Himadri Mandal

8610122976 • mandalhimadri06@gmail.com • quirtt.github.io
quirtt • quirtt

Education

Indian Statistical Institute, Kolkata **August 2023 — present**
Statistics Undergraduate, Ongoing 3rd year — 83.4% cumulative (2nd year)

Programming experience

Python: PyTorch, Numpy, Pandas, Bash, R, Next.js + TailwindCSS, LaTeX, Linux (Arch Linux on i3)

Work experience

Visiting Student Researcher: SCDLDS, Ashoka University *May 2025 — Ongoing*

Selected fellowships and awards

ISI K. Outstanding Performance: awarded January 2024, received ₹1500

ISI Kolkata B.Stat. Entrance: awarded August 2023, ranked 11th (out of 50,000 students)

Indian Olympiad Qualifier for Mathematics, KV: awarded February 2021, rank 4 in my region (out of >100,000 students all over India)

Atlas Fellowship India Finalist: awarded September 2022, received 1000$, top 200 (out of 4000 students) in a rationality fellowship

IISc Enumeration Finalist: awarded October 2022

CMI Tessellate Finalist: awarded October 2021

Research

Limits of Wtd Emp Samplers: w/ Sarvesh R. Iyer, Sandeep Juneja, et. al. ***Poster, AISTATS 2026***
arXiv. Second Author. We discuss the asymptotic efficiency of an estimator obtained by exponentially tilting the empirical distribution. We provide a sharp characterization of how much one can accurately tilt distributions given a certain number of samples. We delineate asymptotic regimes under which $R_{n,\theta}$ accurately estimates X_θ.

Active Simple Hypothesis Testing: w/ Sushant Vijayan, **TIFR** ***Submitted to IEEE ISIT 2026***
Paper. This work presents various methods for obtaining improved bounds on the optimal minimax error exponent in the Active Simple Hypothesis Testing (ASHT) problem with a given sample budget using a recent PDE characterization, and results from Stability analysis from theory of Viscosity solutions in PDEs. The ASHT problem is a difficult problem studied by diverse communities like Information Theory, Bandits and Operations Research, but a clean characterization of this optimal exponent has not been available until now.

Guarantees for Diffusion on Tilted Samples: w/ Sandeep Juneja, et. al. ***Submitted to ICML 2026***
Paper. First Author. We prove that it is possible to exponentially twist samples accurately using diffusion, under suitable conditions. Such optimization problems arise naturally in finance, weather and climate modelling and many other domains. We support these by extensive experimentation.

Polynomial Sparsity Bounds: w/ A. Chaudhury, V. Bhargava (**IIT Bombay**) *May 2025 — Ongoing*
On-going research project on improving the $d^2 \log(n)$ bound in the Sparsity Bounds conjecture. Building on this Paper. This problem has fundamental connections to Polynomial Identity Testing, an important theoretical CS problem. Recently, we managed to obtain the first n-independent bound for a difficult, special case.

Circuit Phenomenology Using Sparse Autoencoders: w/ David Udell, **SPAR** *June 2024 — July 2024*

Preprint. Sparse autoencoders enable interpretable representations of model activations, aiding mechanistic interpretability by uncovering causal circuits. We went through the literature, independently implemented circuit discovery for GPT-2-small and ended up finding big errors in the implementation of the algorithm in the latest paper by David Bau. We solved those issues, and experiemented with newer ideas, improving circuit discovery in GPT-2-small.

Selected Courses/Workshops

exSPLORe2026 Workshop: **SCDLDS**, Ashoka University *January 2026*

A workshop in Statistics, Probability, Learning and Optimization Research had two days of basic tutorials for graduates and advanced undergraduates, offering 3-hour sessions each on the basics of probability, learning, and optimization, led by leading experts. The following days had research workshops and talks from both academia experts and industry leaders.

DS: Probability and Optimization Methods: ICTS, Bangalore *August 2025*

Highly rigourous sessions on Game Theoretic MARL, Reinforcement Learning, Learning Theory, Diffusion and Flow Matching, Optimal Transport, Data Assimilation, Posterior Sampling, and Theory for Overparamterized Learning for faculties, graduate students and advanced undergraduates.

exSPLORe2025 Workshop: **SCDLDS**, Ashoka University *January 2025*

This workshop in Statistics, Probability, Learning and Optimization Research had two days of basic tutorials for graduates and advanced undergraduates, offering 3-hour sessions each on the basics of probability, learning, and optimization, led by leading experts. The following three days had a research workshop, with talks by leaders in the field.

CaMLAB: Cambridge AI Safety Hub *April 2024*

Course to build ML engineering fundamentals for AI Safety research. Includes basics of PyTorch, training and tuning GPTs and ResNets, interpreting models with TransformerLens, and an introduction to RL, RLHF.

Deep Learning: ISI Kolkata *January 2024 — March 2024*

A winter course on deep learning covering Autoencoders, CNNs, GANs, GNNs, Diffusion models, RNNs, Attention mechanics, Transformers, etc.

Measure Theory: Maths Club, ISI Kolkata *December 2023 — February 2024*

Introductory course on Measure Theory which helped me understand all the details in our Probability courses. Also independently studied Measure Theoretic Probability from **Rick Durette**.

Projects

Theoretical

Independence Is Almost Dependence:

Blog. Article. My independently discovered proof to a theorem: given two independent random variables X, Y you can come up with two new random variables U, V which have the same marginals and ϵ-close joints but are deterministically dependent.

Axiom of Choice, the Zorn's Lemma and their equivalence:

Blog. An expository article on Axiom of Choice, Zorn's Lemma and their equivalence.

Empirical/Programming

Ponderings on OthelloGPT:

Blog. Mechanistic Interpretability project. OthelloGPT is a GPT model trained on Othello games to predict all the possible legal moves. I look into how the model computes how a certain cell is blank.

ORIGAMI:

Repo. Implements arXiv:2303.17062, AISTATS 2023. A paper on dimensionality reduction of the support to improve computational efficiency in downstream decision making.

More experience

Volunteer and outreach

Owner: awas *October 2020 — October 2022*

Served as the **organizer and mentor** for daily math problem solving sessions, philosopical debates, and programming discussions for over two years on **Discord**. Mentored smart math enthusiasts, from all over India, learn hard math, and guided a few of them to get into the Indian training camp of IMO.

Owner: LML: Learning Machine Learning, ISI Kolkata *January 2024 — March 2025*

Organized extensive discussions on and around Machine Learning (Theory & Applied), Artificial Intelligence, and Statistics papers for more than a year.

Core Member: Theory CS Club, ISI Kolkata *March 2025 — Ongoing*

Engaged in extensive discussions on research papers on and around Theoretical CS focussing on Complexity Theory.