TAIYI WANG

CONTACT

+44 7788 475798
+86 13810893171

Taiyi.Wang@cl.cam.ac.uk
wty0107@gmail.com

Cambridge,
United Kingdom
CB1 2GB

https://kevinwty0107.github.io/

https://www.linkedin.com/in/taiyi-wang-3438121a2/

ACADEMIC ACTIVITITES

- Reviewer for ICLR 2025(**Highlighted Reviewer**), 2026
- Reviewer for NeurIPS 2022, 2024(**Top Reviewer**), 2025
- Reviewer for ICML 2025, 2026
- Program Committee for Euro MLSys, 2022, 2023, 2024, 2025
- Reviewer for SIGMOD 2026

PROFESSIONAL SUMMARY

Diligent researcher with a robust academic and industry background in applied reinforcement learning and LLM agent development. Drawing on my extensive lectureship and TA experience from my senior PhD years, I effectively co-supervised students, combining advanced research with practical, hands-on lessons to empower young teams and deliver real-world results. I owned a successful AI-centric startup before.

INTERESTS AND EXPERTISE

Reinforcement Learning, LLM-Agent, Scalable Machine Learning, ML-enhanced System Optimization

EDUCATION

January 2026

Doctor of Philosophy in Computer Science
University of Cambridge

June 2021

Master of Engineering in Computer Science
Johns Hopkins University

July 2019

Bachelor of Science in Physics, minor in math
Peking University

WORKING EXPERIENCE

July 2025 – Dec 2025

Student Researcher
Google DeepMind, London

- Autonomous Agents Team, led by Edward Grefenstette

- Developed RLFT/post-training pipeline, authored patent and paper

June 2024 - November 2024

Research Scientist Intern

Noah's Ark Research Center UK, London

- RL x LLM Agent Group
- Developed large-scale RLFT framework for VLM-based mobile agent

September 2022 - May 2024

Co-founder and Chief Technology Officer

Powersense Technology Limited, Cambridge

November 2018 - May 2019

Research and Development Engineer

Baidu, Beijing

- Big Data Group

September 2018 - November 2018

Research and Development Intern

Sohu.com Ltd, Beijing

- Advertising Strategy Group

HIGHLIGHTED RESEARCH EXPERIENCE

Research Assistant, Center for Data Science Peking University and Princeton University (2019 Fall)
Instructed by: Prof. Weinan E (Princeton), Zaiwen Wen (BICMR)
Project: Building Reinforcement Learning Simulator for Continuous Control

Visiting Scholar, CUSP Big Data Interaction (BDI) Lab, NYU (2018 Summer)
Instructed by: Prof. Huy T. Vo
Project: Simulating Taxi Ride Sharing at Scale

Research Assistant, Intelligent Transportation Research Center, EECS, MIT (2018 Spring)
Instructed by: Prof. Berthold K.P. Horn, Dr. Yajun Fang

Project: Building Bilateral Control Simulation Platform (BCSP)

AWARDS

- Student Award, ACM SIGMOD 2025 (International Conference on Management of Data)
- Pillman and Cody Award, University of Cambridge
- Runner-up in Shenzhen Innovation and Entrepreneurship Competition, Global Final, Shenzhen, China
- Runner-up in the Chris Abell postdoc Business Plan Competition, Cambridge, United Kingdom
- Finalist (F Awards, Top 1%), Mathematical Contest in Modeling (MCM)
- Excellent Graduate Student Award in School of Physics, Peking University
- Excellent Graduation thesis, Peking University
- Special Award, the 5th Youth Physics Tournament, Peking University
- Third Prize, Freshman Scholarship, Peking University

PUBLICATIONS

---2026----

- Taiyi Wang, Sian Gooding, Florian Hartmann, Oriana Riva, Edward Grefenstette. "A Subgoal-driven Framework for Improving Long-Horizon LLM Agents", under review
- Taiyi Wang, Youhe Jiang, Fangcheng Fu, Guoliang He, Eiko Yoneki. "OServe: Accelerating LLM Serving via Spatial-Temporal Workload Orchestration", under review

---2025----

- Taiyi Wang, Zhihao Wu, Jianehng Liu, et al. "DistRL: An Asynchronous Distributed Reinforcement Learning Framework for On-Device Control Agents", the Thirteenth International Conference on Learning Representations (ICLR'25)
- Taiyi Wang, Liang Liang, Guang Yang, Thomas Heinis, Eiko Yoneki. "A New Paradigm in Tuning Learned Indexes: A Reinforcement Learning-Enhanced Approach", the International Conference on Management of Data (SIGMOD'25)
- Youhe Jiang, Fangcheng Fu, Xiaozhe Yao, Taiyi Wang, Bin Cui, Ana Klimovic, Eiko Yoneki. "ThunderServe: High-performance and Cost-efficient LLM Serving in Cloud Environments", the Eighth Annual

Conference on Machine Learning and Systems (MLSys'25)

- Taiyi Wang, Jianheng Liu, Bryan Lee, Zhihao Wu, Yu Wu. "OCMDP: Observation-Constrained Markov Decision Process", International Joint Conference on Neural Networks (IJCNN'25)
- Youhe Jiang*, Fangcheng Fu*, Taiyi Wang*, Eiko Yoneki. "LiveServe: Efficient LLM Serving via Workload-aware Scheduling and Workload-adaptive Switching", under review
- Taiyi Wang, Zakir Singh, Eiko Yoneki. "AutoIndexer: A Reinforcement Learning-Enhanced Index Advisor Towards Scaling Workloads", under review
- Wenxuan Li*, Taiyi Wang*, Eiko Yoneki. "HiBO: Hierarchical Bayesian Optimization via Adaptive Search Space Partitioning", under review

---2024----

- Taiyi Wang, Eiko Yoneki. "IA2: Leveraging Instance-Aware Index Advisor with Reinforcement Learning for Diverse Workloads", the 4th Workshop on Machine Learning and Systems of EuroSys, 2024 (EuroMLSys'24)
- Taiyi Wang, Eiko Yoneki. "Enhancing Generalization through Task Vector Fusion in Deep Reinforcement Learning for Database Optimization", European Conference on Computer Systems (EuroSys), Poster, 2024
- George-Octavian Bărbulescu, Taiyi Wang, Zak Singh, Eiko Yoneki. "Learned Graph Rewriting with Equality Saturation: A New Paradigm in Relational Query Rewrite and Beyond", under review

---2023 and Before (Selected)----

- Ali Hadian, Behzad Ghaffari, Taiyi Wang, Thomas Heinis. "COAX: Correlation-Aware Indexing on Multidimensional Data with Soft Functional Dependencies", DBML workshop of the 39th IEEE International Conference on Data Engineering (DBML@ICDE), 2023
- Hao Sun, Taiyi Wang. "Toward Causal-Aware RL: State-Wise Action-Refined Temporal Difference", NeurIPS 2022 Workshop DeepRL, 2022
- Hao Sun, Ziping Xu, Meng Fang, Zhenghao Peng, Taiyi Wang, Bolei Zhou. "Constrained MDPs can be Solved by Early-Termination with Recurrent Models", NeurIPS 2022 Workshop FMDM, 2022
- Hao Sun, Ziping Xu, Taiyi Wang, Meng Fang, Bolei Zhou. "Supervised Q-Learning can be a Strong Baseline for Continuous Control", NeurIPS 2022 Workshop FMDM, 2022
- Taiyi Wang, Jiahao Shi. "Solving Maximal Stable Set problem via Deep Reinforcement Learning", International Conference on Communication Technology, Computational Engineering and Artificial Intelligence (ICAART'20)