

Saheli Das

📍 Kolkata, India (UTC+5:30) | ✉️ sahelidas0506@gmail.com | 📞 +91 99039 05541

🌐 saheli.tech | github.com/saheli56 | [in linkedin.com/in/saheli-das](https://in.linkedin.com/in/saheli-das) | leetcode.com/u/saheli56

Summary

AI/ML Engineer specializing in autonomous agent architectures, time-series machine learning systems, and deterministic workflow orchestration. Experienced in developing semantic caching layers, AST analysis tooling, low-latency API engines, and production ML pipelines with measurable business impact.

Education

JIS College of Engineering

Bachelor of Technology in Information Technology

2024 – 2028

CGPA: 8.55/10 (till 4th semester)

Technical Skills

- **Languages:** Python, C++, TypeScript, JavaScript, SQL, Shell
- **Machine Learning & Data:** PyTorch, TensorFlow, Scikit-learn, LightGBM, NumPy, Pandas, OpenCV, OCR (Tesseract)
- **Agent & System Design:** Autonomous Agents, Semantic Caching, Function Calling, Vector Embeddings, WebSockets, Fastify
- **Backend & Frameworks:** FastAPI, Flask, React, Next.js, Node.js, REST APIs, Git, GitHub Actions

Projects

Do It For Me (DIFM) | Autonomous Web Agent & Workflow Automation Platform

[GitHub](#)

- Architected a multi-model autonomous web action agent combining an accessibility tree parser (**@difm/a11y-tree**) with a Fastify/WebSocket server and Chrome MV3 extension.
- Engineered an intent compiler and prompt compressor reducing token overhead by **35%**, featuring dynamic tree compaction (40–60 nodes) and zero-token deterministic heuristic fast-paths.
- Implemented a 4-tier risk policy guardrail (Tier 0–3) with rate-limit semaphores, OCR document bill extraction, and automated end-to-end task scheduling across 25+ automated test suites.

Autonomous AI Revenue Recovery Engine | Self-Healing Payment & Churn Remediation System

[GitHub](#)

- Architected an autonomous 5-stage closed-loop pipeline (Detector, Diagnoser, Strategist, Executor, Auditor) resolving payment drops, checkout friction, and subscription mandate declines.
- Built a hybrid diagnoser combining deterministic error mappings (**<1ms latency**) with LLM fallback, utilizing rate-limit semaphores and in-memory semantic caching to eliminate 429 errors.
- Enforced 12 priority policy rules and safety guardrails, achieving **41.2% gross recovery** on a ~₹7.62L portfolio with **574.2x net ROI** and **100% safety precision**.

Store Demand Forecasting System | Temporal ML Pipeline for Retail Inventory

[GitHub](#)

- Engineered a time-series forecasting pipeline using lag, rolling-statistical, and calendar features to predict retail store demand.
- Benchmarked LightGBM, Random Forest, Tweedie regression, and two-stage forecasting models against statistical baselines using time-aware cross-validation.
- Reduced RMSE by **15%** versus weekly-median baseline on dense demand distributions while ensuring strict zero-leakage temporal feature engineering.

Achievements

Competitive Hackathons

6× Podium Finisher

- 1× Winner (1st Place) – IEM HackOasis 2.0
- 1× Cybersecurity Track Winner – InnoFusion 3.0
- 4× Runner-Up – Hexafalls, HackInverse 0.5, MetaMorph 2.0, CodeVoyage 2026