

PROFESSIONAL EXPERIENCE

Assistant Professor in Statistics and Operations Research University of North Carolina at Chapel Hill	Chapel Hill, NC, USA 2025.7-present
Visiting Researcher in Computer Science University of Chicago Host: Dr. Haifeng Xu	Chicago, IL, USA 2023.5-2025.6
Student Researcher, Intern Meta Research, Google Research, TikTok AML Lab	San Francisco Bay Area, CA, USA 2021.5-2024.7
Full-time Algorithm Engineer Alibaba Group, Taobao Recommendation Team	Beijing, China 2017.8-2019.7

EDUCATION

University of Virginia Ph.D. in Computer Science, Co-advised by Dr. Haifeng Xu and Dr. Hongning Wang – Thesis: Towards a Responsible Platform Economy – Strategic Behavior, Learning Dynamics, and Social Good	Charlottesville, VA, USA 2019.8 –2025.5
Peking University M.S. in Computational Mathematics, Advisor: Dr. Tiejun Li B.S. in Information and Computing Science, double major in Philosophy	Beijing, China 2013.9–2016.7 2009.9–2013.7

HONORS AND AWARDS

- UVA Copenhaver Bicentennial Graduate Research Award (\$12k). 2024
- UVA Graduate Teaching Award. 2021
- Outstanding Graduate Student Award of Peking University. 2016
- Bronze Medalist in Team Contest of Applied and Computational Mathematics, Shing-Tung Yau College Student Mathematics Contests (Ranked top 6 nation-wide). 2012
- Gold Medalist in Chinese Mathematics Olympics (Ranked top 40 nation-wide). 2009

WORKING PAPERS

*Equal contribution; authors listed in alphabetical order.

6. *Y. Chen, *Y. He, *M. Jordan, ***F. Yao**, “How Sampling Shapes LLM Alignment: From One-Shot Optima to Iterative Dynamics”, preprint at arXiv:2602.12180.
5. L. Luo, **F. Yao**, E. Manzoor, N. Yang, “Platform Design When Creators Train Their AI Substitutes”, *under review*.
4. Y. Wu, J. Tang, J. Liu, H. Xu, **F. Yao**, “Do AI Overviews Benefit Search Engines? An Ecosystem Perspective”, preprint at arXiv:2601.22493.

3. B. Shen, Z. Cheng, N. Gong, **F. Yao**, Y. Dong, “CREDIT: Certified Defense of Deep Neural Networks against Model Extraction Attacks”, *under review*.
2. *S. Ahmadi, *A. Blum, *H. Xu, ***F. Yao**, “Strategic Filtering for Content Moderation: Free Speech or Free of Distortion?”, *under review*.
1. Y. He, **F. Yao**, Y. Yu, X. Qiu, M. Li, H. Xu, “The Complexity of Tullock Contests”, preprint at arXiv:2412.06444.

JOURNAL ARTICLES

- J3. *J. Wu, *H. Xu, and ***F. Yao**, “Uncoupled Bandit Learning towards Rationalizability: Benchmarks, Barriers, and Algorithms”, *under major revision at JMLR*. (Supersedes C3.)
- J2. *R. Sundaram, *A. Vullikanti, *H. Xu, and ***F. Yao**, “Pac-Learning for Strategic Classification”, *Journal of Machine Learning Research, JMLR*, 2023. (Supersedes C1.)
- J1. **F. Yao**, F. Li, T. Li, “Mean Field Study of a Propagation-Turnover Lattice Model for the Dynamics of Histone Marking”, *Science China Physics, Mechanics & Astronomy* 60, 1-15, 2017.

CONFERENCE PUBLICATIONS

- C14. H. Kiyohara, **F. Yao**, S. Dean, “Policy Design for Two-sided Platforms with Participation Dynamics”, *ICML*, 2025.
- C13. Y. Yu, **F. Yao**, S. Pan, “Beyond Self-Interest: How Group Strategies Reshape Content Creation in Recommendation Systems?”, *ICML*, 2025.
- C12. ***F. Yao**, *Y. Cheng, E. Wei, and H. Xu, “Single-Agent Poisoning Attacks Suffice to Ruin Multi-Agent Learning”, *ICLR*, 2025.
- C11. Y. Cheng, **F. Yao**, X. Liu, and H. Xu, “Learning from Imperfect Human Feedback: a Tale from Corruption-Robust Dueling”, *ICLR*, 2025.
Selected for **NSF poster awards** (15 out of 116 accepted papers) at the Midwest Machine Learning Symposium, 2024.
- C10. **F. Yao**, Y. Liao, J. Liu, S. Nie, Q. Wang, H. Xu, H. Wang, “Unveiling User Satisfaction and Creator Productivity Trade-Offs in Recommendation Platforms”, *Neurips*, 2024.
- C9. **F. Yao**, Y. Liao, M. Wu, C. Li, Y. Zhu, J. Yang, J. Liu, Q. Wang, H. Xu, and H. Wang, “User Welfare Optimization in Recommender Systems with Competing Content Creators”, *KDD*, 2024.
- C8. **F. Yao**, C. Li, D. Nekipelov, H. Wang, and H. Xu, “Human vs. Generative AI in Content Creation Competition: Symbiosis or Conflict?”, *ICML*, 2024.
- C7. **F. Yao**, C. Li, K. Sankararaman, Y. Liao, Y. Zhu, Q. Wang, H. Wang, and H. Xu, “Rethinking Incentives in Recommender Systems: Are Monotone Rewards Always Beneficial?”, *Neurips*, 2023.
- C6. **F. Yao**, C. Li, D. Nekipelov, H. Wang, and H. Xu, “How Bad is Top- K Recommendation under Competing Content Creators?”, *ICML*, **Oral presentation**, 2023.
- C5. M. Wu, **F. Yao**, and H. Wang, “An End-to-End Solution for Spatial Inference in Smart Buildings”, *BuildSys*, **Best Paper Nomination**, 2023.

C4. **F. Yao**, C. Li, D. Nekipelov, H. Wang, and H. Xu, “Learning from a Learning User for Optimal Recommendations”, **ICML**, 2022.

Also selected for **spotlight presentation** (5 out of 38 accepted posters) at the ICML 2023 Workshop on Interactive Learning with Implicit Human Feedback.

C3. *J. Wu, *H. Xu, and ***F. Yao**, “Multi-Agent Learning for Iterative Dominance Elimination: Formal Barriers and New Algorithms”, **COLT**, 2022.

C2. **F. Yao**, C. Li, D. Nekipelov, H. Wang, and H. Xu, “Learning the Optimal Recommendation from Revealed Preferences”, **AAAI**, 2022.

C1. *R. Sundaram, *A. Vullikanti, *H. Xu, and ***F. Yao**, “Pac-Learning for Strategic Classification”, **ICML**, **Oral presentation**, 2021.

INVITED TALKS

- Northwestern University, IDEAL workshop on LLMs and Strategic Agents, “How Sampling Shapes Preference Alignment”. 2026.3
- HKUST Guangzhou, FTEC seminar, “How Sampling Shapes Preference Alignment”. 2026.2
- INFORMS Annual Meeting, Atlanta, “Rethinking the role of GenAI in content creation ecosystem”. 2025.10
- UNC Chapel Hill, Aligning AI with Society Workshop – invited Panel Discussion. 2025.9
- UChicago, Data Science Research Day, “Steering Multi-Agent Learning via Single-Agent Poisoning”. 2025.5
- UNC Chapel Hill, “Towards a Sustainable Content Creation Ecosystem: From Theory to Practice”. 2025.1
- The Hong Kong University of Science and Technology (Guangzhou), “Understanding and Optimizing Multi-Agent Content Ecosystems”. 2025.1
- Carnegie Mellon University, guest lecture, “Modeling Competition-Driven Content Ecosystems”. 2024.11
- INFORMS Annual Meeting, Seattle, “Optimizing Competition-Driven Content Ecosystems”. 2024.10
- Cornell University, ESIF Economics and AI+ML Meeting, “Human v.s. GenAI Competition”. 2024.8
- George Mason University, seminar talk, “Understanding Competition-Driven Content Ecosystems”. 2024.6
- Northwestern University, Midwest Workshop on Control and Game Theory, “Understanding Competition-Driven Content Ecosystems”. 2024.4
- Mila & Vector Institute, seminar talk, “Understanding Competition-Driven Content Ecosystems”. 2024.4
- Cornell University, seminar talk, “Understanding Competition-Driven Content Ecosystems”. 2024.2
- Meta Research, “How Bad is Top-K Recommendation under Competing Content Creators?”. 2023.8
- Uber Research, “Learning from a Learning User for Optimal Recommendations”. 2022.6

TEACHING

- **STOR 565** at UNC Chapel Hill Spring 2026
Machine Learning
- **STOR 323** at UNC Chapel Hill Fall 2025
Ethics of Machine Learning in Society

PROFESSIONAL SERVICE

- **Workshop (Co-)organizer:** Collaboration and Evolution of Foundation and Specialized Models, ICMR 2024
- **Program Committee:** ICML (2022-2026), ICLR (2023-2026), Neurips (2022-2025), AAAI (2021-2025), IJCAI (2021-2024), KDD (2021-2024)
- **Journal Reviewing Activities:** INFORMS Journal on Computing, ACM Transactions on Economics and Computation, Journal of Computational and Graphical Statistics (JCGS), Autonomous Agents and Multi-Agent Systems (JAAMAS), Operational Research, IEEE/ACM Transactions on Networking, Transactions on Intelligent Systems and Technology, etc.