

Reinforcement Learning and Optimal Control

by

Dimitri P. Bertsekas

Massachusetts Institute of Technology

DRAFT TEXTBOOK

This is a draft of a textbook that is scheduled to be finalized in 2019, and to be published by Athena Scientific. It represents “work in progress,” and it will be periodically updated. It more than likely contains errors (hopefully not serious ones). Furthermore, its references to the literature are incomplete. Your comments and suggestions to the author at dimitrib@mit.edu are welcome. The date of last revision is given below.

May 14, 2019

WWW site for book information and orders

<http://www.athenasc.com>



Athena Scientific, Belmont, Massachusetts

Athena Scientific
Post Office Box 805
Nashua, NH 03060
U.S.A.

Email: info@athenasc.com
WWW: <http://www.athenasc.com>

Publisher's Cataloging-in-Publication Data

Bertsekas, Dimitri P.
Reinforcement Learning and Optimal Control
Includes Bibliography and Index
1. Mathematical Optimization. 2. Dynamic Programming. I. Title.
QA402.5 .B465 2019 519.703 00-91281

ISBN-10: 1-886529-39-6, ISBN-13: 978-1-886529-39-7

ABOUT THE AUTHOR

Dimitri Bertsekas studied Mechanical and Electrical Engineering at the National Technical University of Athens, Greece, and obtained his Ph.D. in system science from the Massachusetts Institute of Technology. He has held faculty positions with the Engineering-Economic Systems Department, Stanford University, and the Electrical Engineering Department of the University of Illinois, Urbana. Since 1979 he has been teaching at the Electrical Engineering and Computer Science Department of the Massachusetts Institute of Technology (M.I.T.), where he is currently the McAfee Professor of Engineering.

His teaching and research have spanned several fields, including deterministic optimization, dynamic programming and stochastic control, large-scale and distributed computation, and data communication networks. He has authored or coauthored numerous research papers and seventeen books, several of which are currently used as textbooks in MIT classes, including “Dynamic Programming and Optimal Control,” “Data Networks,” “Introduction to Probability,” and “Nonlinear Programming.”

Professor Bertsekas was awarded the INFORMS 1997 Prize for Research Excellence in the Interface Between Operations Research and Computer Science for his book “Neuro-Dynamic Programming” (co-authored with John Tsitsiklis), the 2001 AACC John R. Ragazzini Education Award, the 2009 INFORMS Expository Writing Award, the 2014 AACC Richard Bellman Heritage Award, the 2014 INFORMS Khachiyan Prize for Lifetime Accomplishments in Optimization, the 2015 MOS/SIAM George B. Dantzig Prize, and the 2018 INFORMS John von Neumann Theory Prize. In 2001, he was elected to the United States National Academy of Engineering for “pioneering contributions to fundamental research, practice and education of optimization/control theory, and especially its application to data communication networks.”

ATHENA SCIENTIFIC
OPTIMIZATION AND COMPUTATION SERIES

1. Abstract Dynamic Programming, 2nd Edition, by Dimitri P. Bertsekas, 2018, ISBN 978-1-886529-46-5, 360 pages
2. Dynamic Programming and Optimal Control, Two-Volume Set, by Dimitri P. Bertsekas, 2017, ISBN 1-886529-08-6, 1270 pages
3. Nonlinear Programming, 3rd Edition, by Dimitri P. Bertsekas, 2016, ISBN 1-886529-05-1, 880 pages
4. Convex Optimization Algorithms, by Dimitri P. Bertsekas, 2015, ISBN 978-1-886529-28-1, 576 pages
5. Convex Optimization Theory, by Dimitri P. Bertsekas, 2009, ISBN 978-1-886529-31-1, 256 pages
6. Introduction to Probability, 2nd Edition, by Dimitri P. Bertsekas and John N. Tsitsiklis, 2008, ISBN 978-1-886529-23-6, 544 pages
7. Convex Analysis and Optimization, by Dimitri P. Bertsekas, Angelia Nedić, and Asuman E. Ozdaglar, 2003, ISBN 1-886529-45-0, 560 pages
8. Network Optimization: Continuous and Discrete Models, by Dimitri P. Bertsekas, 1998, ISBN 1-886529-02-7, 608 pages
9. Network Flows and Monotropic Optimization, by R. Tyrrell Rockafellar, 1998, ISBN 1-886529-06-X, 634 pages
10. Introduction to Linear Optimization, by Dimitris Bertsimas and John N. Tsitsiklis, 1997, ISBN 1-886529-19-1, 608 pages
11. Parallel and Distributed Computation: Numerical Methods, by Dimitri P. Bertsekas and John N. Tsitsiklis, 1997, ISBN 1-886529-01-9, 718 pages
12. Neuro-Dynamic Programming, by Dimitri P. Bertsekas and John N. Tsitsiklis, 1996, ISBN 1-886529-10-8, 512 pages
13. Constrained Optimization and Lagrange Multiplier Methods, by Dimitri P. Bertsekas, 1996, ISBN 1-886529-04-3, 410 pages
14. Stochastic Optimal Control: The Discrete-Time Case, by Dimitri P. Bertsekas and Steven E. Shreve, 1996, ISBN 1-886529-03-5, 330 pages

Contents

1. Exact Dynamic Programming	
1.1. Deterministic Dynamic Programming	p. 2
1.1.1. Deterministic Problems	p. 2
1.1.2. The Dynamic Programming Algorithm	p. 7
1.1.3. Approximation in Value Space	p. 12
1.2. Stochastic Dynamic Programming	p. 14
1.3. Examples, Variations, and Simplifications	p. 17
1.3.1. Deterministic Shortest Path Problems	p. 19
1.3.2. Discrete Deterministic Optimization	p. 21
1.3.3. Problems with a Terminal State	p. 24
1.3.4. Forecasts	p. 26
1.3.5. Problems with Uncontrollable State Components	p. 28
1.3.6. Partial State Information and Belief States	p. 33
1.3.7. Linear Quadratic Optimal Control	p. 37
1.3.8. Systems with Unknown Parameters - Adaptive Control	p. 41
1.4. Reinforcement Learning and Optimal Control - Some Terminology	p. 44
1.5. Notes and Sources	p. 46
2. Approximation in Value Space	
2.1. General Issues of Approximation in Value Space	p. 6
2.1.1. Methods for Computing Approximations in Value Space	p. 7
2.1.2. Off-Line and On-Line Methods	p. 8
2.1.3. Model-Based Simplification of the Lookahead Minimization	p. 9
2.1.4. Model-Free Q-Factor Approximation in Value Space	p. 10
2.1.5. Approximation in Policy Space on Top of Approximation in Value Space	p. 13
2.1.6. When is Approximation in Value Space Effective?	p. 14

2.2. Multistep Lookahead	p. 16
2.2.1. Multistep Lookahead and Rolling Horizon	p. 16
2.2.2. Multistep Lookahead and Deterministic Problems	p. 18
2.3. Problem Approximation	p. 20
2.3.1. Enforced Decomposition	p. 21
2.3.2. Probabilistic Approximation - Certainty Equivalent Control	p. 27
2.4. Rollout	p. 34
2.4.1. On-Line Rollout for Deterministic Discrete Optimization	p. 36
2.4.2. Stochastic Rollout and Monte Carlo Tree Search	p. 47
2.4.3. Rollout with an Expert	p. 56
2.5. On-Line Rollout for Deterministic Infinite-Spaces Problems - Optimization Heuristics	p. 58
2.5.1. Model Predictive Control	p. 59
2.5.2. Target Tubes and the Constrained Controllability Condition	p. 66
2.5.3. Variants of Model Predictive Control	p. 70
2.6. Notes and Sources	p. 72
 3. Parametric Approximation	
3.1. Approximation Architectures	p. 2
3.1.1. Linear and Nonlinear Feature-Based Architectures	p. 2
3.1.2. Training of Linear and Nonlinear Architectures	p. 9
3.1.3. Incremental Gradient and Newton Methods	p. 10
3.2. Neural Networks	p. 23
3.2.1. Training of Neural Networks	p. 26
3.2.2. Multilayer and Deep Neural Networks	p. 30
3.3. Sequential Dynamic Programming Approximation	p. 35
3.4. Q-factor Parametric Approximation	p. 36
3.5. Parametric Approximation in Policy Space by Classification	p. 39
3.6. Notes and Sources	p. 41
 4. Infinite Horizon Reinforcement Learning	
4.1. An Overview of Infinite Horizon Problems	p. 3
4.2. Stochastic Shortest Path Problems	p. 6
4.3. Discounted Problems	p. 16
4.4. Exact and Approximate Value Iteration	p. 21
4.5. Policy Iteration	p. 25
4.5.1. Exact Policy Iteration	p. 26
4.5.2. Optimistic and Multistep Lookahead Policy Iteration	p. 30
4.5.3. Policy Iteration for Q-factors	p. 32
4.6. Approximation in Value Space - Performance Bounds	p. 34

- 4.6.1. Limited Lookahead Performance Bounds p. 36
- 4.6.2. Rollout p. 39
- 4.6.3. Approximate Policy Iteration p. 43
- 4.7. Simulation-Based Policy Iteration with Parametric
- Approximation p. 46
- 4.7.1. Self-Learning and Actor-Critic Systems p. 46
- 4.7.2. A Model-Based Variant p. 48
- 4.7.3. A Model-Free Variant p. 50
- 4.7.4. Implementation Issues of Parametric Policy Iteration . . p. 52
- 4.8. Q-Learning p. 55
- 4.8.1. Optimistic Policy Iteration with Parametric Q-Factor
- Approximation - SARSA and DQN p. 56
- 4.9. Additional Methods - Temporal Differences p. 58
- 4.10. Exact and Approximate Linear Programming p. 69
- 4.11. Approximation in Policy Space p. 71
- 4.11.1. Training by Cost Optimization - Policy Gradient,
- Cross Entropy, and Random Search Methods p. 75
- 4.11.2. Expert Supervised Training p. 84
- 4.11.3. Approximate Policy Iteration, Rollout, and
- Approximation in Policy Space p. 86
- 4.12. Notes and Sources p. 89
- 4.13. Appendix: Mathematical Analysis p. 93
- 4.13.1. Proofs for Stochastic Shortest Path Problems p. 93
- 4.13.2. Proofs for Discounted Problems p. 99
- 4.13.3. Convergence of Exact and Optimistic
- Policy Iteration p. 99
- 4.13.4. Performance Bounds for One-Step Lookahead p. 101
- 4.13.5. Performance Bounds for Rollout p. 104
- 4.13.6. Performance Bounds for Approximate Policy
- Iteration p. 107

5. Aggregation

- 5.1. Aggregation Frameworks p.
- 5.2. Classical and Biased Forms of the Aggregate Problem p.
- 5.3. Bellman’s Equation for the Aggregate Problem p.
- 5.4. Algorithms for the Aggregate Problem p.
- 5.5. Some Examples p.
- 5.6. Spatiotemporal Aggregation for Deterministic Problems p.
- 5.7. Notes and Sources p.

- References** p.

- Index** p.

Preface

Turning to the succor of modern computing machines, let us renounce all analytic tools.

Richard Bellman [Bel57]

From a teleological point of view the particular numerical solution of any particular set of equations is of far less importance than the understanding of the nature of the solution.

Richard Bellman [Bel57]

In this book we consider large and challenging multistage decision problems, which can be solved in principle by dynamic programming (DP for short), 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*, and *neuro-dynamic programming*. We will use primarily the most popular name: reinforcement learning.

Our subject has benefited greatly from the interplay of ideas from optimal control and from artificial intelligence. 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 within the maze of competing ideas that constitute the current state of the art.

There are two general approaches for DP-based suboptimal control. The first is *approximation in value space*, where we approximate in some way the optimal cost-to-go function with some other function. The major alternative to approximation in value space is *approximation in policy space*, whereby we select the policy by using optimization over a suitably restricted class of policies, usually a parametric family of some form. In

some schemes these two types of approximation may be combined, aiming to capitalize on the advantages of both.

While we provide a substantial treatment of approximation in policy space, most of the book is focused on approximation in value space. Here, the control at each state is obtained by optimization of the cost over a limited horizon, plus an approximation of the optimal future cost. The latter cost, which we generally denote by \tilde{J} , is a function of the state where we may be. It may be computed by a variety of methods, possibly involving simulation and/or some given or separately derived heuristic/suboptimal policy. The use of simulation often allows for implementations that do not require a mathematical model, a major idea that has allowed the use of DP beyond its classical boundaries.

We discuss selectively four types of methods for obtaining \tilde{J} :

- (a) *Problem approximation*: Here \tilde{J} is the optimal cost function of a related simpler problem, which is solved by exact DP. Certainty equivalent control and enforced decomposition schemes are discussed in some detail.
- (b) *Rollout and model predictive control*: Here \tilde{J} is the cost function of some known heuristic policy. The needed cost values to implement a rollout policy are often calculated by simulation. While this method applies to stochastic problems, the reliance on simulation favors deterministic problems, including challenging combinatorial problems for which heuristics may be readily implemented. Rollout may also be combined with adaptive simulation and Monte Carlo tree search, which have proved very effective in the context of games such as backgammon, chess, Go, and others.

Model predictive control was originally developed for continuous-space optimal control problems that involve some goal state, e.g., the origin in a classical control context. It can be viewed as a specialized rollout method that is based on a suboptimal optimization for reaching a goal state.

- (c) *Parametric cost approximation*: Here \tilde{J} is chosen from within a parametric class of functions, including neural networks, with the parameters “optimized” or “trained” by using state-cost sample pairs and some type of incremental least squares/regression algorithm. Approximate policy iteration and its variants are covered in some detail, including several actor-critic schemes. These involve policy evaluation with temporal difference-based training methods, and policy improvement that may rely on approximation in policy space.
- (d) *Aggregation*: Here the cost function \tilde{J} is the optimal cost function of some approximation to the original problem, called aggregate problem, which has fewer states. The aggregate problem can be formulated in a variety of ways, and may be solved by using exact DP

techniques. Its optimal cost function is then used as \tilde{J} in a limited horizon optimization scheme. Aggregation may also be used to provide local improvements to parametric approximation schemes that involve neural networks or linear feature-based architectures.

We have adopted a gradual expository approach, which proceeds along four directions:

- (1) *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.
- (2) *From finite horizon to infinite horizon problems:* We first discuss finite horizon exact and approximate DP methodologies, which are intuitive and mathematically simple in Chapters 1-3. We then progress to infinite horizon problems in Chapters 4 and 5.
- (3) *From deterministic to stochastic models:* We often discuss separately deterministic and stochastic problems. The reason is that deterministic problems are simpler and offer special advantages for some of our methods.
- (4) *From model-based to model-free implementations:* Reinforcement learning methods offer a major potential benefit over classical DP approaches, which were practiced exclusively up to the early 90s: they can be implemented by using a simulator/computer model rather than a mathematical model. In our presentation, we first discuss model-based implementations, and then we identify schemes that can be appropriately modified to work with a simulator.

After the first chapter, each new class of methods is introduced as a more sophisticated or generalized version of a simpler method introduced earlier. Moreover, we illustrate some of the methods by means of examples, which should be helpful in providing insight into their use, but may also be skipped selectively and without loss of continuity.

The mathematical style of this book is somewhat different from the one of the author's DP books [Ber12], [Ber17], [Ber18a], and the 1996 neuro-dynamic programming (NDP) research monograph, written jointly with John Tsitsiklis [BeT96]. While we provide a rigorous, albeit short, mathematical account of the theory of finite and infinite horizon DP, 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).

Several of the methods that we present are often successful in practice, but have less than solid performance properties. This is a reflection of the

state of the art in the field: there are no methods that are guaranteed to work for all or even most problems, but there are enough methods to try on a given problem with a reasonable chance of success in the end. To aid in this process, we place primary emphasis on developing proper intuition into the inner workings of each type of method. Still, however, it is important to have a foundational understanding of the analytical principles of the field and of the mechanisms underlying the central computational methods. To quote a statement from the preface of the NDP monograph [BeT96]: “It is primarily through an understanding of the mathematical structure of the NDP methodology that we will be able to identify promising or solid algorithms from the bewildering array of speculative proposals and claims that can be found in the literature.”

Another statement from a recent NY Times article [Str18], in connection with DeepMind’s remarkable AlphaZero chess program, is also worth quoting: “What is frustrating about machine learning, however, is that the algorithms can’t articulate what they’re thinking. We don’t know why they work, so we don’t know if they can be trusted. AlphaZero gives every appearance of having discovered some important principles about chess, but it can’t share that understanding with us. Not yet, at least. As human beings, we want more than answers. We want insight. This is going to be a source of tension in our interactions with computers from now on.”[†] To this we may add that human insight can only develop within some structure of human thought, and it appears that mathematical reasoning with algorithmic models is the most suitable structure for this purpose.

Dimitri P. Bertsekas

May 2019

[†] The two 1957 Bellman quotations at the beginning of this preface also express this tension, although the first of these, while striking and widely cited, is admittedly taken a little out of context. Bellman’s fascinating and informative autobiography [Bel84] contains a lot of information on the origins of DP (selected quotations from this autobiography have been compiled by Dreyfus [Dre02]). Among others, Bellman states that “In order to make any progress, it is necessary to think of approximate techniques, and above all, of numerical algorithms. Finally, having devoted a great deal of time and effort, mostly fruitless, to the analysis of many varieties of simple models, I was prepared to face up to the challenge of using dynamic programming as an effective tool for obtaining numerical answers to numerical questions.” He goes on to attribute his motivation to work on numerical DP to the emergence of the (then primitive) digital computer, which he calls “Sorcerer’s Apprentice.”