Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics)

By Warren B. Powell

Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) By Warren B. Powell

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.

Download Approximate Dynamic Programming: Solving the Curse ...pdf

Read Online Approximate Dynamic Programming: Solving the Cur ...pdf

Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics)

By Warren B. Powell

Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) By Warren B. Powell

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.

Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) By Warren B. Powell Bibliography

- Sales Rank: #727448 in Books
- Published on: 2011-09-27
- Original language: English
- Number of items: 1
- Dimensions: 9.30" h x 1.60" w x 6.10" l, 2.20 pounds
- Binding: Hardcover
- 606 pages

Download and Read Free Online Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) By Warren B. Powell

Editorial Review

From the Author

This book is a major revision of the first edition, with seven new or heavily revised chapters. This edition starts clearly from a foundation in reinforcement learning (using classical RL notation and concepts), but continues to build a bridge to the types of high-dimensional problems familiar to operations research. The book uses "a" for discrete actions, but switches to "x" for vector-valued decisions, common in the operations research community but largely ignored in computer science.

The book begins with an overview of a wide array of problems (in chapter 2), with an introduction to classical Markov decision processes in chapter 3 (minor changes from the first edition). Chapter 4 is a completely rewritten introduction to reinforcement learning using classical concepts, with one major exception. It now provides an extended overview of the concept of the post-decision state variable, which is used throughout the book (because it avoids the imbedded expectation within the min/max operator). The RL community is very familiar with the use of state-action pairs (in Q-factors), which is a clumsy form of post-decision state. Chapter 4 gives a series of examples where the post-decision state can (but does not always) provide significant benefits.

Chapter 5 is a minor revision of the old chapter 5, providing an in-depth discussion of how to model a dynamic program.

The book brings together different fields within stochastic optimization by identifying (in chapter 6) four fundamental classes of policies: 1) myopic policies (which ignore the future), 2) lookahead policies (which optimize over a short horizon to determine the decision to be made now), 3) policy function approximations (analytic functions that return an action given a state), and 4) policies based on value function approximations. Policy function approximations (a little-used term that I am promoting) and value function approximations both require some sort of method for approximating a function, of which there are three basic classes: lookup tables, parametric models, and nonparametric models.

Of course, you can build hybrid policies by mixing and matching.

Chapter 7 provides an in-depth treatment of methods for optimizing policy function approximations (widely known as "policy search" in the RL literature). In addition to classical material from stochastic search, it also includes a description of the knowledge gradient concept, which we developed as part of research to develop insights into the exploration-exploitation problem.

Chapters 8, 9 and 10 provide a layered presentation of how to go about designing policies based on value function approximations. Chapter 8 is an overview of popular methods for approximating functions. Chapter 9 discusses methods for obtaining an update to a value function approximation for a fixed policy, and Chapter 10 describes the complex problem of approximating a value function approximation while simultaneously optimizing over policies.

Chapter 11 provides an introduction to the fundamentals of updating value function approximations based on stochastic approximation methods, with an overview of stepsize formulas. This is a revision of the old chapter 6, somewhat streamlined but with new insights into how to design stepsize rules for different algorithmic strategies (Q-learning, approximate value iteration, LSTD, LSPE) and a new, optimal stepsize

rule designed specifically for methods based on bootstrapping (TD(0), Q-learning, approximate value iteration).

Chapter 12 is a completely rewritten chapter on the exploration vs. exploitation problem, including a new algorithm for using the knowledge gradient idea (designed originally for pure learning problems) in the presence of a physical state.

Chapters 13 and 14 retain the material from the first edition on designing value function approximations in the context of high-dimensional resource allocation problems which exploit convexity.

Warren Powell

About the Author

WARREN B. POWELL, PhD, is Professor of Operations Research and Financial Engineering at Princeton University, where he is founder and Director of CASTLE Laboratory, a research unit that works with industrial partners to test new ideas found in operations research. The recipient of the 2004 INFORMS Fellow Award, Dr. Powell has authored more than 160 published articles on stochastic optimization, approximate dynamicprogramming, and dynamic resource management.

Users Review

From reader reviews:

Linda Long:

Information is provisions for those to get better life, information presently can get by anyone with everywhere. The information can be a know-how or any news even restricted. What people must be consider when those information which is inside the former life are challenging be find than now could be taking seriously which one is suitable to believe or which one the resource are convinced. If you obtain the unstable resource then you understand it as your main information you will have huge disadvantage for you. All those possibilities will not happen with you if you take Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) as your daily resource information.

Donna Bledsoe:

A lot of people always spent all their free time to vacation or maybe go to the outside with them friends and family or their friend. Do you know? Many a lot of people spent many people free time just watching TV, or even playing video games all day long. If you want to try to find a new activity that is look different you can read a book. It is really fun for you personally. If you enjoy the book that you simply read you can spent all day long to reading a publication. The book Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) it is very good to read. There are a lot of folks that recommended this book. They were enjoying reading this book. When you did not have enough space to deliver this book you can buy the actual e-book. You can m0ore quickly to read this book through your smart phone. The price is not very costly but this book provides high quality.

Shelia Tonn:

That reserve can make you to feel relax. This specific book Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) was colorful and of course has pictures around. As we know that book Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) has many kinds or category. Start from kids until young adults. For example Naruto or Investigator Conan you can read and believe you are the character on there. Therefore not at all of book are usually make you bored, any it offers you feel happy, fun and loosen up. Try to choose the best book for you personally and try to like reading that will.

Russell Fielder:

A lot of reserve has printed but it takes a different approach. You can get it by online on social media. You can choose the best book for you, science, comedian, novel, or whatever simply by searching from it. It is referred to as of book Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics). You can include your knowledge by it. Without departing the printed book, it might add your knowledge and make anyone happier to read. It is most crucial that, you must aware about book. It can bring you from one place to other place.

Download and Read Online Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) By Warren B. Powell #YJ3RC2GW50I

Read Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) By Warren B. Powell for online ebook

Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) By Warren B. Powell Free PDF d0wnl0ad, audio books, books to read, good books to read, cheap books, good books, online books, books online, book reviews epub, read books online, books to read online, online library, greatbooks to read, PDF best books to read, top books to read Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) By Warren B. Powell books to read online.

Online Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) By Warren B. Powell ebook PDF download

Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) By Warren B. Powell Doc

Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) By Warren B. Powell Mobipocket

Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition (Wiley Series in Probability and Statistics) By Warren B. Powell EPub