


Deep Learning and Artificial Intelligence WS 2018/19

Exercise 14: Pong Competition

Exercise 14-1 Pong Competition

Wrapping up this lecture, you can demonstrate your learned knowledge in a (voluntary) competition. The challenge is to train a (Deep) Reinforcement Learning agent that learns to play pong. Pong is a famous computer game in which the goal is to hit the ball with a paddle such that the opponent misses the ball and loses (see figure below). While playing, the ball becomes faster and faster.

The environment is implemented in Python. As basis for this competition, the implementation of the university of Toronto¹ is used. Download the sources from the lecture web-site.



Rules:

- No GPU: Your agent is not allowed to use a GPU.
- Only RL agents are valid submissions.
- You may not implement physical formulas (such as the point of impingement) to simplify the problem setting.

Hint: The environment gives the agent 50 milliseconds of time to choose an action at each time step. If he takes longer, he just doesn't move at all in that step.

Submission:

- Deadline: Thu, Feb 28th 2019 11:59 p.m.
- Platform: Upload your solution to the exercise sheet *Pong Competition* in Uniworx.
- Format: Upload a ZIP file that contains your pre-trained parameters and a python file that contains the implementation of your agent. Use the file "sample_submission.py" to create an instance of your agent in the "get_agent" function. Name the file "team_pseudonymous.py" where you replace "pseudonymous" by a pseudonymous name for your team.

Competition:

- After the final exam we will publish the competition results of your agents on the lecture web page. We use your pseudonymous name in the ranking.

¹<http://www.cs.toronto.edu/~guerzhoy/niftypong/>