universität freiburg

WS 25/26 Seminar
Robot Learning

Akshay L Chandra

Robot Learning Lab

16 October 2025

ellis UNIT FREIBURG

robot learning

Agenda

I. Organization: Enrollment, important dates and evaluation.

II. Robot Learning Lab: Our research interests and publications.

III. Topics: Seminar Papers.

?. Questions.

I.
Organization

Enrollment, important dates and evaluation criteria

Seminar

Objectives

- Learn to read and understand **scientific literature**.
- Familiarize with the **State-of-the-Art** (SOTA) in the field.
- Discover **limitations**, propose **improvements** and potential **future** work.
- Build knowledge from **related work,** prior and follow-ups.
- Improve **presentation** skills.
- Develop abilities for **synthesis** (diagram drawing, summarizing main ideas, …).

TL;DR:

Show us that you have a **solid** grasp of your topic.

Enrollment Procedure

Select **3 papers** in decreasing order of preference.

Register for the seminar in HISinOne.

Students selected based on HISinONE Priority.

Students assigned papers based on their preferences

Fill in our **Google Form**

By **20.10.2025**

Please check the course website for more information:

https://rl.uni-freiburg.de/teaching/ws25/robotlearning/

Important Dates

Event	Date	Time
Lecture 1: Introduction *	16.10.2025	16:00
HISinOne Registration + Paper Selection	19.10.2025	
Place Allocation	23.10.2025	
Paper Assignment	24.10.2025	
Supervisor Meeting	-	
Lecture 2: How to Make a Presentation *	12.01.2026	14:00
Lecture 3: Block Seminar Presentations *	06.02.2026	9:00 - 17:00
Paper Summary Submission	20.02.2026	< 23:59

* **Mandatory in-person attendance**

Evaluation Criteria

Evaluation	Due Date
Seminar Presentation	06.02.2026
Paper Summary	20.02.2026

- **Presentation:** at most **20 min**.
- **Summary:** at most **7 pages** excluding bibliography and figures.
- **Final Grade:**
  - Presentation (slides & delivery) + Summary + **Seminar Participation**.

II.

Robot Learning Lab

Our research interests and publications

Autonomous Robotics

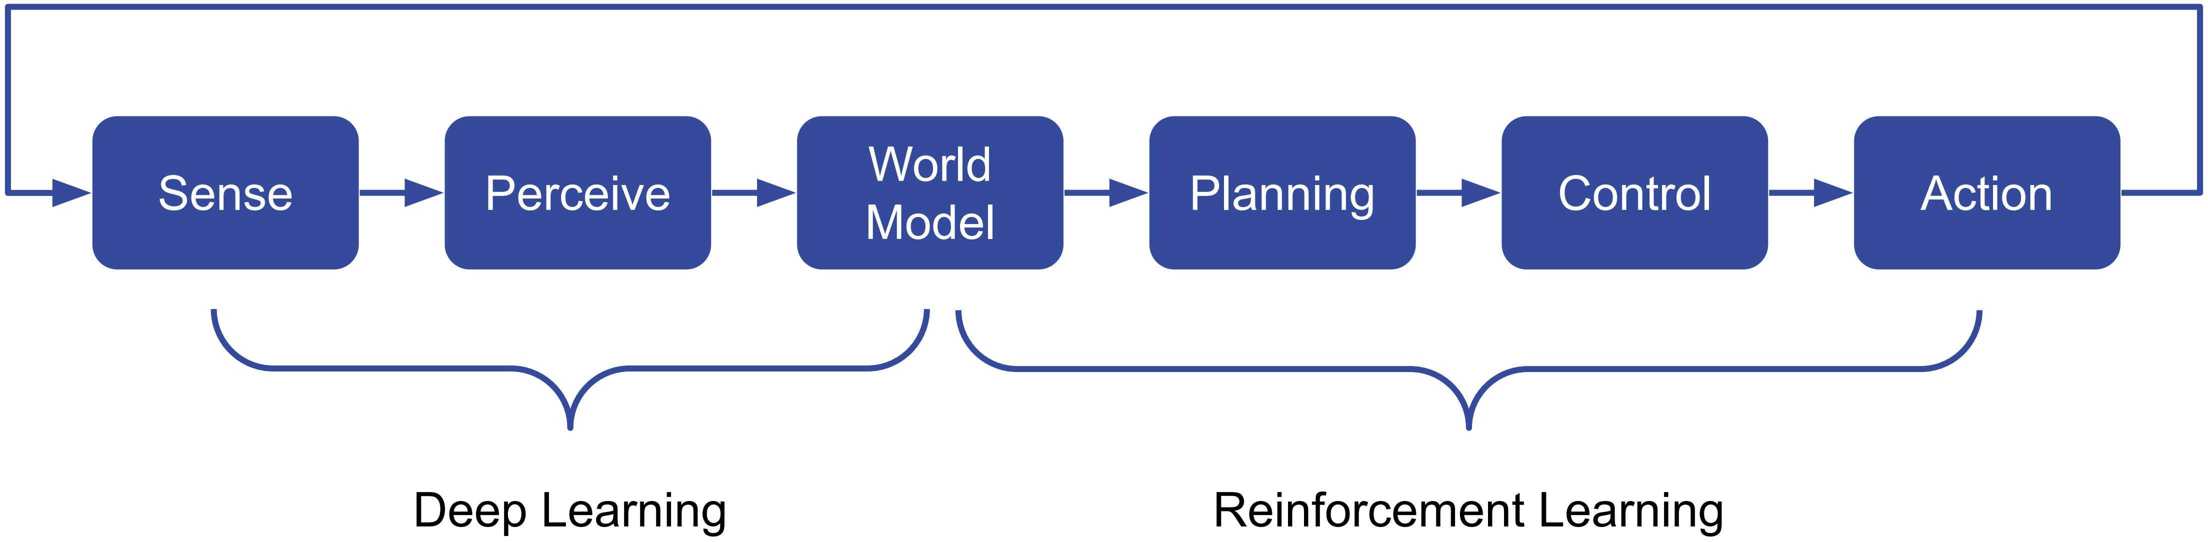

Can we **learn** certain parts of this pipeline?

Robot Learning

Learning …

- … models of robots, tasks or environments
- … deep hierarchies/representations from sensor and motor representations to task representations
- … plans and control policies
- … methods for probabilistic inference from multi-modal data
- … structured spatio-temporal representations, e.g. low-dim. embeddings of Movements

How can we ensure **autonomous operation** of embodied AI systems?

Robot Learning Lab

Research Areas

Perception
- Recognition
- Depth Estimation
- Motion Estimation

State Estimation
- Tracking & Prediction
- SLAM
- Registration

Motion Planning
- Hierarchical Learning
- Reinforcement Learning
- Learning from demonstration

Responsible Robotics
- Fairness
- Explainability & Privacy
- Practical Ethics

Mobile Manipulation
- Whole-Body Motion
- Long-Horizon Reasoning
- Planning for Sensing

Human-Robot Interaction
- Socially-Compliant Behavior
- Human-Robot Collaboration
- Behavior Adaptation & Safety

Learning Fundamentals
- Socially-Supervised Learning
- Continual & Interactive Learning
- Multimodal & Multitask Learning

Many Seminal Works

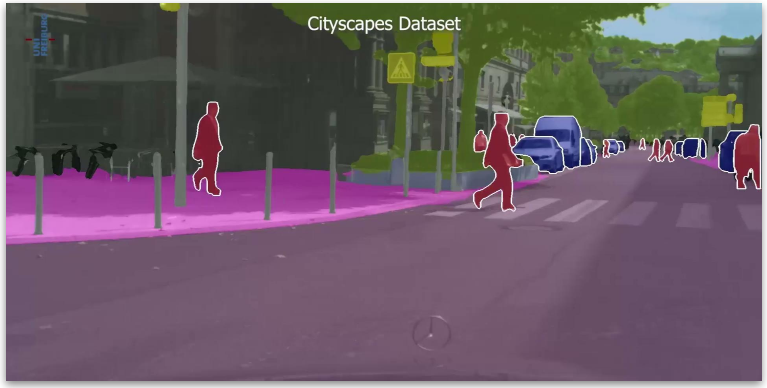

Scene Understanding

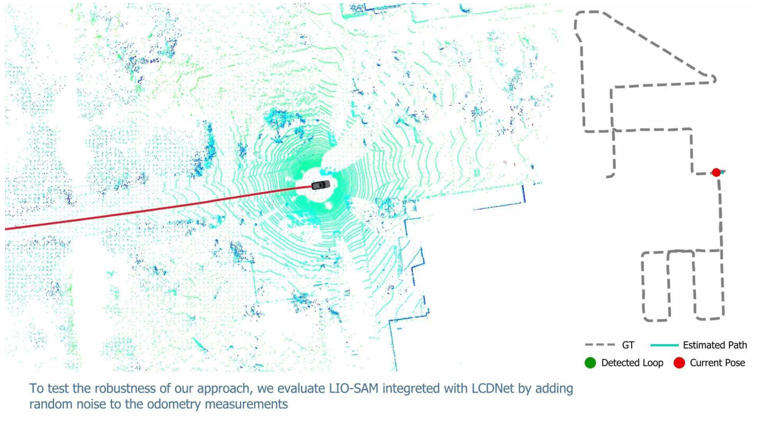

Simultaneous Localization and Mapping

Motion Planning

Learning from Demonstrations

Robotic Perception — Mobility

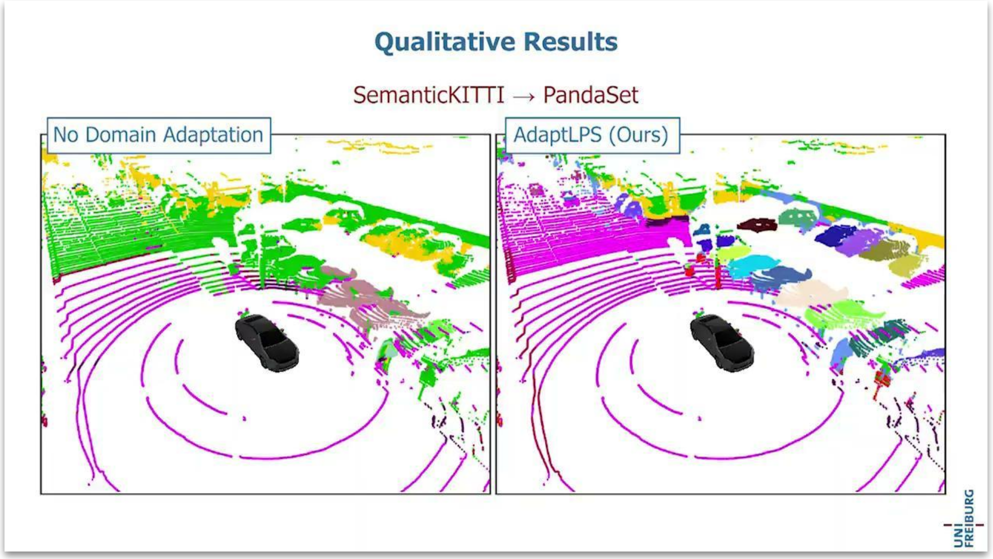

Unsupervised LiDAR Domain Adaptation
Besic, Gosala, Cattaneo, Valada
RA-L ‘22

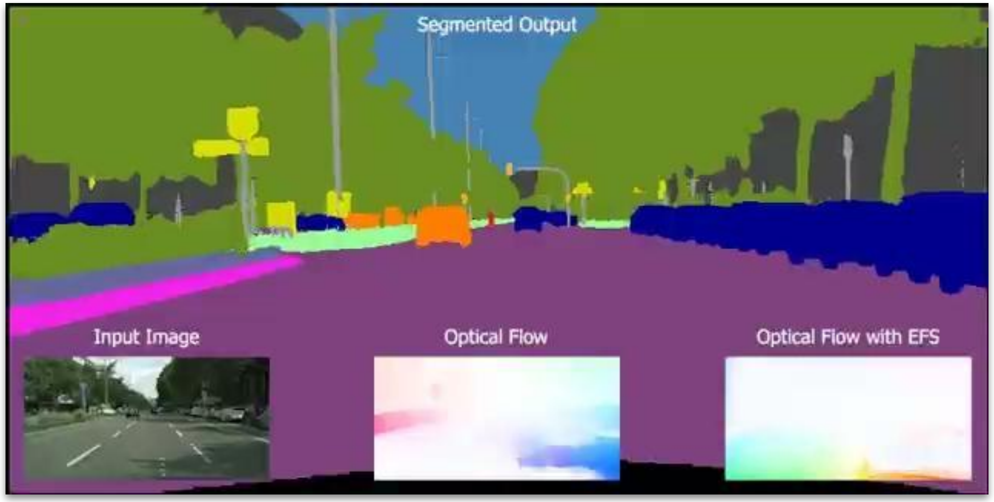

Semantic Motion Segmentation
Vertens, Valada, Burgard
ICRA ‘17

Robotic Perception — Manipulation

Single-Shot Reconstruction

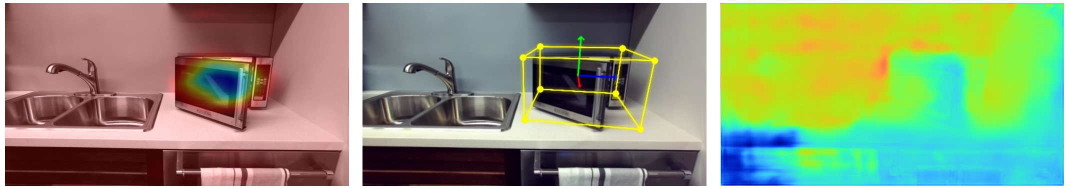

Heatmap | Poses and Bounding Box | Auxiliary Depth Prediction

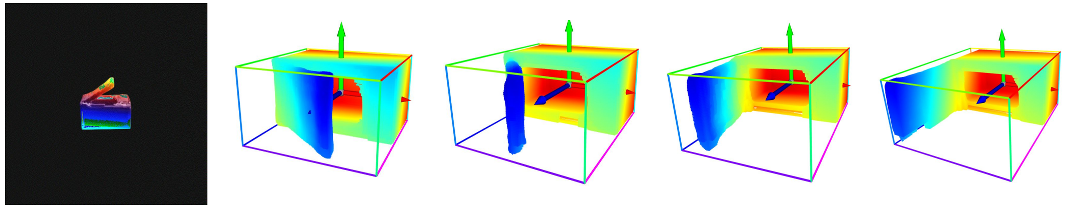

Category and Joint Agnostic Reconstruction of ARTiculated Objects
Heppert, et al
CVPR ‘23

Learning scale-invariant compact representations for mobile manipulation

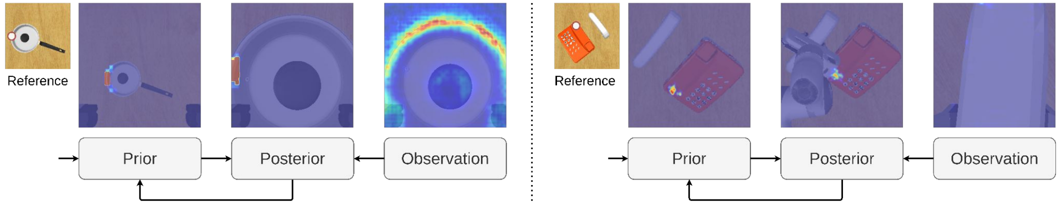

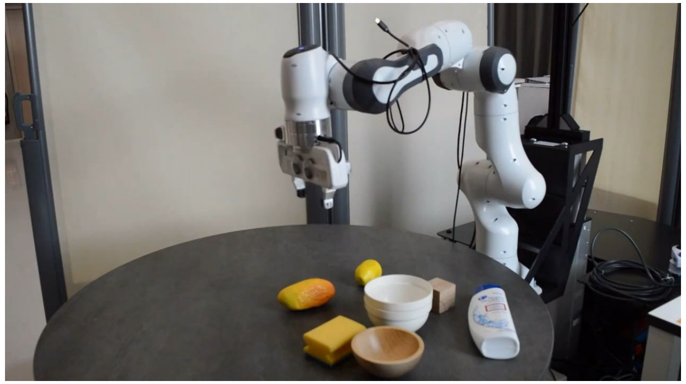

Bayesian Scene Keypoints for Deep Policy Learning in Robotic Manipulation
von Hartz, et al
RA-L ‘23

Mapping and Localization

Continual SLAM

Vödisch, Cattaneo, Burgard, Valada

ISSR ‘22

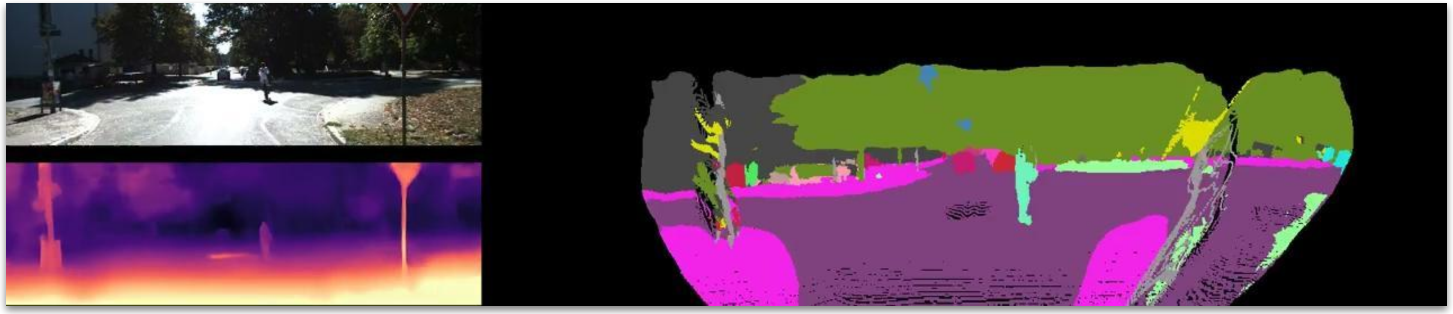

Continual Depth Estimation and Segmentation

Vödisch, Petek, Burgard, Valada

RSS ‘23

III.

Topics

Seminar Papers

Supervisor: Julia Hindel

Modality-Incremental Learning with Disjoint Relevance Mapping Networks for Image-based Semantic Segmentation

https://arxiv.org/abs/2411.17610

- In autonomous driving, deep learning now plays a crucial role in environment perception.
- More diverse sensors →more safety and robust perception
- This paper addresses the issue of continual learning with large domain shifts, e.g. different sensor modalities.
- They propose **Relevance Mapping Networks** to deal with catastrophic forgetting

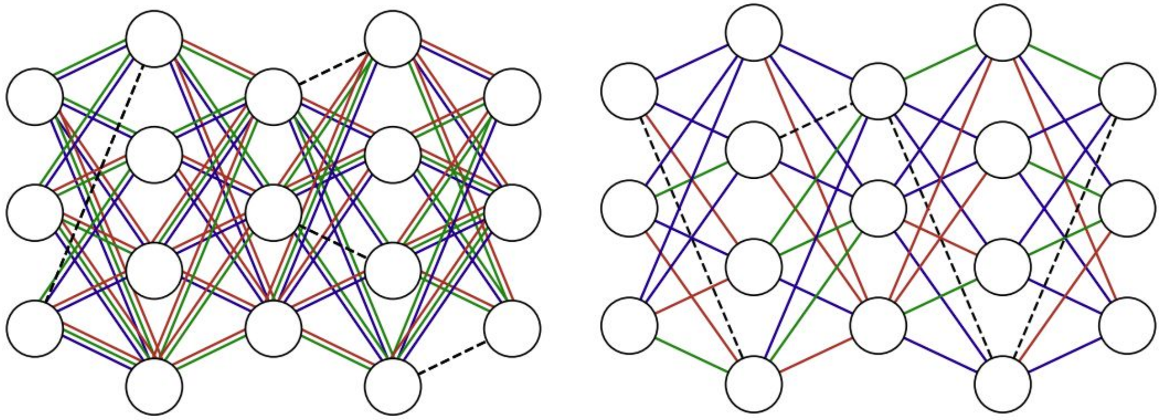

(a) Relevance Mapping Network (RMN)

(b) Disjoint Relevance Mapping Network (DRMN)

Modality A — Modality B — Modality C — Unused connection

Supervisor: Julia Hindel

Fed-EC: Bandwidth-Efficient Clustering-Based Federated Learning for Autonomous Visual Robot Navigation

https://arxiv.org/abs/2411.04112

- This paper proposes a federated learning framework that is deployed with vision based autonomous robot navigation in diverse outdoor environments
- The server:
  - clusters the robots
  - learns an aggregated model which is then shared with the robots
- They report 23x communication reduction, better transferability to new clients.

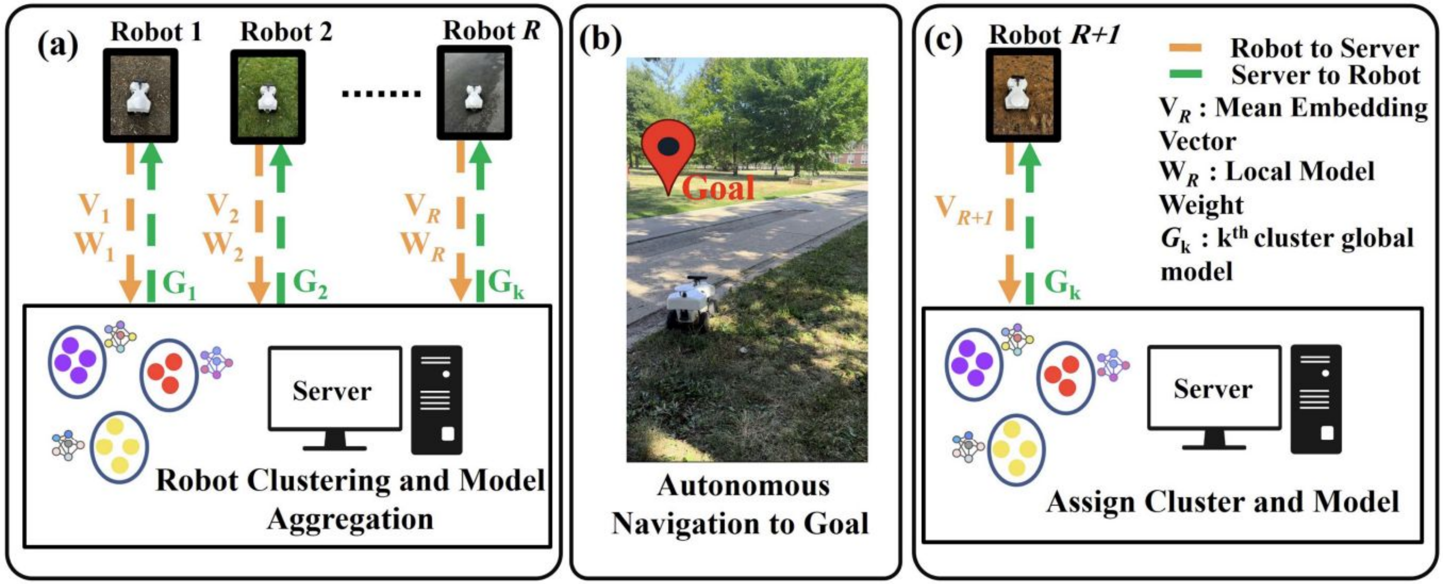

Supervisor: Julia Hindel

Privacy Risks in Reinforcement Learning for Household Robots

https://arxiv.org/abs/2306.09273

- This papers proposes an attack on the reinforcement learning training processes (value-based or gradient-based algorithms)
- They utilize **gradient inversion** to reconstruct states, actions and supervisory signals.
- Very relevant for federated learning techniques that solely utilize gradients computed on private user data to optimize models

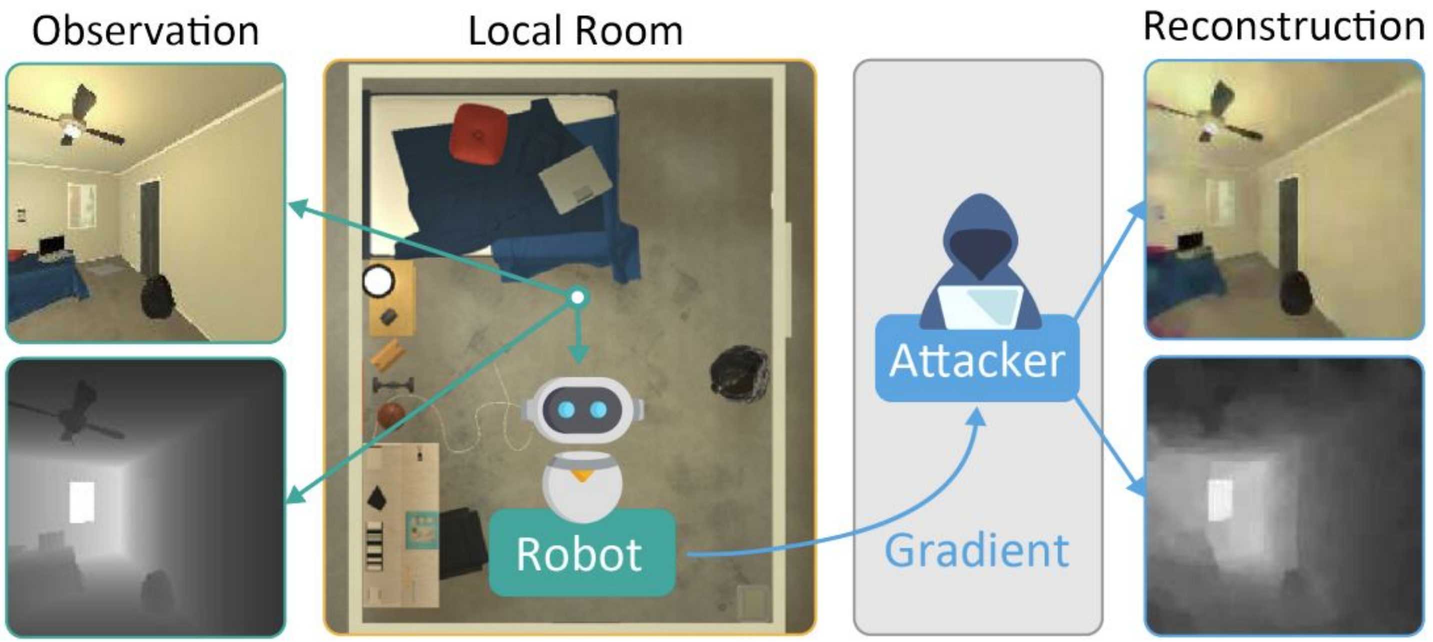

Supervisor: Julia Hindel

Continual Learning of Unsupervised Monocular Depth from Videos

https://arxiv.org/abs/2311.02393

- One of the first works to address continual learning for depth estimation
- They propose a rehearsal-based dual-memory method, which utilizes spatio-temporal consistency for continual learning in depth estimation

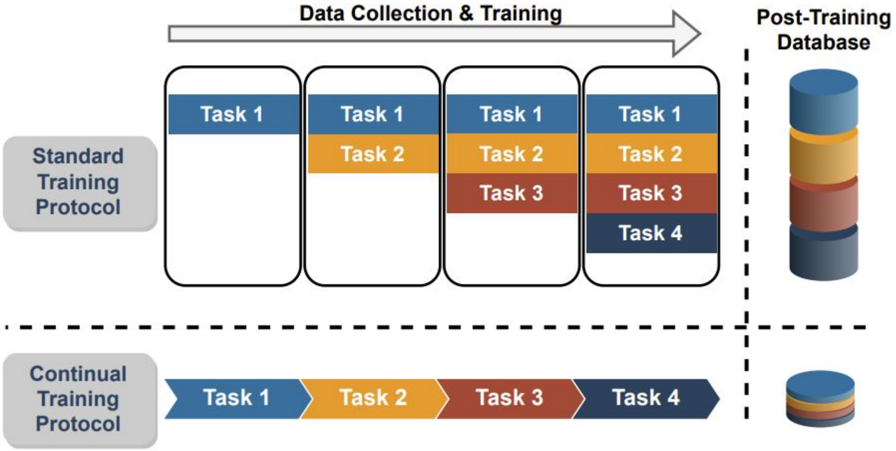

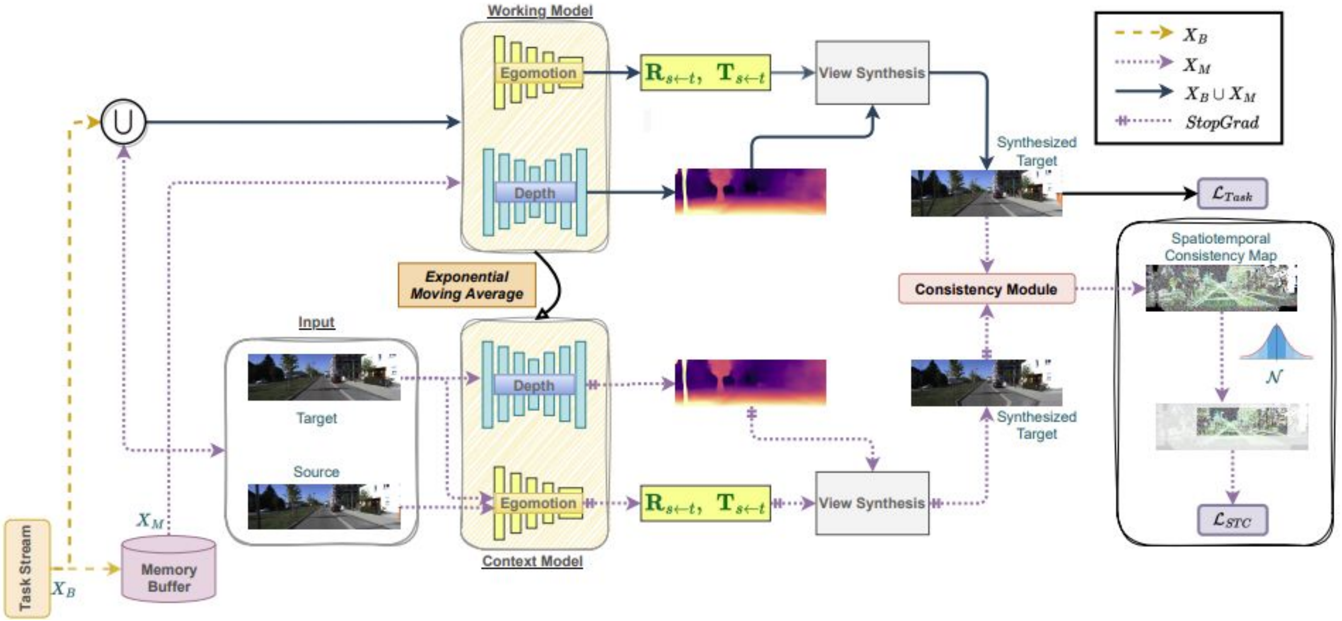

Supervisor: Julia Hindel

Reducing Non-IID Effects in Federated Autonomous Driving with Contrastive Divergence Loss

https://arxiv.org/abs/2303.06305

- Autonomous driving data is generally non-IID.
  - Significant distribution shift
  - Accumulation of errors
- This has many negative effects on the convergence of the learning process
- This paper proposes a “contrastive divergence” loss to mitigate these effects in a federated learning setting.

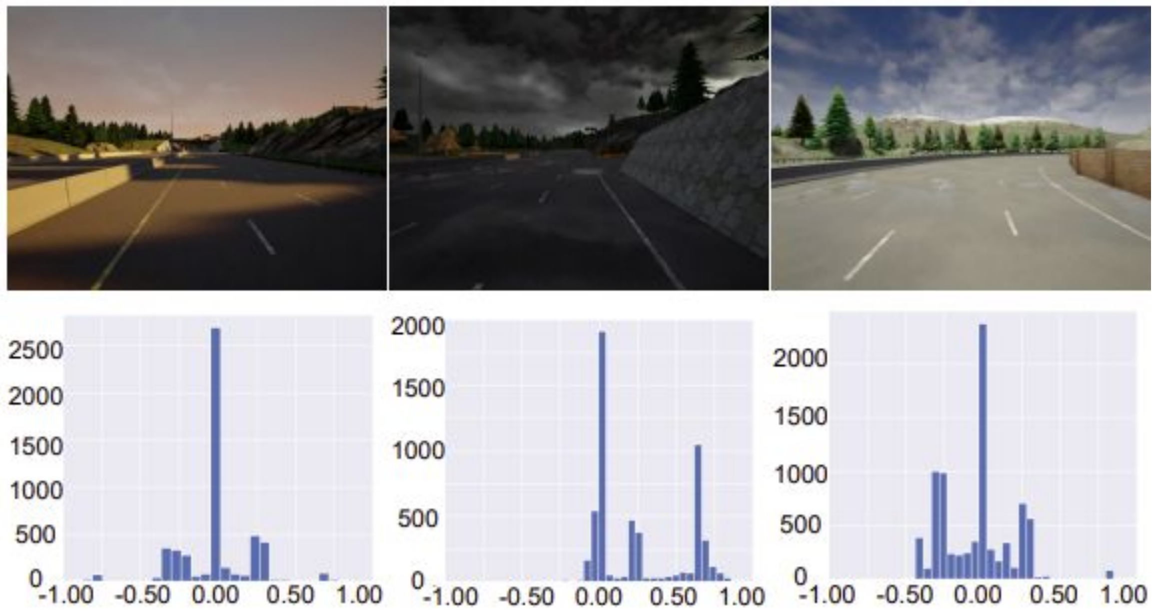

Supervisor: Martin Büchner

HOMER: Learning In-the-Wild Mobile Manipulation via Hybrid Imitation and Whole-Body Control

https://arxiv.org/abs/2506.01185

- HOMER which combines a **hybrid IL agent** with a fast, kinematics-based **whole-body controller** for sample-efficient, generalizable mobile manipulation in-the-wild
  - Data Collection (Whole-Body Teleop)
  - Data Annotation
  - Novel Policy Architecture

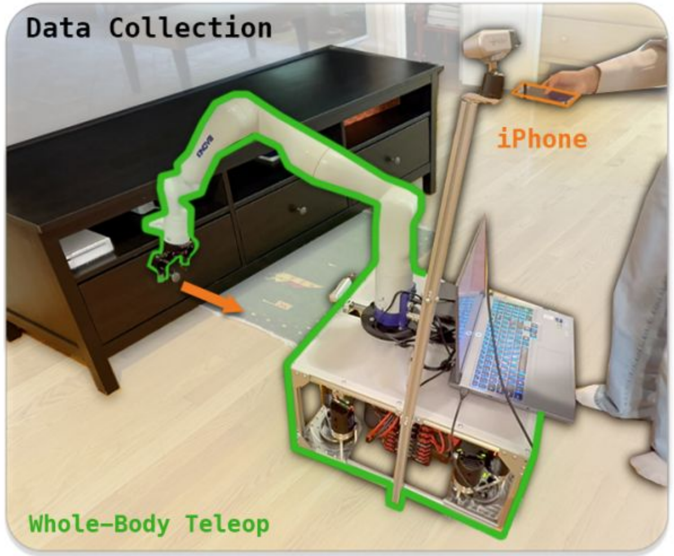

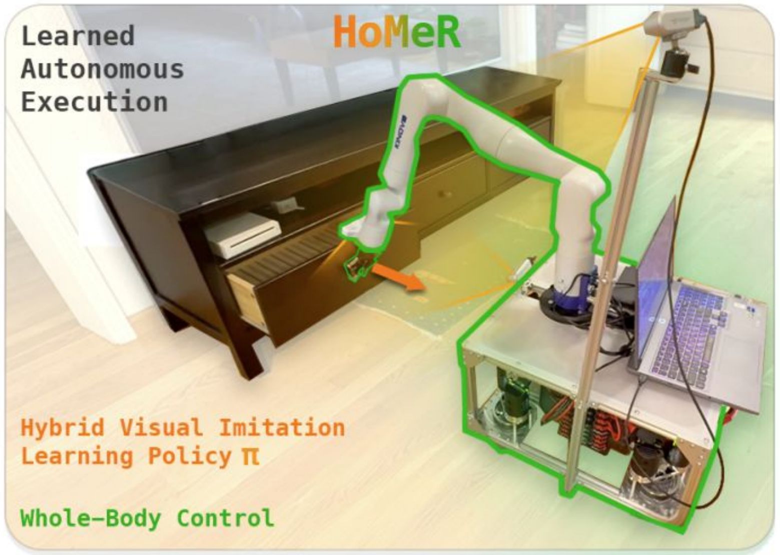

Supervisor: Martin Büchner

VGGT-SLAM: Dense RGB SLAM Optimized on the SL(4) Manifold

https://arxiv.org/abs/2505.12549

- VGGT-SLAM is a dense RGB SLAM system
  - Requires uncalibrated cameras with no assumptions on the camera motion
  - No assumptions on scene structure
  - Optimising over the SL(4) manifold
  - Better than VGGT!

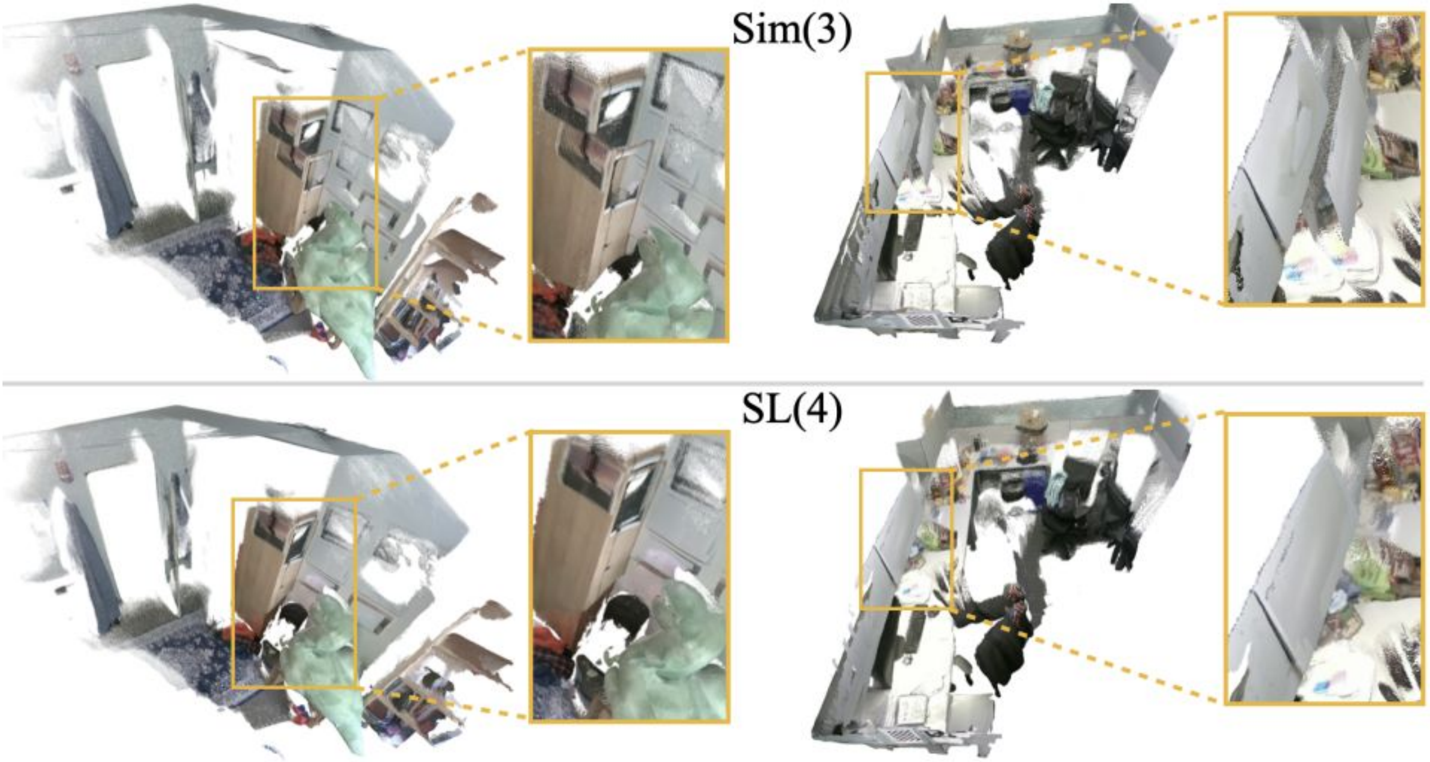

Supervisor: Martin Büchner

Planning from Point Clouds over Continuous Actions for Multi-object Rearrangement

https://arxiv.org/abs/2509.04645

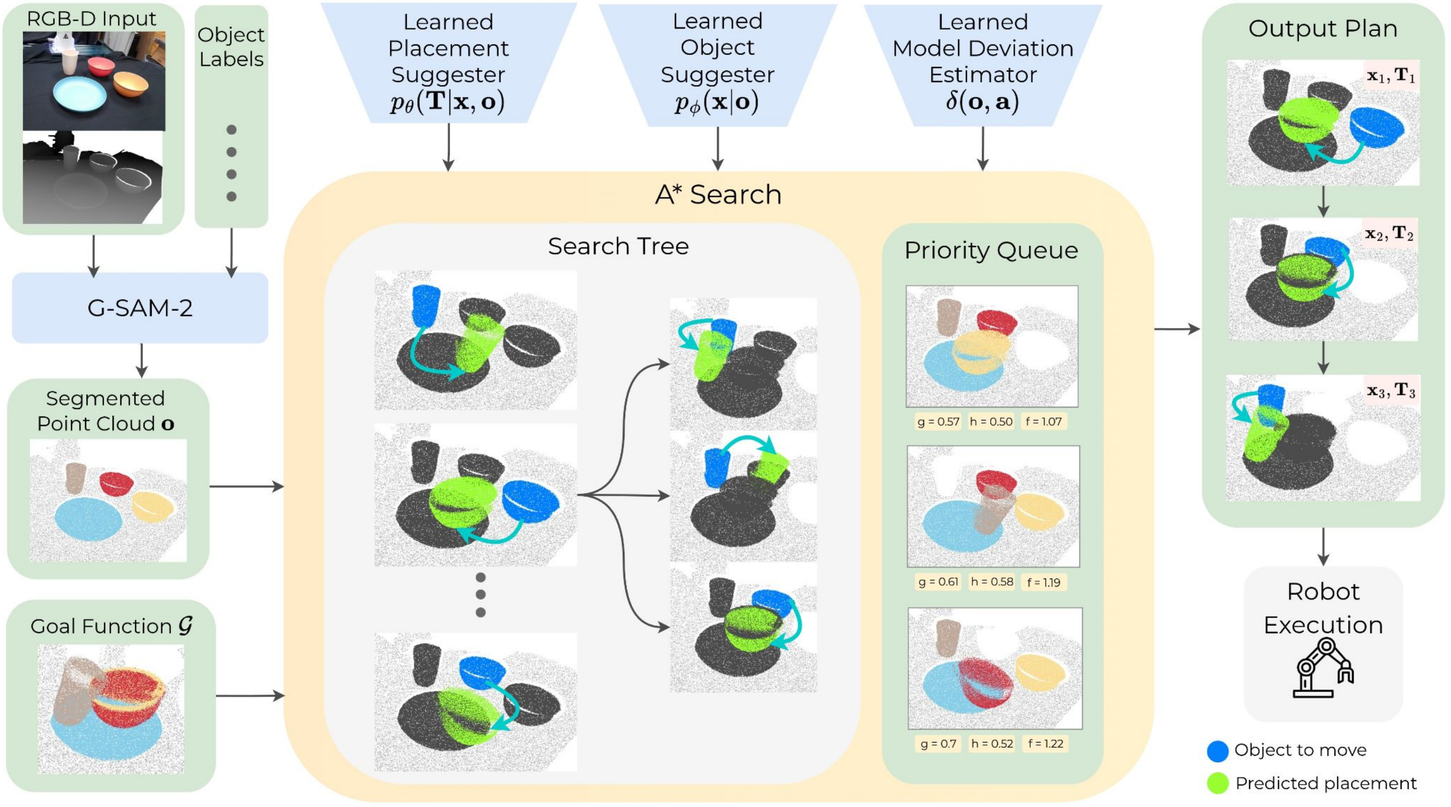

Supervisor: Martin Büchner

ScrewSplat: An End-to-End Method for Articulated Object Recognition

https://arxiv.org/abs/2508.02146

Supervisor: Martin Büchner

Steerable Scene Generation with Post Training and Inference-Time Search

https://arxiv.org/abs/2505.04831

- This work generates large-scale scene data using procedural models that approximate realistic environments for robotic manipulation
  - Trains a Unified Diffusion-based Generative Model
  - Goal-directed scene synthesis that respects physical feasibility

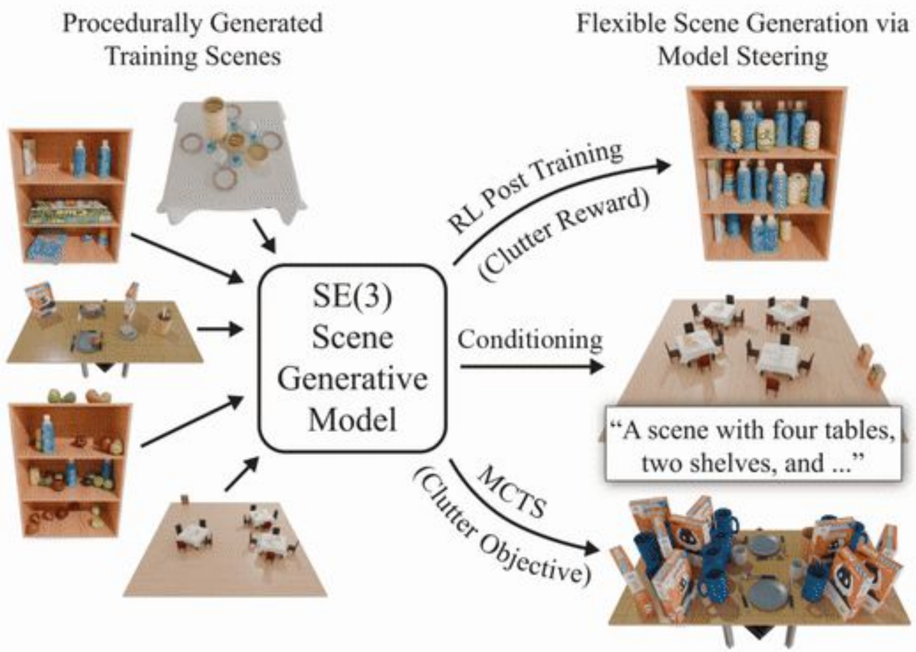

Steer scene generation toward task-specific goals (e.g., higher clutter than seen during pretraining)

Supervisor: Iman Nematollahi

V-JEPA 2: Self-Supervised Video Models Enable Understanding, Prediction and Planning

https://arxiv.org/abs/2506.09985

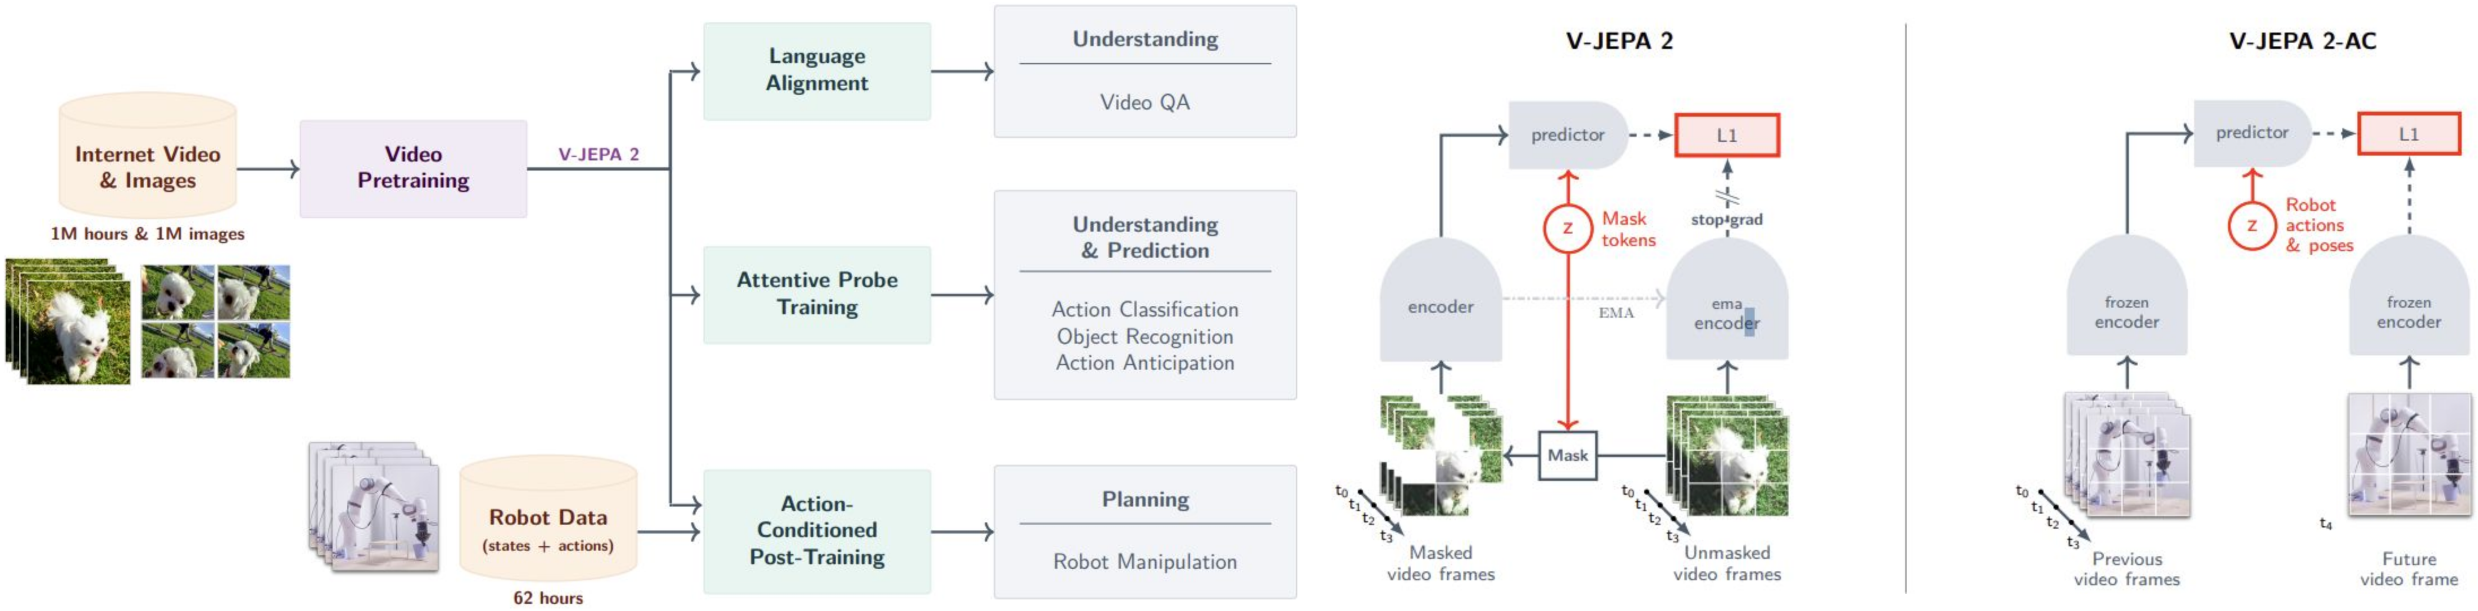

Supervisor: Iman Nematollahi

Unified World Models: Coupling Video and Action Diffusion for Pretraining on Large Robotic Datasets

https://arxiv.org/abs/2504.02792

- Unified World Models (UWM) combine **video** and **action** diffusion in one transformer

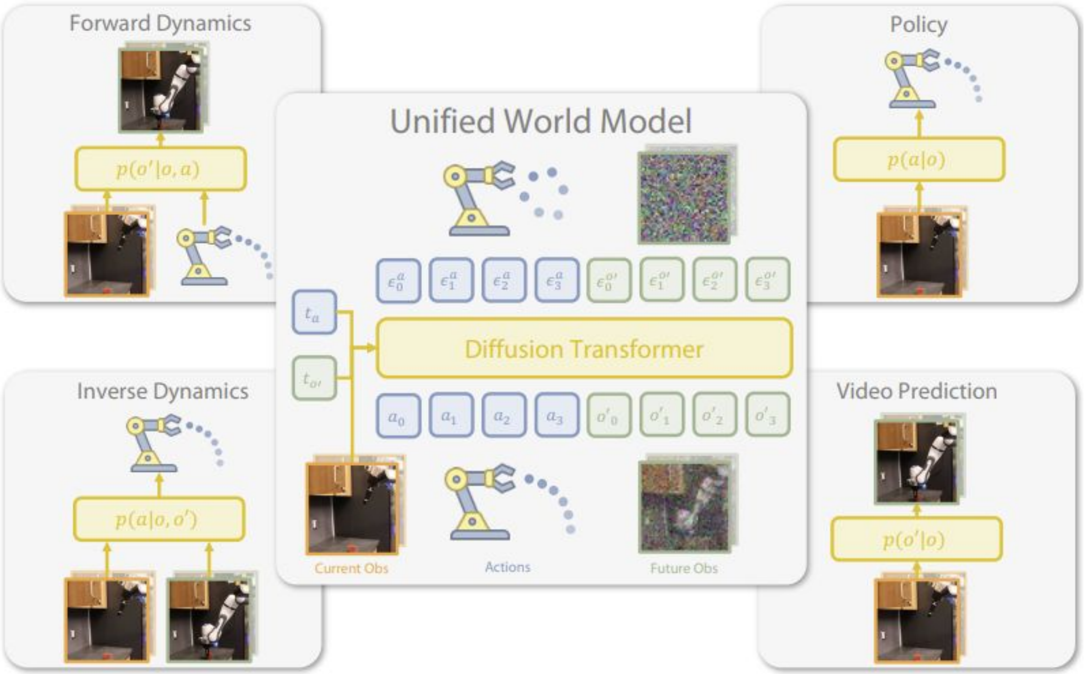

- Flexible Inference

Supervisor: Iman Nematollahi

SENSEI: Semantic Exploration Guided by Foundation Models to Learn Versatile World Models

https://arxiv.org/abs/2503.01584

- A framework for guiding the **intrinsically motivated exploration** of model-based agents through foundation models without assuming access to expert data, high-level actions, or perfect environment captions.

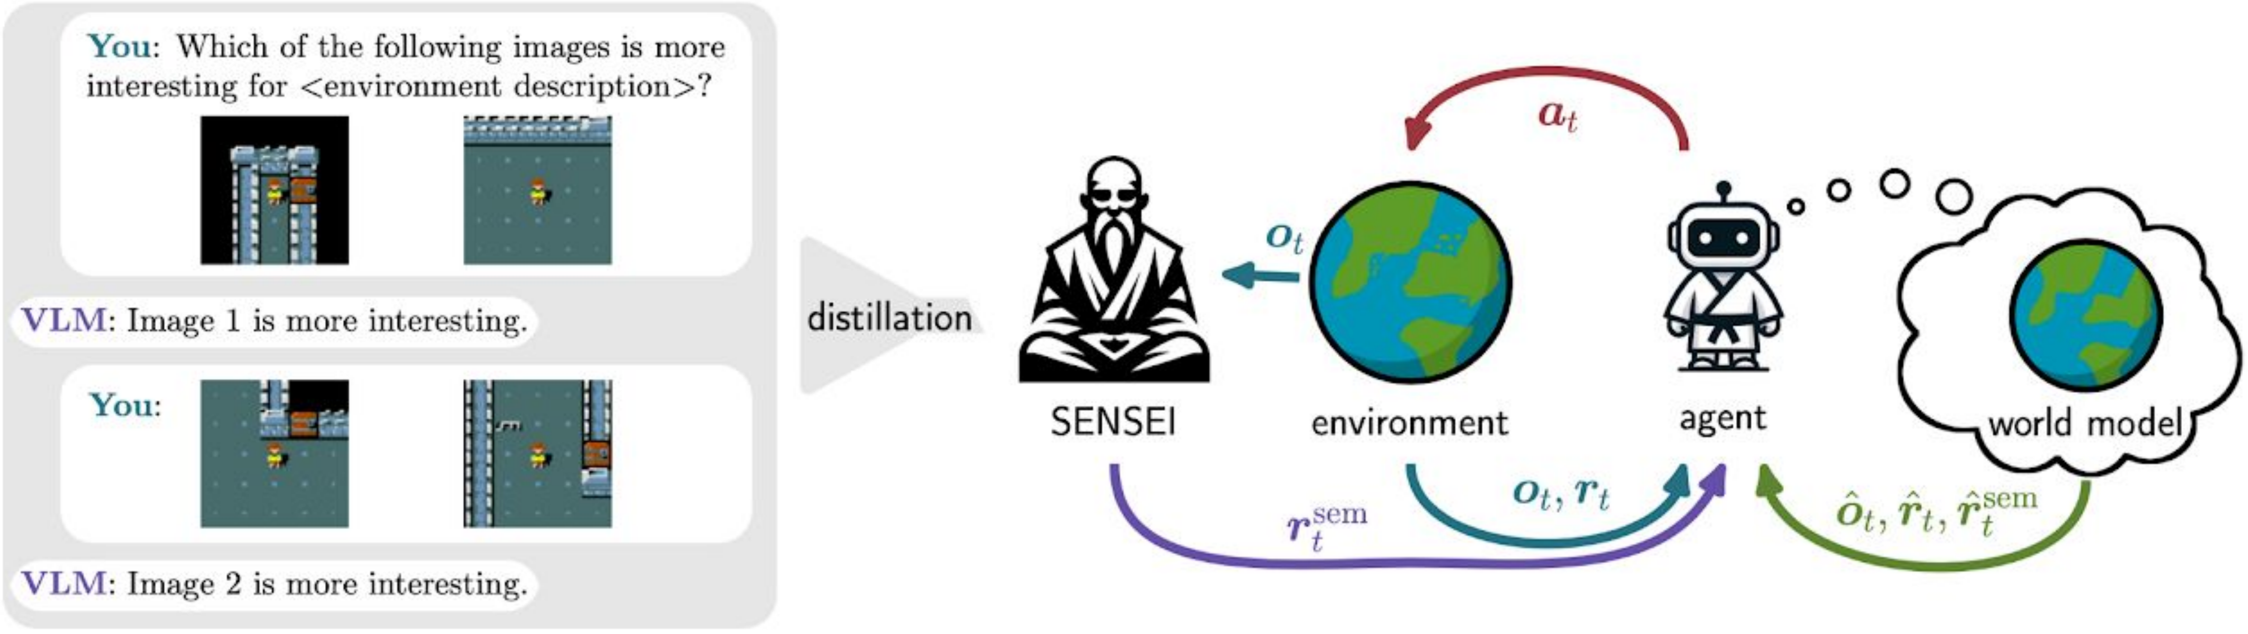

Supervisor: Iman Nematollahi

Continual Reinforcement Learning by Planning with Online World Models

https://arxiv.org/abs/2507.09177

- Continual reinforcement learning refers to a naturalistic setting where an agent needs to endlessly evolve, by trial and error.
- However, catastrophic forgetting is a known well-documented issue.
- This papers plans with online world models to mitigate the problem.
  - Key insight: a unified dynamics model
  - Provides theoretical guarantees of “no-regret” under mild assumptions

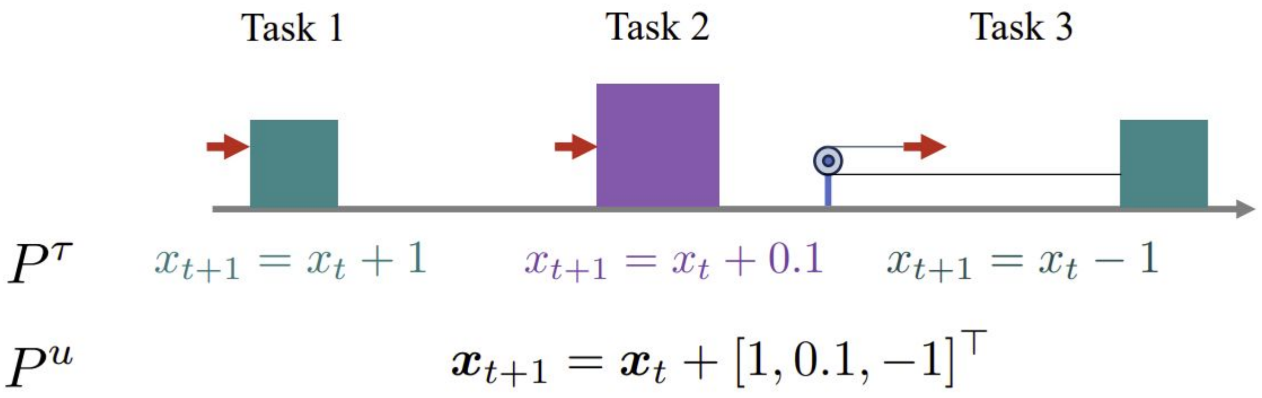

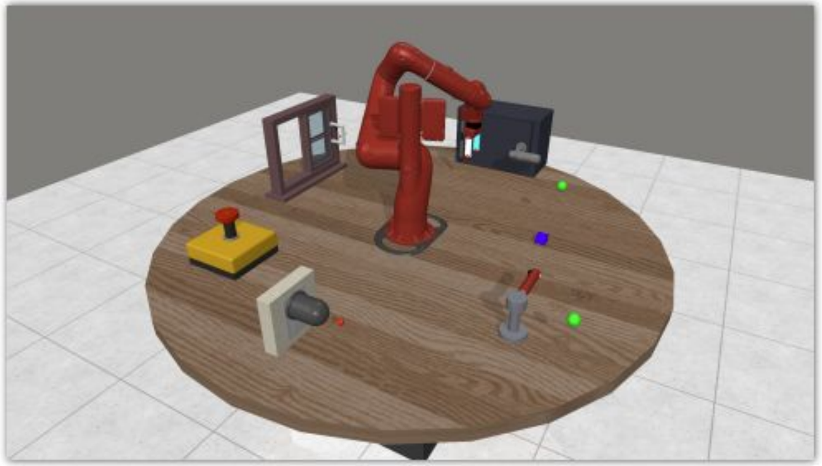

Supervisor: Iman Nematollahi

Mobi-π: Mobilizing Your Robot Learning Policy

https://arxiv.org/abs/2505.23692

- This work introduces policy mobilization
- Mobi-π finds the optimal robot pose that leads to an in-distribution viewpoint for executing a policy π
- Achieves task success without the need to collect additional demonstrations

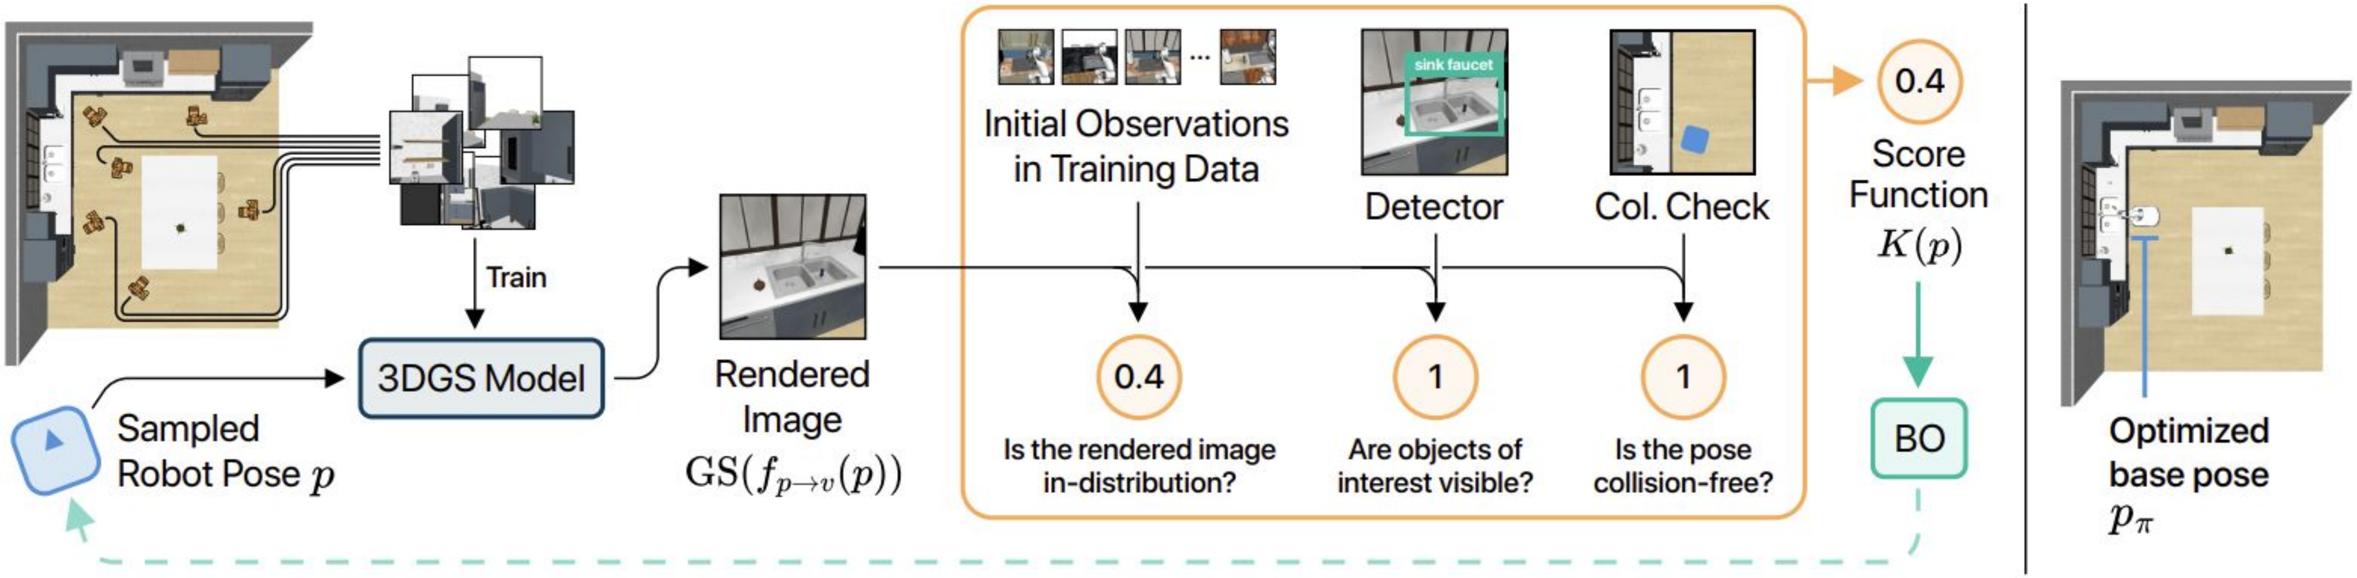

Supervisor: Akshay L Chandra

Steering Your Diffusion Policy with Latent Space Reinforcement Learning

https://arxiv.org/abs/2506.15799

- Diffusion Policies are great!
  - Multi-modality, stable training and more.
- However, imitation learners have problems.
- To improve them beyond BC:
  - Use Reinforcement Learning.
  - [This Work] Simply steer them in the latent space!
- Highly sample efficient.
- Only requires black-box access to the diffusion policy.

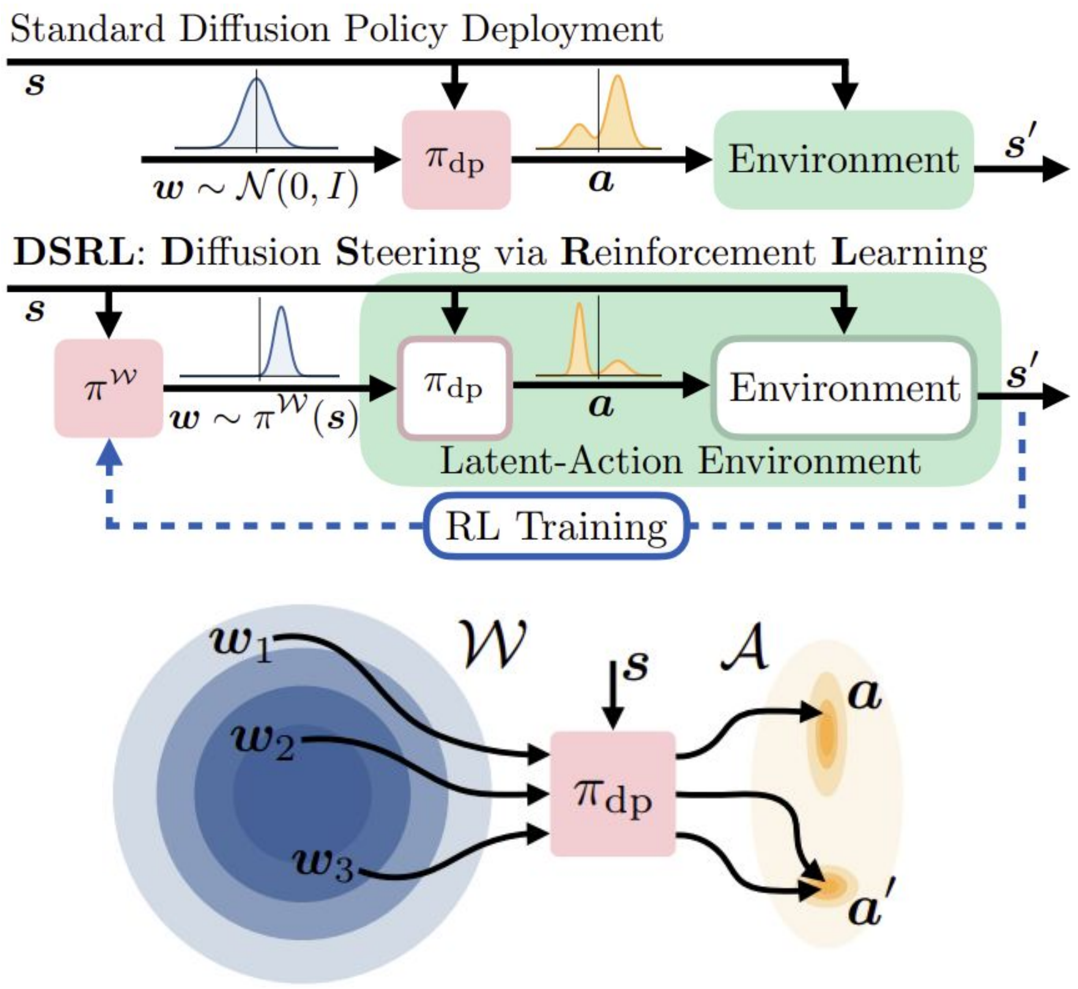

Supervisor: Akshay L Chandra

VLA-RL: Towards Masterful and General Robotic Manipulation with Scalable Reinforcement Learning

https://arxiv.org/abs/2505.18719

- VLAs are great!
- At the end of the day, they are merely imitation learners.
- To improve:
  - Use Reinforcement Learning
  - With Sparse Rewards
- This work: Explores a wide variety of recipes for fine-tuning VLAs with RL
  - Trajectory generation
  - Value functions
  - LoRA of Action-head

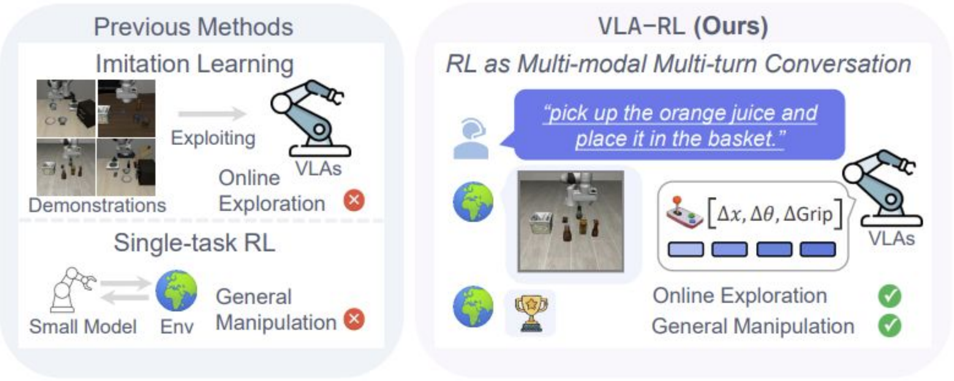

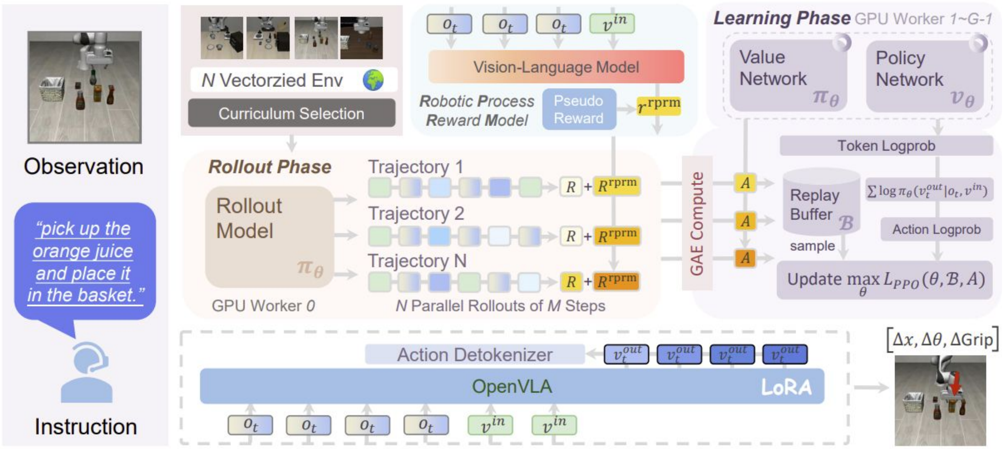

Supervisor: Akshay L Chandra

RoboMonkey: Scaling Test-Time Sampling and Verification for Vision-Language-Action Models

https://arxiv.org/abs/2506.17811

- VLAs are great!
- Can we squeeze the best out of them before fine-tuning them?

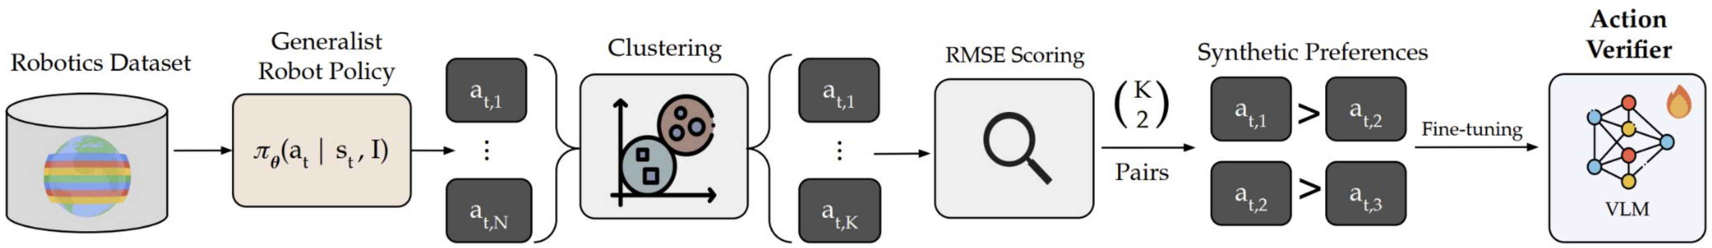

- This work: Test-Time Scaling!

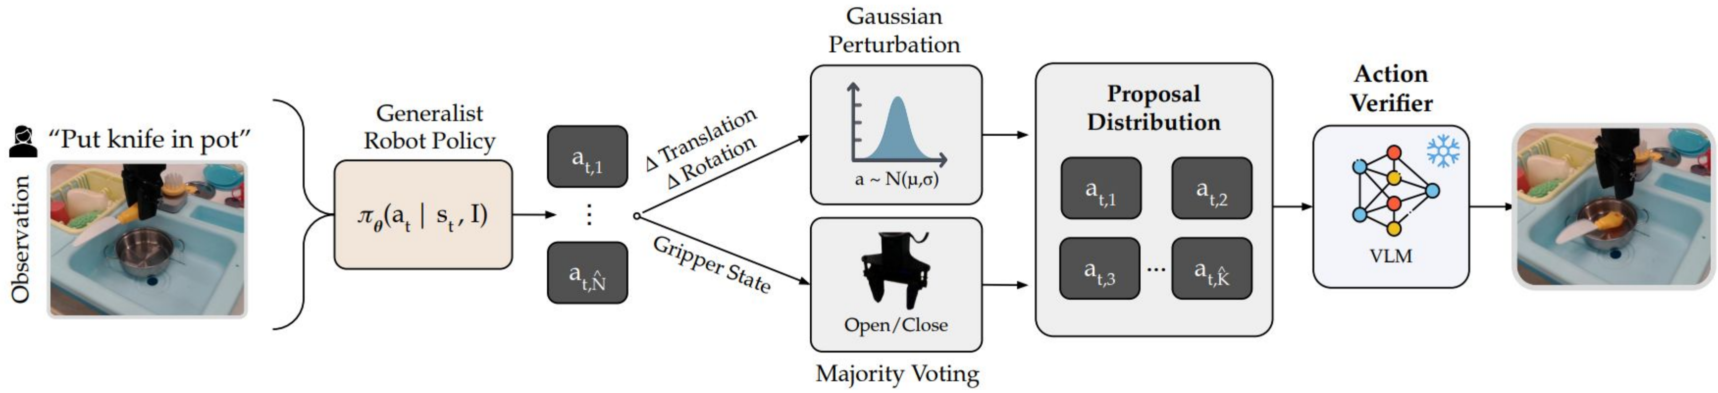

Supervisor: Akshay L Chandra

Deep Reactive Policy: Learning Reactive Manipulator Motion Planning for Dynamic Environments

https://arxiv.org/abs/2509.06953

- Motion planning is great!
  - We have really fast IK solvers now!
  - But they cannot generalise!
- This work: Neural motion policies that can generalise in complex or dynamic settings
  - A transformer trained on **10 million** generated expert trajectories across diverse simulation scenarios!

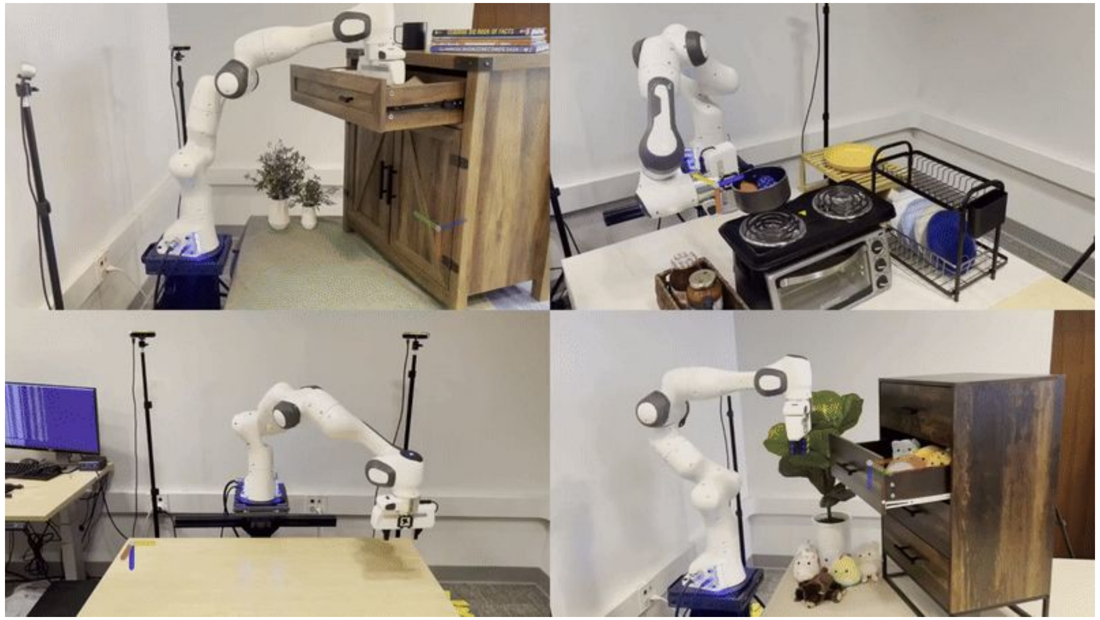

Supervisor: Akshay L Chandra

ReinFlow: Fine-tuning Flow Matching Policy with Online Reinforcement Learning

https://arxiv.org/abs/2505.22094

4.1 A Two-Layer "Diffusion Policy MDP"

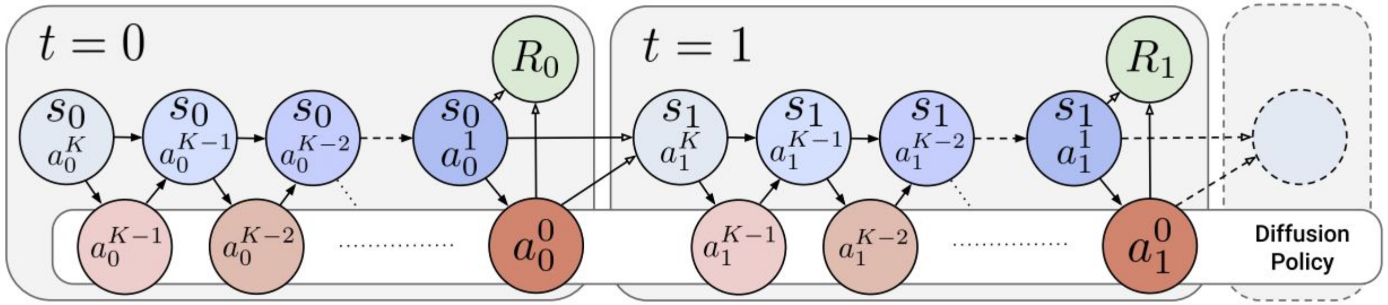

Source: Ren et al., DPPO, 2025

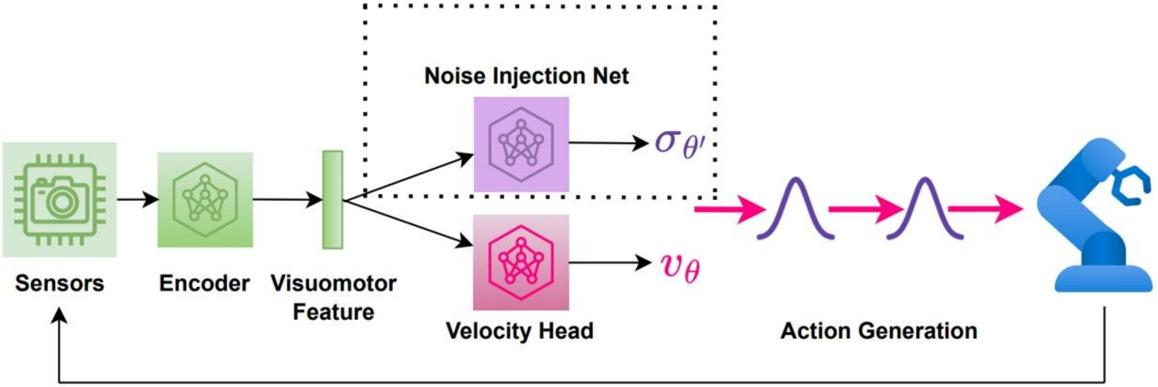

Questions?

Open Positions

- We have multiple **MSc Project** and **Thesis** topics related to many directions of robot learning.

Please check our website for information on how to apply:

https://rl.uni-freiburg.de/open-positions

Questions or Comments

Akshay L Chandra

Robot Learning Lab

chandra@cs.uni-freiburg.de