



UNIVERSITÀ
DEGLI STUDI DI MILANO-BICOCCA

SYLLABUS DEL CORSO

Reinforcement Learning

2627-2-FDS02Q042

Aims

This course aims to develop students' knowledge, skills, and competences in reinforcement learning (RL), preparing them for advanced research and applications in data-driven decision-making. Upon completion, students will:

- Acquire a deep understanding of RL theories and algorithms (DdD 1)
- Apply RL methods to model and solve sequential decision problems (DdD 2)
- Develop critical judgment in evaluating RL approaches and their suitability (DdD 3)
- Communicate technical concepts and findings clearly in oral and written form (DdD 4)
- Foster autonomous learning to stay current with evolving RL research and applications (DdD 5)

Contents

- **Contents:**
- Fundamentals of Reinforcement Learning and Markov Decision Processes (MDPs)
- Dynamic programming techniques
- Monte Carlo methods and Temporal Difference learning
- Model-free methods: Q-learning and SARSA
- Exploration vs exploitation trade-offs
- Policy Gradient approaches (introduction)
- Function approximation and an introduction to deep reinforcement learning
- Practical RL applications in Data Science

Detailed program

Introduction to RL and MDPs
Dynamic Programming
Monte Carlo Methods
Temporal Difference Learning
Q-learning and SARSA
Exploration vs Exploitation
Policy Gradient Methods (Intro)
Function Approximation
Deep RL Overview (Guest lecture and videos)
RL Applications in Data Science (Project proposal presentations)

Prerequisites

Prerequisites

- Probability and statistics fundamentals (DdD 1)
- Machine learning basics (DdD 1)
- Programming in Python, including familiarity with scientific libraries (DdD 2)
- Calculus and linear algebra knowledge (DdD 1)

Teaching form

Teaching Form

Lectures (focus on knowledge and understanding - DdD 1): 21 hours (18 lecture based, 3 interactive mode)

Coding labs and exercises (applying knowledge - DdD 2): 27 hours (18 lecture based, 9 interactive mode)

Up to 6 hours can be online if needed.

Group discussions and case studies (developing judgment - DdD 3)

Project presentations (communication skills - DdD 4)

Encouragement of autonomous study through research paper reviews and self-directed projects (learning skills - DdD 5)

Textbook and teaching resource

Textbook and Teaching Resources

- Main textbook: Sutton, R. S., & Barto, A. G. (2018). Reinforcement Learning: An Introduction (2nd Edition). Available online: <http://incompleteideas.net/book/the-book-2nd.html> (DdD 1, DdD 5)
- Additional Textbook: Shie Mannor, Yishay Mansour and Aviv Tamar, Reinforcement Learning: Foundations, Cambridge University Press, 2026 (DdD 1, DdD 5)
- Supplementary academic papers and online tutorials (DdD 5)
- Jupyter notebooks and coding environments (DdD 2)

Semester

I semester

Assessment method

Final exam (written/oral) on theory and applications (70%) – evaluates knowledge, judgment, and communication (DdD 1, 3, 4)

the exam consist of exercices to evaluate the ability in applying the concept on specific problems and open questions to evaluate knowledge and argumentaton ability on the subject.

Group project and final report (30%) – develops communication and autonomous learning skills (DdD 4, 5)

No partial exams will be held.

Office hours

By appointment

Sustainable Development Goals

INDUSTRY, INNOVATION AND INFRASTRUCTURE
