Sneha B Gawali

+91-8762621757 • sneha.gawali19@gmail.com • LinkedIn • GitHub • Bangalore

Experience

GRAG AI Factory Bangalore, India
AI Research Analyst Intern Apr 2026 – Aug 2026

- Engineered a clinical ML pipeline on de-identified screening records; built scalable curation protocols addressing missing attributes, class imbalances, and cross-registry noise.
- Trained gradient-boosted and deep learning classifiers with stratified k-fold evaluation (Recall, ROC-AUC); packaged inference as containerized FastAPI/Docker endpoints for healthcare triage.

Rooman Technology Remote, India
Machine Learning Intern Feb 2025 – May 2025

- Built end-to-end ML pipelines (preprocessing, feature engineering, benchmarking) and deployed regression, classification, and clustering models via REST APIs to AWS Cloud; translated outputs into operational insights for technical and business stakeholders.

Projects

DisputeFlow — Payment-Dispute Decision-Support Platform Python | FastAPI | LangGraph | XGBoost | PostgreSQL | Docker *2026*

- Built a LangGraph-orchestrated workflow for chargeback/dispute triage: intake classification, evidence-gap analysis, calibrated XGBoost win-probability scoring, policy-constrained LLM strategy, and human-in-the-loop escalation for high-risk cases.
- Enforced strict evidence guardrails so the LLM cannot infer unsupported facts (a shipping address is never treated as delivery proof); added hash-chained, HMAC-signed audit logging and a governed outcome-feedback loop for controlled retraining.

SentryEval — Adversarial Evaluation Platform for LLM Agents Python | Multi-LLM APIs | Streamlit | CLI *2026*

- Built a multi-model red-teaming harness testing LLMs across four threat surfaces — direct/indirect prompt injection, jailbreaks, and an original context-confusion surface — with side-by-side comparison across providers.
- Reported Wilson-score 95% confidence intervals on every attack-success metric and gated results on Cohen's kappa judge agreement ($\kappa \geq 0.6$); engineered a reproducibility contract with SHA-256 hash-pinned datasets, deterministic seeding, and a structured "model graveyard" of transcripts for failure review.

IncidentRAG — Evidence-Grounded Incident-Triage Agent Python | FastAPI | Qdrant | Neo4j | AWS *2026*

- Built a hybrid-retrieval RAG pipeline (BM25 + Qdrant dense search + reciprocal-rank fusion + cross-encoder reranking) with Neo4j dependency-graph traversal for service blast-radius analysis during incidents.
- Required every root-cause claim to cite retrieved evidence via a grounding-verification pass, paired with a risk classifier and approval gate that blocks high-risk remediation; instrumented with OpenTelemetry and shipped staged AWS deployment assets (EKS, Helm, Terraform).

Technical Skills

Languages: Python, SQL, C++, JavaScript/TypeScript
LLM & Agent Engineering: LangChain, LangGraph, RAG, tool calling, prompt engineering, structured outputs, guardrails
ML & Evaluation: PyTorch, Scikit-learn, XGBoost, feature engineering, model calibration, Wilson CI
AI Security & Observability: Prompt-injection defense, red-teaming, audit trails, OpenTelemetry
Backend, Data & Cloud: FastAPI, REST APIs, PostgreSQL, Docker, AWS, Git, GitHub, Streamlit

Education

Maratha Mandal's Engineering College, VTU Belagavi, India
B.E. in Robotics and Artificial Intelligence, CGPA: 8.2/10 2022 – 2026

Leadership & Achievements

- Founder, Women in AI/Robotics Engineering Collective — designed ML and embedded systems workshop curricula, led project mentorship, and built a structured hackathon participation pipeline to grow female representation in competitive AI.
- Finalist, Hack[AI]thon (Sphere Hive × VRIF, VTU Belagavi) — national Data-Centric AI Challenge focused on dataset quality enhancement and model training optimization via 3LC.ai.
- Publication: NeuroDrive hybrid BCI coupling EEG classification with real-time computer vision for hands-free vehicle navigation, *IJCRT*, Nov 2025.