Zain Syed

346-290-3918 | zain.alabideen.syed@gmail.com | linkedin.com/in/zain-syed1 | github.com/zain-syed01

Education

University of Houston Houston, TX

Bachelor of Science in Computer Science *Expected May 2028*

- **GPA:** 3.51 | **Honors:** Dean's List (Fall 2024)
- **Relevant Coursework:** Data Structures & Algorithms, Database Systems, Computer Organization & Architecture, Data Science & Statistical Learning, Discrete Mathematics

Professional Experience

IoT Cyber Defense Extern February 2026 – April 2026

Extern (partnered with Hydroficient) *Remote*

- Detected sensor anomalies (pressure drifts, flow surges) at a 0.757 F1 score, 0.736 Precision, and 0.780 Recall by deploying an unsupervised Isolation Forest model on IoT telemetry.
- Hardened a live MQTT data pipeline to a 100% attack rejection rate at +0.4 ms latency overhead by engineering a four-layer cryptographic defense architecture (mTLS, freshness, sequence counters, signing).
- Accelerated incident review in a simulated smart water IoT environment by building a real-time dashboard that translated model/security detections into actionable alerts.
- Documented threat analysis methods and defense results for mentor review, with clear evaluation notes for reproducibility.

Project Experience

Multi-Agent Equity Research Terminal | *LangGraph, FastAPI, ChromaDB, PostgreSQL* September 2026

- Architected a multi-agent orchestration pipeline using LangGraph to concurrently execute dedicated LLM agents for quantitative financial analysis and semantic news retrieval, synthesizing findings through an arbiter agent.
- Prevented hallucinated investment recommendations by engineering a deterministic guardrail node that programmatically validates explicit rating schemas and citations of core valuation metrics.
- Eliminated redundant LLM calls and minimized latency by deploying a 24-hour TTL Neon PostgreSQL cache layer and persistent ChromaDB vector store, containerized via Docker on Render and Vercel.

Predictive Maintenance Classifier | *Python, Pandas, scikit-learn, Matplotlib* October 2026

- Improved failure detection on the AI4I 2020 predictive-maintenance dataset (10,000 rows), raising failure-class F1 from 0.18 to 0.72 on a held-out test set, by comparing Logistic Regression and Random Forest with stratified splits and 5-fold cross-validation.
- Caught about 59% of rare failures (∼3.4% of rows) at 0.93 precision by evaluating models on failure recall/F1 instead of overall accuracy under class imbalance.
- Identified the strongest predictors of failure labels (torque, rotational speed) by running permutation importance on held-out failure-class F1.

AI-Powered Applicant Tracking System (ATS) | *Python, Gemini API, Streamlit, FastAPI* May 2026

- Automated candidate profile analysis to generate targeted feedback by integrating the Gemini API into a custom Python-based screening tool.
- Achieved accurate context retrieval against target job descriptions by implementing CRUD operations and a Retrieval-Augmented Generation (RAG) architecture using ChromaDB.
- Deployed the decoupled application for reliable public access by hosting the FastAPI backend on Render and deploying the interactive frontend on Streamlit Community Cloud.

Technical Skills

Languages: Python, SQL (PostgreSQL), JavaScript, TypeScript, Bash, HTML/CSS
ML / Data: Pandas, scikit-learn, Matplotlib, Isolation Forest, LangGraph, ChromaDB, Gemini API
Frameworks & Tools: FastAPI, Streamlit, Next.js, React, SQLModel, Pytest, Docker, Kubernetes, Linux, Git, Jira
Cloud: AWS, Azure, Terraform, GitHub Actions (CI/CD), Render, Vercel
Certifications: CompTIA A+, Google Cybersecurity Certificate (Coursera)