

1. Record Nr.	UNINA9910874683203321
Autore	Sanghi Nimish
Titolo	Deep Reinforcement Learning with Python : RLHF for Chatbots and Large Language Models // by Nimish Sanghi
Pubbl/distr/stampa	Berkeley, CA : , : Apress : , : Imprint : Apress, , 2024
ISBN	9798868802737 9798868802720
Edizione	[2nd ed. 2024.]
Descrizione fisica	1 online resource (0 pages)
Disciplina	005.133
Soggetti	Python (Computer program language) Natural language processing (Computer science) Artificial intelligence - Computer programs
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Includes index.
Nota di contenuto	Chapter 1: Introduction to Reinforcement Learning -- Chapter 2: The Foundation – Markov Decision Processes -- Chapter 3: Model Based Approaches -- Chapter 4: Model Free Approaches -- Chapter 5: Function Approximation and Deep Reinforcement Learning -- Chapter 6: Deep Q-Learning (DQN) -- Chapter 7: Improvements to DQN -- Chapter 8: Policy Gradient Algorithms -- Chapter 9: Combining Policy Gradient and Q-Learning -- Chapter 10: Integrated Planning and Learning -- Chapter 11: Proximal Policy Optimization (PPO) and RLHF -- Chapter 12: Introduction to Multi Agent RL (MARL) -- Chapter 13: Additional Topics and Recent Advances.
Sommario/riassunto	Gain a theoretical understanding of the most popular libraries in deep reinforcement learning (deep RL). This new edition focuses on the latest advances in deep RL using a learn-by-coding approach, allowing readers to assimilate and replicate the latest research in this field. New agent environments ranging from games, and robotics to finance are explained to help you try different ways to apply reinforcement learning. A chapter on multi-agent reinforcement learning (MARL) covers how multiple agents can be trained, while another chapter focuses on the widely used deep RL algorithm, proximal policy optimization (PPO). You'll see how reinforcement learning with human

feedback (RLHF) has been used to fine-tune Large Language Models (LLMs) to chat and follow instructions. An example of this is the OpenAI ChatGPT offering human like conversational capabilities. You'll also review the steps for using the code on multiple cloud systems and deploying models on platforms such as Hugging Face Hub. The code is in Jupyter Notebook, which can be run on Google Colab, and other similar deep learning cloud platforms, allowing you to tailor the code to your own needs. Whether it's for applications in gaming, robotics, or Generative AI, Deep Reinforcement Learning with Python will help keep you ahead of the curve.

---