

Manish Mishra

SUMMARY

AI Engineer leading end-to-end GenAI system development across business, data, and application layers, delivering enterprise AI solutions for Big Pharma client on Azure and Databricks. Specializing in multi-agent systems, RAG, and data platform design, with strong focus on data governance, explainability and traceability.

EXPERIENCE

Aganitha Cognitive Solutions | Full Stack Data Scientist II

Feb. 2023 – Present

- Led a team of 4 engineers to deliver 5+ GenAI prototypes with a 75% production conversion rate. Driving pilots for enterprise GenAI applications for R&D and Tech Ops Teams, including:
 - * AI-powered scientific study report generation system leveraging protocols, ELNs, and templates to produce compliant reports, reducing authoring time by 80%.
 - * GxP-compliant deviation lifecycle assistant for deviation drafting, investigation, and audit-ready reporting, reducing lifecycle time by 40%
- Developed multiple production ready CoE GenAI solutions including:
 - * AskLLM: plug-and-play framework enabling LLM integration with React-based frontend and configurable data contexts.
 - * LegalLLM: RAG-based contract intelligence platform for contract comparison, compliance evaluation, and automated contract generation.
- Modernized PheWeb (PheWAS browser) by migrating Flask → FastAPI with DuckDB integration, achieving 60% reduction in query latency. Designed parquet-based data architecture for large-scale GWAS datasets (1,000+ phenotypes), enabling efficient querying and incremental updates.

Tata Consultancy Services | System Engineer

Jun 2021 – Jul 2022

- Designed and implemented backend solutions for the Global Banking Product (GBP), focusing on Loan Product Management:
 - * Implemented interest rate control system with floor and cap limits for loan products.
 - * Developed an automated invoice generation system to streamline loan processing workflows.

TECHNICAL SKILLS

- **AI/ML Development:** PyTorch, Scikit-learn, Transformers, MLFlow
- **Gen AI:** Langchain, MCP, Google ADK, Autogen, Langgraph, BAML
- **Backend:** Python, FastAPI, Pydantic, Typer, Pandas, PostgreSQL, DuckDB
- **Cloud & DevOps:** Linux, Docker, Azure (AI Foundry, OpenAI, AI Search, Web Apps, Container Apps), AWS(EKS, Batch, Bedrock), Databricks, Terraform

EDUCATION

Indira Gandhi Institute of Technology, Sarang

Jun 2017 – Jun. 2021

- * BTech, Electrical Engineering
- * GPA: 8.30
- * Publication: A Novel Approach for Identifying Social Media Posts Indicative of Depression.