Shubhankar Tiwari

Software Engineer | Backend Systems & Applied AI

📱 +91 9910608519
✉ tiwarishubhankar@gmail.com
🌐 shubhankar-tiwari.vercel.app
in linkedin.com/in/shubhankar-tiwari-514040165/

Professional Summary

Software Engineer with **4+ years of experience** in **Java, Spring Boot, Microservices**, and **cloud-native deployments on OpenShift**. Skilled in building scalable, high-availability backend systems and REST APIs for enterprise fintech, with ownership of **CI/CD pipelines**, cross-functional regression testing, release validation, and production incident triage. Kaggle **Notebooks Expert (Rank #2,441 / 59,663)** with applied AI research in LLM behavioral reliability evaluation and production quality-gating for AI outputs in financial services.

Experience

Apr 2022 – Present
Software Engineer I A, *Bank of America*, Hyderabad, India
- Develop and maintain client-facing fintech microservices using **Java, Spring Boot, REST APIs, SQL** on **OpenShift**.
- Own backend service delivery: design, implementation, **CI/CD integration**, deployment validation, and production stabilization.
- Lead **production incident triage** — deployment failures, SSL misconfigurations, database connectivity, service dependency conflicts; reduced recurring escalations.
- Perform **cross-functional manual regression testing** during release cycles — validating end-to-end system behaviour across service boundaries before production promotion.
- Improve release stability through proactive validation, pipeline debugging, load-test support, and root cause analysis.
- Contribute to **OpenShift onboarding** and deployment readiness, strengthening infrastructure resilience.

Aug 2021 – Mar 2022
Product Development Trainee, *HighRadius Technologies*
Built full-stack modules for **Order-to-Cash and Treasury Systems** using **Java, JavaScript, Python**. Contributed to Agile sprint delivery focusing on backend logic, API integration, and production readiness.

Projects

Jan 2026 – Present
LLM Reliability Evaluation Platform, *Python, Monte Carlo Sampling, HuggingFace, Sentence Embeddings*
Built a **behavioral consistency evaluation system** for open-source LLMs; ran **240 inference calls** on Qwen 2.5-7B across temperature $\times$ top-p grid (n=5 per prompt), computing instability, semantic dispersion, and entropy. Key finding: **top-p dominates reliability at 3.1$\times$ the effect of temperature** — optimal config T=0.1, top-p=0.5. Benchmark dataset published on Kaggle.

Sep – Dec 2025
Conservative Auto-Regeneration Policy for AI-Generated Financial Narratives, *Python, ROUGE-L, Sentence Embeddings, Bootstrap Calibration*
Designed a **production quality-gating pipeline** for AI outputs in financial services — **conservative AND policy** triggers auto-regeneration only when both embedding similarity AND ROUGE-L fall below thresholds calibrated via **bootstrap percentile sweep**. Built priority-scored human review queue with audit-grade provenance tracking and PII sanitization.

Jun – Aug 2020
Song Recommender System, *Python, K-Means, VADER, Scikit-learn*
Co-built a **workout playlist recommender** using K-Means clustering on BPM and VADER sentiment analysis to match songs to workout intensity. Co-authored with Tanish Maheshwari.

Education

2018 – 2022
B.Tech, Computer Science and Engineering, *SRM Institute of Science and Technology*, **94.4%**

Certifications & Recognition

Kaggle — **Notebooks Expert** — Rank #2,441 / 59,663 (personal best #707); 34 notebooks, 10 bronze medals; ML, Deep Learning & Computer Vision certified (2025)

2022 — REST API — Etrain

Technical Skills

Backend — Java, Spring Boot, Microservices, REST APIs, SQL, OpenShift, Jenkins, CI/CD, Distributed Systems

AI / ML — Python, PyTorch, HuggingFace Transformers, Scikit-learn, VADER, Sentence Embeddings, ROUGE-L, LLM Evaluation

Practices — Production Incident Triage, Regression Testing, Load Testing, Root Cause Analysis, Agile, Release Validation

Languages

English — Professional
Hindi — Native