

Lingzhi Yuan

lingzhiy@umd.edu | (+1) 305-216-0051 | lingzhiyxp.github.io | GitHub | LinkedIn | Google Scholar

EDUCATION

University of Maryland, College Park

Aug. 2025 – Present

Ph.D. in Computer Science (expected May 2030) | GPA: 4.00/4.00

Advisor: Prof. Furong Huang

Zhejiang University

Sep. 2021 – Jun. 2025

B.Eng. in Automation | GPA: 3.96/4.00 | Rank: 1/58

Minor in Intensive Training Honors Program (ITP), Chu Kochen Honors College (Top 40/5400)

EXPERIENCE

Furong Huang Lab, University of Maryland

Aug. 2025 – Present

Research Assistant | Self-improving agents, trustworthy ML, and multi-agent systems, advised by Prof. Furong Huang.

UNIUBI AI, Hangzhou, China

Mar. 2025 – Jun. 2025

Student Intern | Instruction following and data augmentation for vision-language models.

Secure Learning Lab, The University of Chicago

Mar. 2024 – Dec. 2024

Research Assistant | Trustworthy ML and multimodal safety, advised by Prof. Bo Li.

Nanyang Technological University, Singapore

Aug. 2023 – Sep. 2023

Summer School Participant | Machine Learning and Its Applications program, led by Prof. Kezhi Mao.

SELECTED PUBLICATIONS

* denotes equal contribution

Meta²: Precomputing Reusable Knowledge for Online Harness Evolution in Long-Horizon Agents

Lingzhi Yuan*, Yiguang Cai*, Weize Liu, Feijiang Han, Yongkyun Lee, Udari Madhushani Sehwag, Furong Huang
Preprint

FlowBank: Query-Adaptive Agentic Workflow Optimization through Precompute-and-Reuse

Lingzhi Yuan, Chenghao Deng, Fangxu Yu, Souradip Chakraborty, Mohammad Rostami, Furong Huang
Conference on Neural Information Processing Systems (NeurIPS), 2026

PromptGuard: Soft Prompt-Guided Unsafe Content Moderation for Text-to-Image Models

Lingzhi Yuan*, Xinfeng Li*, Chejian Xu, Guan hong Tao, Xiaojun Jia, Yihao Huang, Wei Dong, Yang Liu, Bo Li
IEEE Transactions on Information Forensics and Security (TIFS)

TSRBench: A Comprehensive Multi-task Multi-modal Time Series Reasoning Benchmark for Generalist Models

Fangxu Yu, Xingang Guo, Lingzhi Yuan, Haoqiang Kang, Hongyu Zhao, Lianhui Qin, Furong Huang, Bin Hu, Tianyi Zhou
International Conference on Machine Learning (ICML), 2026

MMDT: Decoding the Trustworthiness and Safety of Multimodal Foundation Models

Chejian Xu*, Jiawei Zhang*, Zhaorun Chen*, Chulin Xie*, Mintong Kang*, Zhuowen Yuan*, Zidi Xiong*, Chenhui Zhang, Lingzhi Yuan, Yi Zeng, Peiyang Xu, Chengquan Guo, Andy Zhou, Jeffrey Ziwei Tan, Zhun Wang, Alexander Xiong, Xuandong Zhao, Yu Gai, Francesco Pinto, Yujin Potter, Zhen Xiang, Zinan Lin, Dan Hendrycks, Dawn Song, Bo Li
International Conference on Learning Representations (ICLR), 2025

Hypothesis-Conditioned Evaluation for Efficient Harness Optimization

Lingzhi Yuan*, Zhi Zhang*, Arash A. Amini, Furong Huang
Under review

LatentShield: Composable Latent Multi-Agent Systems with Decision-Boundary Monitoring

Baicheng Chen*, Anirudh Satheesh*, Lingzhi Yuan, Drew Gjerstad, Andrew Mendz, Ami Britti Antoni Gnanaraj, Mingyu Kang, Benjamin Leinwand, Andrew Adams Schoen, Furong Huang
Under review

SELECTED PROJECTS

Online Harness Evolution for Long-Horizon AgentsJul. 2026 – Oct. 2026

Project Lead | Advisor: Prof. Furong Huang, University of Maryland

A precompute-and-reuse paradigm that learns reusable evolution knowledge offline to guide task-local harness edits within a run.

- Formulated harness evolution by edit timing, available evidence, and affected executions, distinguishing offline optimization from pre-run, within-run, and after-run adaptation.
- Learned a textual Meta-Skill offline from contrastive rollouts (no evolver / Meta-Skill-guided / verifier-informed) that guides an online evolver to keep or edit the harness at checkpoints, without test-time verifier access.
- Achieved the best performance among evaluated methods on Terminal-Bench 2.1 (81.00%) and DeepSWE 1.1 (44.44%) with GPT-5.6-Terra; released as a preprint.

Query-Adaptive Agentic Workflow OptimizationNov. 2025 – May 2026

Project Lead | Advisor: Prof. Furong Huang, University of Maryland

A portfolio-based framework that precomputes a bank of complementary agentic workflows offline and routes each query to the best one at inference time.

- Showed workflow complementarity: discarded task-level workflows retain reuse value, and query-level generation gains can be recovered by cheaper precomputed workflows.
- Designed a three-stage pipeline: complementarity-oriented MCTS search, coverage-aware portfolio curation, and GNN-based query–workflow matching.
- Best average score across five benchmarks, improving over the strongest automated and handcrafted baselines by 4.26% and 14.92% (relative); published at NeurIPS 2026.

Soft Prompt-Guided Defense for Text-to-Image ModelsJun. 2024 – Dec. 2024

Project Lead | Advisor: Prof. Bo Li, University of Chicago / UIUC

An effective and efficient framework to mitigate NSFW generation across unsafe categories in T2I models.

- Surveyed safety protection techniques for T2I models and identified key open challenges.
- Built a multi-category moderation framework achieving state-of-the-art NSFW moderation with high time efficiency; published in IEEE TIFS.

Comprehensive Assessment of Trustworthiness in Multimodal ModelsMar. 2024 – Oct. 2024

Research Assistant | Advisor: Prof. Bo Li, University of Chicago / UIUC

Evaluated the safety and trustworthiness of multimodal foundation models.

- Applied red-teaming algorithms (MMP, GCG for T2I; Attack-VLM for I2T) on surrogate models to build an adversarial robustness benchmark.
- Evaluated state-of-the-art models (e.g., DALL-E 3, GPT-4o) and analyzed their vulnerability under adversarial conditions; published at ICLR 2025.

SKILLS

ProgrammingPython, C + + , MATLAB

ToolsPyTorch, Transformers, Diffusers, OpenCV, Git, Linux/Unix

LanguagesTOEFL 105 (Reading 30, Listening 27, Speaking 23, Writing 25)

HONORS AND AWARDS

University of Maryland Graduate School Dean’s Fellowship2025

Infineon Power Semiconductor Scholarship2023, 2024

Zhejiang University Second Class Scholarship2022, 2023

Nitori International Scholarship2022

SERVICES

Teaching Assistant, CMSC 320: Introduction to Data Science, University of MarylandFall 2025

Journal Reviewer, International Journal of Computer Vision (IJCV)

Conference Reviewer, International Conference on Learning Representations (ICLR)2027