Aditya Soni

MSc Student · University of Alberta · Edmonton, AB

adityasoni25@gmail.com · dasyud.github.io · Google Scholar · GitHub

Research Interests

My research focuses on building reliable reinforcement learning algorithms that behave predictably and robustly across the distribution shifts, non-stationarity, and deployment conditions that arise in real-world sequential decision-making problems.

Education

University of Alberta *Edmonton, AB*
Master of Science in Computing Science Sep 2025 – Present

- Advisor: Martha White | GPA: 4.0 / 4.0

Birla Institute of Technology and Science, Pilani *Hyderabad, India*
Bachelor of Engineering in Electronics and Instrumentation Aug 2019 – Jul 2023

Experience

Microsoft Research *Bangalore, India*
Research Fellow Jul 2023 – Jun 2025
Mentors: Mayukh Das, Ching-An Cheng

- **Offline RL for Robust Bandwidth Estimation in RTC.** Designed and implemented a post-deployment policy shaping algorithm that makes offline RL policies robust to non-stationarity, deployed as a bandwidth estimator in RTC applications including Microsoft Teams. Achieved up to 18% improvement in call video quality score over IQL baseline.
- **Intelligent Overclocking in Azure Datacenters.** Designed offline policies for fulfilling service overclocking requests in Microsoft Azure datacenters, formulating the task as a constrained bilevel policy optimization problem.

Microsoft Research *Bangalore, India*
Research Intern Jan 2023 – Jun 2023
Mentors: Mayukh Das, Alok Gautam Kumbhare, Pulkit Misra

- **Fine-tuning Server Parameters for Workload Performance and Sustainability.** Applied reinforcement learning to tune server parameters, reducing power consumption by 11.3% while maintaining workload latency and throughput.

Publications

Workshop Papers

[W.1] **A. Soni**, M. Das, P. Misra, and C. Bansal. Intelligent Overclocking for Improved Cloud Efficiency. *Cloud Intelligence / AIOps Workshop, ASPLOS*, 2024. [Paper]

Preprints

[C.1] **A. Soni**, M. Das, A. Parayil, S. Ghosh, S. Shandilya, C.-A. Cheng, V. Gopal, S. Khairy, G. Mittag, Y. Hosseinkashi, and C. Bansal. Streetwise Agents: Empowering Offline RL Policies to Outsmart Exogenous Stochastic Disturbances in RTC. *Preprint.* [Paper]

Talks

Intelligent Overclocking for Improved Cloud Efficiency *April 2024*
Cloud Intelligence / AIOps Workshop, ASPLOS

Teaching

University of Alberta, Edmonton, AB

Teaching Assistant Machine Learning II (CMPUT 467/567) *Winter 2026*

Service and Outreach

Reviewer

- Reinforcement Learning Conference (RLC): 2026

Contributor

- PyTorch: ONNX export functionalities