

Bowen Jin

Room 1117, Siebel Center for Computer Science, 201 N. Goodwin Ave, Urbana, IL 61801

Email: bowenj4@illinois.edu | Homepage: <https://peterjin.me> | Phone: 217-819-1796

RESEARCH INTERESTS

Large Language Models • Agents • Reinforcement Learning • Information Retrieval • Multimodality

My research centers on **LLM Post-Training**, approached from three key perspectives: **augmentation**, **alignment**, and **reasoning/agents**:

- **Augmentation** – Enhancing foundation models with external data sources to improve their capabilities [1], [2].
- **Alignment** – Developing methods to better align LLMs with human preferences and values [3].
- **Reasoning & Agents** – Unlocking the reasoning and agentic capabilities of LLMs through reinforcement learning [4], [5].

I invent **Search-R1**, one of the first open-source frameworks to enable LLMs' reasoning and agentic capability through *reinforcement learning* ([3k+](#) [Github stars](#)).

EDUCATION

University of Illinois at Urbana-Champaign 2021.08 - 2026 (expected)

Ph.D. Student in Computer Science. Advisor: Prof. Jiawei Han

Thesis: Foundation Models and External Data: Learning, Augmentation, and Reasoning.

University of Illinois at Urbana-Champaign 2021.08 - 2023.12

M.S. in Computer Science. Advisor: Prof. Jiawei Han

Thesis: Language Models as Semantic Indexers.

Tsinghua University 2017.09 - 2021.07

B.S. in Electrical Engineering & Statistics. Advisor: Prof. Yong Li

Outstanding Graduates (Top 1%).

RESEARCH AND INDUSTRY EXPERIENCE

Apple 2025.05 - 2025.09

- Research Intern. Apple Foundation Model Team. *Mentor: Dr. TJ Collins and Dr. Zirui Wang.*
- Project: Improve reasoning and agentic behavior of Apple Foundation Model.

Google 2024.05 - 2025.05

- Student Researcher. Google Cloud Research Group. *Mentor: Dr. Jinsung Yoon and Sercan O. Arik.*
- Project: RAG and Long-context LLMs (ICLR'25), LLM Alignment (ICML'25) and agentic RL for LLMs (COLM'25).

Amazon.com Inc 2023.05 - 2023.12

- Research Intern. Query Understanding Group. *Mentor: Dr. Xianfeng Tang.*
- Project: Semantic ID and Generative retrieval (ICML'24).

Microsoft 2022.05 - 2022.08

- Research Intern. Information and Data Science Group.
- Project: Dense Retrieval for Heterogeneous Data.

Microsoft 2020.09 - 2021.03

- Research Intern. Social Computing Group. *Mentors: Dr. Zheng Liu and Dr. Xing Xie.*
- Project: Knowledge-empowered News Recommendation.

HONORS AND AWARDS

2025	OpenAI's Researcher Access Program
2024	NeurIPs Top Reviewer
2024	Yunni & Maxine Pao Memorial Fellowship, The Grainger College of Engineering, UIUC
2024	Apple PhD Fellowship (20 recipients globally)
2024	WSDM 2024 NSF Student Travel Grant
2023	KDD 2023 Student Travel Grant
2021	Outstanding Graduates, Tsinghua University (Top 1%)
2021	"Star of Tomorrow" Award, Microsoft Research
2020	China National Scholarship, Tsinghua University (Top 1%)
2019	China National Scholarship, Tsinghua University (Top 1%)
2019	Honorable Mention (top 15.35%), Mathematical Contest in Modeling
2018	China National Scholarship, Tsinghua University (Top 1%)
2017	First Prize, National Olympiad in Mathematics in Provinces

SELECTED PUBLICATIONS

"*" indicates equal contribution. [Google Scholar] [Github]

First author papers:

Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan O. Arik, Dong Wang, Hamed Zamani, Jiawei Han.
Search-R1: Training LLMs to Reason and Leverage Search Engines with Reinforcement Learning.
The 2nd conference on language modeling (COLM 2025). [3k+](#) [Github stars](#)

Bowen Jin, Jinsung Yoon, Priyanka Kargupta, Sercan O. Arik, Jiawei Han.
An Empirical Study on Reinforcement Learning for Reasoning-Search Interleaved LLM Agents.
arxiv 2025.

Bowen Jin, Jinsung Yoon, Zhen Qin, Ziqi Wang, Wei Xiong, Yu Meng, Jiawei Han, Sercan O. Arik.
LLM Alignment as Retriever Optimization: An Information Retrieval Perspective.
The 42nd International Conference on Machine Learning (ICML 2025).

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

Bowen Jin, Jinsung Yoon, Jiawei Han and Sercan O. Arik.
Long-Context LLMs Meet RAG: Overcoming Challenges for Long Inputs in RAG.
The 13rd International Conference on Learning Representations (ICLR 2025).

Bowen Jin, Ziqi Pang, Bingjun Guo, Yu-Xiong Wang, Jiaxuan You, Jiawei Han.
InstructG2I: Synthesizing Images from Multimodal Attributed Graphs.
2024 Conference on Neural Information Processing Systems (NeurIPS 2024).

Bowen Jin, Gang Liu, Chi Han, Meng Jiang, Heng Ji, Jiawei Han.
Large Language Models on Graphs: A Comprehensive Survey.
IEEE Transactions on Knowledge and Data Engineering (IEEE TKDE 2024).

Bowen Jin, Chulin Xie, Jiawei Zhang, Kashob Kumar Roy, Yu Zhang, Zheng Li, Ruirui Li, Xianfeng Tang, Suhang Wang, Yu Meng, Jiawei Han.
Graph Chain-of-Thought: Augmenting Large Language Models by Reasoning on Graphs.
The 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024 Findings).

Bowen Jin, Hansi Zeng, Guoyin Wang, Xiusi Chen, Tianxin Wei, Ruirui Li, Zhengyang Wang, Zheng Li, Yang Li, Hanqing Lu, Suhang Wang, Jiawei Han, and Xianfeng Tang.
Language Models as Semantic Indexers.
The 41st International Conference on Machine Learning (ICML 2024).

Yu Zhang*, Xiusi Chen*, **Bowen Jin***, Sheng Wang, Shuiwang Ji, Wei Wang, Jiawei Han.
A Comprehensive Survey of Scientific Large Language Models and Their Applications in Scientific Discovery.
The 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024).

Kerui Zhu*, Bo-Wei Huang*, **Bowen Jin***, Yizhu Jiao, Ming Zhong, Kevin Chang, Shou-De Lin, Jiawei Han.
Investigating Instruction Tuning Large Language Models on Graphs.
The 1st conference on language modeling (COLM 2024).

Bowen Jin, Yu Zhang, Qi Zhu, and Jiawei Han.
Heterformer: Transformer-based Deep Node Representation Learning on Heterogeneous Text-Rich Networks.
The 29th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD 2023).

Bowen Jin, Wentao Zhang, Yu Zhang, Yu Meng, Xinyang Zhang, Qi Zhu, and Jiawei Han.
Patton: Language Model Pretraining on Text-rich Networks.
The 61st Annual Meeting of the Association for Computational Linguistics (ACL 2023).

Bowen Jin, Yu Zhang, Yu Meng, and Jiawei Han.

Edgeformers: Graph-Empowered Transformers for Representation Learning on Textual-Edge Networks.
The 11st International Conference on Learning Representations (ICLR 2023).

Bowen Jin, Wentao Zhang, Yu Zhang, Yu Meng, Han Zhao, and Jiawei Han.

Learning Multiplex Embeddings on Text-rich Networks with One Text Encoder.
2023 NeurIPs New Frontiers in Graph Learning Workshop (NeurIPs GLFrontiers 2023).

Bowen Jin, Chen Gao, Xiangnan He, Depeng Jin, and Yong Li.

Multi-behavior Recommendation with Graph Convolutional Networks.
The 43rd ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2020).

Other papers:

Zihao Li, Lecheng Zheng, **Bowen Jin**, Dongqi Fu, Baoyu Jing, Yikun Ban, Jingrui He, Jiawei Han
Can Graph Neural Networks Learn Language with Extremely Weak Text Supervision?
The 63rd Annual Meeting of the Association for Computational Linguistics (ACL 2025).

Yu Zhang, Yanzhen Shen, SeongKu Kang, Xiushi Chen, **Bowen Jin**, Jiawei Han.

“Why Should I Review This Paper?” Unifying Semantic, Topic, and Citation Factors for Paper-Reviewer Matching.
The 2025 ACM Web Conference (WWW 2025).

Jimeng Shi, **Bowen Jin**, Jiawei Han, Sundararaman Gopalakrishnan, Giri Narasimhan.

CoDiCast: Conditional Diffusion Model for Global Weather Forecasting with Uncertainty Quantification.
The 34th International Joint Conference on Artificial Intelligence (IJCAI 2025).

SeongKu Kang, **Bowen Jin**, Wonbin Kweon, Yu Zhang, Dongha Lee, Jiawei Han, Hwanjo Yu.

Improving Scientific Document Retrieval with Concept Coverage-based Query Set Generation.
The 18th ACM International Conference on Web Search and Data Mining (WSDM 2025).

Sizhe Zhou, Yu Meng, **Bowen Jin**, Jiawei Han.

Grasping the Essentials: Tailoring Large Language Models for Zero-Shot Relation Extraction.
The 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024).

Ran Xu, Wenqi Shi, Yue Yu, Yuchen Zhuang, **Bowen Jin**, May Dongmei Wang, Joyce C. Ho, Carl Yang

RAM-EHR: Retrieval Augmentation Meets Clinical Predictions on Electronic Health Records.
The 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024).

Tianxin Wei, **Bowen Jin**, Ruirui Li, Hansi Zeng, Zhengyang Wang, Jianhui Sun, Qingyu Yin, Hanqing Lu, Suhang Wang, Jingrui He, and Xianfeng Tang.

Towards Universal Multi-Modal Personalization: A Language Model Empowered Generative Paradigm.
The 12th International Conference on Learning Representations (ICLR 2024).

Hansi Zeng, Chen Luo, **Bowen Jin**, Sheikh Muhammad Sarwar, Tianxin Wei, Hamed Zamani.

Scalable and Effective Generative Information Retrieval.
The 2024 ACM Web Conference (WWW 2024).

SeongKu Kang, Shivam Agarwal, **Bowen Jin**, Dongha Lee, Hwanjo Yu, Jiawei Han.

Improving Retrieval in Theme-Specific Applications using a Corpus Topical Taxonomy.
The 2024 ACM Web Conference (WWW 2024).

Wei Hu, **Bowen Jin**, Minhao Jiang, Sizhe Zhou, Zhaonan Wang, Jiawei Han, Shaowen Wang.

Geospatial Topological Relation Extraction from Text with Knowledge Augmentation.
Proceedings of the 2024 SIAM International Conference on Data Mining (SDM 2024).

Zhaonan Wang, **Bowen Jin**, Wei Hu, Minhao Jiang, S Kang, Z Li, Sizhe Zhou, Jiawei Han, Shaowen Wang.

Geospatial Knowledge Hypercube.
32nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL 2024).

Yu Zhang, **Bowen Jin**, Xiushi Chen, Yanzhen Shen, Yunyi Zhang, Yu Meng, and Jiawei Han.

Weakly-supervised Multi-label Classification of Full-Text Scientific Papers.
The 29th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD 2023).

Yu Zhang, **Bowen Jin**, Qi Zhu, Yu Meng, and Jiawei Han.

The Effect of Metadata on Scientific Literature Tagging: A Cross-Field Cross-Model Study.
The 2023 ACM Web Conference (WWW 2023).

Pengcheng Jiang, Shivam Agarwal, **Bowen Jin**, Xuan Wang, Jimeng Sun, and Jiawei Han.

Text-augmented Open Knowledge Graph Completion via Pretrained Language Models.
The 61st Annual Meeting of the Association for Computational Linguistics (ACL 2023 Findings).

CONFERENCE TUTORIALS AND INVITED TALKS

Conference Tutorials

Bowen Jin, Yu Zhang, Yunyi Zhang, and Jiawei Han.

Integrating Textual and Graph Data: Advancing Knowledge Discovery with Semantic and Structural Insights.
SIAM International Conference on Data Mining (SDM 2025).

Xinze Li, Yushi Bai, **Bowen Jin**, Fengbin Zhu, Liangming Pan and Yixin Cao.

Long Context vs. RAG: Strategies for Processing Long Documents in LLMs.

The 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2025).

Bowen Jin, Yu Zhang, Sha Li, and Jiawei Han.

Bridging Text Data and Graph Data: Towards Semantics and Structure-aware Knowledge Discovery.

The 17th ACM International Conference on Web Search & Data Mining (WSDM 2024).

Invited Talks

Training Search Agents with Reinforcement Learning

Apple Foundation Model Team, Jun 2025.

Empirical Studies of Training Reasoning-Search Interleaved LLM Agents

BAAI Community, May 2025.

Training Search Agents with Reinforcement Learning

Jina.AI, Mar 2025.

Training LLMs to Reason and Leverage Search Engines with Reinforcement Learning

UIUC-NLP Seminar, Mar 2025.

Language Language Model Alignment as Information Retrieval

Google Cloud AI Research, Dec 2024.

Long-context Retrieval-augmented Generation

UIUC-NLP Seminar, Dec 2024.

Long-context LLM and Retrieval-augmented Generation

BuzzRobot Community, Nov 2024.

Multimodal Learning on Graphs

Emory University, Oct 2024.

PROFESSIONAL SERVICES

Conference Program Committee Member

ICLR 2025; WWW 2025; ICML 2025; NeurIPs 2025; EMNLP 2025;

ICLR 2024; WWW 2024; SDM 2024; ICML 2024; ACL 2024; COLM 2024; NeurIPs 2024 (**top reviewer**);

WSDM 2023; KDD 2023; NeurIPs 2023;

Journal Reviewer

IEEE Transactions on Knowledge and Data Engineering (TKDE);

ACM Transactions on Information System (TOIS);

IEEE Transactions on Big Data (TBD);

Student Volunteer

KDD 2023

TEACHING AND MENTORING

Teaching

Special Topics in NLP for Science (CSCE 689), TAMU

Spring 2025

Guest Instructor

Foundation Model (CS 396), Northwestern University

Spring 2025

Guest Instructor

Introduction to Data Mining (CS 412), UIUC

Spring 2024

Lead Teaching Assistant

Data Mining Principles (CS 512), UIUC

Fall 2023

Guest Instructor

Research Mentoring

Bingjun Guo	B.S. student at UIUC CS → M.S. student at UIUC CS (1 paper in NeurIPs)
Xianrui Zhong	M.S. student at UIUC CS (1 paper in COLM)
Yanzhen Shen	B.S. student at UIUC CS → M.S. at Stanford (1 paper in COLM)
Sizhe Zhou	M.S. student at UIUC CS → MLE at Adobe (1 paper in EMNLP)
Runchu Tian	M.S. student at UIUC CS (1 paper in submission)
Yi Fang	M.S. student at NYU → Ph.D. student at Virginia Tech (1 paper in CVPR)
Jiacheng Shen	B.S. student at NYU → Ph.D. student at NYU (1 paper in CVPR)
Kashob Kumar Roy	Ph.D. student at UIUC CS (1 paper in ACL)
Bo-Wei Huang	M.S. student at UIUC CS → Microsoft (1 paper in COLM)
Kerui Zhu	Ph.D. student at UIUC CS (1 paper in COLM)
Sayar Ghosh Roy	M.S. at UIUC CS → MLE at Apple (1 paper in COLING)
Pengcheng Jiang	M.S. at UIUC CS → Ph.D. student at UIUC CS (1 paper in ACL)
Wei Hu	Ph.D. student at UIUC GIS (1 paper in SDM)
Wentao Zhang	M.S. at UIUC CS → Ph.D. student at University of Waterloo (1 paper in ACL)