

ANKAN DEBNATH

+91 7595988801 | [Portfolio](#) | ankandebnath12b@gmail.com | [GitHub](#) | [LinkedIn](#)

EDUCATION

- PES University, Bangalore** 2022 - 2026
B.Tech, Computer Science and Engineering
- Recipient of **Distinction award**, Computer Science, PES University
- Pathfinder Higher Secondary Public School, Kolkata** 2020 - 2022
12th WBCHSE

TECHNICAL SKILLS

- **Languages:** Python, Go, JavaScript/TypeScript (Node.js, React, vue), C++, C
- **Databases:** PostgreSQL, MySQL, MongoDB, Redis (caching, queues), Prisma/ORM; schema design & indexing.
- **Data Science & ML Libraries:** Data pipelines / ETL, Pandas; ML/fraud detection concepts, LLMs & RAG: embeddings, vector stores (FAISS/Weaviate), agentic workflows, ChatGPT, Hugging Face, LangChain, PySpark
- **Cloud & DevOps:** AWS (EC2, S3, Lambda), Docker, basic Kubernetes; CI/CD: GitHub Actions; Terraform basics
- **Backend & Architecture:** RESTful APIs, gRPC, Webhooks, Kafka, microservices design, idempotency & retries.
- **Tools & Workflow:** Git, Postman, Linux, Prometheus, Grafana, ELK stack, Sentry

EXPERIENCE

- Fullstack-AI/ML Intern – Transaction Analyst Pvt. Ltd.** 2026(March - Present)
- Worked on the eKYC platform, contributing to the integration of AI/ML features into the workflow.
 - Built and improved frontend UI components and web interface elements for a smoother user experience.
 - Created and maintained technical documentation for application features, processes, and frontend implementation.
- Research Intern - Centre of Cognitive Computing & Computational Intelligence, PES UNIVERSITY** 2024(jun - aug)
- Built ML pipeline for fake news detection across 5 Indian languages, achieving 87% accuracy on 5K+ articles
 - Engineered translation model processing 50K+ daily articles, improving detection coverage by 60%
 - Developed evaluation framework across 8 linguistic patterns, boosting precision from 72% to 87%
 - Published 15K+ labeled dataset on Kaggle for research collaboration

PROJECTS

- 1) Real-time speech-to-text (STT) processing system** *Python, Faster-Whisper, WebRTC VAD, ECAPA-TDNN, INT8 quantization*
- Productionized a low-latency STT pipeline using Faster-Whisper with INT8 quantization, achieving 1–2s end-to-end transcription latency for real-time commercial workflows.
 - Integrated speaker verification (ECAPA-TDNN) and WebRTC VAD to generate speaker-attributed transcripts with ~500ms identification latency, supporting secure and compliant deployments.
 - Implemented event-driven architecture with Docker, CI/CD, and observability (latency, WER, throughput) to ensure reliable, monitorable ML services for downstream business applications. [Link](#)
- 2) Longview-AI -- AI Agent with human-like memory** *Streamlit, AutoGen, Zep Cloud, Ollama, Qwen 3, Secure Secrets, CI/CD*
- Built a production-ready RAG conversational service (LLM + vector store) exposing a FastAPI-style scoring & retrieval API and Dockerized deployment for integration with CRM and sales apps.
 - Implemented session-linked memory, secure secrets handling, and access controls to meet regulatory & data-isolation needs for fintech use-cases, enabling safe multi-tenant usage.
 - Delivered business metrics and monitoring: reduced average manual reply time by 65% and achieved a suggested-reply acceptance rate of 74%; added logging, usage analytics, and retraining hooks to support product iteration. [Link](#)
- 3) NomadIQ - multi-agent hotel booking crew** *Python, CrewAI, DeepSeek-R1 | Ollama, OpenAI GPT-4, Browserbase, Playwright, Docker*
- Engineered a real-time pricing pipeline (scrapers → Playwright → ETL → feature store) feeding a pricing model that produced cost-saving recommendations and dynamic price suggestions consumed by downstream apps.
 - Exposed model outputs via REST API and built stakeholder dashboards (Power BI / Tableau) for revenue ops and pricing teams to visualize lift, adoption, and A/B test results.
 - Focused on efficiency & compliance: implemented secure API management, local LLM support for PII-sensitive flows, and monitoring that reduced pricing rollback incidents by 40%. [Link](#)
- 4) Developing Adaptive Self-Driving Cars Using Liquid Neural Networks and Reinforcement Learning**
CARLA, Liquid Neural Networks, DDPG, CNN
- Led end-to-end experiment lifecycle in deterministic CARLA simulations: model training, automated CI runs, deterministic validation suites, and telemetry logging to detect regressions early.
 - Improved control robustness by integrating uncertainty-aware perception (Bayesian segmentation) into the control loop; reduced instability/edge failures by ~20% during validation.
 - Packaged experiments into reproducible pipelines (Docker + orchestration) with clear retraining triggers and performance monitoring to enable safe model rollouts and rollback policies.