

TIAN XU

Address: 163 Xianlin Avenue, Nanjing, China

E-mail: xut@lamda.nju.edu.cn

Website: www.lamda.nju.edu.cn/xut/

Education

Nanjing University

Ph.D. in Computer Science, School of Artificial Intelligence, Advisor: Yang Yu.

Sep. 2019 – Present

Nanjing, China

Northwestern Polytechnical University

B.S. in Mathematics and Applied Mathematics, School of Natural and Applied Sciences.

Sep. 2015 – June 2019

Xi'an, China

GPA: 92.1/100

Research Interests

Reinforcement learning theory and algorithms.

Experience

The Chinese University of Hong Kong, Shenzhen

Visiting Student, Advisor: Zhi-Quan (Tom) Luo.

Jun. 2021 – Sep. 2021

Shenzhen, China

Publications and preprints

* indicates equal contribution.

- Ziniu Li, Congliang Chen, **Tian Xu**, Zeyu Qin, Jiancong Xiao, Ruoyu Sun, Zhi-Quan Luo. Preserving Diversity in Supervised Fine-Tuning of Large Language Models. In Proceedings of the 13th International Conference on Learning Representations (ICLR), 2025.
- **Tian Xu***, Zhilong Zhang*, Ruishuo Chen, Yihao Sun, Yang Yu. Provably and Practically Efficient Adversarial Imitation Learning with General Function Approximation. In *Advances in Neural Information Processing Systems 38 (NeurIPS)*, 2024.
- Ziniu Li, **Tian Xu**, Yushun Zhang, Zhihang Lin, Yang Yu, Ruoyu Sun, Zhi-Quan Luo. ReMax: A Simple, Effective, and Efficient Reinforcement Learning Method for Aligning Large Language Models. In *Proceedings of the Forty-first International Conference on Machine Learning (ICML)*, 2024.
- Ziniu Li*, **Tian Xu***, Yang Yu. Policy Optimization in RLHF: The Impact of Out-of-preference Data. In *Tiny Paper Track of the 12th International Conference on Learning Representations (ICLR)*, 2024.
- Fan-Ming Luo, **Tian Xu**, Xingchen Cao, Yang Yu. Reward-Consistent Dynamics Models are Strongly Generalizable for Offline Reinforcement Learning. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*, 2024. **Spotlight (acceptance rate < 5%)**
- Ziniu Li*, **Tian Xu***, Yang Yu, Zhi-Quan Luo. Data Selection in Imitation Learning: Theoretical Justifications and Algorithms. In *Advances in Neural Information Process System 36 (NeurIPS)*, 2023. **Spotlight (acceptance rate < 5%)**
- Chengxing Jia*, Fuxiang Zhang*, **Tian Xu**, Jing-Cheng Pang, Zongzhang Zhang, Yang Yu. Model Gradient: Unified Model and Policy Learning in Model-based Reinforcement Learning. *Frontiers of Computer Science*, 2023.
- **Tian Xu***, Ziniu Li*, Yang Yu, Zhi-Quan Luo. Provably Efficient Adversarial Imitation Learning with Unknown Transitions. In *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence (UAI)*, 2023. **Oral Presentation (acceptance rate 16/778 = 2.0%)**
- Fan-Ming Luo, **Tian Xu**, Hang Lai, Xiong-Hui Chen, Weinan Zhang, Yang Yu. A Survey on Model-based Reinforcement Learning. *SCIENCE CHINA Information Sciences*, 2023.
- **Tian Xu***, Ziniu Li*, Yang Yu, Zhi-Quan Luo. Understanding Adversarial Imitation Learning in Small Sample Regime: A Stage-coupled Analysis. Submitted to *TPAMI*. *arXiv*: 2208.01899, 2022.
- Ziniu Li*, **Tian Xu***, Yang Yu, Zhi-Quan Luo. Rethinking ValueDice: Does It Really Improve Performance? In *Proceedings of the 10th International Conference on Learning Representations (ICLR) (Blog Track)*, 2022.
- Ziniu Li*, **Tian Xu***, Yang Yu. A Note on Target Q-learning For Finite MDPs with a Generative Oracle. Submitted to *Artificial Intelligence*. *arXiv*: 2203.11489, 2022.

- Jing-Cheng Pang, **Tian Xu**, Shengyi Jiang, Yu-Ren Liu, Yang Yu. Reinforcement Learning With Sparse-Executing Actions via Sparsity Regularization. Submitted to *IEEE Transactions on Neural Networks and Learning Systems (TNNLS)*. *arXiv*: 2105.08666, 2022.
- **Tian Xu***, Ziniu Li*, Yang Yu. More Efficient Adversarial Imitation Learning Algorithms with Known and Unknown Transitions. Ecological Theory of RL Workshop in NeurIPS 2021. *arXiv*: 2106.10424v2, 2021.
- **Tian Xu**, Ziniu Li, Yang Yu. Error Bounds of Imitating Policies and Environments for Reinforcement Learning. In *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2021.
- **Tian Xu**, Ziniu Li, Yang Yu. Error Bounds of Imitating Policies and Environments. In *Advances in Neural Information Processing System 33 (NeurIPS)*, 2020.

Talks

- **A Deeper Look at BC and Adversarial based Methods**. DAI 2020, Imitation Learning: Single & Multi-Agent Workshop.

Service

- Conference Reviewer: NeurIPS (2022-), ICML (2022-), ICLR (2024-), AISTATS (2024-), UAI (2021-), RLC (2024-), EWRL (2022-).
- Journal Reviewer: IEEE Transactions on Emerging Topics in Computational Intelligence (TETCI).

Awards

- 2nd winner on KDD CUP 2020 Reinforcement Learning Competition Track.
- National Scholarship for Graduates, 2020.
- National Scholarship for Undergraduates, 2016, 2017.
- The Mathematical Contest in Modeling, M Prize, 2017.