

Terrance Luangrath

571.332.8954 | tksluangrath@gmail.com | [linkedin.com/in/terranceluangrath](https://www.linkedin.com/in/terranceluangrath) | [tksluangrath.github.io](https://github.com/tksluangrath)

EDUCATION

University of Virginia

Master of Science in Data Science

Charlottesville, VA

Aug. 2026

James Madison University

Bachelor of Science in Mathematics, Minor in Data Analytics

Harrisonburg, VA

May 2024

EXPERIENCE

Graduate Data Science Capstone Team Member

Smurfit WestRock (Sponsor)

Apr. 2026 – Present

Charlottesville, VA

- Enabled auditable, repeatable review of packaging-sustainability opportunities across a 12-market by 16-technology category matrix by contributing to a 4-agent pipeline (Ingest → Normalize → Score → Publish) built for a corporate sponsor
- Enforced data quality upstream by building a Python ingestion agent that captures evidence from internal and public sources into structured records, validating metadata requirements across 200+ entries and rejecting incomplete ones before they reach scoring

Junior Data Analyst

ACTFORE

Jun. 2024 – May 2025

Reston, VA

- Maintained 99%+ data accuracy across 1,500+ records processed daily by applying consistent validation checks in compliance with data-privacy standards

PROJECTS

Noise-Robust Keyword Spotting | [GitHub](#)

Jan. – May 2026

- Reduced model training time by optimizing the PyTorch training runner for improved GPU utilization, supporting a 4-person team's comparison of CNN and Audio Spectrogram Transformer architectures for noise-robust keyword spotting
- Tuned spectral-gating parameters across a 3x3 grid (ω , δ) by building an audio preprocessing pipeline and sweep evaluation script, improving noise-robust keyword classification on the Google Speech Commands and MUSAN corpora

Fraud Detection ML Pipeline | [GitHub](#)

Aug. – Sep. 2025

- Delivered a live HTTPS scoring API accessible via standard API queries by provisioning an Azure ML training job, container registry, and Azure Container Apps deployment
- Achieved 94% recall and 0.99 ROC-AUC across 6.3M+ transactions by training and evaluating a scikit-learn model on severely imbalanced data (0.13% positive), with MLflow-logged experiments and Pydantic-validated inputs

Course Recommendation Agent | [GitHub](#)

Jan. – Feb. 2026

- Built a solo end-to-end LLM agent that autonomously calls the right tool for a task by routing semantic search, skill-gap analysis, and timeline estimation through a LangChain agent backed by a locally-run Llama 3.1 model
- Delivered personalized learning paths across 2,759 courses and 1,754 skills by blending content-based filtering (sentence-transformer embeddings + ChromaDB) with collaborative filtering, containerized end-to-end via Docker Compose

TECHNICAL SKILLS

Programming: Python (Pandas, NumPy, scikit-learn, XGBoost, PyTorch, TensorFlow, FastAPI), SQL (PostgreSQL, SQLite), R, TypeScript

Machine Learning & AI: LLMs, Retrieval-Augmented Generation (RAG), LangChain, Prompt Engineering, Autonomous Agents, Supervised/Unsupervised Learning, Deep Learning

Tools & Infrastructure: Git, Jupyter, Docker, Azure ML, Azure Container Apps, Hugging Face, ETL and Data Pipelines