

Fraud Detection in Financial Industries

Using Machine Learning as a Tool

Prepared By

Name: **ELSEDEEG AHMED**

ID No.: **251002546**

Name: **MOHAMED ABDELSATAR**

ID No.: **251001550**

**Submitted in fulfillment of the requirements for the degree Master of Science in
Artificial Intelligence**

At

College of Informatics

Midocean University

Supervisor Name: **Dr. Khaled Eskaf**

The Digital Revolution in Finance Creates a New Battlefield

Financial industries in the GCC and Saudi Arabia are rapidly digitizing, with mobile-based lending platforms handling high volumes of applications. This growth creates new vulnerabilities.

Traditional manual and rule-based systems are inefficient against dynamic fraud patterns, including identity spoofing, synthetic data manipulation, and falsified claims.



Rapid Digital Lending
Growth (GCC/KSA)




Increasingly
Sophisticated Fraud


- **The Challenge:** High-volume, noisy user data and complex feature interactions.
- **The Stakes:** Financial losses, erosion of institutional trust, and failure to meet regulatory obligations from bodies like the Saudi Central Bank (SAMA).

Our Mission: To Build a High-Accuracy, Transparent Fraud Detection Framework

How can financial institutions deploy a **high-accuracy, transparent machine-learning-driven fraud flagging tool** adaptable to Saudi mobile **consumer financing**, while ensuring **generalization reliability** and **regulatory maturity compliance**?



Black Box Model




Glass Box Framework

Key Objectives:

1. Develop and evaluate three distinct ML pipelines (Logistic Regression, Random Forest, XGBoost).
2. Propose a stacking framework governed by Explainable AI (XAI) for audit-level accountability.
3. Use XAI (SHAP & LIME) to bridge the gap between accuracy and auditor interpretability.
4. Ensure model maturity aligns with SAMA's governance and cybersecurity frameworks.

Previous Research Reveals a Performance Gap Between Simulated and Real-World Data

A review of existing literature shows **high performance on synthetic data**, but **more modest results on real-world loan applications**. This highlights the unique complexity of loan fraud compared to transaction-level fraud.



High Accuracy on Synthetic/Credit Card Data

- **Example 1:** 99.79% (Random Forest) on PaySim simulated data.
- **Example 2:** 97.92% (Naïve Bayes) on European credit card data.
- **Qualifier:** Does not capture the complexity of real lending environments.




Realistic Accuracy on Real Loan Data


- **Example:** 78.2% (Random Forest) in a loan fraud detection study by Kumar et al. (2022).
- **Qualifier:** Reflects the inherent challenges of limited behavioral signals and heterogeneous applicant profiles.

Identifying the Gaps: Our Opportunity to Contribute

Our analysis of prior work identified several critical gaps that this research addresses directly, creating a framework that is both academically rigorous and operationally relevant for the region.



Our End-to-End Blueprint for a Governed Detection System



Assembling Our Toolkit: Three Complementary Modeling Approaches


Logistic Regression (LR)



Role:
The Explainable Baseline

Strength:
Provides native interpretability and an efficient linear boundary, establishing a strong performance benchmark.


Random Forest (RF)



Role:
The Robust Ensemble

Strength:
Highly robust to noisy features. Consistently performs well in fraud studies and offers a strong balance of accuracy and audit-readiness.

XGBoost (XGB)




Role:
The High-Performance Specialist

Strength:
Superior at detecting complex, non-linear patterns and feature interdependencies common in sophisticated fraud schemes.

A Clear Progression Towards Higher Accuracy

Benchmarking the three models on the isolated test set revealed a significant and progressive improvement in performance, validating the increasing complexity of the learners.

F1-Score Comparison



Metric	LR	RF	XGBoost
Accuracy	92%	95%	100%
F1-score	91.6%	95%	1.00
ROC-AUC	96.5%	98.4%	1.00

A Perfect Score. But Perfection Demands Scrutiny.

The fast-tuned XGBoost pipeline achieved perfect separability on the isolated test segment, with zero misclassifications. This confirms the correctness of the implementation and the high quality of the preprocessed data.

XGBoost Confusion Matrix

	Predicted Fraud	Predicted Legitimate
Actual Fraud	458 (TP)	0 (FN)
Actual Legitimate	0 (FP)	422 (TN)

Accuracy:
100%


Precision:
1.00

Recall:
1.00

F1-Score:
1.00


The Governance Imperative: Beyond Perfect Test Accuracy

A 100% test score on a single financial corpus is an indicator of correct implementation, not a guarantee of real-world infallibility. Academic integrity and responsible deployment demand a discussion of generalization risk. True model reliability is found in the balance between performance and transparency.




Opening the Black Box: How XAI Delivers Auditable Transparency

Explainable AI frameworks like SHAP are essential for tracing prediction accountability. They quantify how much each feature contributed to a model's decision, providing the 'why' behind the 'what.' This aligns with SAMA requirements for transparency and decision reasoning.



An Integrated Framework: From Powerful Algorithm to Accountable System

By combining a high-performance model with a robust explainability layer, we create a complete fraud detection framework. This system mitigates misclassification costs (costly False Negatives vs. customer-impacting False Positives) and builds the institutional trust necessary for deployment.



High-Accuracy Model
(XGBoost)



Explainability Layer
(SHAP/LIME)



Regulator-Ready Framework
(SAMA-Aligned)

The final framework is designed for operational feasibility, improves prediction traceability, and enables regulator-aligned interpretation.

Conclusion: A New Paradigm for Governed Fraud Detection

This research successfully designed, built, and validated an end-to-end framework for financial fraud detection that prioritizes both accuracy and accountability.



Validated a high-performance ML pipeline on real, nearly-balanced Saudi loan data.



Demonstrated the critical need for governance and transparency over a singular focus on raw accuracy.




Pioneered an XAI-integrated framework for auditable, regulator-ready fraud detection suitable for the GCC financial industry.



Established that a stacked ensemble architecture (RF + XGBoost) governed by XAI offers the most realistic and scalable path forward.

The Road Ahead: A Privileged Roadmap for Future Work

This thesis provides a durable foundation. The next evolution layers focus on scaling the framework for production environments and enhancing its governance capabilities.



Building the Future of Finance on a Foundation of Trust

This thesis establishes a durable foundation for GCC consumer financing fraud detection engines using supervised learning stacking architectures, with **Explainability** governance penetrating reliability. The work privileges **transparency over perfection** and confirms Machine Learning's viability as a **business resilience tool**, not a confidence illusion.

