

Ayush Srivastava

miracneroid.in | [linkedin](#) | [github](#)

ayushsrivastava@gmail.com | +91-9621293445

Bengaluru, India

SUMMARY

AI/ML Engineer with production experience building LLM-integrated systems, RAG pipelines, agentic AI workflows, and scalable backend microservices using Python, FastAPI, and Docker. Proven track record delivering real-time inference services and ML pipelines for enterprise applications. SAP Certified Generative AI Developer; published researcher in deep learning.

EXPERIENCE

- AI/ML Engineer** Bengaluru, India
Merida Tech Minds | Python, FastAPI, LLMs, RAG, Async Python, Docker, ML Pipelines, MLOps May 2026 – Present
 - Architecting scalable **AI microservices** and distributed backend systems with FastAPI for intelligent **workflow automation** and employee analytics across enterprise clients.
 - Building production **ML pipelines** for predictive analytics and **intelligent document processing (IDP)**, integrating **LLM-integrated systems** with async task orchestration for reliable at-scale execution.
 - Developing **agentic AI workflows** and **real-time inference** services applying **RAG** and **prompt engineering** patterns; optimizing model-serving pipelines to reduce inference latency for enterprise deployments.
 - Designing **REST APIs** adhering to microservices principles and collaborating cross-functionally to ship AI-powered features into full-stack business applications using **MLOps** practices.
- Backend & AI Engineering Intern** Remote
Google Cloud Generative AI — AICTE | Python, Vertex AI, ETL, GCP Jul 2024 – Sep 2024
 - Built Python-based **ETL pipelines** processing 10,000+ samples on **Vertex AI**; improved pipeline efficiency by **35%** through optimized **GenAI workflow** orchestration.
 - Designed and deployed **GCP-hosted AI data pipelines** integrating cloud-native services for scalable generative AI workloads.
- Web Development Intern** Remote
TripFox Travellers | FastAPI, PostgreSQL Mar 2024 – Jun 2024
 - Built backend APIs supporting 1000+ concurrent users; optimized DB queries improving data retrieval by **40%**.

PROJECTS

- AdSync AI — Agentic Landing Page Personalization Engine** [Preview](#) | [GitHub](#)
Stack: FastAPI, React, Supabase, LLM APIs (Groq, Gemini, OpenAI), Docker, Netlify
 - Engineered multi-stage **agentic pipelines** processing 100+ dynamic personalisation workflows with **multi-provider LLM** orchestration across Groq, Gemini, and OpenAI.
 - Built FastAPI microservices handling 50+ real-time **AI orchestration** requests; deployed production architecture on Supabase Edge Functions and cloud infrastructure.
- MultiPDF Chat — RAG Backend System** [GitHub](#)
Stack: Python, FastAPI, Vector Embeddings, Semantic Search
 - Built **RAG** backend with **vector embeddings** and semantic retrieval across 5+ PDFs; reduced query response latency by **30%** via optimized **embedding retrieval** and similarity search.
- Emotify — Real-Time Emotion & Activity Recognition** [GitHub](#)
Stack: Python, TensorFlow, OpenCV, YOLO, CNN
 - Achieved **78% accuracy** and **20+ FPS** real-time inference using CNN + YOLO pose estimation; reduced inference complexity by 40% through modular backend architecture. *Published research paper.*

TECHNICAL SKILLS

- AI & ML:** LLM APIs, RAG Pipelines, Agentic AI, Prompt Engineering, Vector Embeddings, TensorFlow, OpenCV, MLOps, Intelligent Document Processing
- Backend:** Python, FastAPI, Django, REST APIs, GraphQL, Async Programming, Microservices, JWT Auth, MVC
- Databases:** Sql, SQLAlchemy, PostgreSQL, MySQL, MongoDB, Supabase, Redis
- DevOps & Cloud:** Docker, GCP, AWS, CI/CD, Git, Linux, Vertex AI, Netlify, Vercel
- Frontend:** React.js, Next.js **Languages:** Python, SQL, Java, JavaScript

EDUCATION

- Kalinga Institute of Industrial Technology (KIIT)** Bhubaneswar, India
Bachelor of Technology – Computer Science and Engineering — CGPA: 7.4 2022 – 2026
 - Relevant Coursework: Data Structures, Algorithms, Operating Systems, DBMS, Computer Networks

CERTIFICATIONS & LEADERSHIP

- SAP Certified** – Generative AI Developer (2026) [\[Certificate\]](#)
- Founder & President** – Coding Ninjas KIIT Chapter; led 10+ technical initiatives impacting 2000+ students.