

Yuchen Zhuang

(+1) 404-751-7641 | yczhuang@gatech.edu | <https://night-chen.github.io/> | 756 W Peachtree St NW, Atlanta, GA 30308

EDUCATION

Computer Science and Engineering, Georgia Institute of Technology

Machine Learning Ph.D., ECE Master of Science

Ph.D. Advisor: [Dr. Chao Zhang](#)

Research Interest: Text Mining; Applied Machine Learning; Information Extraction

Atlanta, Georgia

Aug 2019-May 2025 (expected)

School of Information Science and Engineering, Southeast University (SEU)

Bachelor of Engineering in Information Technology

Bachelor Advisor: [Dr. Chuan Zhang](#)

Research Interest: Mathematical Modeling in Multi-Scale Communication; Molecular Computing

Nanjing, China

Sept 2015-July 2019

INDUSTRY EXPERIENCE

Rufus Group, Amazon, Palo Alto, California

Applied Scientist Intern, Host: Jingfeng Yang, Haoming Jiang, Qing Ping, Bing Yin

Publication: Improving Fundamental Agent Capabilities of LLMs through Continual Pre-Training (NAACL'25)

May 2024-Aug 2024

Data-Driven Media Systems Group, Adobe, San Jose, California

Research Scientist Intern, Manager: Saayan Mitra

Mentor: Xiang Chen, Tong Yu, Victor Soares Bursztyn, Ryan Rossi, Somdeb Sarkhel

Publication: ToolChain*: Efficient Action Space Navigation in Large Language Models with A* Search (ICLR'24)

May 2023-Aug 2023

Personalization Group, Amazon, Seattle, Washington

Applied Scientist Intern, Manager: Tong Zhao, Xin Shen, Mentor: Yan Zhao, Chaosheng Dong, Ming Wang

Publication: G-STO: Sequential Main Shopping Intention Detection via Graph-Regularized Stochastic Transformer (CIKM'23)

May 2022-Aug 2022

SELECTED PUBLICATIONS

(The full publication list can be found [this link](#). * denotes equal contributions)

- Yuchen Zhuang**, Jingfeng Yang, Haoming Jiang, Xin Liu, Kewei Cheng, Sanket Lokegaonkar, Yifan Gao, Qing Ping, Tianyi Liu, Binxuan Huang, Zheng Li, Zhengyang Wang, Pei Chen, Ruijie Wang, Rongzhi Zhang, Nasser Zalmout, Priyanka Nigam, Bing Yin, Chao Zhang. [Hephaestus: Improving Fundamental Agent Capabilities of Large Language Models through Continual Pre-Training](#). In Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL 2025).
- Haorui Wang, Marta Skreta, Cher Tian Ser, Wenhao Gao, Lingkai Kong, Felix Strieth-Kalthoff, Chenru Duan, **Yuchen Zhuang**, Yue Yu, Yanqiao Zhu, Yuanqi Du, Alan Aspuru-Guzik, Kirill Neklyudov, Chao Zhang. [Efficient Evolutionary Search Over Chemical Space with Large Language Models](#). In The Thirteenth International Conference on Learning Representations (ICLR 2025).
- Yuchen Zhuang**, Haotian Sun, Yue Yu, Qifan Wang, Chao Zhang, Bo Dai. [HYDRA: Model Factorization Framework for Black-Box LLM Personalization](#). In Thirty-eighth Conference on Neural Information Processing Systems (NeurIPS 2024).
- Lingkai Kong, Haorui Wang, Wenhao Mu, Yuanqi Du, **Yuchen Zhuang**, Yifei Zhou, Yue Song, Rongzhi Zhang, Kai Wang, Chao Zhang. [Aligning Large Language Models with Representation Editing: A Control Perspective](#). In Thirty-eighth Conference on Neural Information Processing Systems (NeurIPS 2024).
- Haotian Sun*, **Yuchen Zhuang***, Wei Wei, Chao Zhang, Bo Dai. [BBox-Adapter: Lightweight Adapting for Black-Box Large Language Models](#). In The Forty-first International Conference on Machine Learning (ICML 2024), (Spotlight).
- Yuchen Zhuang**, Xiang Chen, Tong Yu, Saayan Mitra, Victor Bursztyn, Ryan A. Rossi, Somdeb Sarkhel, Chao Zhang. [ToolChain*: Efficient Action Space Navigation in Large Language Models with A* Search](#). In The Twelfth International Conference on Learning Representations (ICLR 2024).
- Yue Yu*, **Yuchen Zhuang***, Jieyu Zhang*, Yu Meng, Alexander Ratner, Ranjay Krishna, Jiaming Shen, Chao Zhang. [Large Language Model as Attributed Training Data Generator: A Tale of Diversity and Bias](#). In Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS 2023).
- Yuchen Zhuang**, Yue Yu, Kuan Wang, Haotian Sun, Chao Zhang. [ToolQA: A Dataset for LLM Question Answering with External Tools](#). In Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS 2023).
- Haotian Sun*, **Yuchen Zhuang***, Lingkai Kong, Bo Dai, Chao Zhang. [AdaPlanner: Adaptive Planning from Feedback with Language Models](#). In Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS 2023).

10. Wenqi Shi, **Yuchen Zhuang**, Yuanda Zhu, Henry J. Iwinski, J. Michael Wattenbarger, May D. Wang. Retrieval-Augmented Large Language Models for Adolescent Idiopathic Scoliosis Patients in Shared Decision-Making. In *ACM Conference on Bioinformatics, Computational Biology, and Health Informatics (BCB 2023)*. **(Best Paper Award)**.
11. **Yuchen Zhuang**, Xin Shen, Yan Zhao, Chaosheng Dong, Ming Wang, Jin Li, Chao Zhang. [G-STO: Sequential Main Shopping Intention Detection via Graph-Regularized Stochastic Transformer](#). In proceedings of *ACM International Conference on Information and Knowledge Management (CIKM 2023)*.
12. **Yuchen Zhuang**, Yue Yu, Lingkai Kong, Xiang Chen, Chao Zhang. [DyGen: Learning from Noisy Labels via Dynamics-Enhanced Generative Modeling](#). In proceedings of *29TH ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2023)*.
13. Yue Yu, **Yuchen Zhuang**, Rongzhi Zhang, Yu Meng, Jiaming Shen, Chao Zhang. [ReGen: Zero-Shot Text Classification via Training Data Generation with Progressive Dense Retrieval](#). In proceedings of *Findings of The 61st Annual Meeting of the Association for Computational Linguistics (ACL-Findings 2023)*.
14. Lingkai Kong, Jiaming Cui, Haotian Sun, **Yuchen Zhuang**, B Aditya Prakash, Chao Zhang. [Autoregressive Diffusion Model for Graph Generation](#). In proceedings of *The Fortieth International Conference on Machine Learning (ICML 2023)*.
15. **Yuchen Zhuang**, Yinghao Li, Junyang Zhang, Yue Yu, Yingjun Mou, Xiang Chen, Le Song, Chao Zhang. [ReSel: N-ary Relation Extraction from Scientific Text and Tables by Learning to Retrieve and Select](#). In proceedings of *The 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022)*.
16. Lingkai Kong, Jiaming Cui, **Yuchen Zhuang**, Rui Feng, B Aditya Prakash, Chao Zhang. [End-To-End Stochastic Programming with Energy-based Model](#). In proceedings of *Advances in Neural Information Processing Systems (NeurIPS 2022)*, **(Oral)**.
17. Lingkai Kong, Haoming Jiang, **Yuchen Zhuang**, Jie Lyu, Tuo Zhao, Chao Zhang. [Calibrated Language Model Fine-tuning for In- and Out-of-Distribution Data](#). In proceedings of *The 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)*.

PREPRINTS & MANUSCRIPTS

1. Changhao Li*, **Yuchen Zhuang***, Rushi Qiang, Haotian Sun, Hanjun Dai, Chao Zhang, Bo Dai. [Matryoshka: Learning to Drive Black-Box LLMs with LLMs](#).

HONORS AND AWARDS (SELECTED)

- JP Morgan PhD Fellowship 2024 *July, 2024*
- NeurIPS 2023 Scholar Award *Oct. 2023*
- KDD Student Travel Award, ACM SIGKDD *May. 2023*
- Best Paper Award, ACM Conference on Bioinformatics, Computational Biology, and Health Informatics. *Sep. 2023*
- Second Prize of Excellent Undergraduate Student Graduation Thesis in Jiangsu Province. *Jun, 2020*
- Most Influential Graduate Award Nomination **(20/4000)**, Southeast University *Jun, 2019*
- Advanced Individual in Scientific Research, Southeast University *Jun. 2019*
- Qingyun Sun Innovation Scholarship, Southeast University *May2018, May 2019*
- International Collaboration Symposium on Information, Production & Systems **Excellent Paper Award** *Dec, 2017*
- First Prize, National High School Mathematical League *Feb 2014*
- First Prize, National Olympiad in Informatics in Provinces *Feb 2014*

SERVICES

- Teaching Assistant, CS 7641 Machine Learning* *Fall 2020*
- Teaching Assistant, CSE 8803 Deep Learning for Text Data* *Fall 2021*
- Teaching Assistant, GT Big Data Bootcamp* *Spring 2021,2022,2023*
- Teaching Assistant, GT Natural Language Processing Bootcamp* *Spring 2023*
- Conference Area Chair, ACL 2024, EMNLP 2024, NAACL 2025.*
- Conference Program Committee, ARR (2023-), EMNLP (2022-), ICLR (2024-), NeurIPS (2023-), KDD (2021-), ACL (2023-), AAI (2023-), ICML (2021-), AISTATS (2024), SDM (2024).*