Jiacheng Lin

(217) 979-4293 • linjc1216@gmail.com • Homepage • Linkedin

Education

University of Illinois Urbana-Champaign, Department of Computer Science — Urbana, IL
Ph.D. in Computer Science — 08/2023 – 05/2027 (expected)
Tsinghua University, Department of Automation — Beijing, China
M.Sc. in Automation — 09/2020 – 06/2023
Tsinghua University, Department of Automation — Beijing, China
B.Eng. in Automation — 08/2016 – 06/2020

Research Interest

Foundation Models • Reinforcement Learning • Agentic AI • Post-Training

My research focuses on **Foundation Models** (e.g., large language models (LLMs), multi-modal models), **Reinforcement Learning** (RL), and **Agentic AI**, with applications in healthcare, biomedicine, and recommendation systems. I am particularly interested in how these models can be used for **reasoning, retrieval, and decision-making** in complex, real-world scenarios.

Recently, I led a survey paper about **Adaptation of Agentic AI**. We categorized existing agent papers into four paradigms. Feel free to check the survey **here**.

Experience

Keiji AI Inc. — Remote
Machine Learning Research Intern, advised by Dr. Zifeng Wang and Prof. Jimeng Sun — 08/2025 – Present
Project: Developed a Data Science Assistant/Cursor for clinical data analysis and visualization.

Stores Foundation AI (SFAI), Amazon — Palo Alto, CA
Applied Scientist Intern, advised by Dr. Hyokun Yun, and Lihong Li — 05/2025 – 08/2025
Project: Continual SFT; Bridging SFT and RL; Shopping and Recommendation Agents.

Search Explore, Amazon — Palo Alto, CA
Applied Scientist Intern, advised by Dr. Kun Qian, Sheng Wang and Danai Koutra — 05/2024 – 08/2024
Project: LLMs for Knowledge Graph Augmented Text Embedding.

AI4Science Asia Group, Microsoft Research — Beijing, China
Research Intern, advised by Dr. Lijun Wu and Tao Qin — 11/2022 – 07/2023
Project: Multi-modal Scientific Pre-Training Models.

ML & DRL Group, Microsoft Research Asia — Beijing, China
Research Intern, advised by Dr. Lijun Wu and Tao Qin — 03/2022 – 06/2022
Project: Drug Discovery.

Selected Papers and Projects [Google Scholar]

DeepRetrieval: Hacking Real Search Engines and Retrievers with LLMs via RL (COLM 2025) ☆ 685

- Developed DeepRetrieval, a RL-based **search agent** that optimizes retrieval performance by interacting with the downstream search engine and using verifiable reward (e.g., Recall, NDCG) without supervised queries.
- Achieved state-of-the-art performance on real-world retrieval tasks, boosting recall@10 from 24.68% to 65.07% on scientific (PubMed) paper retrieval and from 32.11% to 63.18% on clinical trial retrieval.

Rec-R1: Bridging Generative Large Language Models and User-Centric Recommendation Systems via Reinforcement Learning (TMLR 2025) ☆ 121

- Developed Rec-R1, a general RL **agentic framework** that directly optimizes LLM4Rec framework using feedback from recommendation systems.
- Achieved substantial improvements across product search and sequential recommendation, and product re-ranking.

TALR: SFT Doesn't Always Hurt General Capabilities: Revisiting Domain-Specific SFT in LLMs

- **Challenged prior papers' claims** that domain-specific SFT severely degrades general capabilities, showing that using smaller learning rates can largely mitigate performance loss.
- Proposed TALR (Token-Adaptive Loss Reweighting), a method that further reduces general-performance degradation by **adaptively down-weighting hard tokens** during fine-tuning.
- Conducted token-level analysis revealing that TALR induces a **curriculum-like learning dynamic**.

Data Science Agent of Keiji AI

- Developed an intelligent agent capable of interacting with complex clinical data, generating analysis plans, and visualizing results seamlessly.
- Streamlined the workflow for biomedical researchers and clinicians by integrating data processing, analysis, and visualization into a unified system.

Selected Awards

- Outstanding Graduates (top 5%)
- Comprehensive Excellence Award (top 10%)
- Intel Fellowship for Graduates of Tsinghua University
- Stars of Tomorrow Internship Award of Microsoft Research Asia
- NeurIPS 2025 Top Reviewer

Skills

Programming Languages: Python, Matlab, C/C++, C#, Bash
Software Libraries and Tools: PyTorch, NumPy, Scikit-Learn, PyTorch Geometric, Git, LaTeX, Docker
Hobbies: Music production (hip-hop/EDM), running, basketball, tennis

Invited Talks and Coverage

Invited Talks:

LLMs, RL, and the Rise of Self-Evolving Agents in Real-World Healthcare
Machine Learning Seminar at UIUC, Oct 2025

From LLMs to Self-Evolving Agents: Practice and Lessons Learned Unity, Oct 2025

Unlocking the Future of Drug Discovery: Towards Building Multi-modal Biomedical Foundation Models NCSA Talk Session at UIUC, Dec 2023

Coverage:

VentureBeat. Title: s3: The new RAG framework that trains search agents with minimal data. May, 2025

Other Selected Projects

TrialPanorama: Developing Large Language Models for Clinical Research Using One Million Clinical Trials

- Introduced TrialPanorama, a large-scale structured resource that aggregates 1.6M clinical trial records from 15 global registries and links them with biomedical ontologies and literature.
- Trained domain-specific LLMs with **Cascade SFT** and **Cascade RL**, which significantly **suppresses catastrophic forgetting** and outperforms closed-source LLMs on eight clinical trial tasks.

Selected Publications

"*" indicates equal contribution.

General Domains

Jiacheng Lin, Tian Wang, Kun Qian.
Rec-R1: Bridging Generative Large Language Models and User-Centric Recommendation Systems via Reinforcement Learning.
TMLR 2025.

Pengcheng Jiang*, **Jiacheng Lin***, Lang Cao, Runchu Tian, SeongKu Kang, Zifeng Wang, Jimeng Sun, Jiawei Han.
DeepRetrieval: Hacking Real Search Engines and Retrievers with Large Language Models via Reinforcement Learning.
COLM 2025.

Pengcheng Jiang, Xueqiang Xu, **Jiacheng Lin**, Jinfeng Xiao, Zifeng Wang, Jimeng Sun, Jiawei Han.
s3: You Don't Need That Much Data to Train a Search Agent via RL.
EMNLP 2025.

Jiacheng Lin, Meng Xu, Zhihua Xiong, Huangang Wang.
CAMBranch: Contrastive Learning with Augmented MILPs for Branching.
ICLR 2024.

Ziyi Chen, Xiaocong Yang, **Jiacheng Lin**, Chenkai Sun, Jie Huang, Kevin Chen-Chuan Chang.
Cascade Speculative Drafting for Even Faster LLM Inference.
NeurIPS 2024, 2023.

Pengcheng Jiang, **Jiacheng Lin**, Zifeng Wang, Jimeng Sun, Jiawei Han.
GenRES: Rethinking Evaluation for Generative Relation Extraction in the Era of Large Language Models.
NAACL 2024.

Jiacheng Lin, Jialin Zhu, Huangang Wang, Tao Zhang.
Learning to branch with Tree-aware Branching Transformers.
Knowledge-Based Systems, 2022.

Jiacheng Lin, Kun Qian, Haoyu Han, Nurendra Choudhary, Tianxin Wei, Zhongruo Wang, Sahika Genc, Edward W. Huang, Sheng Wang, Karthik Subbian.
GT2VEC: Large Language Models as Multi-Modal Encoders for Text and Graph-Structured Data.
KDD 2025 Workshop SKnow-LLM, 2024.

Healthcare Domain

Jiacheng Lin, Hanwen Xu, Addie Woicik, Jianzhu Ma, Sheng Wang.
Pisces: A cross-modal contrastive learning approach to synergistic drug combination prediction.
RECOMB 2023.

Jiacheng Lin, Lijun Wu, Jinhua Zhu, Xiaobo Liang, Yingce Xia, Shufang Xie, Tao Qin, Tie-Yan Liu.
R2-DDI: relation-aware feature refinement for drug–drug interaction prediction.
Briefings in Bioinformatics, 2023.

Hanwen Xu*, **Jiacheng Lin***, Addie Woicik, Zixuan Liu, Jianzhu Ma, Sheng Zhang, Hoifung Poon, Liewei Wang, Sheng Wang.
Pisces: A multi-modal data augmentation approach for drug combination synergy prediction.
Cell Genomics, 2025.

Xiaoyang Chen, Keyi Li, Xuejian Cui, Zian Wang, Qun Jiang, **Jiacheng Lin**, Zhen Li, Zijing Gao, Rui Jiang.
EpiAgent: Foundation model for single-cell epigenomic data.
Nature Methods, 2025.

Professional Services

Conference Program Committee Member
ICLR 2026, NeurIPS 2025, ICLR 2025, AAAI 2025, NAACL 2025, COLING 2025, ICRA 2025, RECOMB 2024

Journal Reviewer
ACM Transactions on Knowledge Discovery from Data (TKDD);
ACM Transactions on Information System (TOIS);
Public Library of Science One (PLOS ONE);
The Computer Journal;