

ARYAN

📍 Kurukshetra, Haryana, India | 📞 +91 9992416236

✉️ aryanupadhyay2301@gmail.com | 🌐 aryanupadhyay.com | 🔄 [Aryanupadhyay23](https://github.com/Aryanupadhyay23) | 📺 [aryanupadhyay23](https://www.linkedin.com/in/aryanupadhyay23)

PROFILE

AI/ML engineer with hands-on experience building production-grade ML, Deep Learning, GenAI, RAG, and MLOps systems, along with backend development using Python, FastAPI, and MySQL. Experienced in developing scalable APIs, deploying machine learning solutions, and building end-to-end AI applications.

TECHNICAL SKILLS

Languages: Python, C, C++

Machine Learning: Supervised Learning, Unsupervised Learning

Deep Learning: ANN, RNN, CNN, LSTM, GRU, Transfer Learning, Image Classification, Object Detection

MLOps: Git, GitHub, MLflow, DVC, Docker, Redis, AWS (EC2, ECR, ECS, S3, CodeDeploy), Weights & Biases, GitHub Actions, CI/CD, DagsHub

GenAI & Agentic AI: LangChain, LangGraph

Frameworks & Libraries: FastAPI, Flask, Pydantic, Streamlit, Pandas, NumPy, Matplotlib, Plotly

EXPERIENCE

UNIFYITC AND SERVICES (OPC) Pvt Ltd

July 2026 – Present

Software Developer

FleetControl CRM & Automated Regulatory Compliance Reporting Engine

- Developed scalable REST APIs using FastAPI, Python, Pydantic, and AsyncIO for fleet compliance, permits, taxation, subscriptions, and analytics modules.
- Designed and optimized MySQL stored procedures and complex SQL queries to aggregate large datasets and calculate real-time business KPIs, compliance statuses, and renewal metrics.
- Implemented asynchronous database operations using asyncmy connection pooling, improving concurrent request handling and database efficiency.
- Developed complex backend business logic for regulatory compliance and filing workflows, including date-based rules, reporting periods, status calculations, and automated compliance tracking.

PROJECTS

(Live: Mon–Fri, 9:00 AM – 5:00 PM)

YouTube RAG Intelligence

[🔗 Source Code](#) | [Live Link](#)

- Built an AI video assistant that allows users to paste any YouTube URL to ask questions, explore transcripts, and generate clear, structured summaries.
- Implemented a Corrective RAG workflow with LangGraph that refines user questions, verifies retrieved video context, and falls back to web search when answers aren't in the video.
- Combined semantic vector search (ChromaDB + Gemini embeddings) and keyword matching (BM25) to accurately locate relevant moments in the video transcript.
- Added timestamp-aware chunking to ground AI answers with clickable YouTube timestamps, allowing users to jump directly to the exact moment in the video.
- Created a Map-Reduce summarization feature to distill long video transcripts into concise key points and executive overviews.
- Developed a FastAPI backend for real-time streaming responses, containerized the app with Docker, and automated deployment to AWS EC2 using AWS CodeDeploy and GitHub Actions.

Conversational Research Paper RAG

[🔗 Source Code](#) | [Live Link](#)

- Developed a conversational research assistant that allows users to upload multiple academic papers and explore complex topics through continuous multi-turn dialogue.
- Built a hybrid retriever combining BM25 keyword matching and ChromaDB vector search to accurately locate relevant passages, extracting text and data tables to support cross-paper analysis.
- Optimized document ingestion using PyMuPDF for fast PDF table extraction and multi-threaded parallel batching to quickly process and index large research documents.
- Implemented verification logic in LangGraph that checks whether uploaded documents contain sufficient evidence before answering, falling back to Tavily web search when information is missing.
- Integrated source attribution with exact page badges, quote highlights, and interactive page previews so users can visually verify claims against the original document pages.
- Developed an asynchronous FastAPI backend delivering real-time streaming responses to a Streamlit frontend, con-

tainerized with Docker and deployed to AWS EC2 using GitHub Actions and AWS CodeDeploy.

Text Sentiment Analysis

 [Source Code](#) | [Live Link](#)

- End-to-end MLOps backend for a 3-class Twitter sentiment classifier (Positive / Negative / Neutral); powers a YouTube Chrome Extension returning real-time sentiment predictions via a Flask REST API.
- Built a 7-stage DVC pipeline; selected LightGBM via Optuna HPO across 8 models achieving 91.1% accuracy and 91.1% weighted F1; tracked experiments on [MLflow](#) via DagsHub.
- Built a full GitHub Actions CI/CD pipeline; runs pytest suites for model loading, signature, performance, and API validation before promoting the Docker image to AWS ECR; deployed to AWS EC2 via AWS CodeDeploy with model aliased @production only on all tests passing.

EDUCATION

Kurukshetra University, Kurukshetra
Master of Computer Applications (MCA) — Online

2026 – Present

Kurukshetra University, Kurukshetra
Bachelor of Computer Applications (BCA)

2023 – 2026

CERTIFICATIONS

- Advanced RAG — CampusX [\[View Certificate\]](#)
- Mathematics for AI/ML — Krish Naik [\[View Certificate\]](#)