Yi Fang

Ph.D. Student, Department of Computer Science, Virginia Tech
yif@vt.edu | LinkedIn | Google Scholar | Homepage

RESEARCH INTERESTS

Scalable compositional intelligence: developing capable and efficient foundation models, enabling them to communicate and collaborate, and in turn using them to improve and build foundation models.

- **Diverse and Efficient Foundation Models:** Foundation models across modalities, from LLMs to graph, vision, and biological foundation models, and their efficiency with long contexts.
- **Communication and Collaboration:** Principles and mechanisms that let LLM agents communicate and collaborate, with each other and with other foundation models.
- **Foundation Models that Build Foundation Models:** LLM agents that build new foundation models, from architecture and data to training.

EDUCATION

Virginia Tech Sept 2025 – Present
Ph.D. in Computer Science
Advisor: Dr. Xuan Wang
Student member, Sanghani Center for Artificial Intelligence and Data Analytics

Wuhan University Sept 2017 – Jun 2021
B.S. in Computer Science

SELECTED PUBLICATIONS

Full list on Google Scholar

Cache2LoRA: Understanding and Improving Context Internalization
Yi Fang*, Gaurav Srivastava, Zhenyu Bi, Sirui Ding, Xuan Wang*
Under review, 2026

Joint, Decoupled, or Agent: When Each Multimodal Alignment Paradigm Wins
Yi Fang*, Gaurav Srivastava, Sirui Ding, Xuan Wang*
Under review, 2026

Toward the Science of Multi-Agent Communication: A Comprehensive Survey
Yi Fang*, Jun Hou, Priya Pitre, Gaurav Srivastava, Sirui Ding, Xuan Wang*
Under review, 2026

BioArc: Discovering Optimal Neural Architectures for Biological Foundation Models
Yi Fang*, Haoran Xu, Jiaxin Han, Sirui Ding, Yizhi Wang, Yue Wang, Xuan Wang*
International Conference on Machine Learning (**ICML 2026**)

GraphGPT-O: Synergistic Multimodal Comprehension and Generation on Graphs
Yi Fang*, Bowen Jin, Jiacheng Shen, Sirui Ding, Qiaoyu Tan, Jiawei Han*
IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR 2025**)

UniGLM: Training One Unified Language Model for Text-Attributed Graphs
Yi Fang*, Dongzhe Fan, Sirui Ding, Ninghao Liu, Qiaoyu Tan*
ACM International Conference on Web Search and Data Mining (**WSDM 2025**)

GAugLLM: Improving Graph Contrastive Learning for Text-Attributed Graphs with Large Language Models
Yi Fang*, Dongzhe Fan, Daochen Zha, Qiaoyu Tan*
ACM SIGKDD Conference on Knowledge Discovery and Data Mining (**KDD 2024**)

GRANTS AND AWARDS

NSF NAIRR Pilot Award, Project Lead 2026
Efficient Multimodal Long Context Reasoning via Parameter Generation. One-year allocation of 20,000 GPU hours on NCSA Delta.

TUTORIALS AND INVITED TALKS

Tutorial: Towards The Science of Multi-Agent Communication Dec 2026
NeurIPS 2026, Atlanta, GA (upcoming). Organizer and speaker, with Yu Zhang (Texas A&M University) and Xuan Wang (Virginia Tech).

Invited Talk: KAAI Workshop 2026 2026
Foundation Models & Agentic Methods for Astrophysical Simulation, Carnegie Mellon University.

PROFESSIONAL SERVICE

- **Area Chair:** ACL Rolling Review (ARR) 2026.
- **Conference Reviewer:** ICLR 2027; NeurIPS 2026; ICML 2026; ACL Rolling Review (ARR) 2025, 2026; ACML 2025; KDD 2024.
- **Workshop Reviewer:** ACL 2025 Student Research Workshop (SRW).
- **Journal Reviewer:** ACM Computing Surveys (CSUR).
- **Student Volunteer:** KDD 2024.

RESEARCH EXPERIENCE

Research Intern, Alibaba Sept 2024 – Dec 2024

- Improved multimodal representations for retrieval and recommendation, and deployed the model for online A/B testing in production fine-ranking.

Research Intern, UIUC (Remote) Jun 2024 – Nov 2024

- Developed GraphGPT-O for joint multimodal graph comprehension and generation; first-author paper at **CVPR 2025**.

Research Assistant, NYU Shanghai Sept 2023 – Sept 2024

- Developed language- and multimodal-based methods for graph representation learning; first-author papers at **KDD 2024** and **WSDM 2025**, and a co-authored paper at **NAACL 2025**.

SKILLS

Programming: Python.
Machine Learning: PyTorch, Hugging Face Transformers, PEFT, vLLM, DeepSpeed, PyG, LangChain.
Tools: Linux, Git, Docker, Slurm, Weights & Biases.
Languages: Mandarin (native), English.