Aditya Modi

1045 La Avenida St.
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. In general, my interests range from active learning to bandits, reinforcement learning and causal inference. Overall, my aim is to devise novel methods and advance our understanding of building AI agents which can augment and/or compete with human-level decision making, safely and reliably.

Education

Sept '16-Nov '21 **PhD, Computer Science**, *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 **Applied Research Scientist**, *Microsoft*, Mountain View, CA.
Working on applied research problems on topics ranging from bandits and RL to causal inference to improve the advertising products at MS.
Current focus: Offline RL and combinatorial contextual bandits for post-auction models; causally-faithful independently-improvable ML pipeline design.

July-Oct 2018 **Research Intern**, *Microsoft Research*, Redmond.
Optimizing modular software pipelines via Reinforcement Learning
Mentors: Debadeepta 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 2022 **On the Statistical Efficiency of Reward-free Exploration in Non-linear Reinforcement Learning**.
Jinglin Chen*, **Aditya Modi***, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal [arxiv]
Conference on Neural Information Processing Systems (NeurIPS), 2022.

ICML 2022 **Multi-task Learning of Linear Control Systems under Instability**.
Aditya Modi, Ziping Xu, Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari [link coming soon]
ICML Workshop on Complex Feedback in Online Learning, 2022.

ALCoS 2022 **Joint Learning-Based Stabilization of Multiple Unknown Linear Systems**.
Mohamad Kazem Shirani Faradonbeh, **Aditya Modi** [arxiv]
IFAC Workshop on Adaptive Learning and Control Systems (ALCOS), 2022.

Automatica **Joint Learning of Linear Time-Invariant Dynamical Systems**.
Aditya Modi, Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, George Michailidis [arxiv]
Accepted provisionally to Automatica (IFAC journal).

JMLR **Model-Free Representation Learning and Exploration in Low-rank MDPs**.
Aditya Modi*, Jinglin Chen*, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal [arxiv]
Accepted to Journal of Machine Learning Research (JMLR) with minor revisions.

ICML 2020 **Clinician-in-the-Loop Decision Making: Reinforcement Learning with Near-Optimal Set-Valued Policies**.
Shengpu Tang, **Aditya Modi**, Michael Sjoding, Jenna Wiens [link]
International Conference on Machine Learning (ICML), 2020.

UAI 2020 **No-regret Exploration in Contextual Reinforcement Learning**.
Aditya Modi and Ambuj Tewari [link]
Conference on Uncertainty in Artificial Intelligence (UAI), 2020
Abridged version accepted to ICML 2019 wkshp on RL for Real Life and RLDM 2019.

AISTATS 2020 **Sample Complexity of Reinforcement Learning with Linearly Combined Model Ensembles**.
Aditya Modi, Nan Jiang, Ambuj Tewari, Satinder Singh [link]
International Conference on Artificial Intelligence and Statistics (AISTATS), 2020.

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 [link]
AAAI Conference on Artificial Intelligence (AAAI), 2020
Invited poster at *ICML 2019 Workshop on Reinforcement Learning for Real Life*

ALT 2018 **Markov Decision Processes with Continuous Side Information**.
Aditya Modi, Nan Jiang, Satinder Singh, Ambuj Tewari [link]
International Conference on Algorithmic Learning Theory (ALT) 2018

Scholastic Achievements

2019-20 NeurIPS 2019, 2020, 2022 and ICML 2020 best reviewer award.

2019 Microsoft Research PhD Fellowship Finalist (University Nominee).

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

Program Committee/reviewer

Conference PC	AAAI Conference on Artificial Intelligence	2019
	Conference on Artifical Intelligence and Statistics (AISTATS)	2019-23
	Conference on Algorithmic Learning Theory (ALT)	2020
	International Conference on Machine Learning (ICML)	2019-23 (2020*)
	Conference on Neural Information Processing Systems (NeurIPS)	2019-22 ('19,'20,'22*)
	Conference on Uncertainty in AI (UAI)	2022-23
	International Conference on Learning Representations (ICLR)	2022-23
	Conference on Lifelong Learning Agents (CoLLAs)	2022
	* **Top reviewer award**	
Journal Reviews	IEEE Transactions on Information Theory	2022
	Journal of Machine Learning Research (JMLR)	2022
Workshop 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
	Reinforcement Learning for Real Life Workshop, NeurIPS	2022

Fall '20 Long term participant in Simons Institute' (UC Berkeley) program on Theory of Reinforcement Learning

April 2018 Participant in 2nd Center for Human-Compatible AI (CHAI) annual workshop.

2017, 2018 Co-organizer, Statistical Machine Learning Reading group, Univ. of Michigan.

Teaching/Mentoring experience

(Co-)Mentoring	Eddy Hudson (PhD Candidate, UT Austin), Applied Scientist Intern, Microsoft Ads	Jun-Aug '22
	Ying Fan (PhD Candidate, UW-Madison), Research Intern, Microsoft Research	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.