Karlsruhe Institute of Technology

Autonome Lernende Roboter (ALR) and Intelligent Sensor-Actuator-Systems Laboratory (ISAS)

Prof. Gerhard Neumann

Prof. Dr.-Ing. Uwe D. Hanebeck

Project Type

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

Supervisors

- Tai Hoang
- tai.hoang@kit.edu
- Markus Walker
- markus.walker@kit.edu

Difficulty

Algorithmic

Math

Application

MPC on Graph Neural Simulators

Description

Robotic control problems often have a *structured* state space: the system state lives on a mesh or graph (e.g., deformable objects, cloth, soft bodies, fluids). Model Predictive Control (MPC) is a natural tool for such settings because it iteratively refines control inputs by rolling out a dynamics model and reoptimizing in closed-loop. In practice, however, high-fidelity physics simulators can be too slow for online optimization, and analytic models are often hard to derive or calibrate. Recent Graph Neural Simulators (GNSs) provide a promising alternative: they learn dynamics directly from data while exploiting mesh/graph structure, and they can generate fast multi-step rollouts with message passing. In this project, we use Information-preserving Graph Neural Simulators (IGNS) as the learned world model [3]. IGNS is designed for accurate long-range information propagation with efficient inference, which is especially important for MPC, where many rollouts must be evaluated per control step. The student will implement MPC and evaluate oracle model vs. learned-model trajectories, and optionally add a TD-MPC2-style terminal value for better long-horizon behavior [1].

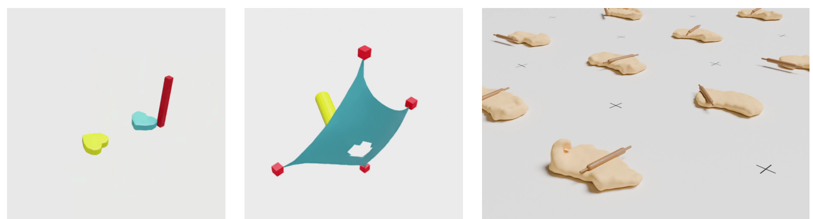

Figure 1: Rigid-Pushing, Cloth-Hanging, and Rolling-Flat tasks from [2, 4].

Tasks

- MPC baselines. Implement receding-horizon MPC, and verify correctness with oracle dynamics.
- World model learning. Collect trajectories on the benchmark tasks (Fig. 1) and train IGNS as a multi-step rollout model [3]. Evaluate short- and long-horizon prediction accuracy.
- Closed-loop control with learned dynamics. Replace oracle dynamics with IGNS in MPC. Compare performance under matched compute budgets (e.g., horizon length, sample count, optimizer iterations) and analyze robustness to model error.
- Scaling and extensions. Run the same pipeline across the tasks in Fig. 1 (Rewarped, HEPi). Optional: add a TD-MPC2-style terminal value to reduce horizon and improve long-horizon behavior [1].

References

[1] Nicklas Hansen, Hao Su, and Xiaolong Wang. Td-mpc2: Scalable, robust world models for continuous control, 2024.

[2] Tai Hoang, Huy Le, Philipp Becker, Vien Anh Ngo, and Gerhard Neumann. Geometry-aware RL for manipulation of varying shapes and deformable objects. In *The Thirteenth International Conference on Learning Representations*, 2025.

[3] Tai Hoang, Alessandro Trenta, Alessio Gravina, Niklas Freymuth, Philipp Becker, Davide Bacciu, and Gerhard Neumann. Improving long-range interactions in graph neural simulators via hamiltonian dynamics. In *The Fourteenth International Conference on Learning Representations*, 2026.

[4] Eliot Xing, Vernon Luk, and Jean Oh. Stabilizing reinforcement learning in differentiable multiphysics simulation. *International Conference on Learning Representations (ICLR)*, 2025.