

Faizan Khan

(929) 548-2656 | faizanakhan2003@gmail.com | [linkedin.com/in/faizan-khan234](https://www.linkedin.com/in/faizan-khan234) | github.com/jellyfishing2346

EDUCATION

CUNY Brooklyn College

B.S. in Computer Science, Minor in Data Science

Brooklyn, NY

Expected May 2026

Relevant Coursework: Data Structures, Analysis of Algorithms, Data Tools and Algorithms, Machine Learning

TECHNICAL SKILLS

Languages: Python, JavaScript/TypeScript, Go, Java, C++, Rust, SQL, Bash

AI & LLMs: Anthropic Claude API, OpenAI, LangChain, RAG pipelines, prompt engineering, model fine-tuning, embeddings, Pinecone, PyTorch, TensorFlow

Full Stack: React, Next.js, Node.js/Express, FastAPI, Flask, RESTful APIs, distributed microservices, Docker, Kubernetes

Cloud & Data: AWS (Lambda, EC2, S3, EventBridge, CloudWatch), GCP, CI/CD pipelines, PostgreSQL, MongoDB, Redis, ETL pipelines, Git, Linux

EXPERIENCE

Brooklyn College

Research Assistant

Brooklyn, NY

Jun 2024 – May 2026

- Designed and shipped a production LLM pipeline for NERIS (National Emergency Response Information System), a federal platform tracking firefighter reports across the US — using the Anthropic Claude API and model fine-tuning to extract structured data from unstructured narratives and automate a manual supervisor review process end-to-end.
- Built Python ETL pipelines with automated validation, schema enforcement, and anomaly detection that cut manual processing time by 40% and reduced data quality incidents by ~30%, delivering production-grade datasets for downstream research and AI workloads.

InZone Inc.

Early-stage startup building engagement tools for student athletes

Remote

Full Stack Software Engineering Intern

Oct 2025 – Dec 2025

- Owned end-to-end delivery of a distributed microservices backend in Node.js/Express serving 5,000+ daily active users, leading architecture decisions from system design through production monitoring and performance tuning.
- Built and shipped consumer-facing iOS features in Swift and Dart, containerized multi-service applications with Docker and GCP achieving 99.9% uptime, and mentored 4 engineers on API design and coding best practices.

AutoLake LLC

B2B data platform providing automated data lake infrastructure for analytics teams

San Francisco, CA (Remote)

Software Engineering Intern

Jul 2025 – Sep 2025

- Drove a 25% increase in system throughput by standardizing RESTful API contracts across teams, and identified and remediated 15+ critical security vulnerabilities (XSS, CSRF) through systematic OWASP-aligned production audits.

PROJECTS

NERIS LLM Pipeline | Python, Anthropic Claude API, FastAPI, PostgreSQL

2025

- Designed an agentic LLM pipeline for a federal firefighter reporting system — engineered prompt workflows, fine-tuned the model on representative training data, and deployed a production API replacing a manual review process end-to-end.

Stellar Search | TypeScript, Python, FastAPI, PostgreSQL, Pinecone, open-source LLMs

2025

- Built a full-stack semantic search platform powered by open-source LLMs and RAG pipelines — designed distributed TypeScript and Python microservices delivering sub-200ms query latency across 100K+ indexed documents.

AI Content Detector | TypeScript, Anthropic Claude API, Manifest V3

2025

- Shipped a production Chrome Extension integrating the Anthropic Claude API, applying prompt engineering and validation logic to deliver real-time AI content classification to active daily users.

DenCity | React, Node.js, Express, PostgreSQL, Docker, AWS

2026

- Led end-to-end architecture for a real-time distributed data platform at a hackathon — React frontend, Node.js/Express microservices, AWS-deployed PostgreSQL with zero PII retention by design, shipped in 24 hours.

LEADERSHIP

Brooklyn College Computer Science Club

Student Treasurer

Brooklyn, NY

Jun 2023 – May 2026

- Scaled active membership to 100+ students by managing the club budget and coordinating technical workshops on AI engineering, full-stack development, and systems design.