

Summary

Research Interests Machine Learning (Recommendation Systems, Learning to Rank, Algorithm Design, etc.), Differential Privacy, Computational Social Choice, and Quantum Computation

Experience 1+ years of industry experience in personalization, recommendation systems, LLM, theoretical computer science, differential privacy, and adversarial robustness

Academia Published 20+ papers, filed 2 patents, and received 250+ citations to date.
Reviewed 50+ papers for top AI conferences or journals.

Education

01/2018 – 05/2023 **Ph.D. in Computer Science**, Rensselaer Polytechnic Institute (RPI) Troy, NY
Thesis: *Group Decision Makings from Partial Preferences* [\[Link\]](#)

08/2015 – 05/2018 **M.Eng. in Material Engineering**, Rensselaer Polytechnic Institute (RPI) Troy, NY
Focus on optics and polymer physics, received Presidential Graduate Research Fellowship

08/2010 – 05/2014 **B.S. in Mathematics and Physics**, Tsinghua University Beijing, China
Minor in Computer Technology, and in Academic Talent Program

Experience

07/2023 – Present **Research Software Engineer at Google Core ML** Mountain View, CA
Design and Research More Efficient Transformers for Recommendation Systems

05/2022 – 08/2022 **Research Intern at Google Core ML** Mountain View, CA
Project: Unbiased Learning to Rank with a More Accurate Position Bias Estimator

- Significantly more accurate ($\sim 8\%$) on predictions *v.s.* state-of-the-art without any cost
- Proposed a novel probabilistic model to more accurately estimate position bias and designed an unbiased recommendation system based on it.
- Implemented the data pipeline and machine learning algorithm for the proposed unbiased recommendation system. Integrated all into Google codebase. (Python, TensorFlow, C/C++)

05/2019 – 08/2019 **Visiting Scholar at MIT-IBM Watson AI Lab** Cambridge, MA
and
Project: Certifiably Robust Interpretation via Rényi Differential Privacy

08/2018 – 12/2018

- Delivered one academic paper (on top AI journal and conference) and two patents.
- Significantly more robust ($\sim 12\%$) and more accurate *v.s.* state-of-the-art without any cost. Note that we improved two trade-off properties at the same time.
- Designed the first-ever algorithm with theoretically guaranteed top- k robustness against ℓ_∞ -norm attacks, where the theoretical proofs use Rényi differential privacy.
- Implemented the proposed robust algorithm and tested its robustness, accuracy, and computational efficiency on various tasks, including neural network interpretation, image classification, and objective detection. (PyTorch, TorchRay, TensorFlow, MATLAB)

Skills

Implementation Python, C/C++, MATLAB, TensorFlow, PyTorch, TorchRay, \LaTeX

Design / Theory Algorithm design, Time & Sample-complexity analysis, User-modeling, MCMC, Statistics, Privacy analysis, Robustness analysis, Model identifiability, Smoothed analysis

Review Services

Journal Information Sciences, TMLR, ACM ToIS, Sankhya B

Conference NeurIPS (20, 21, 22 & 23), ICML (22, 23 & 24), ICLR (23 & 24), AAAI (21 & 22), IJCAI-22

Selected Publications

- TMLR* **Smoothed Differential Privacy** [PDF]
Ao Liu, Yu-Xiang Wang, and Lirong Xia
- UAI-23* **Accelerating Voting by Quantum Computation** [PDF]
Ao Liu, Qishen Han, Lirong Xia, and Nengkun Yu
- AIJ* and **Certifiably Robust Interpretation via Rényi Differential Privacy** [Link] [ArXiv]
AAAI-23 (oral) *Ao Liu*, Xiaoyu Chen, Sijia Liu, Lirong Xia, and Chuang Gan
- AAAI-22* **The Semi-Random Likelihood of Doctrinal Paradoxes** [PDF]
Ao Liu and Lirong Xia
- UAI-20* (oral) **How Private Are Commonly-Used Voting Rules?** [PDF]
Ao Liu, Yun Lu, Lirong Xia, and Vassilis Zikas
- AAAI-19* (oral) **Near-Neighbor Methods in Random Preference Completion** [PDF]
Ao Liu, Qiong Wu, Zhenming Liu, and Lirong Xia
- AAAI-19* (oral) **Learning Plackett-Luce Mixture from Partial Preferences** [PDF]
Ao Liu, Zhibing Zhao, Chao Liao, Pinyan Lu, and Lirong Xia
- ETRA-19* (oral) **Differential Privacy for Eye-Tracking Data** [PDF]
Ao Liu, Lirong Xia, Andrew Duchowski, Reynold Bailey, Kenneth Holmqvist, and Eakta Jain
- AAAI-23* (oral) **Differentially Private Condorcet Voting** [PDF]
Zhechen Li, *Ao Liu*, Lirong Xia, Yongzhi Cao, and Hanpin Wang
- IJCAI-22* (oral) **Learning Mixtures of Random Utility Models with Features from Incomplete Preferences**
Zhibing Zhao, *Ao Liu*, and Lirong Xia [PDF]
- JAIR* and **Learning to Design Fair and Private Voting Rules** [PDF]
IJCAI-23 (oral) Farhad Mohsin, *Ao Liu*, Pin-Yu Chen, Francesca Rossi, and Lirong Xia
- ETRA-20 Adjunct* **Let It Snow: Adding Pixel Noise to Protect the Users Identity** [Link]
Brendan John, *Ao Liu*, Lirong Xia, Sanjeev Koppal, and Eakta Jain
- J. Polym. Sci. B: Polymer Physics* **Simulation of Pulse Responses of Lithium Salt-Doped Poly-Ethyleneoxide** [Link]
Ao Liu, F. Zeng, Y. Hu, S. Lu, W. Dong, X. Li, C. Chang, and D. Guo
- *View All Publications* at [Personal Website] [Google Scholar]

Patents

- US Patent **Certifiably Robust Interpretation** [PDF]
Ao Liu, Sijia Liu, Bo Wu, Lirong Xia, Qi Cheng Li, and Chuang Gan
- US Patent **Interpretation Maps with Guaranteed Robustness** [PDF]
Ao Liu, Sijia Liu, Abhishek Bhandwadar, Chuang Gan, Lirong Xia, and Qi Cheng Li

Awards and Teaching

- 09/2019–05/2022 **RPI-IBM AI Horizon Scholarship**
- 09/2016–05/2017 **RPI Presidential Graduate Research Fellowship** [Certificate]
- 01/2023–05/2023 **Teaching Assistant** of CSCI 4150: Introduction to AI
- 08/2017–12/2017 **Teaching Assistant** of MATH 1020: Calculus II
- 04/2021 **Guest Lecture** at CSCI 4967/6967: Economics and Computation