

Fangcong Yin

fy666@nyu.edu
fangcong-yin-2.github.io

EDUCATION

PhD Ph.D. in Computer Science, New York University, 2028 (Expected)
PhD Ph.D. in Computer Science, The University of Texas at Austin, *Transferred to NYU*
BS Bachelor of Science in Information Science, *Summa Cum Laude*, Cornell University, 2023
BS Bachelor of Science in Computer Science, University of Notre Dame, *Transferred to Cornell*

RESEARCH AREAS

Large language model post-training and agentic RL: Multi-turn reinforcement learning, efficient fine-tuning.

Long context language modeling: Long context reasoning, long horizon agents, agentic context management.

WORK EXPERIENCES

May 2026–Nov 2026 **Research Scientist Intern, Meta**
Created a novel multi-turn RL method to jointly learn context compaction and long-horizon agentic tasks with self-rewarded credit assignment.
Improved LLMs on long horizon agentic search and coding tasks over strong RL baselines.

May 2025–Sep 2025 **Research Scientist Intern, Amazon**
Developed efficient cost-aware tool-use agent systems for reasoning on charts and visualizations.
Achieved state-of-the-art accuracy on chart reasoning benchmarks with reduced token costs

SELECTED PUBLICATIONS

* Equal Contribution

2026 Liyan Tang*, **Fangcong Yin***, Greg Durrett. “Visually Grounded Self-Reflection for Vision-Language Models via Reinforcement Learning.” In submission to ICLR 2026.

2025 **Fangcong Yin**, Zeyu Leo Liu, Liu Leqi, Xi Ye, and Greg Durrett. “Learning Composable Chains-of-Thought.” *Findings of EMNLP 2026; Workshop on Foundations of Reasoning in Language Models (Oral)*, *NeurIPS 2025*.

2024 **Fangcong Yin**, Xi Ye, and Greg Durrett. “LoFiT: Localized Fine-tuning on LLM Representations.” *NeurIPS 2024; Workshop on Foundation Model Interventions (Oral)*, *NeurIPS 2024*.

SKILLS

Programming Languages: Python, C++, Java, JavaScript, Matlab

Frameworks: PyTorch, veRL, vLLM, Harbor, trl

Others: Reinforcement Learning, Fine-tuning, Interpretability, Efficient LLM Training and Inference

PUBLICATIONS AND PREPRINTS

* Equal Contribution † Mentoring Role

- 2026 **Liyan Tang***, **Fangcong Yin***, Greg Durrett. “Visually Grounded Self-Reflection for Vision-Language Models via Reinforcement Learning.” In submission to ICLR 2026.
- 2026 Songtao Wang, Quang Hieu Pham, **Fangcong Yin**, Xinpeng Wang, Jocelyn Qiaochu Chen, Greg Durrett, and Xi Ye. “Detecting and Suppressing Reward Hacking with Gradient Fingerprints.” *COLM 2026*.
- 2026 Manas Mehta, **Fangcong Yin**[†], Greg Durrett. “Randomized YaRN Improves Length Generalization for Long-Context Reasoning.” *Findings of EMNLP 2026*.
- 2026 Omatharv Bharat Vaidya, Connor Thomas Jerzak, Zayne Rea Sprague, **Fangcong Yin**, Nhat Ho. “When Many Answers Are Valid, Voting Fails: Symbolic Verification for Best-of-K Causal Reasoning in LLMs.” *Preprint*.
- 2026 Xi Ye*, Wuwei Zhang*, **Fangcong Yin**, Howard Yen, and Danqi Chen. “DySCO: Dynamic Attention-Scaling Decoding for Long-Context LMs.” *Preprint*.
- 2025 **Fangcong Yin**, Zeyu Leo Liu, Liu Leqi, Xi Ye, and Greg Durrett. “Learning Composable Chains-of-Thought.” *Findings of EMNLP 2026; Workshop on Foundations of Reasoning in Language Models (Oral), NeurIPS 2025*.
- 2025 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. “ChartMuseum: Testing Visual Reasoning Capabilities of Large Vision-Language Models.” *NeurIPS 2025*.
- 2025 Femi Bello, Anubrata Das, Fanzhi Zeng, **Fangcong Yin**, and Liu Leqi. “Linear Representation Transferability Hypothesis: Leveraging Small Models to Steer Large Models.” *Preprint*.
- 2025 Wuwei Zhang, **Fangcong Yin**, Howard Yen, Danqi Chen, and Xi Ye. “Query-Focused Retrieval Heads Improve Long-Context Reasoning and Re-ranking.” *EMNLP 2025*.
- 2025 Xi Ye, **Fangcong Yin***, Yinghui He*, Joie Zhang*, Howard Yen*, Tianyu Gao, Greg Durrett, and Danqi Chen. “LongProc: Benchmarking Long-Context Language Models on Long Procedural Generation.” *COLM 2025*.
- 2025 Xinyu Zhao, **Fangcong Yin**, and Greg Durrett. “Understanding Synthetic Context Extension via Retrieval Heads.” *ICML 2025*.
- 2025 Zayne Sprague, **Fangcong Yin**, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. “To CoT or Not to CoT? Chain-of-Thought Helps Mainly on Math and Symbolic Reasoning.” *ICLR 2025*.
- 2024 **Fangcong Yin**, Xi Ye, and Greg Durrett. “LoFiT: Localized Fine-tuning on LLM Representations.” *NeurIPS 2024*.
- 2023 **Fangcong Yin** and Marten van Schijndel. “Linguistic Compression in Single-Sentence Human-Written Summaries.” *Findings of EMNLP 2023*.

RESEARCH EXPERIENCES

New York University (Advisor: Greg Durrett)

- 2026– **Graduate Research Assistant**
Study post-training and reinforcement learning methods for long-horizon LLM agents, including agentic context management and credit assignment over multi-turn trajectories.

TAUR Lab, The University of Texas at Austin (Advisor: Greg Durrett)

- 2023–25 **Efficient LLM Adaptation via Localization**
Developed localized, parameter-efficient fine-tuning that matches LoRA with 200x fewer parameters (*NeurIPS 2024*).
Investigated efficient adaptation of large LMs by steering their representations using smaller LMs (*Preprint*).
- 2024–25 **Limits of Chain-of-Thought Reasoning**
Ran controlled experiments comparing chain-of-thought prompting against tool augmentation across 14 LLMs and 20 reasoning tasks (*ICLR 2025*).
Improved compositional generalization of LLMs by augmenting fine-tuning data with composable chain-of-thought traces (*Findings of EMNLP 2026*).
- 2024–25 **Interpretable Long-Context LLM Reasoning**
Benchmarked the limits of frontier long-context LLMs on long-form procedural generation (*COLM 2025*).
Used interpretability methods to predict the transferability of synthetic-data fine-tuning for long-context tasks via attention heads (*ICML 2025*).
Improved long-context reasoning and re-ranking by leveraging query-focused attention heads (*EMNLP 2025*).

Cornell NLP Group, Cornell University

- 2022–23 **Linguistic Compression in Summarization** (Advisor: Marten van Schijndel)
Probed LLMs to characterize how human-written and model-generated summaries differ in linguistic compression (*Findings of EMNLP 2023*).
- 2021–23 **Movie Summarization Benchmark** (Advisor: Claire Cardie)
Created a multi-reference, scene-to-scene fine-grained movie summarization dataset for long-form summarization.

Human Language Technology Center of Excellence, Johns Hopkins University

- 2022 **Visiting Researcher**
Participated in the Summer Camp for Applied Language Exploration (SCALE) workshop on authorship identification.

Data Mining Towards Decision Making Lab, University of Notre Dame

- 2020–21 **Undergraduate Research Assistant** (Advisor: Meng Jiang)
Experimented with diverse question generation for document retrieval augmentation and built benchmark datasets for evaluating generation diversity.

RWTH Aachen University

- 2020 **Research Assistant** (Advisors: Elma Kerz and Daniel Wiechmann)
Applied linear mixed-effects models to correlate text readability features with eye-tracking measures during human reading.

GRANTS AND AWARDS

- 2025 Spotlight Presentation, Workshop on Foundations of Reasoning in Language Models, NeurIPS 2025
- 2024 Spotlight Presentation, Workshop on Foundation Model Interventions, NeurIPS 2024
- 2023 Merrill Presidential Scholar, Cornell University

2022	Honorable Mention, Computing Research Association (CRA) Outstanding Undergraduate Researcher Award
2022	Cornell Engineering Learning Initiatives Undergraduate Research Grant (\$2,500). Linguistic Influences on Automatic Summarization Strategies.
2022	Wood Excellence in Engineering Education Research Award (\$3,000). Linguistic Influences on Automatic Summarization Strategies.
2022	Cornell Engineering Learning Initiatives Undergraduate Research Grant (\$2,100). MovieRecap Dataset Creation and Evaluation.
2021–22	College of Engineering Dean’s List, Cornell University
2019–21	College of Engineering Dean’s List, University of Notre Dame

SERVICE

Conference Reviewing

ACL Rolling Review (December 2024; April 2024; May 2026)

EMNLP 2024; ACL 2026; ICLR 2025; COLM 2025, COLM 2026; NeurIPS 2025; ICML 2026

Workshop Reviewing

Workshop on Foundation Model Interventions, NeurIPS 2024

The 4th Workshop on Processing and Evaluating Event Representations (PEER 2025)

TEACHING EXPERIENCES

The University of Texas at Austin

2025 Head Teaching Assistant, Natural Language Processing (CS 388). Fall.

Cornell University

2023 Teaching Assistant, Language and Information (CS 4300). Spring.

2022 Teaching Assistant, Natural Language Processing (CS 4740). Fall.

2022 Teaching Consultant, Object-Oriented Programming and Data Structures (CS 2110). Spring.

Updated September 2026