

Xiaocong Yang

Ph.D. Student, CS @ UIUC | Founder, [AI Interpretability @ Illinois](#)

[Personal Website](#) | [Blog Posts](#) | [Google Scholar](#) | [LinkedIn](#) | xy51@illinois.edu

Education

Ph.D. in Computer Science University of Illinois Urbana-Champaign	2025.8 - 2029.5 (expected)
- Founder and current leader of <i>AI Interpretability @ Illinois</i> research group.	
M.S. in Computer Science University of Illinois Urbana-Champaign	2023.8 - 2025.5
B.S. in Economics Tsinghua University	2018.9 - 2023.6

Leadership & Teaching

Founder & Group Leader AI Interpretability @ Illinois	2025.8 - present
- Founded and lead UIUC's cross-lab research group on AI interpretability, with 4 faculty advisors, 20+ student researchers, and 5 concurrent research projects at this moment with bi-weekly advising meetings. - Secured 45k+ GPU-hours on National Center for Supercomputing Applications (NCSA) Delta through ACCESS program as the PI, and \$20k+ computing credits from Microsoft Azure and AWS. - Established external collaborations with faculties and Ph.D. students from Stanford University, University of Oxford, CU Boulder, etc.	

Instructor CS 591 Biologically Plausible AI, UIUC	Spring 2026
<i>Independently delivered 2 class sessions on AI Interpretability in the Era of LLMs: Architecture, Behavior and Beyond</i> [Section 1] [Section 2]	

Instructor CS 591 Biologically Plausible AI, UIUC	Fall 2025
<i>Independently delivered 4 class sessions on A Tutorial on Interpretable and Biologically Plausible LLM</i> [Section 1] [Section 2] [Section 3] [Section 4]	

Publication

- Xiaocong Yang.** Generative Interpretability via Scalable Neuro-Symbolic Models, **ACM AI Leadership Summit 2026** [\[link\]](#)
- Himanshu Udipi, **Xiaocong Yang**, ChengXiang Zhai. Globally Optimal Training of Spiking Neural Networks via Parameter Reconstruction, in submission to TMLR [\[link\]](#)
- Xiaocong Yang.** LLM2IR: Simple Unsupervised Contrastive Learning Makes Long-Context LLM Great Retriever. M.S. Thesis [\[link\]](#)
- Xiaocong Yang**, Jiacheng Lin, Ziqi Wang, Chengxiang Zhai. *Learning by Analogy: Enhancing Few-Shot Prompting for Math Word Problem Solving with Computational Graph-Based Retrieval*, preprint [\[link\]](#)
- Ziyi Chen, **Xiaocong Yang**, Jiacheng Lin, Chenkai Sun, Kevin Chen-Chuan Chang, Jie Huang. *Cascade Speculative Drafting for Even Faster LLM Inference*, **NeurIPS 2024** [\[link\]](#)
- Xiaocong Yang**, James Huang, Wenxuan Zhou, Muhao Chen. *Parameter-efficient Tuning with Special Token Adaptation*, **EACL 2023** [\[link\]](#)
- Xingcheng Yao, Yanan Zheng, **Xiaocong Yang**, Zhilin Yang. *NLP from Scratch Without Large-Scale Pretraining: A Simple and Efficient Framework*, **ICML 2022** [\[link\]](#)
- Leixian Shen, Enya Shen, Yuyu Luo, **Xiaocong Yang**, Xuming Hu, Xiongshuai Zhang, Zhiwei Tai, Jianmin Wang. *Towards natural language interfaces for data visualization: A survey*, **TVCG 2022** [\[link\]](#)
- Hao Zhou, Pei Ke, Zheng Zhang, Yuxian Gu, Yinhe Zheng, Chujie Zheng, Yida Wang, Chen Henry Wu, Hao Sun, **Xiaocong Yang**, Bosi Wen, Xiaoyan Zhu, Minlie Huang, Jie Tang. *Eva: An open-domain chinese dialogue system with large-scale generative pre-training*, preprint [\[link\]](#)

Research

Research Assistant TIMAN Group, University of Illinois Urbana-Champaign	2024.3 - present
Advisor: Chengxiang Zhai Donald Biggar Willett Professor, UIUC	

- Proposed and published several papers ranging from AI Interpretability, Biologically Plausible AI and Information Retrieval (publication #6 - #9).
- Co-instructed CS591 Biologically Plausible AI starting from Fall 2025 and continuing.
- Founded and led AI Interpretability @ Illinois research group under his supervision.

Research Assistant | FORWARD Data Lab, University of Illinois Urbana-Champaign **2023.8 - 2024.1**
 Advisor: [Kevin Chenchuan Chang](#) | Professor, UIUC Urbana-Champaign

- Proposed *Cascade Speculative Drafting*, where we vertically and horizontally cascade multiple draft models of different size. Our approach is mathematically proved to be more efficient than standard speculative decoding and achieves up to **81%** additional speedup in the experiments.
- Paper is published at NeurIPS 2024.

Visiting Researcher | LUKA Group, University of Southern California **2022.4 - 2022.9**
 Advisor: [Muhao Chen](#) | Assistant Professor, Information Science Institute, USC Los Angeles

- Proposed and led research on *PASTA*, a model finetuning method that is more efficient than all of existing baselines which tunes only **0.029%** of total parameters to match the performance of full finetuning.
- Paper is published at EACL 2023.

Research Assistant | IIS, Tsinghua University **2021.7 - 2023.1**
 Advisor: [Zhilin Yang](#) | Assistant Professor, Institute for Interdisciplinary Information Sciences, Tsinghua University | Beijing
 Founder & CEO, Moonshot AI

- Contributed to *Task-driven Language Modeling*, proposing a task-specific pretraining framework that reduced by **100x** the total training cost and size of corpus and democratizing the development of large LLMs.
- Analyzed the functionality of “vertical attention heads” in LLMs.
- Paper is published at ICML 2022.

Research Assistant | Conversational AI (CoAI) Group **2021.1 - 2021.6**
 Advisor: [Huang Minlie](#) | Associate Professor | Department of Computer Science and Technology, Tsinghua University Beijing

- Contributed to the development of the open-domain dialogue system *EVA* and the corresponding paper.
- Designed the interface and client-server communication framework for the deployment of the dialog system, tested rule-based and model-based methods to filter out toxic content in the generated text.

Internship

Applied Research Scientist | Pinterest **2026.5-2026.8**
 Supervisor: [Weijie \(Jenny\) Jiang](#) | Senior Machine Learning Engineer, Pinterest Core Engineering

- Proposed and implemented to apply Multi-gate Mixture-of-Experts (MMoE) architecture to the ranking model used in the recommendation system at Pinterest, and obtained up to **3.0%** improvement on multiple tasks including click and save compared to baseline model.

AI Scientist Intern | Intuit **2025.5-2025.8**
 Supervisor: [Ruocheng Guo](#) | Staff Research Scientist, Intuit AI Research Mountain View

- Raised RestBench tool-call success by **+5 percentage points** by designing tool description improvement pipeline for agentic LLM flows; scoped and led end-to-end study at Intuit.

Invited Talk

Guest Speaker | Citadel GQS PhD Colloquium **2026.4**
 New York

- Invited to represent AI Interpretability @ Illinois and give a presentation about generative interpretability of LLMs and our current work to Citadel's core quantitative research teams.