

Contents

| | | |
|---|-------------|------------|
| <i>Preface</i> | <i>page</i> | <i>xi</i> |
| <i>Acknowledgments</i> | | <i>xii</i> |
| <i>List of Symbols</i> | | <i>xiv</i> |
| Part I Background in Control and Reinforcement Learning | | 1 |
| 1 Introduction | | 3 |
| 1.1 What Is This Book About? | | 3 |
| 1.2 The Origins of Control | | 4 |
| 1.3 Formalization and Examples of a Control Problem | | 5 |
| 1.4 Simple Control Algorithms | | 6 |
| 1.5 Classical Theory: Optimal and Robust Control | | 8 |
| 1.6 The Need for a New Theory | | 9 |
| 1.7 Bibliographic Remarks | | 12 |
| 2 Dynamical Systems | | 13 |
| 2.1 Examples of Dynamical Systems | | 13 |
| 2.2 Solution Concepts for Dynamical Systems | | 17 |
| 2.3 Intractability of Equilibrium, Stabilizability, and Controllability | | 20 |
| 2.4 Bibliographic Remarks | | 22 |
| 2.5 Exercises | | 23 |
| 3 Markov Decision Processes | | 25 |
| 3.1 Reinforcement Learning | | 25 |
| 3.2 Markov Decision Processes | | 28 |
| 3.3 The Bellman Equation | | 34 |

| | | |
|----------|---|-----------|
| 3.4 | The Value Iteration Algorithm | 36 |
| 3.5 | Bibliographic Remarks | 38 |
| 3.6 | Exercises | 39 |
| 4 | Linear Dynamical Systems | 40 |
| 4.1 | Approximating General Dynamics via LTV Systems | 41 |
| 4.2 | Stabilizability of Linear Systems | 41 |
| 4.3 | Controllability of Linear Dynamical Systems | 44 |
| 4.4 | Quantitative Definitions | 46 |
| 4.5 | Bibliographic Remarks | 48 |
| 4.6 | Exercises | 48 |
| 5 | Optimal Control of Linear Dynamical Systems | 50 |
| 5.1 | The Linear Quadratic Regulator | 50 |
| 5.2 | Optimal Solution of the LQR | 51 |
| 5.3 | Infinite-Horizon LQR | 53 |
| 5.4 | Robust Control | 55 |
| 5.5 | Bibliographic Remarks | 56 |
| 5.6 | Exercises | 56 |
| | Part II Basics of Online Control | 57 |
| 6 | Regret in Control | 59 |
| 6.1 | Online Convex Optimization | 60 |
| 6.2 | Regret for Control | 62 |
| 6.3 | Expressivity of Control Policy Classes | 64 |
| 6.4 | Bibliographic Remarks | 69 |
| 6.5 | Exercises | 70 |
| 7 | Online Nonstochastic Control | 71 |
| 7.1 | From Optimal and Robust to Online Control | 72 |
| 7.2 | The Online Nonstochastic Control Problem | 73 |
| 7.3 | The Gradient Perturbation Controller | 75 |
| 7.4 | Bibliographic Remarks | 79 |
| 7.5 | Exercises | 80 |
| 8 | Online Nonstochastic System Identification | 82 |
| 8.1 | The Role of System Identification in Online Control | 82 |
| 8.2 | A Nonstochastic Approach to System Identification | 83 |
| 8.3 | Nonstochastic System Identification | 84 |
| 8.4 | From Identification to Nonstochastic Control | 87 |
| 8.5 | Summary and Key Takeaways | 87 |

| | | |
|--|--|-----|
| 8.6 | Bibliographic Remarks | 88 |
| 8.7 | Exercises | 89 |
| Part III Learning and Filtering | | 91 |
| 9 | Learning in Unknown Linear Dynamical Systems | 93 |
| 9.1 | Learning in Dynamical Systems | 93 |
| 9.2 | Online Learning of Dynamical Systems | 94 |
| 9.3 | Classes of Prediction Rules | 95 |
| 9.4 | Efficient Learning of Linear Predictors | 98 |
| 9.5 | Summary | 99 |
| 9.6 | Bibliographic Remarks | 100 |
| 9.7 | Exercises | 100 |
| 10 | Kalman Filtering | 102 |
| 10.1 | Observable Systems | 103 |
| 10.2 | Optimality of the Kalman Filter | 104 |
| 10.3 | Bayes-Optimality under Gaussian Noise | 108 |
| 10.4 | Conclusion | 109 |
| 10.5 | Bibliographic Remarks | 110 |
| 10.6 | Exercises | 111 |
| 11 | Spectral Filtering | 113 |
| 11.1 | Spectral Filtering in One Dimension | 114 |
| 11.2 | Spectral Predictors | 116 |
| 11.3 | Online Spectral Filtering | 120 |
| 11.4 | Conclusion | 121 |
| 11.5 | Bibliographic Remarks | 122 |
| 11.6 | Exercises | 123 |
| Part IV Online Control with Partial Observation | | 125 |
| 12 | Policy Classes for Partially Observed Systems | 127 |
| 12.1 | Linear Observation Action Controllers | 128 |
| 12.2 | Linear Dynamic Controllers | 130 |
| 12.3 | Disturbance Response Controllers | 131 |
| 12.4 | Summary | 134 |
| 12.5 | Bibliographic Remarks | 135 |
| 12.6 | Exercises | 136 |

| | | |
|-----------|--|------------|
| 13 | Online Nonstochastic Control with Partial Observation | 138 |
| 13.1 | The Gradient Response Controller | 139 |
| 13.2 | Extension to Time-Varying Systems | 142 |
| 13.3 | Conclusion | 143 |
| 13.4 | Bibliographic Remarks | 143 |
| 13.5 | Exercises | 144 |
| | <i>References</i> | 147 |
| | <i>Index</i> | 155 |