Reinforcement Learning

Oliver Wallscheid

Lecture

1. Introduction to Reinforcement Learning
2. Markov Decision Processes
3. Dynamic Programming
4. Monte Carlo Methods
5. Temporal-Difference Learning
6. n-Step Bootstrapping
7. Planning and Learning with Tabular Methods
8. Reinforcement Learning in Finite State and Action Spaces
9. Function Approximation with Supervised Learning
10. On-Policy Prediction with Function Approximation
11. Value-Based Control with Function Approximation
12. Eligibility Traces
13. Policy Gradient Methods
14. Further Contemporary RL Algorithms
15. Course Completion and Outlook

Tutorial

1. Python Tutorial
2. Markov Decision Processes
3. Dynamic Programming
4. Monte Carlo Methods
5. Temporal-Difference Learning
6. n-Step Methods

7. Learning and Planning
8. Function Approximation with Supervised Learning
9. Function Approximators in Value Prediction
10. Function Approximators in Control
11. Eligibility Traces
12. Policy Gradients
13. DDPG & PPO

Materials

May 7, 2025