

AAKASH SANGANI

Data Scientist · ML and LLM Systems

Mumbai, India · aakash.sangani2410@gmail.com · linkedin.com/in/aakash-sangani-3149a11a7 · github.com/aakash2410 · aakash2410.github.io

SUMMARY

Data scientist, 3+ years across pharma analytics, fintech experimentation and AI governance. Builds and evaluates ML and LLM systems: edge inference, retrieval, Bayesian inference, fairness audits. Author of two open-source evaluation tools and an open agentic-failure benchmark.

TECHNICAL SKILLS

ML and LLM: RAG, agentic AI, LLM evaluation and red-teaming, MCP servers, Bayesian inverse RL, LSTM, XGBoost, SHAP, AIF360, scikit-learn, TensorFlow

Edge and inference: INT8 quantisation, ONNX, parallel batched inference, sentence-transformers, vector search, ASR / MT / TTS pipelines

Data and infrastructure: Python, SQL, pandas, NumPy, PostgreSQL, Apache Spark, A/B testing, AWS, Docker, Ansible, CI/CD gates, SARIF

Governance: EU AI Act, NIST AI RMF, ISO 42001, GDPR/DPDP, HIPAA, OWASP ASI, MITRE ATLAS

EXPERIENCE

AI Consultant

Feb 2026 – present

Kalpa Impact

Mumbai

- Designed the inference pipeline and owned the benchmarking for **Suno Sutra**, the open-source offline AI device from Digital India Bhashini (MeitY) and Current AI: batching and scheduling tradeoffs across three models under a hard RAM ceiling cut latency to first audio from ~8s to 2–3s and peak memory by 40%, alongside the team's INT8/ONNX conversion.
- Built and packaged **The Specialist**, a grounded offline retrieval module for that device: multilingual-e5-small embeddings with a deterministic TF-IDF fallback, two refusal gates, and a source citation on every answer.
- Designed the evaluation framework for the **VYOMA Innovation Challenge** (₹80 lakh, Bhashini × Current AI × Kalpa Impact): jury rubric, judging matrix and scoring thresholds. Verified edge-AI claims submitted to the Gates Foundation; wrote a section of *AGI Governance* with the Principal Scientific Adviser's Office on AI incident reporting.

Business Technology Solutions Associate → Senior Associate

Sep 2023 – Jan 2026

ZS Associates

Pune

- Launched an agentic AI analytics product for natural-language querying of sales-performance metrics.
- Owned incentive-compensation product logic for a Fortune 100 pharma client at 100% accuracy. Led a team of 4; top 1% performer firm-wide, promoted to Senior Associate within 8–10 months.

Product Analyst Intern

Jul 2022 – Feb 2023

Bright Money

Bangalore

- Built A/B testing infrastructure and SQL dashboards; launched a referral programme that saved \$120K in customer-acquisition cost over 5 months.

OPEN-SOURCE PROJECTS

Inverse-Objective Regulator

github.com/aakash2410/Inverse-Objective-Regulator

- Black-box auditor that infers a deployed agent's revealed objective from its tool-call traces via Bayesian inverse RL and scores divergence from the declared purpose, emitting JSON and SARIF mapped to the EU AI Act, OWASP, MITRE ATLAS and NIST AI RMF.
- Open 8-agent benchmark of agentic failure modes. A deterministic cross-encoder critic ranked the planted objective top-2 on every seed for 6 of 8 agents (mean $r = 0.55$), deterministically and with no API cost, against a frontier LLM-as-judge spanning 0.08–0.43 across agents. Sole author of the paper draft, not yet on arXiv.

rag-compliance-auditor

github.com/aakash2410/MCP_RAG_Compliance

- Sole-authored package on PyPI. MCP server auditing RAG pipelines against the EU AI Act, GDPR/DPDP, HIPAA, SEBI/RBI and NIST AI RMF across 7 dimensions; 8 MCP tools, RS256-signed audit reports, runs as a CI gate with SARIF.

EDUCATION AND PUBLICATION

B.Tech, Information Technology, D. J. Sanghvi College of Engineering, University of Mumbai, 2019–2023. CGPA 8.90 / 10, full merit scholarship. Book chapter: Sangani, A., Doshi, P., & Sawant, V. "Unusual Social Media Behavior Detection Using Distributed Data Stream Mining," Auerbach (CRC Press), 2022–23.