

PRIT CHHODAVADIYA

prитpatel.connect@gmail.com | [Github](#) | [linkedin](#) | [Portfolio](#) | +91 8905007676 | Ahmedabad, India

PROFESSIONAL SUMMARY

Final-year B.Sc. Data Science student with demonstrated expertise in architecting and deploying production-grade AI systems across Generative AI, MLOps, and Agentic AI. Proficient in the full ML engineering lifecycle — from data pipeline design and model development to REST API deployment, containerization, and CI/CD automation. Seeking a junior AI/ML or MLOps engineering role to contribute measurable impact from day one.

TECHNICAL SKILLS

Languages Python · SQL

ML / DL Frameworks PyTorch · TensorFlow · Scikit-learn · XGBoost · HuggingFace Transformers

GenAI / LLM LangChain · LangGraph · RAG (Hybrid: Dense + BM25 + RRF) · Pinecone · ChromaDB · Sentence Transformers

MLOps & Deployment Docker · GitHub Actions CI/CD · MLflow · Weights & Biases · FastAPI · FastMCP · Redis · MongoDB · PostgreSQL

Data Engineering Pandas · NumPy · MySQL

NLP fine-tuning · Token Classification · Sentiment Analysis · BIO Tagging

PROJECTS

Neural Hybrid RAG System [↗ Live on Hugging Face Spaces](#)

- Architected to **solve the core failure of pure vector-only RAG — poor exact-match recall on private documents**. Enables organisations to **query internal knowledge bases with production-grade accuracy** that naive implementations cannot achieve.
- Engineered a **Hybrid Retrieval pipeline** fusing e5-base-v2 dense search with a custom BM25 variant via Reciprocal Rank Fusion; parallelized embedding batches using Synchronous Programming and PDF ingestion using ProcessPoolsr; deployed via **FastAPI, SSE streaming, Docker, Pinecone gRPC, Redis, MongoDB**, and LangSmith observability — reducing embedding latency by **89% (88s → 8s)** and PDF extraction by **90% (52s → under 5s)**, achieving **P50 ~5.25s** in production.
- Delivered **complete end-to-end ownership** from architecture through production deployment — live on Hugging Face Spaces, demonstrating **enterprise-grade RAG system delivery** shipped independently within a single sprint cycle.

Customer ChurnGuard — End-to-End MLOps Pipeline [↗ Live on AWS EC2](#)

- Designed an **internal automation tool** to eliminate a labour-intensive manual churn reporting workflow — replacing it with a **fully automated SaaS micro-product** delivering real-time scoring and explainable predictions via a single REST API.
- Trained **XGBoost and Random Forest classifiers** with Optuna tuning; integrated **SHAP explainability** per prediction; governed model lifecycle via **MLflow and W&B**; enforced CI/CD gates with multi-stage Docker — delivering an **88% reduction in analyst reporting time** with full reproducibility and audit traceability.
- Shipped a **fully autonomous production platform** with zero manual steps from code commit to deployed API — demonstrating **complete end-to-end MLOps ownership** across pipeline architecture, model versioning, and automated client-facing reporting.

FinAgent — Multi-Agent Financial Analysis Platform [↗ Live on Hugging Face Spaces](#)

- Built to automate end-to-end Stock analysis using an **agentic AI pipeline** — eliminating manual research workflows by orchestrating specialised agents for **market data, news search, risk detection, and investment memo generation**.
- Architected a **multi-agent LangGraph workflow** with conditional agent routing (market → search → analyst → risk mitigation → reporter); exposed agent tools via **FastMCP server**; streamed live pipeline progress via SSE; persisted agent state using **SQLite-backed LangGraph checkpoints**; deployed as single-container and split-container Docker modes with GitHub Actions CI/CD.
- Delivered a **production-ready agentic AI platform** capable of full stock analysis with live streaming — handling mid-analysis **risk detection with automatic routing** before generating the final investment memo. Live on Hugging Face Spaces.

EDUCATION

Bachelor of Science in Data Science Expected: May 2026

Silver Oak University, Ahmedabad, India

Relevant coursework: Machine Learning · Data Preprocessing · Data Visualization · Statistics · Database Management Systems