FANGCONG YIN

fangcongyin@utexas.edu
https://fangcong-yin-2.github.io

EDUCATION

New York University 2026 - 2028 (Expected)
Ph.D. in Computer Science

The University of Texas at Austin 2023 - 2025 (Transferred to NYU)
Ph.D. in Computer Science GPA: *3.88*

Cornell University 2021 - 2023
Bachelor of Science in Information Science, System, and Technology
(Summa Cum Laude; Merrill Presidential Scholar) GPA: *4.03*

University of Notre Dame 2019 - 2021
Bachelor of Science in Computer Science GPA: *3.97*

PUBLICATIONS

1. **Fangcong Yin**, Zeyu Leo Liu, Liu Leqi, Xi Ye, and Greg Durrett. 2025. Learning Composable Chains-of-Thought. *Foundations of Reasoning in Language Models Workshop at the 39th Conference on Advances in Neural Information Processing Systems (FoRLM @ NeurIPS 2025).*

2. Liyan Tang, Grace Kim, Xinyu Zhao, Thom Lake, Wenxuan Ding, **Fangcong Yin**, Prasann Singhal, Manya Wadhwa, Zeyu Leo Liu, Zayne Sprague, Ramya Namuduri, Bodun Hu, Juan Diego Rodriguez, Puyuan Peng, and Greg Durrett. 2025. ChartMuseum: Testing Visual Reasoning Capabilities of Large Vision-Language Models. *Proceedings of the 39th Conference on Advances in Neural Information Processing Systems (NeurIPS 2025).*

3. Femi Bello, Anubrata Das, Fanzhi Zeng, **Fangcong Yin**, Liu Leqi. 2025. Linear Representation Transferability Hypothesis: Leveraging Small Models to Steer Large Models. *Preprint.*

4. Wuwei Zhang, **Fangcong Yin**, Howard Yen, Danqi Chen, and Xi Ye. 2025. Query-Focused Retrieval Heads Improve Long-Context Reasoning and Re-ranking. *Proceedings of the 2025 Conference on Empirical Methods for Natural Language Processing (EMNLP 2025).*

5. Xi Ye, **Fangcong Yin**, Yinghui He, Joie Zhang, Howard Yen, Tianyu Gao, Greg Durrett, and Danqi Chen. 2025. LongProc: Benchmarking Long-Context Language Models on Long Procedural Generation. *Second Conference on Language Modeling (COLM 2025).*

6. Xinyu Zhao, **Fangcong Yin**, and Greg Durrett. 2025. Understanding Synthetic Context Extension via Retrieval Heads. *Forty-second International Conference on Machine Learning (ICML 2025).*

7. Zayne Sprague, **Fangcong Yin**, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. 2025. To CoT or Not to CoT? Chain-of-thought Helps Mainly on Math and Symbolic Reasoning. *Proceedings of the 13th International Conference on Learning Representations (ICLR 2025).*

8. **Fangcong Yin**, Xi Ye, and Greg Durrett. 2024. LoFiT: Localized Fine-tuning on LLM Representations. *Proceedings of the 38th Conference on Advances in Neural Information Processing Systems (NeurIPS 2024).*

9. **Fangcong Yin** and Marten van Schijndel. 2023. Linguistic Compression in Single-Sentence Human-Written Summaries. *Findings of the 2023 Conference on Empirical Methods for Natural Language Processing (EMNLP 2023-Findings).*

RESEARCH EXPERIENCE

Applied Scientist Intern, Amazon Web Services New York, NY
Efficient Tool-Use Agent for Chart Reasoning May 2025 – Sep 2025

- Develop efficient tool-use agent systems for reasoning on charts and visualizations that achieve state-of-the-art accuracy on chart reasoning benchmarks with reduced token costs

TAUR Lab, The University of Texas at Austin (Advisor: Greg Durrett) Austin, TX
Efficient LLM Adaptation via Localization Aug 2023 – Present

- Develop localized, parameter-efficient fine-tuning methods that achieve comparable performance to LoRA with 200x fewer parameters **(Paper accepted to NeurIPS 2024)**
- Investigate efficient adaptations of large LMs by steering their representations using smaller LMs (Paper submitted to NeurIPS 2025)

Limits of Chain-of-Thought Reasoning May 2024 – Present

- Conduct controlled experiments to compare the performance of 14 LLMs between chain-of-thought prompting and tool augmentation to investigate the limits of chain-of-thought on 20 reasoning tasks (**Paper accepted to ICLR 2025**)

- Improve the compositional generalization of LLMs by augmenting fine-tuning data with composable chain-of-thought traces (Paper submitted to NeurIPS 2025)

Interpretable Long-Context LLM Reasoning July 2024 – Present

- Benchmark the limits of frontier long-context LLMs in generating long-form outputs of procedures (**Paper accepted to COLM 2025**)
- Utilize interpretability methods to predict the transferability of fine-tuning with synthetic data for long-context tasks via attention heads (**Paper accepted to ICML 2025**)
- Improve long-context reasoning and re-ranking by leveraging query-focused attention heads (**Paper accepted to EMNLP 2025**)

Cornell NLP Group, Cornell University Ithaca, NY

Linguistic Compression in Summarization (Advisor: Marten van Schijndel)Apr 2022 – June 2023

- Probe LLMs to explore the difference between human-written and model-generated summaries in terms of linguistic compression **(Paper accepted to EMNLP 23-Findings)**

Movie Summarization Benchmark (Advisor: Claire Cardie) Sep 2021 – June 2023

- Create a multi-reference scene-to-scene fine-grained movie summarization dataset for long-form summarization

Human Language Technology Center of Excellence, Johns Hopkins University Remote

Visiting Researcher May 2022 – August 2022

- Attend Summer Camp for Applied Language Exploration (SCALE) 2022 workshop on authorship identification

Data Mining Towards Decision Making Lab, University of Notre Dame Notre Dame, IN

Undergraduate Research Assistant (Advisor: Meng Jiang) Aug 2020 – Aug 2021

- Experiment with diverse question generation for document retrieval augmentation
- Build benchmark datasets to evaluate the diversity of natural language generation

RWTH Aachen University Aachen, Germany (Remote)

Research Assistant (Advisor: Elma Kerz & Daniel Wiechmann) May 2020 – Aug 2020

- Leverage linear mixed effects models to correlate text readability features with eye-tracking measures during human reading

GRANTS & AWARDS

- Spotlight Presentation, The Foundation Model Interventions Workshop at NeurIPS 2024.
- Honorable Mention, Computing Research Association's (CRA) Outstanding Undergraduate Researcher Award for 2023 Dec 2022

- Cornell Engineering Learning Initiatives Undergraduate Research Grant. *Linguistic Influences on Automatic Summarization Strategies.* ($ 2500) Fall 2022
- Wood Excellence Engineering Edu Research Award. *Linguistic Influences on Automatic Summarization Strategies.* ($ 3000) Summer 2022
- Cornell Engineering Learning Initiatives Undergraduate Research Grant. *MovieRecap Dataset Creation and Evaluation.* ($ 2100) Spring 2022
- College of Engineering Dean's List, Cornell University Fall 2021 – Fall 2022
- College of Engineering Dean's List, University of Notre Dame Fall 2019 – Spring 2021

SERVICES

Conference Reviewer: ACL Rolling Review (December 2024; April 2024), EMNLP 2024, ICLR 2025, COLM 2025, NeurIPS 2025
Workshop Reviewer: NeurIPS 2024 Workshop on Foundation Model Interventions, The 4th Workshop on Processing and Evaluating Event Representations (PEER2025)
PhD Admission Reviewer: The University of Texas at Austin (2023)

TEACHING EXPERIENCE

Cornell University Ithaca, NY
Teaching Assistant, Language and Information (CS 4300) Spring 2023
Teaching Assistant, Natural Language Processing (CS 4740) Fall 2022
Teaching Consultant, Object-Oriented Programming and Data Structures (CS 2110) Spring 2022

WORK EXPERIENCE

Innovation and Automation Lab, Marmon Holdings, Inc. Chicago, IL
Intern Software Developer, Innovation and Automation Lab Dec 2020 – Aug 2021

- Develop a web application that automatically extracts hand measurements from images
- Customize object detection model to detect beacon lights from real-time videos in a manufacturing plant
- Create a full-stack native mobile application to detect scan medical trackers and upload their broadcasting data to database for a medical Internet of Things project

SKILLS

Programming Languages: Python, C++, Java, JavaScript, Matlab
Frameworks: PyTorch, Scikit-Learn, Angular, React
Languages: English, Mandarin Chinese, Japanese, German