

DEBDEEP SANYAL

 [Email](#)  [LinkedIn](#)  [Portfolio Site](#)  [Google Scholar](#)

Research engineer working on foundation-model training and evaluation. Research spans time-series representations and LLM training behaviour, including policy optimization, adversarial training, and unlearning.

Publications and Preprints

- [1] **Debdeep Sanyal**, Aakash Sen Sharma, Dhruv Kumar, Saurabh Deshpande, Murari Mandal. [Policy Optimization Prefers The Path of Least Resistance](#). **EMNLP**, main conference, 2026.
- [2] **Debdeep Sanyal**, Umakanta Maharana, Yash Sinha, Hong Ming Tan, Mohan Kankanhalli, Murari Mandal. [OrgAccess: A Benchmark for Role-Based Access Control in Organization Scale LLMs](#). **EMNLP**, main conference, 2026.
- [3] **Debdeep Sanyal**, Agniva Maiti, Umakanta Maharana, Dhruv Kumar, Ankur Mali, C. Lee Giles, Murari Mandal. [Investigating Pedagogical Teacher and Student LLM Agents: Genetic Adaptation Meets Retrieval-Augmented Generation Across Learning Styles](#). **EMNLP**, main conference, 2025.
- [4] **Debdeep Sanyal**, Manodeep Ray, Murari Mandal. [AntiDote: Bi-level Adversarial Training for Tamper-Resistant LLMs](#). **AAAI**, 2026.
- [5] **Debdeep Sanyal**, Murari Mandal. [Agents Are All You Need for LLM Unlearning](#). **COLM**, 2025.
- [6] Aakash Sen Sharma, **Debdeep Sanyal**, Manodeep Ray, Vivek Srivastava, Shirish Karande, Murari Mandal. [The Realignment Problem: When Right becomes Wrong in LLMs](#). **ICML**, 2026.
- [7] Aakash Sen Sharma*, **Debdeep Sanyal***, Priyansh Srivastava, Shirish Karande, Mohan Kankanhalli, Murari Mandal. [Nine Ways to Break Copyright Law and Why Our LLM Won't: A Fair Use Aligned Generation Framework](#). **Findings of EMNLP**, 2025.
- [8] **Debdeep Sanyal**, Aaryan Nagpal, Dhruv Kumar, Murari Mandal, Saurabh Deshpande. [time2time: Causal Intervention in Hidden States to Simulate Rare Events in Time Series Foundation Models](#). **NeurIPS workshop, BERT2S**, 2025.
- [9] Aaryan Nagpal, **Debdeep Sanyal**, Murari Mandal, Dhruv Kumar, Saurabh Deshpande. [Mix, Don't Pick: Why Synthetic Corpus Composition Matters for Time Series Foundation Model Pretraining](#). **ICML workshop, FMSD**, 2026.
- [10] **Debdeep Sanyal**, Manya Pandey, Dhruv Kumar, Saurabh Deshpande, Murari Mandal. [Confidence is Not Competence](#). *arXiv Preprint*, 2025.

*Equal contribution.

Research Experience

Research Engineer | *Birla AI Labs*

April 2026 – Present

- Owned end-to-end training, evaluation, and inference for a **180M-parameter decoder-only time-series foundation model**, developed over five months.
- Engineered a distributed training pipeline with deterministic sampling, distributed data scheduling, dense forecasting supervision, and checkpoint recovery; pretrained on approximately **31.5B non-padding patch-token exposures** across **8 NVIDIA H100 GPUs** using PyTorch DDP.
- Designed and evaluated post-training recipes through controlled experiments and frequency-horizon diagnostics, outperforming **TimesFM 2.5** and most first-generation models across **97 GIFT-Eval tasks** while approaching second-generation models.
- Leading pretraining and architecture experimentation for **Bharat-TSFM**, a planned **500M-parameter foundation model** under the **IndiaAI initiative**, targeting sovereign compute.

Research Assistant | *Birla AI Labs*

July 2025 – March 2026

- Authored **time2time**, the lab's first paper, on causal interventions in hidden states of time-series foundation models.
- Ran architecture ablations to inform design choices for the lab's current foundation model.

Education

Kalinga Institute of Industrial Technology

2022–2026

B.Tech, Computer Science and Engineering

CGPA: 8.28/10

Technical Skills

Model development: PyTorch, Hugging Face Transformers.

Programming: Python, Java, C.