



Prof. R. Wattenhofer

Democratic Deep Reinforcement Learning

Deep reinforcement learning algorithms have had a lot of success in recent years in playing games at super-human performance. However, humans still beat these algorithms in games that require long term reasoning and many relational inferences. As humans living in democracies we rely hard to solve problems often on the wisdom of the crowd by voting for solutions.

In this thesis we want to investigate whether several different instantiations of deep reinforcement learning algorithms can come to better solutions collectively.

Requirements: Knowledge in Deep Learning, or solid background in Machine Learning. Implementation experience is an advantage. You should be able to read and understand the first 12 chapters of the "Deep Learning Book" by Goodfellow et al. (available for free online from MIT press). If you are interested in the topic but new to deep learning we expect you to complete an introductory deep learning course before applying for the thesis, such as Andrew Ng's coursera course (use the free trial!)¹ or this Udacity course². Further, experience with reinforcement learning or bandit problems is an advantage.



Interested? Please contact us for more details!

Contacts

- Oliver Richter: richtero@ethz.ch, ETZ G63
- Gino Brunner: brunnegi@ethz.ch, ETZ G63

¹<https://www.coursera.org/specializations/deep-learning>

²<https://classroom.udacity.com/courses/ud730>