

ZHINING LIU

Email: zhining.liu@outlook.com ◇ Homepage: <https://zhiningliu.com>

Palo Alto, CA

🔗 Github:ZhiningLiu1998 🎓 GoogleScholar 🌐 LinkedIn

My recent research focuses on building **self-improving agentic AI for industrial R&D at scale**. More broadly, I study **data-centric trustworthy AI**, including reliable long-context LLM/VLM reasoning and data-efficient learning with noisy context and limited data. I have authored 30+ papers at top venues including ICML/NeurIPS/ICLR, with 1,000+ citations, and maintain open-source projects with 4,000+ GitHub stars and 140k+ downloads.

EDUCATION

- **University of Illinois Urbana-Champaign** Urbana, IL
Ph.D. in Computer Science 08/2022 - 08/2026
Advisor: Prof. [Hanghang Tong](#)
- **Jilin University** Changchun, China
M.Eng. in Computer Science - Advisor: Prof. [Yi Chang](#) 08/2019 - 06/2022
B.Sc. in Computer Science - Tang-Aoqing Honors Program (Top ~0.5% at matriculation) 08/2015 - 06/2019

EXPERIENCE

- **Amazon.com, Inc.**
Applied Scientist 07/2026 - Present, Palo Alto, CA
– Build self-improving agent harnesses for industrial R&D at scale.
Applied Scientist Intern @Ads 05/2025 - 12/2025, Palo Alto, CA
– Improved VLM reasoning reliability, resulting in papers at ICLR'26, ICML'26, and ACL'26.
– Mentor: Dr. [Hui Liu](#)
Applied Scientist Intern @Rufus 05/2024 - 08/2024, Palo Alto, CA
– Improved RAG-enhanced LLM reasoning, resulting in a paper at ACL'25 Main.
– Mentor: Dr. [Rana Ali Amjad](#)
Applied Scientist Intern @Search 05/2023 - 08/2023, Seattle, WA
– Developed multi-task learning methods for user query understanding and partially ordered entity ranking.
– Mentors: Dr. [Minmin Shen](#), Dr. [Vassilis N. Ioannidis](#), Dr. [Changhe Yuan](#)
- **University of Illinois Urbana-Champaign**
Research Assistant @Computer Science Dept. 08/2022 - 08/2026, Urbana, IL
- **Microsoft Research Asia**
Research Intern @Machine Learning Group 08/2018 - 06/2019, Beijing, China
– Explore large-scale, highly-imbalanced classification for intelligent logistics. Our proposed solution was adopted by an industry-leading logistics company and published on ICDE'20 and NeurIPS'20.
– Mentors: Dr. [Wei Cao](#), Dr. [Jiang Bian](#)

PUBLICATIONS

First/Co-first Author

1. **Zhining Liu**, Ziyi Chen, Hui Liu, Chen Luo, Xianfeng Tang, Suhang Wang, Jingying Zeng, Zhenwei Dai, Zhan Shi, Tianxin Wei, Hanqing Lu, Benoit Dumoulin, Hanghang Tong. *Seeing but Not Believing: Probing the Disconnect Between Visual Attention and Answer Correctness in VLMs*. (ICLR 2026).
2. **Zhining Liu***, Xiao Lin*, Ze Yang*, Gaotang Li, Ruizhong Qiu, Shuke Wang, Hui Liu, Haotian Li, Sumit Keswani, Vishwa Pardeshi, Huijun Zhao, Wei Fan, Hanghang Tong. *MORALISE: A Structured Benchmark for Moral Alignment in Visual Language Models*. (ICML 2026).

3. **Zhining Liu***, Tianyi Wang*, Xiao Lin, Penghao Ouyang, Gaotang Li, Ze Yang, Hui Liu, Sumit Keswani, Vishwa Pardeshi, Huijun Zhao, Wei Fan, Hanghang Tong. *Do VLMs Have a Moral Backbone? A Study on the Fragile Morality of Vision-Language Models*. (**ACL Findings 2026**).
4. **Zhining Liu**, Zihao Li, Ze Yang, Tianxin Wei, Jian Kang, Yada Zhu, Hendrik Hamann, Jingrui He, Hanghang Tong. *CLIMB: Class-imbalanced Learning Benchmark on Tabular Data*. (**NeurIPS 2025**). [\[Paper\]](#) [\[Code\]](#)
5. **Zhining Liu**, Rana Ali Amjad, Ravinarayana Adkathimar, Tianxin Wei, Hanghang Tong. *SelfElicit: Your Language Model Secretly Knows Where the Relevant Evidence is*. (**ACL Main 2025**). [\[Paper\]](#) [\[Code\]](#)
6. **Zhining Liu**, Ze Yang, Xiao Lin, Ruizhong Qiu, Tianxin Wei, Yada Zhu, Hendrik Hamann, Jingrui He, Hanghang Tong. *Breaking Silos: Adaptive Model Fusion Unlocks Better Time Series Forecasting*. In 42nd International Conference on Machine Learning (**ICML 2025**). [\[Paper\]](#) [\[Code\]](#)
7. **Zhining Liu**, Ruizhong Qiu, Zhichen Zeng, Yada Zhu, Hendrik Hamann, Hanghang Tong. *AIM: Attributing, Interpreting, Mitigating Data-encoded Unfairness*. In 30th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (**KDD 2024**). [\[Paper\]](#) [\[Code\]](#)
8. **Zhining Liu**, Zhichen Zeng, Ruizhong Qiu, Hyunsik Yoo, David Zhou, Zhe Xu, Yada Zhu, Kommy Welde-mariam, Jingrui He, and Hanghang Tong. *Class-Imbalanced Graph Learning without Class Rebalancing*. In 41st International Conference on Machine Learning (**ICML 2024**). [\[Paper\]](#) [\[Code\]](#)
9. **Zhining Liu***, Eunice Chan*, Ruizhong Qiu, Yuheng Zhang, Ross Maciejewski, Hanghang Tong. *Group Fairness via Group Consensus*. In ACM Conference on Fairness, Accountability, and Transparency (**FAccT 2024**). [\[Paper\]](#)
10. **Zhining Liu***, Hangting Ye*, Wei Cao, Amir Mohammad Amiri, Jiang Bian, Yi Chang, Jon D. Lurie, Jim Weinstein, Tie-Yan Liu. *Web-based Long-term Spine Treatment Outcome Forecasting*. In 29th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (**KDD 2023**). [\[Paper\]](#)
11. **Zhining Liu***, Hangting Ye*, Xinyi Shen, Wei Cao, Shun Zheng, Xiaofan Gui, Huishuai Zhang, Yi Chang, Jiang Bian. *UADB: Unsupervised Anomaly Detection Booster*. In IEEE 39th International Conference on Data Engineering (**ICDE 2023**). [\[Paper\]](#) [\[Code\]](#)
12. **Zhining Liu**, Pengfei Wei, Jing Jiang, Wei Cao, Jiang Bian, and Yi Chang. *MESA: Boost Ensemble Imbalanced Learning with Meta-sampler*. In Proceedings of the 33rd Advances in Neural Information Processing Systems (**NeurIPS 2020**). [\[Paper\]](#) [\[Code\]](#)
13. **Zhining Liu**, Wei Cao, Zhifeng Gao, Jiang Bian, Hechang Chen, Yi Chang, and Tie-Yan Liu. *Self-paced ensemble for highly imbalanced massive data classification*. In Proceedings of the IEEE 36th International Conference on Data Engineering (**ICDE 2020**). [\[Paper\]](#) [\[Code\]](#)

Co-author Papers

1. Yanjun Zhao, Ruizhong Qiu, Tianxin Wei, Yuanchen Bei, **Zhining Liu**, Lingjie Chen, Ismini Lourentzou, Hanghang Tong, Jingrui He. *ReContext: Recursive Evidence Replay as LLM Harness for Long-Context Reasoning*. **EMNLP Findings 2026**. [\[Paper\]](#) [\[Code\]](#)
2. Xiao Lin, Zhicheng Tang, Weilin Cong, Mengyue Hang, Kai Wang, Yajuan Wang, Zhichen Zeng, Ting-Wei Li, Hyunsik Yoo, **Zhining Liu**, Xuying Ning, Ruizhong Qiu, Wen-Yen Chen, Shuo Chang, Rong Jin, Huayu Li, Hanghang Tong. *Mixture of Sequence: Theme-Aware Mixture-of-Experts for Long-Sequence Recommendation*. **WWW 2026 (Oral)**.
3. Yuanchen Bei, Tianxin Wei, Xuying Ning, Yanjun Zhao, **Zhining Liu**, Xiao Lin, Yada Zhu, Hendrik Hamann, Jingrui He, Hanghang Tong. *Mem-Gallery: Benchmarking Multimodal Long-Term Conversational Memory for MLLM Agents*. **ACL Main 2026**.
4. Chengming Cui, Tianxin Wei, Ziyi Chen, Ruizhong Qiu, Zhichen Zeng, **Zhining Liu**, Xuying Ning, Duo Zhou, Jingrui He. *AdaFuse: Adaptive Ensemble Decoding for Large Language Models*. **ACL Main 2026**.
5. Hyunsik Yoo, Ting-Wei Li, SeongKu Kang, **Zhining Liu**, Charlie Xu, Qilin Qi, Hanghang Tong. *Continual Low-Rank Adapters for LLM-based Generative Recommender Systems*. **ICLR 2026**.
6. Zihao Li, Xiao Lin, **Zhining Liu**, Jiaru Zou, Ziwei Wu, Lecheng Zheng, Dongqi Fu, Yada Zhu, Hendrik Hamann, Hanghang Tong, Jingrui He. *Language in the Flow of Time: Time-Series-Paired Texts Weaved into a Unified Temporal Narrative*. **ICLR 2026**.

7. Qi Yu, Zhichen Zeng, Yuchen Yan, **Zhining Liu**, Baoyu Jing, Ruizhong Qiu, Ariful Azad, Hanghang Tong. *PLANETALIGN: A Comprehensive Python Library for Benchmarking Network Alignment*. **ICLR 2026**.
8. Zihao Li, Zhichen Zeng, Xiao Lin, Feihao Fang, Yanru Qu, Zhe Xu, **Zhining Liu**, Xuying Ning, Tianxin Wei, Ge Liu, Hanghang Tong, Jingrui He. *Flow Matching Meets Biology and Life Science: A Survey*. **npj Artificial Intelligence 2026**.
9. Baoyu Jing, Sanhorn Chen, Lecheng Zheng, Boyu Liu, Zihao Li, Jiaru Zou, Tianxin Wei, **Zhining Liu**, Zhichen Zeng, Ruizhong Qiu, Xiao Lin, Yuchen Yan, Qi Yu, Dongqi Fu, Jingrui He, Hanghang Tong. *TSAQA: Time Series Analysis Question and Answering Benchmark*. **GEM 2026**. [\[Paper\]](#) [\[Code\]](#)
10. Ruizhong Qiu, Hanqing Zeng, Yinglong Xia, Yiwen Meng, Ren Chen, Jiarui Feng, Dongqi Fu, Qifan Wang, Jiayi Liu, Jun Xiao, Xiangjun Fan, Benyu Zhang, Hong Li, **Zhining Liu**, Hyunsik Yoo, Hanghang Tong. *ReMix: Reinforcement Routing for Mixtures of LoRAs in LLM Finetuning*. **LLA Workshop at ICLR 2026**. [\[Paper\]](#)
11. Dongqi Fu, Yada Zhu, **Zhining Liu**, Lecheng Zheng, Xiao Lin, Zihao Li, Liri Fang, Katherine Tieu, Onkar Bhardwaj, Kommy Weldemariam, Hanghang Tong, Hendrik Hamann, Jingrui He *ClimateBench-M: A Multi-Modal Climate Data Benchmark with a Simple Generative Method*. **CIKM 2025**.
12. Zihao Li, Feihao Fang, Xitong Zhang, Jiaru Zou, **Zhining Liu**, Wei Xiong, Ziwei Wu, Baoyu Jing, Jingrui He *Not All Voices Are Rewarded Equally: Probing and Repairing Reward Models across Human Diversity*. **EMNLP Findings 2025**.
13. Yunzhe Li, Juntong Wang, Hari Sundaram, **Zhining Liu**. *LLM-RecG: A Semantic Bias-Aware Framework for Zero-Shot Sequential Recommendation*. **RecSys 2025**.
14. Wenxuan Bao, Zhichen Zeng, **Zhining Liu**, Hanghang Tong, Jingrui He. *Matcha: Mitigating Graph Structure Shifts with Test-Time Adaptation*. **ICLR 2025**. [\[Paper\]](#)
15. Xiao Lin, **Zhining Liu**, Dongqi Fu, Ruizhong Qiu, Hanghang Tong. *BackTime: Backdoor Attacks on Multivariate Time Series Forecasting*. In Proceedings of the 38th Advances in Neural Information Processing Systems (**NeurIPS 2024 Spotlight**) [\[Paper\]](#)
16. Zhichen Zeng, Ruizhong Qiu, Zhe Xu, **Zhining Liu**, Yuchen Yan, Tianxin Wei, Lei Ying, Jingrui He, Hanghang Tong. *Graph Mixup on Approximate Gromov-Wasserstein Geodesics*. In 41st International Conference on Machine Learning (**ICML 2024**). [\[Paper\]](#)
17. Hyunsik Yoo, Zhichen Zeng, Jian Kang, Ruizhong Qiu, David Zhou, **Zhining Liu**, Fei Wang, Charlie Xu, Eunice Chan, Hanghang Tong. *Ensuring User-side Fairness in Dynamic Recommender Systems*. In 2024 ACM Web Conference (**Web 2024**). [\[Paper\]](#)
18. Zhichen Zeng, Boxin Du, Si Zhang, Yinglong Xia, **Zhining Liu**, Hanghang Tong. *Hierarchical Multi-Marginal Optimal Transport for Network Alignment*. In 38th AAAI Conference on Artificial Intelligence (**AAAI 2024**). [\[Paper\]](#)
19. Guo, Kai, Xiaofeng Cao, **Zhining Liu**, and Yi Chang. *Taming Over-Smoothing Representation on Heterophilic Graphs*. **Information Sciences (2023)**: 119463. [\[Paper\]](#)
20. Warmesley, Dana, Alex Waagen, Jiejun Xu, **Zhining Liu**, and Hanghang Tong. *A Survey of Explainable Graph Neural Networks for Cyber Malware Analysis*. In Proceedings of the 2022 IEEE International Conference on Big Data (**Big Data 2022**). [\[Paper\]](#)
21. Zhepei Wei, Yue Wang, Jinnan Li, **Zhining Liu**, Erxin Yu, Yuan Tian, Xin Wang, Yi Chang. *Towards Unified Representations of Knowledge Graph and Expert Rules for Machine Learning and Reasoning*. **AAACL Main 2022**. [\[Paper\]](#)

Preprints and Submissions

1. Xuying Ning, Katherine Tieu, Dongqi Fu, Tianxin Wei, Zihao Li, Yuanchen Bei, Jiaru Zou, Mengting Ai, **Zhining Liu**, et al. *Code as Agent Harness*. arXiv preprint arXiv:2605.18747 (2026). [\[Paper\]](#) [\[Code\]](#) (**Hugging Face Daily #1 Paper**)
2. Tianxin Wei*, Ting-Wei Li*, **Zhining Liu***, Xuying Ning, Ze Yang, Jiaru Zou, Zhichen Zeng, Ruizhong Qiu, Xiao Lin, Dongqi Fu, Zihao Li, Mengting Ai, Duo Zhou, Wenxuan Bao, Yunzhe Li, Gaotang Li, Cheng Qian, Yu Wang, Xiangru Tang, Yin Xiao, Liri Fang, Hui Liu, Xianfeng Tang, Yuji Zhang, Chi Wang, Jiaxuan You, Heng Ji, Hanghang Tong, Jingrui He. *Agentic Reasoning for Large Language Models*. arXiv preprint arXiv:2601.12538 (2026). [\[Paper\]](#) [\[Code\]](#) (**Hugging Face Weekly #1 Paper**)

3. Zhichen Zeng, Mengyue Hang, Xiaolong Liu, Xiaoyi Liu, Xiao Lin, Ruizhong Qiu, Tianxin Wei, **Zhining Liu**, Siyang Yuan, Chaofei Yang, Yiqun Liu, Hang Yin, Jiyan Yang, Hanghang Tong. *Hierarchical LoRA MoE for Efficient CTR Model Scaling*. arXiv preprint arXiv:2510.10432 (2025). [\[Paper\]](#)
4. Xiao Lin, Zhichen Zeng, Tianxin Wei, **Zhining Liu**, Yuzhong Chen, Hanghang Tong. *CATS: Mitigating Correlation Shift for Multivariate Time Series Classification*. arXiv preprint arXiv:2504.04283 (2025). [\[Paper\]](#)
5. Tianxin Wei, Zhichen Zeng, Ruizhong Qiu, **Zhining Liu**, Xuying Ning, Xinrui He, Wenxuan Bao, Qi He, Xianfeng Tang, Hanghang Tong, Jingrui He. *Inference Scaling of LLM Ensembling: Bridging Token Spaces with Token Translation*. In submission.
6. Lingjie Chen, Ruizhong Qiu, Siyu Yuan, **Zhining Liu**, Tianxin Wei, Hyunsik Yoo, Zhichen Zeng, Deqing Yang, Hanghang Tong. *WAPITI: A Watermark for Finetuned Open-Source LLMs*. arXiv preprint arXiv:2410.06467 (2024). [\[Paper\]](#)
7. Yuchen Yan, Yuzhong Chen, Huiyuan Chen, Xiaoting Li, Zhe Xu, Zhichen Zeng, Lihui Liu, **Zhining Liu**, Hanghang Tong. *TheGCN: Temporal Heterophilic Graph Convolutional Network*. arXiv preprint arXiv:2412.16435 (2024). [\[Paper\]](#)
8. **Zhining Liu**, Pengfei Wei, Zhepei Wei, Boyang Yu, Jing Jiang, Wei Cao, Jiang Bian, Yi Chang. *Handling Inter-Class and Intra-Class Imbalance in Class-Imbalanced Learning*. arXiv preprint arXiv:2111.12791 (2021). [\[Paper\]](#)
9. **Zhining Liu**, Jian Kang, Hanghang Tong, Yi Chang. *IMBENS: Ensemble Class-Imbalanced Learning in Python*. arXiv preprint arXiv:2111.12776 (2021). [\[Paper\]](#) [\[Code\]](#)

(*indicates equal contribution)

SELECTED OPEN-SOURCE PROJECTS

4,000+ GitHub stars and 200,000+ downloads in total by August 2026

- (400+ stars) **IMBENS: Class-imbalanced Ensemble Learning in Python** 2022
 - 3K+ downloads/month Python toolbox for class-imbalanced ensemble learning
 - Links: [\[Github\]](#) [\[Documentation\]](#) [\[Gallery\]](#) [\[PyPI\]](#)
- (1,500+ stars) **Awesome Imbalanced Learning** 2022
 - Curated list of imbalanced learning papers, codes, libraries, and other learning resources.
 - Links: [\[English\]](#) [\[Chinese\]](#)
- (600+ stars) **Awesome Machine Learning Resources** 2020
 - Curated list of awesome resources across various machine learning topics.
 - Links: [\[English\]](#) [\[Chinese\]](#)

SELECTED AWARDS & HONORS

- C.W. Gear Outstanding Graduate Student (top 0.2%), University of Illinois 2026
- C.L. and Jane W.-S. Liu Award (2 recipients; top 0.3%), University of Illinois 2025
- Top 10 Graduate Student (University's Highest Honor; top 0.04%), Jilin University 2022
- National Scholarship (Top 0.2% Nationally), Ministry of Education of China 2020
- National Scholarship (Top 0.2% Nationally), Ministry of Education of China 2019
- National Endeavor Scholarship (Top 2% Nationally), Ministry of Education of China 2017
- Academic Performance Scholarship, Jilin University 2019,2020,2021
- Undergraduate Research Scholarship, Jilin University 2018
- Outstanding Student Leader Service Award, Jilin University 2016,2017,2018,2019

TALKS

- 06/2025: **Fidelity Investments** - Your Language Model Secretly Knows the Evidence [\[Slides\]](#)

PROFESSIONAL SERVICES

- **Conference Program Committee/Reviewer**
 - The Fourteenth International Conference on Learning Representations (ICLR) 2026
 - The 63rd Annual Meeting of the Association for Computational Linguistics (ACL) 2025
 - Forty-Second International Conference on Machine Learning (ICML) 2025
 - ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD) 2024
 - IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2022
- **Journal Reviewer**
 - Transactions on Machine Learning Research (TMLR) 2025
 - IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) 2024-2025
 - IEEE Transactions on Knowledge and Data Engineering (TKDE) 2019-2024
 - IEEE Transactions on Cybernetics (TCyb) 2020
 - Elsevier Neurocomputing (Neurocomputing) 2019-2020
- **External Reviewer**
 - ACM International Conference on Information and Knowledge Management (CIKM) 2022
 - ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD) 2020
 - IEEE International Conference on Data Mining (ICDM) 2019
- **Conference Organizing Volunteer**
 - ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR) 2020

MENTORING EXPERIENCE

- **Mr. Tianyi Wang (M.S. Student, University of Illinois CS)** 2025 - Now
 - Improving VLM reasoning and reflection efficiency. Paper published in **ACL 2026**.
- **Mr. Haopeng Ouyang (Senior, University of Illinois CS)** 2025 - Now
 - VLM moral value alignment, jailbreaking, and defense. Paper published in **ACL 2026**.
- **Ms. Shuke Wang (Junior, University of Illinois ECE)** 2024 – 2025
 - Mentoring her on research in dynamic fairness and LLM moral value alignment evaluation. We produced an extensive multi-modal moral benchmark. The paper is currently under submission.
- **Mr. Xiao Lin (Ph.D. student, University of Illinois CS)** 2024 – 2025
 - Guided academic writing, including structuring arguments, refining storytelling, and designing experiments. Helped him build a solid foundation in academic writing and scientific communication. Our collaboration led to a first-authored paper published in **NeurIPS 2024 (Spotlight)**.
- **Ms. Eunice Chan (Ph.D. student, University of Illinois CS)** 2023 – 2024
 - Closely guided her throughout the entire research cycle of a project on group algorithmic fairness, including ideation, experimental design, code implementation, and paper writing. Helped her gain independence in conducting research and strengthened her ability to communicate effectively. Our collaboration resulted in a co-first-authored paper published in **ACM FAccT 2024**.
- **Mr. Hangting Ye (B.Sc. then Ph.D. student, Jilin University)** 2021 – 2022
 - Provided extensive guidance on research in intelligent healthcare and anomaly detection, assisting with problem formulation, experiment design, and manuscript preparation. Supported his transition from undergraduate to Ph.D.-level research by refining his analytical and problem-solving skills. We made two co-first-authored publications in **IEEE ICDE 2023** and **ACM SIGKDD 2023**.

TEACHING EXPERIENCE

- University of Illinois CS 512 Data Mining Principle (Lead Teaching Assistant) Spring, 2024

[CV compiled on Monday 24th August, 2026]