Srija Mukhopadhyay

✉ srija.mukhopadhyay@research.iiit.ac.in
sri-ja.github.io
sri-ja ajirs Google Scholar

Education

Nov 2021 – Present **International Institute of Information Technology Hyderabad**.
B.Tech in Computer Science and MS (by Research) in Computational Linguistics
- **CGPA:** 8.94/10.00 (Total); 9.28/10.00 (MS)
- **Dean's Academic List:** 2021–22 M (Top 5%), 2022–23 M (Top 10%), 2023–24 S (Top 15%)
- **Dean's Research List:** 2024–2025 (Top 5%)

Publications

2025 **PrivacyBench: A Conversational Benchmark for Evaluating Privacy in Personalized AI**.
S. Mukhopadhyay, S. Reddy, S. Muthukumar, J. An, P. Kumaraguru
Preprint

2025 **InterChart: Benchmarking Visual Reasoning Across Decomposed Charts**.
A. Iyengar*, **S. Mukhopadhyay***, A. Qidwai*, S. Singh, D. Roth, V. Gupta
AACL 2025 Website

2025 **MapIQ: Evaluating Multimodal Large Language Models for Map QA**.
V. Srivastava, F. Lei, **S. Mukhopadhyay**, V. Gupta, R. Maciejewski
COLM 2025 Paper

2025 **PRAISE: Enhancing Product Descriptions with LLM-Driven Structured Insights**.
A. Qidwai*, **S. Mukhopadhyay***, P. Khatiwada*, D. Roth, V. Gupta
ACL 2025 (System Demonstrations) Website

2025 **MAPWise: Evaluating Vision-Language Models for Advanced Map Queries**.
S. Mukhopadhyay, A. Rajgaria, P. Khatiwada, M. Shrivastava, D. Roth, V. Gupta
NAACL 2025 (Oral), ***Nominated for Outstanding Paper*** Website

2024 **Unraveling the Truth: Do VLMs really Understand Charts?**.
S. Mukhopadhyay*, A. Qidwai*, A. Garimella, P. Ramu, V. Gupta, D. Roth
EMNLP 2024 Website

Research Experience

Jan 2025 – Present **Precog Research Group, IIIT Hyderabad**.
Advisor: Prof. Ponnurangam Kumaraguru
Collaborators: Prof. Saurav Prakash (IIT-M), Prof. Sheetal Kalyani (IIT-M), Prof. Jisun An (Indiana University)
- Investigating fairness constraints in Federated Learning (FL) setups to mitigate algorithmic bias across heterogeneous clients.
- Optimizing the training of Binarized Neural Networks (BNNs) to enable efficient learning and adaptation on resource-constrained devices.
- Evaluating privacy risks in personalized agents to detect secret leakage and ensure alignment with Contextual Integrity frameworks.

Aug 2023 – Present **Language Technologies Research Center (LTRC), IIIT Hyderabad**.
Advisor: Prof. Manish Shrivastava *(MS Thesis Advisor)*

- Explored lateral model compression strategies to test the limits of parameter reduction in large-scale architectures.
- Conducted comparative analysis of diverse Large Language Models (LLMs) to map representational similarities and architectural divergences.
- **Current Research:**
  - Investigating non-autoregressive frameworks, like diffusion models, to align language generation more closely with human cognitive patterns.
  - Systematically diagnosing the root causes of performance degradation in long-context models to understand the limits of effective context utilization.

May 2025 – June 2025 **Whitebox AI Safety Fellowship**.
Mentor: Kyle Reynoso

- Conducted mechanistic interpretability research to diagnose the phenomenon of emergent misalignment in Large Language Models (LLMs) using cross-encoders.
- Analysed other diverse tasks that might lead to misaligned behaviors in models.
- Was recognized as a Top 3 Project in the cohort.

Jan 2024 – Aug 2025 **Cognitive Computation Group (UPenn) / CORAL Lab (ASU)**.
Advisors: Prof. Dan Roth (UPenn), Prof. Vivek Gupta (ASU)
Collaborators: Dr. Aparna Garimella (Adobe), Pritika Ramu (Adobe), Prof. Ross Maciejewski (ASU)

- Investigated the robustness and consistency of Vision-Language Models (VLMs) in interpreting complex visual information, specifically charts, graphs, and geospatial maps.
- Benchmarked the robustness of VLM reasoning capabilities on different kinds of charts, including multiple charts, identifying failure points.
- Developed "MAPWise" and "MapIQ" benchmarks to evaluate multimodal models on advanced geospatial queries, counterfactual reasoning, and map-reading literacy.

Jan 2024 – May 2024 **Software Engineering Research Center (SERC), IIIT Hyderabad**.
Advisor: Prof. Karthik Vaidhyanathan

- Investigated automated code summarization techniques to extract functional natural language descriptions from code snippets, benchmarking on the APPS dataset.
- Explored the integration of Abstract Syntax Trees (ASTs) to enrich model inputs with hierarchical structural context, aiming to enhance the semantic accuracy of generated summaries.

Industry Experience

Aug 2025 – Present **Enterpret**.
Role: Machine Learning Research Intern

- Enhancing the taxonomy retrieval system to extract actionable insights from large-scale customer feedback data.
- Leveraging Large Language Models and Retrieval Augmented Generation to improve the granularity and semantic accuracy of feedback classification.

May 2025 – Aug 2025 **Google**.
Role: Software Engineering Intern Team: Google Cloud Storage (GCS) Orion

- Created a topology visualization tool for GCS Orion, enabling engineers to isolate and debug critical network issues across 100,000+ nodes.
- Received Pre-Placement Offer (PPO) for Full-Time Software Engineer role.

May 2024 – July 2024 **Google**.
Role: Software Engineering Intern Team: Android PDF

- Designed and implemented the core image selection pipeline for the Android PDF viewer.
- Received Return Internship Offer for Summer 2025 based on exceptional performance.

May 2023 – July 2023 **Google**.
Role: STEP Intern Team: Google Cloud Storage (GCS) Metering

- Built integrity checks for a high-throughput pipeline processing petabytes of Cloud usage data, ensuring billing accuracy for enterprise customers.
- Converted to standard Software Engineering Intern (Return Offer) for Summer 2024.

Honors and Awards

- **India AI Fellowship (2025):** Fellowship and aid to conduct research on LLM Similarity.
- **Amazon ML Challenge (2024):** Secured **13th Rank** nationally among ∼75,000 participants (Top 0.02%).
- **Google Generation Scholarship (2024):** Awarded for academic excellence and leadership in technology.
- **JPMorgan Chase Quant Research Mentorship (2024):** Selected for the specialized quantitative research track.
- **D. E. Shaw Ascend Educare (2022):** Selected as a Fellow for the prestigious mentorship cohort as one of the youngest participants.

Teaching and Leadership

Teaching Assistantships

Spring 2025, Spring 2024 — **Head Teaching Assistant**, *Introduction to Natural Language Processing (NLP)*.
Instructor: Prof. Manish Shrivastava

Monsoon 2024 — **Teaching Assistant**, *Language and Society*.
Instructor: Prof. Aditi Mukherjee

Monsoon 2023 — **Teaching Assistant**, *Algorithm Analysis & Design*.
Instructor: Prof. Suryajith Chillara

Volunteership

July 2023 — **Team Leader & Mentor**, *International Linguistics Olympiad (IOL)*, Bansko, Bulgaria.
Selected to lead and mentor the Indian National Team at the global finals.
- Trained students on advanced linguistic puzzle-solving techniques during the training camp.
- Helped manage team strategy and logistics during the competition in Bulgaria.

May 2023 – Dec 2024 — **Lead Student Coordinator**, *Panini Linguistics Olympiad (PLO)*, India.
- Orchestrated the national-level selection process for the Indian team, managing logistics for thousands of participants and coordinating with regional centers.
- Helped conduct the Asia Pacific Linguistics Olympiad in India to select the Indian team.

Mar 2023 – Feb 2024 — **Club Coordinator**, *Open Source Developers' Group (OSDG), IIIT Hyderabad*.
- Organized an intra-college beginner hackathon attracting over 150 participants from the first and second years.
- Conducted technical workshops on Linux systems, Open Source contributions, and Google Summer of Code (GSoC) mentorship.