Robotics and Embedded Systems

Master/Bachelor Thesis

Diffusion-based Continual Learning for Robotic Manipulation

Background

Catastrophic forgetting remains a major challenge in robotic lifelong learning. In robotic manipulation, behavior clone combined with data replay is a common strategy for retaining previously acquired policy knowledge. However, such policy models often show limited long-term decision-making capability and suboptimal generalization across sequential tasks. Recently, diffusion models have demonstrated strong representative capability in decision-making tasks. This project aims to investigate the potential of diffusion policies within a continual learning framework for robotic manipulation, with the goal of improving policy retention and overall performance.

Your Tasks

In this project, you will focus on improving both the continual learning simulation platform and the continual learning algorithm. Specifically, you may begin your thesis by concentrating on one of the following tasks:

- **Enhanced Continual Learning Algorithm**: extend our existing continual learning framework by incorporating and optimizing diffusion policies to improve continual learning performance.
- **KUKA-based Simulation Platform**: based on the Continual World and MetaWorld benchmarks, adapt the manipulator to a KUKA robotic arm and develop a KUKA-based continual manipulation simulation platform.
- **Algorithm Evaluation**: evaluate the improved continual learning algorithm on the KUKA-based continual learning simulation environment.

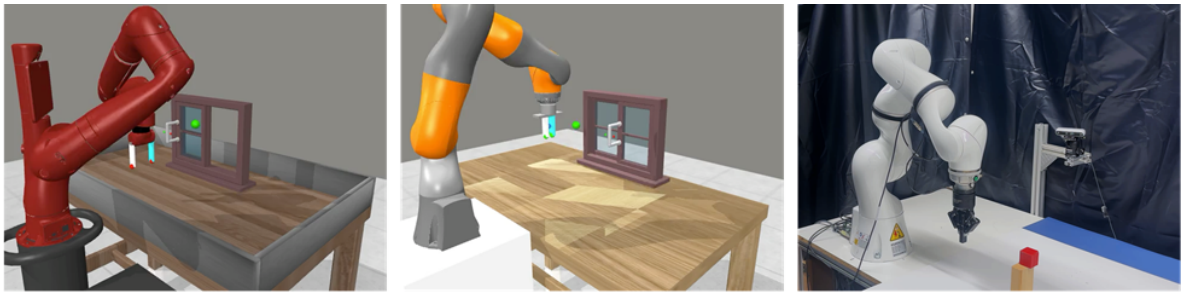

(a) Continual World benchmark (b) KUKA-based simulation scenario (c) real-world scenario setup

Figure 1 Experiment Platform

Requirements

- High self-motivation and passion on research.
- Python programming experiences.
- Six-month working time.
- Existing knowledge about reinforcement learning or supervised learning will be a bonus.

Supervisor: Prof. Alois Knoll

Advisor: Shimin Liu (shimin.liu@tum.de) (Attach your CV and Transcript)

Reference

[1] Chi, Cheng, et al. "Diffusion policy: Visuomotor policy learning via action diffusion." The International Journal of Robotics Research 44.10-11 (2025): 1684-1704.

[2] Chen, Feng, et al. "Stable continual reinforcement learning via diffusion-based trajectory replay." arXiv preprint arXiv:2411.10809 (2024).

[3] Yu, Tianhe, et al. "Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning." Conference on robot learning. PMLR, 2020.

[4] Wołczyk, Maciej, et al. "Continual world: A robotic benchmark for continual reinforcement learning." Advances in Neural Information Processing Systems 34 (2021): 28496-28510.