Syllabus
Reinforcement Learning, 6hp

Issued by the WASP graduate school management group *draft*.

Main field of study

AI/mlx

Course level

PhD Student course

Course offered for

PhD Students in the WASP graduate school

Entry requirements

The participants are assumed to have a background in mathematics corresponding to the contents of the WASP-course "Mathematics and Machine Learning". Other entry requirements for the course:

- Basic knowledge about RL (MDP, tabular RL, value iteration, policy search, Q-learning, SARSA, etc.) corresponding to the optional course segment provided in this course.
- Probability theory (estimation, Monte Carlo, etc.)
- Optimization (mean-squares error, categorical cross entropy loss, dynamic programming)
- Deep learning basics (fully-connected and convolutional layers)
- Programming in Python (Numpy, plotting, deep learning with Python, etc.)

Intended learning outcomes

Knowledge and Understanding

After completed studies, the student shall be able to

- explain and characterize the concept of RL and categorize RL agents
- describe, explain, compare, and characterize different types of basic and advanced reinforcement learning methods
- derive from first principles and explain what the underlying mathematical principles of these reinforcement learning methods are
- restate a control problem as an RL problem

Competences and Skills

After completed studies, the student shall be able to

- analyze and compare results of reinforcement learning methods
- implement (relevant parts of) advanced reinforcement learning algorithms
- apply advanced reinforcement learning algorithms
- read and critically review scientific publications about reinforcement learning
- use OpenAI gym and deep learning libraries to implement (relevant parts of) RL algorithms

Judgment and Approach

After completed studies, the student shall be able to

- discuss and reflect on important and advanced concepts in reinforcement learning
- discuss and reflect on what influences the performance of these methods
- discuss and reflect on when which of these methods applies
- discuss and reflect on scientific publications about reinforcement learning
- propose extensions and modifications to improve the performance of an RL algo-rithm for a specific problem

Course content

The course is organized in the sections described below.

Section 0 – Introduction to Basic RL and Control

This course section is provided for the students who do not fulfill the necessary requirements and is optional, with an exception made for the Quiz (Unit 7). The video lectures are taken from a SMART(er) course in reinforcement learning at Örebro University.

- RL foundations
- Dynamic Programming
- Monte Carlo Methods
- Temporal-difference Learning
- Planning with a Model and Learning
- Public Perception of RL and RL in Media
- Quiz (mandatory): Recap on Control and Reinforcement Learning Basics
- Lab and exercises

Section 1 – Temporal-difference Learning in Continuous Spaces

- RL with function approximation
- Temporal-difference Learning in Continuous Spaces
- Temporal-difference learning for Linear Quadratic (LQ) problem)
- Lab and exercises

Section 2 - Policy Search in Continuous Spaces
- Policy gradient methods
- PG for MDPs with continuous action space
- Improving PG
- Actor-critic methods
- Lab and exercises

Section 3 - Methods with Model-learning
- Model-based Policy Search
- Adaptive control

Section 4 - Other Advanced Topics in RL
- (Deep) RL and Reproducible
- Elephant in the Room: Critical Opinion on Deep RL
- RL from the control perspective
- Multi-objective RL
- Monte Carlo Tree Search
- Guided Policy Search

The course includes three 2-day meetings with intense teaching on-site.

Teaching and working methods

Video lectures: The students watch pre-recorded video lectures.
Self-study reading: The students read assigned papers, book chapters, or other material.
Reflective learning journal writing: The students write entries into a reflective learning journal
Peer feedback and group discussions: The students share their works or results and give each other feedback (within a learning group).
Labs: The students have computer-based sessions and do coding and evaluation in groups. Teaching assistants will provide formative feedback on the lab reports.
Presentations: The students work in groups and present recent RL papers. The stu-ents can select from the list given in this document.
Exercise and questions sessions: To practice the materials covered in each section of the course, the students are provided with theoretical exercises. The answers or possible solutions to exercises are discussed in synchronous sessions (possible at different times due to the number of students). In these sessions, the students can also ask their questions about the course materials or exercises. Exercises are not graded.

Examination

Reflective learning journal hand-in: The students hand in their reflective learning journal and it is graded according to a grading rubric at the end of the course.
Presentation: The students work in groups, present recent RL papers and get graded according to a grading rubric.
Lab: The students do practical work in computer-based sessions and hand in a report. The report will be graded according to a grading rubric.
Quiz: The students do a quiz on the learning platform.

The course allows one single retry for the assessment tasks at a date 6 months after the course concluded. The assessment tasks might be altered for the retry.

Grades

Fail or Pass