

RAVI KANT SHUKLA

Senior AI Engineer

LLM • RAG • Agentic AI • Python • Java • Kafka • AWS

+91 7376986992
ravikantshukla91@gmail.com
linkedin.com/in/ravikantshukla91
github.com/ravikantshukla • ravishukla.in
Lucknow, India • Open to Remote

AI ENGINEERING EXPERIENCE

Senior AI / Backend Engineer — Sierra Circuits, Bangalore

— Sep 2023 – Present

▶ Oneflow — Intelligent PCB Data Extraction (Backend)

- Python ML pipeline parses Gerber & BOM files, auto-extracting 30+ PCB specs; cut form-fill from ~20 min to <2 min — 150+ uploads/day, ~15 hrs/week manual entry eliminated.
- Async ingestion: upload → S3 → Lambda → ML extraction → DB; REST APIs pre-populate order forms on upload with full history and diffing for repeat orders.

▶ PCB Complexity Scoring Engine (Sole Engineer)

- Rules-based AI engine scoring 20+ manufacturability params; replaced 30–45 min manual review with sub-second scoring — 100+ boards/day, ~60% faster order processing.
- AWS Lambda + API Gateway; Terraform + GitHub Actions CI/CD cut deploy time from ~2 hrs to 12 min.

▶ LangChain4j RAG Assistant (Sole Engineer)

- Production RAG pipeline (LangChain4j + Spring Boot + Python embeddings): PCB docs → vector store → LLM Q&A; resolves queries in seconds vs ~10 min manual search.
- Agentic tool-calling over live order data cut support escalations ~40%. Kafka streaming + Prometheus/Grafana across 6 microservices.

Senior Software Engineer — Globallogic, Noida

— Sep 2023

▶ Calabrio (Twilio/Cisco) • Halo Rise (Amazon) • COBA (Ericsson)

- Calabrio: Kafka consumers + Streams REST APIs for real-time CCaaS metrics — 500+ agents. Twilio/Cisco Webex integration.
- Halo Rise: Java microservices for Amazon AI health monitoring; Python Lambda ML pipelines; AWS CDK infra.
- COBA: Monolith → Docker/K8S/Helm migration; OpenTelemetry tracing across 8 services, ~35% faster debugging.

System Engineer | TCS — Lucknow

— Aug 2019 – Feb 2021

- Spring Batch + Quartz processing 10K+ records/day; REST APIs for real-time retrieval; GDPR cookie integration.

Software Engineer | HCL Technologies — Lucknow

— May 2016 – Aug 2019

- REST APIs (Spring MVC, BrightSpot CMS) for Thales; Angular 2 UI; 85%+ coverage via JUnit & Mockito.

PROFILE

9+ yrs backend engineering, 3 AI systems in production: Oneflow (Python ML, 150+/day), complexity scoring engine (~60% faster), LangChain4j RAG assistant (~40% fewer escalations). Available for fully remote Senior AI Engineer roles.

SKILLS

AI / LLM

Python (ML pipelines, embeddings), LangChain4j, Spring AI, RAG, Agentic Tool Calling, Vector DBs (Pinecone, Milvus), OpenAI / Claude / Llama

ML Integration

ML API integration, doc ingestion, embedding pipelines, structured output parsing

Languages

Java 17+ (expert), Python (proficient), TypeScript

Streaming

Kafka 4.0, Kafka Streams, Kafka Connect, Schema Registry

Cloud & DevOps

AWS (Lambda, ECS, S3, API GW, DynamoDB, IAM), Docker, K8S, Terraform, CDK, GitHub Actions

Observability

Prometheus, Grafana, OpenTelemetry, ELK, CloudWatch

Frameworks

Spring Boot 3, Spring WebFlux, Spring Batch, Spring Cloud

Databases

MySQL, DynamoDB, Cassandra, MongoDB, Vector Stores

AI PRODUCTS

Oneflow Python ML → Gerber/BOM extraction → auto-filled forms

Scoring Engine Rules-based AI, 20+ params, AWS Lambda, sole engineer

RAG Assistant LangChain4j + Python, agentic tool-calling, vector store

WEBSITE

ravishukla.in — AI/ML & backend blog, 500+ subscribers

EDUCATION

B.Tech – Electronics Engineering

Azad Institute of Engg & Tech, Lucknow | 2013

CERTIFICATIONS

AWS Certified Developer – Associate

Confluent Certified Kafka Developer

(Target Q3–Q4 2026)