ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Prof. R. Wattenhofer

Practical Work:

Reinforcement Learning for the Jass AI

Deep reinforcement learning has made headlines recently, beating the world best players in complex board games such as chess, Go and Shogi. More so, it also proved itself in card games, where the current state of the game is only partially observable, by beating the top players in no-limit Texas hold'em poker. We are interested in the algorithms behind this success and ask the simple question: How good can we get at *Jassen*? And, can we set the strength of the AI automatically for a better user experience?

In this thesis, we will build on top of our existing AI and try to improve it. The aim is to use the previously developed imitation learning approach to accelerate training in reinforcement learning.

Requirements: Knowledge in reinforcement learning or a solid background in machine learning. Experience with PyTorch (or TensorFlow) and OpenAI Gym is an advantage. We will have weekly meetings to address questions, discuss progress and think about future ideas.

Contact

- Benjamin Estermann: besterma@ethz.ch, ETZ G60.1
- Flint Xiaofeng Fan: xiafan@ethz.ch, ETZ G93