

SIDDHANT GUPTA

+91-7836015759 ◊ writetosiddhant@gmail.com ◊ <https://sidworks01.github.io/>

 [LinkedIn](#)  [Google Scholar](#)  [GitHub](#)  [Hugging Face](#)

EDUCATION

Indian Institute of Technology Roorkee
B.Tech in Mechanical Engineering

2022 - 2026
Overall GPA: **7.73/10**

PUBLICATIONS

- Improving Multilingual Capabilities in LLMs with Cultural Knowledge** [ORAL] [Paper] [Poster] [PPT]
Ram Kadiyala, S Pullakhandam, **Siddhant Gupta**, et al.
Advised by: **Cohere Labs**, Gave the oral [AAACL'25] - Main (Oral)
- Uncovering Cultural Representation Disparities in Vision-Language Model** [Paper] [Poster] [PPT]
Siddhant Gupta, Ram Kadiyala, Jebish Purbey, Srishti Yadav, Alejandro Salamanca
Advised by: **Dr. Desmond Elliott** [AAACL'25] - Findings
- Multilingual Hate Speech Detection and Target Identification in Devanagari Languages** [Paper]
Siddhant Gupta, Siddh Singhal, Azmine Toushik Wasi [COLING'24] - Workshop
- SeqWen at the Financial Misinformation Detection Challenge** [Paper]
Jebish Purbey, **Siddhant Gupta**, Nikhil Manali, Siddhartha Pullakhandam, Drishti Sharma,
Ashay Srivastava, Ram Mohan Rao Kadiyala [COLING'24] - Workshop
- Spot the BlindSpots: Quantifying Fine-Grained LLM Biases in Contact Center Summaries** [Paper]
Kawin Mayilvaghanan, **Siddhant Gupta**, Ayush Kumar
Internship at **Observe.ai** [EMNLP'25] - Industry Track
- DSBC : Data Science task Benchmarking with Context engineering** [Paper] [Poster] [PPT]
Ram Kadiyala, **Siddhant Gupta**, Jebish Purbey, Giulio Martini, Suman Debnath
Advised by: **Dr. Hamza Farooq** [AAACL'25] - Main
- If the Yardstick is Bent: Counterfactual Fairness Evaluation of LLM-Based Contact Center Agent Quality Assurance System** [Paper]
Kawin Mayilvaghanan, **Siddhant Gupta**, Ayush Kumar
Internship at **Observe.ai** [Under Review]
- Robust and Fine-Grained Detection of AI Generated Texts** [Paper]
Ram Kadiyala, Siddhartha Pullakhandam, **Siddhant Gupta**, Jebish Purbey [Under Review]
- Lexical Reranking of Semantic Retrieval (LeSeR) for Regulatory Question Answering** [Paper]
Jebish Purbey, Drishti Sharma, **Siddhant Gupta**, K Murad, S Pullakhandam, R Kadiyala [COLING'24]
- Evaluating Generalization Capabilities of LLM-Based Agents in Mixed-Motive Scenarios Using Concordia** [Paper]
{20+ authors}, **Siddhant Gupta**, et al. [NeurIPS'25]
- Global PIQA: Evaluating Physical Commonsense Reasoning Across 100+ Languages and Cultures**
One of the authors [Paper]

EXPERIENCE

ML Researcher at Massachusetts Institute of Technology (MIT) Nov 2025 - Present
Built a JSON workflow framework translating natural-language tasks into computation graphs.

- Developed a subsystem to structure user insights into datasets and benchmark LLM and multi-agent workflows.

NLP Community Lead at **Cohere Labs**

June 2023 – Present

Led research under **Expedition Aya (2024–25)** on reasoning evaluation, cultural bias in VLMs, speech/ASR, ML-generated text detection, and multilingual LLMs.

- Built a DAG-based open-source agentic framework for **Cohere Labs**; collaborated globally on multi-agent evaluation, A2A orchestration, and shared memory modules in the **ML-Agent** program.
- Spoke at **Cohere ML Summer School** on Transformer & LLMs; organized 50+ NLP focused talks and paper-reading sessions in the **NLP** program.

ML Intern at **Observe.ai**

May 2025 - Nov 2025

Engineered and evaluated bias-detection for LLMs across 18 models and 3,000+ transcripts using new fairness metrics.

- Improved Auto-QA scoring pipelines, reducing misclassification by 12–18% and enhancing output reliability

Research Fellow at **Traversaal.ai**

Jan 2025 – Apr 2025

Benchmarking data-science agents - automated graph generation with ReAct based agent workflow and design

PROJECTS

Cadingo - Design Automation for Mechanical Workflows

Jul 2025 - Present

B.Tech Thesis

- Developed software automating mechanical design workflows through a single prompt and human feedbacks, including custom simulation models and generating product design manufacturing and assembly documentation
- Built workflow modules integrated for engineering pipelines, enabling copilots to streamline iterative design tasks.
- Integrated MCP server for workflow executions with engineering tools like AutoCAD, ANSYS for analysis.

SpeechAya: Multilingual Speech Synthesis via Unified LLM Architecture

Sept 2023 - Nov 2023

- Engineered a multilingual LLM pipeline integrating speech and text modalities using 1000+ hours of audio.
- Optimized tokenization using mHuBERT; achieved 0.112 WER on PolyAI/minds14 dataset with Qwen2-1.5b model

RELEVANT KEY COURSES AND SKILLS

Machine Learning

Stanford CS229 (Machine Learning), Stanford CS231n (Computer Vision), Deep Learning Specialization, Data Mining, Artificial Intelligence

CS & Math

Programming and Data Structures, Linear Algebra, Probability and Statistics

Libraries and Frameworks

PyTorch, TensorFlow, HuggingFace Transformers, PEFT, TRL, BitsAndBytes, FlashAttention, xFormers, Triton, Sklearn, Librosa, NLTK, OpenCV, NumPy, Pandas, Matplotlib, Gradio, ONNX Runtime

POSITION OF RESPONSIBILITIES

- Software head at **Artificial Intelligence and Electronic Society (ArIES), IITR** *May 2023 - Present*
- NLP Lead and Leading Community Projects at **Cohere Labs** *June 2023 - Present*

REFERENCES

- **Sara Hooker** - Ex-Google Research Scientist, Ex-VP at Cohere Labs
- **Suman Debnath** - Technical Lead at Anyscale, Ex- Principal Developer Advocate at Amazon Web Services
- **Marzieh Fadaee** - Head at Cohere Labs