

# Anurag Verma

Full Stack Engineer · Python, FastAPI, LLM Integration & Cloud Infrastructure

Backend systems, LLM integration, cloud infrastructure on Azure · \$30/hr · Available

Noida, India · [github.com/anurag629](https://github.com/anurag629) · [linkedin.com/in/anurag629](https://linkedin.com/in/anurag629)

## SKILLS

---

Python, FastAPI, Django, JavaScript, TypeScript, Bash, Azure, AWS, Docker, CI/CD, GitHub Actions, PostgreSQL, MySQL, MongoDB, Redis, LLM Gateway, RAG, OpenAI API, Azure AI Studio, Hugging Face, REST APIs, Pytest, asyncio, Celery, System Design

## ABOUT

---

Full Stack Engineer with 3+ years building production-grade backend systems, AI-integrated APIs, and developer tooling. Specialised in Python (FastAPI, Django), LLM integration, and cloud infrastructure on Azure. Shipped high-impact features at DataWars — including a multi-provider LLM gateway, a client-side code execution environment, and 300+ data science labs. Open source contributor and active technical writer.

## EXPERIENCE

---

### Full Stack Engineer · DataWars

Jul 2025 – Present

Remote · full time

- Built a multi-provider LLM gateway using async Python (FastAPI, asyncpg), unifying OpenRouter, Anthropic, AWS Bedrock, Exa, and Tavily under an OpenAI-compatible API with per-institution licensing and credit-based billing.
- Designed a database-driven provider resolution layer with 3-tier cascade fallback, enabling enterprise BYO-credential flows (e.g., customer-owned AWS Bedrock accounts) while maintaining unified billing.
- Implemented envelope-encrypted credential storage and upstream HTTP audit logging across Transaction, ProviderResult, and HTTPExchange ledgers for per-request cost reconciliation.
- Designed a browser-based LeetCode-style coding environment for Python and JavaScript, executing code entirely client-side — eliminating costly server infrastructure.
- Transitioned DataWars from server-heavy Docker-based systems to a lightweight client-side architecture, reducing operational costs by 35% and improving scalability.

### Data Science Technical Leader · DataWars

Jun 2023 – Jul 2025

Remote · full time

- Led a 9-member data science team developing 300+ interactive ML/data science educational labs, improving delivery efficiency by 50%.
- Architected a Pytest-based assertion library with 150+ statistical validation checks, reducing data quality issues by 40%.
- Designed automated data validation pipelines and QA frameworks for machine learning projects.
- Mentored team members on Python, statistical methods, and machine learning best practices.

### Volunteer Machine Learning Engineer · IVY

Nov 2022 – May 2023

Remote · freelance

- Contributed open-source ML integrations in TensorFlow, PyTorch, and JAX; merged 4 PRs adding framework-agnostic API adapters used across the IVY codebase.
- Collaborated with cross-functional teams to implement ML model serving and monitoring solutions.

## PROJECTS

---

### **ChatCops — AI-Powered DevOps Automation Bot**

FastAPI · LangChain · Python · GitHub Actions · Docker · PostgreSQL

- A natural language DevOps assistant built with FastAPI and LangChain, enabling engineers to trigger CI/CD pipelines, query logs, and manage deployments via chat.
- Integrated with GitHub Actions, Docker, and PostgreSQL.
- Reduced manual deployment steps by 60%.

### **OMEGA — AI-Powered Educational Animation Platform**

Django · Next.js · Python · Google Gemini · Azure OpenAI · Docker · Manim

- A platform for generating and executing Manim animation scripts using advanced AI models, built with a Django backend and Next.js frontend.
- Integrated Google Gemini and Azure OpenAI for script generation.
- Used Dockerized execution environments for safe, isolated animation rendering.