Karlsruhe Institute of Technology

Intuitive Robots Lab (IRL)
TT-Prof. Rudolf Lioutikov

Project Type

- ☑ Master Thesis
- ☐ Bachelor Thesis
- ☐ Research Project

Supervisors

Hongyi Zhou
hongyi.zhou@kit.edu
Ge Li
ge.li@kit.edu

Difficulty

Algorithmic

Math

Application

Requirements

- ☑ Python and Pytorch
- ☑ Knowledge about robotics
- ☑ Knowledge about ML and RL
- ☑ ROS and C++ (recommended)

Solving Real-World Robot Manipulation Tasks with Deep Reinforcement Learning

Description

Deep reinforcement learning (DRL) has demonstrated remarkable success in solving robot control tasks within simulated environments. However, the challenge arises when transferring these learned skills directly to the real robot hardware, primarily due to the significant disparities between simulation and reality. To address this challenge, an additional step of online fine-tuning on real hardware [1] becomes necessary. Moreover, most DRL approaches generate noisy actions during exploration, which makes collecting samples on real robots even more challenging. Recently, we introduced a method [2] to incorporate Movement Primitives(MPs) into the DRL framework. This integration enables smooth and efficient exploration at the trajectory level, offering potential benefits for real-world fine-tuning.

Figure 1: Franka panda robots

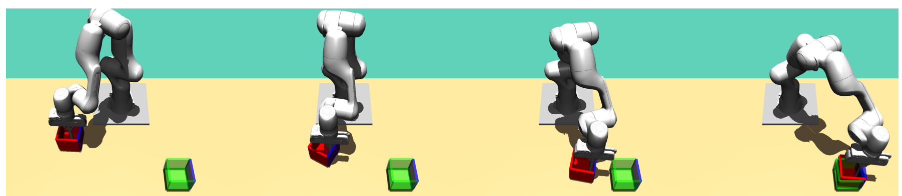

Figure 2: Box-pushing task in simulation

In this thesis, we want to develop a sim-to-real transformation pipeline for Reinforcement Learning with Movement Primitives (MPRL) approach, enhancing robot manipulation and control in real-world scenarios. Specifically, our objective is to deploy a box-pushing task (Fig.2) on a Franka Panda robot platform (Fig.1).

Tasks

The tasks in this project will involve:

- Conducting literature research to gain a comprehensive understanding of MPRL methods and techniques for sim-to-real transformation.
- Developing a pipeline to mitigate the Sim2Real gap for MPRL approach.
- Deploying the RL agent pretrained in simulation onto a real robot and conducting online fine-tuning on the pretrained model.
- Evaluating the algorithm's performance against baseline approaches.

References

[1] Saminda Wishwajith Abeyruwan, Laura Graesser, David B D'Ambrosio, Avi Singh, Anish Shankar, Alex Bewley, Deepali Jain, Krzysztof Marcin Choromanski, and Pannag R Sanketi. i-sim2real: Reinforcement learning of robotic policies in tight human-robot interaction loops. In *Conference on Robot Learning*, pages 212–224. PMLR, 2023.

[2] Fabian Otto, Hongyi Zhou, Onur Celik, Ge Li, Rudolf Lioutikov, and Gerhard Neumann. Mp3: Movement primitive-based (re-) planning policy. *arXiv preprint arXiv:2306.12729*, 2023.