Karlsruhe Institute of Technology

Intuitive Robots Lab (IRL) & Autonomous Learning Robots (ALR)

Prof. Rudolf Lioutikov

Project Type

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

Supervisors

- Moritz Reuss
- moritz.reuss@kit.edu
- Philipp Becker
- philipp.becker@kit.edu

Difficulty

Algorithmic

Math

Application

Joint Representations for Real World Robotic Skill Learning

Description

To solve complex tasks in unstructured real-world settings, robots need to perceive their environment and interact with it. Here, a common approach is to first learn a concise representation of the robot and its environment and use this representation as the basis for acting. To form such representations, the robot usually needs to merge information from various sensor sources, such as cameras, force sensors, and the sensing in its actuators. In recent work[1], we showed that learning a joint representation for all sensors is favorable over using individually learned representations and investigated various ways to do so.

In order to effectively use representations learned from large, uncurated datasets, it is essential to have efficient policies that can successfully tackle long-horizon tasks. To this end, we use a novel diffusion-based imitation learning policy for learning goal-conditioned skills, which has demonstrated its effectiveness in the field of robot policies. Although diffusion models are commonly used as generative models for image generation, we believe that their application in this context can lead to substantial improvements in the quality of learned policies.

In this thesis, we aim at combining these two approaches into an efficient pipeline for imitation-based skill learning from multiple sensors. The main focus will be on the deployment and evaluation of the resulting method on a real robot in a challenging kitchen environment.

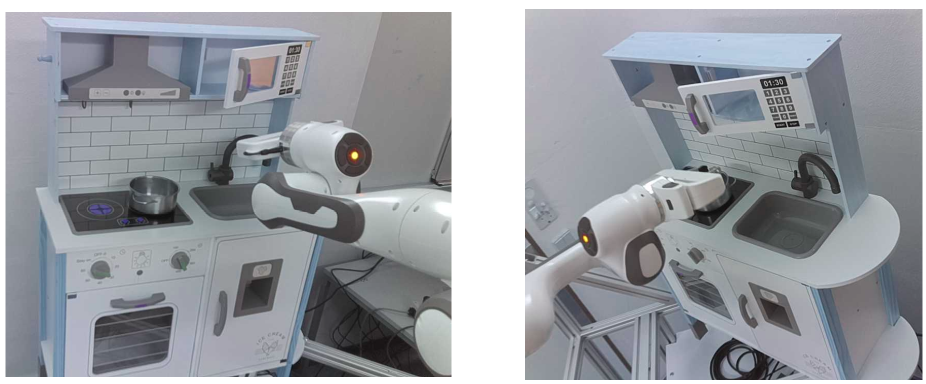

Figure 1: Visual example of the opening cabinet door skill in a real play kitchen environment. The agent has access to two visual representations, which we want to combine to a joint one for effective policy learning.

Tasks

- Getting familiar with the relevant related work, simulation, and the existing real robot pipeline.
- Development of the approach by merging the two existing code bases and preliminary evaluation in simulation.
- Deployment of the approach on the real robot.
- Data collection and large-scale evaluation on the real robot.

References

[1] Philipp Becker, Sebastian Markgraf, Fabian Otto, and Gerhard Neumann. Reinforcement learning from multiple sensors via joint representations, 2023.