

# Yuxiao Qu

608-335-1010 | [yuxiaoq@andrew.cmu.edu](mailto:yuxiaoq@andrew.cmu.edu) | <https://cohenqu.github.io/>

## Education

---

### Carnegie Mellon University

PhD in Machine Learning  
Advisor: Prof. Aviral Kumar & Prof. Ruslan Salakhutdinov

May 2024 - Present  
Pittsburgh, PA

### Carnegie Mellon University

Master of Science in Machine Learning  
Advisor: Prof. Aviral Kumar  
GPA: 4.28/4.33

Aug 2023 - May 2024  
Pittsburgh, PA

### University of Wisconsin - Madison

Bachelor of Science in Computer Science and Mathematics  
Advisor: Prof. Josiah P. Hanna  
GPA: 4.00/4.00

Sep 2020 – Dec 2022  
Madison, WI

### The Chinese University of Hong Kong

Major: Computer Science

Sep 2017 – May 2020  
Hong Kong

## Honors and Scholarships

---

- Best paper award in 2021 IEEE/ACM Asian and South Pacific Design Automation Conference
- 2021 ICPC North Central NA Regional Contest (15th/96)
- 2020 Professor Charles K. Kao Research Exchange Scholarship
- Best project award in 2019 Undergraduate Summer Research Internship

## Publications

---

(\*indicates Equal Contribution)

### 1. Optimizing Test-Time Compute via Meta Reinforcement Finetuning

Yuxiao Qu, Matthew Y. R. Yang, Amrith Setlur, Lewis Tunstall, Edward Emanuel Beeching, Ruslan Salakhutdinov, Aviral Kumar  
Preprint

### 2. Harnessing Webpage UIs for Text-Rich Visual Understanding

Junpeng Liu, Tianyue Ou\*, Yifan Song\*, Yuxiao Qu\*, Wai Lam, Chenyan Xiong, Wenhui Chen, Graham Neubig, Xiang Yue  
The Thirteenth International Conference on Learning Representations (ICLR), 2025

### 3. Recursive Introspection: Teaching Language Model Agents How to Self-Improve

Yuxiao Qu, Tianjun Zhang, Naman Garg, Aviral Kumar  
The Thirty-Eighth Annual Conference on Neural Information Processing Systems (NeurIPS), 2024

### 4. Guided Data Augmentation for Offline Reinforcement Learning and Imitation Learning

Nicholas E. Corrado, Yuxiao Qu, John U. Balis, Adam Labiosa, Josiah P. Hanna  
The first Reinforcement Learning Conference (RLC), 2024

### 5. FLEE-GNN: A Federated Learning System for Edge-Enhanced Graph Neural Network in Analyzing Geospatial Resilience of Multicommodity Food Flows

Yuxiao Qu, Jinmeng Rao, Song Gao, Qianheng Zhang, Wei-Lun Chao, Yu Su, Michelle Miller, Alfonso Morales, Patrick Huber

ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery, 2023

**6. Simulation-Acquired Latent Action Spaces for Dynamics Generalization**

Nicholas Corrado, **Yuxiao Qu**, Josiah P. Hanna

Proceedings of Machine Learning Research (PMLR), 2022

**7. Polite or Direct? Conversation Design of a Smart Display for Older Adults Based on Politeness Theory**

Yaxin Hu, **Yuxiao Qu**, Adam Maus, Bilge Mutlu

ACM CHI Conference on Human Factors in Computing Systems (CHI), 2022

**8. TreeNet: Deep Point Cloud Embedding for Routing Tree Construction**

Wei Li, **Yuxiao Qu**, Gengjie Chen, Yuzhe Ma, Bei Yu

ACM Asian and South Pacific Design Automation Conference (ASPDAC), 2021 (Best Paper Award)

---

## Invited Talks

**Recursive Introspection: Teaching Language Model Agents How to Self-Improve**

ICML Workshop on Large Language Models and Cognition

July 2024

---

## Teaching Experience

**Undergraduate Teaching Assistant**

University of Wisconsin – Madison

COMP SCI 577: Introduction to Algorithms (Honor Session)

Fall 2022

COMP SCI 537: Introduction to Operating Systems

Fall 2022

---

## Work Experience

**Morgridge Institute for Research**

Apr 2023 – Aug 2023

*Research Software Engineer*

Madison, WI

---

## Professional Service

Reviewer for ICML (2024), NeurIPS (2024), AISTATS (2025), ICLR (2025)