Pratyush Mohanty

Bengaluru, India pratyushmohanty2912@gmail.com +91 6263283117 Linkedin Github

RESEARCH INTERESTS

Machine Learning · Natural Language Processing · Large Language Models · Multi-Agent Systems · Distributed Systems

EDUCATION

M.S. Ramaiah Institute of Technology, Bengaluru **CGPA: 8.53 / 10** *Oct 2023 – June 2027*

B.E. Electronics & Communication Engineering

Relevant Coursework: Machine Learning, Deep Learning, Linear Algebra, Probability & Statistics, Signal Processing, Robotics, Data Structures & Algorithms

EXPERIENCE

Research Intern: *Centre for Artificial Intelligence and Robotics (CAIR), DRDO* *September 2026 – Present*

- Engineering a decoupled, two-phase NLP framework for robust sentence segmentation on unformatted text (ASR streams, chat logs, OCR scans, and web remnants).
- Developed a 21-rule **Tier 1 Heuristic Classifier** achieving **0.22 ms latency** (**4,400 samples/sec**) and **98.3% precision** on clean text to triage pristine inputs with zero external NLP overhead.
- Evaluated a **Tier 2 XGBoost** classifier via **GroupKFold**; diagnosed shortcut learning (*Clever Hans* effect) where shallow n-gram models failed out-of-domain transfer (F1 dropped to 20%), proving the need for contextual representations.
- Curated multi-domain benchmarking corpora across 5 noise distributions (**LibriSpeech**, **ICDAR**, **WNUT-17**), enforcing strict holdouts and MinHash deduplication to eliminate target leakage.
- Currently integrating lightweight contextual models (**DeBERTa-v3**) into a cascading router and developing Phase 2 token-level sequence segmenters to replace brittle punctuation-dependent splitters.

ClawSwarm: *Independent Research Project* GitHub

- Designed **ClawSwarm AI**, an open-source multi-agent framework that dynamically learns agent specializations via 6D Exponential Moving Average (EMA) skill vectors, avoiding the rigid, hardcoded roles of frameworks like AutoGen and CrewAI.
- Formulated task routing as a multi-armed bandit problem using a **Skill-Adaptive ϵ-greedy policy** with entropy-based decay to optimize dynamic task-to-agent matching.
- Engineered a **heartbeat fault-detection and leader-election protocol** (2s polling) that isolates crashed or degrading nodes, reassigns in-flight tasks, and updates routing priors via failure penalties.
- Conducted a **50-seed simulation study** (1,000 tasks/run), achieving a **63% reduction in cumulative regret** relative to skill-agnostic baselines (paired Wilcoxon signed-rank $p < 0.001$) and characterizing exploration costs against greedy baselines.
- Implemented a **cross-model LLM-as-judge evaluation pipeline** (using distinct model providers) to mitigate self-preference evaluation bias; open-sourced the codebase with reproducible benchmarks.

Multimodal Deepfake Detection: *Undergraduate Research Project* GitHub

Advisor: Dr. S. Sethu Selvi (Professor, Dept. of E&C Engineering)

- Implemented an audio-visual deepfake detector in a 4-member team, based on the **Synchronization-Aware Feature Fusion (SAFF)** and **Cross-Modal Graph Attention Network (CM-GAN)** approach (*IEEE Trans. Multimedia*, 2024). The model uses an **EfficientNet-B0** visual encoder and a CNN audio encoder on **Mel-spectrograms**, with a cosine-similarity synchronization score, transformer-based cross-modal attention, and gated fusion for real/fake classification.
- Trained and tested on **LAV-DF** (100K+ English videos). The combined SAFF + CMGAN model reached **94.57% accuracy, 96.39% F1, and 0.979 AUC**, compared with 92.36% and 93.96% accuracy for the SAFF-only and CM-GAN-only variants.
- Evaluated the LAV-DF-trained model on Hindi videos (**HAV-DF**, 67K+) without fine-tuning: **80.12% accuracy, 85.49% F1, 0.844 AUC**. The drop from in-dataset results shows the model is sensitive to language and dataset shift.
- Worked on the preprocessing pipeline (uniform frame sampling, 16 kHz audio extraction, 128-bin Mel-spectrograms) and on the model comparison experiments.

PROJECTS

NexusAI — *TypeScript · Node.js · Redis · NLP · MongoDB* GitHub

- Built an asynchronous email-to-schema extraction system using LLM zero-shot extraction with a trained Naive Bayes fallback for low-confidence predictions, with a Redis/BullMQ event-driven queue to decouple model inference latency from client requests.

TECHNICAL SKILLS

Languages: Python, C++, JavaScript, TypeScript, SQL

ML / Research: PyTorch, scikit-learn, XGBoost, Hugging Face Transformers (DeBERTa), EfficientNet, Graph Attention Networks, Multimodal Fusion, Multi-Armed Bandits, LLM-as-Judge Evaluation, Statistical Testing (Wilcoxon signed-rank)

Systems & Infra: FastAPI, asyncio, Redis, Docker, MongoDB, PostgreSQL, Git, GitHub Actions (CI/CD), Linux

Research Tools: LaTeX, Weights & Biases, Hugging Face Transformers, FAISS, Jupyter

ACHIEVEMENTS & RECOGNITION

- **LeetCode:** Rating **1755**; **400+ problems** solved across competitive programming platforms.
- **HackAWar Hackathon** — **3rd Runner-Up** (top 5 overall); **Whackiest Hackathon** — **Top 10 Finalist & Theme Winner**.
- **JEE Main** – 96th percentile among 1.2 million candidates