

# Yao Liu

yao.liu.chn@gmail.com ♦ <http://www.yao-liu.com>

## RESEARCH AREAS

---

Reinforcement Learning, Large Language Models

## WORK EXPERIENCE

---

### Senior Applied Scientist

Aug 2022 - Now

*Amazon (AWS AI → Amazon Store Foundational AI)*

*Seattle, WA*

- RLHF training of LLMs for shopping (Amazon Rufus).
- Built large-scale RLHF training algorithm and infrastructure and launched Amazon's first RLHF-trained LLM as a core contributor (Amazon Titan).
- Foundational research: reinforcement learning, post-training of LLMs, and LLM-based agents.

### Research Scientist

Aug 2021 - Aug 2022

*ByteDance*

*Bellevue, WA*

- I led several launches of a bandit-based retrieval algorithm in TikTok and Douyin (Chinese version of TikTok) video and ads recommendation.

## EDUCATION

---

### Ph.D. in Computer Science, Stanford University

*June 2021*

Dissertation: Adaptive and Efficient Batch Reinforcement Learning Algorithms

Committee: Emma Brunskill (Advisor), Tengyu Ma, Stefan Wager, Benjamin Van Roy

### B.S. in Machine Intelligence, Peking University

*July 2016*

GPA: 3.75/4.0, Rank: 1/46, Summa Cum Laude.

## PUBLICATIONS

---

1. Siliang Zeng, **Yao Liu**, Huzefa Rangwala, George Karypis, Mingyi Hong, Rasool Fakoor  
From Demonstrations to Rewards: Alignment Without Explicit Human Preferences  
In Submission
2. Zhepeng Cen, **Yao Liu**, Siliang Zeng, Pratik Chaudhar, Huzefa Rangwala, George Karypis, Rasool Fakoor  
Bridging the Training-Inference Gap in LLMs by Leveraging Self-Generated Tokens  
TMLR
3. Ke Yang, **Yao Liu**, Sapana Chaudhary, Rasool Fakoor, Pratik Chaudhari, George Karypis, Huzefa Rangwala  
AgentOccam: A Simple Yet Strong Baseline for LLM-Based Web Agents  
ICLR 2025
4. Jesse Zhang, Minh Heo, Zuxin Liu, Erdem Biyik, Joseph J Lim, **Yao Liu**, Rasool Fakoor  
EXTRACT: Efficient Policy Learning by Extracting Transferrable Robot Skills from Offline Data  
CoRL 2024
5. Kavosh Asadi\*, **Yao Liu**\*, Shoham Sabach\*, Yin Ming\*, Rasool Fakoor  
Learning the Target Network in Function Space  
ICML 2024

6. Zuxin Liu, Jesse Zhang, Kavosh Asadi, **Yao Liu**, Ding Zhao, Shoham Sabach, Rasool Fakoor  
TAIL: Task-specific Adapters for Imitation Learning with Large Pretrained Models  
(ICLR 2024) 2024 International Conference on Learning Representations
7. **Yao Liu**, Pratik Chaudhari, Rasool Fakoor  
Budgeting Counterfactual for Offline RL  
(NeurIPS 2023) In Advances in Neural Information Processing Systems 36
8. Kavosh Asadi, Shoham Sabach, **Yao Liu**, Omer Gottesman, Rasool Fakoor  
TD Convergence: An Optimization Perspective  
(NeurIPS 2023) In Advances in Neural Information Processing Systems 36
9. Sherry Ruan, Allen Nie, William Steenbergen, Jiayu He, JQ Zhang, Meng Guo, **Yao Liu**, Kyle Dang Nguyen, Catherine Wang, Rui Ying, James Landay, Emma Brunskill  
Reinforcement learning tutor better supported lower performers in a math task  
In Machine Learning Journal
10. **Yao Liu**, Dipendra Misra, Miro Dudik, Robert E. Schapire  
Provably Sample-Efficient RL with Side Information about Latent Dynamics  
(NeurIPS 2022) In Advances in Neural Information Processing Systems 35
11. **Yao Liu**, Yannis Flet-Berliac, Emma Brunskill  
Offline Policy Optimization with Eligible Actions  
(UAI 2022) In Proceedings of the 38th Conference on Uncertainty in Artificial Intelligence
12. **Yao Liu**, Adith Swaminathan, Alekh Agarwal, Emma Brunskill  
Provably Good Batch Off-Policy Reinforcement Learning Without Great Exploration  
(NeurIPS 2020) In Advances in Neural Information Processing Systems 33
13. **Yao Liu**, Pierre-Luc Bacon, Emma Brunskill  
Understanding the Curse of Horizon in Off-Policy Evaluation via Conditional Importance Sampling  
(ICML 2020) In Proceedings of the 37th International Conference on Machine Learning
14. Omer Gottesman, Joseph Futoma, **Yao Liu**, Sonali Parbhoo, Leo Anthony Celi, Emma Brunskill, Finale Doshi-Velez  
Interpretable Off-Policy Evaluation in Reinforcement Learning by Highlighting Influential Transitions  
(ICML 2020) In Proceedings of the 37th International Conference on Machine Learning
15. **Yao Liu**, Adith Swaminathan, Alekh Agarwal, Emma Brunskill  
Off-Policy Policy Gradient with State Distribution Correction  
(UAI 2019, **Oral**) In Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence
16. Omer Gottesman, **Yao Liu**, Scott Sussex, Emma Brunskill, Finale Doshi-Velez  
Combining Parametric and Nonparametric Models for Off-Policy Evaluation  
(ICML 2019, **Oral**) In Proceedings of the 36th International Conference on Machine Learning
17. **Yao Liu**, Omer Gottesman, Aniruddh Raghu, Matthieu Komorowski, Aldo Faisal, Finale Doshi-Velez, Emma Brunskill  
Representation Balancing MDPs for Off-Policy Policy Evaluation  
(NeurIPS 2018) Advances in Neural Information Processing Systems 31
18. **Yao Liu**, Emma Brunskill  
When Simple Exploration is Sample Efficient: Identifying Sufficient Conditions for Random Exploration to Yield PAC RL Algorithms  
(EWRL 2018) In the 14th European Workshop on Reinforcement Learning
19. **Yao Liu**, Philip S. Thomas, Emma Brunskill

Model Selection for Off-Policy Policy Evaluation  
(RLDM 2017) In the 3rd Multidisciplinary Conference on Reinforcement Learning and Decision Making, 2017. Extended abstract.

20. **Yao Liu**, Zhaohan Daniel Guo, Emma Brunskill  
PAC Continuous State Online Multitask Reinforcement Learning with Identification  
(AAMAS 2016) In Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems
21. Tong Lin, **Yao Liu**, Bo Wang, Liwei Wang, Hongbin Zha  
Local Orthogonality Preserving Alignment for Nonlinear Dimensionality Reduction  
Journal of Computer Science and Technology, 2016, Vol. 31, Issue 3: Page 512-524.

## PROFESSIONAL SERVICES

---

Journal Reviewing: Biometrika, Journal of Machine Learning Research (JMLR), Transactions on Pattern Analysis and Machine Intelligence (TPAMI), Machine Learning, Artificial Intelligence, Transactions on Machine Learning Research.

Conference Reviewing: Conference on Neural Information Processing Systems, International Conference on Machine Learning, International Conference on Learning Representations, AISTATS, The Conference on Uncertainty in Artificial Intelligence, AAAI Conference on Artificial Intelligence

Guest editor for Machine Learning Journal.

## AWARDS

---

School of Engineering Fellowship, Stanford University	Sept 2017
Undergraduate Research Award, School of EECS, Peking University	Nov 2015
First Award, National Mathematical Olympiad in Provinces (A country-wide high school mathematics competition)	Sept 2011