

Yusen Zhang

450 Computer Science Bldg.
500 W 120th St, New York, NY 10027
yz5296@columbia.edu
www.yuszh.com

RESEARCH INTERESTS	Scalable, efficient, trustworthy, collaborative large language model systems. Multi-Agent LLMs. Human-centric large language model evaluation and applications.
EDUCATION	<i>Ph.D. Computer Science and Engineering</i> 2021 - 2025 Penn State University, State College, PA Advisor: Dr. Rui Zhang
	<i>M.Sc. Computer Science (Data Science Track)</i> 2019 - 2021 Emory University, Atlanta, GA Advisor: Dr. Jinho D. Choi
RESEARCH EXPERIENCE	<i>Postdoctoral Research Scientist</i> Oct. 2025 - Present Columbia University, New York, NY Project: Agent and Multiagent Large Language Models Advisor: DAPLab Faculty (Dr. Zhou Yu, Dr. Eugene Wu)
	<i>Student Researcher</i> Aug. 2023 - May 2024 Google, Sunnyvale, CA Project: Multi-agent Large Language Models Advisor: Dr. Ruoxi Sun, Dr. Sercan Arik
	<i>Applied Scientist Intern</i> Jun. 2022 - Aug. 2022 Amazon, New York, NY Project: Long Dialogue Summarization Advisor: Dr. Xiaofei Ma, Dr. Wei Xiao
	<i>Research Intern</i> Feb. 2022 - Jun. 2022 Microsoft, Redmond, WA Project: Controllable Summarization Advisor: Dr. Yang Liu
	<i>Research Assistant</i> Jan. 2021 - Dec. 2021 Yale University, New Haven, CT Project: Long Text Summarization Advisor: Dr. Dragomir R. Radev
	<i>Research Intern</i> Aug. 2017 - Jun. 2018 Microsoft Research Asia, Beijing, China Project: Intent Classification and Name Entity Recognition Advisor: Dr. Zijia Lin

PUBLICATIONS & PREPRINTS * indicates equal contribution (**citation: 1214, h-index: 16**)

- [P6]. Xi Li, Shu Zhao, Xiaohan Zou, Fei Zhao, Fuxiao Liu, **Yusen Zhang**, Cheng Han, Yushun Dong, Jiaqi Wang. “Evolving Safety Landscape of Multi-modal Large Language Models: A Survey of Emerging Threats and Safeguards”
- [P5]. Nan Zhang, Eugene Kwek, **Yusen Zhang**, Muyu Pan, Suhang Wang, Prasenjit Mitra, Rui Zhang. “QuantLRM: Quantization of Large Reasoning Models via Fine-Tuning Signals”
- [P4]. Jiaqi Wang, **Yusen Zhang**, Xi Li, Simeng Han, Zining Zhu, Lingjuan Lyu, Yaqing Wang, Ruiquan Huang, Arman Cohan, Rui Zhang, Sheng Li, Xin (Eric) Wang. “Survey on Neural Network Parameter Acquisition: from Optimization to Generation”
- [P3]. Simeng Han, Tianyu Liu, Chuhan Li, Xuyuan Xiong, **Yusen Zhang**, Arman Cohan. “HybridMind: Meta Selection of Natural Language and Symbolic Language for Enhanced LLM Reasoning”
- [P2]. Ryo Kamoi, **Yusen Zhang**, Nan Zhang, Sarkar Snigdha Sarathi Das, Rui Zhang. “Training Step-Level Reasoning Verifiers with Formal Verification Tools”.
- [P1]. Jiaqi Wang*, **Yusen Zhang***, Xi Li*. “NeuroGen: Neural Network Parameter Generation via Large Language Models”.
- [C24]. Nan Zhang, **Yusen Zhang**, Prasenjit Mitra, Rui Zhang. “When Reasoning Meets Compression: Benchmarking Compressed Large Reasoning Models on Complex Reasoning Tasks”. ICLR 2026.
- [C23]. **Yusen Zhang**, Sarkar Das, Rui Zhang. “Verbosity \neq Veracity: Verbosity Compensation Behavior of Large Language Models”. UncertainNLP @ EMNLP 2025.
- [C22]. Ryo Kamoi, **Yusen Zhang**, Sarkar Snigdha Sarathi Das, Ranran Haoran Zhang, Rui Zhang. “VisOnlyQA: Large Vision Language Models Still Struggle with Visual Perception of Geometric Information”. COLM 2025.
- [C21]. **Yusen Zhang**, Wenliang Zheng, Aashrith Madasu, Peng Shi, Ryo Kamoi, Hao Zhou, Zhuoyang Zou, Shu Zhao, Sarkar Snigdha Sarathi Das, Vipul Gupta, Xiaoxin Lu, Nan Zhang, Ranran Haoran Zhang, Avitej Iyer, Renze Lou, Wenpeng Yin, Rui Zhang. “HRScene: How Far Are VLMs from Effective High-Resolution Image Understanding?”. ICCV 2025.
- [C20]. Wenliang Zheng, Sarkar Snigdha Sarathi Das, **Yusen Zhang**, Rui Zhang. “GREATERPROMPT: A Unified, Customizable, and High-Performing Open-Source Toolkit for Prompt Optimization”. ACL 2025 Demo Track.
- [C19]. Xiaoxin Lu, Ranran Haoran Zhang, **Yusen Zhang**, Rui Zhang. “Enhance Multimodal Consistency and Coherence for Text-Image Plan Generation”. ACL 2025 Findings.
- [C18]. Xi Li, **Yusen Zhang**, Renze Lou, Chen Wu, Jiaqi Wang. “Chain-of-Scrutiny: Detecting Backdoor Attacks for Large Language Models”. ACL 2025 Findings.
- [C17]. Renze Lou, Hanzi Xu, Sijia Wang, Jiangshu Du, Ryo Kamoi, Xiaoxin Lu, Jian Xie, Yuxuan Sun, **Yusen Zhang**, Jihyun Janice Ahn, Hongchao Fang, Zhuoyang Zou, Wenchao Ma, Xi Li, Kai Zhang, Congying Xia, Lifu Huang, Wenpeng Yin. “AAAR-1.0: Assessing AI’s Potential to Assist Research”. ICML 2025, **Best Paper Award at the 2nd AI4Research workshop @ AAAI 2025**.

- [C16]. Sarkar Snigdha Sarathi Das, Ryo Kamoi, Bo Pang, **Yusen Zhang**, Caiming Xiong, Rui Zhang. “GReaTer: Gradients over Reasoning Makes Smaller Language Models Strong Prompt Optimizers”. ICLR 2025.
- [C15]. Haoyuan Li, **Yusen Zhang**, Rui Zhang, Snigdha Chaturvedi. “Coverage-based Fairness in Multi-document Summarization”. NAACL 2025, **SAC Award of Summarization**.
- [C14]. **Yusen Zhang**, Ruoxi Sun, Yanfei Chen, Tomas Pfister, Rui Zhang*, Sercan Ö. Arik*. “Chain of Agents: Large Language Models Collaborating on Long-Context Tasks”. NeurIPS 2024.
- [C13]. Ryo Kamoi, **Yusen Zhang**, Nan Zhang, Jiawei Han, Rui Zhang. “When Can LLMs Actually Correct Their Own Mistakes? A Critical Survey of Self-Correction of LLMs”. TACL 2024. **Oral at EMNLP 2024**.
- [C12]. Ryo Kamoi, Sarkar Snigdha Sarathi Das, Renze Lou, Jihyun Janice Ahn, Yilun Zhao, Xiaoxin Lu, Nan Zhang, **Yusen Zhang**, Ranran Haoran Zhang, Sujeeth Reddy Vummanthala, Salika Dave, Shaobo Qin, Arman Cohan, Wenpeng Yin, Rui Zhang. “Evaluating LLMs at Detecting Errors in LLM Responses”. COLM 2024.
- [C11]. **Yusen Zhang**, Nan Zhang, Yixin Liu, Alexander Fabbri, Junru Liu, Ryo Kamoi, Xiaoxin Lu, Caiming Xiong, Jieyu Zhao, Dragomir Radev, Kathleen McKeown, Rui Zhang. “Fair Abstractive Summarization of Diverse Perspectives”. NAACL 2024.
- [C10]. Nan Zhang, **Yusen Zhang**, Wu Guo, Prasenjit Mitra, Rui Zhang. “FaMe-Summ: Investigating and Improving Faithfulness of Medical Summarization”. EMNLP 2023.
- [C9]. **Yusen Zhang**, Jun Wang, Zhiguo Wang, Rui Zhang. “XSemPLR: Benchmarking Cross-Lingual Semantic Parsing in Multiple Natural Languages and Meaning Representations”. ACL 2023.
- [C8]. **Yusen Zhang**, Yang Liu, Ziyi Yang, Yuwei Fang, Yulong Chen, Dragomir Radev, Chenguang Zhu, Michael Zeng, Rui Zhang. “MACSum: Controllable Summarization with Mixed Attributes”. TACL 2023. **Oral at ACL 2023**.
- [C7]. **Yusen Zhang**, Ansong Ni, Ziming Mao, Chen Henry Wu, Chenguang Zhu, Budhaditya Deb, Ahmed Hassan Awadallah, Dragomir R. Radev, Rui Zhang. “Summ^N: A Multi-Stage Summarization Framework for Long Input Dialogues and Documents”. ACL 2022.
- [C6]. Ziming Mao*, Chen Henry Wu*, Ansong Ni, **Yusen Zhang**, Rui Zhang, Tao Yu, Budhaditya Deb, Chenguang Zhu, Ahmed Hassan Awadallah, Dragomir R. Radev. “DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization”. ACL 2022.
- [C5]. Ansong Ni, Zhangir Azerbayev, Mutethia Mutuma, Troy Feng, **Yusen Zhang**, Tao Yu, Ahmed Hassan Awadallah, Dragomir Radev. “SummerTime: Text Summarization Toolkit for Non-experts”. EMNLP 2021 Demo Track.
- [C4]. **Yusen Zhang***, Ansong Ni*, Tao Yu, Rui Zhang, Chenguang Zhu, Budhaditya Deb, Asli Celikyilmaz, Ahmed Hassan Awadallah, Dragomir Radev. “An Exploratory Study on Long Dialogue Summarization: What Works and What’s Next”. EMNLP 2021 Findings.
- [C3]. Chang Shu*, **Yusen Zhang***, Xiangyu Dong, Peng Shi, Tao Yu, Rui Zhang. “Logic-Consistency Text Generation from Semantic Parses”. ACL 2021 Findings.

- [C2]. **Yusen Zhang***, Xiangyu Dong*, Shuaichen Chang, Tao Yu, Peng Shi, Rui Zhang. “Did You Ask a Good Question? A Cross-Domain Question Intention Classification Benchmark for Text-to-SQL”. Intex-Sempar @ EMNLP 2020.
- [C1]. Hui Chen, Zijia Lin, Guiguang Ding, Jianguang Lou, **Yusen Zhang**, Borje Karlsson. “GRN: Gated Relation Network to Enhance Convolutional Neural Network for Named Entity Recognition”. AAAI 2019

HONORS AND AWARDS	SAC Award, NAACL 2025, ACL Committee	2025
	Best Paper Award, the 2nd AI4Research workshop at AAAI 2025	2025
	Outstanding Reviewer, EMNLP 2024, ACL Committee	2024
	Dr. Tse-Yun Feng Graduate Student Award, CSE Dept., Penn State University	2022
	The 28 th place (19 teams to Final), ACM-ICPC North America Champions	2020
	The 3 rd place, ACM-ICPC Southeast American Regional Contest	2019
	“Star of Tomorrow” Excellence Award, Microsoft Research Asia	2018

SERVICES	Workshop Organizer (Co-Chair), WMAC 2026 @ AAAI-26
	Review committee, AAAI 2026, ICML 2025, ICLR 2025, EMNLP 2022,2023,2024,2025, ACL 2023,2025 CoNLL 2023,2024,2025 Neurips 2024,2025, NAACL 2025, ACL Rolling Review Nov. 2021
	Program Committee of InTex-SemPar @ EMNLP 2020, AI4Research @ AAAI 2025, NewSumm @ EMNLP 2025
	Student Volunteer @ NAACL 2024

TALKS AND PRESENTATIONS	Invited Lecture. “Multiagent System in the Era of LLMs: Testbeds, Analyses, Methods“. Auburn University, Nov. 2025
	Invited talk. “HRScene: How Far Are VLMs from Effective High-Resolution Image Understanding?”. NICE Academic Group, July 2025.
	Invited talk. “Multi-Agent Systems in the Era of LLM: Foundation, Testbeds, and Beyond“. Columbia University, July 2025.
	Invited talk and poster. “MACSum: Controllable Summarization with Mixed Attributes“. MASC-SLL. 2023 Apr. 2023.
	Invited talk. “MACSum: Controllable Summarization with Mixed Attributes“. Westlake University NLP Group. Feb. 2023.

MENTORED STUDENTS	Tianle Zhou Position: Columbia Undergrad
	Wenliang Zheng (ICCV 2025, ACL 2025 Demo) Position: PSU Undergrad → Research Intern @ Microsoft Research
	Aashrith Madasu (ICCV 2025) Position: PSU Master

PROFESSIONAL REFERENCE Rui Zhang, Ph.D.
Assistant Professor,
Computer Science and Engineering Department,
Penn State University.
Address: W329 Westgate Building, University Park, PA 16802
Email: r.zhang@psu.edu
Relationship: Ph.D. Advisor