QIRUN DAI

(+1) 530-709-2773 ⋄ qirundai@uchicago.edu ⋄ qrdai.github.io

EDUCATION

The University of Chicago, IL   *Sept. 2025 - Present*
Ph.D. in Computer Science   **Advisor:** Prof. Chenhao Tan

Fudan University, Shanghai   *Sept. 2021 - June 2025*
B.Eng. in Artificial Intelligence (Honor Class)   **GPA:** 3.83/4.00; **Rank:** 1/25
Honor Thesis: Building a Trustworthy Clinical Reasoning Agent by Attributing LLM Generations to
Training Data and Context Synchronously   **Advisor:** Prof. Jiaqing Liang

University of California, Davis   *Sept. 2023 - Dec. 2023*
UCEAP Exchange Student in Computer Science   **GPA:** 4.00/4.00

RESEARCH INTERESTS

I am broadly interested in data-centric topics in post-training for large language models (LLMs).

- Extending the generalization and transfer of reasoning behaviors of LLMs: from math-centered paradigms to versatile domains (e.g., argumentative writing) and formalisms (e.g., causal reasoning).

- Data curation, selection, attribution, and data-efficient supervision algorithms.

PUBLICATIONS

* denotes equal contributions, and † denotes equal advising.

Executable Counterfactuals: Improving LLMs' Causal Reasoning Through Code
Aniket Vashishtha[*], **Qirun Dai**[*], Hongyuan Mei, Amit Sharma[†], Chenhao Tan[†], Hao Peng[†].
NeurIPS 2025 Workshop on Foundations of Reasoning in LMs; Under Review at ICLR 2026

The Best Instruction-Tuning Data are Those That Fit
Dylan Zhang, **Qirun Dai**, Hao Peng.
NeurIPS 2025 (Spotlight)

Improving Influence-based Instruction Tuning Data Selection for Balanced Learning of Diverse Capabilities
Qirun Dai, Dylan Zhang, Jiaqi W. Ma, Hao Peng.
*Findings of **EMNLP 2025**; **ICLR 2025** Workshop on DATA-FM*

A Survey of Data Attribution: Methods, Applications, and Evaluation in the Era of Generative AI
Junwei Deng[*], Yuzheng Hu[*], Pingbang Hu[*], Ting-Wei Li[*], Shixuan Liu[*], Jiachen T. Wang, Dan Ley, **Qirun Dai**, Benhao Huang, Jin Huang, Cathy Jiao, Hoang Anh Just, Yijun Pan, Jingyan Shen, Yiwen Tu, Weiyi Wang, Xinhe Wang, Shichang Zhang, Shiyuan Zhang, Ruoxi Jia, Himabindu Lakkaraju, Hao Peng, Weijing Tang, Chenyan Xiong, Jieyu Zhao, Hanghang Tong, Han Zhao, Jiaqi W. Ma.
Preprint, 2025

Demonstration Distillation for Efficient In-Context Learning
Tong Chen, **Qirun Dai**, Zhijie Deng, Dequan Wang.
Preprint, 2024

RESEARCH EXPERIENCES

The University of Chicago   *May 2025 - Present*
Research Assistant, Department of Computer Science
Advisor: Prof. **Chenhao Tan**. External Collaborator: Prof. **Hao Peng**.

- Proposed a computational framework that establishes code-based verifiable counterfactual reasoning with three cognitive steps: abduction, intervention, and prediction. Along with a scalable data synthesis pipeline, our framework identified the failure of both SOTA LLMs and long-CoT distillation-based SFT in performing true counterfactual reasoning, but showed the promise of RLVR under extensive exploration. [**In Submission to ICLR 2026**]

University of Illinois Urbana-Champaign   *Apr. 2024 - Aug. 2025*
Research Intern, Siebel School of Computing and Data Science
Advisor: Prof. **Hao Peng**, and Prof. **Jiaqi W. Ma**.

- Identified the poor performance of data selection methods built upon gradient-based influence estimation, when instruction tuning LLMs for multiple diverse capabilities. Attributed this problem to an inherent bias in cross-task influence, and proposed an iterative data selection algorithm with instance-level normalization as the solution. [**Findings of EMNLP 2025**]

- Further proposed a response selection framework that aligns instruction tuning data with the pretrained distribution of the target model, revealing that the on-policy property is not exclusive to RL data, but can also boost efficiency and generalization simultaneously in SFT. [**NeurIPS 2025 (Spotlight)**]

Shanghai Jiao Tong University   *Jan. 2023 - Mar. 2024*
Research Intern, Qingyuan Research Institute and Shanghai AI Lab
Advisor: Prof. **Dequan Wang**.

- Developed an agentic framework that iteratively compresses in-context learning (ICL) demonstrations with three LLM agents targeting different granularity: Distillist, Generalist and Specialist. [**Submitted to ICLR 2024**]

- Further proposed an ICL demonstration selection method built on sparse retrieval techniques, targeting knowledge- and reasoning-intensive QA tasks where pretrained dense embeddings can be easily confused due to a lack of strong reasoning or domain-specific knowledge.

TEACHING EXPERIENCES

Introduction to Computer Systems   **Fudan University**
Teaching Assistant with Prof. Jiaqing Liang   *Fall Semester 2024*

AWARDS

Deans' Honors List
University of California, Davis, 2023

Outstanding Student Award
Fudan University, 2022 - 2023

Panasonic Scholarship
Fudan University, 2021 - 2022

SKILLS

Programming: Python, C/C++, CUDA, MATLAB, R, HTML/CSS/JavaScript, Java
Languages: English (fluent), Chinese (native)