

VIKRAM MARKALI

+1 (347) 822-2120 | vrm9190@nyu.edu | in/vikram-markali | github.com/Vikram739 | vikrammarkali.me | New York, USA

EDUCATION

New York University (NYU) New York, USA
Master of Science in Computer Science Aug 2025 - May 2027
Coursework: Artificial Intelligence, Data Structures & Algorithms, Big Data, Machine Learning, Database Systems, Software Engineering.

University of Pune (SPPU) Pune, India
Bachelor of Engineering in Information Technology (GPA: 3.82) Aug 2021 - Jun 2024
Coursework: Data Structures, Algorithms, Object-Oriented Programming, Operating Systems, DBMS, Computer Networks.

EXPERIENCE

Software Engineer Intern New York, USA
NYU's Endless Frontier Labs (EFL) Apr 2026 - Present

- Architected an **agentic AI sourcing platform** using **Python, RAG pipelines, and LLM orchestration** to scrape, score, and surface high-growth startups; reduced BD research time by **60%**.
- Built **LangChain** workflows with **FAISS vector retrieval** for thesis-aligned startup analysis, delivering structured insights to investment managers across **30+** portfolio founders.
- Engineered **FastAPI and Flask microservices** in **Python** automating sourcing, onboarding, and cohort management workflows end-to-end on a scalable **RESTful** architecture.

Freelance Software Engineer New York, USA (Remote)

Visnity AI (Startup) | Capstone Project Aug 2025 - Dec 2025

- Built AI-powered gig job-matching platform using **React, Python (FastAPI), and SQL** with scalable RESTful API architecture targeting instant-earning gig jobs.
- Implemented **LangChain LLM pipeline** for candidate-job matching with **Python**-based skill extraction; boosted match accuracy by **38%** with **HTML/JS** frontend integration.
- Designed **RESTful microservices** covering job listings, user profiles, and payment reconciliation with **Firebase** and **PostgreSQL** as data backends.

Software Engineer Pune, India

Worldline Global Services Pvt Ltd Oct 2024 - Jul 2025

- Architected **15+ RESTful APIs** using **Java (Spring Boot)** for payment gateway integrations, reducing merchant onboarding time by **40%** across distributed production systems.
- Engineered **Java** payment processing microservices achieving **99.9% uptime**; reduced transaction latency by **35%** and scaled throughput to **10,000+ TPS** for tap-and-pay flows.
- Optimized **SQL** queries and instrumented services with monitoring and alerting; resolved **8+ bottlenecks** and delivered zero-defect bi-weekly releases via Agile/Scrum.

Software Engineer Intern Pune, India

Maverick Labs Pvt Ltd May 2023 - May 2024

- Built **Java REST API** GPS tracking services for Tracko, Transporto, and InfraMachines; improved location accuracy by **40%** across **4,500+ concurrent users**.
- Optimized **PostgreSQL** queries with **Redis** caching, cutting data retrieval time by **50%**; refactored codebase using OOP principles, reducing server response time by **35%**.

TECHNICAL SKILLS

Languages: Python | Java | JavaScript | TypeScript | SQL | HTML/CSS | C/C++

AI/ML & Data: LangChain | RAG Pipelines | LLM Orchestration | FAISS Vector DB | TensorFlow | PyTorch | Tableau

Frameworks & Tools: Spring Boot | React | FastAPI | Node.js | REST API | Docker | AWS | Firebase | Git | CI/CD

Databases & Infra: MySQL | PostgreSQL | Redis | Linux | Agile/Scrum | System Design | Web Scraping

ACADEMIC PROJECTS

CourseBid Exchange Aug 2025 - Dec 2025

React, TypeScript, Java (Spring Boot), Spring AI, AWS, JWT/OAuth, SQL, Node.js

- Built full-stack course bidding platform with **Java (Spring AI)** LLM bid optimization, **JWT/OAuth** auth, and **SQL**-backed data layer serving **1,500+ active students**.
- Scaled **Java (Spring Boot) + AWS** backend with **Redis** session caching for **1,500+ concurrent users**, maintaining p95 latency under **300ms** with full CI/CD coverage.

OrbiNasaSense - Spacecraft Orbital Watch System Sept 2025 - Present

Python, TensorFlow, PyTorch, LSTM, Transformers, FastAPI, Firebase

- Built real-time anomaly detection for NASA spacecraft telemetry using **Python, TensorFlow, and PyTorch**; **Transformer** model achieved best F1 on long-range dependency patterns.
- Engineered automated **Python** data pipeline with multi-scale normalization and feature engineering, benchmarking **LSTM, Transformer, and Isolation Forest** models end-to-end without manual intervention.