UPEKA DILUKSHA *BSc (Hons.) Computer Engineering Undergraduate*

Email: dilukshaupeka@gmail.com | Phone: +94 77-5922852 | Portfolio: diluksha-upeka.me |
LinkedIn: upeka-diluksha | Github: upeka-diluksha

Built and deployed production GraphRAG and RAG systems processing 500+ page documents with semantic retrieval and graph-based entity extraction. Researching TGNN-LLM integration for real-time anomaly detection in Software Defined Vehicular Networks. Seeking AI/ML engineering roles focused on LLM orchestration, RAG pipelines, and agentic AI systems.

TECHNICAL SKILLS

Languages: Python, TypeScript / JavaScript, C/C++, SQL
AI & Machine Learning: LlamaIndex, LangChain, GraphRAG, PyTorch, Pandas, Scikit-learn, HuggingFace Transformers, Pinecone, Gemini API, OpenAI API, RAG, GraphRAG, Prompt Engineering, Vector Databases, Graph Neural Networks, Embedding Models
Backend & DevOps: Node.js, Express.js, FastAPI, Docker, CI/CD, Neo4j, MongoDB, MySQL
Frontend: React, Next.js, React Native, TailwindCSS

EXPERIENCE

Software Engineering Intern - Capricon Solution 2025 May – 2025 Nov
- Engineered and maintained backend REST APIs for an enterprise-scale POS system using an MVC architecture, optimizing database schemas to handle 1000+ daily transactions.
- Designed and optimized complex MySQL relational databases, applying indexing strategies to maintain high throughput data retrieval under heavy load.

Editorial Committee Lead, IEEE Student Branch, University of Ruhuna 2024 Jan – 2024 Dec
- Led the editorial committee for the IEEE student branch, producing technical publications and coordinating content strategy across events.

PROJECT EXPERIENCE

NeuroSpace - Multi-Modal Agentic AI Platform (Github) 2026 – Present
Python, FastAPI, Neo4j, Next.js, Docker, MinIO
- Built an end-to-end GraphRAG engine capable of extracting entities from 50+ page PDFs and 1+ hour MP4 video transcripts using local embeddings.
- Engineered a hybrid query pipeline combining vector search with graph traversal, reducing AI hallucinations.
- Modelled a Neo4j knowledge graph with 1000+ nodes and relationships, enabling graph traversal across multimodal document and video corpora.
- Developed a responsive Next.js frontend focused on UX-oriented implementation, creating an intuitive human-in-the-loop interface for users to interact with and validate multi-modal graph traversal results.

Temporal GNN-LLM Network Security - Final Year Research Project (Ongoing) 2026 – Present
Python, PyTorch, LLMs, Graph Neural Networks
- Supervised by Dr. P.A.D.S.N. Wijesekara, focusing on applied Deep Learning for network security.
- Developing an event-driven control flooding attack mitigation scheme for Software Defined Vehicular Networks (SDVN), modeling complex node interactions and routing behaviors applicable to logistics optimization.
- Integrating temporal graph neural networks (TGNN) with agentic LLM workflows to identify and neutralize network anomalies in real-time.

ContextIQ - Document Intelligence RAG Platform (Live) 2026 Jan – 2026 March
Python, Streamlit, Gemini, Pinecone, LangChain
- Built a cloud-deployed RAG system that ingests and embeds 500+ page documents into Pinecone, managing vector embeddings for semantic retrieval.
- Engineered and iterated on prompt designs via the Gemini API, evaluating LLM outputs to optimize context retrieval and ensure system reliability across multi-turn queries.

Salary Estimator - End-to-End MLOps Pipeline (Live) 2026 Jan – 2026 Feb
Python, Flask, Docker, GitHub Actions, Render
- Created a production ML system via an automated GitHub Actions CI/ CD pipeline to ensure fast feedback loops and system reliability.
- Containerized the Flask application with Docker and deployed to Render.

EDUCATION

BSc (Hons.) in Engineering - Computer Engineering 2021 – Present
University of Ruhuna - Faculty of Engineering, Expected Graduation: August 2026,

G.C.E. Advanced Level (A/L) - Physical Science Stream 2019
Mahinda College Galle, Result: AAB - District Rank 126

LICENSES & CERTIFICATIONS

Generative AI with Large Language Models: Issued by DeepLearning.AI
Supervised Machine Learning: Issued by DeepLearning.AI
Machine Learning in Production: Issued by DeepLearning.AI
Generative AI: Prompt Engineering Basics: Issued by IBM