Karlsruhe Institute of Technology

Autonome Lernende Roboter (ALR)
Prof. Gerhard Neumann

Project Type

- [ ] Master Thesis
- [x] Bachelor Thesis
- [ ] Research Project

Supervisors

Onur Celik
celik@kit.edu

Difficulty

Algorithmic

Math

Application

Reinforcement Learning for Robot Manipulation Tasks

Description

Reinforcement Learning (RL) has shown increasing performance in the last years and allowed agents to solve sophisticated learning tasks [2, 3, 4]. The RL community commonly uses OpenAi's gym environments [1] (see Fig. 2) to benchmark the algorithms. While these environments provide a unified, generalized and easy-to-handle opportunity to easily compare to other baselines, they usually do not consider realistic robot motion and robot manipulation behavior.
To provide a dimensionally correct test environment compared to our robot-lab setup, we have created a simulation framework (SF) (see Fig. 1), in which custom environments can be built. This SF is tweaked to provide a best possible, physically realistic simulation.
In this thesis we will first create a sophisticated robot manipulation environment in our SF. Examples for these environments are a robot arm that learns to open a drawer, place shaped objects into appropriately shaped holes, pushes a box to a certain position and orientation and more creative scenes :). As a second step, this work will involve testing recent RL algorithms on the created environment to show that they are suitable to learn difficult robot manipulation tasks. Further add-ons like transferring to the real robot can be considered depending on the progress of the thesis. In this project the student will get familiar with realistic robot manipulation tasks and directly apply state of the art methods in deep RL, such as SAC and PPO.

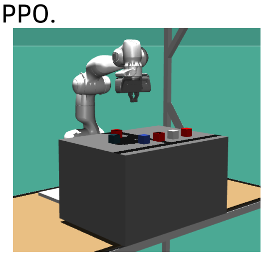

Figure 1: Scene manipulation in our SF.

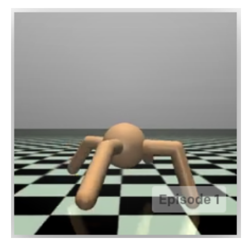

Figure 2: Open Ai's Ant Environment [1]

Tasks

The tasks in this project will involve:

- Literature review: Get familiar with RL algorithms with a detailed literature review in recent methods.
- Soft Actor-Critic (SAC): Get known to the Soft Actor-Critic algorithm, which we will use to learn on the newly created environment
- Simulation Framework: Get known to our lab's simulation framework (SF), which we will use to build the robot manipulation environment.
- Evaluating: We will do an intensive evaluation on the environment with recent algorithms (SAC, PPO, ...)
- Transfer to real robot (optional): We might consider transferring the learned policies to the real robot in our lab.

References

[1] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.

[2] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.

[3] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.

[4] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.