N A R E S H K U M A R

P Y T H O N B A C K E N D D E V E L O P E R | G E N A I D E V E L O P E R

nareshkumar04094@gmail.com | Portfolio | LinkedIn | GitHub

PROFESSIONAL SUMMARY

Python Backend & GenAI Developer with **2+ years of professional experience** building AI-powered applications, REST APIs, data pipelines, and LLM-based workflows. Hands-on experience with **Python, FastAPI, LangChain, LangGraph, RAG,** vectorless retrieval, OpenAI API, Google Gemini API, ChromaDB, MySQL, and ClickHouse. Experienced in developing schema-drift-resilient data pipelines, semantic data layers, AI product recommendation systems, deterministic decision engines, and document retrieval workflows. Focused on backend engineering, Generative AI integration, and enterprise data platforms.

TECHNICAL SKILLS

- **Programming** : Python, SQL
- **Backend & APIs** : FastAPI, REST APIs, SQLAlchemy, APScheduler, Streamlit
- **GenAI & LLM** : LangChain, LangGraph, RAG, AI Agents, OpenAI API, Gemini API
- **Databases & Data Platforms** : MySQL, ClickHouse, ChromaDB, Salesforce Integration
- **Data Engineering** : ETL Pipelines, Incremental Data Synchronization, Data Ingestion, Schema Evolution, Batch Processing
- **Data & Machine Learning** : Pandas, NumPy, XGBoost
- **DevOps & Version Control** : Docker, Git, GitHub

PROFESSIONAL EXPERIENCE

Junior Developer

Biz Impact Cloud Solutions (BICS) | Aug 2024 – Present

AI Reporting Agent

- Engineered an **incremental Salesforce-to-ClickHouse ETL pipeline** covering **11 Salesforce objects and approximately 10K records per incremental sync**, with full-sync capability for approximately **500K records.**
- Integrated scheduled data synchronization using **Python, FastAPI, and APScheduler** to automate recurring data refreshes.
- Engineered an automated **schema-drift detection and self-healing workflow** that detects Salesforce metadata changes, generates updated ClickHouse schemas, hydrates isolated shadow staging tables, and performs atomic table swaps using **EXCHANGE TABLES**, with rollback handling and developer notifications.
- Designed and implemented a **ClickHouse semantic data layer** containing Salesforce field descriptions, picklist options, object relationships, and business-purpose metadata to support AI-driven analytical querying.
- Customized and integrated the **PageIndex open-source implementation** as a vectorless retrieval component, avoiding recurring dependency on the PageIndex API.
- Extended document processing beyond the original TXT/MD support to handle **PDF and Word documents** containing enterprise schemas, entity relationships, and operational instructions.

B2B Product Recommendation & Decision Engine

- Developed the backend workflow for a **client-facing B2B AI product recommendation system covering 442 configurable industrial products**, using Python, FastAPI, LangGraph, OpenAI API, and ChromaDB.
- Implemented **ChromaDB-based vector retrieval** to identify relevant products from natural-language customer requirements and product attributes.
- Designed and implemented a **deterministic decision engine** that evaluates product feasibility, applies business constraints, filters unsuitable candidates, and ranks recommendations.
- Integrated product data from **MySQL and Salesforce**, adapting the existing data and retrieval layer to support product intelligence and recommendation workflows.
- Modified and integrated the existing **Streamlit frontend** to support recommendation workflows and backend REST APIs.
- Contributed to the next-phase architecture for evolving the system into an **AI product consultant**, including improved product attribute modeling, scoring dimensions, and decision-making capabilities.

Scheduled Synchronization & Operational KPI Engine

- Built an **incremental Salesforce-to-MySQL synchronization pipeline from scratch**, integrating **9 source objects: 8 Salesforce objects and 1 external database source.**
- Implemented scheduled synchronization using Python, FastAPI, APScheduler, and SQLAlchemy ORM for structured data access and retrieval.
- Developed **FastAPI REST endpoints** exposing business metrics, including orders, accounts, and other operational KPIs, for consumption by a Next.js dashboard.
- Identified and mapped required fields and relationships across multiple source objects to support consistent operational reporting.

AI Product Visualization System

- Provided technical guidance and development support for a **generative AI product visualization system** that uses reference products to visualize customer-specific customizations.
- Supported the development team through **architecture discussions, implementation guidance, troubleshooting, testing, feature validation, and code reviews.**
- Supported the AI workflow involving **LangChain and Google Gemini API**, helping validate functionality against product visualization requirements.

ADDITIONAL PROJECT EXPERIENCE

Medal Price Prediction Application

- Engineered a **FastAPI backend and interactive Streamlit interface** for medal price prediction using **XGBoost, Pandas, and NumPy.**
- Exposed prediction functionality through REST endpoints and developed an interactive interface for submitting inputs and viewing prediction results.

SFTP Document Synchronization Pipeline

- Engineered a **Python-based SFTP pipeline** to synchronize and preserve original source documents before processing by the existing OpenAI embedding workflow.
- Automated source-document transfer and storage to retain original files for downstream document retrieval and RAG applications.

EDUCATION

B.E. (Computer Science and Engineering) [2019 - 2023]
Panimalar Institute of Technology, Affiliated to Anna University,
CGPA: 8.51/10 - (85.1%)