# Reinforcement Learning course

## Presentation of the course and Class info


Mario Martin

Mario Martin - UPC


12/2/2026

# Reinforcement Learning Course

## Instructor

- Mario Martín (UPC Professor)
- email: mmartin@cs.upc.edu
- Office: Omega Building, room 202 (second floor).

## Materials

Web for the course on
https://www.cs.upc.edu/~mmartin/RL-MAI.html

# Logistics of the course

- Lectures will be face to face on Fridays 11:00-14:00 (room B4002).
- We will not change the room for the Labs. Usually Labs will consist in understanding and running some notebooks.

# Syllabus

**First part:**

1. Basic concepts of Reinforcement Learning
2. Model based methods for policy learning
3. Basic RL algorithms: Model free methods
    - Monte Carlo, Q-learning and Temporal Differences: TD($\lambda$)
4. Function approximation in RL and Deep Learning
5. Policy gradient methods
    - REINFORCE and Actor-Critic methods
    - TRPO and PPO
    - DDPG, TD3, SAC

# Syllabus

**Second part:**

6. Inverse Reinforcement Learning and RLHF
7. Monte Carlo Tree Search and Alfa-Zero algorithms
8. Sample efficiency I: Model Based Reinforcement Learning (MBRL)
9. Sample efficiency II: Sparse Rewards problem
    - Exploration, Curricular Learning, Hierarchical Learning, Hindsight Experience Replay (HER)
10. Sample efficiency III: Multiple task and life-long learning
    - Transfer learning, Multi-task learning, Meta-learning
11. Multi-agent RL
    - RL in Cooperative, Competitive and Mixed cases problems

# Resources

- Slides: On the web page of the curse (Also in the old page, with some support videos).
- Books:
  - Sutton and Barto, *An Introduction to Reinforcement Learning*, 2nd Edition (2018). MIT Press. Available here.
  - Miguel Morales *Grokking Deep Reinforcement Learning* Manning, 2020.
- Recommended courses with materials:
  - (Basic) David Silver's course Reinforcement Learning, 2015.
  - (Advanced) Sergey Levine's course CS285 (Berkeley): Deep Reinforcement Learning, Fall 2021.
- Notebooks: See Lab section in web page of the course.

# Course Evaluation

The evaluation of the course will consist on three parts:

1. A test questionnaire about the topics of the course
2. Implementation of a domain and/or a reinforcement learning algorithm (f.i. in the OpenGym framework and python).
3. A research paper about the current state of the art of one topic related to RL.

Final grade will be resulting of this formula:

$$\text{Grade} = 0,3 * \text{Test} + 0,3 * \text{ Implementation} + 0,4 * \text{Paper}$$

# Course Evaluation: Quiz

- Test will be about basics of RL that will include the topics covered in the first part of the course (see syllabus)
- The test will done the **April 10th**.

# Course Evaluation: Implementation

- Work could consist in implementing a non trivial algorithm, environment or a exploration technique explained in class
- You are free to choose the task and the algorithm
- The implementation must be in the python language and one of the platforms proposed.
- You will have to write short report about the implementation and results.
- The deadline for delivering this work will be **May 11th 2026**.

# Course Evaluation: Paper

- You will have to write brief paper about the current state of the art of the research in a advanced topic covered in the second part of the course.
  1. Choose one topic of your interest and collect relevant bibliography on it
  2. Choose 2-3 related relevant interesting papers on the topic
  3. You will have to summarize the approaches presented on those papers and compare and criticise them the relation of the topic with other areas
- Send me and e-mail with the papers you selected. Each student should work on different papers and will be assigned using the first-to-choose-first-to-assign policy. You can work on the topic only *after my approval*.
- The deadline for delivering this report is **June 19th 2026**.

# Lab classes

- Any `OS` is Ok.
- I will assume knowledge on `python` and `Pytorch`
- Don't worry. We will install the needed packages to play with different algorithms and environments
- In case you don't have a GPU in your laptop don't worry because we will use also `Google Colab`

# Caveat: Applications of RL

- Some real world successful applications in RL:
  - Robotics: walking, manipulation, etc.
  - Autonomous driving.
  - Cooperation in a multi-agent team of marine UV.
  - Control of plasma in fusion reactors.
  - Adaptive optics in extra large telescopes.
  - Control of water or energy.
  - Medicine: Control of ventilation in Covid times or medication in Sepsis.
  - Design of chips (claimed superhuman level).
  - Design of medical drugs.
  - Chatbot (f.i. ChatGPT) training!
  - Design of boats in America's Cup.