

1. Record Nr.	UNISA996547957103316
Autore	Li Shengbo Eben
Titolo	Reinforcement learning for sequential decision and optimal control // Shengbo Eben Li
Pubbl/distr/stampa	Singapore : , : Springer, , [2023] ©2023
ISBN	9789811977848 9789811977831
Edizione	[1st ed. 2023.]
Descrizione fisica	1 online resource (484 pages)
Disciplina	006.31
Soggetti	Reinforcement learning Control theory - Data processing
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	Chapter 1 Introduction of Reinforcement Learning -- Chapter 2 Principles of RL Problems -- Chapter 3 Model-free Indirect RL: Monte Carlo -- Chapter 4 Model-Free Indirect RL: Temporal-Difference -- Chapter 5 Model-based Indirect RL: Dynamic Programming -- Chapter 6 Indirect RL with Function Approximation -- Chapter 7 Direct RL with Policy Gradient -- Chapter 8 Infinite Horizon Approximate Dynamic Programming -- Chapter 9 Finite Horizon ADP and State Constraints -- Chapter 10 Deep Reinforcement Learning -- Chapter 11 Advanced RL Topics.
Sommario/riassunto	Have you ever wondered how AlphaZero learns to defeat the top human Go players? Do you have any clues about how an autonomous driving system can gradually develop self-driving skills beyond normal drivers? What is the key that enables AlphaStar to make decisions in Starcraft, a notoriously difficult strategy game that has partial information and complex rules? The core mechanism underlying those recent technical breakthroughs is reinforcement learning (RL), a theory that can help an agent to develop the self-evolution ability through continuing environment interactions. In the past few years, the AI community has witnessed phenomenal success of reinforcement learning in various fields, including chess games, computer games and robotic control. RL

is also considered to be a promising and powerful tool to create general artificial intelligence in the future. As an interdisciplinary field of trial-and-error learning and optimal control, RL resembles how humans reinforce their intelligence by interacting with the environment and provides a principled solution for sequential decision making and optimal control in large-scale and complex problems. Since RL contains a wide range of new concepts and theories, scholars may be plagued by a number of questions: What is the inherent mechanism of reinforcement learning? What is the internal connection between RL and optimal control? How has RL evolved in the past few decades, and what are the milestones? How do we choose and implement practical and effective RL algorithms for real-world scenarios? What are the key challenges that RL faces today, and how can we solve them? What is the current trend of RL research? You can find answers to all those questions in this book. The purpose of the book is to help researchers and practitioners take a comprehensive view of RL and understand the in-depth connection between RL and optimal control. The book includes not only systematic and thorough explanations of theoretical basics but also methodical guidance of practical algorithm implementations. The book intends to provide a comprehensive coverage of both classic theories and recent achievements, and the content is carefully and logically organized, including basic topics such as the main concepts and terminologies of RL, Markov decision process (MDP), Bellman's optimality condition, Monte Carlo learning, temporal difference learning, stochastic dynamic programming, function approximation, policy gradient methods, approximate dynamic programming, and deep RL, as well as the latest advances in action and state constraints, safety guarantee, reference harmonization, robust RL, partially observable MDP, multiagent RL, inverse RL, offline RL, and so on.
