is written to provide a succinct and efficient introduction to the subject for these students, while offering a sound and fundamental preparation for more advanced courses in linear and nonlinear optimization, and many stochastic models and analyses. It provides relevant analytical tools for this varied audience and will also serve professionals, corporate managers, and technical consultants.

A collection of articles which provide examples that demonstrate the application of dynamic programming to a wide variety of decision problems in agriculture.

*Introduction to sequential decision processes* covers use of dynamic programming in studying models of resource allocation, methods for approximating solutions of control problems in continuous time, production control, more. 1982 edition.

*Introduction to Stochastic Dynamic Programming* presents the basic theory and examines the scope of applications of stochastic dynamic programming. The book begins with a chapter on various finite-stage models, illustrating the wide range of applications of stochastic dynamic programming. Subsequent chapters study infinite-stage models: discounting future returns, minimizing nonnegative costs, maximizing nonnegative returns, and maximizing the long-run average return. Each of these chapters first considers whether an optimal policy need exist—providing counterexamples where appropriate—and then presents methods for obtaining such

policies when they do. In addition, general areas of application are presented. The final two chapters are concerned with more specialized models. These include stochastic scheduling models and a type of process known as a multiproject bandit. The mathematical prerequisites for this text are relatively few. No prior knowledge of dynamic programming is assumed and only a moderate familiarity with probability—including the use of conditional expectation—is necessary.

Learn the science of collecting information to make effective decisions Everyday decisions are made without the benefit of accurate information. Optimal Learning develops the needed principles for gathering information to make decisions, especially when collecting information is time-consuming and expensive. Designed for readers with an elementary background in probability and statistics, the book presents effective and practical policies illustrated in a wide range of applications, from energy, homeland security, and transportation to engineering, health, and business. This book covers the fundamental dimensions of a learning problem and presents a simple method for testing and comparing policies for learning. Special attention is given to the knowledge gradient policy and its use with a wide range of belief models, including lookup table and parametric and for online and offline problems. Three sections develop ideas with increasing levels of sophistication: Fundamentals explores fundamental topics, including adaptive learning, ranking and selection, the knowledge gradient, and bandit problems Extensions and Applications features coverage of linear belief models, subset selection models, scalar function optimization, optimal bidding, and stopping problems Advanced Topics explores complex methods including simulation optimization, active learning in mathematical programming, and optimal continuous measurements Each chapter identifies a specific learning problem, presents the related, practical algorithms for implementation, and concludes with numerous exercises. A related website features additional applications and downloadable software, including MATLAB and the Optimal Learning Calculator, a spreadsheet-based package that provides an introduction to learning and a variety of policies for learning.

Optimization is the act of obtaining the "best" result under given circumstances. In design, construction, and maintenance of any engineering system, engineers must make technological and managerial decisions to minimize either the effort or cost required or to maximize benefits. There is no single method available for solving all optimization problems efficiently. Several optimization methods have been developed for different types of problems. The optimum-seeking methods are mathematical programming techniques (specifically, nonlinear programming techniques). Nonlinear Optimization: Models and Applications presents the concepts in several ways to foster understanding. Geometric interpretation: is used to re-enforce the concepts and to foster understanding of the mathematical procedures. The student sees that many problems can be analyzed, and approximate solutions found before analytical solutions techniques are applied. Numerical approximations: early on, the student is exposed to numerical techniques. These numerical procedures are algorithmic and iterative.

Worksheets are provided in Excel, MATLAB®, and Maple™ to facilitate the procedure. Algorithms: all algorithms are provided with a step-by-step format. Examples follow the summary to illustrate its use and application. Nonlinear Optimization: Models and Applications: Emphasizes process and interpretation throughout Presents a general classification of optimization problems Addresses situations that lead to models

illustrating many types of optimization problems Emphasizes model formulations  
Addresses a special class of problems that can be solved using only elementary calculus Emphasizes model solution and model sensitivity analysis  
About the author: William P. Fox is an emeritus professor in the Department of Defense Analysis at the Naval Postgraduate School. He received his Ph.D. at Clemson University and has taught at the United States Military Academy and at Francis Marion University where he was the chair of mathematics. He has written many publications, including over 20 books and over 150 journal articles. Currently, he is an adjunct professor in the Department of Mathematics at the College of William and Mary. He is the emeritus director of both the High School Mathematical Contest in Modeling and the Mathematical Contest in Modeling.

applicability of the inflating/deflating contour algorithm.

?A research monograph providing a synthesis of old research on the foundations of dynamic programming, with the modern theory of approximate dynamic programming 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 dynamic programming 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 dynamic programming, and 2) The new class of semicontractive models, exemplified by stochastic shortest path problems, where some but not all policies are contractive. The 2nd edition aims primarily to amplify the presentation of the semicontractive models of Chapter 3 and Chapter 4 of the first (2013) edition, and to supplement it with a broad spectrum of research results that I obtained and published in journals and reports since the first edition was written (see below). As a result, the size of this material more than doubled, and the size of the book increased by nearly 40%. The book is an excellent supplement to several of our books: *Dynamic Programming and Optimal Control* (Athena Scientific, 2017), and *Neuro-Dynamic Programming* (Athena Scientific, 1996).

Dynamic Programming Models and Applications Courier Corporation

*Dynamic Programming and Its Applications* provides information pertinent to the theory and application of dynamic programming. This book presents the development and future directions for dynamic programming. Organized into four parts encompassing 23 chapters, this book begins with an overview of recurrence conditions for countable state Markov decision problems, which ensure that the optimal average reward exists and satisfies the functional equation of dynamic programming. This text then provides an extensive analysis of the theory of successive approximation for Markov decision problems. Other chapters consider the computational methods for deterministic, finite horizon problems, and present a unified and insightful presentation of several foundational questions. This book discusses as well the relationship between policy iteration and Newton's method. The final chapter deals with the main factors severely

limiting the application of dynamic programming in practice. This book is a valuable resource for growth theorists, economists, biologists, mathematicians, and applied management scientists.

This textbook covers the fundamentals of optimization, including linear, mixed-integer linear, nonlinear, and dynamic optimization techniques, with a clear engineering focus. It carefully describes classical optimization models and algorithms using an engineering problem-solving perspective, and emphasizes modeling issues using many real-world examples related to a variety of application areas. Providing an appropriate blend of practical applications and optimization theory makes the text useful to both practitioners and students, and gives the reader a good sense of the power of optimization and the potential difficulties in applying optimization to modeling real-world systems. The book is intended for undergraduate and graduate-level teaching in industrial engineering and other engineering specialties. It is also of use to industry practitioners, due to the inclusion of real-world applications, opening the door to advanced courses on both modeling and algorithm development within the industrial engineering and operations research fields.

This book is intended to provide an introductory text of Nonlinear and Dynamic Programming for students of managerial economics and operations research. The author also hopes that engineers, business executives, managers, and others responsible for planning of industrial operations may find it useful as a guide to the problems and methods treated, with a view to practical applications. The book may be considered as a sequel to the author's Linear Programming in Industry (1960, 4th revised and enlarged edition 1974), but it can be used independently by readers familiar with the elements of linear programming models and techniques. The two volumes constitute an introduction to the methods of mathematical programming and their application to industrial optimization problems. The author feels that the vast and ever-increasing literature on mathematical programming has not rendered an introductory exposition superfluous. The general student often tends to feel somewhat lost if he goes straight to the special literature; he will be better equipped for tackling real problems and using computer systems if he has acquired some previous training in constructing small-scale programming models and applying standard algorithms for solving them by hand. The book is intended to provide this kind of training, keeping the mathematics at the necessary minimum. The text contains numerous exercises. The reader should work out these problems for himself and check with the answers given at the end of the book. The text is based on lectures given at the University of Copenhagen.

Dynamic programming is a method of solving multi-stage problems in which decisions at one stage become the conditions governing the succeeding stages. It can be applied to the management of water reservoirs, allowing them to be operated more efficiently. Originally published in 2007, this is one of the few books dedicated solely to dynamic programming techniques used in reservoir management. It presents the applicability of these techniques and their limits on the operational analysis of reservoir systems. The dynamic programming models presented in this book have been applied to reservoir systems all over the world, helping the reader to appreciate the applicability and limits of these models. The book also includes a model for the operation of a reservoir during an emergency situation. This volume will be a valuable reference to researchers in



hydrology, water resources and engineering, as well as professionals in reservoir management.

A single source guide to operations research (OR) techniques, this book covers emerging OR methodologies in a clear, concise, and unified manner. Building a bridge between theory and practice, it begins with coverage of fundamental models and methods such as linear, nonlinear, integer, and dynamic programming, networks, simulation, queuing, inventory, stochastic processes, and decision analysis. The book then explores emerging techniques including multiple criteria optimization, meta heuristics, robust optimization, and complexity and large scale networks. Each chapter gives an overview of a particular methodology, illustrates successful applications, and provides references to computer software availability.

This book presents an authoritative collection of contributions by researchers from 16 different countries (Austria, Chile, Georgia, Germany, Mexico, Norway, P.R. of China, Poland, North Macedonia, Romania, Russia, Spain, Turkey, Ukraine, the United Kingdom and United States) that report on recent developments and new directions in advanced control systems, together with new theoretical findings, industrial applications and case studies on complex engineering systems. This book is dedicated to Professor Vsevolod Mykhailovych Kuntsevich, an Academician of the National Academy of Sciences of Ukraine, and President of the National Committee of the Ukrainian Association on Automatic Control, in recognition of his pioneering works, his great scientific and scholarly achievements, and his years of service to many scientific and professional communities, notably those involved in automation, cybernetics, control, management and, more specifically, the fundamentals and applications of tools and techniques for dealing with uncertain information, robustness, non-linearity, extremal systems, discrete control systems, adaptive control systems and others. Covering essential theories, methods and new challenges in control systems design, the book is not only a timely reference guide but also a source of new ideas and inspirations for graduate students and researchers alike. Its 15 chapters are grouped into four sections: (a) fundamental theoretical issues in complex engineering systems, (b) artificial intelligence and soft computing for control and decision-making systems, (c) advanced control techniques for industrial and collaborative automation, and (d) modern applications for management and information processing in complex systems. All chapters are intended to provide an easy-to-follow introduction to the topics addressed, including the most relevant references. At the same time, they reflect various aspects of the latest research work being conducted around the world and, therefore, provide information on the state of the art.

Humans interact with and are part of the mysterious processes of nature. Inevitably they have to discover how to manage the environment for their long-term survival and benefit. To do this successfully means learning something about the dynamics of natural processes, and then using the knowledge to work

with the forces of nature for some desired outcome. These are intriguing and challenging tasks. This book describes a technique which has much to offer in attempting to achieve the latter task. A knowledge of dynamic programming is useful for anyone interested in the optimal management of agricultural and natural resources for two reasons. First, resource management problems are often problems of dynamic optimization. The dynamic programming approach offers insights into the economics of dynamic optimization which can be explained much more simply than can other approaches. Conditions for the optimal management of a resource can be derived using the logic of dynamic programming, taking as a starting point the usual economic definition of the value of a resource which is optimally managed through time. This is set out in Chapter I for a general resource problem with the minimum of mathematics. The results are related to the discrete maximum principle of control theory. In subsequent chapters dynamic programming arguments are used to derive optimality conditions for particular resources.

This book provides a practical introduction to computationally solving discrete optimization problems using dynamic programming. From the examples presented, readers should more easily be able to formulate dynamic programming solutions to their own problems of interest. We also provide and describe the design, implementation, and use of a software tool that has been used to numerically solve all of the problems presented earlier in the book. Features step-by-step examples based on actual data and connects fundamental mathematical modeling skills and decision making concepts to everyday applicability. Featuring key linear programming, matrix, and probability concepts, "Finite Mathematics: Models and Applications" emphasizes cross-disciplinary applications that relate mathematics to everyday life. The book provides a unique combination of practical mathematical applications to illustrate the wide use of mathematics in fields ranging from business, economics, finance, management, operations research, and the life and social sciences. In order to emphasize the main concepts of each chapter, "Finite Mathematics: Models and Applications" features plentiful pedagogical elements throughout such as special exercises, end notes, hints, select solutions, biographies of key mathematicians, boxed key principles, a glossary of important terms and topics, and an overview of use of technology. The book encourages the modeling of linear programs and their solutions and uses common computer software programs such as LINDO. In addition to extensive chapters on probability and statistics, principles and applications of matrices are included as well as topics for enrichment such as the Monte Carlo method, game theory, kinship matrices, and dynamic programming. Supplemented with online instructional support materials, the book features coverage including: Algebra Skills Mathematics of Finance Matrix Algebra Geometric Solutions Simplex Methods Application Models Set and Probability Relationships Random Variables and Probability Distributions Markov Chains Mathematical Statistics Enrichment in Finite Mathematics An ideal

textbook, " Finite Mathematics: Models and Applications "is intended for students in fields from entrepreneurial and economic to environmental and social science, including many in the arts and humanities. Carla C. Morris, PhD, is Assistant Professor of Mathematics in the Associate in Arts Program at the University of Delaware. A member of The Institute for Operations Research and the Management Sciences and the Mathematical Association of America, Dr. Morris teaches courses ranging from college algebra to calculus and statistics. Robert M. Stark, PhD, is Professor Emeritus in the Departments of Mathematical Sciences and Civil and Environmental Engineering at the University of Delaware. Dr. Stark's teaching and research interests include applied probability, mathematical optimization, operations research, and mathematics education.

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.

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.

An accessible treatment of the modeling and solution of integer programming problems, featuring modern applications and software. In order to fully comprehend the algorithms associated with integer programming, it is important to understand not only how algorithms work, but also why they work. Applied Integer Programming features a unique emphasis on this point, focusing on problem modeling and solution using commercial software. Taking an application-oriented approach, this book addresses the art and science of mathematical modeling related to the mixed integer programming (MIP) framework and discusses the algorithms and associated practices that enable those models to be solved most efficiently. The book begins with coverage of successful applications, systematic modeling procedures, typical model types, transformation of non-MIP models, combinatorial optimization problem models, and automatic preprocessing to obtain a better formulation. Subsequent chapters present



algebraic and geometric basic concepts of linear programming theory and network flows needed for understanding integer programming. Finally, the book concludes with classical and modern solution approaches as well as the key components for building an integrated software system capable of solving large-scale integer programming and combinatorial optimization problems. Throughout the book, the authors demonstrate essential concepts through numerous examples and figures. Each new concept or algorithm is accompanied by a numerical example, and, where applicable, graphics are used to draw together diverse problems or approaches into a unified whole. In addition, features of solution approaches found in today's commercial software are identified throughout the book. Thoroughly classroom-tested, *Applied Integer Programming* is an excellent book for integer programming courses at the upper-undergraduate and graduate levels. It also serves as a well-organized reference for professionals, software developers, and analysts who work in the fields of applied mathematics, computer science, operations research, management science, and engineering and use integer-programming techniques to model and solve real-world optimization problems.

This classic book is an introduction to dynamic programming, presented by the scientist who coined the term and developed the theory in its early stages. In *Dynamic Programming*, Richard E. Bellman introduces his groundbreaking theory and furnishes a new and versatile mathematical tool for the treatment of many complex problems, both within and outside of the discipline. The book is written at a moderate mathematical level, requiring only a basic foundation in mathematics, including calculus. The applications formulated and analyzed in such diverse fields as mathematical economics, logistics, scheduling theory, communication theory, and control processes are as relevant today as they were when Bellman first presented them. A new introduction by Stuart Dreyfus reviews Bellman's later work on dynamic programming and identifies important research areas that have profited from the application of Bellman's theory.

*Economic Modeling and Inference* takes econometrics to a new level by demonstrating how to combine modern economic theory with the latest statistical inference methods to get the most out of economic data. This graduate-level textbook draws applications from both microeconomics and macroeconomics, paying special attention to financial and labor economics, with an emphasis throughout on what observations can tell us about stochastic dynamic models of rational optimizing behavior and equilibrium. Bent Jesper Christensen and Nicholas Kiefer show how parameters often thought estimable in applications are not identified even in simple dynamic programming models, and they investigate the roles of extensions, including measurement error, imperfect control, and random utility shocks for inference. When all implications of optimization and equilibrium are imposed in the empirical procedures, the resulting estimation problems are often nonstandard, with the estimators exhibiting nonregular asymptotic behavior such as short-ranked covariance, superconsistency, and non-Gaussianity. Christensen and Kiefer explore these properties in detail, covering areas including job search models of the labor market, asset pricing, option pricing, marketing, and retirement planning. Ideal for researchers and practitioners as well as students, *Economic Modeling and Inference* uses real-world data to illustrate how to derive the best results using a combination of theory and cutting-edge econometric techniques. Covers identification and estimation of dynamic programming models Treats sources of

error--measurement error, random utility, and imperfect control Features financial applications including asset pricing, option pricing, and optimal hedging Describes labor applications including job search, equilibrium search, and retirement Illustrates the wide applicability of the approach using micro, macro, and marketing examples

Operations Research (OR) began as an interdisciplinary activity to solve complex military problems during World War II. Utilizing principles from mathematics, engineering, business, computer science, economics, and statistics, OR has developed into a full fledged academic discipline with practical application in business, industry, government and military. Currently regarded as a body of established mathematical models and methods essential to solving complicated management issues, OR provides quantitative analysis of problems from which managers can make objective decisions. Operations Research and Management Science (OR/MS) methodologies continue to flourish in numerous decision making fields. Featuring a mix of international authors, Operations Research and Management Science Handbook combines OR/MS models, methods, and applications into one comprehensive, yet concise volume. The first resource to reach for when confronting OR/MS difficulties, this text – Provides a single source guide in OR/MS Bridges theory and practice Covers all topics relevant to OR/MS Offers a quick reference guide for students, researchers and practitioners Contains unified and up-to-date coverage designed and edited with non-experts in mind Discusses software availability for all OR/MS techniques Includes contributions from a mix of domestic and international experts The 26 chapters in the handbook are divided into two parts. Part I contains 14 chapters that cover the fundamental OR/MS models and methods. Each chapter gives an overview of a particular OR/MS model, its solution methods and illustrates successful applications. Part II of the handbook contains 11 chapters discussing the OR/MS applications in specific areas. They include airlines, e-commerce, energy systems, finance, military, production systems, project management, quality control, reliability, supply chain management and water resources. Part II ends with a chapter on the future of OR/MS applications.

The goal of the Encyclopedia of Optimization is to introduce the reader to a complete set of topics that show the spectrum of research, the richness of ideas, and the breadth of applications that has come from this field. The second edition builds on the success of the former edition with more than 150 completely new entries, designed to ensure that the reference addresses recent areas where optimization theories and techniques have advanced. Particularly heavy attention resulted in health science and transportation, with entries such as "Algorithms for Genomics", "Optimization and Radiotherapy Treatment Design", and "Crew Scheduling".

Optimization models play an increasingly important role in financial decisions. This is the first textbook devoted to explaining how recent advances in optimization models, methods and software can be applied to solve problems in computational finance more efficiently and accurately. Chapters discussing the theory and efficient solution methods for all major classes of optimization problems alternate with chapters illustrating their use in modeling problems of mathematical finance. The reader is guided through topics such as volatility estimation, portfolio optimization problems and constructing an index fund, using techniques such as nonlinear optimization models, quadratic programming formulations and integer programming models respectively. The book is based on Master's courses in financial engineering and comes with worked examples, exercises

and case studies. It will be welcomed by applied mathematicians, operational researchers and others who work in mathematical and computational finance and who are seeking a text for self-learning or for use with courses.

DPMax stands for 'dynamic programming to the max'. It highlights the graphical and textual analyses of 2 of the most common dynamic programming algorithms: The Longest Common Subsequence and The Longest/Shortest Paths Using Weights. It takes a brief look at the subjects of optimization and dynamic programming before delving into the core subjects of the book. It is a must-have for bioinformaticians, computer scientists and molecular biologists.

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

Individuals and families make key decisions that impact many aspects of financial stability and determine the future of the economy. These decisions involve balancing current sacrifice against future benefits. People have to decide how much to invest in health care, exercise, their diet, and insurance. They must decide how much debt to take on, and how much to save. And they make choices about jobs that determine employment and unemployment levels. Forward-Looking Decision Making is about modeling this individual or family-based decision making using an optimizing dynamic programming model. Robert Hall first reviews ideas about dynamic programs and introduces new ideas about numerical solutions and the representation of solved models as Markov processes. He surveys recent research on the parameters of preferences--the intertemporal elasticity of substitution, the Frisch elasticity of labor supply, and the Frisch cross-elasticity. He then examines dynamic programming models applied to health spending, long-term care insurance, employment, entrepreneurial risk-taking, and consumer debt. Linking theory with data and applying them to real-world problems, Forward-Looking Decision Making uses dynamic optimization programming models to shed light on individual behaviors and their economic implications.

The trustworthiness of models and optimization is limited because the associated systems might be changing and data about them can be limited, i.e., there is "parametric" uncertainty. This dissertation provides applications and theory related to mitigating the effects of changing systems and data limitations in optimal decision-

making. The primary application considered relates to reducing the maintenance costs associated with cyber security. By selecting optimal policies addressing data limitations, losses from stolen information and maintenance costs can be balanced. The approximated expected savings from implementing the suggested policies at a large Midwestern organization is over \$14M with a discount factor of 0.95 monthly. The dissertation also integrates data and dynamic programming models for project management decision-making that accounts for coordination and planning costs. This facilitates more accurate schedules with significant cost savings. Insights are provided into the choice between traditional planning methods and agile project management methods that reduce planning complexity. In many situations, we find that the so-called optimal approaches are suboptimal because they fail to address sizable coordination and planning costs. Two types of parametric uncertainty are explored here, each of which results in fundamentally different formulations and solution schemes. The first type of uncertainty considered relates to system parameters fluctuating over time randomly. The related models differ from ordinary inhomogeneous approaches because the specific parameters are not known and are assumed to fluctuate with known distributions. Associated decision problems are referred to as "Markov decision processes with random inhomogeneity" and proposed optimal solutions methods. Proof is given that the solution produced by backward induction is optimal for the finite horizon problems, and that the value-iteration-based algorithm gives solutions converging to the infinite horizon solutions, together with results regarding monotonicity property and rate of the convergence. The second type of parametric uncertainty is caused by insufficient data for parameter estimation, i.e., "data-driven" uncertainty. Previous researchers studying data-driven Markov decision processes declare the problem is intractable. Therefore, they propose approximation methods. We prove that their methods can approximate suboptimal solutions by a numerical example. We also provide a dynamic programming algorithm to generate data-driven optimal policies with learning. We do this by demonstrating that the problem is equivalent to partially observable Markov decision processes. Further, by exploiting the structure of the problem and bounds assuming perfect information, we develop a bounding heuristic method for the infinite horizon problems.

[Copyright: f684f2f4590f6481c46e6da061992863](https://doi.org/10.1007/978-1-4939-9928-6)