

Jiancheng (JC) Liu

Email: ljcc (at) openai (dot) com | [Personal Website](#) | [Google Scholar](#) | [LinkedIn](#)

INDUSTRIAL EXPERIENCE

OpenAI

Member of Technical Staff

San Francisco, California

2025 - Present

- Personal AGI core models team.
- Contributed to the [GPT-5 series](#) and [gpt-oss series](#).

IBM Research

Research Scientist Intern

2024

San Jose, California

Taichi Graphics Technology

Research Engineer, Startup Member

2021 - 2022

Beijing, China

Megvii Technology, Inc.

Research Intern

2016 - 2017

Beijing, China

EDUCATION

Michigan State University, East Lansing, Michigan

Master of Science in Computer Science

2023 - 2025

Computer Science and Engineering

Tsinghua University, Beijing, China

Bachelor of Engineering in Computer Science

2015 - 2020

Institute for Interdisciplinary Information Sciences

PEER-REVIEWED PUBLICATIONS

* denotes equal contribution; † denotes an author under my supervision.

1. Z. Di, Z. Zhu, J. Jia, **J. Liu**, Z. Takhirov, B. Jiang, Y. Yao, S. Liu, Y. Liu, "Label Smoothing Improves Machine Unlearning," **ICLR'26**. [[Link](#)]
2. **J. Liu***, **C. Fan***[†], L. Lin*, J. Jia, R. Zhang, S. Mei, S. Liu, "Simplicity Prevails: Rethinking Negative Preference Optimization for LLM Unlearning," **NeurIPS'25**. [[Link](#)]
3. **J. Liu***, Y. Yao*, Y. Gong*, X. Liu, Y. Wang, X. Lin, S. Liu, "Can Adversarial Examples Be Parsed to Reveal Victim Model Information?" **WACV'25**. [[Link](#)]
4. M. Zhang, C. Wang, S. C. Kramer, J. G. Wallwork, S. Li, **J. Liu**, X. Chen, M. D. Piggott, "Towards Universal Mesh Movement Networks," **NeurIPS'24 (Spotlight, acceptance rate ~2%)**. [[Link](#)]
5. J. Jia, **J. Liu**, Y. Zhang, P. Ram, N. Baracaldo, S. Liu, "WAGLE: Strategic Weight Attribution for Effective and Modular Unlearning in Large Language Models," **NeurIPS'24**. [[Link](#)]
6. Y. Zhang, X. Chen, J. Jia, Y. Zhang, C. Fan, **J. Liu**, M. Hong, K. Ding, S. Liu, "Defensive Unlearning with Adversarial Training for Robust Concept Erasure in Diffusion Models," **NeurIPS'24**. [[Link](#)]
7. Y. Zhang, C. Fan, Y. Zhang, Y. Yao, J. Jia, **J. Liu**, G. Zhang, G. Liu, R. Kompella, X. Liu, S. Liu, "UnlearnCanvas: A Stylized Image Dataset to Benchmark Machine Unlearning for Diffusion Models," **NeurIPS'24 D&B**. [[Link](#)]
8. J. Jia, Y. Zhang, Y. Zhang, **J. Liu**, B. Runwal, J. Diffenderfer, B. Kailkhura, S. Liu, "SOUL: Unlocking the Power of Second-Order Optimization for LLM Unlearning," **EMNLP'24**. [[Link](#)]
9. **J. Liu***, **C. Fan***[†], A. Hero, S. Liu, "Challenging Forgets: Unveiling the Worst-Case Forget Sets in Machine Unlearning," **ECCV'24**. [[Link](#)]
10. Y. Zhang*, J. Jia*, X. Chen, A. Chen, Y. Zhang, **J. Liu**, K. Ding, S. Liu, "To Generate or Not? Safety-Driven Unlearned Diffusion Models Are Still Easy To Generate Unsafe Images ... For Now," **ECCV'24**. [[Link](#)]
11. **J. Liu***, **C. Fan***[†], Y. Zhang, E. Wong, D. Wei, S. Liu, "SalUn: Empowering Machine Unlearning via Gradient-based Weight Saliency in Both Image Classification and Generation," **ICLR'24 (Spotlight, acceptance rate ~5%)**. [[Link](#)]
12. **A. Chen***[†], Y. Zhang*, J. Jia, J. Diffenderfer, **J. Liu**, K. Parasyris, Y. Zhang, Z. Zhang, B. Kailkhura, S. Liu, "DeepZero: Scaling up Zeroth-Order Optimization for Deep Model Training," **ICLR'24**. [[Link](#)]
13. **J. Liu***, J. Jia*, P. Ram, Y. Yao, G. Liu, Y. Liu, P. Sharma, S. Liu, "Model Sparsity Can Simplify Machine Unlearning," **NeurIPS'23 (Spotlight, acceptance rate ~3%)**. [[Link](#)]
14. Y. Yao, X. Guo, V. Asnani, Y. Gong, **J. Liu**, X. Lin, X. Liu, S. Liu, "Reverse Engineering of Deceptions on Machine- and Human-Centric Attacks," **Fnt'24**. [[Link](#)]
15. **A. Chen***[†], Y. Zhang*, Y. Zhang*, J. Jia, **J. Liu**, G. Liu, M. Hong, S. Chang, S. Liu, "Selectivity Drives Productivity: Efficient Dataset Pruning for Enhanced Transfer Learning," **NeurIPS'23**. [[Link](#)]
16. X. Wu, R. L. Li, F. L. Zhang, **J. C. Liu**, J. Wang, A. Shamir, S. M. Hu, "Deep Portrait Image Completion and Extrapolation," **TIP'20**. [[Link](#)]
17. Y. Hu, **J. Liu***, A. Spielberg*, J. B. Tenenbaum, W. T. Freeman, J. Wu, D. Rus, W. Matusik, "ChainQueen: A Real-Time Differentiable Physical Simulator for Soft Robotics," **ICRA'19**. [[Link](#)]
18. X. Chen, E. Zhou, Y. Mo, **J. Liu**, Z. Cao, "Delving Deep Into Coarse-to-Fine Framework for Facial Landmark Localization," **CVPR'17 Workshop**. [[Link](#)]

HONORS & AWARDS

| | |
|---|------|
| Graduate Travel Fellowship, Michigan State University | 2024 |
| Scholar Award, NeurIPS | 2023 |
| Cloud Computing Fellowship, Michigan State University | 2023 |
| Xuetang Fellowship (top 3%), Tsinghua University | 2019 |
| Gold Medal (11th place/364 teams), ACM-ICPC Asian Regional Contest | 2017 |
| Gold Medal (7th place/228 teams), ACM-ICPC Asian Regional Contest | 2016 |
| Science and Technology Innovation Fellowship, Tsinghua University | 2016 |
| Gold Medal (4th place/245 teams), China Collegiate Programming Contest (CCPC) | 2015 |
| Outstanding Freshman Fellowship, Tsinghua University | 2015 |
| Gold Medal (13th place/61,176 individuals), National Olympiad in Informatics (NOI), China | 2014 |
| — <i>Fifth-place finalist among 50 awardees in the International Olympiad in Informatics (IOI) Chinese Team Selection</i> | |

ACADEMIC SERVICES

Organizer & Presenter Activities

| | |
|---|------|
| • NeurIPS Workshop — <i>3rd AdvML Frontiers</i> [Link] | 2024 |
| • CVPR Tutorial — <i>Machine Unlearning in Computer Vision</i> [Link] | 2024 |
| • ICML Workshop — <i>2nd AdvML Frontiers</i> [Link] | 2023 |
| • CVPR Tutorial — <i>Reverse Engineering of Deceptions</i> [Link] | 2023 |
| • Taichi Developer Conference — <i>Flexible Run-time Creation and Deletion of Taichi Fields</i> [Link] | 2021 |
| • USTC Summer Course — <i>Advances in Computer Graphics</i> [Link] | 2021 |

Committee Services

- **Conference Reviewer:** ICML, NeurIPS, ICLR, CVPR, ECCV, AISTATS, ICASSP, IJCNN, CPAL
- **Journal Reviewer:** IEEE TNNLS, IEEE TKDE, JCBC

COMMUNITY ENGAGEMENT

Mentorship

| | |
|--|---|
| Mr. Chongyu Fan <i>Undergraduate Student, Huazhong University of Science and Technology</i> — <i>Research topic: Machine Unlearning on foundation models.</i> [Link] — <i>Current graduate student at Michigan State University.</i> | Visiting Student, MSU 2023 - 2024 |
| Mr. Aochuan Chen <i>Undergraduate Student, Tsinghua University</i> — <i>Research topic: Efficient machine learning.</i> [Link] — <i>Current graduate student at The Hong Kong University of Science and Technology.</i> | Research Intern, MSU 2023 |
| Mr. Orion Smith <i>Computer Science teacher, East Lansing High School</i> — <i>Research Experiences for Teachers</i> program for K-12 educators. | Research Intern, MSU 2023 |
| Mr. Jiashi Li <i>Undergraduate Student, Peking University</i> — <i>Developing memory recycling system for Taichi Lang.</i> [Link] — <i>Current machine learning engineer at DeepSeek team of High-Flyer Inc.</i> | Software Engineer Intern, Taichi 2021 - 2022 |

Teaching

- Social, Ethical, and Professional Issues in Computing MSU CSE, 2025 Spring
- Advances in Computer Graphics USTC, 2021 Summer

Contributed Open-Source Projects

| | |
|---|---------------------------------------|
| • Universal Online Judge (uoj.ac) — <i>Online judging platform for competitive programming, 48k+ registered users.</i> | Co-Founder & Lead Front-End Developer |
| • Taichi Lang (taichi-dev/taichi) — <i>Programming language for high-performance computation, 28k+ GitHub stars.</i> | Core Contributor |
| • gpt-oss (openai/gpt-oss) — <i>Open-weight large language model repository by OpenAI, 19k+ GitHub stars.</i> | Core Contributor |