Karlsruhe Institute of Technology

Autonome Lernende Roboter (ALR)
Prof. Gerhard Neumann

Project Type

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

Supervisors

☞ Florian Seligmann
@ florian.seligmann@kit.edu

Difficulty

Algorithmic

Math

Application

Intervention-based Data Collection for Vision-Language-Action Models

Description

Vision-Language-Action (VLA) models have shown immense promise in robotic manipulation, but their performance is heavily bottlenecked by data quality and coverage. To improve robustness, many approaches rely on Dataset Aggregation (DAgger) style imitation learning. However, the standard implementation in VLA pipelines is rarely truly interactive: Operators observe the VLA's failure modes during rollout, reset the environment, and then collect new demonstrations from scratch to cover those edge cases [3]. This "offline" approach is inefficient as it often fails to capture the exact state where the policy originally deviated.

In this project, we aim to transition to an intervention-based data collection paradigm [1, 2]. Instead of starting over, the operator will monitor the live rollout and intervene by pausing the VLA policy the moment it enters a failure state. The operator then takes over via teleoperation to provide targeted, corrective actions before returning control to the policy.

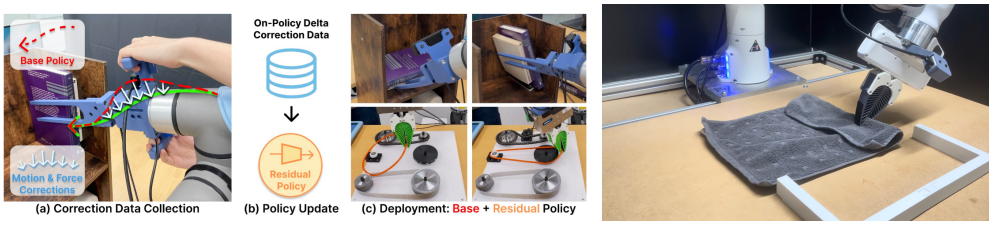

Figure 1: Left: Interactive data collection in CR-DAgger [2]. Right: X-VLA [3] folding a towel on our Franka Emika Panda setup.

This thesis will rigorously evaluate the sample efficiency and final policy performance of this interactive approach on a Franka Emika Panda robot. The core research question is whether targeted, mid-rollout interventions yield a higher success rate than simply scaling up the dataset with more non-targeted data or traditional, from-scratch failure demonstrations.

Tasks

The tasks in this project will involve:

- Implement the VLA baseline: Deploy an existing VLA architecture on our robotic setup and establish the standard "offline" data collection pipeline.
- Intervention infrastructure: Develop the system architecture to seamlessly pause the autonomous VLA policy mid-rollout, switch to human teleoperation, and accurately record the corrective state-action data.
- Targeted data collection: Design a set of manipulation tasks and collect failure-recovery data using both the traditional reset-and-repeat method and the targeted intervention approach.
- Evaluation and comparison: Fine-tune the VLA on both datasets and systematically comparing the resulting policies.

References

[1] Michael Kelly, Chelsea Sidrane, Katherine Driggs-Campbell, and Mykel J Kochenderfer. Hg-dagger: Interactive imitation learning with human experts. In *ICRA*, 2019.

[2] Xiaomeng Xu, Yifan Hou, Zeyi Liu, and Shuran Song. Compliant residual DAgger: Improving real-world contact-rich manipulation with human corrections. In *NeurIPS*, 2025.

[3] Jinliang Zheng, Jianxiong Li, Zhihao Wang, Dongxiu Liu, Xirui Kang, Yuchun Feng, Yinan Zheng, Jiayin Zou, Yilun Chen, Jia Zeng, Tai Wang, Ya-Qin Zhang, Jingjing Liu, and Xianyuan Zhan. X-VLA: Soft-prompted transformer as scalable cross-embodiment vision-language-action model. In *ICLR*, 2026.