

Course title: Brief Introduction to Reinforcement Learning

Length: 4 hours

This intensive half-day session gives participants a sharp and motivating introduction to Reinforcement Learning (RL). The course focuses on building strong intuition, demonstrating real-world relevance, and showing how RL fits into the modern AI landscape. Ideal for professionals, students, and teams who want to quickly understand what RL is—and why it matters.

After this session, participants will be able to:

- Explain the core concepts of RL (agent, environment, policy, rewards, value).
- Describe how RL differs from supervised and unsupervised learning.
- Understand the intuition behind exploration vs. exploitation.
- Identify common real-world applications where RL creates value.
- Evaluate whether RL is relevant for their own projects or domain.
- Interpret simple RL diagrams and algorithm workflows.

Course title: Reinforcement Learning Explained

Length: 10 × 4 hours

The full journey, from fundamentals to modern, practical RL. This comprehensive 10 half-days course follows and expands on the content from the book Reinforcement Learning Explained. Designed for individuals and teams who want a deep, structured understanding of RL, it blends theory, intuitive explanations, hands-on examples, and implementation insights. Participants will learn classical and modern RL methods and understand when and how to use them effectively.

The course delivers the greatest value when participants complete the exercises between sessions. True RL knowledge can only be achieved by hands on training.

After completing the course, participants will be able to:

- Understand all core RL concepts, from value functions to policy optimization.
- Implement fundamental RL algorithms (e.g., TD learning, Q-learning, Monte Carlo methods).
- Explain modern approaches such as deep RL, policy gradients, actor–critic, and MCTS.
- Analyze and compare different RL strategies for specific problem types.
- Design RL environments and reward structures for real applications.
- Debug common RL issues such as instability and sub-optimal policies.
- Evaluate RL performance and interpret training curves and metrics.
- Understand the computational and practical challenges of deploying RL.