

1. Record Nr.	UNINA9910467009503321
Autore	Lapan Maxim
Titolo	Deep reinforcement learning hands-on : apply modern RL methods, with deep Q-networks, value iteration, policy gradients, TRPO, AlphaGo Zero and more // Maxim Lapan
Pubbl/distr/stampa	Birmingham, England : , : Packt Publishing, , 2018
ISBN	1-78883-930-7
Edizione	[1st edition]
Descrizione fisica	1 online resource (1 volume) : illustrations
Disciplina	006.31
Soggetti	Reinforcement learning Electronic books.
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	"Expert insight."
Nota di bibliografia	Includes bibliographical references and index.
Sommario/riassunto	This practical guide will teach you how deep learning (DL) can be used to solve complex real-world problems. About This Book Explore deep reinforcement learning (RL), from the first principles to the latest algorithms Evaluate high-profile RL methods, including value iteration, deep Q-networks, policy gradients, TRPO, PPO, DDPG, D4PG, evolution strategies and genetic algorithms Keep up with the very latest industry developments, including AI-driven chatbots Who This Book Is For Some fluency in Python is assumed. Basic deep learning (DL) approaches should be familiar to readers and some practical experience in DL will be helpful. This book is an introduction to deep reinforcement learning (RL) and requires no background in RL. What You Will Learn Understand the DL context of RL and implement complex DL models Learn the foundation of RL: Markov decision processes Evaluate RL methods including Cross-entropy, DQN, Actor-Critic, TRPO, PPO, DDPG, D4PG and others Discover how to deal with discrete and continuous action spaces in various environments Defeat Atari arcade games using the value iteration method Create your own OpenAI Gym environment to train a stock trading agent Teach your agent to play Connect4 using AlphaGo Zero Explore the very latest deep RL research on topics including AI-driven chatbots In Detail Recent developments in

reinforcement learning (RL), combined with deep learning (DL), have seen unprecedented progress made towards training agents to solve complex problems in a human-like way. Google's use of algorithms to play and defeat the well-known Atari arcade games has propelled the field to prominence, and researchers are generating new ideas at a rapid pace. Deep Reinforcement Learning Hands-On is a comprehensive guide to the very latest DL tools and their limitations. You will evaluate methods including Cross-entropy and policy gradients, before applying them to real-world environments. Take on both the Atari set of virtual games and family favorites such as Connect4. The book provides an introduction to the basics of RL, giving you the know-how to code intelligent learning agents to take on a formidable array of practical tasks. Discover how to implement Q-learning on 'grid world' environments, teach your agent to buy and trade stocks, and find out how natural language models are driving the boom in chatbots. Style and approach Deep Reinforcement Learning Hands-On explains the art of building self-learning agents using algorith...
