Guowei Xu

xgw23@mails.tsinghua.edu.cn xugw-kevin.github.io Kevin_GuoweiXu XuGW-Kevin

Looking for an onsite visiting research internship starting from 2026 spring

Education

Institute for Interdisciplinary Information Sciences, Tsinghua University *Sep. 2023 – Present*
BS in Computer Science and Technology

- GPA: 3.98/4.00 (15 A+, 15 A)

Research Interests

My research interests include but not limit to language models (LLMs/VLMs) and reinforcement learning.

Publications & Preprints

[1] metaTextGrad: Automatically optimizing language model optimizers arXiv 2025 (in submission)
Guowei Xu, Mert Yuksekgonul, Carlos Guestrin, James Zou

[2] LLaVA-CoT: Let Vision Language Models Reason Step-by-Step ICCV 2025
Guowei Xu*, Peng Jin*, Ziang Wu*, Hao Li, Yibing Song, Lichao Sun, Li Yuan

[3] MENTOR: Mixture-of-Experts Network with Task-Oriented Perturbation for Visual Reinforcement Learning ICML 2025
Suning Huang*, Zheyu Zhang*, Tianhai Liang, Yihan Xu, Zhehao Kou, Chenhao Lu, **Guowei Xu**, Zhengrong Xue, Huazhe Xu

[4] ACE : Off-Policy Actor-Critic with Causality-Aware Entropy Regularization **ICML 2024 (Oral)**
Tianying Ji*, Yongyuan Liang*, Yan Zeng, Yu Luo, **Guowei Xu**, Jiawei Guo, Ruijie Zheng, Furong Huang, Fuchun Sun, Huazhe Xu

[5] DrM: Mastering Visual Reinforcement Learning through Dormant Ratio Minimization **ICLR 2024 (Spotlight)**
Guowei Xu*, Ruijie Zheng*, Yongyuan Liang*, Xiyao Wang, Zhecheng Yuan, Tianying Ji, Yu Luo, Xiaoyu Liu, Jiaxin Yuan, Pu Hua, Shuzhen Li, Yanjie Ze, Hal Daumé III, Furong Huang, Huazhe Xu

[6] Can Pre-Trained Text-to-Image Models Generate Visual Goals for Reinforcement Learning? NeurIPS 2023
Jialu Gao*, Kaizhe Hu*, **Guowei Xu**, Huazhe Xu

Research Experiences

Research Assistant (LLM Optimization/Reasoning) *Stanford, CA*
Advisor: Prof. James Zou, Stanford University *Jun. 2024 – Present*

- Optimize LLM optimizer via meta-learning and align it with specific tasks, achieving an average performance improvement of 6% on benchmarks including BBH, MMLU, and GPQA.

Research Assistant (VLM Reasoning) *Shenzhen, China*
Advisor: Prof. Li Yuan, Peking University *Jul. 2024 – Mar. 2025*

- Introduce LLaVA-CoT, the first visual language model capable of spontaneous and systematic reasoning, and further improve its performance via SWIRES, an inference-time scaling algorithm, achieving an average 9.4% performance gain over the base model on benchmarks including MMBench, MMStar, MMVet, MathVista, AI2D and HallusionBench.

Research Assistant (Reinforcement Learning) *Beijing, China*
Advisor: Prof. Huazhe Xu, Tsinghua University *Jan. 2023 - Jun. 2024*

- Discover dormant ratio, a critical intrinsic model metric that affects the sample efficiency of reinforcement

learning.

- Propose the DrM reinforcement learning algorithm, which outperforms the best existing baselines by 65%, 35%, and 75% in terms of asymptotic performance on the DeepMind Control Suite, MetaWorld, and Adroit benchmarks, respectively. Subsequently co-develop two additional RL algorithms, ACE and MENTOR: ACE outperforms the long-standing SOTA algorithm SAC by 40% across 29 diverse continuous control tasks; MENTOR is the first RL algorithm to enable real-world (non-simulated) robot learning.

Professional Services

- **Reviewers:** ICLR 2025, ICML 2025, NeurIPS 2025, CoRL 2023 workshop.
- **Teaching Assistant:** Deep Reinforcement Learning (Graduate Course), Natural Language Processing.

Honors and Awards

Sparking Program *May. 2025*
The most prestigious and selective academic organization for students at Tsinghua University

National Scholarship *Nov. 2024*
Top 0.5% undergarduate students in China

Freshman First-Class Scholarship, Tsinghua University *Sep. 2023*
Top scholarship for freshman in Tsinghua University

Overall winner (best total score), 52nd International Physics Olympiad (IPhO) *Jul. 2022*
The participation is open for the 5 best physics students of a country. 371 students from more than 70 countries participated at IPhO2022.

Selected Courses

CATEGORY	COURSE	GRADES
Research	Students Research Training	A+
Artificial Intelligence	Machine Learning	A+
	Advanced Computer Graphics	A+
	Computer Vision	A+
	Natural Language Processing	A+
	Embodied Artificial Intelligence	A+
	Introduction to Large Language Model Applications	A+
	Deep Learning	A
	Artificial Intelligence: Principles and Techniques	A
Computer Science	Introduction to Programming in C/C++	A+
	Mathematics for Computer Science	A
	Introduction to Computer Systems	A
	Algorithm Design	A
	Theory of Computation	A
	Quantum Computer Science	A