Syllabus
Reinforcement Learning, 6hp

Issued by the WASP graduate school management group: September 02, 2024

Main field of study

AI/MLX

Course level

PhD Student course

- AS track: elective
- AI track: prioritized (elective)
- Joint curriculum: advanced

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:

- Probability theory (estimation, Monte Carlo, etc.)
- Optimization (mean-squares error, categorical cross entropy loss, dynamic programming)
- Deep learning basics (back-propagation, fully connected, convolutional layers)
- Programming in Python (Numpy, plotting, deep learning with Python, etc.)
- Basic knowledge about RL (MDP, tabular RL, value iteration, policy search, Q-learning, SARSA, etc.) corresponding to module 1 of this course, in case module combination 2+3+4 is selected instead of module combination 1+2+3.

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, and
- 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 established software, frameworks, and libraries to implement (relevant parts of) RL algorithms and environments.

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 to a given scenario or problem,
- discuss and reflect on scientific publications about reinforcement learning,
- propose extensions and modifications to improve the performance of an RL algorithm for a specific problem.

Course content

The course is organized into 4 sequential modules that build on each other. Students take 3 of these modules in sequence, i.e. either module combination 1+2+3 or module combination 2+3+4. Combination 1+2+3 is for students without prior knowledge of RL while combination 2+3+4 is for students with prior knowledge corresponding to module 1.

Module 1 – Introduction to Basic RL and Control

- RL foundations
- Dynamic Programming
- Monte Carlo Methods
- Tabular temporal-difference learning
- Planning with a Model and Learning
- Public Perception of RL and RL in Media
- Control and Reinforcement Learning Basics
- Basic RL with function approximation
- Basic policy gradient methods
- Lab and exercises

Module 2 – Deep RL and control-based methods part 1

- Deep temporal-difference learning in discrete actions
- Deep temporal-difference learning with continuous actions
- Temporal-difference learning for Linear Quadratic (LQ) problem)
- Deep policy gradient methods
- Maximum entropy RL
- Lab and exercises

Module 3 – Deep RL and control-based methods part 2

- Deep actor-critic methods
- Model-based policy search
- Monte Carlo tree search
- RL with constraints
- Critical reflection about RL research
- Lab and exercises

Module 4 – Advanced topics in RL

- Selected advanced methods, e.g., multiple objectives, hierarchical RL, multiple agents, uncertainty, transfer.
- Outlook on RL research
- Lab and exercises

The course includes four 2-day on-campus meetings which are aligned to the four modules. Students attend the three meetings aligned with their selected module combination.

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 alone or in groups. Teachers or teaching assistants will provide formative feedback.

Presentations: The students work in groups and present results or reading material.

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 reading material or results and get graded according to a grading rubric.

Lab: The students do practical work in computer-based sessions and show or present results and get graded according to a grading rubric. Depending in the number of students this might be done in form of a report.

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