

EDUCATION

Harvard University Ph.D., Computer Science, Advisor: Milind Tambe	2020 – Current Cambridge, MA
Birla Institute of Technology and Science, Pilani BE (Hons.), Computer Science – GPA: 9.02/10, Merit Scholarship, Graduated with Distinction	2013 – 2017 Rajasthan, India

EXPERIENCE

Harvard University Graduate Research Assistant, Advisor: Prof. Milind Tambe – Conducting research in machine learning and sequential decision-making, especially as it applies to challenges in public health.	2020 – Current Cambridge, MA
Google Research India Research Intern, Advisor: Aparna Taneja – Authored a methods paper that will allow the scaling of machine learning-based mobile health interventions to millions of beneficiaries.	Summer 2023 Bangalore, India
ARMMAN Research Intern, Advisor: Neha Madhiwalla – Authored a workshop paper on ‘low-listenership prediction’ for the NGO’s Kilkari program, one of the largest mobile health programs in the world.	Spring 2023 Mumbai, India
Singapore Management University Research Engineer, Advisor: Prof. Pradeep Varakantham – Authored two first-author research papers that use Reinforcement Learning (RL) to address Sequential Decision Making problems that underlie societal challenges in Transportation and Security.	2018 – 2020 Singapore
Microsoft Research India Research Intern, Advisor: Dr. Colin Scott and Dr. Bill Thies – Helped build an Android app to augment local peer-to-peer file transfer like Bluetooth (a substitute to the internet for media acquisition in low resource communities) by creating a barter economy. Helped pilot the application in a village in Bihar, India along with my advisor and local partners from the region.	Spring 2017 Bangalore, India
Microsoft Research India Research Intern, Advisor: Dr. Sundararajan Sellamanickam – Investigated the ‘explainability’ of Recurrent Neural Networks in terms of compositional linguistic structures like ‘and’ and ‘but’ for the task of Sentiment Analysis in English.	Fall 2016 Bangalore, India

PUBLICATIONS

* indicates equal contribution

Work in Progress

- [WiP2] N. Boehmer*, Y. Nair*, **S. Shah***, L. Janson, A. Taneja, and M. Tambe, “Evaluating the Effectiveness of Index-Based Treatment Allocation”, *In Submission*, 2024.
- [WiP1] S. Verma, Y. Zhao, **S. Shah**, N. Boehmer, A. Taneja, and M. Tambe, “Group Fairness in Predict-Then-Optimize Settings for Restless Bandits”, *In Submission*, 2024.

Rigorously Reviewed Conference Publications

- [C10] **S. Shah**, A. Suggala, M. Tambe, and A. Taneja, “Efficient Public Health Intervention Planning Using Decomposition-Based Decision-Focused Learning”, *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2024.
- [C9] **S. Shah**, A. Perrault, B. Wilder, and M. Tambe, “Leaving the Nest: Going Beyond Local Loss Functions for Predict-Then-Optimize”, *Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI)*, 2024.
- [C8] K. Wang, S. Verma, A. Mate, **S. Shah**, A. Taneja, N. Madhiwalla, A. Hegde, and M. Tambe, “Decision-Focused Learning in Restless Multi-Armed Bandits with Application to Maternal and Child Care Domain”, *Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI)*, 2023.
- [C7] **S. Shah**, B. Wilder, A. Perrault, and M. Tambe, “Decision-Focused Learning without Differentiable Optimization: Learning Locally Optimized Decision Losses”, *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [C6] **S. Shah**, M. Lowalekar, and P. Varakantham, “Joint Pricing and Matching for City-Scale Ride-Pooling”, in *International Conference on Automated Planning and Scheduling (ICAPS)*, 2022.
- [C5] K. Wang, **S. Shah**, H. Chen, A. Perrault, F. Doshi-Velez, and M. Tambe, “Learning MDPs from Features: Predict-Then-Optimize for Sequential Decision Problems by Reinforcement Learning”, *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [C4] J. A. Killian, A. Biswas, **S. Shah**, and M. Tambe, “Q-Learning Lagrange Policies for Multi-Action Restless Bandits”, in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD)*, 2021.
- [C3] N. Raman, **S. Shah**, and J. Dickerson, “Data-Driven Methods for Balancing Fairness and Efficiency in Ride-Pooling”, in *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI)*, 2021.
- [C2] **S. Shah**, M. Lowalekar, and P. Varakantham, “Neural Approximate Dynamic Programming for On-Demand Ride-Pooling”, in *Proceedings of 34rd AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- [C1] **S. Shah**, A. Sinha, P. Varakantham, A. Perrault, and M. Tambe, “Solving Online Threat Screening Games using Constrained Action Space Reinforcement Learning”, in *Proceedings of 34rd AAAI Conference on Artificial Intelligence (AAAI)*, 2020.

Workshop Papers

- [W1] **S. Shah**, S. Verma, A. Mahale, K. M. Sudan, A. Hegde, A. Taneja, and M. Tambe, “Preliminary results in low-listenership prediction in one of the largest mobile health programs in the world”, in *Autonomous Agents for Social Good Workshop (AAMAS)*, 2023.

Demonstrations

- [D1] A. Kumar, **S. Shah**, M. Lowalekar, P. Varakantham, A. Ottley, and W. Yeoh, “FairVizARD: A Visualization System for Assessing Fairness of Ride-Sharing Matching Algorithms”, in *International Conference on Automated Planning and Scheduling (ICAPS)*, 2021.

PROFESSIONAL SERVICE

- **Conference PC Member:** AAAI ('23, '24), IJCAI ('23, '24), NeurIPS ('23), ICLR ('23), ICML ('24), EAAMO ('22)
- **Workshop PC Member:** Autonomous Agents for Social Good at AAMAS ('20, '21), AI for Social Good at Harvard CRCS ('20)
- **Top Reviewer** recognition for NeurIPS '23

TEACHING

- **Head Teaching Fellow**, Harvard University Fall 2021
CS 120: Algorithms and their Limitations
 - This was the first offering of the course. Helped design the course, managed undergraduate course assistants and course logistics, lead section, held office hours, designed and graded problem sets.