

Adith Swaminathan

Research Scientist
Machine Learning and Inference Research Team
Netflix Research

SKYPE: adith.swaminathan
E-MAIL: adith387@gmail.com
PHONE: +1 (607) 442 2517

Research Interests

Machine Learning to improve Decision-Making for User-Interactive Systems
Reinforcement Learning • Causal Reasoning • Machine Learning Applications

Education

2011-2017

Cornell University

Ph.D. in Computer Science

- Graduate Minor in Applied Math
- Advisor: Thorsten Joachims
- Thesis: Counterfactual Evaluation and Learning from Logged User Feedback
- Grade Point Average: 4.187 / 4.3
- Committee: Thorsten Joachims, Johannes Gehrke, Éva Tardos, Robert Kleinberg

2014

Cornell University

M.Sc. in Computer Science

2006-2010

Indian Institute of Technology, Bombay

B.Tech. in Computer Science & Engineering

- Thesis: Learning Many-Many Matchings
- Advisor: Soumen Chakrabarti
- Cumulative Performance Index : 9.32 / 10

Awards

2020

Best Paper at ICLR Workshop on Trustworthy Machine Learning

2017

Nominated by Cornell University for The ACM Doctoral Dissertation Award

2017

Best Paper at The 10th ACM International Conference on Web Search and Data Mining

2009

OP Jindal Scholarship for excellence in leadership

2006

CBSE Merit Scholarship for excellence in All India Engineering Entrance Examinations

Industry Experience

2024-

Research Scientist, Machine Learning and Inference Research Team at Netflix Research

2023-2024

Principal Researcher, Augmented Learning and Reasoning Group at Microsoft Research

- Designed learning algorithms for AI systems using rich feedback
- Developed techniques to personalize multi-stakeholder, teachable AI copilots

2020-2023

Senior Researcher, Augmented Learning and Reasoning Group at Microsoft Research @ Redmond

- Pioneered Causal RL algorithms for resource allocation, which can save Azure \$50M+ annually
- Modularized ML pipelines using new causal ML insights, and earned Bing Ads \$60M+ annually

2017-2020

Researcher, Reinforcement Learning group at Microsoft Research AI

- Founded the Reinforcement Learning group in Redmond
- Designed RL-powered videogame AI shipped with Xbox games like Forza Motorsport 6 and Minecraft; spotlighted and reviewed by Bill Gates and Satya Nadella
- Causal ML for slate recommendation, used by Netflix, Spotify, etc. and available in Personalizer

2015

Research Intern, Machine Learning group at Microsoft Research NYC

- Developed off-policy estimators that evaluate policies for contextual combinatorial bandits
- Mentor: Miroslav Dudík and John Langford

- 2013 Research Intern, Computer Human Interactive Learning group at Microsoft Research Redmond
- Created interpretable text classification models that are teachable by novice users
 - Mentor: David Grangier and Carlos Garcia
- 2012 Research Intern, Search Labs at Microsoft Research Silicon Valley Center
- Explored methods to automatically augment textbooks with relevant high quality videos
 - Mentor: Krishnaram Kenthapadi, and Rakesh Agrawal, Sreenivas Gollapudi and Anitha Kannan
- 2010-2011 Strategist, Tower Research Capital
- Conceptualized and implemented several profitable automated high frequency trading systems
 - Wrote high performance protocol implementations for many electronic trading exchanges
 - Created critical monitoring, alerting and broadcasting systems
- 2009 Software Development Engineer (Testing) Intern, Microsoft R&D India Development Center
- With the Office Mobile team, in charge of Excel Mobile testing and performance benchmarking
 - Automated stress-tests on my own initiative that improved test team productivity many-fold

Teaching

- 2019 Summer Graduate School on Mathematics of Machine Learning
- Co-organized with Sébastien Bubeck, Anna Karlin
- 2017-2019 Microsoft AI-611 Advanced Projects
- Instructor, multi-part lectures on Bandits and Reinforcement Learning
 - Mentored 10 product teams to execute high-impact ML projects
- 2017 Microsoft edX course on Reinforcement Learning
- Instructor, Bandits module
- 2016 Tutorial at The 39th International ACM SIGIR Conference
- Counterfactual Evaluation and Learning for Search, Recommendation and Ad Placement
 - Co-taught with Thorsten Joachims
- 2014-2015 Teaching Assistant for CS 4780/5780/6780: Machine Learning, Advanced Machine Learning
- Was awarded both years for *exemplary performance as a Graduate Teaching Assistant*
- 2012 Teaching Assistant for CS 2110: Object-Oriented Programming and Data Structures
- *Homework Czar*, in charge of creating, grading and quality control of all homeworks

Publications

JOURNALS

- 2021 Recommendations as Treatments
- Thorsten Joachims, Ben London, Yi Su, *Adith Swaminathan*, Lequn Wang
 - AI Magazine by The Association for Advancement of Artificial Intelligence
- 2015 Batch Learning from Logged Bandit Feedback through Counterfactual Risk Minimization
- *Adith Swaminathan*, Thorsten Joachims
 - Journal of Machine Learning Research Special Issue in memory of Alexey Chervonenkis

CONFERENCES

- 2024 Trace is the New AutoDiff: Unlocking Efficient Optimization of Computational Workflows
- Ching-An Cheng, Allen Nie, *Adith Swaminathan*
 - In proceedings of The 38th Annual Conference on Neural Information Processing Systems
- 2024 How to Solve Contextual Goal-Oriented Problems with Offline Datasets?
- Ying Fan, Jingling Li, *Adith Swaminathan*, Aditya Modi, Ching-An Cheng
 - In proceedings of The 38th Annual Conference on Neural Information Processing Systems
- 2024 On Overcoming Miscalibrated Conversational Priors in LLM-based ChatBots
- Christine Herlihy, Jennifer Neville, Tobias Schnabel, *Adith Swaminathan*
 - In proceedings of The 40th Conference on Uncertainty in Artificial Intelligence
- 2023 Hindsight Learning for MDPs with Exogenous Inputs
- Sean R. Sinclair, Felipe Vieira Frujeri, Ching-An Cheng, Luke Marshall, Hugo Barbalho, Jingling Li, Jennifer Neville, Ishai Menache, *Adith Swaminathan*
 - In proceedings of The 40th International Conference on Machine Learning

- 2022 Improving Long-Term Metrics for Recommendation using Short-Horizon Reinforcement Learning
- Bogdan Mazouze, Paul Mineiro, Pavithra Srinath, Reza Sharifi Sedeh, Doina Precup, *Adith Swaminathan*
 - In proceedings of The European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases
- 2021 Heuristic-Guided Reinforcement Learning
- Ching-An Cheng, Andrey Kolobov, *Adith Swaminathan*
 - In proceedings of The 35th Annual Conference on Neural Information Processing Systems
- 2020 Provably Good Batch Off-Policy Reinforcement Learning Without Great Exploration
- Yao Liu, *Adith Swaminathan*, Alekh Agarwal, Emma Brunskill
 - In proceedings of The 34th Annual Conference on Neural Information Processing Systems
- 2020 Active Learning for ML-Enhanced Database Systems
- Lin Ma, Bailu Ding, Sudipto Das, *Adith Swaminathan*
 - In proceedings of The ACM SIGMOD/PODS International Conference on Management of Data
- 2020 Learning Calibratable Policies using Programmatic Style-Consistency
- Eric Zhan, Albert Tseng, Yisong Yue, *Adith Swaminathan*, Matthew Hausknecht
 - In proceedings of The 37th International Conference on Machine Learning
- 2020 Working Memory Graphs
- Ricky Loynd, Roland Fernandez, Asli Celikyilmaz, *Adith Swaminathan*, Matthew Hausknecht
 - In proceedings of The 37th International Conference on Machine Learning
- 2020 Metareasoning in Modular Systems: On-the-Fly Configuration using Reinforcement Learning
- Aditya Modi, Debadepta Dey, Alekh Agarwal, *Adith Swaminathan*, Besmira Nushi, Sean Arist, Eric Horvitz
 - In proceedings of The 34th AAAI Conference on Artificial Intelligence
- 2019 Off-Policy Policy Gradient with Stationary Distribution Correction
- Yao Liu, *Adith Swaminathan*, Alekh Agarwal, Emma Brunskill
 - In proceedings of The Conference on Uncertainty in Artificial Intelligence
- 2019 A Distillation Approach to Data Efficient Individual Treatment Effect Estimation
- Maggie Makar, *Adith Swaminathan*, Emre Kiciman
 - In proceedings of The 33rd AAAI Conference on Artificial Intelligence
- 2018 Deep Learning with Logged Bandit Feedback
- Thorsten Joachims, *Adith Swaminathan*, Maarten de Rijke
 - In proceedings of The 6th International Conference on Learning Representations
- 2017 Unbiased Learning-to-Rank with Biased Feedback
- Thorsten Joachims, *Adith Swaminathan*, Tobias Schnabel
 - **Best Paper Award**
 - In proceedings of The 10th ACM International Conference on Web Search and Data Mining
- 2017 Off-Policy Evaluation for Slate Recommendation
- *Adith Swaminathan*, Akshay Krishnamurthy, Alekh Agarwal, Miroslav Dudík, John Langford
 - In proceedings of The 31st Annual Conference on Neural Information Processing Systems
- 2016 Unbiased Comparative Evaluation of Ranking Functions
- Tobias Schnabel, *Adith Swaminathan*, Peter Frazier, Thorsten Joachims
 - **Best Presentation Award**
 - In proceedings of The 2nd International Conference on the Theory of Information Retrieval
- 2016 Counterfactual Evaluation and Learning for Search, Recommendation and Ad Placement
- *Adith Swaminathan*, Thorsten Joachims
 - Tutorial and in proceedings of The 39th International ACM SIGIR Conference on Research and Development in Information Retrieval
- 2016 Recommendations as Treatments: Debiasing Learning and Evaluation
- Tobias Schnabel, *Adith Swaminathan*, Ashudeep Singh, Navin Chandak, Thorsten Joachims
 - In proceedings of The 33rd International Conference on Machine Learning
- 2015 The Self-Normalized Estimator for Counterfactual Learning
- *Adith Swaminathan*, Thorsten Joachims
 - In proceedings of The 29th Annual Conference on Neural Information Processing Systems
- 2015 Counterfactual Risk Minimization: Learning from Logged Bandit Feedback
- *Adith Swaminathan*, Thorsten Joachims
 - In proceedings of The 32nd International Conference on Machine Learning

- 2014 Mining Videos from the Web for Electronic Textbooks
- Rakesh Agrawal, Maria Christoforaki, Sreenivas Gollapudi, Anitha Kannan, Krishnaram Kenthapadi, *Adith Swaminathan*
 - In proceedings of The 12th International Conference on Formal Concept Analysis
- 2013 Beyond Myopic Inference in Big Data Pipelines
- Karthik Raman, *Adith Swaminathan*, Johannes Gehrke, Thorsten Joachims
 - In proceedings of The ACM SIGKDD Conference on Knowledge Discovery and Data mining
- 2012 Temporal Corpus Summarization using Submodular Word Coverage
- Ruben Sipos, *Adith Swaminathan*, Pannaga Shivaswamy, Thorsten Joachims
 - In proceedings of The International Conference on Information and Knowledge Management

WORKSHOPS

- 2024 LLF-Bench: Benchmark for Interactive Learning from Language Feedback
- Ching-An Cheng, Andrey Kolobov, Dipendra Misra, Allen Nie, *Adith Swaminathan*
 - In ICLR Workshop on LLM Agents
- 2024 Combining Open-box Simulation and Importance Sampling for Tuning Large-Scale Recommenders
- Kaushal Paneri, Michael Munje, Kailash Singh Maurya, *Adith Swaminathan*, Yifan Shi
 - In RecSys Workshop on Causality, Counterfactuals and Sequential Decision-Making
- 2023 The Importance of Directional Feedback for LLM-based Optimizers
- Allen Nie, Ching-An Cheng, Andrey Kolobov, *Adith Swaminathan*
 - In NeurIPS Workshop on Foundation Models for Decision Making
- 2023 Interactive Robot Learning from Verbal Correction
- Huihan Liu, Alice Chen, Yuke Zhu, *Adith Swaminathan*, Andrey Kolobov, Ching-An Cheng
 - In CoRL Workshop on Language and Robot Learning
- 2023 Towards Modular Machine Learning Pipelines
- Aditya Modi, Jivat Neet Kaur, Maggie Makar, Pavan Mallapragada, Amit Sharma, Emre Kıcıman, *Adith Swaminathan*
 - In ICML Workshop on Localized Learning
- 2020 Improved Image Wasserstein Attacks and Defenses
- J. Edward Hu, *Adith Swaminathan*, Hadi Salman, Greg Yang
 - **Best Paper Award**
 - In ICLR Workshop on Trustworthy ML
- 2019 Customizing Scripted Bots: Sample-Efficient Imitation Learning for Human-like Behavior in Minecraft
- Brian Broll, Matthew Hausknecht, Dave Bignell, *Adith Swaminathan*
 - In AAMAS Workshop on Adaptive and Learning Agents
- 2018 Multi-Preference Actor Critic
- Ishan Durugkar, Matthew Hausknecht, *Adith Swaminathan*, Patrick MacAlpine
 - In NeurIPS Workshop on Deep Reinforcement Learning
- 2016 Large-scale Validation of Counterfactual Learning Methods: A Test-Bed
- Damien Lefortier, *Adith Swaminathan*, Xiaotao Gu, Thorsten Joachims, Maarten de Rijke
 - In NeurIPS Workshop on Inference and Learning of Hypothetical and Counterfactual Interventions in Complex Systems
- 2015 Unbiased Ranking Evaluation on a Budget
- Tobias Schnabel, *Adith Swaminathan*, Thorsten Joachims
 - In proceedings of The 24th International Conference on World Wide Web

PRE-PRINTS

- 2024 AutoAttacker: An LLM-guided System to Implement Automatic Cyber-attacks
- Jiachen Xu, Jack W. Stokes, Geoff McDonald, Xuesong Bai, David Marshall, Siyue Wang, *Adith Swaminathan*, Zhou Li
- 2019 NAIL: A General Interactive Fiction Agent
- Matthew Hausknecht, Ricky Loynd, Greg Yang, *Adith Swaminathan*, Jason D. Williams
 - **Winner** of the CIG 2018 Text-Based Adventure AI Competition

PATENTS

- 2024 Updating Computational Workflow using Trace Feedback
 - Ching-An Cheng, Allen Nie, *Adith Swaminathan*, Ahmed Awadallah
- 2023 Cybersecurity Engine in a Security Management System
 - Jack W. Stokes, Geoff McDonald, David Marshall, Siyue Wang, Jiacen Xu, *Adith Swaminathan*
- 2020 Leveraging Query Executions to Improve Index Recommendations
 - Bailu Ding, Sudipto Das, Surajit Chaudhuri, Vivek R Narasayya, Ryan Marcus, Lin Ma, *Adith Swaminathan*
- 2017 Computational Query Modeling and Action Selection
 - Alekh Agarwal, Miroslav Dudík, Akshay Krishnamurthy, John Langford, *Adith Swaminathan*

Invited Talks

- 2023 Learning Algorithms for LLM-Augmented Systems
 - Impromptu Talk and Fireside Chat with Soumen Chakrabarti
 - At Indian Symposium on Machine Learning; Mumbai, India
- 2022 Hindsight Learning for Adaptive Systems
 - At Toyota Technological Institute Beyond RL Workshop; Chicago, U.S.A.
- 2021 Designer-friendly Machine Learning and Reinforcement Learning for Videogames
 - At AI and Gaming Research Summit; Redmond, U.S.A.
- 2020 Doing the Best with the Data we Have
 - At Rossi Seminar at Booth School; Chicago, U.S.A.
- 2020 Evaluating and Optimizing Recommender Systems with Logged User Behavior
 - At Emory University; Atlanta, U.S.A.
- 2019 Some Recent Advances in Learning from User Behavior
 - At Indian Institute of Technology Bombay; Mumbai, India
- 2019 Bandits, Reinforcement Learning and Causal Inference
 - At INFORMS Annual Meeting; Seattle, U.S.A.
- 2019 Counterfactual Reasoning for Recommendation
 - At Next Generation Recommender Systems Workshop; Mountain View, U.S.A.
- 2019 Data Efficient Causal Effect Estimation
 - At ICERM Tripods Workshop on Machine Learning in Health Research; Providence, U.S.A.
- 2018 Fixing Biases in Extreme Classification
 - At The Dagstuhl Seminar on Extreme Classification; Wadern, Germany
- 2018 BanditNet: Deep Learning from Logged Bandit Feedback
 - At California Institute of Technology; Pasadena, U.S.A.
 - At Spotify Machine Learning Day; Stockholm, Sweden
 - At Workshop on Personalization, Recommendation and Search; Los Gatos, U.S.A.
- 2017 Building Recommenders and Search Engines by Re-using Logged User Feedback
 - Keynote**
 - At Workshop on Machine Learning Methods for Recommender Systems; Houston, U.S.A.
- 2017 Causal Data Mining
 - At INFORMS Annual Meeting; Houston, U.S.A.
- 2016 Preferences and Counterfactuals
 - At The 28th European Conference on Operational Research; Poznan, Poland
- 2016 Counterfactual Risk Minimization for Interactive Learning Systems
 - At The Search Engines Amsterdam Meetup; Amsterdam, The Netherlands

Community Activities

- 2018-2020 Co-organized RecSys REVEAL Workshops on Reinforcement Learning for Recommendation
- 2016-2019 Co-organized NeurIPS, ICML Workshops on Causality and Machine Learning
- 2019 Co-organized ICML 2019 Workshop on Real-world Sequential Decision Making
- 2017 Co-organized CrowdAI Challenge on Counterfactual Policy Learning for Display Advertising
- 2023-2025 Senior Area Chair: AAAI

2021-2025 Area Chair: ICML, NeurIPS, ICLR
 2014-2020 Reviewer: NeurIPS, KDD, ICML, AISTATS, AAAI
 2018-2020 Reviewer: ICLR, WSDM, JAIR, ALT, WIML
 2015-2016 Reviewer: ICTIR, EVIA, DIR

Other Activities

2022 Created and co-organized the Microsoft Redmond Campus Photo Challenge to foster inclusion among interns during the COVID-19 pandemic
 2017-2022 Organized the Microsoft Research Reinforcement Learning Reading Group
 2017-2018 Created the Microsoft Research AI Seminar and the Distinguished Lecture Series
 2012-2015 Board Member and Webmaster of the Society for the Promotion of Indian Classical Music And Culture Amongst Youth (SPICMACAY) Cornell Chapter
 2009 Placed 3rd in the Computer Science Engineering Association Programming Contest, IIT Bombay
 2007-2008 Coordinator of the Pro-Nites concerts at Mood Indigo, IIT Bombay's cultural festival
 2007-2017 Trained Carnatic singer (a style of Indian Classical Music)
 2006-2007 Represented IIT Bombay in Mumbai Division Football Association League (Soccer) and in the Inter-IIT Tournament, and won the Institute Football League twice
 2004 Won the Division 2 title in the Mumbai School Sports Association Soccer tournament

Graduate Coursework

Cornell University

- Data-Sparse Matrix Computations
- Algorithmic Game Theory
- Analysis of Algorithms
- Bayesian Statistics and Data Analysis
- Optimal Learning
- Advanced Topics in Machine Learning
- Machine Learning Theory
- Counterfactual Machine Learning
- Structure of Information Networks
- Advanced Topic Modeling
- Advanced Language Technologies
- Advanced Database Systems
- Advanced Programming Languages

IIT Bombay

- Foundations of Machine Learning
- Information Retrieval & Mining for the Web
- New Trends in Information Technology
- Organization of Web Information
- Machine Learning: Theory and Methods
- Statistical Relational Learning
- Advanced Machine Learning
- Decision Analysis and Game Theory
- Program Analysis

References

Paul Bennett

- Director of Research
- Spotify, Boston
- email: paulb@alumni.cmu.edu

Jennifer Neville

- Partner Research Manager
- Microsoft Research, Redmond
- email: jenneville@microsoft.com