

Xinyuan Li

✉ xinyuanli2327@outlook.com Website GitHub Scholar

Education

Zhejiang University

B.E. in Automation Engineering; Minor in ACEE Honor Class of CKC

Sept 2023 - Present
Hangzhou, Zhejiang

Research Experience

MLL Lab, Northwestern University

Research Intern Advisor: Prof. Manling Li

Jul 2025 - Present

The Cost of Thinking: Real-Time VLA Agents

Apr 2026 - Present

- Led development of a **15-task VLA benchmark** across Atari, action games, continuous control, and embodied simulation.
- Co-designed **latency-aware VLA training with RL teachers**; led implementation and outperformed baselines in real time.
- Validated latency simulation on five GPU platforms; Temporal profiles achieved the **lowest mean real-simulation return gap**.
- Mapped fixed-delay generalization and fixed/stochastic/mixed-profile robustness, including companion-task co-training; identified task-dependent, asymmetric transfer.
- Ablated latency prompts, context length, observation stride, fine-tuning scope, and memory (frame stacks, ghost trails, mosaic, KV cache, WanOFT); timing probes showed ordered context supports decodable action timing.

Geometric Consistency in Unified Multimodal Models

Jul 2025 - Apr 2026

- Designed an end-to-end study spanning test-time view scaling, five-task CO3D fine-tuning, interleaved reasoning data, and geometry-aware training.
- Built multi-view spatial supervision with quality control and PCA-based camera-trajectory calibration; retained 6,697 of 20,286 CO3D sequences across 51 categories in the final documented filter.
- Implemented contrastive/geometric objectives and diagnostic probes in UMMs(Bagel).

Open Source Contributions

starVLA [GitHub, 3.6k+ stars]

Jun 2026 - Jul 2026

- Diagnosed a gradient-accumulation bug in Accelerate-based training that discarded earlier microbatch gradients, reducing each accumulated update to its final microbatch. [Issue]
- Submitted a six-file fix across three VLA/VLM trainers to align optimization, scheduling, logging, evaluation, and checkpointing with true Accelerate/DeepSpeed update boundaries, with a focused regression test. [PR]

Publications & Preprints

[1] The Cost of Thinking: Evaluating Latency Sensitivity in Real-Time Foundation Model Agents

Xinyuan Li*, Talha Chafekar*, Zihan Wang*, Zeju Li*, Qineng Wang, Ruohan Zhang, Zhiding Yu, Jiajun Wu, Manling Li

*Equal contribution.

Under Review at ICLR 2027

[2] ACTUAL: Learning Action Chunks under Variable Inference Latency

Xinyuan Li, Manling Li

Preprint

Contribution:

- Proposed and implemented ACTUAL, an asynchronous action-chunk RL framework that accounts for variable delays between observation and action execution during policy training.
- Designed value estimates using the state at execution start and each action's planned execution count, with rewards accumulated and future returns discounted over the actual execution duration; implemented in FastTD3.

[3] MindEdit-Bench: Benchmarking Object-Level Counterfactual Spatial Reasoning in VLMs from In-the-Wild Photos

Leyuan Yu*, Xiao Tang*, Minghao Liu*, Xinyuan Li, Xiaokai Bai, Sheng Zhou, Qunshu Lin, Weihao Xuan, Naoto Yokoya

*Equal contribution.

arXiv preprint [Paper] [Dataset]

Contribution: Built and ran a unified evaluation pipeline spanning API-served VLMs and locally hosted VLMs via vLLM; evaluated all 15 models on 1,003 human-verified questions; helped refine benchmark questions; created result visualizations; and contributed to manuscript writing.

Selected Projects

CMU 15-213: Introduction to Computer Systems

Jan 2026 - Feb 2026

- Studied machine-level execution, memory hierarchies, process control, and dynamic memory management through core CS:APP labs. Applied these concepts to reason about program correctness and performance. [\[Code\]](#)

CMU 11-868: Large Language Model Systems

Jul 2025 - Sep 2025

- Built a full-stack LLM system from scratch, including auto-differentiation, CUDA kernels for Transformer components, distributed training, and inference optimizations. [\[Code\]](#)

Mathematical Foundations of Reinforcement Learning

Apr 2025 - May 2025

- Studied RL convergence and optimality, implemented Value Iteration, Policy Iteration, and Q-Learning, and wrote theory-focused notes. [\[Notes\]](#)