Yuqing Yang

E-mail: yyang063@usc.edu ✻ *Phone:* +1-213-696-4881
Website: https://ayyyq.github.io/
Address: 1246 W 30th St, Los Angeles, California, 90007, USA

Education

University of Southern California *Los Angeles, USA*
PhD student in Computer Science *August 2024 - Present*
Advisor: Professor Robin Jia
Research Focus: Self-improving LLMs and agents; trustworthy human–AI collaboration; efficient LLMs and agents.

Fudan University *Shanghai, China*
Master of Science in Computer Science *September 2021 - June 2024*
Advisor: Professor Xipeng Qiu
Research Focus: Information extraction, hallucination, reasoning.

University of Chinese Academy of Sciences *Beijing, China*
Bachelor of Engineering in Computer Science *September 2017 - June 2021*
GPA: 3.94/4.00, Rank: 3/112

Publications

* *indicates equal contribution.*

Unsupervised Tool Evolution for Coding Agents in Data-Lake Analytics
- Yuqing Yang, Yurong Liu, Deqing Fu, Manav Garg, Sravan Babu Bodapati, Fatma Ozcan, Alon Y. Halevy, Jiani Zhang
- Under Review

Rewards Beyond Test Passing: Towards Generalizable Fidelity in LLM Debugging
- Yunkai Zhan, Miaosen Chai, Shangshang Wang, Jike Zhong, Yuqing Yang, Deqing Fu, Serina Chang, Wang Bill Zhu
- Under Review

Self-Evolving LLM Memory Extraction Across Heterogeneous Tasks
- Yuqing Yang, Tengxiao Liu, Wang Bill Zhu, Taiwei Shi, Linxin Song, Robin Jia
- arXiv preprint 2026

Can Coding Agents Optimize Algorithms Autonomously?
- Tengxiao Liu, Yuqing Yang, Xi Ye, Danqi Chen
- Blog 2026

Memory Retrieval for Changing Preferences
- Yuehan Qin, Li Li, Linxin Song, Wei Yang, Jiate Li, Yuqing Yang, Yue Zhao
- **NeurIPS 2026**

When LLMs Read Tables Carelessly: Measuring and Reducing Data Referencing Errors
- Yuqing Yang, Qi Zhu, Zhen Han, Boran Han, Zhengyuan Shen, Shuai Wang, Vassilis N. Ioannidis, Huzefa Rangwala
- **ACL 2026 Oral**

When Do LLMs Admit Their Mistakes? Understanding the Role of Model Belief in Retraction

- Yuqing Yang, Robin Jia
- **CoLM 2026 Oral & KnowFM@ACL 2026 Best Paper**

Weak-to-Strong Reasoning
- Yuqing Yang, Yan Ma, Pengfei Liu
- **EMNLP Findings 2024**

Alignment for Honesty
- Yuqing Yang, Ethan Chern, Xipeng Qiu, Graham Neubig, Pengfei Liu
- **NeurIPS 2024**

OlympicArena: Benchmarking Multi-discipline Cognitive Reasoning for Superintelligent AI
- Zhen Huang, Zengzhi Wang, Shijie Xia, Xuefeng Li, Haoyang Zou, Ruijie Xu, Run-Ze Fan, Lyumanshan Ye, Ethan Chern, Yixin Ye, Yikai Zhang, Yuqing Yang, Ting Wu, Binjie Wang, Shichao Sun, Yang Xiao, Yiyuan Li, Fan Zhou, Steffi Chern, Yiwei Qin, Yan Ma, Jiadi Su, Yixiu Liu, Yuxiang Zheng, Shaoting Zhang, Dahua Lin, Yu Qiao, Pengfei Liu
- **NeurIPS Datasets and Benchmarks 2024**

BeHonest: Benchmarking Honesty of Large Language Models
- Steffi Chern*, Zhulin Hu*, Yuqing Yang*, Ethan Chern, Yuan Guo, Jiahe Jin, Binjie Wang, Pengfei Liu
- arXiv preprint 2024

Full Parameter Fine-tuning for Large Language Models with Limited Resources
- Kai Lv, Yuqing Yang, Tengxiao Liu, Qinghui Gao, Qipeng Guo, Xipeng Qiu
- **ACL 2024 Oral**

Plan, Verify and Switch: Integrated Reasoning with Diverse X-of-Thoughts
- Tengxiao Liu, Qipeng Guo, Yuqing Yang, Xiangkun Hu, Yue Zhang, Xipeng Qiu, Zheng Zhang
- **EMNLP 2023**

CoLLiE: Collaborative Training of Large Language Models in an Efficient Way
- Kai Lv*, Shuo Zhang*, Tianle Gu, Shuhao Xing, Jiawei Hong, Keyu Chen, Xiaoran Liu, Yuqing Yang, Honglin Guo, Tengxiao Liu, Yu Sun, Qipeng Guo, Hang Yan, Xipeng Qiu
- **EMNLP Demo 2023**

An AMR-based Link Prediction Approach for Document-level Event Argument Extraction
- Yuqing Yang*, Qipeng Guo*, Xiangkun Hu, Yue Zhang, Xipeng Qiu, Zheng Zhang
- **ACL 2023**

DORE: Document ordered relation extraction based on generative framework
- Qipeng Guo*, Yuqing Yang*, Hang Yan, Xipeng Qiu, Zheng Zhang
- **EMNLP 2022 Findings**

Uncertain local-to-global networks for document-level event factuality identification
- Pengfei Cao, Yubo Chen, Yuqing Yang, Kang Liu, Jun Zhao
- **EMNLP 2021**

Experience

SystemsResearch@Google *Sunnyvale, USA*
Mentor: Jiani Zhang *May 2026 - August 2026*

- Worked on tool-centered evolution for data-lake analytics agents.

Amazon Bedrock Core Science *Santa Clara, USA*
Mentor: Qi Zhu *May 2025 - August 2025*

- Enhanced LLM robustness for table tasks by identifying failure modes and optimizing data referencing accuracy.

GAIR Lab, Shanghai Jiao Tong University *Shanghai, China*
Advisor: Professor Pengfei Liu *June 2023 - June 2024*

- Conducted research on LLM alignment, focusing on improving knowledge manipulation and reasoning capabilities.

Amazon AWS AI *Shanghai, China*
Mentors: Qipeng Guo, Professor Zheng Zhang, Professor Yue Zhang *July 2021 - June 2023*

- Worked on information extraction, specializing in relation extraction and event extraction.

Institute of Automation, Chinese Academy of Sciences *Beijing, China*
Advisors: Professor Yubo Chen, Professor Jun Zhao *September 2020 - June 2021*

- Worked on information extraction, with a primary emphasis on event factuality prediction.

Awards

- **Annenberg Fellowship**, University of Southern California, 2024.
- **Outstanding Graduate of Fudan University**, Fudan University, 2024.
- **Xiaomi Scholarship**, Fudan University, 2023.
- **Outstanding Graduate of Beijing**, University of Chinese Academy of Sciences, 2021.
- **Outstanding Graduate of UCAS**, University of Chinese Academy of Sciences, 2021.

Skills

LLM Stack	SFT, DPO, RLVR, Inference, Interpretability
English Proficiency	TOEFL: 108
Teaching Assistant	Probability Theory and Mathematical Statistics (Fudan, Fall 2022)