



Master Thesis

Active Perception *with* Viewpoint Memory *for* Robotic Manipulation

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

Dr.-Ing. Yuli Wu
Group Leader

Physical AI Lab
Chair of Imaging and Computer Vision
RWTH Aachen University
Kopernikusstr. 16
52074 Aachen, Germany

<https://pai.rwth-aachen.de/>

*We are inviting applications for the **Master Thesis Topic** in the **Physical AI Lab** at the Chair of Imaging and Computer Vision, Faculty of Electrical Engineering and Information Technology, RWTH Aachen University. Our group develops generalist robots that autonomously solve real-world long-horizon tasks.*

Introduction of the Topic

Robots fail as soon as what they need is out of view: behind them, on a shelf above, or on the hidden side of an object. *Active perception* means moving the camera or the object to find out, and remembering what was seen once it leaves the frame. On the ActiveArena benchmark [1], the best-performing Vision-Language-Action (VLA) policy stores one frame per action chunk, uniform in time, and still loses about a quarter of its success under new backgrounds, lighting, and clutter. For a robot that moves its own camera, memory may be better organized by where it looked than by when, covering the workspace with fewer frames and exposing what remains unexplored. You will build viewpoint-indexed memory into the ActiveArena VLA baselines, train it at scale, and run the ablations needed to find out.

[1] <https://leeibo.github.io/ActiveArena/>

Your Profile

- Experience of at least one of the following areas, and ideally several: machine learning, computer vision, language models, robotics, diffusion models, reinforcement learning, or related fields
- Strong programming skills, with the ability to rapidly prototype, implement, and evaluate research ideas
- High motivation, scientific curiosity, and the ability to work both independently and collaboratively
- Interest in embodied AI and spatial reasoning; experience with robot simulators (e.g., SAPIEN, MuJoCo, Isaac) or large-scale GPU training is a plus

What We Offer

- An open research question rather than a reimplementing task
- Close supervision and scientific mentoring; strong support for publications
- A ready-to-run active-perception benchmark with released demonstrations, baseline models, and a fixed evaluation protocol
- Access to extensive on-premises computing infrastructure, including around 120 GPUs and additional high-end GPU clusters available on demand

Application

Duration of the master thesis is 6 months. Please send your application to the following email address yuli.wu@lfb.rwth-aachen.de using the subject “**ActivePercept_{Name}**” and include your CV and academic transcripts.