

# YUDA SONG

(858)353-5390 ◊ yudas@cs.cmu.edu ◊ [yudasong.github.io](https://yudasong.github.io) ◊ [github.com/yudasong](https://github.com/yudasong)

## EDUCATION

---

**Carnegie Mellon University** August 2022 -  
Ph.D. in Machine Learning  
*Advisors:* Aarti Singh and J. Andrew Bagnell

**Carnegie Mellon University** August 2020 - December 2021  
M.S. in Machine Learning  
*Advisor:* Kris Kitani

**University of California, San Diego** September 2016 - June 2020  
B.S. in Computer Science, B.S. in Mathematics  
Summa Cum Laude  
*Advisor:* Sicun Gao

## RESEARCH INTEREST

---

I study the theory of interactive decision making in settings where agents operate in complex, high-dimensional environments, including those powered by foundation models. My research aims to identify what aspects of learning and interaction are fundamentally changed by powerful models and what constraints are imposed by the environment and data. By developing principled algorithms informed by problem structure, I focus on methods that are both theoretically grounded and effective in large-scale, real-world decision-making systems.

## HONORS

---

Two Sigma PhD Fellowship Runner-up 2025  
Neurips Outstanding Reviewer 2022

## WORK EXPERIENCE

---

**FAIR Paris** May 2025 - August 2025  
Student Researcher  
*Mentor:* Remi Munos

**Amazon NYC** May 2024 - December 2024  
Student Researcher  
*Mentors:* Udaya Ghai and Dean Foster

**Microsoft Research NYC** May 2023 - August 2023  
Student Researcher  
*Mentors:* Akshay Krishnamurthy and Dylan Foster

## PREPRINT

---

**Yuda Song\***, Lili Chen\*, Fahim Tajwar, Remi Munos, Deepak Pathak, J. Andrew Bagnell, Aarti Singh, Andrea Zanette, “Expanding the Capabilities of Reinforcement Learning via Text Feedback”, in *arXiv*, 2026. <https://arxiv.org/abs/2602.02482>.

Fahim Tajwar\*, Guanning Zeng\*, Yueer Zhou, **Yuda Song**, Daman Arora, Yiding Jiang, Jeff Schneider, Ruslan Salakhutdinov, Haiwen Feng, Andrea Zanette, “Maximum Likelihood Reinforcement Learning”, in *arXiv*, 2026. <https://arxiv.org/abs/2602.02710>.

## PUBLICATION

---

**Yuda Song**, Dhruv Rohatgi, Aarti Singh, Drew Bagnell, “To Distill or Decide? Understanding the Algorithmic Trade-off in Partially Observable Reinforcement Learning”, in *Conference on Neural Information Processing Systems (NeurIPS)*, 2025. <https://arxiv.org/abs/2510.03207>.

**Yuda Song**, Julia Kempe, Remi Munos, “Outcome-Based Exploration for LLM Reasoning”, in *Conference on Neural Information Processing Systems (NeurIPS) ARLET Workshop*, 2025. <https://arxiv.org/abs/2509.06941>.

Zhaoyi Zhou, **Yuda Song**, Andrea Zanette, “Accelerating Unbiased LLM Evaluation via Synthetic Feedback”, in *International Conference on Machine Learning (ICML)*, 2025. <https://arxiv.org/abs/2502.10563>.

**Yuda Song**, Hanlin Zhang, Udaya Ghai, Carson Eisenach, Sham M. Kakade, Dean Foster, “Mind the Gap: Examining the Self-Improvement Capabilities of Large Language Models”, in *International Conference on Learning Representations (ICLR)*, 2025. <https://arxiv.org/abs/2412.02674>.

**Yuda Song**, Gokul Swamy, Aarti Singh, J. Andrew Bagnell, Wen Sun, “The Importance of Online Data: Understanding Preference Fine-tuning via Coverage”, in *Conference on Neural Information Processing Systems (NeurIPS)*, 2024. <https://arxiv.org/abs/2406.01462>.

**Yuda Song**, Drew Bagnell, Aarti Singh, “Hybrid Reinforcement Learning from Offline Observation Alone”, in *International Conference on Machine Learning (ICML)*, 2024. <https://arxiv.org/abs/2406.07253>.

**Yuda Song**, Lili Wu, Dylan J. Foster, Akshay Krishnamurthy, “Rich-Observation Reinforcement Learning with Continuous Latent Dynamics”, in *International Conference on Machine Learning (ICML)*, 2024. <https://arxiv.org/abs/2405.19269>.

Yifei Zhou\*, Ayush Sekhari\*, **Yuda Song**, Wen Sun, “Offline Data Enhanced On-Policy Policy Gradient with Provable Guarantees”, in *International Conference on Learning Representations (ICLR)*, 2024. <https://arxiv.org/abs/2311.08384>.

Alekh Agarwal\*, **Yuda Song**\*, Wen Sun\*, Kaiwen Wang\*, Mengdi Wang\*, Xuezhou Zhang\*, “Provable Benefits of Representational Transfer in Reinforcement Learning”, in *Conference on Learning Theory (COLT)*, 2023. <https://arxiv.org/abs/2205.14571>.

Anirudh Vemula, **Yuda Song**, Aarti Singh, Drew Bagnell, Sanjiban Choudhury, “The Virtues of Laziness in Model-based RL: A Unified Objective and Algorithms”, in *International Conference on Machine Learning (ICML)*, 2023. <https://arxiv.org/abs/2303.00694>.

**Yuda Song**\*, Yifei Zhou\*, Ayush Sekhari, J. Andrew Bagnell, Akshay Krishnamurthy, Wen Sun, “Hybrid RL: Using Both Offline and Online Data Can Make RL Efficient”, in *International Conference on Learning Representations (ICLR)*, 2023. <https://arxiv.org/abs/2210.06718>.

Chengzhuo Ni, **Yuda Song**, Xuezhou Zhang, Zihan Ding, Chi Jin, Mengdi Wang, “Representation Learning for General-sum Low-rank Markov Games”, in *International Conference on Learning Representations (ICLR)*, 2023. <https://arxiv.org/abs/2210.16976>.

Xuezhou Zhang, **Yuda Song**, Masatoshi Uehara, Mengdi Wang, Alekh Agarwal, Wen Sun, “Efficient Reinforcement Learning in Block MDPs: A Model-free Representation Learning Approach”, in *International Conference on Machine Learning (ICML)*, 2022. <https://arxiv.org/abs/2202.00063>.

**Yuda Song**, Ye Yuan, Wen Sun, Kris Kitani, “Online No-regret Model-Based Meta RL for Personalized Navigation”, in *Learning for Dynamics & Control Conference (L4DC)*, 2022. <https://arxiv.org/abs/2204.01925>.

Ye Yuan, **Yuda Song**, Zhengyi Luo, Wen Sun, Kris Kitani, “Transform2Act: Learning a Transform-and-Control Policy for Efficient Agent Design”, in *International Conference on Learning Representations*

(ICLR), 2022. <https://arxiv.org/abs/2110.03659>.

**Yuda Song**, Wen Sun, “PC-MLP: Model-based Reinforcement Learning with Policy Cover Guided Exploration”, in *International Conference on Machine Learning (ICML)*, 2021. <https://arxiv.org/abs/2107.07410>.

**Yuda Song**, Aditi Mavalankar, Wen Sun, Sicun Gao, “Provably Efficient Model-based Policy Adaptation”, in *International Conference on Machine Learning (ICML)*, 2020. <https://arxiv.org/abs/2006.08051>.

## TEACHING EXPERIENCE

---

### Lecturer

- CMU 10734: Foundations of Autonomous Decision Making under Uncertainty (Fall 2024, Fall 2025)

### Guest Lecturer

- Cornell CS6789: Foundations of Reinforcement Learning (Fall 2024)
- CMU 17740: Algorithmic Foundations of Interactive Learning (Fall 2024)

### Teaching Assistant

- UCSD CSE291: Topics in Search and Optimization (Winter 2020)
- UCSD CSE154: Deep Learning (Fall 2019)
- UCSD CSE150: Introduction to AI: Search and Reasoning (Winter 2019, Spring 2020)
- UCSD CSE30: Computer Organization and Systems Programming (Spring 2019, Winter 2018)
- UCSD CSE11: Introduction to CS & OOP (Fall 2018)

## SERVICE

---

### Reviewer

- Conference: COLT (2025-), ICML (2021-), NeurIPS (2021-), ALT (2024-), ICLR (2022-), AAAI (2021-2022)
- Journal: Transactions on Machine Learning Research, Journal of Machine Learning Research, IEEE Transactions on Signal Processing