Karlsruhe Institute of Technology

Autonome Lernende Roboter (ALR)
Prof. Gerhard Neumann

Project Type

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

Supervisors

Gerhard Neumann
gerhard.neumann@kit.edu
Hanna Ziesche
Hanna.Ziesche@de.bosch.com
Vien Ngo
AnhVien.Ngo@bosch.com

Difficulty

Algorithmic

Math

Application

Learning Complex Manipulations of Cluttered Scenes with Deep RL

Description

In this project, we aim to develop and optimize a hierarchical policy based on movement primitives, that allows a robot to achieve a long-term goal in visual grasping and manipulation.

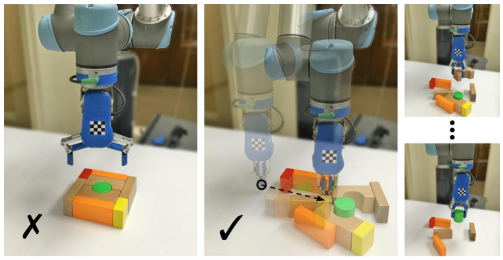

Figure 1: Example of pushing and grasping synergies [2].

Reinforcement learning from visual inputs for robotic manipulation is a topic of growing importance both for the scientific community and for industrial research. One popular approach, first proposed in [2], relies on a discretization of the action space according to the pixel space and a finite number of discrete rotations. Due to the chosen action representation the approach is however limited to simple policies of linear pushes and top-down grasps. Furthermore, the method proposed by [2] is limited to short time horizons and is thus not applicable to tasks that require reasoning over extended time horizons. The goal of this master thesis will be to find methods to address these two short-comings by means of learned motion primitives.

Tasks

Your task will be to extend the approach proposed in [2] to complex, structured policies, parametrized as learnable motion primitives (e.g. Probabilistic Movement Primitives or Dynamic Movement Primitives). You will work in a provided simulation environment and learn the parameters of the motion primitives by means of policy gradient. Due to the known instabilities of policy gradient, you will be using recently developed Trust Region Layers [1]. In details, your task will comprise

- Extending the formulation of the Trust Region Layers by a discrete part.
- Incorporating the necessary algorithmic changes into an existing Python framework.
- Learning a structured policy using motion primitives in simulation with a policy gradient approach.
- Thorough experimentation with the implemented method in simulation and assessment of algorithmic choices.

Special Feature

The thesis will be hosted by the Bosch Center of Artificial Intelligence, i.e. you will be given a 6 month contract by Bosch to work on your thesis.

References

[1] Fabian Otto, Philipp Becker, Ngo Anh Vien, Hanna Carolin Ziesche, and Gerhard Neumann. Differentiable trust region layers for deep reinforcement learning. In *International Conference on Learning Representations*, 2021.

[2] Andy Zeng, Shuran Song, Stefan Welker, Johnny Lee, Alberto Rodriguez, and Thomas A. Funkhouser. Learning synergies between pushing and grasping with self-supervised deep reinforcement learning. In *IEEE International Conference on Intelligent Robots and Systems (IROS) 2018*, pages 4238–4245.