

Dynamic Programming and Optimal Control

Table of Contents:

Volume 1: 3rd Edition

1. The Dynamic Programming Algorithm

1. Introduction
2. The Basic Problem
3. The Dynamic Programming Algorithm
4. State Augmentation and Other Reformulations
5. Some Mathematical Issues
6. Dynamic Programming and Minimax Control
7. Notes, Sources, and Exercises

2. Deterministic Systems and the Shortest Path Problem

1. Finite-State Systems and Shortest Paths
2. Some Shortest Path Applications
 1. Critical Path Analysis
 2. Hidden Markov Models and the Viterbi Algorithm
3. Shortest Path Algorithms
 1. Label Correcting Methods - A* Algorithm
 2. Branch-and-Bound
 3. Constrained and Multiobjective Problems
4. Notes, Sources, and Exercises

3. Deterministic Continuous-Time Optimal Control

1. Continuous-Time Optimal Control
2. The Hamilton-Jacobi-Bellman Equation
3. The Pontryagin Minimum Principle
 1. An Informal Derivation Using the HJB Equation

2. A Derivation Based on Variational Ideas
3. The Minimum Principle for Discrete-Time Problems
4. Extensions of the Minimum Principle
 1. Fixed Terminal State
 2. Free Initial State
 3. Free Terminal Time
 4. Time-Varying System and Cost
 5. Singular Problems
5. Notes, Sources, and Exercises

4. Problems with Perfect State Information

1. Linear Systems and Quadratic Cost
2. Inventory Control
3. Dynamic Portfolio Analysis
4. Optimal Stopping Problems
5. Scheduling and the Interchange Argument
6. Set-Membership Description of Uncertainty
 1. Set-Membership Estimation
 2. Control with Unknown-but-Bounded Disturbances
7. Notes, Sources, and Exercises

5. Problems with Imperfect State Information

1. Reductions to the Perfect Information Case
2. Linear Systems and Quadratic Cost
3. Minimum Variance Control of Linear Systems
4. Sufficient Statistics and Finite-State Markov Chains
5. Sequential Hypothesis Testing
6. Notes, Sources, and Exercises

6. Suboptimal and Adaptive Control

1. Certainty Equivalent and Adaptive Control
 1. Caution, Probing, and Dual Control
 2. Two-Phase Control and Identifiability
 3. Certainty Equivalent Control and Identifiability
 4. Self-Tuning Regulators

2. Open-Loop Feedback Control
3. Limited Lookahead Policies and Applications
 1. Performance Bounds of Limited Lookahead Policies
 2. Computational Issues in Limited Lookahead
 3. Problem Approximation -- Enforced Decomposition
 4. Aggregation
 5. Parametric Cost-to-Go Approximation
4. Rollout Algorithms
 1. Discrete Deterministic Problems
 2. Q-Factors Evaluated by Simulation
 3. Q-Factor Approximation
5. Model Predictive Control and Related Methods
 1. Rolling Horizon Approximations
 2. Stability Issues in Model Predictive Control
 3. Restricted Structure Policies
6. Additional Topics in Approximate DP
 1. Discretization
 2. Other Approximation Approaches
7. Notes, Sources, and Exercises

7. Introduction to Infinite Horizon Problems

1. An Overview
2. Stochastic Shortest Path Problems
3. Discounted Problems
4. Average Cost Problems
5. Semi-Markov Problems
6. Notes, Sources, and Exercises

8. Appendix A: Mathematical Review

9. Appendix B: On Optimization Theory

10. Appendix C: On Probability Theory

11. Appendix D: On Finite-State Markov Chains

12. Appendix E: Kalman Filtering

13. Appendix F: Modeling of Stochastic Linear Systems

14. Appendix G: Formulating Problems of Decision Under Uncertainty

Volume 2: 4th Edition

1. Discounted Problems - Theory

1. Minimization of Total Cost - Introduction
 1. The Finite-Horizon DP Algorithm
 2. Shorthand Notation and Monotonicity
 3. A Preview of Infinite Horizon Results
 4. The Finite-Horizon DP Algorithm
 5. Randomized and History-Dependent Policies
2. Discounted Problems - Bounded Cost per Stage
3. Scheduling and Multiarmed Bandit Problems
4. Discounted Continuous-Time Problems
5. The Role of Contraction Mappings
 1. Sup-Norm Contractions
 2. Discounted Problems - Unbounded Cost per Stage
6. General Forms of Discounted Dynamic Programming
 1. Basic Results Under Contraction and Monotonicity
 2. Discounted Dynamic Games
7. Notes, Sources, and Exercises

2. Discounted Problems - Computational Methods

1. Markovian Decision Problems
2. Value Iteration
 1. Monotonic Error Bounds for Value Iteration
 2. Variants of Value Iteration
 3. Q-Learning
3. Policy Iteration
 1. Policy Iteration for Costs
 2. Policy Iteration for Q-Factors
 3. Optimistic Policy Iteration
 4. Limited Lookahead Policies and Rollout
4. Linear Programming Methods
5. Methods for General Discounted Problems
 1. Limited Lookahead Policies and Approximations
 2. Generalized Value Iteration

- 3. Approximate Value Iteration
- 4. Generalized Policy Iteration
- 5. Generalized Optimistic Policy Iteration
- 6. Approximate Policy Iteration
- 7. Mathematical Programming
- 6. Asynchronous Algorithms
 - 1. Asynchronous Value Iteration
 - 2. Asynchronous Policy Iteration
 - 3. Policy Iteration with a Uniform Fixed Point
- 7. Notes, Sources, and Exercises

3. Stochastic Shortest Path Problems

- 1. Problem Formulation
- 2. Main Results
- 3. Underlying Contraction Properties
- 4. Value Iteration
 - 1. Conditions for Finite Termination
 - 2. Asynchronous Value Iteration
- 5. Policy Iteration
 - 1. Optimistic Policy Iteration
 - 2. Approximate Policy Iteration
 - 3. Policy Iteration with Improper Policies
 - 4. Policy Iteration with a Uniform Fixed Point
- 6. Countable State Spaces
- 7. Notes, Sources, and Exercises

4. Undiscounted Problems

- 1. Unbounded Costs per Stage
 - 1. Main Results
 - 2. Value Iteration
 - 3. Other Computational Methods
- 2. Linear Systems and Quadratic Cost
- 3. Inventory Control
- 4. Optimal Stopping
- 5. Optimal Gambling Strategies
- 6. Nonstationary and Periodic Problems
- 7. Notes, Sources, and Exercises

5. Average Cost per Stage Problems

1. Finite-Spaces Average Cost Models
 1. Relation with the Discounted Cost Problem
 2. Blackwell Optimal Policies
 3. Optimality Equations
2. Conditions for Equal Average Cost for all Initial States
3. Value Iteration
 1. Single-Chain Value Iteration
 2. Multi-Chain Value Iteration
4. Policy Iteration
 1. Single-Chain Policy Iteration
 2. Multi-Chain Policy Iteration
5. Linear Programming
6. Infinite-Spaces Problems
 1. A Sufficient Condition for Optimality
 2. Finite State Space and Infinite Control Space
 3. Countable States -- Vanishing Discount Approach
 4. Countable States -- Contraction Approach
 5. Linear Systems with Quadratic Cost
7. Notes, Sources, and Exercises

6. Approximate Dynamic Programming - Discounted Models

1. General Issues of Simulation-Based Cost Approximation
 1. Approximation Architectures
 2. Simulation-Based Approximate Policy Iteration
 3. Direct and Indirect Approximation
 4. Monte Carlo Simulation
 5. Simplifications
2. Direct Policy Evaluation - Gradient Methods
3. Projected Equation Methods for Policy Evaluation
 1. The Projected Bellman Equation
 2. The Matrix Form of the Projected Equation
 3. Simulation-Based Methods
 4. LSTD, LSPE, and TD(0) Methods
 5. Optimistic Versions
 6. Multistep Simulation-Based Methods

- 7. A Synopsis
- 4. Policy Iteration Issues
 - 1. Exploration Enhancement by Geometric Sampling
 - 2. Exploration Enhancement by Off-Policy Methods
 - 3. Policy Oscillations - Chattering
- 5. Aggregation Methods
 - 1. Cost Approximation via the Aggregate Problem
 - 2. Cost Approximation via the Enlarged Problem
 - 3. Multistep Aggregation
 - 4. Asynchronous Distributed Aggregation
- 6. Q-Learning
 - 1. Q-Learning: A Stochastic VI Algorithm
 - 2. Q-Learning and Policy Iteration
 - 3. Q-Factor Approximation and Projected Equations
 - 4. Q-Learning for Optimal Stopping Problems
 - 5. Q-Learning and Aggregation
 - 6. Finite Horizon Q-Learning
- 7. Notes, Sources, and Exercises

7. Approximate Dynamic Programming - Nondiscounted Models and Generalizations

- 1. Stochastic Shortest Path Problems
- 2. Average Cost Problems
 - 1. Approximate Policy Evaluation
 - 2. Approximate Policy Iteration
 - 3. Q-Learning for Average Cost Problems
- 3. General Problems and Monte Carlo Linear Algebra
 - 1. Projected Equations
 - 2. Matrix Inversion and Iterative Methods
 - 3. Multistep Methods
 - 4. Extension of Q-Learning for Optimal Stopping
 - 5. Equation Error Methods
 - 6. Oblique Projections
 - 7. Generalized Aggregation
 - 8. Deterministic Methods for Singular Linear Systems

9. Stochastic Methods for Singular Linear Systems
4. Approximation in Policy Space
 1. The Gradient Formula
 2. Computing the Gradient by Simulation
 3. Essential Features for Gradient Evaluation
 4. Approximations in Policy and Value Space
5. Notes, Sources, and Exercises

8. Appendix A: Measure-Theoretic Issues in Dynamic Programming

1. A Two-Stage Example
2. Resolution of the Measurability Issues