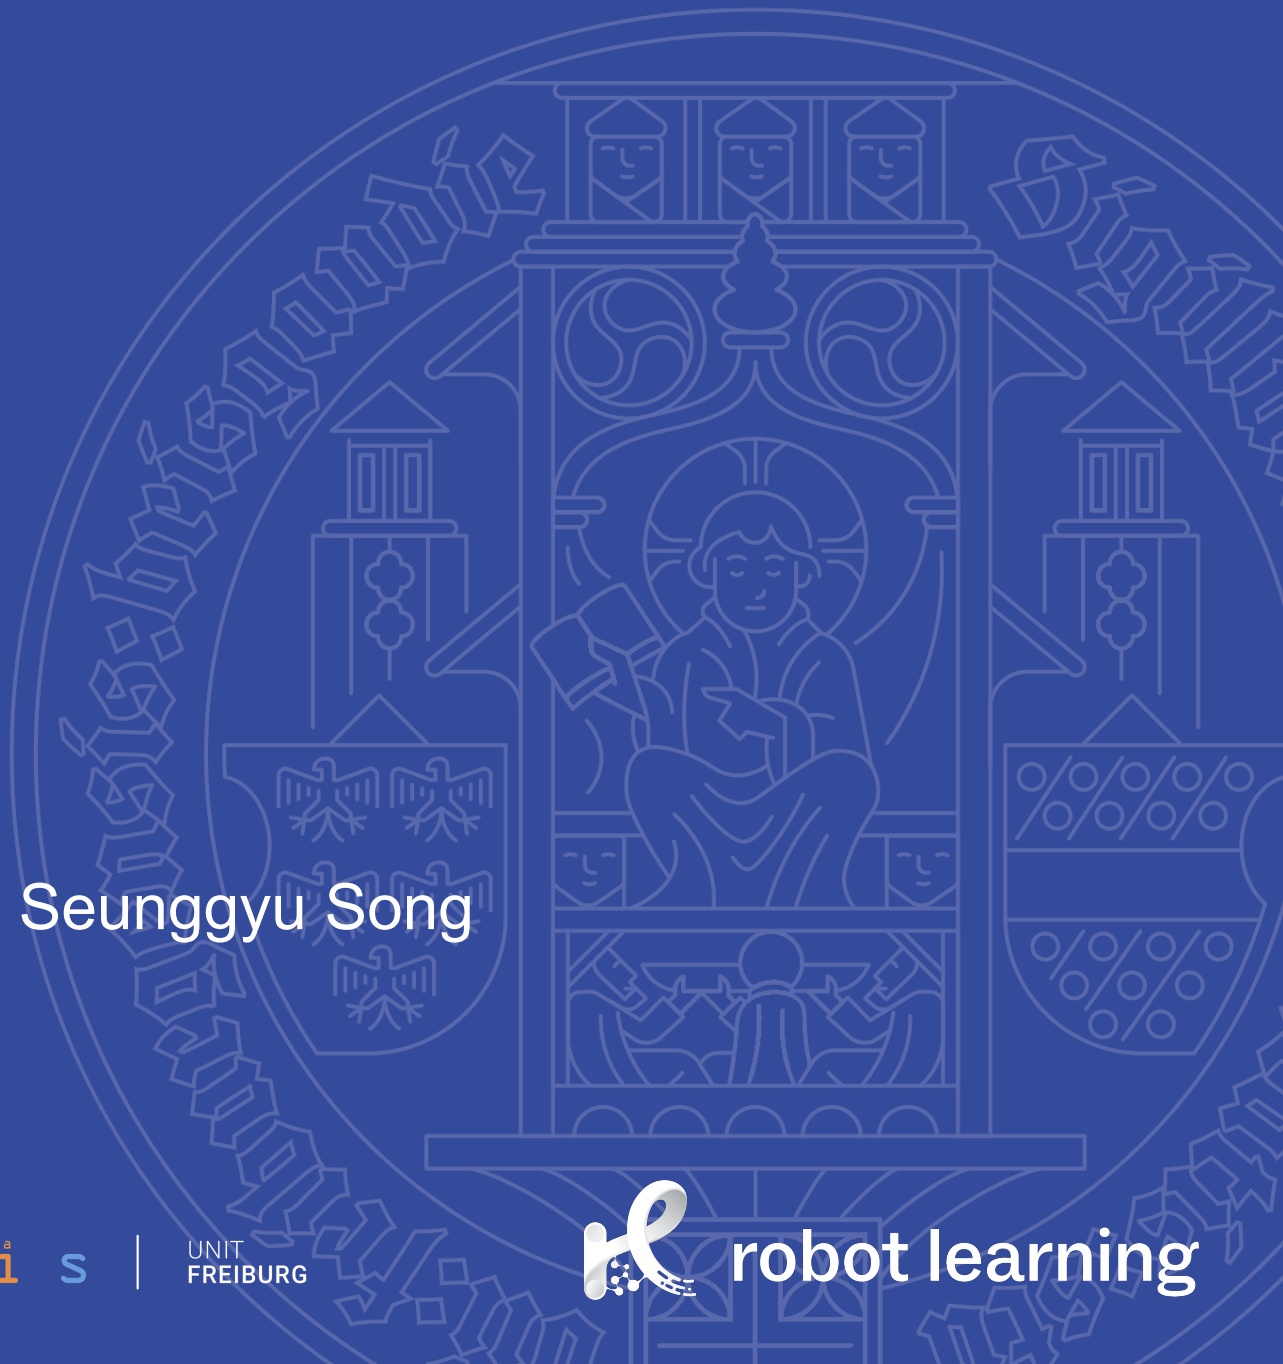
universität freiburg

WS 25/26 Seminar
Robot Learning

Prof. Dr. Abhinav Valada
Yao Lu, Akshay Chandra, Bahey Tharwat, Seunggyu Song

Robot Learning Lab

16 October 2025



ellis | UNIT FREIBURG

robot learning

Agenda

I. Organization: Content, 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.



DIE NATURWISSENSCHAFTEN

23. Jahrgang 29. November 1935 Heft 48

Die gegenwärtige Situation in der Quantenmechanik.

Von E. Schrödinger, Oxford.

Inhaltsübersicht.

§ 1. Die Physik der Modelle.
§ 2. Die Statistik der Modellvariablen in der Quantenmechanik.
§ 3. Beispiele für Wahrscheinlichkeitsvoraussagen.
§ 4. Kann man der Theorie ideale Gesamtheiten unterlegen?
§ 5. Sind die Variablen wirklich verwaschen?
§ 6. Der bewußte Wechsel des erkenntnistheoretischen Standpunktes.
§ 7. Die ψ-Funktion als Katalog der Erwartung.
§ 8. Theorie des Messens, erster Teil.
§ 9. Die ψ-Funktion als Beschreibung des Zustandes.
§ 10. Theorie des Messens, zweiter Teil.
§ 11. Die Aufhebung der Verschränkung. Das Ergebnis abhängig vom Willen des Experimentators.
§ 12. Ein Beispiel.
§ 13. Fortsetzung des Beispiels: alle möglichen Messungen sind eindeutig verschränkt.
§ 14. Die Änderung der Verschränkung mit der Zeit. Bedenken gegen die Sonderstellung der Zeit.
§ 15. Naturprinzip oder Rechenkunstgriff?

§ 1. Die Physik der Modelle.

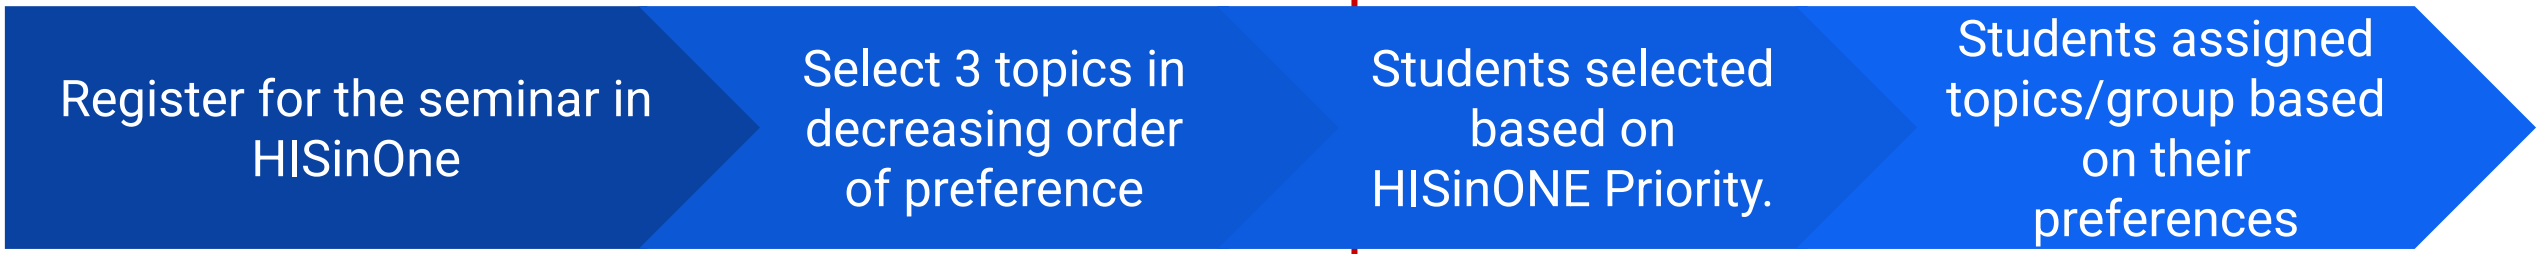
In der zweiten Hälfte des vorigen Jahrhunderts war aus den großen Erfolgen der kinetischen Gastheorie und der mechanischen Theorie der Wärme ein Ideal der exakten Naturbeschreibung hervorgewachsen, das als Krönung jahrhundertelangen Forschens und Erfüllung jahrtausendealter Hoffnung einen Höhepunkt bildet und das klassische heißt. Dieses sind seine Züge.

Von den Naturobjekten, deren beobachtetes Verhalten man erfassen möchte, bildet man, gestützt auf die experimentellen Daten, die man besitzt, aber ohne der intuitiven Imagination zu wehren, eine Vorstellung, die in allen Details genau ausgearbeitet ist, *viel* genauer als irgendwelche Erfahrung in Ansehung ihres begrenzten Umfangs je verbürgen kann. Die Vorstellung in ihrer absoluten Bestimmtheit gleicht einem mathematischen Gebilde oder einer geometrischen Figur, welche aus einer Anzahl von *Bestimmungsstücken* ganz und gar berechnet werden kann; wie z. B. an einem Dreieck eine Seite und die zwei ihr anliegenden Winkel, als Bestimmungsstücke, den dritten Winkel, die anderen zwei Seiten, die drei Höhen, den Radius des eingeschriebenen Kreises usw. mit bestimmen. Von einer geometrischen Figur unterscheidet sich die Vorstellung ihrem Wesen nach bloß durch den wichtigen Umstand, daß sie auch noch in der *Zeit* als vierter Dimension ebenso klar bestimmt ist wie jene in den drei Dimensionen des Raumes. Das heißt es handelt sich (was ja selbstverständlich ist) stets um ein Gebilde, das sich mit der Zeit verändert, das verschiedene *Zustände* annehmen kann; und wenn ein Zustand durch die nötige Zahl von Bestimmungsstücken bekannt gemacht ist, so sind nicht nur alle anderen Stücke in diesem Augenblick mit gegeben (wie oben am Dreieck erläutert), sondern ganz ebenso alle Stücke, der genaue Zustand, zu jeder bestimmten späteren Zeit; ähnlich wie die Beschaffenheit eines Dreiecks an der Basis seine Beschaffenheit an der Spitze bestimmt. Es gehört mit zum inneren Gesetz des Gebildes, sich in bestimmter Weise zu verändern, das heißt, wenn es in einem bestimmten Anfangszustand sich selbst überlassen wird, eine bestimmte Folge von Zuständen kontinuierlich zu durchlaufen, deren jeden es zu ganz bestimmter Zeit erreicht. Das ist seine Natur, das ist die Hypothese, die man, wie ich oben sagte, auf Grund intuitiver Imagination setzt.

Natürlich ist man nicht so einfältig zu denken, daß solchermaßen zu erraten sei, wie es auf der Welt wirklich zugeht. Um anzudeuten, daß man das nicht denkt, nennt man den präzisen Denkbehelf, den man sich geschaffen hat, gern ein *Bild* oder ein *Modell*. Mit seiner nachsichtslosen Klarheit, die ohne Willkür nicht herbeizuführen ist, hat man es lediglich darauf abgesehen, daß eine ganz bestimmte Hypothese in ihren Folgen geprüft werden kann, ohne neuer Willkür Raum zu geben während der langwierigen Rechnungen, durch die man Folgerungen ableitet. Da hat man gebundene Marschroute und errechnet eigentlich nur, was ein kluger Hans aus den Daten direkt herauslesen würde! Man weiß dann wenigstens, wo die Willkür steckt und wo man zu bessern hat, wenn's mit der Erfahrung nicht stimmt: in der Ausgangshypothese, im Modell. Dazu muß man stets bereit sein. Wenn bei vielen verschiedenartigen Experimenten das Naturobjekt sich wirklich so benimmt wie das Modell, so freut man sich und denkt, daß unser Bild in den wesentlichen Zügen der Wirklichkeit gemäß ist. Stimmt es bei einem neuartigen Experiment oder bei Verfeinerung der Meßtechnik nicht mehr, so ist nicht gesagt, daß man sich *nicht* freut. Denn im Grunde ist das die Art, wie allmählich eine immer bessere Anpassung des Bildes, das heißt unserer Gedanken, an die Tatsachen gelingen kann.

Die klassische Methode des präzisen Modells hat den Hauptzweck, die unvermeidliche Willkür in den Annahmen sauber isoliert zu halten, ich möchte fast sagen wie der Körper das Keimplasma, für den historischen Anpassungsprozeß an die fortschreitende Erfahrung. Vielleicht liegt der

Enrollment Procedure



Fill in our **Google Form**

By **26.10.2026**

Please check the course website for more information:

https://rl.uni-freiburg.de/teaching/ws26/robotlearning/

Important Dates

Event	Date	Time
Lecture 1: Introduction *	22.10.2026	13:00 - 14:30
HISinOne Registration + Paper Selection	26.10.2026	
Place Allocation	29.10.2026	
Topic/Group Assignment	30.10.2026	
Supervisor Meeting	contact your supervisor	
Lecture 2: How to Make a Presentation *	07.12.2026	13:00 - 14:30
Lecture 3: Block Seminar Presentations *	04.02.2027	9:00-17:00
Paper Summary Submission	11.02.2027	< 23:59

* **Mandatory in-person attendance**

Evaluation Criteria

Evaluation	Due Date
Seminar Presentation	04.02.2027
Paper Summary	11.02.2027

- **Presentation:** at most **30 min** per **group**. (10 min per group member)
- **Summary:** at most **7 pages** excluding bibliography and figures.
- **Final grade:**
  - Literature Review (relevance and justification of additional paper(s)) **(group)**
  - Presentation (slides & delivery) **(group)**
  - Seminar Participation **(individual)**
  - Summary **(individual)**

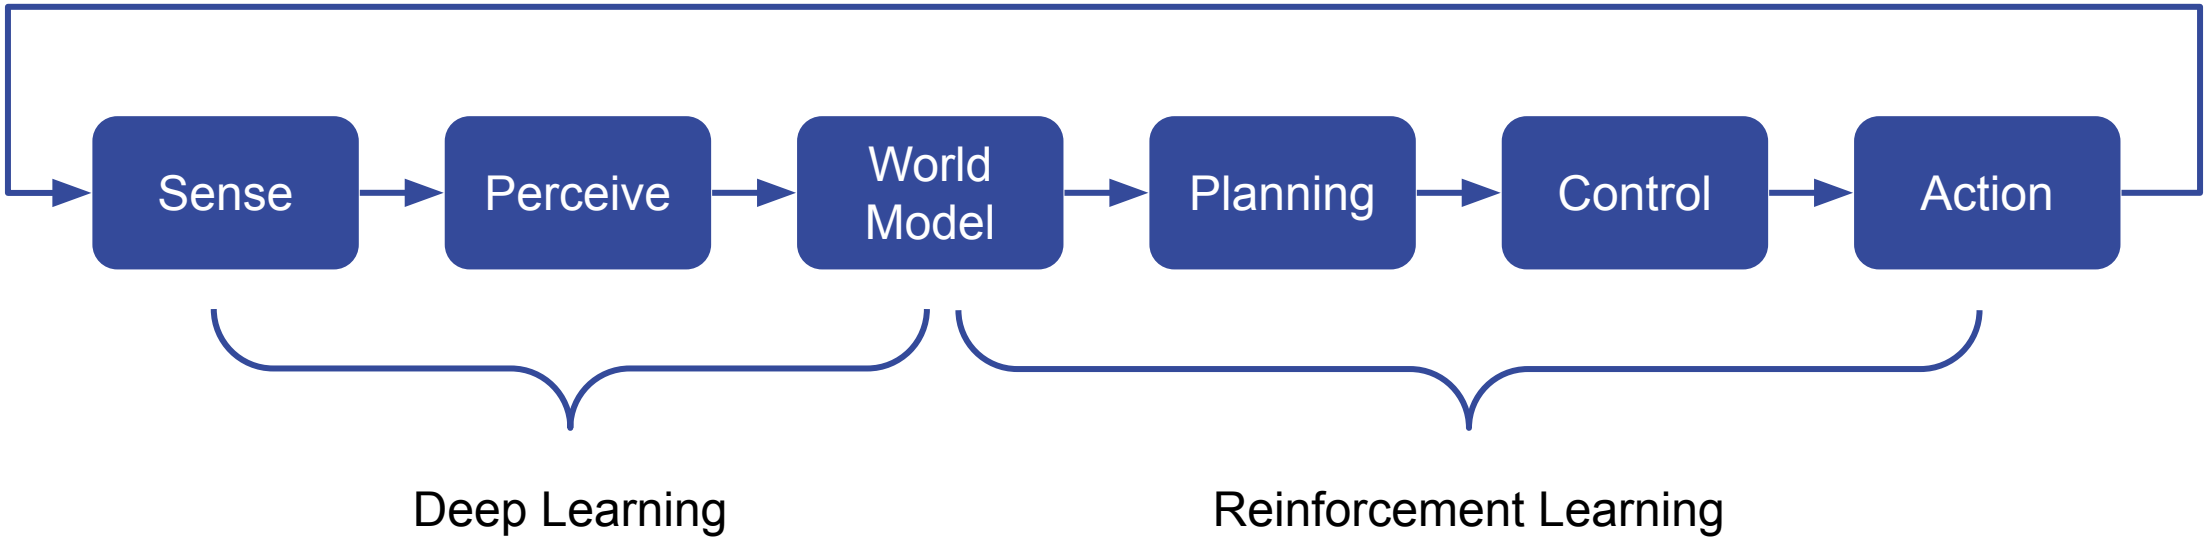
II.

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?

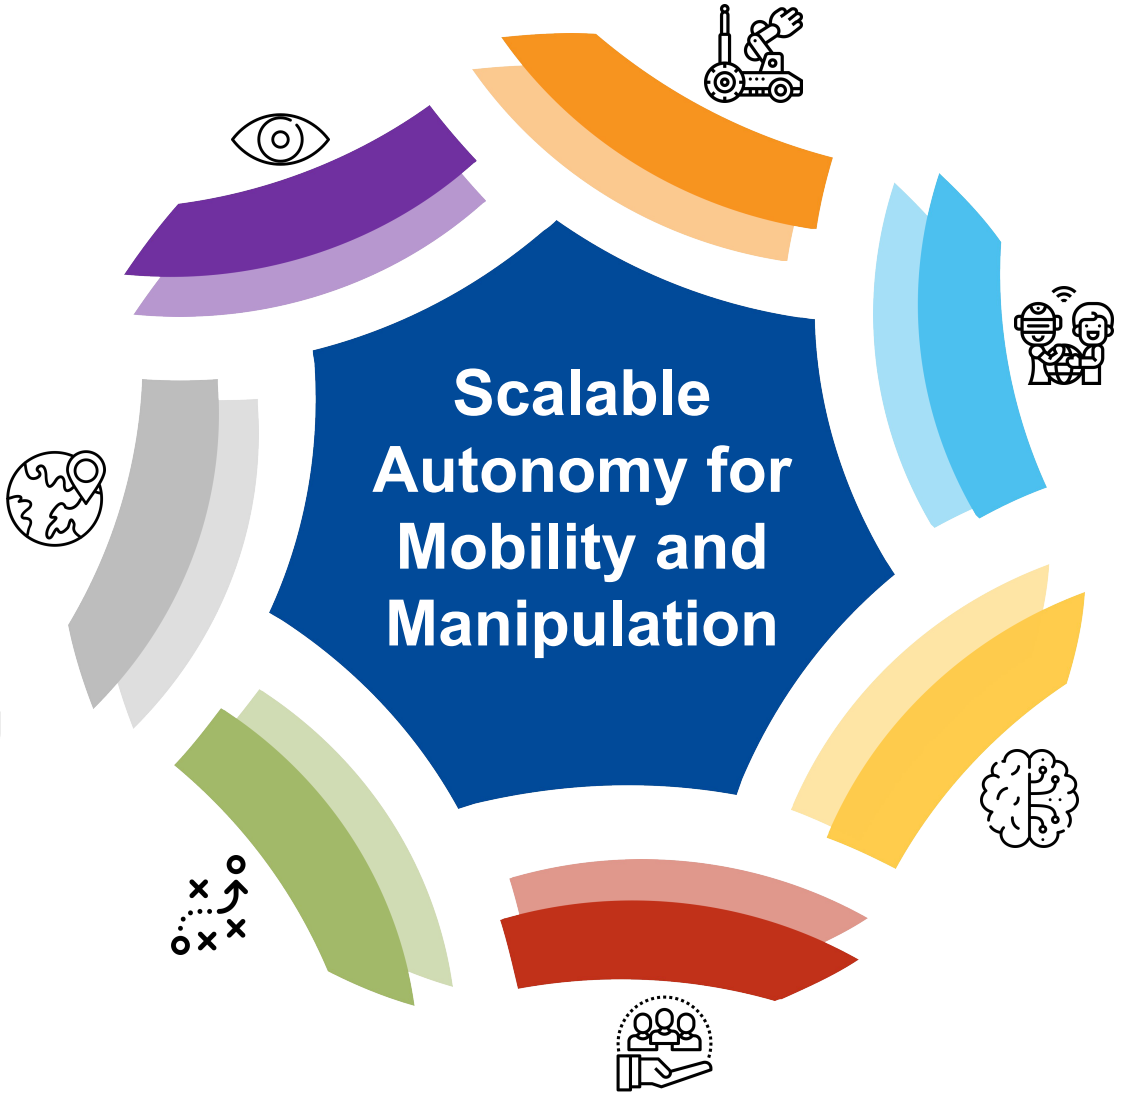
Research Areas

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

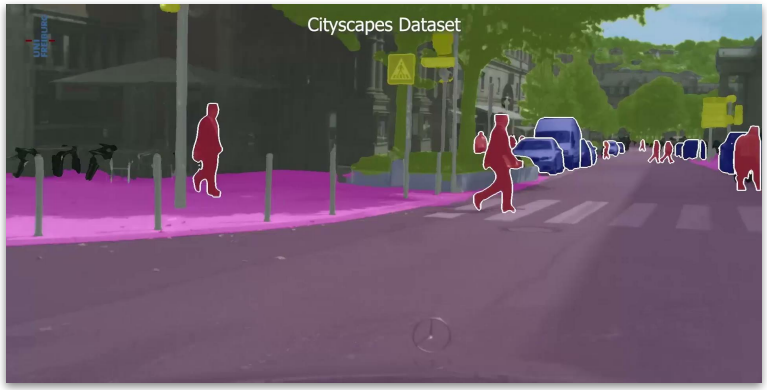


Mobile Manipulation
- Whole-Body Motion
- Long-Horizon Reasoning
- Planning for Sensing

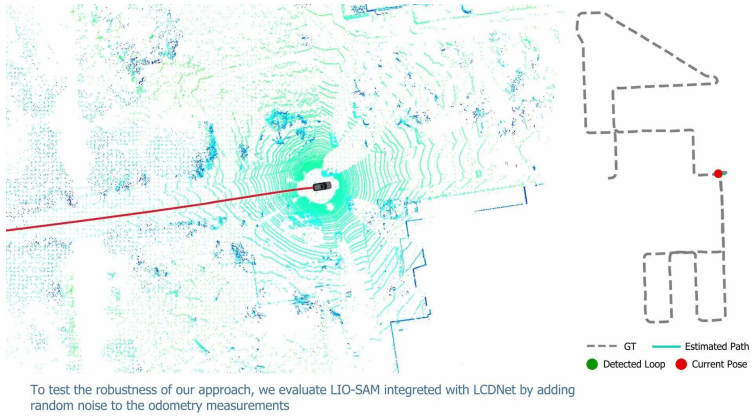
Human-Robot Interaction
- Socially-Compliant Behavior
- Human-Robot Collaboration
- Behavior Adaptation & Safety

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

Many Seminal Works

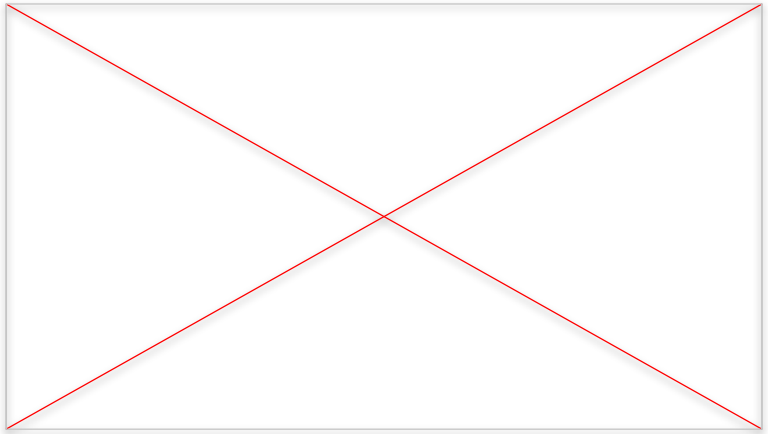


Scene Understanding

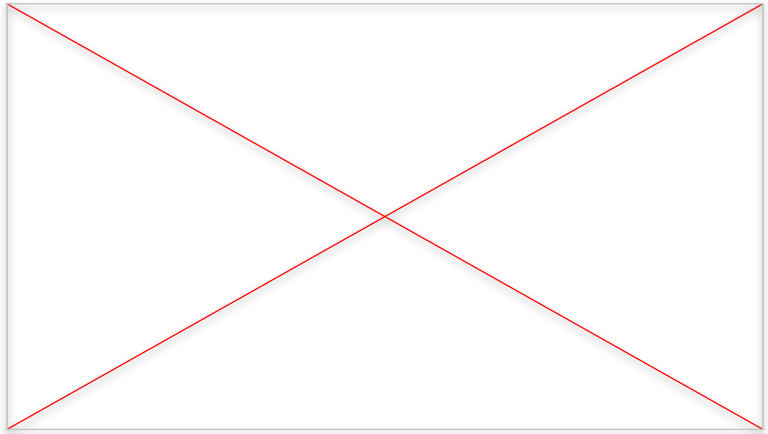


To test the robustness of our approach, we evaluate LIO-SAM integreted with LCDNet by adding random noise to the odometry measurements

Simultaneous Localization and Mapping

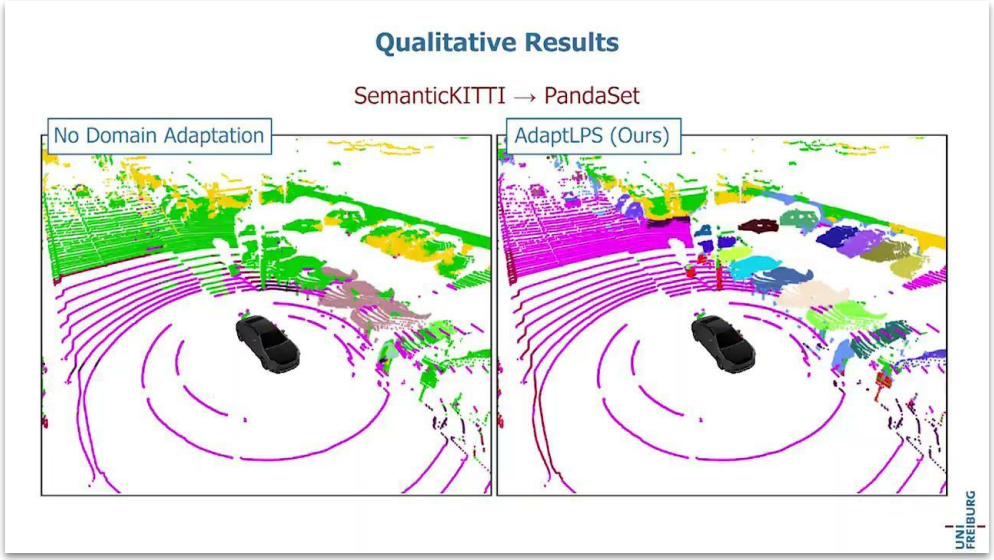


Motion Planning



Learning from Demonstrations

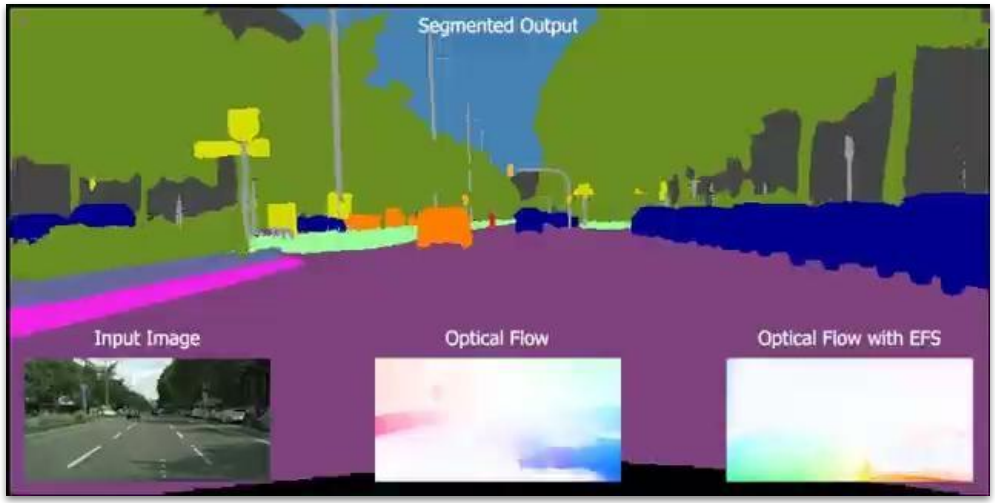
Robotic Perception — Mobility



Unsupervised LiDAR Domain Adaptation

Besic, Gosala, Cattaneo, Valada

RA-L ‘22



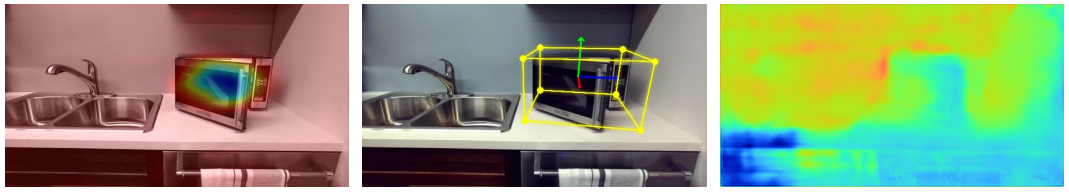
Semantic Motion Segmentation

Vertens, Valada, Burgard

ICRA ‘17

Robotic Perception — Manipulation

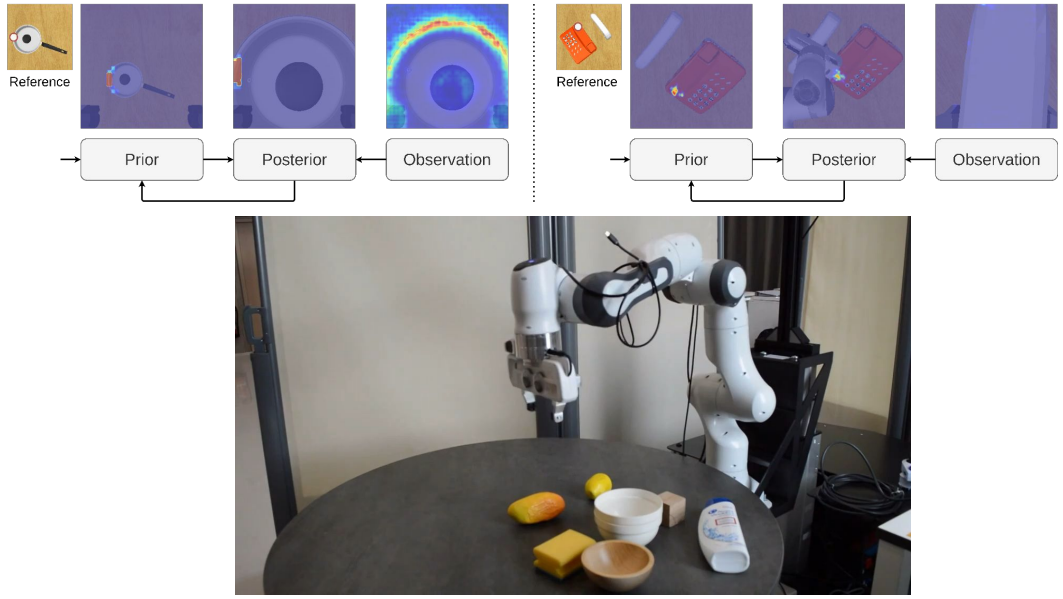
Single-Shot Reconstruction



Category and Joint Agnostic Reconstruction of ARTiculated Objects

Heppert, et al
CVPR ‘23

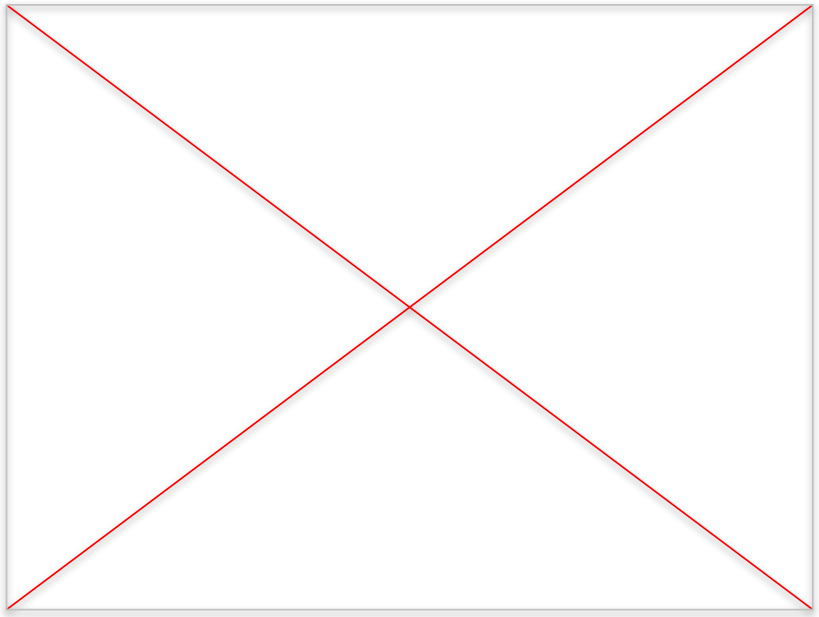
Learning scale-invariant compact representations for mobile manipulation



Bayesian Scene Keypoints for Deep Policy Learning in Robotic Manipulation

von Hartz, et al
RA-L ‘23

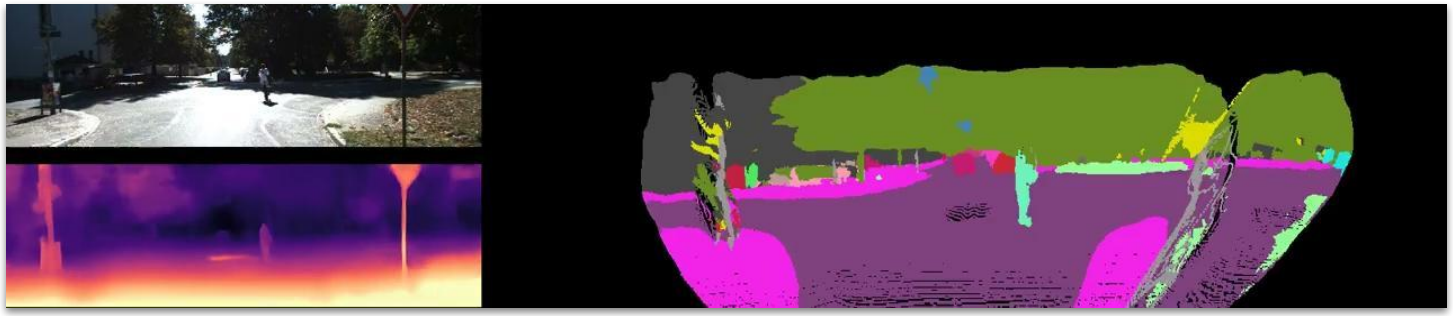
Mapping and Localization



Continual SLAM

Vödisch, Cattaneo, Burgard, Valada

ISSR ‘22



Continual Depth Estimation and Segmentation

Vödisch, Petek, Burgard, Valada

RSS ‘23

III.

Topics

Seminar Papers

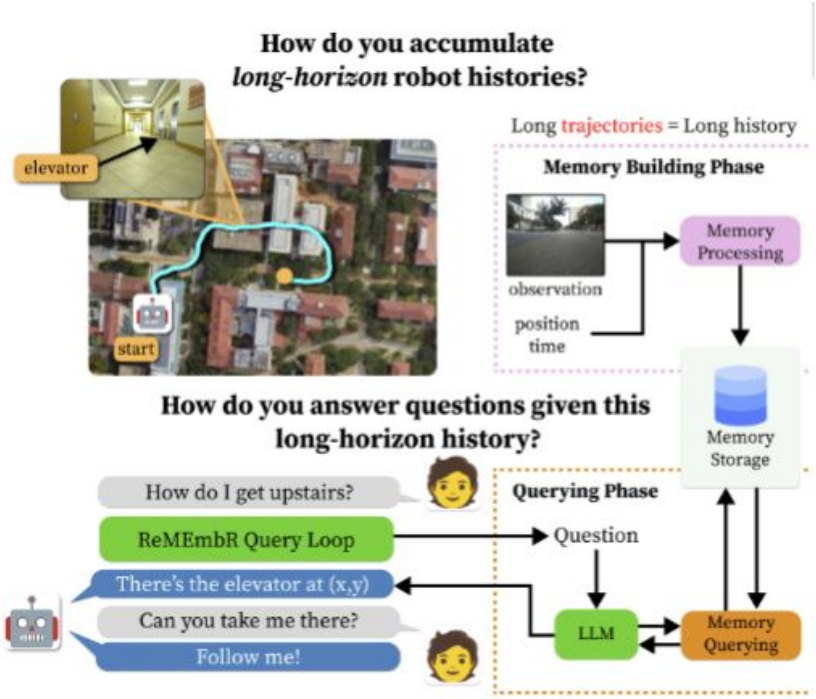


Supervisor: Yao Lu

Spatial-Temporal Memory for Robotics

https://arxiv.org/abs/2410.11989 , publisher
https://arxiv.org/html/2409.13682v1 , publisher

- **Motivation**: Robots operate in changing environments. A single observation cannot capture past object locations, events, or evolving spatial relationships.
- **Paper 1:** *DovSG (RA-L 2025):* maintains a dynamic open-vocabulary 3D scene graph and updates local object relationships as the environment changes.
- **Paper 2:** *ReMEmbR (ICRA 2025):* stores timestamped, localized video descriptions in a retrieval memory for long-horizon spatial and temporal queries.



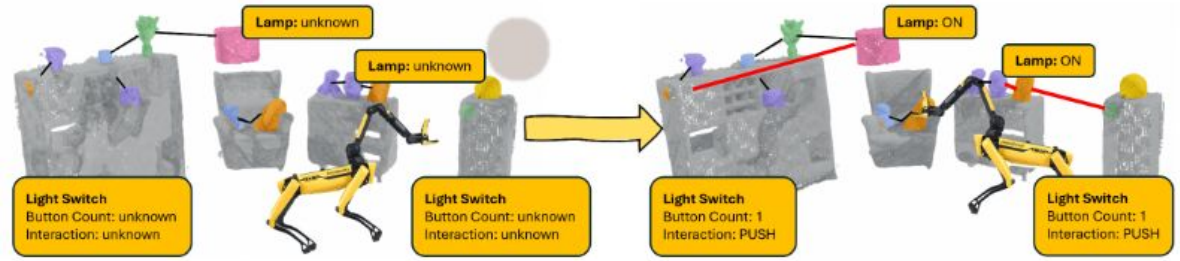
Example usecase for memory [Anwar et al. ICRA 2025]

Supervisor: Yao Lu

Interactive Perception and Learning Functional State Transitions

https://arxiv.org/abs/2409.11870 , publisher
publisher

- **Motivation**: How can a robot discover what its action change the environment?
- **Paper 1:** *SpotLight (Humanoids 2025): Focuses on discovering functional relationships through physical robot interaction.*
- **Paper 2:** *Learning Causal Relationships of Object Properties and Affordances (RAL 2022):* Focuses on learning causal relationships through demonstrations and self-supervised interventions.



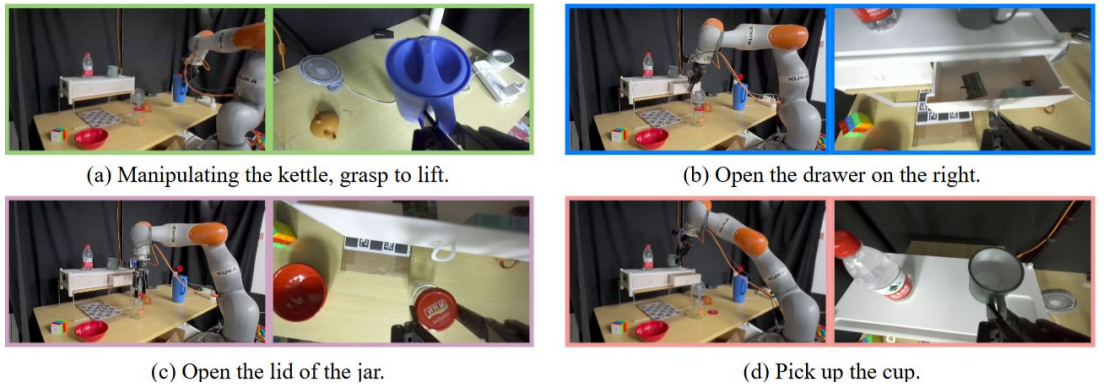
Example of relating a lightswitch to a light [Engelbracht et al. Humanoids 2025]

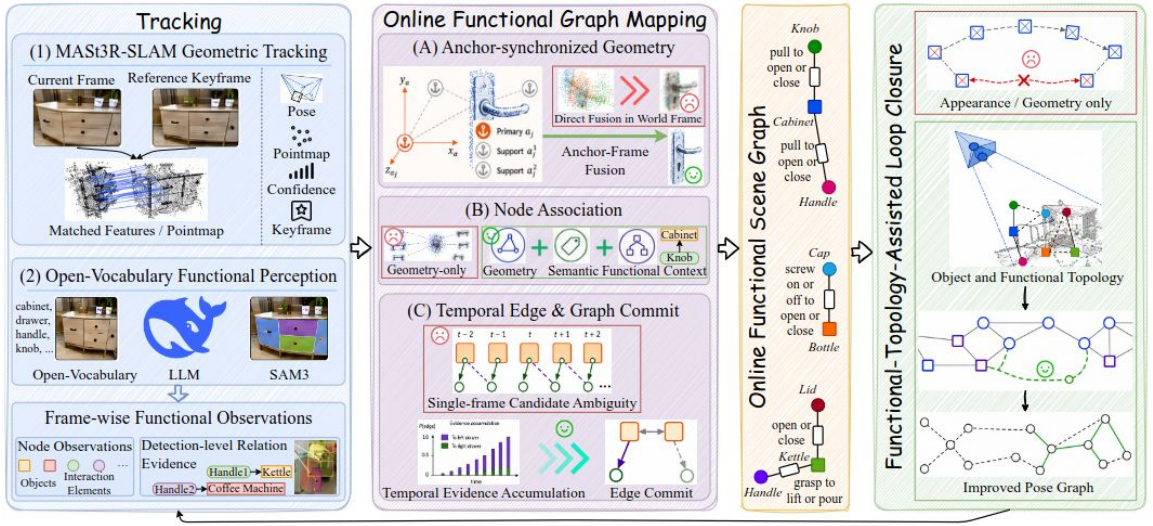
Supervisor: Seunggyu Song

Online Functional Scene Graph Mapping

https://arxiv.org/pdf/2503.19199
https://arxiv.org/pdf/2609.07497

- **Motivation**: Can we build functional scene graphs online from camera input alone? *(without given poses)*
- **Paper 1:** *OpenFunGraph(CVPR 2025):* Addresses how to encode interactive elements and functional knowledge between objects.
- **Paper 2:** *Functional-SLAM (CoRL 2026):* Builds functional scene graphs *online* by estimating camera poses simultaneously.



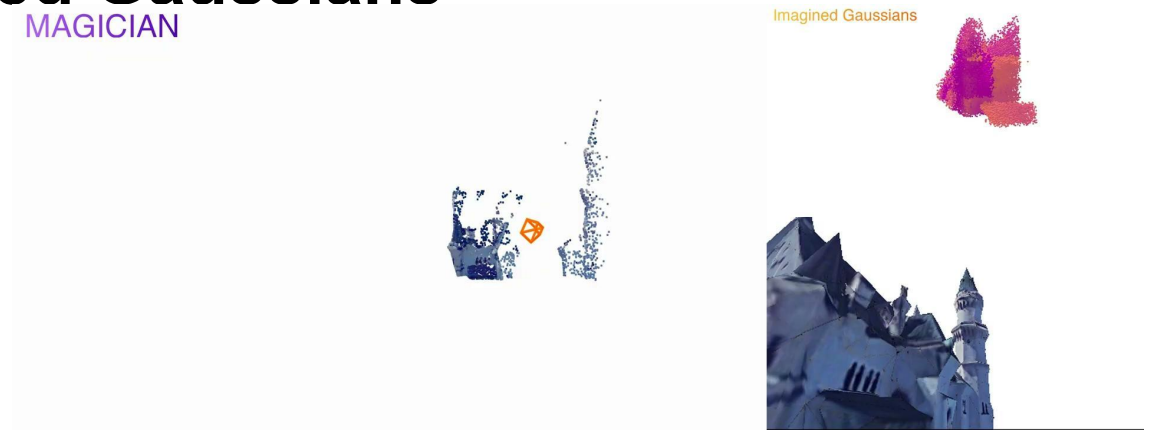


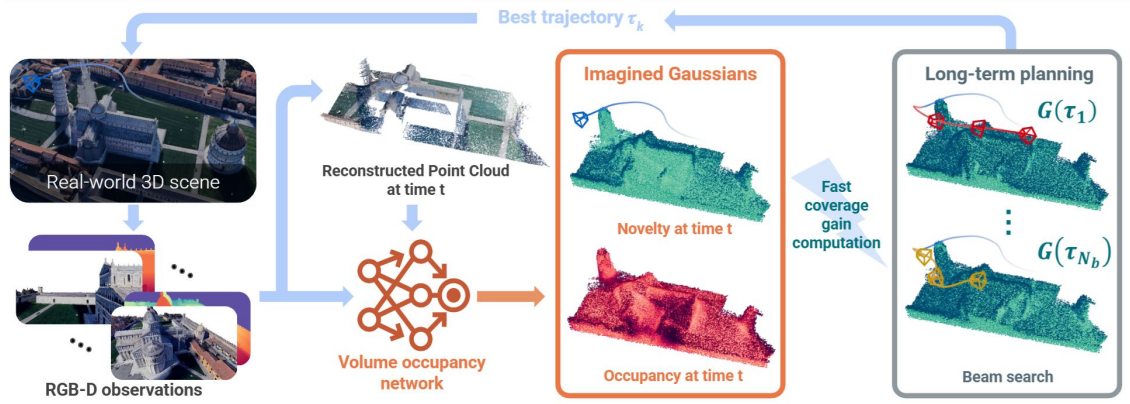
Supervisor: Seunggyu Song

Long Term Active Mapping with Imagined Gaussians

https://arxiv.org/abs/2308.04079
https://arxiv.org/pdf/2603.22650

- **Motivation**: Can a robot efficiently build a map by imagining?
- **Paper 1:** *3DGS(ACM 2023):* Represents a scene as a set of 3D Gaussians, enabling real-time rendering for novel view synthesis.
- **Paper 2:** *MAGICIAN (CVPR 2026):* Renders “Imagined Gaussians” built from predicted occupancy to rapidly estimate coverage gain, enabling long-term planning for active mapping.





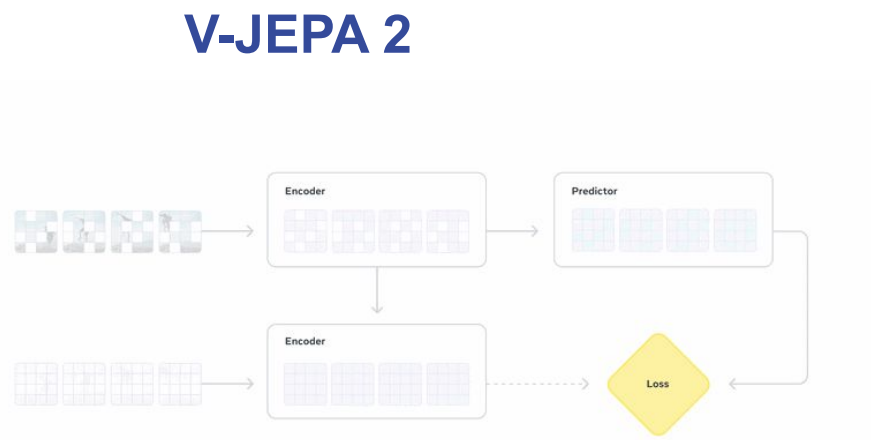
Supervisor: Akshay L Chandra

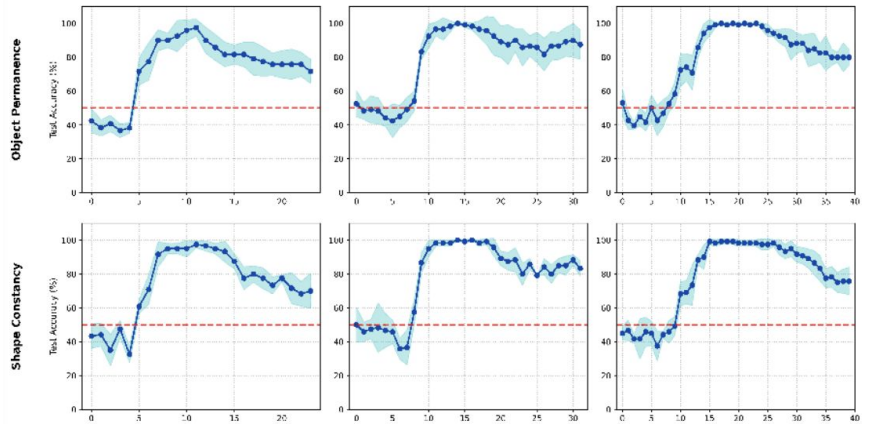
Physics Interpretability in World Models

- **Motivation:**
  - Despite rapid progress in video modeling, it remains unclear whether (and how) video world models represent physical information internally.
  - Physical reasoning in world models is not formally defined, but broadly refers to the ability to infer object properties, dynamics, and interactions, such as solidity, continuity, and causality.
- **Paper 1:** Interpreting Physics in Video World Models (Joseph et al., 2026)
- **Paper 2:** What Can Latent World Models Know? Physical Information in Multimodal Predictive Representations (Tan et al., 2026)







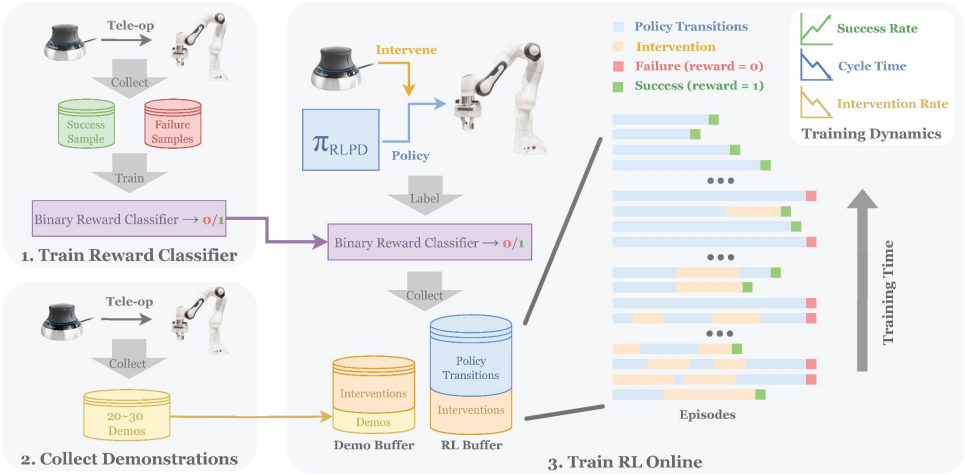


Supervisor: Akshay L Chandra

Post-Training of Robot Policies in Real-World

- **Motivation:**
  - Reinforcement learning (RL) fine-tuning of pretrained policies offers a promising path to fully leveraging pretrained priors much like how GPT-⅔ evolved into ChatGPT.
  - However, post traininig (RL) in the real-world is limited by the well-known practical deployment problems. This challenge of efficiently and reliably learning new tasks in non-trivial.
- **Paper 1:** HIL-SERL: Precise and Dexterous Robotic Manipulation via Human-in-the-Loop Reinforcement Learning (Luo et al., 2024)
- **Paper 2:** EXPO-FT: Sample-Efficient Reinforcement Learning Finetuning for Vision-Language-Action Models(Dong et al., 2026)

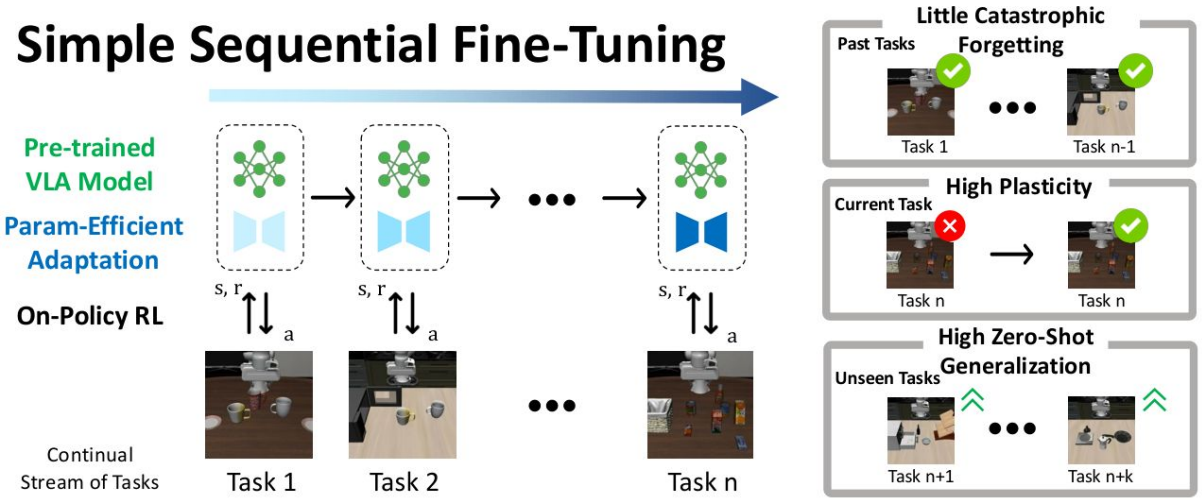
HIL-SERL



Supervisor: Bahey Tharwat

Continual Learning of Robot Policies

- **Motivation:**
  - **Adaptation to Environments:** Robots must continuously acquire new skills as tasks.
  - **Catastrophic Forgetting:** Learning new tasks should not compromise previously acquired skills.
  - **Efficient Knowledge Transfer**: Robots should reuse previously learned knowledge to acquire new skills.

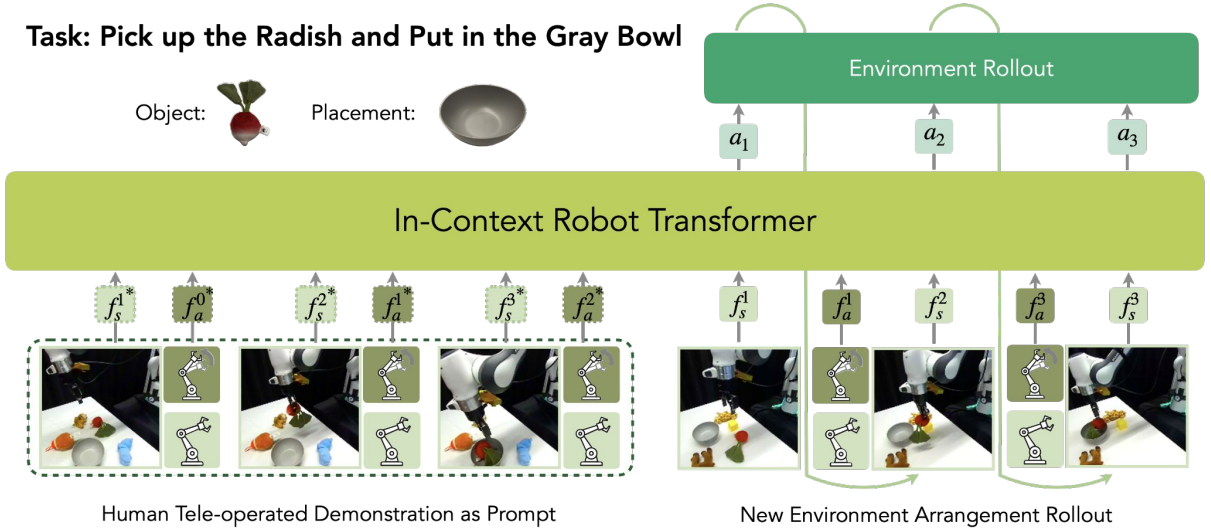


- **Paper 1:** Simple Recipe Works: VLAs are Natural Continual Learners (Jiaheng et al., 2026)
- **Paper 2:** CLARE: Continual Learning via Autonomous Adapter Routing and Expansion (Ralf et al., 2026)

Supervisor: Bahey Tharwat

In-Context Learning of Robot Policies

- **Motivation:**
  - **Rapid Task Adaptation**: Robots should adapt to new tasks through in-context examples without requiring costly model retraining.
  - **Generalization to Unseen Tasks:** Robot policies should leverage in-context information to generalize effectively to novel tasks and environments.



- **Paper 1:** In-Context Imitation Learning via Next-Token Prediction (Max et al., 2024)
- **Paper 2:** RoboTTT: Context Scaling for Robot Policies (Jiang et al., 2026)

Supervisor: Yao Lu

Text Adaption of Self-Supervised Vision Foundation Models

https://arxiv.org/abs/2412.16334 , publisher
https://arxiv.org/abs/2411.19331 , publisher

- **Motivation**: self-supervised vision foundation models learn outstanding visual features, but have no language grounding -> unusable for open-vocabulary task.
- **Paper 1:** *dino.txt* (CVPR 2025): systematic application of the LiT (Locked Image Tuning) strategy on DINOv2
- **Paper 2:** *Talk2DINO* (ICCV 2025): in contrast to dino.txt, Talk2DINO focuses on adapting embeddings from other foundation model (CLIP) to DINO vision space



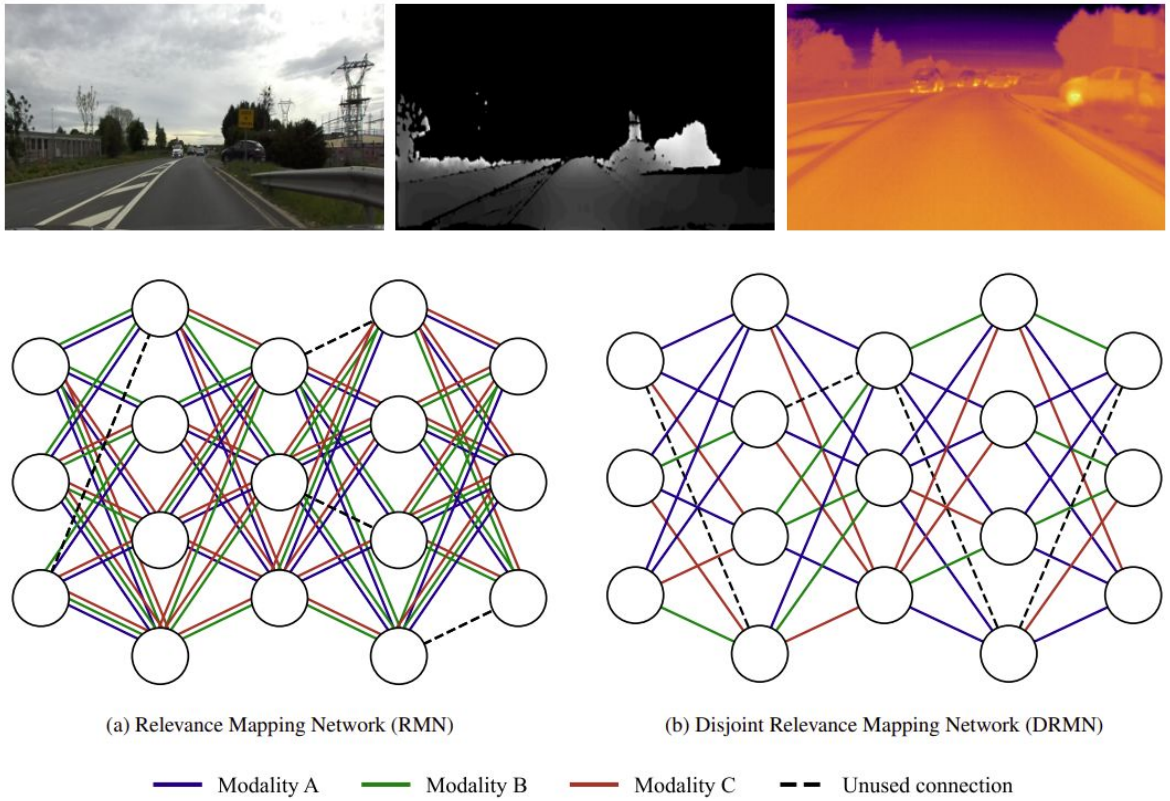
Example output of adding language grounding to DINO [Jose et al. CVPR 2025]

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

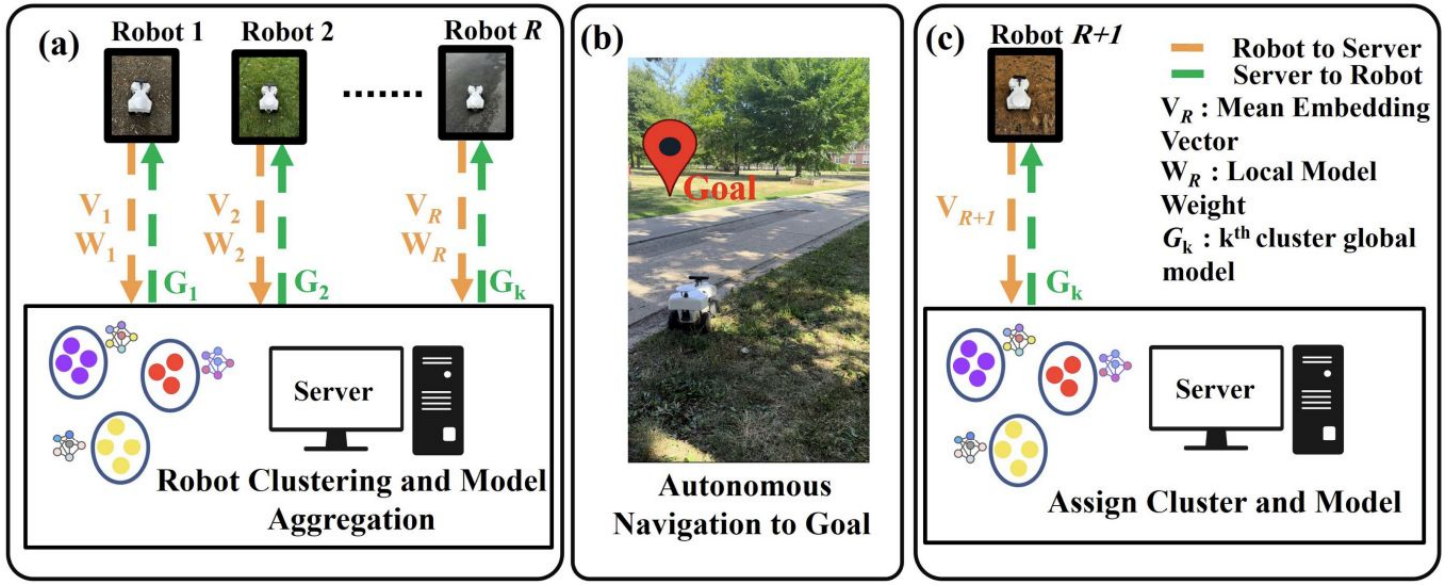


Supervisor: Julia Hindel

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.

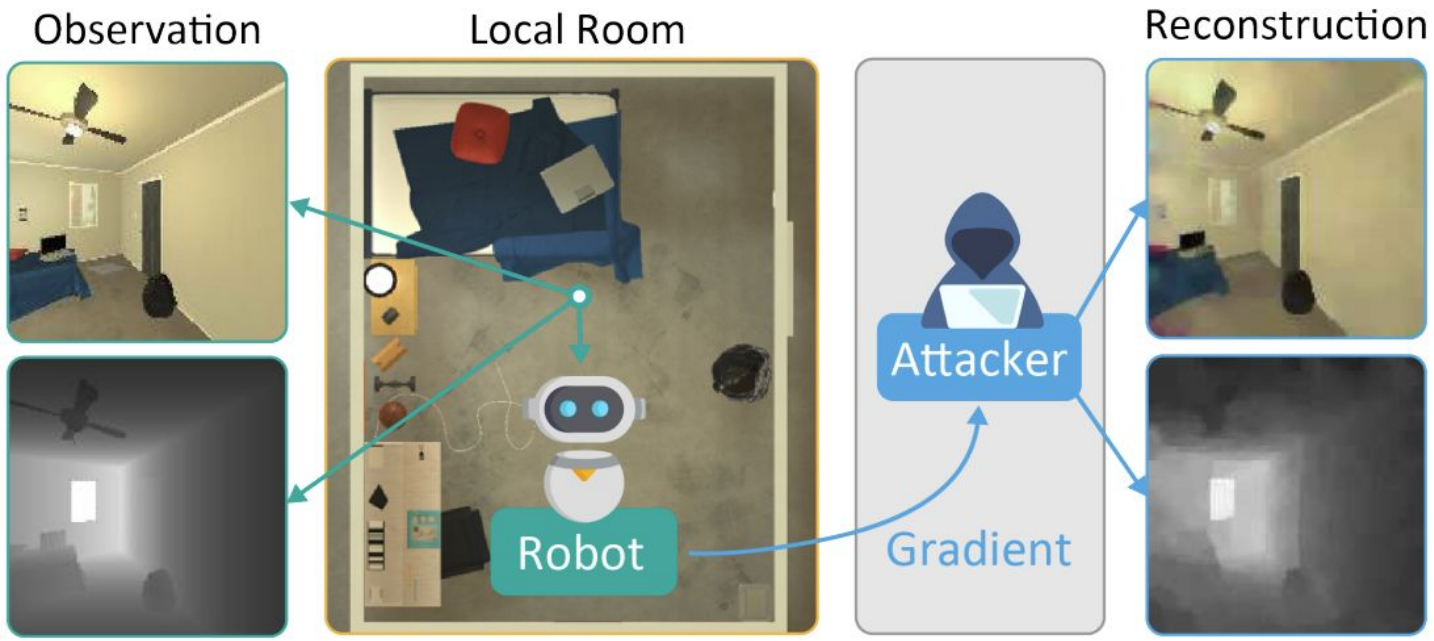


Supervisor: Julia Hindel

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

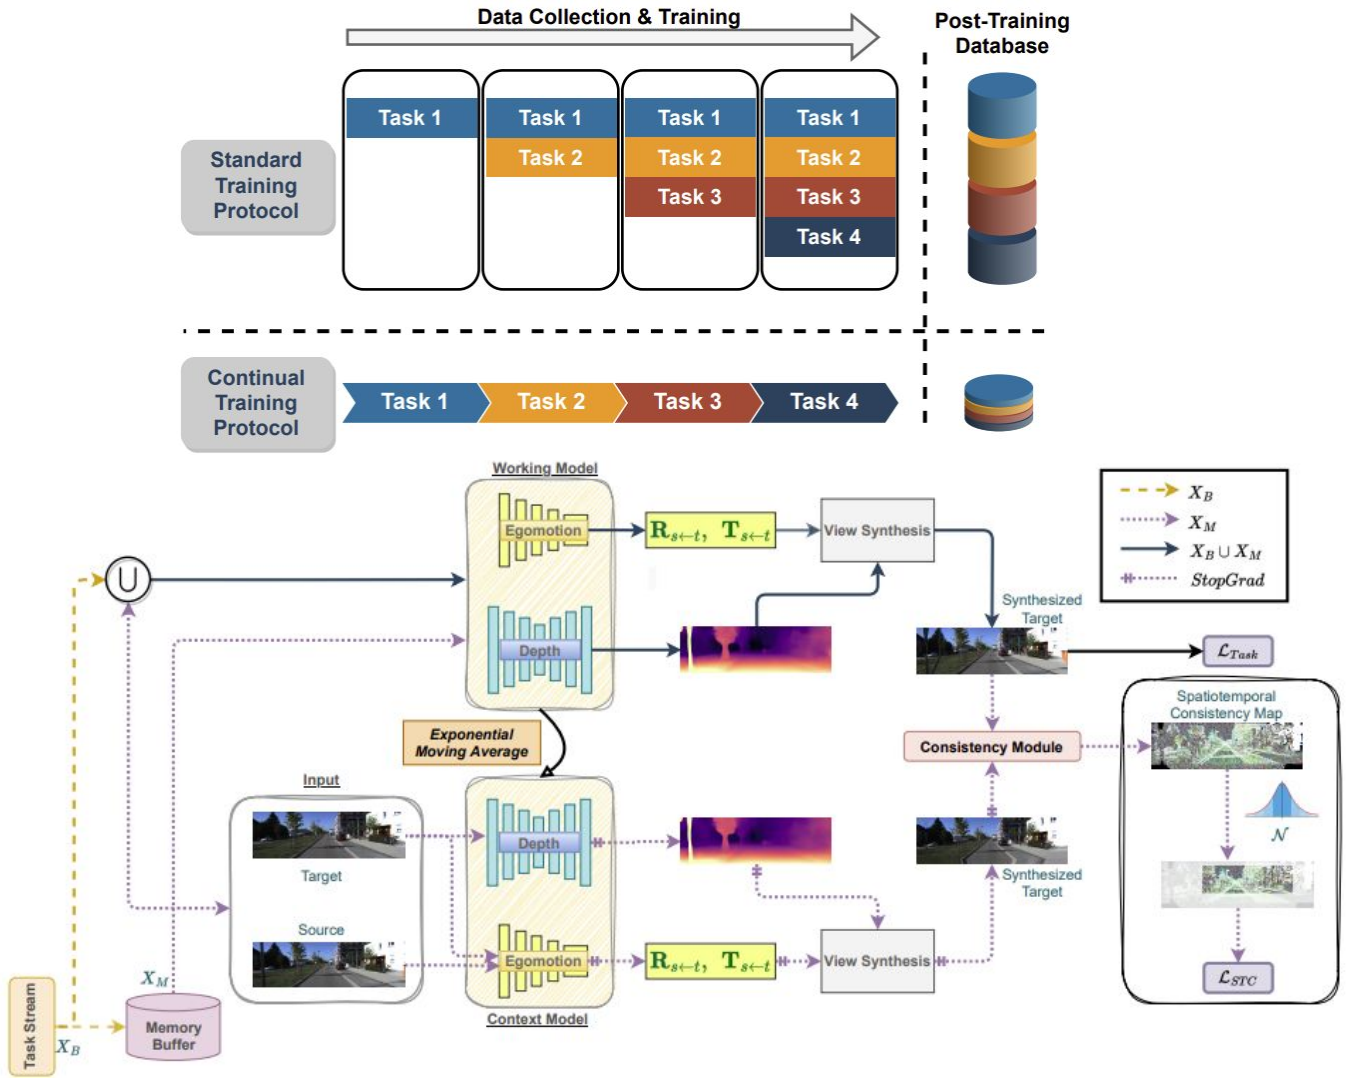


Supervisor: Julia Hindel

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

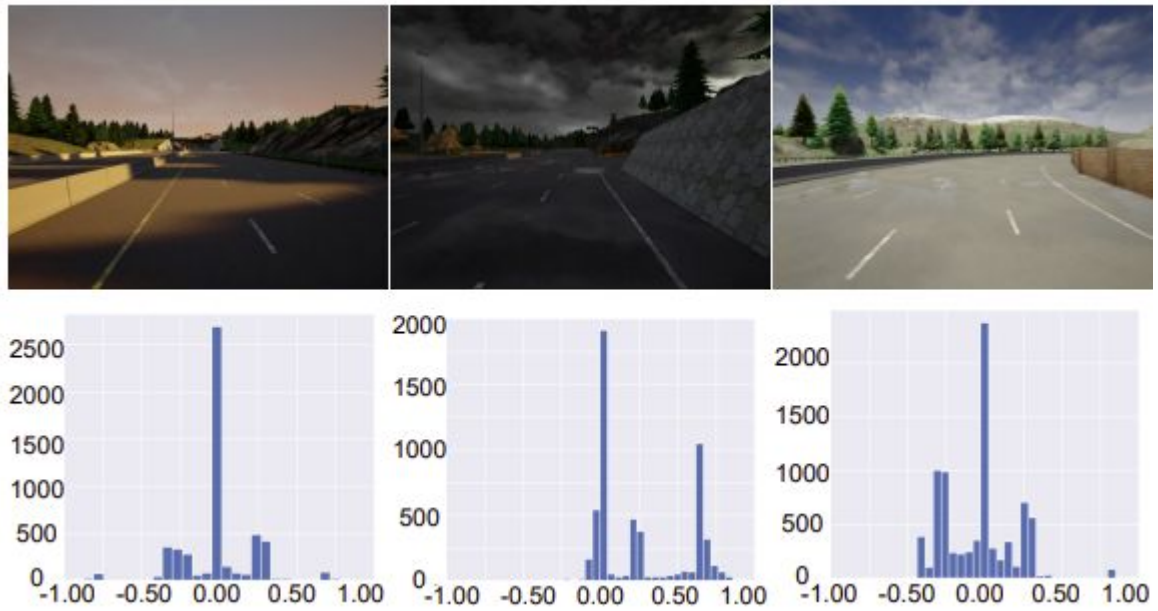


Supervisor: Julia Hindel

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.

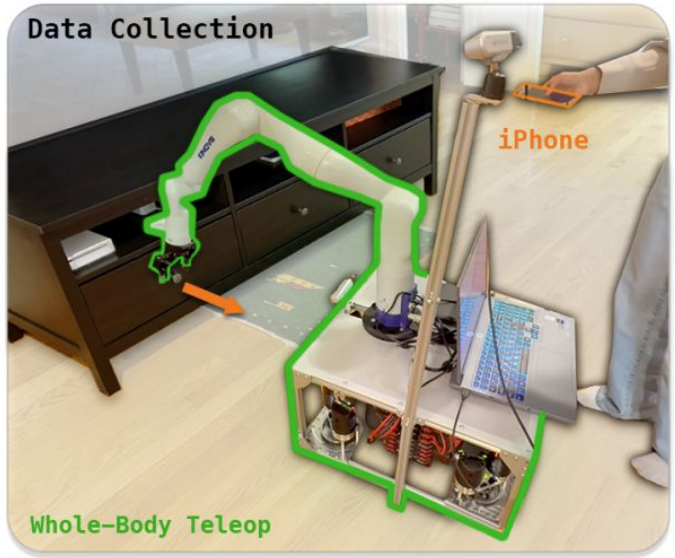


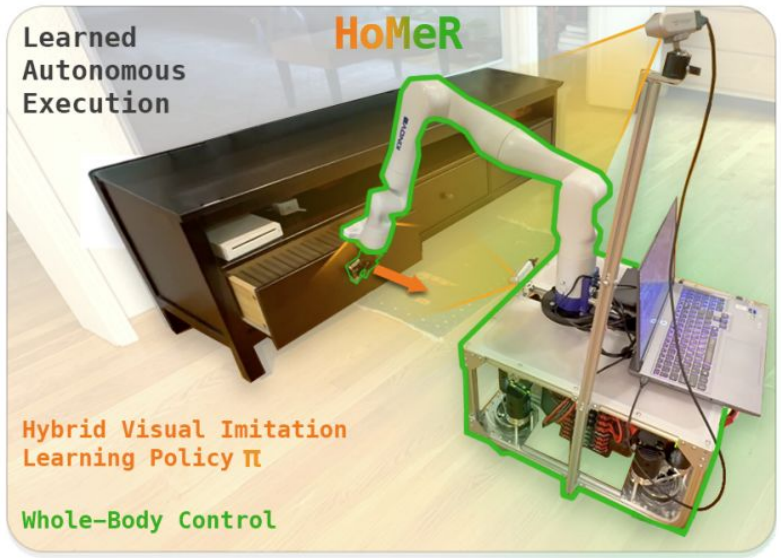
Supervisor: Martin Büchner

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



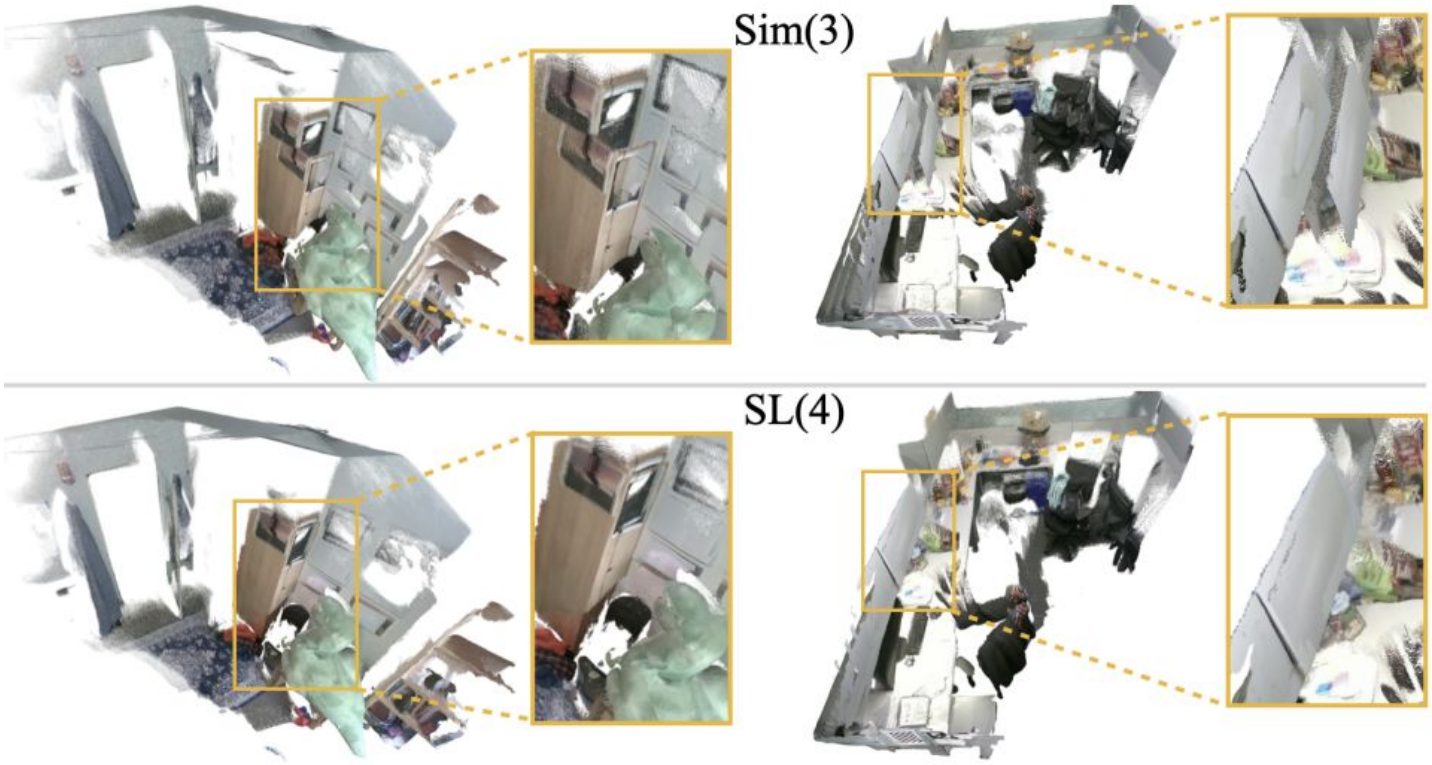


Supervisor: Martin Büchner

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

https://arxiv.org/abs/2505.12549

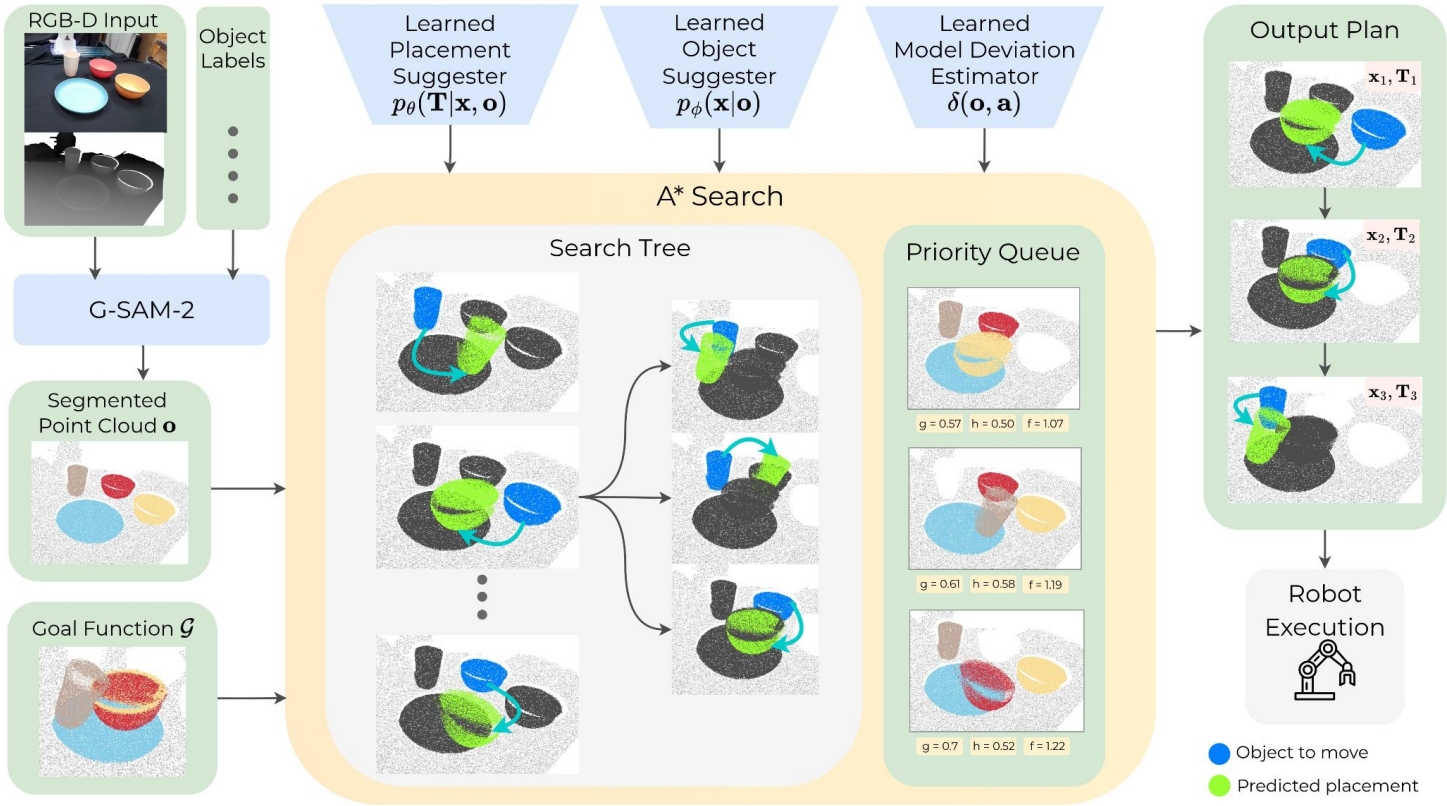
- VGGT-SLAM is a dense RGB SLAM system
  - Requires uncalibrated cameras with no assumptions on the camera motion
  - No assumptions on scene structure
  - Optimising over the SL(4) manifold
  - Better than VGGT!



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

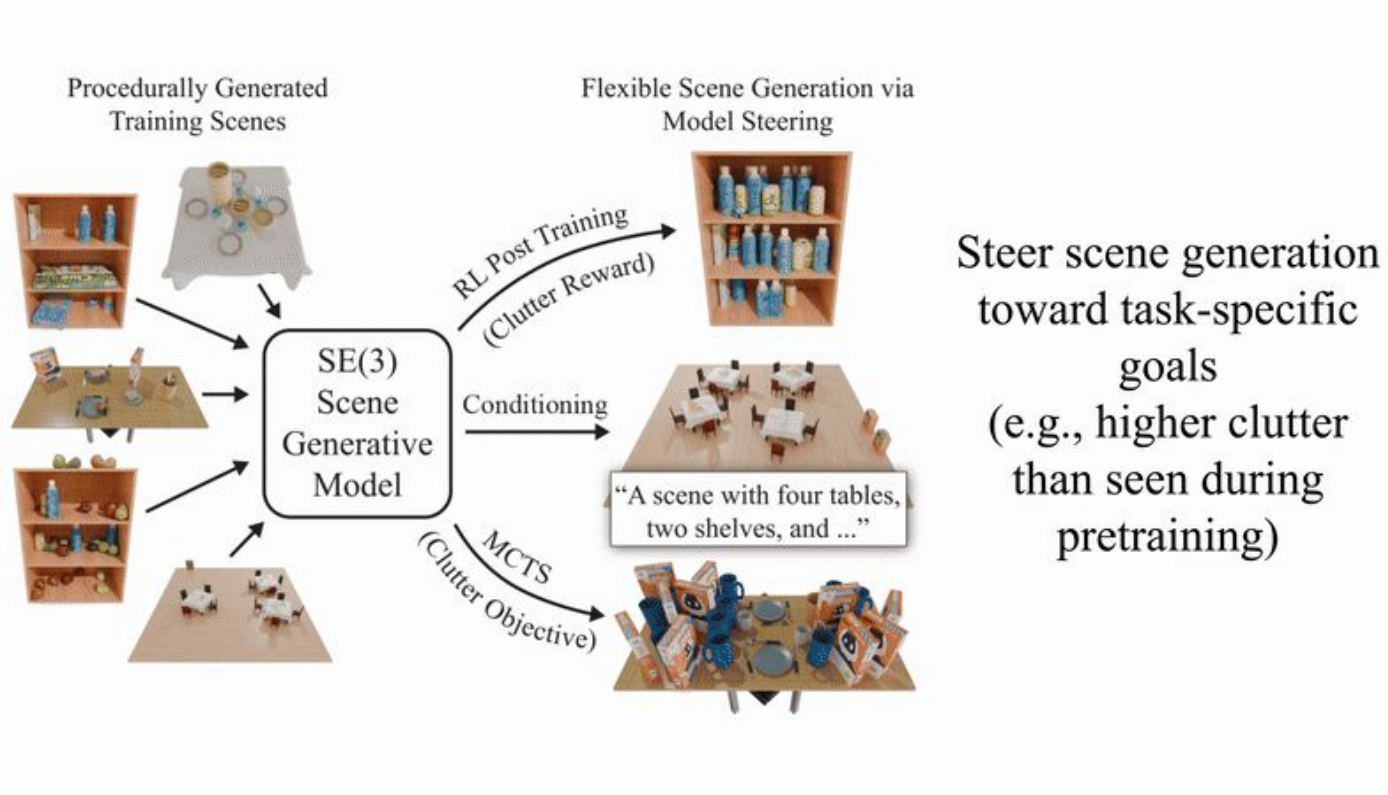


Supervisor: Martin Büchner

Steerable Scene Generation with Post Training and Inference-Time Search

https://arxiv.org/abs/2505.04831

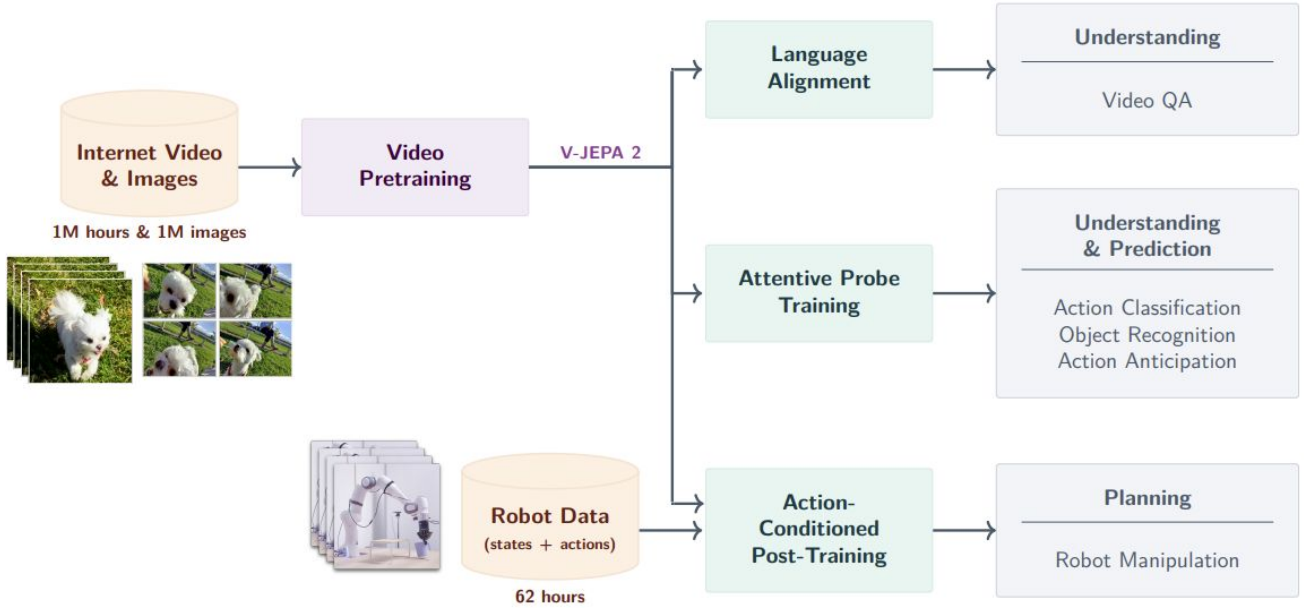
- This work generates large-scale scene data using procedural models that approximate realistic environments for robotic manipulation
  - Trains a Unified Diffusion-based Generative Model
  - Goal-directed scene synthesis that respects physical feasibility

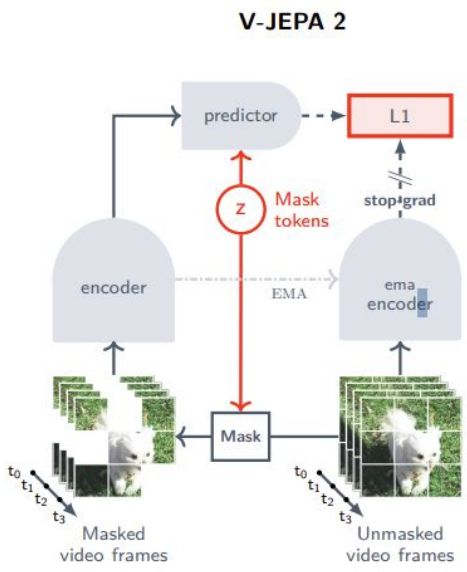


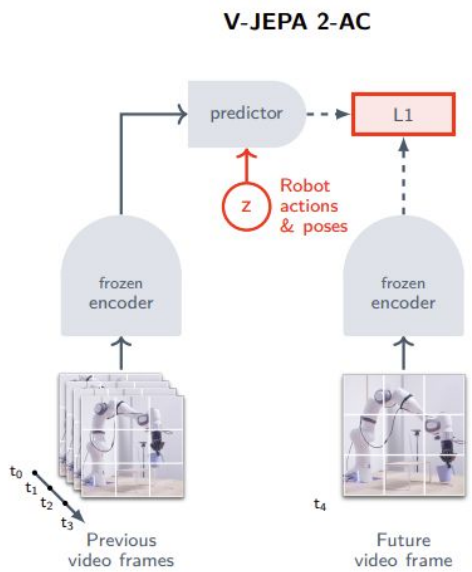
Supervisor: Iman Nematollahi

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

https://arxiv.org/abs/2506.09985





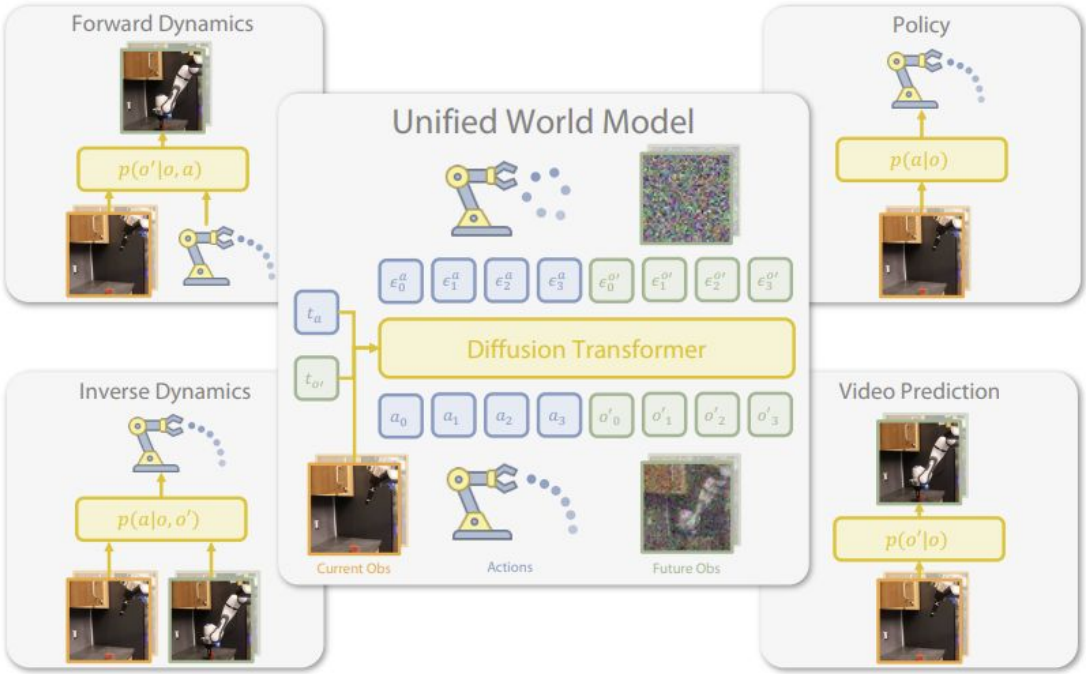


Supervisor: Iman Nematollahi

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

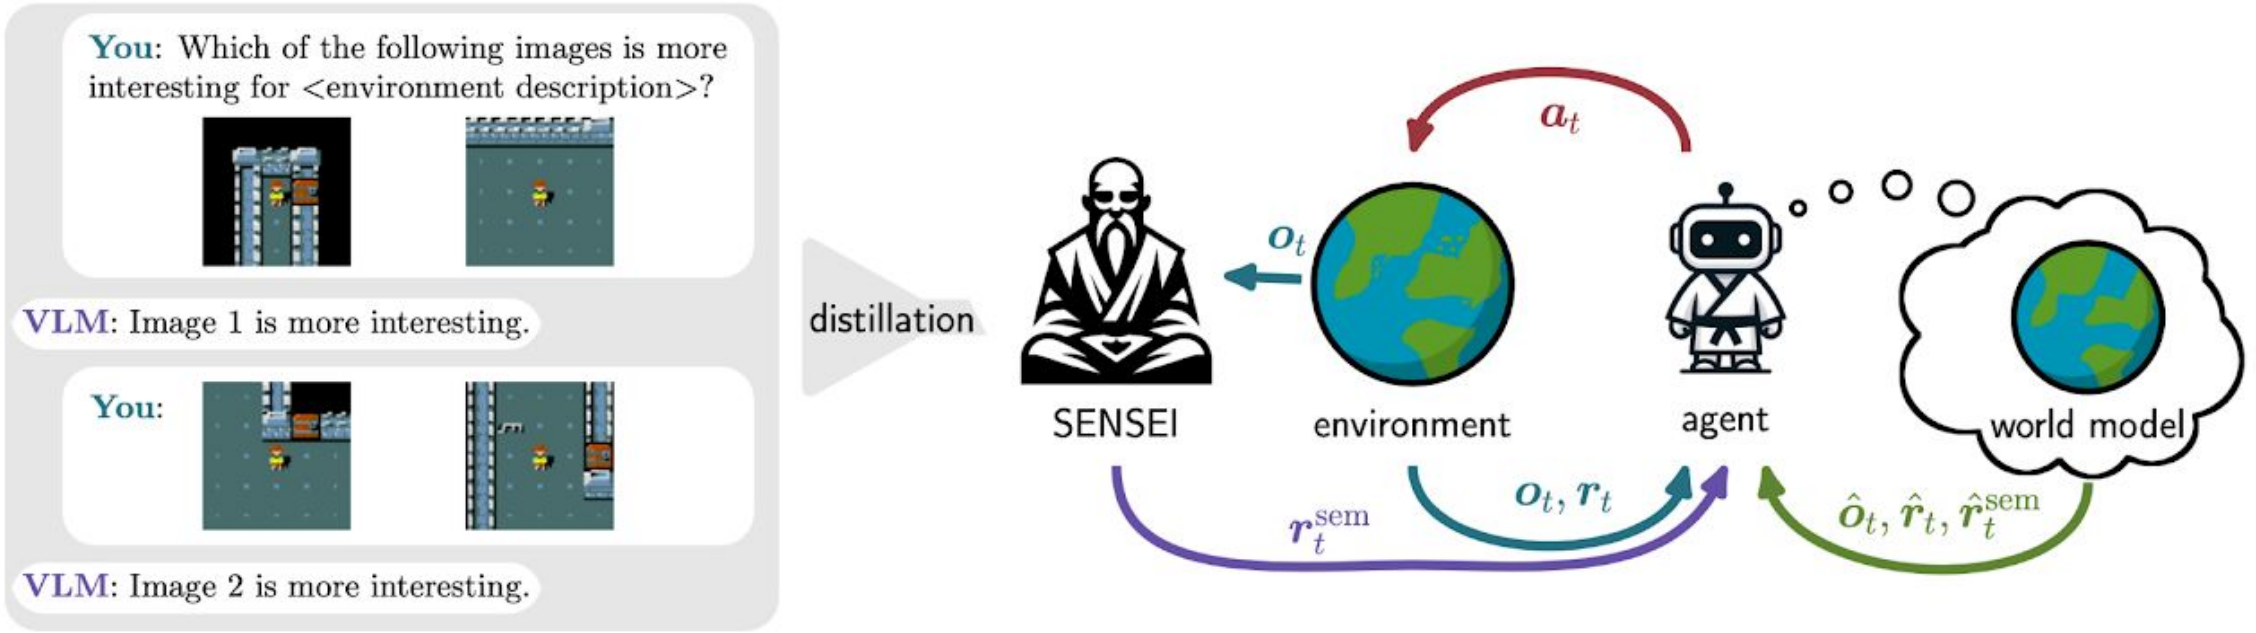


Supervisor: Iman Nematollahi

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.

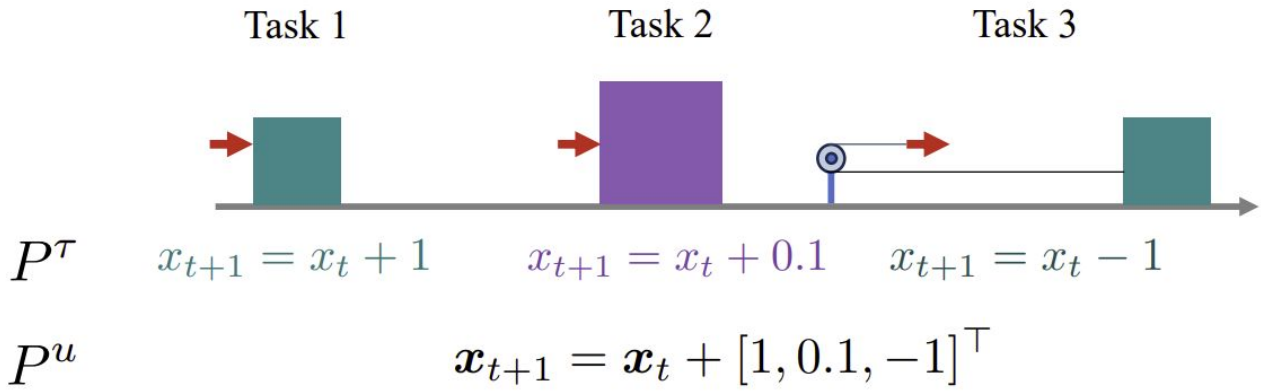


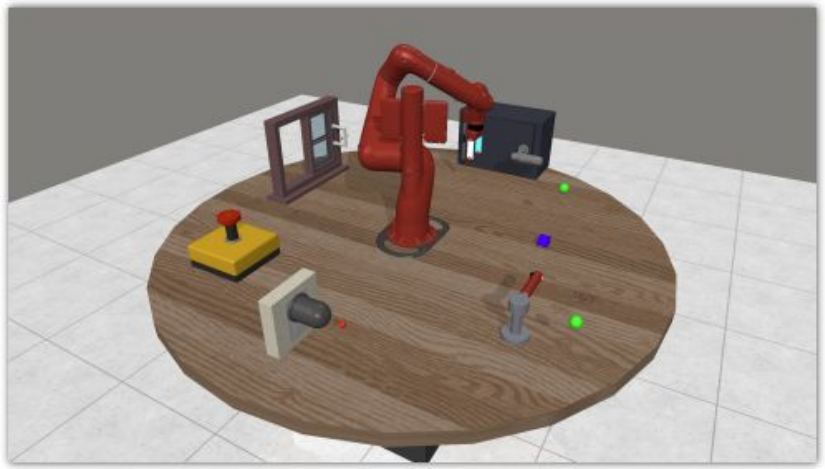
Supervisor: Iman Nematollahi

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



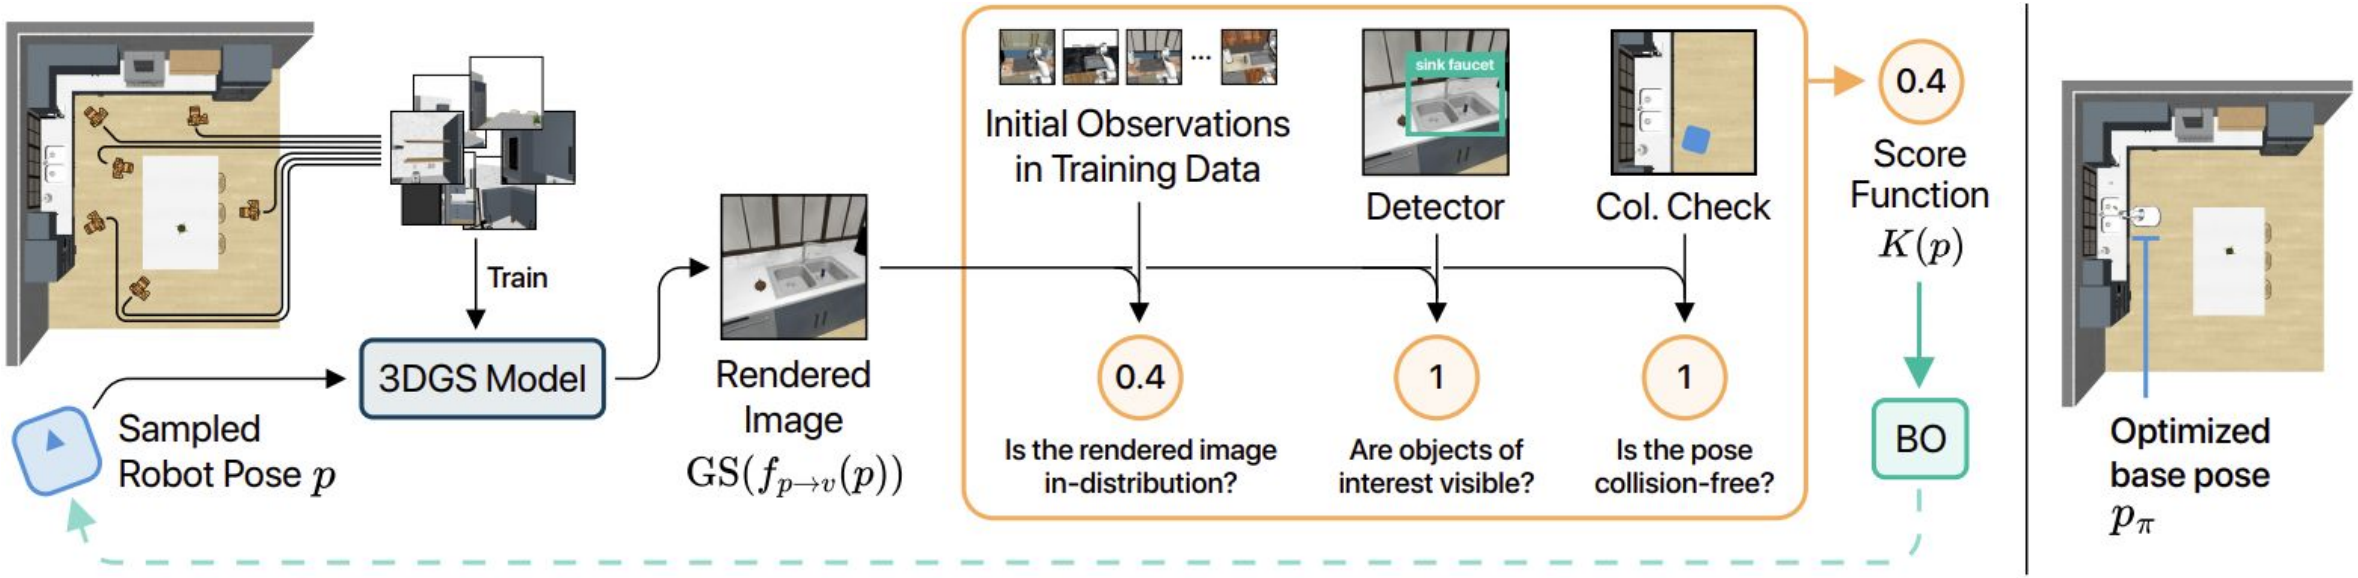


Supervisor: Iman Nematollahi

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



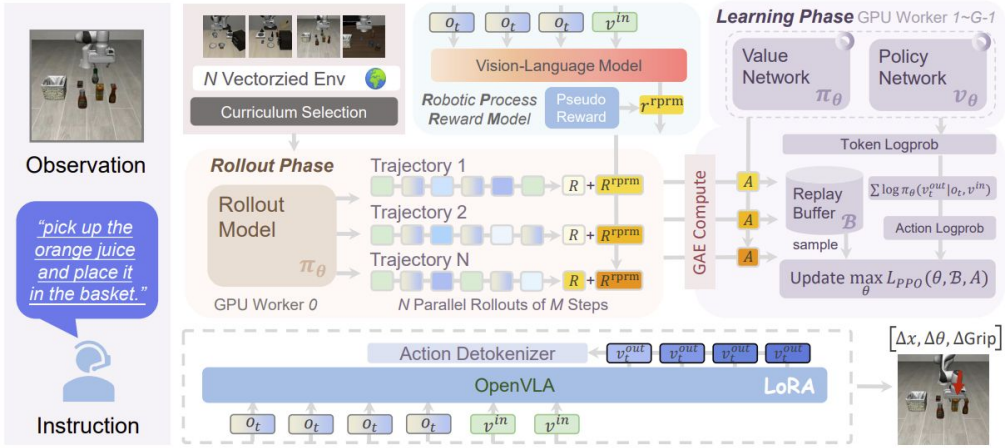
Supervisor: Akshay L Chandra

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



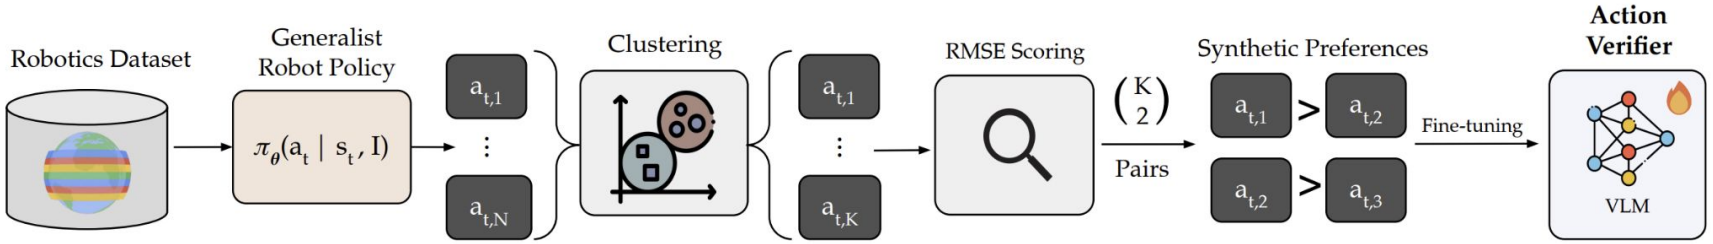


Supervisor: Akshay L Chandra

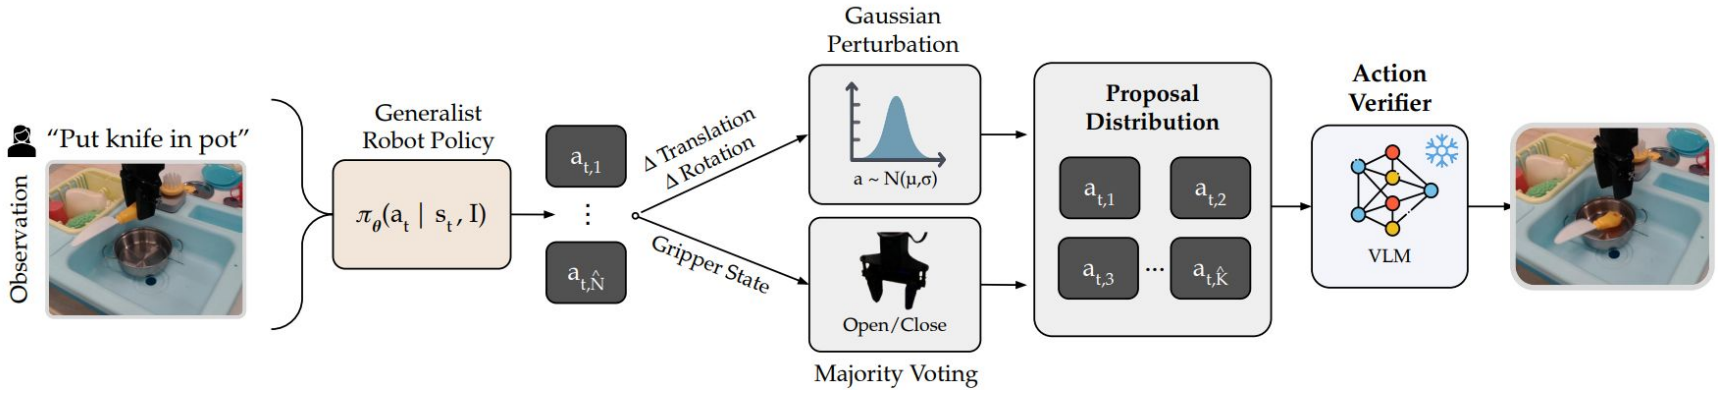
RoboMonkey: Scaling Test-Time Sampling and Verification for Vision-Language-Action Models

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!

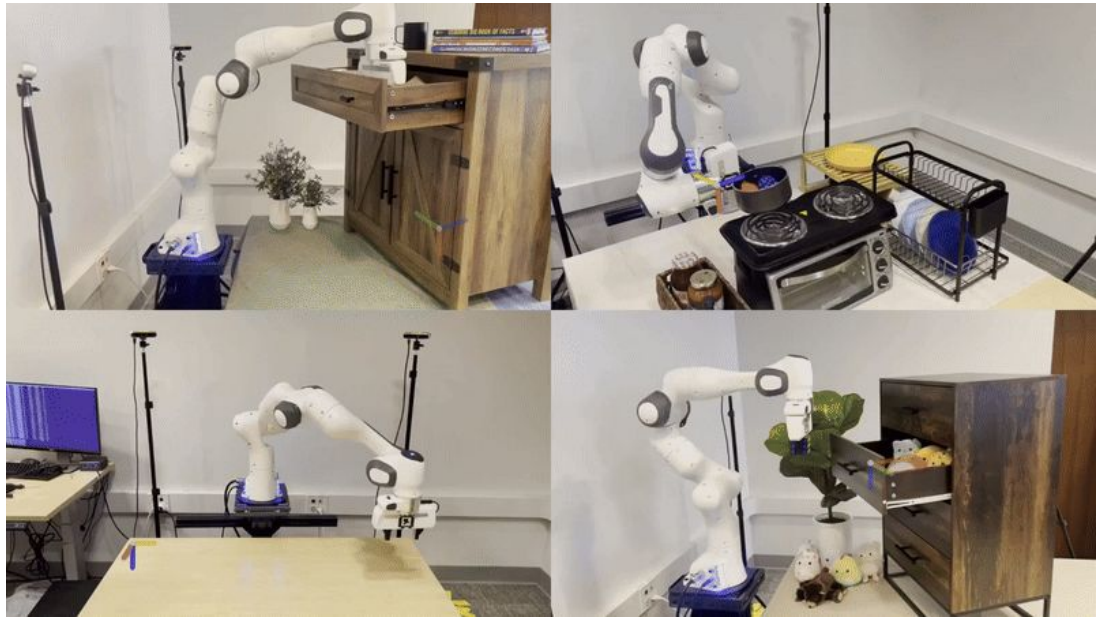


Supervisor: Akshay L Chandra

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

https://arxiv.org/abs/2509.06953

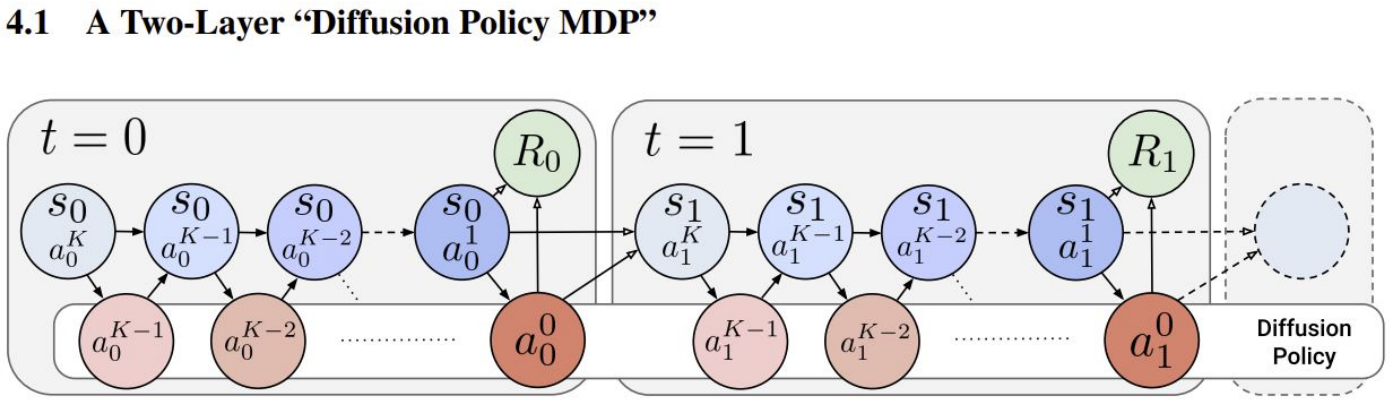
- Motion planning is great!
  - We have really fast IK solvers now!
  - But they cannot generalise!
- This work: Neural motion policies that can generalise in complex or dynamic settings
  - A transformer trained on **10 million** generated expert trajectories across diverse simulation scenarios!



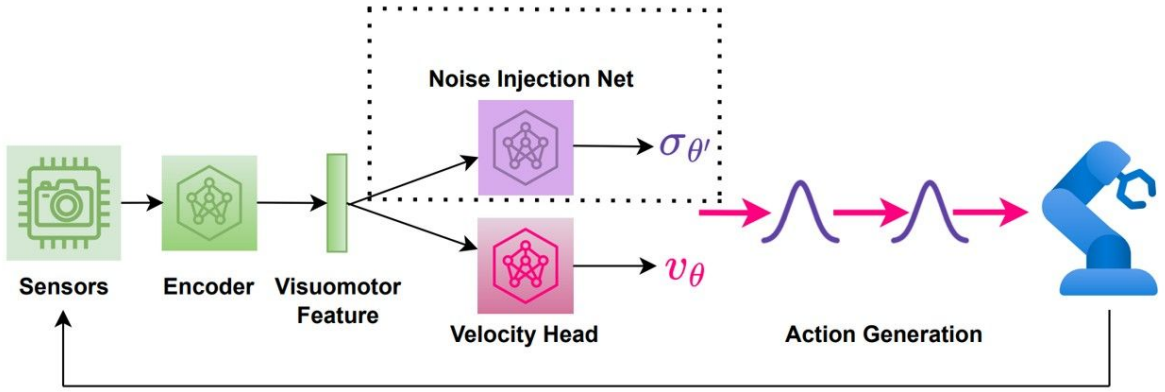
Supervisor: Akshay L Chandra

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

https://arxiv.org/abs/2505.22094



Source: Ren et al., DPPO, 2025



Questions?



Announcement

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

Yao Lu

Robot Learning Lab

lu@cs.uni-freiburg.de