Robert Lewis

Email: contact@robertlewis.dev

robbylew robbylew

Education

Georgia Institute of Technology (Georgia Tech) Atlanta, GA
M.S. in Computer Science **Expected June 2026**

- **Relevant Coursework**: Database Systems Concepts and Design, Database System Implementation, System Design for Cloud Computing, Data Analytics and Security, Computer Networks

University of California, Los Angeles (UCLA) Los Angeles, CA
B.S. in Statistics and Data Science **Achieved June 2023**

- **Relevant Coursework**: Probability Theory, Deep Learning, Neural Networks, Bayesian Statistics, Computer Vision, Data Warehousing, Data Mining, Data Visualization, Data Structures & Algorithms

Technical Skills

Programming & Scripting: C, Python, SQL, Java, Bash

Cloud & MLOps: AWS, Azure, GCP, Docker, Kubernetes, Terraform, Airflow, Jenkins, GitLab CI/CD

Data Engineering & Pipelines: Kafka, Hadoop, Spark, Hive, Jenkins, dbt, Apache Beam

Databases & Storage: MySQL, PostgreSQL, MongoDB, Redis, S3, Elasticsearch, Cassandra

Projects

Twitter Sentiment Dashboard

- Developed a real-time analytics tool to gauge sentiment on fast-moving Twitter data, extracting **5,000+** tweets per minute via the **Twitter API**.
- Fine-tuned a **BERT-based Transformer** model with domain-specific data using **Hugging Face** and containerized the entire pipeline with **Docker** for reliable cross-environment deployments.

LLM-Driven Document Summarization Platform

- Aimed to streamline knowledge retrieval across **millions** of enterprise documents for faster decision-making.
- Engineered a scalable system by orchestrating **LLM** prompts with **LangChain**, leveraging **Pinecone** for semantic indexing, and containerizing the solution with **Docker**.
- Deployed on **AWS Fargate** for serverless scaling, reducing response latency by **30%**.

End-to-End MLOps Pipeline for Image Classification

- Implemented a full **Kubeflow**-driven pipeline handling data ingestion, model training, and validation; provisioned **GPU**-enabled **Kubernetes** clusters using **Terraform**.
- Reduced manual intervention by **40%**, cut inference latency by **30%**, and maintained cost-efficient scalability for large-volume image processing.

Certifications

CompTIA A+

CompTIA Network+