Kumbirai Muhwati

New York, NY 10305
KMuhwati@my.harrisburgu.edu | linkedin.com/in/kumbirai-muhwati
Portfolio | GitHub Profile | National Cyber League Report

PROFESSIONAL SUMMARY

Data scientist specializing in security data science: machine learning for fraud and threat detection, privacy-preserving and federated learning, secure LLM fine-tuning, and auditable data pipelines for AI. Holds an M.S. in Cybersecurity from Yeshiva University (CGPA 3.966). As part-time AI Lead Researcher at UsHealthCloud, leads in-enclave LLM fine-tuning for oncology research on Trusted Execution Environments, with compute support from NVIDIA, Google, and Lambda. Author of five research papers on federated learning, deepfake detection, fraud and intrusion detection, and of nine public end-to-end data science projects with tested code and measured results, two of them with live demos. Fluent in Python and SQL, also working in Rust. Nominated Cybersecurity Innovator of the Year (2026 Cybersecurity Excellence Awards).

EDUCATION

Yeshiva University, *Katz School of Science and Health*, New York, NY **CGPA: 3.966**
Master of Science in Cybersecurity, August 2025

- **Honors:** Science and Technology Fellows Scholarship ($14,390).
- **Relevant Coursework:** Architecture of Secure Operating Systems, Network/Data Security, Cybersecurity Audit, Digital Forensics, Capstone Project.

Africa University **CGPA: 3.6**
Bachelor of Science Honors in Computer Information Systems, June 2021

TECHNICAL SKILLS

- **Languages & Frameworks**: Python (Fluent), SQL, Rust, C#, FastAPI, Django, Flask, Next.js, SQL Server, pytest, Git.
- **Machine Learning & Statistics**: Supervised learning (LightGBM, gradient boosting, random forest, scikit-learn), anomaly and fraud detection, model evaluation and calibration, Explainable AI (SHAP, XAI), causal inference and A/B testing (uplift modeling, EconML, off-policy evaluation), probabilistic forecasting, learning to rank.
- **AI & LLMs**: Empirical ML research, LLM fine-tuning, retrieval-augmented generation (RAG: hybrid BM25 + dense retrieval, reranking, vector databases), embeddings, LLM training-data pipelines (MinHash deduplication, benchmark decontamination, PII redaction, tokenization), Model Context Protocol (MCP), deepfake detection.
- **Security Data Science & Privacy**: Federated learning, differential privacy, Trusted Execution Environments (Intel SGX), threat detection, network traffic analysis, fraud analytics, data-leakage audits, PII detection, Zero Trust Architecture, NIST 800-53/171, HIPAA, risk management.
- **Data Engineering & MLOps**: pandas, Polars, NumPy, Parquet, DuckDB, dbt, Dagster, Apache Iceberg, LanceDB, streaming and change data capture, Docker, Kubernetes, CI/CD (GitHub Actions), model monitoring (drift).
- **Cloud, Infrastructure & Compute**: NVIDIA GPU ecosystems (H200, B200, Blackwell, H100, A100), Vast.ai, Lambda GPU Cloud, AWS (EC2, CloudTrail, IAM), Google Cloud Platform.
- **Cyber Operations**: Offensive security, cryptography, web application exploitation (SQLi, XSS), scanning and reconnaissance, digital and blockchain forensics.

PROFESSIONAL EXPERIENCE

UsHealthCloud, New York, USA January 2025 – Present
AI Lead Researcher (Part-Time)

- **Privacy-Preserving ML Research:** Lead the research program “Artificial Intelligence-Driven

Cybersecurity Framework for In-Enclave Large Language Model Fine-Tuning," fine-tuning LLMs on sensitive oncology research data inside Trusted Execution Environments so the data stays protected from state-sponsored espionage.

- **Secure ML Architecture:** Designed a privacy-preserving machine learning architecture that combines Trusted Execution Environments (TEEs) with the Model Context Protocol (MCP), enabling healthcare institutions to collaborate on AI models without exposing patient data.
- **ML Infrastructure & Compute:** Scale model training and experiments on NVIDIA H100/A100 GPUs with Docker and Vast.ai, secured through acceptance into the NVIDIA Inception Program and the Google for Startups Cloud Program, plus $7,500 in Lambda GPU Cloud credits.
- **Quantified Impact:** Developed a framework projecting $500M in reduced data-breach costs and a 30% reduction in clinical downtime for partner hospitals.

HSOC Cyber, New York, USA September 2025 – September 2026
Cyber Security Intern

- **Security Analytics & Alert Triage:** Analyzed security alerts and event data to separate genuine threats from false positives, escalating potential incidents to senior analysts.
- **Incident Data & Reporting:** Created structured incident documentation covering findings, evidence, and remediation steps, ensuring thorough and consistent reporting.
- **Vulnerability & Risk Assessment:** Supported vulnerability assessments by gathering system data and verifying security-patch status to inform risk prioritization.

Skicc Tech, Zimbabwe February 2022 – October 2023
Technical Lead (GridBox Innovation)

- **Applied AI Development:** Played a pivotal role in developing GridBox, an AI-powered solar mini-grid system.
- **Data & System Performance:** Reduced data latency and improved overall system uptime by approximately 20%.
- **Funding:** Helped secure $20,000 in competitive funding from the Milken-Motsepe Foundation.

NMB Bank, Zimbabwe February 2021 – December 2023
Digital Banking Systems Engineer

- **Risk & Customer Analytics:** Analyzed banking data to support credit-risk assessment, customer segmentation, and operational insights.
- **Data Quality Automation:** Reduced regulatory-reporting errors by 35% by automating and standardizing data-validation checks.
- **Risk Monitoring & Audit Analytics:** Built dashboards and automated data checks to monitor key risk indicators (KRIs) and data quality; led internal audits and supported external regulatory audits with timely, accurate data reporting.

DATA SCIENCE PROJECTS

Nine end-to-end projects with tested code, CI, and documented results at github.com/kmuhwati and kmuhwati.github.io; the first two run live in the browser.

- **Phishing URL Detector Evaluated on Live and Unseen Data** (Python, LightGBM, JavaScript, React). Live demo.
    - **Result:** Showed that the public PhiUSIIL benchmark rewards a shortcut (a 0.999-AUC model flags 99.9% of real web pages), then built an in-browser model that caught 84% of live OpenPhish links while flagging 2.1% of unseen real pages, with per-feature reasons for every score.
- **Small-Count Suppression Privacy Audit** (Python, SciPy linear programming, differential

privacy, React). Live demo.
    - **Result:** Recovered 216 of 1,733 hidden census counts (12.5%) from published totals alone; showed that complementary suppression still leaks through linear programming, while differential privacy at ε = 1 costs 0.77 people of error per cell.
- **Mobile-Money Fraud Detection with a Real-Time Scoring API** (Python, LightGBM, SHAP, FastAPI, Docker).
    - **Result:** Audited 6.36M transactions for label leakage, then deployed a rule plus LightGBM model that caught 99.7% of fraud (98.5% of fraud value) at 64% alert precision, with SHAP reason codes, drift monitoring, and 7 ms scoring.
- **LLM Pretraining Data Refinery** (Python, Dagster, MinHash LSH, fastText, tiktoken).
    - **Result:** Refined 339,259 raw Common Crawl pages into 46,552 deduplicated, benchmark-decontaminated documents (62.1M tokens), redacting 26,457 PII matches; every rejected page is traceable to the rule that removed it.
- **Live Wikipedia Lakehouse with a Continuously Fresh RAG Index** (Apache Iceberg, DuckDB SQL, LanceDB, streaming CDC).
    - **Result:** Streamed 157,412 live events into exactly-once Iceberg tables (0 duplicates after crash tests, vs. 3,359 with a conventional offset file); change data capture cut re-embedding by 96.5% and made edits searchable in 17 seconds (median).
- **Evidence Retrieval for Financial Question Answering (RAG)** (BM25, embeddings, LightGBM LambdaRank).
    - **Result:** Raised the share of FinanceBench questions whose evidence page reaches the LLM from 7% to 58% (recall@5) across 84 SEC filings, at about 3 ms per query on CPU.
- **Uplift Modeling and Profit-Based Targeting on a Randomized Experiment** (EconML, LightGBM, off-policy evaluation).
    - **Result:** Analyzed a 64,000-customer, three-arm A/B test with six uplift learners and doubly robust policy evaluation, finding the $0.27-per-contact break-even above which targeting pays.
- **Probabilistic Demand Forecasting and Inventory Decisions (Walmart M5)** (LightGBM, quantile regression).
    - **Result:** Beat AutoETS by 19% on WRMSSE across 8,230 series; critical-ratio ordering earned 3.7% more profit, about $113k a year for one department across ten stores.
- **Analytics Engineering for a 99k-Order Marketplace (Olist)** (SQL, dbt, DuckDB).
    - **Result:** Built 18 dbt models with 43 data tests; showed a late delivery raises 1-star review risk by 46 percentage points and that 63 sellers cause 13% of late orders.

RESEARCH & PUBLICATIONS

- "A Federated Learning Approach to Secure AI-Based Patient Outcome Prediction Across Hospitals" (Accepted August 2025).
    - **Key Finding:** Achieved <3% performance degradation compared to centralized training while preserving privacy.
- "Mitigating Deepfake-based Impersonation and Synthetic Data Risks in Remote Healthcare Systems" (Published September 2025).
    - **Key Finding:** >95% detection accuracy for deepfakes; eliminated session compromise in <2 seconds.
- "Machine learning and Blockchain approaches for enhancing fraud prevention in financial transactions" (Published July 2025).
    - **Key Finding:** Random Forest model achieved 99.9% accuracy/precision.
- "Mitigating Ransomware in the Energy and Healthcare Sectors through Layered Defense Strategies" (Accepted June 2025).
    - **Key Finding:** Layered defense delayed lateral movement by more than 20 minutes and postponed encryption by 45 minutes.
- "The Role of AI in Strengthening Privacy and Security in the Era of Cyber Crime" (Accepted

May 2025).

- **Key Finding:** Gradient Boosting Classifier achieved 91.3% accuracy in identifying attack traffic with an AUC score of 1.00.

CERTIFICATIONS & MEMBERSHIP

- Certificate of Cloud Security Knowledge v5
- National Cyber League
- ISC2 Cyber Pathway
- ISC2 Capture the Flag
- AI in Healthcare Specialization
- NVIDIA Certified Associate
- Certificate of Competence in Zero Trust
- Red Hat System Administration
- ISACA Membership
- Cyberbit-SOC Lab Program

AWARDS & INDUSTRY RECOGNITION

- Cybersecurity Innovator of the Year Nominee: 2026 Cybersecurity Excellence Awards.
- Judge: 2025 Global Recognition Awards (evaluated 20 shortlisted nominees) & 2025 Globee® Awards.
- Participant: (ISC)² Long Island 4th Annual Capture the Flag Competition (March 2025).
- **94th Percentile (Rank 38th/620):** National Cyber League (NCL) Fall 2025 Individual Game.
  - *Perfect Scores (300/300):* Web Application Exploitation, Scanning & Reconnaissance.
  - *High Proficiency:* Cryptography (90.5%), Password Cracking (80%).

VOLUNTEERING & AFFILIATION

- **Member:** ISACA (since March 2025).
- **Facilitator:** ISC2 New Jersey Chapter Spring 2025 Cyber Pathways Event.
- **Support:** 2025 CyberForce Competition.
- **IT Support:** St. Paul's Memorial Church (digital literacy training for senior parishioners).