

Luke Anglin

AI Engineer | Norfolk, VA, USA | +17574781566 | lukeanglin@gmail.com

Professional Summary

I am an AI/ML Engineer with three years of experience in deploying large-scale cloud-based production software, specializing in analytics and data-driven solutions. I have a strong background in full-stack development and a proven track record of designing and architecting maintainable codebases. My expertise includes programming languages like Python, Go, and C++, as well as database querying languages like SQL. I have a strong understanding of data modeling, and have led the migration of large banking clients to modern cloud architectures in Azure and Snowflake, enabling more efficient querying on large datasets and reducing reliance on multi-million dollar on-premise infrastructure.

I have extensive experience with data pipeline development, including building, maintaining, and optimizing ETL/ELT pipelines. My experience includes multithreading Kafka operations to enhance system efficiency through concurrency and leveraging tools like Jenkins and CI/CD for efficient workflows.

I have worked with LLMs and led the development of agentic systems in order to find savings through automation. Projects range from structured output extraction on long electronic medical records to agentic React-based web UIs for company executives. I have also led company-wide classes on the daily usage of AI to help financial teams automate tedious tasks.

Additionally, I have deployed and maintained various machine learning models, ranging from basic regressions on large claims datasets to deep neural networks and complex diffusion models built on PyTorch. I am proficient in concurrent and parallel processing and researched distributed deep learning protocols at the University of Virginia.

Experience

Adobe

Forward Deployed AI Engineer

San Jose, CA, USA

Sep '25 - Present

- Managed AWS operations for image and video generation solutions for top Fortune companies
- Learned intrinsics of diffusion models and how to optimize model performance

Elevance Health

AI Engineer

Norfolk, VA

Mar '24 - Sep '25

- Engineered an API for an LLM-based, structured output extraction tool which scanned electronic medical records using OCR and computer vision (for handwriting), reducing dependencies on an external consulting firm. Estimated savings of ~\$2 million.
- Owned a reinforcement learning-based bid model for Medicare, reducing part of the financial team's workload by ~2 weeks per year.
- Built and deployed neural networks and temporal machine learning models using PyTorch and CUDA, capable of predicting member cost deciles with 95% accuracy.
- Directed team development of a multimodal Retrieval-Augmented Generation (RAG) applications for M&A teams, saving over 200 hours per project, outfitted with a recommendation engine for future questions.
- Designed and implemented a voice-responsive AI agent using React, JS, OpenAI API, and MySQL databases to automate outreach for members with unpaid balances.
- Led company-wide classes for company executives and finance teams on how to best use AI to automate daily tasks.

EY

AI Technology Consultant

Charlotte, NC

Sep '23 - Mar '24

- Architected and developed LLM-based applications using LangChain and LangGraph, resulting in a reduction of client team sizes and achieving ~\$4 million in savings.
- Designed and implemented various machine learning (ML) models. Models included gradient boosting, tree-based, ensemble, time-series, and neural networks, some for regression and some for classification. Crucial classification models boasted accuracy of above 90%.
- Led the migration of large banking clients to modern cloud architectures (Azure and Snowflake), enabling more efficient querying of large datasets and reducing reliance on multi-million dollar on-premise SQL server infrastructure.
- Managed large DBMS systems, including both relational MySQL databases and non-relational databases like Redis.
- Wrote high-performance, modern C++ code to improve slow and large-scale legacy financial systems.

Tesla

Data Science Intern

Fremont, CA

May '22 - Aug '22

- Multithreaded Kafka operations to enhance system efficiency through concurrency.
- Utilized a suite of tools including Jenkins, Splunk, Postman, and Grafana to analyze logs and debug complex issues, ensuring system reliability and performance.
- Worked on large, high-performance Golang repositories, leveraging CI/CD and source control to drive efficient mobile routing.

University of Virginia

Deep Learning Research Assistant

Charlottesville, VA

Sep '21 - Sep '22

- Employed distributed learning techniques to build and optimize models on resource-constrained IoT systems, enabling on-device learning without centralized data collection, and speeding up training by approximately 55% across GPUs.
- Built and refined deep learning models using Tensorflow, with a specific focus on optimizing GPU usage and familiarizing myself with GPU architecture.
- Built out machine learning and statistical models in high-performance, modern C++.

Elevance Health

Artificial Intelligence Team Intern

Palo Alto, CA

May '20 - Aug '20

- Performed health care industry competitive analysis of machine learning capabilities and examined algorithmic bias.

Education

University of Virginia

Bachelor of Science | Computer Science | GPA: 3.9

May '23

- Received a full academic scholarship
- Graduated with a 3.9 GPA in three years
- Participated in the UVA Robotics Club

Skills

API Development · AWS · Azure · C · C++ · CI/CD · Computer Science · Data Pipelines · Django · Docker · ETL · FastAPI · Flask · Full-Stack Development · GCP · German · Go · GPU/TPU · GraphQL · HTML/CSS · IOS · Java · Jenkins · JS · Kafka · Kotlin · Kubernetes · LLMs · Machine Learning · Mathematics · NumPy · Pandas · Python · Pytorch · R · Rust · Sci-Kit Learn · Security · Snowflake · Spanish · Statistics · TypeScript · Unreal Engine