Aziz Ahmed

Chesapeake, VA | 757-969-9242 | ahmedazizswe@gmail.com | linkedin.com/in/nishataziz | github.com/nishat09

Education

Old Dominion University — Norfolk, VA
Master of Science in Cybersecurity, AI Security Concentration — *Expected Dec. 2027*

BRAC University — Dhaka, Bangladesh
Bachelor of Science in Computer Science and Engineering — *Mar. 2025*

Experience

Lead Developer — Jul. 2025 – Present
VIAstra, AI-Assisted Cervical Cancer Screening (Independent Project) — *Chesapeake, VA*

- Turned my BRAC research into a two-stage clinical support system: image screening plus risk prediction
- Fine-tuned ResNet-50 with transfer learning to classify VIA screening images as positive or negative
- Developed an ensemble model on patient clinical data that flags VIA-negative patients at risk of turning positive
- Built the end-to-end pipeline in Python with PyTorch, TensorFlow, and Keras, from ingestion to inference
- Tested with a partner clinic during research; now improving accuracy and robustness for clinical use

Software Engineer Intern — Jun. 2025 – Aug. 2025
Urmi Group — *Remote*

- Built software features in C++ from business requirements and technical specifications
- Worked across the full lifecycle: requirements analysis, design, development, testing, deployment, and maintenance
- Optimized application components for speed and responsiveness across devices and browsers
- Set up GitHub Actions CI/CD with automated tests and review gates; containerized services with Docker
- Resolved issues with senior developers, tracing pipeline failures to their root cause

Artificial Intelligence Research Assistant — Oct. 2023 – May 2024
BRAC University — *Dhaka, Bangladesh*

- Researched deep learning for cervical cancer screening from VIA images, the work that later became VIAstra
- Curated raw clinical image datasets into standardized, model-ready formats
- Built preprocessing and feature-extraction pipelines that raised model F1-score by 14%
- Ran ablation studies to test which data and model choices mattered
- Documented the pipeline for an internal publication and the lab's open-source prototype

Cybersecurity Intern (VAPT) — Jun. 2023 – Aug. 2023
Arena Web Security — *Dhaka, Bangladesh*

- Performed penetration tests on 3 web applications, from OSINT reconnaissance to exploitation
- Exploited OWASP Top 10 flaws: SQL injection login bypass, XSS, local file inclusion, and web shell upload
- Used Kali Linux, Burp Suite, SQLMap, OWASP ZAP, Acunetix, and Nmap for scanning and API fuzzing
- Wrote a 50+ page assessment report with step-by-step attack evidence, severity ratings, and prioritized fixes

Projects

Clinical NLP Pipeline | *Python, HuggingFace, Stanza, LLMs, PostgreSQL, Plotly*

- Built a pipeline that turns clinical notes from 1,000+ patient records into structured, coded data
- Extracted conditions and medications with Stanza clinical NER and mapped them to ICD-10-CM and RxNorm
- Benchmarked rules against an LLM at every stage, keeping the LLM as a fallback to protect patient data
- Made each entity traceable to its source stage, with a Plotly dashboard to monitor data quality

Semantic Document Search | *TypeScript, Next.js, Gemini Embeddings, Docker, Railway*

- Built a semantic search engine using Gemini embeddings and a cosine-similarity ranker written from scratch
- Set the relevance threshold from measured scores, so queries about the snake Python skip the language
- Cached embeddings by SHA-256 hash of the dataset and batched API calls to stay within rate limits
- Prevented duplicate indexing with a promise-based lock; returned clear status codes for every error case
- Wrote 29 unit and integration tests; shipped as a multi-stage, non-root Docker image on Railway

BurnOut: AI Health & Fitness | *Python, LangChain, OpenAI API, DataStax Vector DB*

- Built an agentic LangChain assistant that creates personalized health and fitness plans from user context
- Implemented RAG: user context is stored in a vector database, retrieved per query, and passed to the LLM

Personal Project

AgentWarden: Runtime Security for LLM Agents | *Python, FastAPI, scikit-learn, Ollama, SQLite*

- Built a security layer that screens tool outputs before an LLM agent reads them and tool calls before they run
- Combined an ML injection classifier, rule-based detectors, and taint tracking to block untrusted data flows
- Red-teamed 3 open models with poisoned web pages, emails, and PDFs
- Attacks succeeded in 4 of 21 runs without protection and 0 of 21 with AgentWarden

Technical Skills

Languages: Python, TypeScript, JavaScript, C++, SQL
Machine Learning: PyTorch, TensorFlow, Keras, scikit-learn, HuggingFace, Stanza
ML Methods: Computer Vision, NLP, Transfer Learning, Ensemble Models
Generative AI: LLMs, RAG, Embeddings, Semantic Search, AI Agents, LangChain
LLM Platforms: Gemini API, OpenAI API, Ollama, Vector Databases
Backend & Data: Next.js, Node.js, Express.js, FastAPI, REST APIs, PostgreSQL, SQLite, Plotly
DevOps & Testing: Docker, GitHub Actions (CI/CD), Railway, Git, Unit and Integration Testing
Security: Penetration Testing, OWASP Top 10, Kali Linux, Burp Suite, SQLMap, OWASP ZAP, Acunetix, Nmap

Certifications

Oracle Cloud Infrastructure Generative AI Professional
Google Cybersecurity Professional Certificate (In Progress)