Sreejeet Maity

Raleigh, North Carolina, U.S.A

Education

North Carolina State University — Raleigh, NC, U.S.A
Ph.D. Electrical Engineering — Aug 2023 - Present

Indian Institute of Science, Bangalore — Bangalore, India
M.Tech Robotics and Autonomous Systems — Aug 2021 - June 2023

Jadavpur University — Kolkata, India
B.E Electrical Engineering — Aug 2017 - July 2021

Research Interests

My research develops provably robust reinforcement learning algorithms with finite-sample guarantees under uncertainty and adversarial corruption. I focus on corruption-tolerant distributed and federated RL, together with tight minimax lower bounds characterizing the fundamental limits of robust learning. My broader interests include violation-constrained RL under model uncertainty and robust LLM post-training through adaptive adversarial curricula and distributionally robust optimization.

Experience

North Carolina State University — Raleigh, NC, U.S.A
Graduate Research/Teaching Assistant — Aug 2023 - Present
Ph.D. Advisor: Dr. Aritra Mitra.

- **Finding Robust Optimal Policy from Corrupted and Correlated Observations.**
  - Showed that vanilla Q-learning is provably fragile under state and reward corruption, and designed a robust variant that employs robust Bellman targets to achieve near-optimal performance.
  - Established non-asymptotic, finite-time convergence, with tightness certified by a matching minimax lower bound. Extended these guarantees to the *asynchronous* setting under temporal dependence, relaxing strong distributional assumptions on reward statistics.
  - Results disseminated across ICML 2026, NeurIPS 2025, and CDC 2024, and CDC 2026. CDC 2026 extended version under review at Automatica.

- **Robust Policy Evaluation under Adversarial Influences and Markovian Data.**
  - Established the first finite-time theory for robust `TD` learning under time-correlated (Markovian) noise and function approximation: (i) proved divergence of vanilla `TD` under adaptive corruption; (ii) designed `Robust-TD` with clean-case rates up to a small additive bias; and (iii) derived a near-tight lower bound certifying inevitability of the corruption term.
  - This research is published in AISTATS 2025.

- **Scalable Hybrid Algorithms for Multi-Agent Federated Learning Setting.**
  - Developed adversarially robust, communication-efficient RL for federated multi-agent settings; introduced a novel hybrid model-free and model-based class of algorithms with finite-time guarantees.
  - Extended to fully decentralized MARL with Byzantine tolerance and logarithmic communication, preserving collaborative speedups; and proved conditions under which the corruption bias diminishes with sample size.
  - Papers (-1, -2) published at ACC 2026. Extensions under review at IEEE TAC, TSP.

Columbia University, New York
Collaborator: Prof. James Anderson, Leonardo Felipe Toso. — New York, NY, U.S.A
Research Collaboration — May 2026 - Present

- **Byzantine-Robust Personalized Federated TD Learning with Shared Representations.**
  Developing Byzantine-robust representation learning for heterogeneous RL agents interacting with distinct MDPs and temporally dependent data. The framework combines a shared low-dimensional value-function representation with personalized local parameters while avoiding heterogeneity bias.

Selected First-Authored Publications

- *Sreejeet Maity*†, Aritra Mitra. **Corruption-Tolerant Optimal Asynchronous Q-Learning**, **International Conference on Machine Learning (ICML 2026)**.
[Summary] [Paper] [Poster] [Slides] [Code]

- *Sreejeet Maity*†, Aritra Mitra. **Adversarially-Robust TD-Learning with Markovian Data: Finite-Time Rates and Fundamental Limits**, **International Conference on Artificial Intelligence and Statistics (AISTATS 2025)**.
[Summary] [Paper] [Poster] [Slides] [Code]

■ *Sreejeet Maity*†, Aritra Mitra. **Robust Asynchronous Q-Learning under Reward and State Corruption via Batching**,
IEEE Conference on Decision and Control (CDC 2026).
[Summary] [Paper] [Poster] [Slides] [Code]

■ *Sreejeet Maity*†, Aritra Mitra. **Robust Q-Learning under Corrupted Rewards**,
IEEE Conference on Decision and Control (CDC 2024).
[Summary] [Paper] [Poster] [Slides] [Code]

■ *Sreejeet Maity*†, Aritra Mitra. **Robust Federated Q-Learning with Almost No Communication**,
2026 American Control Conference (ACC 2026).
[Summary] [Paper] [Poster] [Slides] [Code]

■ *Sreejeet Maity*†, Feng Zhu, Robert Heath, Aritra Mitra. **Variance-Reduced Q-Learning over Static and Time-Varying Networks**,
2026 American Control Conference (ACC 2026).
[Summary] [Paper] [Poster] [Slides] [Code]

Works Under Review

■ *Sreejeet Maity*†, A. Mitra. **Learning from Unreliable Trajectories: Adversarially-Robust Federated Q-Learning**, under review at IEEE Transactions on Automatic Control (TAC).

■ *Sreejeet Maity*†, F. Zhu, R. Heath, A. Mitra. **Decentralized Q-Learning with Asynchronous Sampling and Partial Coverage**, under review at IEEE Transactions on Signal Processing (TSP).

■ *Sreejeet Maity*†, L. Toso, J. Anderson, A. Mitra. **Value Function Representation for Adversarially Robust Federated TD Learning**, under review at Transactions on Machine Learning Research (TMLR).

■ *Sreejeet Maity*†, A. Mitra. **Batch-Robust Asynchronous Q-Learning under Corrupted Feedback with Markovian Data with Near-Optimal Rates**, in preparation for Automatica.

Accepted Workshop Presentations

1. **Multi-Agent Robust FRL with Sparse Communication** NCSU Robotics Symposium 2026.
2. **Corruption-Tolerant Agnostic Q-Learning** NeuRIPS 25-Reliable ML Workshop.
3. **Robust Federated RL with Byzantine Agents** Applied AI Symposium 2025, NESCW 2026.
4. **Theoretical Limits of Robust TD Learning** New York RL Workshop (NYRL 2025), Amazon.
5. **Finite-Time Theory for Robust RL** Northeast Systems and Control Symposium (NESCW 2025).
6. **Adversarially-Robust Deep Q-Network for Algorithmic Trading** MLSS 2025, NCSU.
7. **Robust Algorithms for Adversarial Reinforcement Learning** Applied AI Symposium 2024.

Invited Talks/Seminars

■ **Towards Finite-Time Rates for Adversarially-Robust RL** CORAL Seminar, NCSU.

Current Directions

■ **Violation-Constrained Reinforcement Learning under Model Uncertainty.**
We develop safe RL algorithms that maximize long-term reward while explicitly controlling cumulative constraint violations under uncertain transition dynamics. Our framework permits necessary exploration subject to a finite violation budget, thereby balancing learning efficiency with operational safety. We establish finite-time guarantees for policy suboptimality and cumulative constraint violation, and characterize the fundamental tradeoffs among exploration, robustness, and safety.

■ **Byzantine-Robust Representation Learning for Heterogeneous Reinforcement Learning**
We study collaborative representation learning among heterogeneous reinforcement learning agents, each interacting with a distinct Markov decision process under a fixed policy and generating temporally dependent data. Our approach decomposes each agent's value function into a shared low-dimensional representation and a personalized local parameter, allowing agents to learn common structure while accommodating differences in transition dynamics, rewards, policies, and value functions. We aim to develop finite-time guarantees under Byzantine clients without incurring a persistent bias due to agent heterogeneity.

■ **Robust Budgeted Recalibration under Partial Observability**
We study when a forecasting or learning system should be recalibrated when its latent model mismatch is only partially observed and recalibration opportunities are limited by a finite budget. Using a scalar belief-state formulation, we derive the exact Bellman recursion and establish that the optimal policy has a threshold structure in the posterior mean mismatch. We then develop a conservative robust policy for adversarially corrupted observations and prove finite-sample excess-loss guarantees that depend explicitly on

the robust estimation radius and the time the belief process spends near the optimal recalibration boundary. We also investigate how the optimal threshold varies with the remaining recalibration budget and posterior uncertainty.

■ **Adaptive Adversarial Curriculum Learning for Robust LLM Post-Training**
We study curriculum design for robust large language model post-training, where the training distribution is adaptively reweighted to expose the model to increasingly challenging and failure-inducing examples. We develop a principled framework that combines adaptive adversarial example selection with distributionally robust optimization, while balancing task difficulty, diversity, and learning progress. Our goal is to establish finite-sample guarantees for worst-case performance and characterize when adaptive curricula provide robustness improvements over static training mixtures.

Projects

■ `Federated MARL-GYM` : We introduce a custom multi-agent reinforcement learning environment built with Gymnasium and Pygame, designed for evaluating federated RL (FRL) algorithms. The environment models a grid world where multiple agents navigate to accomplish spatially distributed tasks, like reaching delivery points.

Skills

`Programming Languages`: Python, `MatLab`, `Simulink`, C++, R.
`Softwares and Libraries`: Tensorflow, `PyTorch`, Scikit-learn, Numpy, Pandas, Gymnasium, MuJoCo.
`Mathematical Skills`: Linear Algebra, Probability Theory, Robust Statistics, Stochastic Optimization.
`Research Skills`: Reinforcement Learning, Statistical Learning Theory, Optimization, Control Theory.

Relevant Courses

`Learning Theory`: Theoretical Foundations of Large-Scale Machine Learning, Reinforcement Learning, Machine Learning for Signal Processing, Theory and Applications of Bayesian Learning, Physics Modelling with Neural Networks, Deep Learning and Neural Networks, Distributed Learning.
`Mathematics`: Analysis, Probability and Stochastic Process, High-Dimensional Statistics, Stochastic Models and Applications, Convex Optimization for Data Science, Detection and Estimation Theory.
`Control Theory`: Safety-Critical Control of Robotic Systems, Dynamics of Linear Systems, Networked and Distributed Control, Formal Analysis for Control Theory.

Academic/Professional Services

■ Served as the Head Teaching Assistant for ECE 516: Systems and Control Engineering and ECE 308: Elements of Control Systems at the Department of Electrical and Computer Engineering, North Carolina State University, during the Spring and Fall of 2025.

■ Served as a reviewer for 50+ papers in multiple flagship control/ ML venues, including the American Control Conference (ACC), the IEEE Conference on Decision and Control (CDC), Indian Control Conference (ICC), Learning for Dynamics and Control (L4DC), Annual Conference on Neural Information Processing Systems(NeuRIPS), Journal of Machine Learning Research (JMLR), Transactions in Machine Learning Reserach (TMLR), IEEE Transactions in Automatic Control (TACON), Transactions on Control Systems Technology (TCST) , Transactions in Signal and Information Processing over Networks (TSIPN), and Transactions in Signal Processing (TSP).

Awards

`ACC 2026 Travel Award` — `May 2026`
`L4DC 2025 Student Support Grant` — `May 2025`
`NESCW 2025 Student Support Grant` — `May 2025`
`IEEE CDC 2024 Student Support Award` — `August 2024`
`NC State ECE Student Research Support Award` — `August 2024`
`College of Engineering Graduate Merit Award` — `August 2023-24, 2024-25`