

Ziping Xu

Assistant Professor
School of Data and Information Sciences
University of North Carolina at Chapel Hill

(734)-276-8463
✉ zipingxu@fas.harvard.edu
📄 <https://zipingxu.github.io>

Current Position

2025 – Now **Assistant Professor, University of North Carolina, Chapel Hill**, (From July 1st, 2025).
School of Data and Information Sciences

Education

2018–2023 **Ph.D. in Statistics, University of Michigan**, Ann Arbor, MI.
Advisor: Prof. Ambuj Tewari
Thesis: On the Benefits of Multitask Learning: A Perspective Based on Task Diversity

2014–2018 **B.S. in Data Science, Peking University**, Beijing, China.
Advisor: Prof. Song Xi Chen
Thesis: Meteorological Change and Impacts on Air Pollution: Results from North China

Research Interest

My primary research interests lie in sequential decision-making in health applications, with a focus on mobile health through digital interventions. I work on **theories and methodologies** in sequential decision-making algorithms, such as Reinforcement Learning (RL), to improve their understanding and applicability in health problems. I am also actively involved in implementing RL in **real mobile health products and clinical trials**.

Publications and Preprints

(* denotes equal contribution)

- 2026 Yongyi Guo and **Ziping Xu**. Statistical inference for misspecified contextual bandits. *arXiv preprint*, 2026.
- 2026 S. Ghosh, B. T. Gullapalli, Daiqi Gao, Asim H. Gazi, Anna L. Trella, **Ziping Xu**, Kelly Zhang, et al. Reproducible workflow for online artificial intelligence in digital health. *Philosophical Transactions of the Royal Society A*, 2026.
- 2026 Asim H. Gazi, S. Ghosh, Yongyi Guo, Daiqi Gao, **Ziping Xu**, Inbal Nahum-Shani, and Susan Murphy. Key concepts in online learning and decision making for personalizing just-in-time adaptive interventions. *Annals of Behavioral Medicine*, 2026.
- 2026 Asim H. Gazi, Daiqi Gao, S. Ghosh, **Ziping Xu**, Anna L. Trella, Predrag Klasnja, and Susan A. Murphy. Digital twins for just-in-time adaptive interventions (JITAs): Framework for optimizing and continually improving JITAs. *Journal of Medical Internet Research*, 2026.
- 2026 Daiqi Gao, **Ziping Xu**, Aseel Rawashdeh, Predrag Klasnja, and Susan Murphy. Active measuring in reinforcement learning with delayed negative effects. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2026.
- 2025 Lutong Zou, **Ziping Xu**, Daiqi Gao, and Susan Murphy. Causal directed acyclic graph-informed reward design. In *The Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*, 2025.

- 2025 Anna L Trella, Walter Dempsey, Asim H Gazi, **Ziping Xu**, Finale Doshi-Velez, and Susan A Murphy. Non-stationary latent auto-regressive bandits. *Reinforcement Learning Journal*, 2025.
- 2025 **Ziping Xu**, Hinal Jajal, Sung Won Choi, Inbal Nahum-Shani, Guy Shani, Alexandra Psihogios, Pei-Yao Hung, and Susan Murphy. Reinforcement learning on aya dyads to enhance medication adherence. In *International Conference on Artificial Intelligence in Medicine (AIME)*, 2025.
- 2024 **Ziping Xu**, Iris Yan, and Susan Murphy. An adaptation of RLSVI with explicit action sampling probabilities. In *RLC 2024 Deployable RL Workshop*, 2024.
- 2024 **Ziping Xu**, Zifan Xu, Runxuan Jiang, Peter Stone, and Ambuj Tewari. Sample efficient myopic exploration through multitask Reinforcement Learning with diverse tasks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024.
- 2024 **Ziping Xu**, Kelly Zhang, and Susan Murphy. The fallacy of minimizing cumulative regret in the sequential task setting. *arXiv*, 2024.
- 2024 Kevin Tan* and **Ziping Xu***. A natural extension to online algorithms for hybrid RL with limited coverage. *Reinforcement Learning Journal*, volume 1, 2024.
- 2024 Yongyi Guo*, **Ziping Xu***, and Susan Murphy. Online learning in bandits with predicted context. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2024.
- 2024 Quanfu Fan, Yilai Li, Yuguang Yao, John Cohn, Sijia Liu, **Ziping Xu**, Seychelle Vos, and Michael Cianfrocco. CryoRL: Reinforcement Learning enables efficient cryo-EM data collection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 7892–7902, 2024.
- 2023 **Ziping Xu**, Quanfu Fan, Yilai Li, Emma Rose Lee, John Maxwell Cohn, Ambuj Tewari, Seychelle M Vos, and Michael Cianfrocco. Coupling semi-supervised learning with reinforcement learning for better decision making—an application to cryo-EM data collection. In *NeurIPS 2023 AI for Science Workshop*, 2023.
- 2022 **Ziping Xu** and Ambuj Tewari. On the statistical benefits of curriculum learning. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*. PMLR, 2022.
- 2022 **Ziping Xu**, Eunjae Shim, Ambuj Tewari, and Paul Zimmerman. Adaptive sampling for discovery. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, 2022.
- 2022 Hao Sun, **Ziping Xu**, Taiyi Wang, Meng Fang, and Bolei Zhou. Supervised q-learning for continuous control. In *Deep Reinforcement Learning Workshop NeurIPS 2022*, 2022.
- 2022 Hao Sun, **Ziping Xu**, Zhenghao Peng, Meng Fang, Bo Dai, and Bolei Zhou. Mopa: a minimalist off-policy approach to safe-RL. In *Deep Reinforcement Learning Workshop NeurIPS 2022*, 2022.
- 2022 Eunjae Shim, Joshua A Kammeraad, **Ziping Xu**, Ambuj Tewari, Tim Cernak, and Paul M Zimmerman. Predicting reaction conditions from limited data through active transfer learning. *Chemical Science*. Royal Society of Chemistry, 2022.
- 2022 Aditya Modi, **Ziping Xu**, Mohamad KS Faradonbeh, and Ambuj Tewari. Big control actions help multitask learning of unstable linear systems. In *ICML 2022 Complex feedback in online learning Workshop*, 2022.
- 2022 Yilai Li, Quanfu Fan, **Ziping Xu**, Emma Rose Lee, John Cohn, Veronique Demers, Ja Young Lee, Lucy Yip, Michael A. Cianfrocco, and Seychelle M. Vos. Optimized path planning surpasses human efficiency in cryo-EM imaging. *eLife*. Cold Spring Harbor Laboratory, 2022.
- 2021 **Ziping Xu** and Ambuj Tewari. Representation learning beyond linear prediction functions. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, 2021.
- 2021 **Ziping Xu**, Amirhossein Meisami, and Ambuj Tewari. Decision making problems with funnel structure: A multi-task learning approach with application to email marketing campaigns. In *In Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 127–135. PMLR, 2021.

- 2021 Yangyi Lu*, **Ziping Xu***, and Ambuj Tewari. Bandit algorithms for precision medicine. *arXiv preprint arXiv:2108.04782*, 2021.
- 2020 **Ziping Xu** and Ambuj Tewari. Reinforcement learning in factored MDPs: Oracle-efficient algorithms and tighter regret bounds for the non-episodic setting. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, 2020.
- 2020 **Ziping Xu**, Song Xi Chen, and Xiaoqing Wu. Meteorological change and impacts on air pollution: Results from north China. *Journal of Geophysical Research: Atmospheres*, volume 125, page e2020JD032423. Wiley Online Library, 2020.
- 2020 Tarun Gogineni, **Ziping Xu**, Exequiel Punzalan, Runxuan Jiang, Joshua Kammeraad, Ambuj Tewari, and Paul Zimmerman. Torsionnet: A Reinforcement Learning approach to sequential conformer search. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- 2019 **Ziping Xu** and Ambuj Tewari. Worst-case regret bound for perturbation based exploration. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- 2018 Wenling Liu, **Ziping Xu**, and Tianan Yang. Health effects of air pollution in china. *International journal of environmental research and public health*, volume 15, page 1471. MDPI, 2018.
- 2017 Shuyi Zhang, Bin Guo, Anlan Dong, Jing He, **Ziping Xu**, and Song Xi Chen. Cautionary tales on air-quality improvement in Beijing. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, volume 473, page 20170457. The Royal Society Publishing, 2017.

Fellowships & Awards

- 2026 **NSF Travel Award** (2026 ICSA Applied Statistics Symposium)
- 2026 **NSF Travel Award** (9th International Workshop in Sequential Methodologies, IWSM)
- 2023 **Departmental Outstanding Dissertation Award** (Awarded for outstanding dissertation writing)
- 2022 **Rackham Predoctoral Fellowship** (Awarded to outstanding doctoral students)
- 2020 **Outstanding Graduate Student Instructor Awards** (Awarded to outstanding instructors)
- 2016 **Lee Wai Wing Scholarship** (Awarded to outstanding first year students)

Funding

- 2026 Principal Investigator (UNC–Chapel Hill; lead), National Science Foundation, FDT-BioTech program. “Collaborative Research: FDT-BioTech: Mechanistic-Neural Foundations for Digital Twins of Parkinson’s Disease Symptom Dynamics Under Levodopa Treatment.” \$200,000 (\$800,000 total collaborative), Sept. 2026–Aug. 2029.

Invited Talks

- June 2026 “A Diffusion-Model Subpopulation Digital Twin for Mobile Health Deployment: A Case Study on the HeartSteps Intervention”, 2026 ICSA Applied Statistics Symposium
- June 2026 “The Fallacy of Minimizing Cumulative Regret in the Sequential Task Setting”, 9th International Workshop in Sequential Methodologies (IWSM), American University, Washington, DC

Teaching Experiences

Teaching Fellow

- | | | |
|-------------|--|---|
| Winter 2026 | DATA 521, Foundations of AI | University of North Carolina at Chapel Hill |
| Winter 2024 | STATS 234, Sequential Decision Making | Harvard University |
| | Graduate Student Instructor | |
| Winter 2020 | STATS 503, Statistical Learning II | University of Michigan |
| Fall 2019 | STATS 415, Introduction to Data Mining | University of Michigan |
| Winter 2019 | STATS 250, Introduction to Statistics and Data Analysis | University of Michigan |

Mentoring Experience2024–Now **Lutong Zou**

Topic: Reward Design through Causal DAG

2023–2024 **Iris Yan**

Topic: Approximate Sampling Probability for Online Decision Making

2022 **Zhiyu Yuan** (Next Step: CMU, MS in Computer Engineer)

Topic: Curriculum Learning for Reinforcement Learning

2022 **Ruanxuan Jiang** (Next Step: Software Engineer Citadel)

Topic: Theory of Multitask Reinforcement Learning

2020–2021 **Anh Tuan (Alan) Tran** (Next Step: University of Toronto, MS in Computer Science)

Topic: Sample Efficient Reinforcement Learning

Professional Experience2023 – now Postdoctoral Research Fellow at **Department of Statistics, Harvard University**

Advisor: Dr. Susan A. Murphy, Professor, Department of Statistics and Computer Science

Summer 2022 Research Intern at **MIT-IBM Watson AI Lab**

Advisor: Dr. Quanfu Fan

Summer 2020 Research Intern at **Adobe Inc.**

Advisor: Dr. Amirhossein Meisami

Professional Activities

Organizing Committee

2021 Representative Michigan Student Symposium for Interdisciplinary Statistical Sciences

Journal Paper Reviewer

2026 Science

2026 ACM Health

2026 Stochastic Systems

2023 Statistical Science

2023 Proceedings of the National Academy of Sciences of the United States of America (PNAS)

2022 Journal of the American Statistical Association (JASA)

Conference Paper Reviewer

2020 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)

2021, 2022 International Conference on Artificial Intelligence and Statistics (AISTATS)

2021 International Conference on Algorithmic Learning Theory (ALT)

2023 International Conference on Learning Representations (ICLR)

2022–2024 International Conference on Machine Learning (ICML)

2021–2026 Conference on Neural Information Processing Systems (NeurIPS)

2026 Reinforcement Learning Conference (RLC) & Reinforcement Learning Journal (RLJ)

2026 European Workshop on Reinforcement Learning (EWRL)