



Aalto University
School of Electrical
Engineering

Reinforcement Learning course Exercise arrangements

Reinforcement Learning course staff

Dept. of Electrical Engineering and Automation

September 7, 2020

Exercise structure

- ▶ Six assignments published and returned on MyCourses (first period)
- ▶ Published each week with a two week deadline
- ▶ Programming tasks and open-ended questions
- ▶ 50% of the total course grade

Assignments

1. Setup + introduction
2. Value iteration
3. Grid-based Q-learning
4. Q-learning with function approximation
5. Policy gradient
6. Actor-critic

Rules

- ▶ Individual work
- ▶ Share ideas, not answers
- ▶ Hard deadlines, except well-justified cases

- ▶ Follow submission instructions (available on MyCourses)

Exercise sessions

- ▶ Organized remotely through Zoom
- ▶ Three sessions a week: Monday 12:15, Tuesday 12:15, Wednesday 10:15
- ▶ Format: Q&A sessions
- ▶ Attendance not compulsory
- ▶ Join whenever you want to—no registration so far (may be changed)

Organization

- ▶ General questions answered in the main room
- ▶ Debugging in breakout rooms
- ▶ Mute your mic when not speaking

Slack channel

- ▶ Main communication tool for the course
- ▶ Good place to ask questions
- ▶ Separate channels for each exercise (#exercise-1, etc)
- ▶ Extra points for activity—up to 10% bonus for each assignment
- ▶ Link on MyCourses

Teaching assistants

Main:

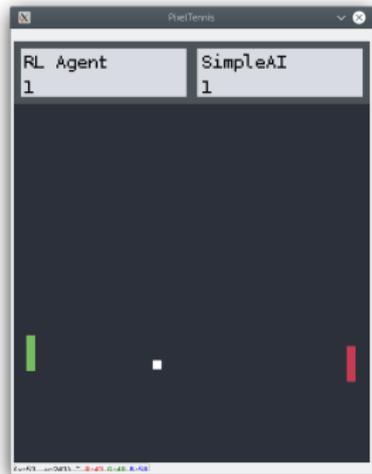
- ▶ Karol Arndt
- ▶ Oliver Struckmeier
- ▶ David Blanco Mulero

Helpers:

- ▶ Aleksii Ikkala
- ▶ Alessandro Amici
- ▶ Luca Scotton

Course Project

- ▶ Pong from pixels, optionally your own
- ▶ Second period (released week 43)
- ▶ Play against a bot and agents submitted by other students
- ▶ 30% of the course grade
- ▶ Last year's environment on Intelligent Robotics Group's GitHub



Today

- ▶ Setup time!
- ▶ Goal: get the provided Exercise 1 code to run on your machine
- ▶ Requires you to install OpenAI Gym, PyTorch, and a few others (see the instructions)
- ▶ Works best on Linux, should be fine on Mac
- ▶ Gym support is experimental on Windows

- ▶ Alternative: remote connection to Aalto servers (see the instructions)