Chair of Imaging and Computer Vision

RWTH AACHEN UNIVERSITY

Univ.-Prof. Dr.-Ing. Volkmar Schulz
Head of Institute

Marian Frei
Research Scientist

ITC Cubes - Office 105
Kopernikusstr. 16
52074 Aachen

marian.frei@lfb.rwth-aachen.de
11.09.2026

Master Theses in Robot Learning

Three open topics on real hardware

Foundation models are moving from language into robotics. We run a bimanual platform to find out what that actually buys on real hardware.

Chair of Imaging and Computer Vision · RWTH Aachen University

Robotics and Foundation Models

For most of its history robot manipulation was engineered task by task, with perception, planning and control built by hand for one setting at a time. Vision Language Action models are changing that. Pretrained on thousands of hours of robot data across many embodiments, they carry general manipulation priors and transfer to new tasks with comparatively little data. The shift that language models brought to text is now under way in robotics.

What is unsettled is almost everything around the model: which data is actually worth collecting and how progress should be measured, how far a policy can be pushed on hard initial states and towards human speed, and what a language model with the right tools can do on hardware it has never seen. Those three questions are the three topics below, and none can be settled in simulation. We run a bimanual manipulator as a testbed, with a trained policy, a dataset and a learned progress model in place, currently on deformable objects and with further task families in preparation.

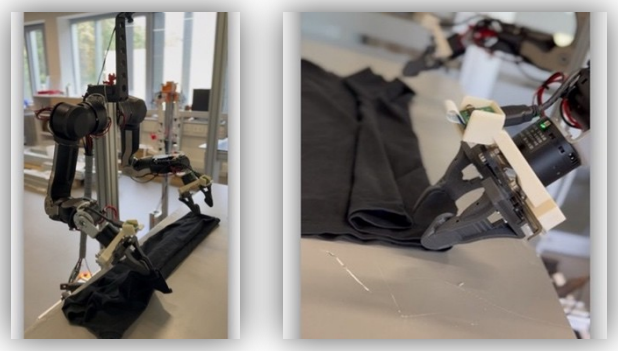

The bimanual platform in operation

Open Topics

Each topic exists as a detailed one page description on request. Common to all three: you start from a system that already runs, everything is measured on the physical robot, and evaluation counts as a contribution in its own right.

01 Scaling Laws for Robot Data: Predicting Success Before Training

Scaling laws made language models predictable; robot learning has no such law yet. The first ones appeared last year, on single arm rigid object tasks, and how demonstration quality enters or how they behave for bimanual deformable manipulation is open. You build automated evaluation trustworthy enough to see a law at all, then measure ours: how success grows with data quantity, quality and coverage.

Keywords: data scaling laws, autonomous evaluation, data curation, imitation learning

02 From Crumpled to Folded, and Fast: Vision Language Action Policies

Folding is reliable from a flat shirt and degrades sharply once the shirt starts crumpled and has to be spread first, and even then our setup needs 40 seconds where a human needs 10. You work both axes, recovery from hard initial states and action representations that can be rescaled in time, then measure where success breaks as difficulty and speed rise.

Keywords: VLA models, imitation learning, action representations, deformable manipulation

Chair of Imaging and Computer Vision | RWTH AACHEN UNIVERSITY

03 Language Models as Robot Operators

A language model has never seen our hardware, yet with the right tools it can read a scene, plan a fold and explain why a grasp failed, which makes the scaffold around the model the real research object. You design that scaffold, compare frontier against local models with success rates rather than demos, and then turn the agent into a data source that generates demonstrations for training a policy.

Keywords: LLM agents, tool design, code as policy, demonstration generation

Your Profile

You are a Master's student in Computer Science, Electrical Engineering, Mechatronics, Robotics, Mathematics or a comparable field.

- **Required:** Python, a modern deep learning framework and a solid grounding in machine learning.
- **Useful per topic:** data curation, evaluation methodology and applied statistics (01), generative modelling such as diffusion or flow matching and real time control (02), LLM agents, tool calling and vision foundation models (03). None of it is a prerequisite.
- **Matters most:** you want to run the full loop yourself, from training through deployment to evaluation on the robot, and you can stay with an experiment that runs for days.

What We Offer

- **A system that already runs:** a bimanual platform, a curated demonstration dataset, a trained policy and a learned progress model, plus a reward model and recorded rollout footage where the topic needs them
- **Compute and workshop:** dedicated strong GPU servers and HPC GPU hours, so you train real models rather than toy versions, plus rapid prototyping and 3D printing
- **Supervision:** close guidance and regular exchange among students, PhD candidates and postdocs, freedom to pursue your own ideas, and a scope shaped together with you
- **Publication:** theses that produce solid results regularly grow into papers, and we write them with you

How to Apply

Send a short message saying which topic draws you and why, along with your CV and transcript. If none of the three fits exactly, we shape one with you.

Marian Frei · marian.frei@lfb.rwth-aachen.de