

VIJAYA KARTHIK

9363615606 | karthikvj398@gmail.com | [LinkedIn](#) | [GitHub](#)

Professional Summary

Highly disciplined Generative AI Engineer with a robust foundation in machine learning operations (MLOps), responsible AI frameworks, and advanced mathematics. Combines rigorous algorithmic problem-solving capabilities with specialized, enterprise-grade training to build and deploy scalable LLM applications. Leverages a strong analytical background to bridge complex technical architectures with ethically grounded, production-ready AI solutions.

Technical Skills

- **Programming Languages:** Python
- **Generative AI & ML Frameworks:** PyTorch, TensorFlow, Hugging Face, LangChain, Llamaindex
- **Cloud Platforms & MLOps:** AWS (Bedrock, SageMaker), Docker, MLflow, GitHub Actions
- **Core Competencies:** Retrieval-Augmented Generation (RAG), LLM Fine-Tuning, Responsible AI, Data Structures & Algorithms
- **Developer Tools & Libraries:** Git, FastAPI, Pandas, NumPy

Technical Projects

Multimodal AGENTIC RAG Integrated Knowledge AI Assistant | *Python, LangChain, Qdrant, Hugging Face, FastAPI, Docker* | March 2026

- **Architected a production-ready Multimodal RAG pipeline** to ingest and process diverse data formats (text, PDFs, images, audio) using PyMuPDF, OpenCV, and Whisper, establishing a highly scalable knowledge extraction engine.
- **Engineered semantic search and vector indexing** by generating embeddings with Sentence Transformers and storing 1,000+ document chunks in a Qdrant vector database, achieving sub-second vector search latency.
- **Integrated open-source Hugging Face LLMs** via LangChain to synthesize retrieved context, developing a custom evaluation pipeline that validated an 85–90% relevant context retrieval accuracy.
- **Enforced software engineering best practices** by structuring the system with a clean modular architecture (separating configuration, ingestion, and testing modules) and deploying the backend API using FastAPI and Docker.

Professional Development

- **Independent Technical Upskilling & Generative AI Engineering** | *June 2025 – Present*
 - Dedicated full-time focus to mastering the modern AI stack, culminating in multiple enterprise-grade certifications and the architecting of production-ready LLM pipelines and fine-tuned models.

- **Independent Study & Competitive Examination Preparation** | 2021 – 2024
 - Undertook rigorous, multi-year preparation for the highly competitive Indian Civil Services (UPSC), developing advanced research, analytical thinking, and complex information synthesis capabilities.

Advanced Technical Certifications

- **IBM Generative AI Engineering Professional Certificate** | *Coursera* * **MLOps (Machine Learning Operations) Specialization** | *Duke University (Coursera)* * **Responsible AI for Developers** | *Google Cloud (Coursera)* * **Mathematics for Generative AI** | *Udemy (Krish Naik)* * **Data Structures and Algorithms** | *Udemy (Elshad Karimov)*

Education

- **Master of Arts in English** | *Bishop Heber College* | 2021
-