

ABINESH LINGESWARAN

Sacramento, CA

716-427-4190

ABINESH.LING@GMAIL.COM

ABINESH LINGESWARAN

ABINESH-L

A Senior Software Engineer with 5+ years of experience in Applied AI/ML and software engineering, specializing in production-grade AI systems, RAG architectures, and LLM integration for enterprise data analysis. Proficient in hybrid retrieval systems, deep learning, NLP, and Generative AI, with a track record of building end-to-end ML pipelines from prototype to production. Proven leadership as a Director overseeing cross-functional teams in Agile environments. Trained 50+ students and career-switchers in ML, software engineering, and DSA through Sayur, a non-profit focused on tech education for underserved communities.

Professional Experience

HCLTech America | Senior Software Engineer

Sacramento, CA (June 2024 - Present)

- Architected an **LLM-powered analysis system** for **NVMe** semiconductor test data, leveraging **domain-adapted fine-tuning (LoRA/QLoRA)** on millions of test log entries to automate **narrative generation** and **visual analysis**, achieving **85%+ narrative quality** and reducing manual analysis from hours to **6 minutes**.
 - Combined **SQL precision queries**, **BM25 sparse retrieval**, and **dense semantic search** with **Reciprocal Rank Fusion (RRF)** to build a 3-layer hybrid retrieval architecture for **context-aware narrative generation**.
 - Engineered an **LLM guardrails and observability layer** with a deterministic **precision engine**, **hallucination detection**, and automated **quality scoring**, enforcing **100% numerical accuracy**, after resolving a critical **1000x unit conversion** error across the analysis pipeline.
 - Built an **agentic multi-LLM pipeline** with **self-healing fallbacks** and autonomous **provider routing** across **Dataiku** and **OpenAI**, reducing manual intervention to zero.
- Worked on the development of a machine learning-driven dashboard for testers, **enabling error forecasting**, **anomaly detection**, and actionable insights into root causes, optimizing test analysis and outcome reliability.
 - Created an ARIMA-based hybrid model with **Python UDFs in Snowpark ML** for **time series forecasting**, predicting error occurrences with **84%** accuracy.
 - Enhanced **anomaly detection** with **Snowpark ML**, using **autoencoder-based UDFs** to flag complex outliers in large-scale testing data (2.5 million entries per day), with integrated **feature analysis tools** that provide root cause insights, empowering testers with clear diagnostics for troubleshooting.
- Enhanced and optimized an **internal Flask + D3.js tester dashboard tool** for analyzing millions of daily test entries, resolving critical bugs & building new analysis features.
 - Resolved a browser crash caused by a **200MB+ bulk data load** and reduced page load time from **20+ minutes to under 10 seconds** by redesigning the Snowflake query architecture to fetch data on demand.
 - Identified and resolved a critical **data visualization bug** across **millions of test records** by implementing **per-record status validation logic** in the rendering pipeline, where incorrect status indicators were causing teams to **misdiagnose results**, improving **debugging accuracy** platform-wide.
 - Built a **FastAPI + Streamlit** clustering tool with **deep-linking**, **auto-zoom navigation**, **CI/CD pipelines**, and Excel-like drill-down modals. Integrated a Snowflake-backed analysis module with pre-computed pattern tables and a **27-test unit test suite**, actively used by teams to pre-screen issues before deep debugging.

Danlaw Inc | Software Engineer

Novi, MI (March 2023 - April 2024)

- Improved accuracy by **18%** for a **Drowsiness Detection System** using a hybrid metric combining **Eye Aspect Ratio (EAR)**, **Percentage of Closure eye time (PERCLOS)**, and **Gaze score**, integrating **GANs** for face synthesis and **GPT-3** to refine feature extraction.
- Worked on Algorithm Development for **Multi-step Trajectory Prediction** using **Time-series Analysis** & Sensor Fusion for a Radar
 - Designed an algorithm by fusing a **Unscented Kalman Filter (UKF)** and **Spatio-Temporal Graph Convolutional Network (ST-GCN)** for trajectory prediction of on-road vehicles using Radar data, thus enhancing prediction accuracy by 16%.

CIET, Feminist Pen Foundation | Machine Learning Engineer

India (Remote) (Jan 2021 - Jun 2022)

- Worked on software development of an **NLP-based SaaS tool "TraceX"**, for infection control of Covid-19 in an **Agile** environment (**Jira**)
 - Designed the information retrieval algorithm and built the ML risk estimation model for COVID-19 infection after **ETL processing**, applying feature selection and **hyperparameter tuning** for regularization, achieving a **precision of 0.78** and **AUC score of 0.8**.
- Worked on the development of a Deep Learning Model for **Cyberbullying Detection** & Prevention on Social Media.
 - Designed and trained a **hybrid RNN-LSTM Neural Network** for cyberbullying detection, leveraging BERT data augmentation and GPT-2 synthetic data generation to improve accuracy by 2.4%.

Feminist Pen Foundation | Director of Administration

India (Remote) (Aug 2020 - Jun 2022)

- Managed administrative personnel, budgets, and HR operations across a 17-person cross-functional team, overseeing departmental resource allocation.

Education

Master of Science in Engineering Science (Focus in Robotics and AI) (JUNE 2022)

University at Buffalo

Technical Skills

- Deep Learning Frameworks & Libraries:** TensorFlow, PyTorch, Keras, CUDA,
- ML-Based Cloud Platforms:** GCP AI, AWS Sagemaker, Snowflake, Snowpark ML, Dataiku
- Programming Languages & OS:** Python, SQL, Java, C++, Linux
- Big Data Technologies:** Scala, Apache Spark, NoSQL
- MLOps & Automation Tools:** MLFlow, Git, Weights & Biases, Docker, Kubernetes
- Web Frameworks:** Flask, FastAPI, Streamlit, D3.js, JavaScript

- GenAI Frameworks:** HuggingFace Transformers, LangChain
- Databases:** ChromaDB, FAISS, DuckDB, Parquet
- Libraries:** OpenCV, Sci-Kit, Pandas, Numpy, Scipy
- Data Visualization Tools:** Matplotlib, Seaborn, Tableau
- Agile & Collaboration Tools:** Jira, Slack
- Infrastructure & Tools:** nginx, REST APIs, CI/CD

Awards & Certifications

- Stanford Embark Competition, created TraceX - Finalist
- Wharton Entrepreneurship Acceleration Program