

# Meet Nakum

Generative AI Engineer - LLMs, RAG, Agentic Systems & Full Stack

meet.nakum31@gmail.com | +91 99130 54473 | mjnakum.com | github.com/MJNakum

## Professional Summary

Generative AI Engineer with **3+ years of experience** designing and deploying **LLM-powered applications, RAG pipelines, and agentic AI systems** in production. Proficient in **Python, Java, and AWS**; hands-on with **prompt engineering, vector databases, LLM fine-tuning (QLoRA), and human-in-the-loop evaluation**. Track record of leading teams to ship AI platforms serving **10,000+ users**.

## Professional Experience

**Zeus Learning** - Mumbai

Jan 2023 – Present

Senior Software Engineer - Generative AI & Full Stack

### AI Smart Evaluator - LLM-Powered Assessment Grading Platform

- Served as **lead** for an AI grading platform (Python, Java Spring Boot, React, AWS) automating assessment evaluation at **10,000+ evaluations** scale.
- Engineered a **multi-stage LLM grading pipeline**: sanitized attempt data, applied **PII masking** and **prompt injection guards**, then routed through a **parametric prompt system** with **dynamic per-evaluator personas** - mimicking each evaluator's historical strictness and feedback style to deliver personalized, consistent AI scores.
- Built a **validation and self-correction loop** - failed LLM outputs were automatically re-submitted with corrective instructions, ensuring reliable structured responses.
- Implemented **RLHF-style feedback capture**: evaluator score edits and feedback stored and used to continuously improve grading quality; conducted **A/B testing** after every prompt update to measure output quality.
- Led and mentored a **cross-functional team of 4 engineers**; owned architecture decisions, planning, and delivery execution.

### AI-Powered Chatbot - RAG & Agentic System

- Architected and deployed an enterprise **AI chatbot** using **Python (FastAPI)** with an **agentic workflow** for autonomous multi-turn conversation handling.
- Built a **RAG pipeline** using **vector embeddings, semantic search, and Pinecone**, deflecting **85% of common support queries** from live proctors.
- Designed **prompt engineering and context engineering** strategies to maximize output quality, reduce hallucinations, and enforce structured responses.
- Implemented **human-in-the-loop (HITL)** escalation: per-response feedback collected; repeated unresolved queries auto-flagged and transferred to human agents.
- Owned **CI/CD, Docker, and AWS deployment**; delivered as a standalone commercial product with zero-downtime rollouts.

### AI Proctoring Platform - Real-Time, Event-Driven System

- Designed **scalable Java Spring Boot microservices** supporting **10,000+ concurrent exam candidates** using event-driven architecture (AWS SQS, CloudWatch).
- Built embeddable **React/Angular SDK components** for LMS integrations, enabling **8 new client onboardings in 6 months**.

## Projects

### LLaMA-3 8B Fine-Tuning with QLoRA

2024

- Fine-tuned **LLaMA-3 8B** using **QLoRA** (r=16, targeting q\_proj & v\_proj, ~8-10M trainable parameters) on a domain-specific assessment dataset using **Hugging Face Transformers + PEFT**.
- Benchmarked fine-tuned model against a **RAG + OpenAI** pipeline; selected RAG approach based on systematic evaluation - demonstrating data-driven architectural decision making.

## Technical Skills

<b>AI / Gen AI</b>	LLMs (GPT, Claude, LLaMA), RAG Pipelines, Agentic AI, Prompt & Context Engineering, RLHF, HITL Evaluation, Vector Embeddings, Semantic Search
<b>AI Frameworks</b>	LangChain, LangGraph, LlamaIndex, PEFT (QLoRA/LoRA), Pinecone, FAISS
<b>Languages Frameworks</b>	Python (FastAPI, Django, Flask), Java (Spring Boot), JavaScript / TypeScript
<b>Backend &amp; Cloud</b>	Microservices, REST APIs, AWS (EC2, S3, Lambda, SQS, CloudWatch), Docker, Kubernetes
<b>Databases</b>	MySQL, MongoDB, DynamoDB, Redis, Pinecone
<b>MLOps &amp; DevOps</b>	CI/CD (Jenkins), A/B Testing, Experiment Tracking, Docker, Kubernetes

## Education

**B.Tech in Information Technology**

2019 – 2023

Birla Vishwakarma Mahavidyalaya, Vallabh Vidyanagar

CGPA: 8.0 / 10