

# VIJAY KARANJKAR

vkaranj@asu.edu | linkedin.com/in/vijay-k-2003 | github.com/Vijay-K-2003  
leetcode.com/u/Vijay-K-2003 | Google Scholar | +1 602-812-1888

## SUMMARY

Software Engineer with a 4.0 GPA M.S. in CS, hands-on experience in distributed systems, ML pipelines, and full-stack engineering.

## TECHNICAL SKILLS

**AI & ML:** PyTorch, Scikit-Learn, XGBoost, LightGBM, GNNs, SMOTE, TensorRT, RAG, LLMs, Agentic Workflows  
**Languages:** Python, Rust, Go, Java, TypeScript, C++, SQL, Bash, HTML, CSS  
**Web, Cloud & DevOps:** AWS (EC2, S3, RDS, EKS), Docker, Kubernetes, Terraform, GitLab CI/CD, Kafka, Redis, RabbitMQ  
**Tools & Frameworks:** Git, Linux, Airflow, FastAPI, React, Node.js, Flask, Prometheus, Grafana  
**Data:** Pandas, NumPy, PostGIS, Spark, PyTorch Geometric, SQL Databases, Data Pipelines

## EXPERIENCE

### Bhabha Atomic Research Center (BARC)

Research / Project Intern

Mumbai, India

Feb 2024 – May 2024

- Architected ML inference pipelines via Triton Inference Server and Docker to handle 100k+ requests with 99.9% uptime.
- Optimized neural models with TensorRT, boosting inference speed by 40% and cutting latency to ~0.2s on 1.2M images.
- Built an NLP data pipeline using LLMs and Fairseq, generating 50k+ synthetic samples and improving BLEU score.
- Trained multi-task models for face recognition, emotion, and spoof detection across a labeled dataset of 1.2 million images.

### Tikaj

Full Stack SDE Intern

Remote

Sept 2023 – Dec 2023

- Built event-driven systems using RabbitMQ and MongoDB, scaling sequential data processing for 1,000+ monitored sources.
- Integrated Redis caching to reduce database reads by 60%, achieving sub-0.2s latency for high-traffic security alert APIs.
- Implemented real-time notification pipelines for security events, delivering instant alerts across 1,000+ data watch sources.
- Developed modular microservice components using REST APIs, improving system maintainability and reducing deployment risk.

### Indian Institute of Information Technology Vadodara

Research Intern

Gujarat, India

May 2023 – July 2023

- Built a RAG-based code security analyzer using LLMs to detect OWASP Top 10 vulnerabilities across 5+ code repositories.
- Designed a document ingestion pipeline parsing 10,000+ lines of code into vector embeddings for vulnerability search.
- Integrated retrieval-augmented generation with LLMs to generate actionable fix suggestions for each detected security flaw.
- Evaluated on testing codebases, achieving 82% precision on injection and authentication-related OWASP issues.

### Extent

Frontend Developer Intern

Remote

Dec 2022 – Jan 2023

- Developed 6 pixel-perfect React and TypeScript dashboards for complex user datasets, cutting bug resolution time by 10%.
- Built a responsive React and SASS web application ensuring cross-device compatibility across 3 distinct screen breakpoints.
- Collaborated with the design team to enforce UI/UX consistency standards, reducing design revision cycles by 30%.
- Refactored frontend component architecture to improve reusability, reducing duplication by 25% across 6 dashboards.

## EDUCATION

### Arizona State University

M.S. Computer Science (GPA: 4.00)

Tempe, AZ

Aug. 2024 – May 2026

### IIT Vadodara

B.Tech. Computer Science and Engineering (CPI: 8.64)

Gujarat, India

Sept. 2020 – May 2024

## PROJECTS

### TCP Message Queue | Rust, Tokio, DashMap, Docker, Crates.io

- Built a Rust TCP broker using Tokio, sustaining 10+ Gbps throughput and 100+ concurrent connections with no blocking.
- Engineered zero-copy fan-out via bytes::Bytes and Arc, cutting memory usage by 87% routing 10GB to 8 subscribers.
- Designed a fixed 32-byte binary protocol reducing per-message parse overhead by 40% vs. JSON, with sub-2ms latency.

### Credit Card Fraud Detection | Python, PyTorch, PyTorch Geometric, XGBoost

- Built a ML pipeline on the IEEE-CIS dataset using SMOTE, reaching 97.50% accuracy with LightGBM on 1.1M rows.
- Trained and compared 7 models (XGBoost, LightGBM, RF, MLP) establishing baselines with up to 89.85% ROC-AUC score.
- Architected a GNN on 1.1M+ transaction-identity edges, boosting recall to 84% and cutting False Discovery Rate to 12%.
- Generated \$68.1M in projected savings by detecting fraudulent transactions with GNN, outperforming standard models.

### URL Shortener | Go, Fiber, Redis, Docker, REST API

- Built a high-performance Go/Fiber REST API with Base62 encoding and custom aliases, managing 10,000+ active short links.
- Integrated Redis caching and rate limiting to achieve sub-50ms link resolution latency for 5,000+ concurrent daily users.
- Containerized the service with Docker Compose for isolated deployments, enabling consistent environments across 3 stages.
- Implemented configurable TTL-based link expiration and real-time click analytics tracking for all 10,000+ registered links.