Hybrid Reinforcement Learning for Robot Design

Technical University of Munich

Department of Informatics

Chair of Robotics, Artificial Intelligence and Real-time Systems

Background

In recent years, deep reinforcement learning algorithms have led to great advances in the design and control of autonomous agents[2, 4]. Traditionally, these methods are applied on either a discrete or continuous action space. In the real scenarios though, an agent needs to operate in hybrid, i.e. mixed discrete-continuous action spaces. As traditional methods fail to generalize to hybrid action spaces without a significant loss in performance, hybrid reinforcement learning methods have been developed to bridge the gap between simulated/discretized problems and real world applications.

Description

In this thesis, you will compare, evaluate and improve hybrid reinforcement learning methods such as MP-DQN[1] and Hybrid MPO[3] in the context of modular robot design. Modular robots, opposed to monolithic manipulators, can easily be adapted and optimized for any task at hand within minutes. It is unclear though how robot modules should be designed in order to offer both, generalizability and performance in cricitcal tasks. You will evaluate the potential of hybrid reinforcement learning to solve these design questions.

Tasks

- **Literature research on parameterized markov decision processes (PAMDP) and hybrid reinforcement learning**
- **Adaption and development of a hybrid RL algorithm for the task at hand**
- **Task-based synchronous learning of module design and configuration**
- **Evaluation and comparison to non learning-based appraoches**

If you are interested in this thesis, please send me an E-Mail with CV, transcript of records, links to previously completed projects and a short statement of motivation why you are suited for this topic.

References

[1] **Craig J. Bester, Steven D. James, and George D. Konidaris. Multi-Pass Q-Networks for Deep Reinforcement Learning with Parameterised Action Spaces. arXiv, 2019.**

[2] **David Ha. Reinforcement learning for improving agent design. *Artificial Life*, pages 352–365, Nov 2019.**

[3] **Michael Neunert, Abbas Abdolmaleki, Markus Wulfmeier, Thomas Lampe, Tobias Springenberg, Roland Hafner, Francesco Romano, Jonas Buchli, Nicolas Heess, and Martin Riedmiller. Continuous-Discrete Reinforcement Learning for Hybrid Control in Robotics. In *Proceedings of the Conference on Robot Learning*, Proceedings of the Conference on Robot Learning, pages 735–751. PMLR, 2019.**

[4] **Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. arXiv, 2019.**

Supervisor:
Prof. Dr.-Ing. Matthias Althoff

Advisor:
Jonathan Külz

Research project:
CONCERT

Type:
Master Thesis only

Research area:
Deep Reinforcement Learning, (Modular) Robotics

Programming language:
Python

Required skills:
Experience w. Deep Learning, Clear Coding Style

Language:
English

Date of submission:
10. November 2022

For more information please contact us:

Phone: +49 (89) 289 - 18114

E-Mail: jonathan.kuelz@tum.de