

Dheeraj Rahul Reddy Piduru

Dallas, TX | +1 (945) 398-0493 | dheerajpiduru.9@gmail.com | linkedin.com/in/dheerajpiduru

PROFESSIONAL SUMMARY

AI Engineer with 3+ years of experience across data engineering, Generative AI, and agentic AI solutions. Hands-on experience building RAG applications, AI agent workflows, and conversational analytics solutions using LangGraph, OpenAI API, open-source vector search tools, Python, FastAPI, and Microsoft Fabric. Skilled in tool calling, retrieval optimization, data integration, and AI application evaluation, with experience connecting AI solutions to business workflows using lightweight, no-cost automation tools.

TECHNICAL SKILLS

- **Programming:** Python, SQL, PySpark, REST APIs
- **Azure AI & GenAI:** Azure AI Foundry, Azure OpenAI, Azure AI Search, RAG, Embeddings, Hybrid/Semantic Search, Prompt Engineering
- **Agentic AI:** LangGraph, Semantic Kernel, Tool Calling, Agent Orchestration, Agent Memory, Pydantic
- **AI Evaluation & Monitoring:** Foundry Evaluation, Azure Monitor, Application Insights, Azure AI Content Safety
- **Microsoft Fabric & Data:** Microsoft Fabric, OneLake, Lakehouse, Data Factory, Data Migration, Data Management, Medallion Architecture, ETL/ELT
- **Development & Integration:** FastAPI, Docker, Git, Power Automate

CERTIFICATIONS

- Google Cloud Certified: Professional Machine Learning Engineer
- Microsoft Certified: AI-103 Azure AI Apps and Agents Developer Associate

PROFESSIONAL EXPERIENCE

AI Engineer · Community Dreams Foundation Feb 2026 - Present

- Built a RAG-based AI application using LangGraph, the OpenAI API, and an open-source vector store (FAISS), indexing 50+ internal program documents and achieving ~85% grounded-response accuracy across 30 evaluation questions.
- Developed LangGraph-based agentic workflows with 3 tool integrations, enabling agents to retrieve context, execute functions, and complete multi-step requests while reducing manual handling by ~20%.
- Integrated the AI assistant with a lightweight SQLite database of program records, enabling conversational access to curated organizational data and reducing typical lookup time from ~6 minutes to under 2 minutes.
- Implemented Pydantic validation, free-tier workflow automation (Zapier), conversation memory, and application logging, reducing malformed tool/API responses by ~20% and automating repetitive downstream hand-offs.

Data Engineer · Accenture May 2021 - Oct 2023

- Built ingestion, transformation, and migration pipelines using Azure Data Factory, Azure Data Lake Storage, PySpark, and SQL, processing 1M+ records per pipeline run across structured and semi-structured sources.
- Implemented Bronze, Silver, and Gold data layers using Medallion architecture, reducing failed data-quality checks by ~25% and improving availability of curated datasets for downstream analytics.
- Developed Python-based data services and FastAPI REST APIs, reducing recurring manual data-extraction requests by ~35% and enabling self-service access for analytics applications.
- Prototyped a RAG-based document Q&A solution using Azure OpenAI and Azure Cognitive Search, indexing 100+ documents and achieving ~80% relevant retrieval across 50 test questions.

EDUCATION

Master of Science - Business Analytics & Artificial Intelligence Dec 2025

The University of Texas at Dallas · TX, USA

Bachelor of Technology - Computer Science & Engineering, Mahindra University Jun 2023

Mahindra University · Hyderabad, India

PROJECTS

Azure Agentic RAG Assistant

- Built an agentic RAG assistant using Azure AI Foundry, Azure OpenAI, and Azure AI Search, indexing 50+ documents with chunking, embeddings, and hybrid retrieval to generate grounded, cited answers.
- Implemented LangGraph-based agent orchestration and tool calling, enabling dynamic routing between retrieval, structured lookup, and response-generation tools for multi-step queries.
- Evaluated retrieval and response quality on a 75-question test set, applying Pydantic validation, conversation memory, and prompt guardrails to reach ~87% grounded-answer accuracy.

Microsoft Fabric Conversational Analytics Agent

- Built a Microsoft Fabric Lakehouse pipeline processing 250K+ records through OneLake and Bronze, Silver, and Gold layers using PySpark and SQL.
- Developed an AI analytics agent connected to curated Fabric data, enabling natural-language business queries and controlled data retrieval without requiring users to write SQL.
- Validated query accuracy across 50 predefined analytical questions, achieving ~90% successful execution and reducing average data lookup time from ~5 minutes to under 1 minute.