JANNET AKANKSHA EKKA

janetekka96@gmail.com | +91 90788 02572 | LinkedIn | GitHub | Portfolio | Kolkata, India · remote or relocation

PROFESSIONAL SUMMARY

I build AI and ML systems, test them against the obvious baseline, and keep them working where other people use them. For about four years at Deloitte I was the engineer on the calls with Fortune 500 client teams (AT&T, Hewlett Packard Enterprise), in US, UK and India hours, and led a 6-person QA team. Rank 1 in the UT Austin McCombs Post Graduate Program in AI/ML (GPA 4.09). Since January 2025 I have built and run a live AI trading system alone; when I graded its six AI agents on independent data, all six called direction at a coin flip, so I rebuilt it around written rules and a statistical gate every model change must pass. Between Oct 12 and Nov 15, 2026 I built and submitted 19 working products and a paper to 21 hackathons, each on its sponsor's platform (GitLab Duo, Meta, NVIDIA Nemotron, Fire TV, FalkorDB, Mozaik and others). Python and SQL every day.

TECHNICAL SKILLS

Client Delivery	Demos, working sessions and requirement calls with client teams across US, UK and India hours; scoping a build to a fixed deadline; defect triage with client engineers; weekend on-call; READMEs, technical reports and demo videos for the people who use the system
AI & Agents	Gemini (Vertex AI, AI Studio), Claude, NVIDIA Nemotron, Meta Muse, Google ADK, GitLab Duo flows, Mozaik concurrent agents, MCP (servers written from scratch), tool calling, structured output, RAG with measured retrieval, FalkorDB graph agents (Cypher), llama.cpp
ML & Evaluation	scikit-learn, XGBoost, CatBoost, Random Forest and boosting, SMOTE, Optuna, TensorFlow, PyTorch, Hugging Face; held-out and forward tests, paired bootstrap confidence intervals, ROC-AUC, recall@k and MRR, event studies against base rates, CPCV, Deflated Sharpe, PBO, FDR control, leakage checks
Data	SQL (window functions, joins), BigQuery, DuckDB, Postgres, Firestore, pandas; pipelines over public APIs that rebuild byte for byte from a snapshot; GitHub Actions schedules. Not used yet: dbt, Trino, Databricks
Deployment	Oracle Cloud (free VM, systemd), Cloudflare Workers, Pages and KV, Render, GitHub Actions; GCP: Cloud Run, Compute Engine, Vertex AI, BigQuery, AlloyDB & pgvector; AWS: S3, EC2, RDS; Docker, including offline networks
Engineering	Python, SQL, TypeScript and JavaScript; Kotlin for one shipped Android app. FastAPI, Flask, Streamlit, React, Next.js, Solidity. Test strategy, pytest (3,100+ tests on one system), Selenium, Jenkins

PRODUCTION SYSTEM

Smart Money Trading (SMT): an explainable AI trading agent · Founder & sole engineer | Jan 2025 – Present

Patent pending: Indian provisional application No. 202631090789, filed myself in July 2026. Shortlisted for the Google AI Startup Buildathon India 2026

Live: SMT World · ***Public layer:*** github.com/JannetEkka/smt-apac · WEEX feature

- Runs 24/7 and places real orders on WEEX's exchange-granted simulated account (about $10K). Built alone by specifying the work and directing AI coding agents: **87,000 lines of Python, 3,100+ tests, 368 merged pull requests**, each reviewed by me before merge.
- Graded my own six AI agents on independent market data: **all six called direction at a coin flip**. So written rules now open every trade, the agents are advisors graded on every call, and every explanation passes a counterfactual test.
- A model change reaches the live system only through a CPCV + Deflated Sharpe + PBO + FDR gate. A forecaster that scored **AUC 0.723** on a BTC forward test over 27 non-overlapping windows is still held in shadow because it has not passed.
- When the GCP billing account closed, I **moved the live system to a zero-cost stack in two days**: an Oracle free-tier VM under systemd, Cloudflare Workers and GitHub Actions. It re-adopted its three open positions from committed state.
- Audited 98 AI-written merges line by line and found an order-format bug that had **rejected 1,137 live orders over three months** while every test passed. A fix now counts as done only when the production log shows it.
- Built **SMT World**, the product for users: a Three.js world on Cloudflare Workers where people watch the agent trade and ask it why, answered by Gemini from the live book.

SELECTED BUILDS

1. Proof of Fix: issues close on production evidence, not on merge · GitLab AI Hackathon (Duo Agent Platform) | Oct 2026

Live: the real run on GitLab · ***Code:*** GitLab project · Demo video

- A GitLab Duo flow reads production logs when a bug is filed and writes a proof contract into the issue; after deploy, the pipeline judges the shipped build on its own logs. Real run: a fix that passed the tests and Duo Code Review cut failures only **from 32% to 9%**, so the issue stayed open; the Duo Developer flow's second fix reached **0 failures on 875 requests**.

2. Chithi: official Indian letters explained, every number checked · Meta Global AI Developer Hackathon (Cerebral Valley) | Nov 2026

Live: chithi-oprj.onrender.com · ***Code:*** github.com/JannetEkka/chithi · Demo video

- A phone photo of a bill or notice, explained in Bengali, Hindi or English on SAM 3.1 and Muse Spark. A check with no model requires every date and amount to be in the quote and in a second reading; 28 live explanations, **0 of 114 values shown wrong**.

3. Lore: backtesting an agent against the obvious baseline · Graph Hacks, FalkorDB × WeMakeDevs | Oct 2026 · 4 days

Live: lore-gh5y.onrender.com · ***Code:*** github.com/JannetEkka/lore

- 1,357 of the sponsor's own PRs in a FalkorDB graph; a Gemini agent answers "who should review this?" through 12 Cypher tools and cites the PRs. Backtested on 300 real PRs: a real reviewer in its top three **90%** of the time, against 88% for "the most frequent reviewer". On the single top pick the baseline wins, and the write-up says so.

4. SmartDesk: a measured comparison that kept the baseline · Google Cloud Gen AI Academy APAC, top 101 of 1,500+ | 2026

Code: github.com/JannetEkka/smartdesk

- Built the eval set before touching retrieval: 40 labelled questions, paired bootstrap at 10,000 resamples. A reranker won on the development embedder (MRR@10 +0.094, p=0.04), then fell to +0.028 with a CI spanning zero in production, while search went from 29 ms to 2,382 ms. **The baseline shipped.**

JANNET AKANKSHA EKKA

jannetekka96@gmail.com | +91 90788 02572 | LinkedIn | GitHub | Portfolio | Kolkata, India · remote or relocation

5. Aftershock: does news move real usage? An event study · NEAR Hackathon, Investment Thesis track | Nov 2026

Live: jannetekka.github.io/aftershock · ***Code:*** github.com/JannetEkka/aftershock

❖ Joined 2,305 days of NearBlocks chain data, DefiLlama and GDELT news counts into one snapshot that rebuilds byte for byte. Of 74 news spikes, 8 moved the price (10.8%, against 5.1% of ordinary days) and 10 moved usage (13.5%, against 16.5%): **usage did not follow the headlines**, and the thesis says so.

6. Spona: concurrent agents that match researchers with companies · NTP Novi Sad Hackathon (JigJoy × Mozaik) | Oct 2026

Live: jannetekka.github.io/spona · ***Code:*** github.com/JannetEkka/spona

❖ 12 agents read 1,047 researchers, 10,779 papers and 489 company challenges in English and Serbian (Latin and Cyrillic) and gave every challenge an outcome in **5 min 48 s for $1.52**. A critic checked every claim against its source while the others worked and sent back 33.

7. Lights Out: a Hermes Agent plugin that asks before spending · Nous Research Hermes Agent Plugin Hackathon | Nov 2026

Code: github.com/JannetEkka/lights-out · Merged into the official Hermes plugin catalog (PR #153822)

❖ Measured on 10 tasks × 3 arms: without the plugin the agent ran 13 and 14 billable AWS creates that no person approved; with it, **0**. One command then left 0 resources running, keeping the 2 marked to stay.

Portfolio models (UT Austin PGP, 2024–25): telecom churn, Random Forest 81.48% test accuracy (for a retention budget I would report recall on churners and lift); NBFC loan default; semiconductor yield on 591 sensor signals, where I later found SMOTE applied before the split, so the reported 99.2% is inflated; SVHN digits, 87.1% on 18,000 images; financial-news sentiment.

Also shipped, Oct–Nov 2026: Footfall (Qloo + Nemotron, a $15 spend cap, $0.45 spent) · Aside (Fire TV, spoilers 6 of 16 → 0 of 16) · Klarbrief (Swiss letters offline on a 4-core CPU, 0 wrong CHF amounts on 48) · Orbit (Catalyst, one-day solo build, top 8.2% of 391) · Countersign (41 of 60 planted errors caught) · Pitch In (PayPal) · a Kaggle paper on Gemma 4 coding agents · Wiretrap · SMT Receipts · SMT World: Clock In (Android) · SQL Noir · Afterburn · Recourse (Round 2) · Beauty Within · Ghost Ledger. All at jannetekka.cv

PROFESSIONAL EXPERIENCE

Deloitte Consulting LLP · Test Automation Engineer (Analyst) | Aug 2019 – Jan 2023 · Bangalore

❖ **Client-facing throughout**: demos, working sessions and requirement calls in US, UK and India hours for Fortune 500 accounts including AT&T Consumer Connect, Hewlett Packard Enterprise and The Dental Supply Company; weekend on-call. **Led a 6-member QA engineering team**.

❖ Built Selenium and Katalon Studio automation that **improved test execution efficiency by 83%;** Jenkins-based reporting cut manual analysis time by **75%**.

❖ Wrote Python analysis and Tableau dashboards tracking quality across **343 business components**, cutting test failures by 25%; analysed **50,000+ SAP Hybris transactions** to isolate performance bottlenecks for the client.

❖ Joined as a QA automation intern in Jan 2019; converted to analyst after six months. 2 Spot Awards and an Applause Award.

AutoKorrekt · Lead Frontend Developer (EdTech startup MVP) | May 2023 – Jan 2024

❖ Led the Next.js frontend of an AI answer-evaluation platform for teachers: OCR validation flows, JWT auth, bulk answer-sheet upload, over a backend using **AWS Textract**, Amazon Comprehend, S3, RDS and SageMaker.

JTech Solutions (pre-incorporation) · Founder & AI Engineer | Jan 2025 – Present · Kolkata

❖ Building SMT (above) as the first product, while a full-time family carer and alongside the PGP in AI/ML.

EDUCATION

Post Graduate Program (PGP) in AI/ML | UT Austin McCombs / Great Lakes Executive Learning (Feb 2024 – Mar 2025)

RANK 1 / COHORT TOPPER (Specialization in Deep Learning, NLP, and Data Modeling): GPA 4.09

Bachelor of Technology (B.Tech), Information Technology | Graduated 2019

Kalinga Institute of Industrial Technology University, Bhubaneswar: CGPA 7.17/10

CERTIFICATIONS & ACHIEVEMENTS

- **WEEX AI Wars I: #1 on the preliminary leaderboard** of 230+ teams, then one of 37 finalists in the live AI trading competition; featured in WEEX's coverage.
- **Winner, Best DeFi Application**, OpenServ × Hack2skill hackathon; **top 101 of 1,500+ teams**, Google Cloud Gen AI Academy APAC 2026 Cohort 1 (SmartDesk).
- **Kaggle RSNA Knee Abnormality Detection** (5,900+ teams, not placed): scored two candidate label sets against 58 radiologist labels (macro AUC 0.894 against 0.870) before retraining on the better one.
- **Google Cloud Gen AI Academy** 2025 and APAC 2026 Cohorts 1–3; CPSAT (Selenium), 2022. All on Credly and Google Skills.