


HOZAIFA OWAISI +1 919 836 3095

23 Monticello Dr Beer 19701 – howaisi.h@gmail.com



SUMMARY:

I am a **Certified AI Engineer for Data Scientist & Certified Data Analyst** dedicated and passionate about pursuing my interests in Software Development, Data Science and AI/ML. My drive stems from a deep love for tackling intricate challenges, whether in machine learning, data exploration and engineering, AI, WebSecurity development and operations. I thrive in the process of problem-solving, spending days searching for solutions banging my head against the wall, reaching out for help, experimenting with different approaches, learning from failures, and eventually uncovering the right path. It's the journey, not just the result, that excites me most. I think I am a quick learner, and I am often eager to join and explore new things.

Top Skills: Python, Java, Data Modeling, Exploratory Data Analytics, Machine Learning/Deep Learning, Microservices, AWS, Azure, Node.JS, JavaScript, SQL and NoSQL.

TECHNICAL SKILLS:

Cloud Platforms	: AWS (SageMaker, Lambda, S3), Azure (AI Foundry, ML Studio), GCP (Vertex AI, BigQuery)
Programming Languages:	: Python, Java, JavaScript, TypeScript, SQL, FastAPI, Flask, Django, pyVenv, Conda, ExpressJS, Spring Boot, NextJS
Machine Learning, Deep Learning & Data Science	: PyTorch, TensorFlow, Scikit-Learn, Keras, Ray Tune, MLOps, AWS SageMaker, Azure AI Foundry, Helsinki NLP, Agentic Systems, LangChain, Retrieval-Augmented Generation (RAG), GenAI, LLMs/vLLMs
AI Frameworks & Libraries	: LangChain, Hugging Face Transformers, Helsinki NLP, OpenAI APIs, PyTorch Lightning
MLOps & Experimentation	: MLFlow, SageMaker Pipelines, Azure ML Pipelines, GitHub Actions, Jenkins
Database & Data Processing & Analysis	: DynamoDB, MongoDB, PostgreSQL, SQL Server, AWS S3, Azure Warehouse, Azure Blob Storage & Apache Spark, Pandas, Numpy, Jupyter Notebooks.
Model Deployment & Serving	: FastAPI, Flask, Django, Docker, Kubernetes (K8s), Serverless Framework
Front-End (Prototyping & Visualization)	: ReactJS, AngularJS, Tailwind CSS, HTML, JavaScript, jQuery
Testing & Automation	: PyUnit, JUnit, Postman (REST API Automation)
Version Control & Build Tools	: Git, GitHub, Bitbucket, Maven, Gradle, NPM, Poetry, PIP
Static Code Analysis	: ESLint, Sonar, CheckStyle.
Services Technologies	: Microservices Architecture, SOA Architecture, Rest and SOAP Web Services.
Methodologies	: Agile, Scrum, Waterfall and Pair Programming.
Collaboration Tools:	: Slack, Teams, Discord, Confluence.

EXPERIENCE:

1. AI Institute of South Carolina – in collaboration with UNCP LLM Research and Development Engineer

Jan 2025 – Apr 2025

Project:

Knowledge Editing and Testing Framework for LLMs

Contributed to the research and development of a **Knowledge Editing and Testing Framework for LLMs**, focusing on enhancing generative AI capabilities through custom data integration, model responsiveness, and scalable deployment using MLOps best practices.

- Researched and implemented **LLM knowledge editing techniques** to update large models efficiently without full retraining, improving adaptability and reducing computational costs.
- Designed **custom augmentation pipelines and benchmarking strategies** to evaluate and optimize LLM editing outcomes across different data domains.
- Developed and deployed **multi-scale AI/ML solutions** tailored for generative AI applications, ensuring high scalability and low latency for a diverse user base of researchers.
- Leveraged **Amazon SageMaker** for end-to-end MLOps workflows including automated training, tuning, evaluation, and deployment of models.
- Conducted comparative analysis of **reasoning vs. non-reasoning LLM architectures**, generating insights into performance trade-offs and inference optimization.
- Led **system change management**, including LLM migrations, updates, and deprecation strategies, ensuring seamless transitions aligned with stakeholder goals.

Tech Stack:

Python, PyTorch, TensorFlow, Keras, Scikit-learn, LangChain, Retrieval-Augmented Generation (RAG), AWS SageMaker, AWS Lambda, AWS EKS, Custom Benchmarking Tools, Git, Bitbucket

2. NSA Laboratory for Analytic Sciences (LAS) @ NCSU Full Stack AI Engineer

Aug 2023 – Jan 2025

Project:

Internet Routing Integrity & RPKI Visualization Platform

Contributed as a Full Stack AI Engineer in a national security-focused research initiative, integrating deep learning models and LLM pipelines into a dynamic web platform for visualizing internet routing integrity and organizational risk classification.

- Spearheaded **frontend and backend development** using **Angular, Spring Boot**, and **REST APIs**, building responsive dashboards with real-time data visualizations.
- Designed and implemented **ETL/ELT pipelines** for ingesting and processing large-scale ARIN and sector classification datasets.
- Collaborated with the data engineering team to model, clean, and encode organizational metadata for downstream AI pipelines.
- Developed and deployed **LLM-based classification models** to categorize organizations by **CISA sectors**, enhancing routing integrity analytics.
- Built an interactive **LLM-powered chat interface** for querying internal data using **RAG (Retrieval-Augmented Generation)**.
- Established a **scalable LLM workflow** including data preprocessing, model training, fine-tuning, validation, and deployment using **RayTune** for hyperparameter optimization.
- Integrated a **custom vector database** using **PostgreSQL pgVector** to enable semantically enriched RAG queries.
- Authored comprehensive **technical reports** and presented research findings at the **NSA-hosted conference** on NCSU's campus.

-
- Coordinated sprints, task assignments, and stakeholder updates in an **Agile** team environment using **Slack, Discord, and Confluence**.

Tech Stack:

Angular, Java, Spring Boot, Python, Django, REST APIs, LLMs/vLLMs, RayTune, AWS (EC2, S3), ELT/ETL, Apache Spark, MongoDB, PostgreSQL (pgVector), Git, JUnit, PyUnit, Agile, Slack, Discord.

3. Corvid Technologies @UNCP

Jan 2023 – Dec 2024

DL Research and Development Engineer**Project:**

Model Development for Rapid Blood Clot Detection using DNNs

Worked as a Machine Learning R&D Engineer in a collaborative academic-industry research initiative focused on early detection of blood clot formation using deep learning applied to computational fluid dynamics (CFD) datasets.

- Designed and implemented **deep neural network (DNN)** models using **PyTorch** for blood clot prediction, significantly improving detection accuracy.
- Conducted **hyperparameter tuning** using **Ray Tune**, enhancing model performance through optimized training configurations in a scalable manner.
- Developed and automated **data preprocessing pipelines** to encode and scale complex **CFD datasets**, improving generalization and training convergence.
- Established a **modular ML pipeline** for data ingestion, training, validation, and evaluation, supporting reproducibility and long-term research scalability.
- Led the transition from a local development environment to a **cloud-orchestrated experimentation platform** using **AWS**, streamlining collaboration and resource usage.
- Built and maintained **CI/CD pipelines** for continuous experimentation, integrated with **Jupyter Notebooks** to support iterative model development.
- Explored and implemented techniques for minimizing **inference latency and model deployment time**, improving real-time usability of the system.
- Contributed to **two peer-reviewed journal publications** and delivered **three oral presentations** at academic and client-sponsored events.
- Maintained rigorous **version control** of code, data, and experiments using **Git**, ensuring reproducibility and collaborative transparency.

Tech Stack:

Python, PyTorch, TensorFlow, Scikit-Learn, Ray Tune, Pandas, NumPy, Git, AWS, Jupyter, Apache Spark, DynamoDB, REST APIs, CI/CD GitHub Action.

EDUCATION:**University of North Carolina at Pembroke (UNCP)**

Jan 2023 – May 2025

Pembroke, NC

Bachelor of Science in Computer Science

Concentration: Cyber Security

Personal Info:

DataCamp Profile: <https://www.datacamp.com/portfolio/hozaifa>

GitHub Link: <https://github.com/hozaifaO>

LinkedIn: <https://www.linkedin.com/in/hozaifa-o-a0b580266/>