

YASH DHANKHAR

Generative AI Developer • Full-Stack AI Applications • RAG & LLM Integration

+91 98686 31184 | yash.dhankhar06@gmail.com | [Linkedin](#) | github.com/OnlineBunker | New Delhi, India

PROFESSIONAL SUMMARY

Built and shipped 3 production-grade AI applications as a first-year B.Tech CS & AI student — including a RAG-powered PDF agent with page-level citation retrieval and an agentic churn analysis system with a 3-node LangGraph pipeline, Groq LLM integration, and semantic strategy retrieval. Specialized in full-stack LLM development using Google Gemini API, LangChain, and vector databases (FAISS, ChromaDB), with a consistent pattern of taking projects from idea to deployed product. Complementary skills in supervised ML and embedded systems enable end-to-end ownership across the entire AI stack.

TECHNICAL SKILLS

Generative AI: Google Gemini API, Groq, LangChain, LangGraph, RAG Pipelines, Prompt Engineering, LLM Integration, Vector Embeddings, FAISS, ChromaDB

Languages: Python, JavaScript (ES6+)

Full-Stack Web: React.js, Node.js, Express.js, REST APIs, Tailwind CSS, Streamlit, FastAPI

ML / Data: Scikit-learn, Pandas, NumPy, Logistic Regression, Decision Trees, EDA, Plotly

Tools & Platforms: Git, GitHub, Vercel, Render, Streamlit Cloud, MicroPython, IoT / Embedded Systems

Concepts: Retrieval-Augmented Generation, Agentic AI Pipelines, Document Intelligence, Supervised Learning, Data Visualization

PROJECTS

PDF Intelligence Agent — RAG-Powered Document Q&A

April 2026

Python • LangChain • FAISS / ChromaDB • Google Gemini API • React.js • FastAPI • PDF Parsing

GitHub: github.com/OnlineBunker/lumi | Live: lumi-wine-gamma.vercel.app/

- Built a full-stack AI web application enabling users to upload PDFs up to 20 MB (~270 pages) and query them through a conversational agent that returns answers with exact page-level citations.
- Implemented a complete RAG pipeline: PDF text extraction → chunking (500 tokens, 100-token overlap) → vector embedding with ChromaDB → semantic retrieval → Gemini-powered response generation.
- Engineered source attribution logic that surfaces the precise page number for every answer, increasing response trustworthiness and verifiability for end users.
- Designed a React frontend with real-time streaming responses and a Node.js/Express backend managing document ingestion, vector store, and LLM orchestration.
- Applicable to legal document review, academic research, and enterprise knowledge retrieval — demonstrating real-world GenAI product thinking.

Explainify — Adaptive AI Knowledge Explainer

March 2026

React.js • Node.js • Express.js • Google Gemini API • Tailwind CSS • Vercel • Render

GitHub: github.com/OnlineBunker/explainify | Live: explainify-drab.vercel.app/

- Engineered a full-stack GenAI web application that transforms any concept or text into tailored explanations across three audience levels: Kid, Student, and Expert.
- Built a React frontend with dynamic level-switching UI and an Express.js backend orchestrating multi-turn prompt chains with Gemini API for iterative simplification.
- Implemented prompt engineering patterns to enforce consistent tone, vocabulary, and depth control per audience level, demonstrating applied LLM behavior shaping.
- Deployed frontend on Vercel and backend on Render with independent CI/CD pipelines.

Customer Churn Prediction & Agentic Retention System

May 2025

Python • Scikit-learn • Groq • LangGraph • Pandas • Plotly • Streamlit • Streamlit Cloud

GitHub: github.com/OnlineBunker/churn-prediction | Live: churn-prediction-web.streamlit.app

- Built an end-to-end ML + agentic system combining a churn prediction model (Logistic Regression, ROC-AUC: 0.9037) with a Groq-powered conversational LLM interface for real-time retention strategy recommendations.
- Engineered a 3-node LangGraph agentic pipeline: `analyze_risk_node` for deterministic CRITICAL/HIGH/MEDIUM/LOW tier classification, `retrieve_strategies_node` for semantic top-3 strategy retrieval, and `generate_report_node` for Groq LLM report synthesis with anti-hallucination prompting.

- Identified customer support call frequency as the strongest churn predictor through correlation analysis; preprocessing pipeline (ColumnTransformer: StandardScaler + OneHotEncoder) ensures reproducible inference.
- Deployed as an interactive Streamlit dashboard with Plotly visualizations where users query the integrated LLM conversationally and receive data-grounded, personalized intervention plans.

EXPERIENCE

Embedded Systems & Data Analyst Intern | Rishihood University *Dec 2025 – Jan 2026 • Sonipat, India*

- Designed a MicroPython-based embedded system using ultrasonic sensors to collect real-time institutional food waste data, enabling automated pattern monitoring.
- Engineered temporal features (day-of-week, meal category, time-based aggregations) to structure raw IoT data for supervised ML classification.
- Trained Logistic Regression and Decision Tree classifiers; identified that weekend dinner periods showed 25% lower food waste than weekday averages, directly informing kitchen planning decisions.

Operations Executive | AIRIS Club, Rishihood University *Aug 2025 – Mar 2026 • Sonipat, India*

- Directed end-to-end operational planning for AI workshops, hackathons, and guest speaker sessions for the university's AI student society.
- Managed cross-team coordination and project timelines, improving on-time event delivery and member onboarding efficiency.

EDUCATION

B.Tech – Computer Science & Artificial Intelligence *Aug 2024 – May 2028*

Newton School of Technology, Rishihood University • Sonipat, India

Relevant Coursework: Generative AI, Machine Learning, Data Structures & Algorithms, Data Visualization & Analytics

Class XII – Science (PCM) | Rajkiya Pratibha Vikas Vidyalaya *2023 – 2024*

Delhi, India

ACHIEVEMENTS & EXTRACURRICULARS

- NDA-153 Recommended — recognized among top candidates for the National Defence Academy for demonstrated leadership, physical excellence, and national-level discipline.
- National Yoga Champion — First place at the Nationals Yoga Championship; demonstrates elite-level focus, consistency, and peak performance under pressure.
- National-Level Baseball Player — represented at national competitions, showcasing team coordination, strategic thinking, and athletic excellence.