

ANMOL YADAV

AI/ML Engineering Student

Hyderabad, India | anmolyadav2906@gmail.com | +91 9502863649

[GitHub](#) | [LinkedIn](#) | [Portfolio](#) | [OmniParse Demo](#)

PROFESSIONAL SUMMARY

AI/ML Engineering student with hands-on experience architecting full-stack intelligent applications, asynchronous API backends, and predictive models. Skilled at taking projects from concept to working prototype using modern Python and JavaScript frameworks, with a growing focus on applied machine learning, local LLM inference, and enterprise-grade AI safety workflows.

TECHNICAL SKILLS

- **Programming:** Python, C, TypeScript, JavaScript
- **Web Development:** Next.js, FastAPI, React, Framer Motion, Tailwind CSS, Flask, Streamlit
- **Data Science & ML:** TensorFlow, scikit-learn, Pandas, NumPy, Matplotlib, Logistic Regression
- **AI Engineering:** LangChain, Pydantic, OpenRouter, Ollama (Local LLMs), Prompt Engineering, Semantic ETL, Human-in-the-Loop (HITL) workflows

PROJECTS

OmniParse – Semantic ETL Engine Aug 2026

Next.js, FastAPI, LangChain, OpenRouter, Pydantic, Vercel, Render

- Architected an enterprise-grade Semantic ETL engine with Human-in-the-Loop (HITL) validation, preventing LLM hallucinations via strict boolean guardrails and Pydantic schema enforcement.
- Engineered self-healing execution loops using LangChain Expression Language (LCEL) to recursively debug and correct AI-generated JSON payloads.
- Deployed a full-stack architecture featuring a responsive Framer Motion UI on Vercel, seamlessly integrated with a cloud-hosted, asynchronous FastAPI backend on Render.

Lumina Async Summarizer Jul 2026

FastAPI, React, Framer Motion, Ollama, Tailwind CSS

- Architected a full-stack asynchronous AI text processing application enabling users to summarize massive document batches via local LLM inference (Phi-3).
- Engineered a fault-tolerant, sequential task queue in the FastAPI backend to strictly manage processing flow, preventing Out-Of-Memory (OOM) GPU crashes during heavy loads.
- Designed a minimalist, high-performance React frontend featuring real-time progress polling, dynamic error handling, and smooth physics-based animations.

M.A.D.E. – Autonomous Multi-Agent Data Engine Sep 2026

LangGraph, FastAPI, FAISS, React, Python

- Built an autonomous multi-agent data engine orchestrated with LangGraph, where agents self-heal through reflection loops and every generated step is audited in real time by an AST-based security pass.
- Won the Agentic AI Project Expo 2026. Integrated FAISS-backed retrieval and Human-in-the-Loop (HITL) guardrails, with a React Control Center providing full observability over agent execution.

EDUCATION

Siddhartha Institute of Engineering and Technology 2024 – 2028

B.Tech in Computer Science & Engineering

Hyderabad, India

Sri Chaitanya Junior Kalasala 2021 – 2024

Intermediate Education

Hyderabad, India

TRAININGS & EXTRACURRICULAR ACTIVITIES

- **Certifications:** Generative AI & TensorFlow (FreeCodeCamp), AI Tools Workshop (IIT Hyderabad)
- **Competitions:** Participated in an AI/ML satellite image classification challenge at Mahindra University.