Haotian Lin

Tel.: +1 (412)909-9214 – Email: vlin3@andrew.cmu.edu
Personal Website: https://darthutopian.github.io

EDUCATION

Carnegie Mellon University — Pittsburgh
M.S. in Robotics — Aug. 2024– now

- GPA: 4.0/4.0
- Related Courses: Anvanced Mechine Learning(A), Optimal Control & RL(A), Introduction to Robot Learning(A)

Tsinghua University — Beijing
B.S. in Mechanics & B.E. in Vehicle Engineering — Sep. 2020 – Jun. 2024

- GPA: 3.6/4.0

University of California, Berkeley — Beijing
Visiting Research Intern (Undergraduate) — Jun. 2023 – Sep. 2024

RESEARCH INTERESTS

I am broadly interested in the intersection of reinforcement learning, robotics and foundation models. In specific, I am interested in scalable RL algorithms and improving robotic foundation models with on-the-fly experience. I broadly explore novel algorithms, architectures, and pre-training/post-training paradigms that aligns with my goal. Eventually, I aim to build scalable, self-improving, generalist robotic agents.

RESEARCH EXPERIENCE

LeCAR Lab & Auton Lab (Carnegie Mellon University) — Pittsburgh, PA
Research Assistant (MSR)
Advisor: Guanya Shi (*guanyas@andrew.cmu.edu*) , Jeff Schneider (*jeff4@andrew.cmu.edu*)

Multi-task RL for Robotics in Parallel Simulation — 6/2025–Now

- Study generalizable skill acquisition and task representation learning in multi-task learning, especially for on-policy RL algorithms with massive parallelization.

Self-improving Foundation Models via Residual RL — 5/2025 –10/2025

- Building data-efficient RL pipeline for manipulation in JAX for real-world learning (Yam Arm, Franka) and simulation benchmarks (LIBERO, SimplerEnv, Mimicgen), supports data collection through different teleoperation approaches and self-play, integrates pipeline for VLA SFT and evaluation.
- Leverage expressive policy prior (pi-0, octo, OpenVLA-oft) to warm-start exploration and probing using a residual RL. Employ a hybrid rollout scheme to collect diverse recovery behavior and distilled these curated behavior back into the generalist via SFT.

Improving Model-Based RL for High-Dimensional Continuous Control — 10/2024 –3/2024

- Study structural policy mismatch and value overestimation issue in plan-based MBRL at high-dimensional action space. Provide bounds for value approximation error.
- Propose a minimalist fix to mitigate the issue for TD-MPC family algorithms and boost performance for across a wild variety of continuous control tasks, e.g., humanoid locomotion.

Mechanical System Control Lab (University of California, Berkeley) — Berkeley, CA
Research Assistant
Advisor: Masayoshi Tomizuka (*tomizuka@berkeley.edu*)

Diffusion Model for Multi-Agent Pedestrian Trajectory Prediction — 10/2023 – 4/2024

- Leverage latent diffusion to model complex social interaction between agents. Propose a novel trajectory prediction scheme that consider trajectory prediction as inpainting problem and solved through posterior sampling provide
- Contribute to conference papers admitted to IROS 2024.

Prediction Benchmark — 7/2023 –9/2023

- Evaluate prediction model's ability of interaction modeling and map representation, bridging the gap of concept interpretation and model performance by Shapley Value-based contribution metrics.

THICV (Tsinghua University) Beijing

Research Assistant

Advisor: Jianqiang Wang (*wjqlws@tsinghua.edu.cn*) , Shaobing Xu (*shaobxu@tsinghua.edu.cn*)

GNN-based prediction framework for generalizable scene understanding 10/2022 – 6/2023

- Designing social interaction aware backbone leveraging prior knowledge for autonomous driving, adaptable to downstream tasks including trajectory prediction and intention detection. Designing conditional diffusion model-based traffic simulator that integrates natural language commands.
- Contribute to conference papers admitted to IEEE IV 2023 and IROS 2025.

iDLab (Tsinghua University) Beijing

Research Assistant

(Students Research Training Program)

Advisor: Shengbo Li (*lishbo@tsinghua.edu.cn*)

DRL Algorithm Development for Integrated Decision-Making & Control 10/2022 – 6/2023

- Testing and verification of **GOPS** toolkit in driving scenario.
- Develop offline-to-online RL algorithms for diffusion and flow policies using Q-score alignment.

SELECTED PUBLICATIONS

- Wenli Xiao, **Haotian Lin,** Haoru Xue, Tairan He, Zhengyi Luo, Yuqi Xie, Fengyuan Hu, Andy Peng, Guanya Shi, Jim Fan, Yuke Zhu. "Self-Improving Vision-Language-Action Models with Data Generation via Residual RL" (In Submission)
- **Haotian Lin**, Pengcheng Wang, Jeff Schneider, Guanya Shi. "TD-M (PC) $^ 2$: Improving Temporal Difference MPC Through Policy Constraint". arXiv preprint arXiv:2502.03550. (In Submission)
- Wang Yixiao, Mingxiao Huo, Zhixuan Liang, Yushi Du, Lingfeng Sun, **Haotian Lin**, Jinghuan Shang et al. "Ver: Vision expert transformer for robot learning via foundation distillation and dynamic routing." arXiv preprint arXiv:2510.05213. (In Submission)
- **Haotian Lin**, Yixiao Wang, Mingxiao Huo, Chensheng Peng, Zhiyuan Liu, Masayoshi Tomizuka. "Joint Pedestrian Trajectory Prediction through Posterior Sampling". 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems.
- Zhiyuan Liu, Leheng Li, Yuning Wang, **Haotian Lin**, Hao Cheng, Zhizhe Liu, Lei He, Jianqiang Wang. "Controllable traffic simulation through llm-guided hierarchical chain-of-thought reasoning". 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems.

SKILLS & FOCUS

- **Reinforcement learning**, **Robot Foundation Models.**
- Coding skills: Python (PyTorch, JAX), C++, Matlab.
- Robotic Physical Simulator: IsaacSim, IsaacGym, Mujoco, MJX
- Software: SolidWorks

EXTRACURRICULAR EXPERIENCES

Secondary Award in Xindong Project (Intelligent vehicle design contest, Tsinghua University) 1/2021

Member of Tsinghua Photography Team (official art group in Tsinghua University) 9/2021 – 6/2024

Committee Member of Tsinghua Photography Team 9/2022 – 6/2023

Member of Student Union (Department Of Art) of Xingjian College 9/2021 – 6/2022