ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Prof. R. Wattenhofer

Exploration-Exploitation Trade-off in Deep Reinforcement Learning

In 2015 Google Deepmind published their DQN paper in which they present an algorithm capable of learning to play Atari games at superhuman performance. Since then, research and applications of deep reinforcement learning (DRL) have gained more and more momentum. Results range from algorithms capable of mastering chess within 4 hours[1] to neural networks capable of constructing other neural networks.[2]

Despite all the astonishing results, DRL models still lack sample efficiency, i.e., training speed, and generality, such as training one algorithm for several tasks. In this thesis, we investigate and compare the different approaches presented in the literature to address the exploration-exploitation trade-off.

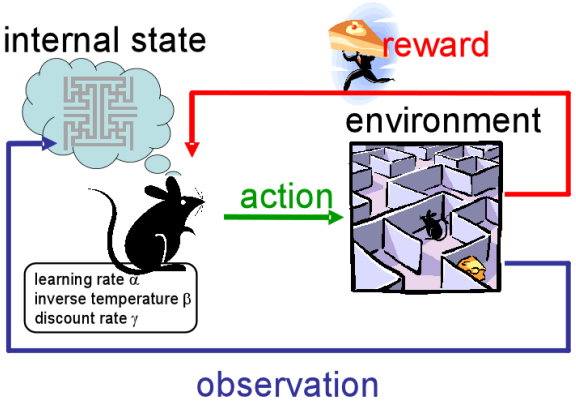

Requirements: Knowledge in Deep Learning, or solid background in Machine Learning. Implementation experience is an advantage. You should be able to read and understand the first 12 chapters of the "Deep Learning Book" by Goodfellow et al. (available for free online from MIT press). If you are interested 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!)[3] or this Udacity course[4].

Contacts

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

[1] https://arxiv.org/abs/1712.01815
[2] https://arxiv.org/abs/1611.01578
[3] https://www.coursera.org/specializations/deep-learning
[4] https://classroom.udacity.com/courses/ud730