Oluseun Morenikeji Olulana

sma.olulana@gmail.com | +1 (774) 696-9152 | Worcester, MA
seunolulana.com | LinkedIn | GitHub: sewen007

PROFILE

Data scientist and Ph.D. researcher specializing in the scientific evaluation of machine learning and large language model (LLM) systems for high-stakes, decision-support contexts. Designs end-to-end pipelines spanning data acquisition, cleaning, modeling, and fairness-aware analysis to produce defensible, contextualized insights. Combines a rigorous mathematical foundation, multidisciplinary collaboration with domain experts, and 4+ years of applied ML engineering. Experienced in evaluating model behavior under uncertainty, communicating findings to technical and non-technical stakeholders, and translating scientific methodology into practical guidance for industry and academic partners.

AREAS OF EXPERTISE

Machine Learning Evaluation · Fairness in AI · Large Language Model Benchmarking · Algorithmic Auditing · Ranking Systems · Experimental Design · Statistical Analysis · Data Acquisition & Cleaning · Prompt Engineering · Retrieval-Augmented Generation · Mentorship & Technical Communication

EDUCATION

Ph.D., Data Science *Expected December 2026*
Worcester Polytechnic Institute (WPI), Worcester, MA
Dissertation focus: Fairness-aware ranking and re-ranking with large language models on tabular decision data; evaluation under missing or inferred sensitive attributes.

M.Sc., Pure & Applied Mathematics
African University of Science & Technology (AUST), Abuja, Nigeria
Thesis: Quadratic Forms with Applications

B.Sc., Mathematics
Federal University of Agriculture, Abeokuta, Nigeria

RESEARCH EXPERIENCE

Graduate Research Assistant — Fairness in ML & LLMs *May 2022 – Present*
Data Science Program, Worcester Polytechnic Institute (WPI)

- Designed and executed large-scale evaluation pipelines benchmarking commercial and open-source LLMs (Gemini, DeepSeek, Llama) on fairness-aware ranking tasks across structured tabular datasets, producing reproducible quantitative comparisons against several baseline strategies.
- Built modular Python frameworks (FAIRER-TabLLM) for structured prompting, ranking serialization, and metric computation, enabling rigorous fairness–utility trade-off analysis suitable for audit-grade reporting.
- Developed methodologies for fair re-ranking under missing or inferred sensitive attributes — addressing a recurring real-world constraint where demographic data is incomplete, restricted, or legally unavailable.
- Translated experimental findings into peer-reviewed publications and conference presentations, communicating technical depth to scientific reviewers and applied insights to broader audiences.
- Mentored undergraduate research assistants in reproducible experimentation, statistical evaluation, and academic writing.

SELECTED RESEARCH PROJECTS

FAIRER-TabLLM: Fairness-Aware Re-Ranking with Large Language Models *WPI, 2024 – Present*

- Built a modular framework that prompts LLMs to reason about equity over tabular candidate data and quantitatively benchmarked it against nine published fairness-aware ranking algorithms — surfacing where LLM-based approaches improve or degrade fairness–utility trade-offs.
- Produced visual and statistical analyses to support transparent communication of model behavior to non-ML stakeholders.

Hidden or Inferred: Fair Learning-to-Rank with Unknown Demographics *WPI, 2023 – 2024*

- Investigated how fair learning-to-rank models behave when sensitive demographic attributes are unavailable at training or inference time — a setting that mirrors real-world data restrictions in hiring, lending, and resource-allocation pipelines.
- Designed and ran controlled experiments comparing two practical strategies: (i) hiding the sensitive attribute entirely and (ii) using a separate classifier to infer it, quantifying how each affects fairness and utility across multiple ranking algorithms and datasets.
- Established empirical guidance for when inference-based approaches are preferable to attribute-blind approaches, and where both fail — published at the AAAI/ACM Conference on AI, Ethics, and Society (AIES 2024).

PetMatch: Image Retrieval with Convolutional Autoencoders *WPI, 2023*

- Contributed to a team effort to train a convolutional autoencoder for image similarity search; deployed a Flask interface for real-time matching, demonstrating end-to-end ML system delivery from data preparation to user-facing application.

Concept-Drift Detection via Multi-Task Learning *WPI, 2022*

- Collaborated with a team to implement PyTorch multi-task models for monitoring model reliability over time — a practical concern for deployed ML systems in regulated or safety-relevant industries.

PUBLICATIONS

- **Olulana, O.**, Murai, F., & Rundensteiner, E. (2026). *Can LLMs Rank Candidates with Missing Sensitive Attributes Fairly?* Transactions on Machine Learning Research (TMLR). [Accepted, to appear]
- **Olulana, O.**, Cachel, K., Murai, F., & Rundensteiner, E. (2026). *Can Large Language Models Rank Tabular Data Fairly?* In 2026 International Conference on Semantic Computing (ICSC) (pp. 226-229). IEEE.
- **Olulana, O.**, Cachel, K., Murai, F., & Rundensteiner, E. (2024). *Hidden or Inferred: Fair Learning-to-Rank with Unknown Demographics.* Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES), vol. 7, pp. 1088–1099.
- Pietrick, A., Romportl, A., Smith, S., **Olulana, O.**, Cachel, K., & Rundensteiner, E. (2022). *Are Fair Learning-to-Rank Models Really Fair? An Analysis Using Inferred Gender.* IEEE MIT Undergraduate Research Technology Conference (URTC).
- **Olulana, O.** & Degla, G. (2013). *Quadratic Forms with Applications.* M.Sc. Thesis, African University of Science and Technology, Abuja, Nigeria.

INDUSTRY EXPERIENCE

Technical Support Engineer (T4 Trainer, Microsoft 365) — Tek Experts (Microsoft Partner) *Sep 2019 – Dec 2021*
Lagos, Nigeria

- Provided enterprise-level diagnosis and remediation for Microsoft 365, Exchange, Active Directory, and Azure AD across global corporate customers — operating at the intersection of customer-facing communication and rigorous technical root-cause analysis.
- Trained and mentored frontline engineers, contributing to standardized troubleshooting playbooks and faster ticket-resolution workflows.

Android Developer — DiamondScripts Ltd. *Mar 2018 – Aug 2018*
Lagos, Nigeria

- Built Android and PHP/MySQL backend systems for retail and logistics clients, including admin dashboards and tracking-data integrations.

Freelance Software / Android Developer *2018 – Present*

- Delivered Android, web, and AI-integrated projects for global clients, including an RFID-scanner Android app with REST API integration and real-time geolocation tracking.

TEACHING, MENTORSHIP & LEADERSHIP

Teaching Assistant — Data Science (Graduate & Undergraduate) *2022 – Present*
Worcester Polytechnic Institute

- Supported instruction in Python, data preprocessing, model evaluation, and ethical decision-making; led office hours and code reviews for 50+ students each semester.

President (formerly Vice President) — African Graduate Student Association (AGSA), WPI *2024 – Present*

- Lead executive team and coordinate community-wide events, including the AGSA Graduation & Hand-Over Celebration; oversee strategy, communications, and stakeholder engagement.

Lead Organizer — AfriDev Hackathon 2025 *Oct 2025*

- Organized a university-wide hackathon focused on applied AI for sustainable development; ran technical workshops on responsible ML and mentored multidisciplinary student teams.

Parliamentarian — NSBE (WPI Chapter) *2025 – Present*

Conference Reviewer — AIES, EAAMO *2025, 2026*

HONORS & AWARDS

- AAAI/ACM AIES Student Travel Award (NSF-sponsored), 2024 — competitive award for AIES Conference attendance.
- AAUW International Doctoral Fellowship, 2022–2023 and 2019–2020 — two-time recipient ($20,000 each), American Association of University Women.
- African Development Bank (AfDB) Scholarship — graduate study in Science and Engineering, AUST.

TECHNICAL SKILLS

Languages: Python, R, SQL, Java (Android), PHP
ML / Data Science: PyTorch, TensorFlow, scikit-learn, Pandas, NumPy, Matplotlib, FastAPI, Flask
LLMs & AI: Gemini, Llama, OpenAI APIs, DeepSeek, prompt engineering, retrieval-augmented generation (RAG), chain-of-thought evaluation
Methods: Experimental design, statistical analysis, algorithmic fairness metrics, ranking evaluation, reproducible research workflows, Git
Certifications: Microsoft Certified: Azure Fundamentals; Udacity Deep Learning Nanodegree