Sanjit Jeevanand

London, UK — +44 7467 338057 — sanjitjeevanand@gmail.com
linkedin.com/in/sanjit-jeevanand — github.com/Sanjit-Jeevanand/
sanjit-ml.com

Education

University College London (UCL) — MSc | Discipline: Computer Science | *Predicted Distinction*
Indian Institute of Technology (IIT) Kharagpur — Dual Degree | Discipline: Engineering | *GPA: 8.0 / 10*

Projects

UK Finance Domain Intelligence System (RAG) — *Production on Google Cloud Run*

- Architected RAG system enabling financial analysts to query **2,000+ pages** across 8 UK-listed companies' annual reports with citation-backed answers, eliminating hours of manual search.
- Engineered deterministic ingestion pipeline with FAISS-based semantic retrieval using SentenceTransformers, achieving **sub-second query response times**.
- Implemented strict source attribution to prevent hallucinations critical for audit and compliance workflows.
- Deployed with FastAPI, Docker, GitHub Actions CI/CD, and serverless autoscaling on GCP Cloud Run.

Online ML Inference & Drift Monitoring Platform — *Production on AWS*

- Built end-to-end ML platform: feature contracts, offline training, versioned deployment, real-time inference, drift detection, and automated retraining.
- Established artifact-driven versioning with auditable promotion gates and rollback, achieving **zero training–serving skew**.
- Deployed low-latency FastAPI inference with statistical drift monitoring (KS, PSI), triggering retraining when shifts exceed thresholds.

Multi-Agent Decision Orchestration System (GenAI) — *Production on Google Cloud Run*

- Built multi-agent LLM system for complex technical decisions (architecture trade-offs, deployment strategies) using orchestrated specialist agents with structured reasoning.
- Architected parallel execution coordinating planner, specialists, critic, and synthesizer with deterministic state management and complete audit trails.
- Developed WebSocket-based real-time UI for live visualization of agent interactions and decision convergence.
- Optimized LLM costs through prompt engineering, achieving **$0.02 per decision cycle** while maintaining quality.

Research & Experience

MIT-Affiliated Research Collaboration *Aug 2024 – Present*

- Developed LLM-assisted pipelines for numeric forecasting from text, implementing preprocessing, feature extraction, and regression modeling.
- Benchmarked classical vs. LLM-based approaches, demonstrating measurable improvements in quantitative reasoning accuracy.

University of Alberta — Machine Learning Intern *May 2024 – Aug 2024*

- Designed ML-driven UI for automated laser imaging using PyQt, applying HCI principles and iterative testing.
- Achieved **25% improvement** in usability scores, reducing operator training time and measurement errors.

Dalhousie University — MITACS Full Stack Development Intern *May 2023 – Aug 2023*

- Engineered backend analytics and personalization algorithms for AI-driven mental health platform (featured in 5+ media outlets).
- Built production data pipelines (Node.js, MongoDB) tracking engagement and intervention effectiveness.

BDO India — Resolution Advisory Intern *Jun 2022 – Jul 2022*

- Developed financial models and risk analysis supporting multi-million-dollar M&A transactions and restructuring.
- Built scenario frameworks quantifying downside risk and recovery outcomes, informing deal structuring.

Achievements & Leadership

President, Naval Architecture & Ocean Engineering Society, IIT Kharagpur *2023–2024*

- Led 40-member society, secured 5+ industry partnerships, organized symposium with 500+ participants.

Project Lead, Computer Graphics Society, IIT Kharagpur *2022–2023*

- Directed cross-functional team developing medieval RPG game, managing design and engineering to deliver production builds.

Winner, UCL Vibe Coding Hackathon *2025*
Top 0.2%, Kharagpur Data Science Hackathon (5,818 participants) *2024*