

Aditya Modi

Mountain View, CA-94043
✉ adityamodi94@gmail.com
📄 adityamodi.github.io

Overview

I'm a machine learning researcher interested in general sequential decision making problems with a focus on the foundations and real-world applications of reinforcement learning. I have extensive experience in both fundamental statistical modeling for various ML applications and in productionizing ML models. In general, my interests range from active learning to bandits, reinforcement learning and causal inference. For applications, my interests lie in most interactive ML settings found in data-driven decision making systems like recommendation systems, AI alignment etc.

Education

- Sept '16-Nov '21 **PhD, Computer Science and Engineering**, *University of Michigan*, Ann Arbor,
Advisors: Satinder Singh and Ambuj Tewari.
Thesis: Provably Efficient Reinforcement Learning Under Linear Model Structures: From Tabular to Feature Based Exploration [\[link\]](#)
- Aug '12- May '16 **Bachelor of Technology**, *Indian Institute of Technology*, Kanpur, GPA – 9.4/10.0.
Major: Computer Science

Professional Experience

- Nov '21-Present **Data and Applied Scientist II**, *Microsoft*, Mountain View, CA.
Applied research in Microsoft Advertising team
Topics: Contextual bandits, reinforcement learning, causal inference, LLMs.
Key Projects: Search result whole page optimization, dynamic ad copy generation, adaptive real-time resource allocation, deployable offline RL, multi-task RL, incorporating transformers/LLM architecture in data-driven decision making pipelines.
- July-Oct 2018 **Research Intern**, *Microsoft Research*, Redmond.
Optimizing modular software pipelines via Reinforcement Learning
Mentors: Debadepta Dey, Eric Horvitz
Worked on the application of contextual bandit, learning to search and policy search methods to input-adaptive parameter/algorithm selection across components in any modular software pipeline. Work published in AAAI 2020.
- Sept-Dec 2016 **Research Assistant**, *University of Michigan*, Ann Arbor.
Data-dependent Importance weighted Active Learning
Advisors: Ambuj Tewari and Barzan Mozafari
Studied the sample complexity of importance-weighted active learning (IWAL) algorithms based on data-dependent complexity measures for bounded loss functions.
- May-July 2015 **Research Intern**, *Microsoft Research*, Bangalore, India.
Active Semi-supervised Performance Evaluation
Advisor: Sundararajan Sellamanickam, Principal Applied Scientist. [\[Report\]](#)
Proposed an estimation method for performance measures of black-box classifiers using scarcely labelled datasets for various non-decomposable performance measures (ROC curve, PR curve, F-measure).

Publications/Preprints

[\(Google scholar profile\)](#)

*Equal contribution

- NeurIPS 2024 **How to Solve Contextual Goal-Oriented Problems with Offline Datasets.** [\[arxiv\]](#)
Ying Fan, Jingling Li, Adith Swaminathan, **Aditya Modi**, Ching-An Cheng
Conference on Neural Information Processing Systems (NeurIPS), 2024
Abridged version accepted to *NeurIPS 2023 Workshop on Generalization in Planning and NeurIPS 2023 Workshop on Goal-Conditioned Reinforcement Learning.*

- Automatica **Joint Learning of Linear Time-Invariant Dynamical Systems.**
Aditya Modi, Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, George Michailidis
Automatica (IFAC journal). [\[link\]](#)[\[arxiv\]](#)
- JMLR 2024 **Model-Free Representation Learning and Exploration in Low-rank MDPs.**
Aditya Modi*, Jinglin Chen*, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal
Journal of Machine Learning Research (JMLR). [\[link\]](#)
- ICML 2023 **Towards Modular Machine Learning Pipelines.**
Aditya Modi, Jivat Kaur, Maggie Makar, Pavan Mallapragada, Amit Sharma, Emre Kiciman, Adith Swaminathan
Localized Learning Workshop, ICML 2023. [\[link\]](#)
- NeurIPS 2022 **On the Statistical Efficiency of Reward-free Exploration in Non-linear Reinforcement Learning.**
Jinglin Chen*, Aditya Modi*, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal
Conference on Neural Information Processing Systems (NeurIPS), 2022. [\[link\]](#)
- ICML 2022 **Multi-task Learning of Linear Control Systems under Instability.**
Aditya Modi, Ziping Xu, Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari
ICML Workshop on Complex Feedback in Online Learning, 2022. [\[link\]](#)
- ALCoS 2022 **Joint Learning-Based Stabilization of Multiple Unknown Linear Systems.**
Mohamad Kazem Shirani Faradonbeh, Aditya Modi
IFAC Workshop on Adaptive Learning and Control Systems (ALCOS), 2022. [\[link\]](#)
- ICML 2020 **Clinician-in-the-Loop Decision Making: Reinforcement Learning with Near-Optimal Set-Valued Policies.**
Shengpu Tang, Aditya Modi, Michael Sjoding, Jenna Wiens
International Conference on Machine Learning (ICML), 2020. [\[link\]](#)
- UAI 2020 **No-regret Exploration in Contextual Reinforcement Learning.**
Aditya Modi and Ambuj Tewari
Conference on Uncertainty in Artificial Intelligence (UAI), 2020
Abridged version accepted to *ICML 2019 wkshp on RL for Real Life and Conference on Reinforcement Learning and Decision Making (RLDM) 2019*. [\[link\]](#)
- AISTATS 2020 **Sample Complexity of Reinforcement Learning with Linearly Combined Model Ensembles.**
Aditya Modi, Nan Jiang, Ambuj Tewari, Satinder Singh
International Conference on Artificial Intelligence and Statistics (AISTATS), 2020. [\[link\]](#)
- AAAI 2020 **Meta-Reasoning in Modular Software Systems via Reinforcement Learning.**
A. Modi, D. Dey, A. Agarwal, A. Swaminathan, B. Nushi, S. Andrist, E. Horvitz
AAAI Conference on Artificial Intelligence (AAAI), 2020
Invited poster at *ICML 2019 Workshop on Reinforcement Learning for Real Life* [\[link\]](#)
- ALT 2018 **Markov Decision Processes with Continuous Side Information.**
Aditya Modi, Nan Jiang, Satinder Singh, Ambuj Tewari
International Conference on Algorithmic Learning Theory (ALT) 2018 [\[link\]](#)

Scholastic Achievements

- 2019-20 NeurIPS 2019, 2020, 2022 and ICML 2020 best reviewer award.
- 2018, 2019 Rackham Travel Grant (ALT '18, ICML '19)
- 2013, 2015 Academic Excellence Award, IIT Kanpur.
- 2014 Ram Parkash Chopra Memorial Scholarship, given for academic excellence, IIT Kanpur.
- 2013-15 Honourable mention in **ACM ICPC Asia Amritapuri** (2014-15, 2013-14) and **Kanpur regionals** (2013-14).
- 2013 O.P. Jindal Engineering and Management scholarship (awarded to select few candidates from top eng. and management institutes in India)

Talks/Presentations

- July 2022 **On the Statistical Efficiency of Reward-free Exploration in Non-linear RL**, *Poster presentation*.
TTI Chicago Workshop on *New Models in Online Decision Making for Real-World Applications*, 2022
- June 2021 **Contextual Reinforcement Learning: Learning optimal intervention policies for a heterogeneous population**, *Contributed talk*.
Canadian Operations Research Society (CORS) annual conference, 2021
- March 2021 **Model-free Representation Learning and Exploration in Low-rank MDPs**, *Invited talk*.
RL Theory virtual seminar series. [Link]
- June 2019 **Meta-Reasoning in Modular Software Systems via Reinforcement Learning**, *Invited poster*.
ICML 2019 Workshop on Reinforcement Learning for Real Life
- March 2019 **Contextual Decision Processes using Generalized Linear Models**, *Speed Oral and poster*.
Mich. Student Symp. on Interdisciplinary Statistical Sciences (MSSISS) 2019
- March 2018 **Markov Decision Processes with Continuous Side Information**, *Oral presentation*.
Mich. Student Symp. on Interdisciplinary Statistical Sciences (MSSISS) 2018

Professional Services and Participation

- Conference PC* AAI Conference on Artificial Intelligence 2019
Conference on Artificial Intelligence and Statistics (AISTATS) 2019-25
Conference on Algorithmic Learning Theory (ALT) 2020
International Conference on Machine Learning (ICML) 2019-24 (2020*)
Conference on Neural Information Processing Systems (NeurIPS) 2019-24 ('19,'20,'22*)
Conference on Uncertainty in AI (UAI) 2022-23
International Conference on Learning Representations (ICLR) 2022-24
Blog post track, International Conference on Learning Representations (ICLR) 2023-24
Conference on Lifelong Learning Agents (CoLLAs) 2022-24
Reinforcement Learning Conference (RLC) 2024
* **Top reviewer award**
- Journal Reviews* IEEE Transactions on Information Theory 2022, 2024
Journal of Machine Learning Research (JMLR) 2024
- Workshop/Symposium PC* Theoretical Foundations of RL (ICML) 2020
Deep Reinforcement Learning workshops (NeurIPS) 2020-22
Workshop on RL Theory (ICML) 2021
European Workshop on Reinforcement Learning (EWRL) 2022-24
Reinforcement Learning for Real Life Workshop (NeurIPS) 2022
Workshop on Interactive Learning with Implicit Human Feedback (ICML) 2023
International Symposium on Artificial Intelligence and Mathematics (ISAIM) 2023
Workshop on Aligning Reinforcement Learning Experimentalists and Theorists (ICML) 2024
- Simons Institute Visiting graduate student, Simons Institute' (UC Berkeley) program on Theory of RL Fall 2020
CHAI Participant in 2nd Center for Human-Compatible AI (CHAI) annual workshop April 2018
U. of Michigan Co-organizer, Statistical Machine Learning Reading group, Univ. of Michigan 2017-18

Teaching/Mentoring experience

- Interns Syomantak Chaudhuri (PhD Candidate, UC Berkeley), Data Scientist Intern, Microsoft Ads May-Aug '24
(co-)mentored Eddy Hudson (PhD Candidate, UT Austin), Applied Scientist Intern, Microsoft Ads Jun-Aug '22
Ying Fan (PhD Candidate, UW-Madison), Research Intern, Microsoft Research Redmond Jun-Aug '22
- Winter 2017 **Graduate Student Instructor**, EECS 445 - Machine Learning, Univ. of Michigan.
- Winter 2016 **Student Mentor**, CS 771 - Machine Learning Techniques, IIT Kanpur.
- Fall 2015 **Teaching Assistant**, ESO 207 - Data Structures and Algorithms, IIT Kanpur.