ATCI: Reinforcement Learning

0. Class info

Mario Martin

CS-UPC

February 14, 2025

Contact information

Advanced Techniques in Computational Intelligence (previously second part of the URL course):

- Mario Martin
- Office 202, second floor Omega Building.
- email: mmartin@cs.upc.edu
- Web page for course: https://www.cs.upc.edu/~mmartin/ATCI-RL.html

Logistics of the course

- Lectures will be face to face on Fridays 11:00-14:00 (room B4002).
- There are videos of the lectures in the web page of the course form last edition (they don't change a lot and I will warn you if they change).
- 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. Basic RL algorithms: Model based methods
   - Value iteration and Policy iteration
   - Asynchronous versions
3. Basic RL algorithms: Model free methods
   - Backups and Role of exploration
   - Monte Carlo, Q-learning and Sarsa
   - Temporal Differences: TD(λ)
4. Function approximation in RL
   - Non-parametric methods and linear methods
   - Deep Learning for RL
5. Policy gradient methods
   - Actor-Critic methods
   - REINFORCE and TRPO
   - 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.
- 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.
- Software: 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:

$$\text{Grade} = 0,20 * \text{Test} + 0,30 * \text{ Implementation} + 0,5 * \text{Paper}$$

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 **approximately** during week **beginning April**.

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 is approx. about **May 11th 2025**.

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. You will have to look for papers related to the topic
  2. You will have to collect relevant bibliography on the topic
  3. You will have to choose the more relevant papers
  4. You will have to summarize the problems described and to summarize the approaches presented on those papers
  5. You will have to situate the problem in the area and to describe the relation of the topic with other areas
- You will find a list of topics in the web page of the curse. Choose the topic and send me and e-mail. Each student should choose a different topic. Topics will be assigned using the first-to-choose-first-to-assign policy.
- You can work also on a topic not in the list *after my approval*.
- The deadline for delivering this report is **June 19th 2025**.

Lab classes

- Any `OS` is Ok.
- I will assume some knowledge of `python`. I will assume also you have `python 3.x` installed in your computer. Install Anaconda distribution if you don't have any installed.
- Don't worry. We will install the different platforms to play with different algorithms and environments. Instructions will be available also in the web page of the course.