

PRAKYATH CHANDRAN

chandranprakyath@gmail.com | +91 7259528079 | LinkedIn | Github | Blog | Portfolio | Bangalore

WORK EXPERIENCE

AI Developer - Intern | Phracta | London

July 2025 – Dec 2025

- Sole AI developer for 3 startup clients; led requirements gathering in client meetings and delivered end-to-end MVPs within 3-4-week cycles
- Built mental health chatbot MVP with compatibility matching using LangChain, OpenAI embeddings, and ChromaDB, implemented Few-Shot prompting for empathetic response generation; delivered for client demo.
- Developed financial analysis tool processing 50+ page institutional data into actionable metrics using RAG pipeline (PDF ingestion, FAISS vector search, GPT-4 synthesis) with FastAPI backend, and achieved <10 second query response
- Tech: Python, FastAPI, LangChain, OpenAI API, Vector DBs (ChromaDB, FAISS), Streamlit, AWS, scikit-learn

Data Scientist Intern | HaiX AI | London

Feb 2025 – Jun 2025

- Benchmarked 5+ LLMs (GPT, Claude, Mistral, LLaMA) by deploying models on server and systematically testing prompt strategies (Few-Shot, Chain-of-Thought) across 100+ prompt variations. Identified 10% cost reduction opportunity using Mistral/Claude over GPT-4 for sentiment tasks.
- Built FastAPI endpoints for both open-source (Hugging Face) and proprietary models, integrated RapidAPI to collect brand engagement data from 7 social platforms
- Evaluated multimodal models (GPT-4V, Vertex AI) for extracting brand mentions from video/reels and compared outputs against in-house model using manual scoring framework
- Tech: Python, FastAPI, Azure OpenAI, Hugging Face, Prompt Engineering, SQL, REST APIs, Postman

KEY PROJECTS

Hope Bridge – AI Mental Health Support Platform | *Python, FastAPI, LangChain, ChromaDB, Docker* [Link](#)

- Built multi-model LLM orchestration with planning logic, querying GPT-4, Claude, Gemini in parallel with quality-based response selection, automatic fallback, and per-request cost tracking.
- Implemented A/B testing framework with statistical significance analysis (Welch's t-test) and prompt version control with rollback
- Designed real-time risk assessment engine (0-100 scoring) with crisis detection and automatic resource injection from 50+ country crisis helpline database
- Created observability stack with Prometheus metrics, and Grafana dashboards and deployed via Docker Compose.

LLM Trading Bot & Financial Document Analyzer | *Python, LangChain, AstraDB, AWS EC2, GCP* [Link](#)

- Developed RAG-based system for querying SEC filings and financial reports (100+ pages) with vector search using DataStax AstraDB
- Deployed on AWS EC2 with session management and real-time chat interface; achieved 2-3 second response times for complex queries.

EDUCATION

MSc Applied AI and Data Science (International Merit Scholarship) | Solent University, Southampton

Bachelor of Computer Applications (First Class) | St Joseph's College, Mysore

SKILLS

AI/ML: LangChain, RAG Systems (Hybrid Search, HyDE, Corrective RAG, Reranking), Prompt Engineering (CoT, ReAct, Few-Shot), Vector Databases (ChromaDB, FAISS, AstraDB), Hugging Face, PyTorch, Speech-to-text

LLMs and Agents: OpenAI API, Claude API, Gemini API, Mistral, Multi-model Orchestration, Agentic AI (Tool Calling, LangGraph, Multi-Agent Systems), A/B Testing, Model Benchmarking, Model Fine-tuning

Backend & Infrastructure: Python, FastAPI, Streamlit, Docker, AWS EC2, MongoDB, Prometheus, Git, REST APIs

Data Analytics: pandas, NumPy, Statistics, SQL, Data Cleaning, Feature Engineering, Tensorflow, Keras