

Sahojit Karmakar

Jalandhar, Punjab 144411

+91-6003441373 • sahojtxd26@gmail.com • [linkedin.com/in/sahojit-karmakar](https://www.linkedin.com/in/sahojit-karmakar) • github.com/Sahojit

TECHNICAL SKILLS

- **Programming Languages:** Python, Rust
- **Technical Skills:** Scikit-learn, XGBoost, PyTorch, HuggingFace, LangChain, Polars, FastAPI, Streamlit, Flask, Docker, GitHub Actions, MySQL, ChromaDB, AWS, Claude API, Context Engineering, RAG, GitHub, VS Code

EXPERIENCE

Data Analyst Intern | Futureense Technologies

Jun'25 – Aug'25

- Analyzed **20K+** marketing leads across Facebook/Google; evaluated CTR, lead quality, and conversion trends using Python, improving ad spend efficiency.
- Cleaned and merged **5K+** campaign rows with 40+ lead attributes; delivered actionable insights for MBA & Tech programs, improving data accuracy by **30%**.
- Built Power BI dashboards tracking campaign KPIs and ROI metrics, enabling the team to optimize targeting strategies weekly.

PROJECTS

[AutonoML – Multi-Agent AutoML Platform](#) | *FastAPI, Streamlit, ChromaDB, Ollama, Python*

Mar'26 – May'26

- Built a **3-agent RAG pipeline** (Planner → Executor → Evaluator) using FastAPI, FAISS, and Ollama (llama3.2); autonomously plans, retrieves, and generates answers grounded in personal study notes with **sub-2s query latency**.
- Engineered a **self-correction loop** where hybrid Evaluator (rule-based + LLM judge) scores responses and retries with targeted failure feedback up to **3 iterations**, significantly reducing hallucinations on domain-specific content.
- Designed a **fully local, modular architecture** with clean separation across FAISS vector search, SQLite memory, and sentence-transformer embeddings.

[AIOps Root Cause Analysis System](#) | *XGBoost, SHAP, FastAPI, PostgreSQL, Docker*

Apr'26 – May'26

- Engineered **96 features** from 800 LEMMA-RCA incident records; XGBoost classifier achieved **99.4% accuracy** across 6 root-cause classes.
- SHAP explainability surfaced top-3 causal features per incident, reducing mean engineer triage time by **60%**.
- Go/No-Go deployment gate with cascade risk detection prevented **~30%** of high-risk deploys from reaching production.

[Model Arbitration Engine](#) | *AWS Lambda, SageMaker, Bedrock, DynamoDB*

Feb'26 – Apr'26

- Built serverless ML routing engine on AWS (API Gateway → Lambda → SageMaker + Bedrock) dynamically selecting among sklearn, XGBoost, and Claude Haiku; achieved **95.3% accuracy** at avg **87ms latency** vs 780ms Bedrock-only baseline.
- Implemented epsilon-greedy routing ($\epsilon=0.1$) with EMA latency tracking and DynamoDB decision ledger, reducing inference cost by **84%** vs routing all requests to Bedrock.
- Benchmarked 3 model tiers: sklearn **43ms/\$0.01/1K**, XGBoost **118ms/\$0.05/1K**, Bedrock **780ms/\$0.25/1K**; engine routes based on real-time cost/latency/accuracy tradeoffs.

ACHIEVEMENTS

- Oracle Cloud Infrastructure 2025 Certified AI Foundations Associate.
- Udemy: Complete Generative AI Course with LangChain & HuggingFace.

EDUCATION

Lovely Professional University

B.Tech – Computer Science & Engineering

Phagwara, Punjab

Aug'23 – Present

Salt Brook Academy

Intermediate – 76%

Dibrugarh, Assam

Aug'21 – May'23