Yuheng Zhao

Department of CSE, HKUST, Hong Kong SAR
https://farrrrland.github.io/
yzhaoep@cse.ust.hk

ABOUT ME

I am a fourth-year Ph.D. candidate in CSE at HKUST, advised by Prof. Wei Wang. My research focuses on scalable systems for reinforcement learning and LLM post-training, with interests in machine learning systems and cloud infrastructure. My work has appeared in top-tier systems venues, including OSDI and NSDI.

EDUCATION

Fudan University | CS | BSc — Sep 2018 - Jun 2022

- Research assistant at Cryptography and Information Security Laboratory, supervised by Prof. Haibin Kan.

The Hong Kong University of Science and Technology | CSE | Ph.D. — Sep 2022 - Jun 2027

EXPERIENCE

Alibaba Group *ROLL* Team - Research Intern - Beijing — Mar 2025 - Present

- Contributed to the development of ***ROLL***, an LLM post-training framework with 3k+ stars on GitHub.
- Proposed ***Tail Batching*** to mitigate the long-tail rollout inefficiency, speeding up on-policy training by up to **2.56×**.
- Proposed ***Cooperative Elasticity***, achieving up to **3.3×** higher throughput for agentic RL training.
- Collaborated closely with MoonCake team to develop solutions for real-world workloads.

Microsoft Research Asia DKI Group - Research Intern - Beijing — Nov 2021 - Jun 2022

- Conducted a large-scale characterization study of Azure private and public cloud workloads, revealing key differences in deployment and utilization patterns with implications for cloud management.
- Received the MSRA "Stars of Tomorrow" Award of Excellence (**Top 10%** of research interns).

University of Chicago, Globus Lab - Research Assistant - Chicago — May 2021 - Oct 2021

- Contributed to *FuncX*, a federated FaaS platform, by implementing a container pool and refining the eviction strategy to mitigate container explosion problem.

PUBLICATIONS

- Wei Gao*, ***Yuheng Zhao****, Tianyuan Wu*, Shaopan Xiong*, Weixun Wang*, Dakai An, Lunxi Cao, Dilxat Muhtar, Zichen Liu, Haizhou Zhao, Ju Huang, Siran Yang, Yongbin Li, Wenbo Su, Jiamang Wang, Lin Qu, Bo Zheng, Wei Wang, "RollArt: Disaggregated Multi-Task Agentic RL Training at Scale", (*Equal contribution) [**OSDI'26**][CCF-A]
- Wei Gao*, ***Yuheng Zhao****, Dakai An, Tianyuan Wu, Lunxi Cao, Shaopan Xiong, Ju Huang, Weixun Wang, Siran Yang, Wenbo Su, Jiamang Wang, Lin Qu, Bo Zheng, Wei Wang, "RollPacker: Taming Long-Tail Rollouts for RL Post-Training with Tail Batching", (*Equal contribution) [**NSDI'26**][CCF-A]
- Tianyuan Wu, Lunxi Cao, Yining Wei, Wei Gao, ***Yuheng Zhao***, Dakai An, Shaopan Xiong, Zhiqiang Lv, Ju Huang, Siran Yang, Yinghao Yu, Jiamang Wang, Lin Qu, Wei Wang, "Weave: Efficient Co-Scheduling for Disaggregated RL Post-Training" [**OSDI'26**][CCF-A]
- Xiaoting Qin, Minghua Ma, ***Yuheng Zhao***, Jue Zhang, Chao Du, Yudong Liu, Anjaly Parayil, Chetan Bansal, Saravan Rajmohan, Íñigo Goiri, Eli Cortez, Si Qin, Qingwei Lin, Dongmei Zhang, "How Different are the Cloud Workloads? Characterizing Large-Scale Private and Public Cloud Workloads" [**DSN'23**][CCF-B]
- ROLL Team, "Reinforcement Learning Optimization for Large-Scale Learning: An Efficient and User-Friendly Scaling Library" [Tech Report]
- Wei Gao*, ***Yuheng Zhao****, Dilxat Muhtar, Dakai An, Xuchun Shang, Tianyuan Wu, Lunxi Cao, Shaopan Xiong, Weixun Wang, Ju Huang, Teng Ma, Siran Yang, Jiamang Wang, Lin Qu, Bo Zheng, Wei Wang, "ROSE: Rollout On Serving GPUs via Cooperative Elasticity for Agentic RL", (*Equal contribution) [In Submission]