

RIJUSMIT BISWAS

New Delhi, India | rijusmit.biswas@gmail.com | +91 8929693930 | linkedin.com/in/rijusmit-biswas | github.com/riju-talk | rijusmit.vercel.app

SUMMARY

Final-year CS undergrad (B.Tech, Sept 2026) working on machine learning under real deployment constraints. I build multi-channel satellite imagery pipelines for crop classification and yield estimation, and measure how far those models degrade when capture conditions change across districts — separating real signal from capture artifacts. I train models under 4-bit quantization inside memory envelopes that match low-end hardware, and build the engineering around the model: evaluation harnesses, calibration, FastAPI inference services, and CI.

EDUCATION

B.Tech, Computer Science & Engineering, IIIT-Delhi | Coursework: Linear Algebra, Probability & Statistics, Optimisation, ML **Nov 2022 to Sept 2026**

EXPERIENCE

Satellite Imagery for Agricultural Prediction — ML (Research) Intern, CSS Lab (Econometrics), IIIT-Delhi **Jan 2025 to May 2025**

- **Pipeline:** built the full computer vision pipeline on Google Earth Engine multi-threaded preprocessing, segmentation mask generation, and fusion with spectral bands to produce multi-channel inputs.
- **Benchmarking:** compared CNN backbones across three downstream tasks from the same imagery phenology date prediction, crop type classification, and yield estimation under one shared evaluation harness to keep the comparison fair. Best configuration reached **0.87 macro-F1** on crop type classification, against a ResNet-18 baseline.
- **Generalization:** quantified performance degradation under varying capture conditions, measuring domain shift across 10+ districts with difference-in-differences and fixed-effects models to separate real signal from capture artifacts. Accuracy dropped **11 points** from in-district to out-of-district evaluation, with cloud-contaminated post-monsoon captures as the main driver.

Long-Context Evaluation & Preference Fine-Tuning — ML Engineer / UG Researcher, MIDAS Lab, IIIT-Delhi **May 2025 to Jan 2026**

- **Training:** built training pipelines across three datasets producing 24 checkpoints with LoRA and QLoRA; merged adapters into base weights and published the artifacts for inference.
- **Constrained training:** trained Qwen3-4B under BitsAndBytes 4-bit quantization within **24 GB** of GPU memory the same envelope that governs memory-constrained and edge deployment.
- **Evaluation:** designed evaluation harnesses for long-context benchmarking; fine-tuned with SFT, DPO, and GRPO and measured behavioral change against held-out baselines. SFT set a **68%** needle-in-a-haystack baseline at 32k context; GRPO added **+12 points** of reasoning-chain accuracy while holding retrieval F1 flat, whereas DPO gained **+8 points** of helpfulness at a **5-point** retrieval cost an explicit alignment tradeoff, not a single headline number. First-author paper under review at AAI 2027; harness and benchmarks open-sourced.

Noisy Real-World Multi-Class Classification — Data Science Intern, Complex Systems Lab, IIIT-Delhi **Aug 2024 to Dec 2024**

- Led a team of five building a 25-class classifier on noisy real-world data, reaching **84.5%** and **74.5%** accuracy on two splits; traced the 10-point gap to distribution shift between splits rather than model capacity. Ran code review and mentored junior members.

Software Engineering Intern, Shipzee **March 2026 to Aug 2026**

- Wrote unit and integration tests covering internal services and third-party APIs; diagnosed live production issues by reading logs and tracing requests across services.

PROJECTS

TURBO — Target-Augmented Shared Fusion for Multimodal Generation | PyTorch, ViT, BART, HuggingFace Transformers

- Independently reimplemented Goel et al. (arXiv:2502.07391): cross attention fusion between Vision Transformer image features and a modified BART text encoder, trained on the MORE+ dataset.
- Owned the full pipeline — feature extraction, target-conditioned representations, teacher-forced training, and BLEU / ROUGE / METEOR evaluation. Reproduced reported BLEU within **2.3 points**; qualitative error analysis showed failures concentrated in attribute binding (color, count, and spatial-relation mismatches).

MineMEETS — Multimodal Embedding & Retrieval Pipeline | CLIP ViT-B/32, Whisper, Pinecone, Docker, GitHub Actions, Pytest

- Built ingestion pipelines generating CLIP ViT-B/32 visual embeddings alongside text and audio behind one unified interface, with 512-dimension validation before upsert and deterministic, idempotent reprocessing.
- Set up the Pinecone vector store with namespace isolation and metadata filtering for similarity search, enabling per-item reindexing and safe rollback.
- Dockerized the service and added GitHub Actions CI running Pytest, Black, Pylint, and MyPy on every pull request; median query latency **85 ms** across **18,000** indexed items.

PUBLICATIONS

- **HorusLLM: A Cross-Evaluation Protocol for Tri-Vector Alignment Interference.** R. Biswas (first author), A. G. Kumar, A. Anand, E. Cambria, Z. Wang, R. R. Shah. *Under review, AAI 2027*. Built the evaluation harness and benchmarks; open-sourced.

TECHNICAL SKILLS

Computer Vision & Geospatial: PyTorch, OpenCV, CNN architectures, segmentation masks, multi-band spectral fusion, Rasterio, GeoPandas, Google Earth Engine

Deep Learning & Efficiency: HuggingFace Transformers, PEFT (LoRA / QLoRA), TRL, Accelerate, BitsAndBytes 4-bit quantization, cross attention fusion, custom training loops

Evaluation & Causal Inference: F1, ROC-AUC and PR-AUC, probability calibration, leakage-aware experiment design, SHAP, automated evaluation harnesses, difference-in-differences, fixed effects

Engineering: Python, C++, SQL (PostgreSQL), FastAPI model serving, Docker, CI/CD (GitHub Actions), Linux, PySpark, ETL pipelines

LEADERSHIP & ADDITIONAL

- Head of Product Development, IIIT-Delhi Student Council (8 engineers); Vice Chair, IEEE Student Branch. ML & DL Specializations (Stanford / DeepLearning.AI). Codeforces Specialist (1403), 2500+ LeetCode.