

Ziyu Ye

ziyu@uchicago.edu · [GitHub](#) · [Homepage](#)

Research Interests

Reinforcement Learning · {Pre, Post}-Training Large Models · Self-Supervised Learning · Variational Diffusion Models

Employment

- 2025.09 - Present **Incoming Research Scientist**, Google DeepMind
- 2024.09 - 2025.06 **Student Researcher**, Google DeepMind
- 2024.06 - 2024.09 **Research Intern**, Google DeepMind
Gemini post-training for AI Alignment and Reasoning
- Related work: [Asymmetric Self-Play for RLHF](#)
- 2023.12 - 2024.06 **Co-Founder**, Eigent AI
Learning and Planning with LLM-based Agents
- Related work: [Reasoning in Reasoning](#), [Agent OWL](#)

Education

- 2025 **Ph.D. in Computer Science**, The University of Chicago
Advisor: Prof. Yuxin Chen
Thesis: Self-Play Methods in Reinforcement Learning for Language Models
- 2020 **M.S. in Computational Analysis and Public Policy**, The University of Chicago
- 2018 **B.S. in Mathematical Economics**, Xi'an Jiaotong University
Honored Graduate of [Tsien Hsue-Shen School](#), [Class of Gifted Young](#)
Advisor: Prof. Yusen Kwoh
Thesis: Algorithmic Approaches to Solving Bounded Rationality

Selected Preprints and Publications

- P1 **Training LLMs via Asymmetric Self-Play: Reward-Guided Prompt Evolving in RLHF**
Ziyu Ye, Rishabh Agarwal, Tianqi Liu, Rishabh Joshi, Sarmishta Velury, Quoc V. Le, Qijun Tan, Yuan Liu
ICML 2025 · NeurIPS LanGame 2024 ([Spotlight](#)) · [Paper](#)
[rl](#) [llm](#)
- P2 **Reasoning in Reasoning: The Hierarchical Framework for Neural Theorem Proving**
Ziyu Ye, Jiacheng Chen, Jonathan Light, Yifei Wang, Jiankai Sun, Mac Schwager, Philip Torr, Guohao Li, Yuxin Chen, Kaiyu Yang, Yisong Yue, Ziniu Hu
In Submission · NeurIPS MATH-AI 2024 · [Paper](#) · [Code](#)
[rl](#) [reasoning](#) [llm](#)
- P4 **Understanding Bias in Deep Anomaly Detection: A Semi-Supervised View with PAC Guarantees**
Ziyu Ye, Yuxin Chen, Heather Zheng
IJCAI 2022 ([Oral](#)) · [Paper](#) · [Code](#)
[learning theory](#)

- P5 **Don't Be Pessimistic Too Early: Look k Steps Ahead in Offline RL**
Chaoqi Wang, Ziyu Ye, Kevin Murphy, Yuxin Chen.
AISTATS 2024 · [Paper](#) · [Code](#)
rl
- P6 **Improving Contextual Bandits via Post-serving Features**
Chaoqi Wang, Ziyu Ye, Zhe Feng, Ashwinkumar Badanidiyuru, Haifeng Xu
NeurIPS 2023 (Spotlight) · [Paper](#) · [Code](#)
rl learning theory
- P7 **When More is Less: Understanding Chain-of-Thought Length in LLMs**
Yuyang Wu, Yifei Wang, Ziyu Ye, Tianqi Du, Stefanie Jegelka, Yisen Wang.
In Submission · [Paper](#)
rl reasoning llm
- P8 **The Price of Sparsity: Memorization and Generalization in Sparse Neural Networks**
Ziyu Ye, Chaoqi Wang, Yuxin Chen
ICML Workshop on Sparsity in Neural Networks 2022 · [Paper](#) · [Code](#)
sparse training
- P9 **Understanding the Role of Equivariance in Self-supervised Learning**
Yifei Wang, Kaiwen Hu, Sharut Gupta, Ziyu Ye, Yisen Wang, Stefania Jegelka
NeurIPS 2024 · [Paper](#) · [Code](#)
learning theory
- P10 **Towards Provably Efficient Quantum Algorithms for Large-Scale Machine Learning Models**
Junyu Liu, Minzhao Liu, Jin-Peng Liu, Ziyu Ye, Yuri Alexeev, Jens Eisert, Liang Jiang
Nature Communications 2023 · [Paper](#) · [Code](#)
learning theory
- P11 **Efficient Online Decision Tree Learning with Active Feature Acquisition**
Arman Rahbar, Ziyu Ye, Yuxin Chen, Morteza Haghir Chehreghani
IJCAI 2023 · [Paper](#)
rl
- P12 **OWL: Optimized Workforce Learning for Generalist Agents**
Meng kang Hu, Yuhang Zhou, Wendong Fan, Yuzhou Nie, Bowei Xia, Tao Sun, Ziyu Ye, Zhaoxuan Jin, Yingru Li, Qiguang Chen, Zeyu Zhang, Yifeng Wang, Qianshuo Ye, Bernard Ghanem, Ping Luo, Guohao Li
In Submission · [Paper](#) · [Code](#)
rl llm

Open-Source Contributions

- CAMEL State-of-the-art LLM-based multi-agent framework
<https://github.com/camel-ai/camel>
- RiR Official PyTorch implementation of hierarchical neural theorem proving
<https://github.com/ziyu-deep/reasoning-in-reasoning>
- OWL A hierarchical training pipeline for generalist agents
<https://github.com/camel-ai/owl>

Teaching

- 2025 CMSC 354 – Machine Learning, Teaching Assistant
Department of Computer Science, The University of Chicago
- 2022 CMSC 353 – Mathematics for Machine Learning, Teaching Assistant
Department of Computer Science, The University of Chicago
- 2020 CAPP 303 – Civic Data and Technology, Teaching Assistant
Irvine B. Harris School of Public Policy, The University of Chicago

Academic Services

- Reviewer ICML, NeurIPS, ICLR, AISTATS, IJCAI, AAAI, UAI, JMLR
- Workshop Organizer [Scaling Environments for Agents](#) at NeurIPS 2025

Selected Awards

- 2012 First Place Award, China Junior Math Olympiad
- 2020 John Crerar Fellowship, The University of Chicago
- 2023 NSF Innovation Corps Award, U.S. National Science Foundation