Shaolong Li

shaolonl@umich.edu shaolonl-16.github.io

EDUCATION

University of Michigan Aug. 2025 – May 2027 (expected)
M.S. in Computer Science and Engineering *Ann Arbor, MI*

Central South University Sep. 2020 – Jun. 2024
B.S. in Computer Science and Technology *Changsha, China*
Overall GPA: 93.76/100 (Rank: 2/235) | **Major GPA:** 93.95/100 (Rank: 1/235)

RESEARCH INTERESTS

My goal is to build robot learning systems that combine deliberate reasoning with efficient, real-time control. I focus on interfaces between high-level agents and low-level policies, asking how to represent task intent and translate it into guidance that policies can reliably follow, and when additional reasoning is worth its latency. I also study visual representations for memory, future prediction, and in-context learning from demonstrations. Beyond vision, I ask how these representations should be combined with language and touch.

PUBLICATIONS

* Equal contribution.

LeanWAM: Learning Latent World Action Models from Frozen Video Models [PDF]
Shaolong Li*, Hongze Fu*, Joyce Chai, and Yinpei Dai
arXiv preprint.

MotionMEM: Self-Supervised Motion Tokens as Memory for Vision-Language-Action Policies [PDF]
Hongze Fu*, **Shaolong Li***, Byeongju Woo, Yinpei Dai, Yunhao Luo, Joyce Chai, and Nima Fazeli
arXiv preprint.

One-Step Flow Policy: Self-Distillation for Fast Visuomotor Policies [PDF]
Shaolong Li, Lichao Sun, and Yongchao Chen
In *European Conference on Computer Vision* (**ECCV**), 2026.

Mixed Sparsity Training: Achieving 4× FLOP Reduction for Transformer Pretraining [PDF]
Pihe Hu*, **Shaolong Li***, Xun Wang, and Longbo Huang
Transactions on Machine Learning Research (**TMLR**), 2025.

Value-Based Deep Multi-Agent Reinforcement Learning with Dynamic Sparse Training [PDF]
Pihe Hu*, **Shaolong Li***, Zhuoran Li, Ling Pan, and Longbo Huang
In *Annual Conference on Neural Information Processing Systems* (**NeurIPS**), 2024.

RESEARCH EXPERIENCE

University of Michigan, SLED Lab Mar. 2026 – Present
Research Assistant, advised by Prof. Joyce Y. Chai *Ann Arbor, MI*

LeanWAM: Learning Latent World Action Models from Frozen Video Models

- Proposed a two-stage framework that compresses future frames into 32 tokens through a frozen video generator and trains a lightweight policy to predict them, eliminating the video generator from deployment.
- Found that the 32 tokens predict actions better than 3,000+ dense future features.
- Outperformed strong VLA and WAM baselines on LIBERO, LIBERO-Plus, and RoboTwin 2.0, while running 12–42× faster than representative WAMs.

MotionMEM: Self-Supervised Motion Tokens as Memory for Vision-Language-Action Policies

- Proposed a self-supervised motion memory for VLA policies that encodes each 33-frame window as one token.
- Raised average success on RoboMME from 30.7% to 37.3% by adding $\leq$96 motion tokens to 512 frame tokens.
- Showed that the compression must be learned: pooled frozen video latents at the same budget trailed by 8–19 points on tasks that depend on event order.

Tsinghua University & Lehigh University Jun. 2025 – Feb. 2026
Research Assistant, advised by Prof. Yongchao Chen and Prof. Lichao Sun *Remote*

One-Step Flow Policy: Self-Distillation for Fast Visuomotor Policies *ECCV 2026*

- Proposed a self-distillation framework for one-step action generation, avoiding iterative sampling and the need for a separately pretrained teacher.
- Achieved 71.6% average success on 56 simulated manipulation tasks in one step, above 100-step diffusion (66.4%) and flow (59.8%) policies, while running over 100× faster.
- Accelerated $\pi_{0.5}$ from 10-step to one-step action generation while improving success on RoboTwin 2.0 tasks (94.7% vs. 91.5% on average).

University of North Carolina at Chapel Hill Sep. 2024 – Dec. 2024
Research Assistant, advised by Prof. Huaxiu Yao *Remote*

Step-Level Preference Optimization for LLM Reasoning

- Adapted Direct Preference Optimization from whole-response to step-level preferences for math reasoning.

Tsinghua University, Institute for Interdisciplinary Information Sciences Jul. 2023 – Jun. 2024
Research Assistant, advised by Prof. Longbo Huang *Beijing, China*

Mixed Sparsity Training: Achieving 4× FLOP Reduction for Transformer Pretraining *TMLR 2025*

- Designed a three-phase sparsity schedule that trains at up to 96% sparsity before restoring dense connectivity.
- Cut GPT-2 pretraining FLOPs by 4× while matching dense zero-shot accuracy (60.5% vs. 60.7% on LAMBADA).

Value-Based Deep Multi-Agent Reinforcement Learning with Dynamic Sparse Training *NeurIPS 2024*

- Proposed a dynamic sparse training framework for value-based multi-agent RL, using multi-step targets to mitigate overestimation and a dual replay buffer to stabilize learning.
- Reduced training and inference FLOPs by up to 20× on SMAC while retaining near-dense performance.

HONORS AND AWARDS

- Outstanding Graduate Mar. 2024
- Ruiwei Scholarship Oct. 2023
- Outstanding Student Dec. 2022
- Yihao Foodstuff Scholarship Nov. 2022
- National Scholarship (Top 0.2%), Ministry of Education of China Dec. 2021
- First Class Scholarship Nov. 2021

SKILLS

Programming: Python, PyTorch, C/C++, Java, MATLAB, HTML/CSS/JavaScript, Shell, TEX, Verilog
Tools: Docker, Git, Linux, Slurm
Languages: English, Chinese