

Sujeet Gund

+91-7620420850 | sujeetgund@gmail.com | sujeetgund.in | github.com/sujeetgund | linkedin.com/in/sujeetgund

SUMMARY

GenAI Engineer specializing in LLM orchestration, agentic AI systems, and RAG pipelines. Experienced in LangGraph multi-agent workflows, MCP server integration, vector retrieval (FAISS, pgvector), and FastAPI backends. Delivered production systems with HITL checkpoints, structured output contracts, and responsible AI guardrails. Deployed on AWS and GCP.

EXPERIENCE

Divam Technologies

Remote, India

GenAI Developer

Apr 2026 – Present

- Led end-to-end development of Public.ai, an AI-powered LinkedIn personal branding platform; sole architect across backend, infrastructure, and integrations, deployed independently on AWS
- Integrated official LinkedIn OAuth API for direct post publishing and built a scheduled posting system using BullMQ for reliable, time-based content delivery at scale
- Built a production LangGraph chat agent with pgvector RAG and PostgresSaver checkpointing for an enterprise client, handling multi-turn agentic conversations with persistent state
- Building and integrating MCP servers via Python SDKs to connect LLMs with external tools and enterprise data sources; applying responsible AI practices including prompt injection defense and Pydantic v2 output validation

EDUCATION

VIT Bhopal University

Bhopal, India

Integrated M.Tech in Artificial Intelligence, CGPA: 9.31/10.0

Expected Graduation: 2028

TECHNICAL SKILLS

Programming: Python, SQL, TypeScript

LLMs & Agentic AI: LangChain, LangGraph, OpenAI API, Prompt Engineering, RAG, HITL, MCP Servers, A2A Protocols, FAISS, Pinecone, pgvector, Hybrid Search

ML/AI: HuggingFace Transformers, scikit-learn, TensorFlow, PyTorch, Pydantic v2, Responsible AI

Backend & Cloud: FastAPI, PostgreSQL, MongoDB, Docker, AWS (EC2, S3, Lambda), GCP (Cloud Run, Vertex AI), Azure OpenAI, Git, Linux

PROJECTS

MailMind: Agentic Email Processing System | *LangGraph, FastAPI, Next.js, PostgreSQL, FAISS* [GitHub](#)

- Built an event-driven email automation agent using LangGraph StateGraph to classify intent and autonomously execute actions (reply, label, forward, escalate) via webhooks in real time
- Engineered HITL checkpoints with interrupt-based pause/resume and PostgresSaver-backed persistent state; integrated FAISS RAG for context-grounded auto-replies with SSE streaming and live monitoring dashboard

Legal Sahayak: Explainable RAG Legal Copilot | *LangChain, FastAPI, Next.js, FAISS* [GitHub](#)

- Built a hybrid FAISS retrieval pipeline across 6+ Indian legal domains using HuggingFace embeddings with semantic + keyword search; enforced structured LLM output contracts via Pydantic v2 schema validation
- Delivered a Dockerized system with confidence scoring, traceable source references, and demographic-aware prompt orchestration; secured via bearer-token API with Next.js 15 frontend

IPDR Graph Engine | *Python, FastAPI, Machine Learning, MongoDB, GCP* [GitHub](#)

- Architected an ML-powered telecom investigation platform achieving **94.16% accuracy** in anomaly detection; ingests large-scale IPDR logs, constructs communication graphs, and generates automated audit reports
- Deployed on GCP Cloud Run (FastAPI) and Vercel (Next.js) with MongoDB, supporting thousands of network records per run with interactive graph visualizations for security analysts

CERTIFICATIONS

Complete Data Science, ML, DL, NLP Bootcamp - Udemy (Dec 2025) | **Applied ML in Python** - U. of Michigan (Dec 2024) | **Supervised ML** - DeepLearning.AI (Jul 2023) | **AWS Cloud Essentials** - AWS (Jun 2023)