



**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

UAE

Supervisor Name: **Dr. Khaled Eskaf**

College of Informatics, Midocean University, Comoros

[2025-Semester 3]

## Acknowledgements

We would like to express our sincere gratitude to Dr. Khaled Eskaf, Supervisor at the College of Informatics, Midocean University, for his invaluable academic guidance, continuous support, and constructive feedback throughout the course of this research. His expertise in artificial intelligence and applied research methodology played a pivotal role in shaping the direction, rigor, and academic quality of this thesis.

We also extend our heartfelt appreciation to our families and close colleagues for their patience, encouragement, and moral support, which were essential to the successful completion of this work, particularly during the intensive phases of data analysis, model development, and documentation.

Our sincere thanks are extended to the executive technical leadership at **Emkan Finance**, specifically the **Chief Information Officer (CIO)**, for providing practical and valuable insights into real-world fraud detection challenges and for supporting the integration of artificial intelligence within an operational financial environment. Their initial approval to implement this module in Emkan's test environment was instrumental in proving the concept and enabling the model to be trained on a substantial volume of real customer atonement data. The exposure to real fraud scenarios, system constraints, and governance considerations significantly enriched the practical relevance of this research and reinforced the importance of aligning AI solutions with institutional resilience, regulatory compliance, and technological governance.

We further acknowledge the contributions of academic mentors and IT professionals who shared valuable perspectives on fraud scenarios, data governance, model reliability, and explainability. Their insights contributed to strengthening the robustness and transparency of the proposed machine learning framework.

Finally, we gratefully acknowledge Midocean University for providing a structured academic environment that fostered discipline, critical thinking, and research excellence, enabling us to complete this study in accordance with high academic and ethical standards.

## Abstract

Financial fraud has rapidly evolved alongside digital transformation, creating increased complexity for institutions that process high-volume financial applications, particularly in the consumer lending domain. Traditional manual screening and rule-based systems struggle with scalability, feature interdependence, and adaptability to dynamic fraud patterns. This study develops and evaluates machine learning-driven approaches to detect fraud in personal loan applications using a real-world tabular dataset consisting of 4,400 loan application records, including 2,110 fraudulent and 2,290 legitimate cases. The dataset is nearly balanced, which helps reduce extreme class imbalance effects commonly found in financial fraud data.

The research proposes a hybrid stacking framework integrating complementary machine learning ensemble learners with external Explainable AI (XAI) frameworks, specifically SHAP and LIME, to enable transparent decision-traceability for auditors and regulatory entities such as the Saudi Central Bank (SAMA). Three predictive pipelines were implemented, benchmarked, and optimized: Logistic Regression (baseline), Random Forest (grid-tuned), and a fast-tuned XGBoost pipeline. Optimization employed systematic hyperparameter estimation, stratified cross-validation, and in-pipeline preprocessing to reduce feature inconsistency, mitigate model bias, and ensure robust generalization-feasibility.

Empirical results demonstrated perfect separability on the held-out test set (accuracy = 100%, precision = 1.00, recall = 1.00, and F1-score = 1.00) for the fast XGBoost pipeline. While these outcomes confirm the correctness of the implementation and the robustness of the preprocessing pipeline, they also naturally open avenues for future academic discussion among researchers regarding potential overfitting risks under strong baseline assumptions. In the context of this graduation project, such results are interpreted as methodological validation rather than conclusive evidence of real-world fraud separability. Accordingly, additional reliability assessments using ROC-AUC and Precision-Recall-AUC curves were conducted to strengthen model accountability, support deployment feasibility, and mitigate misclassification costs in practical financial settings.

The study concludes that stacking diverse fraud detection learners within an ensemble XAI-governed framework creates an operationally feasible fraud-flagging reference for financial institutions, improves prediction traceability, enables regulator-aligned interpretation, and lays a replicable academic foundation for scaling AI-driven fraud mitigation systems in Saudi and GCC financial industries.

## Table of Contents

|                                                                                                                                                                                                                                                                                                                            |           |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <b>Chapter 1 - Introduction .....</b>                                                                                                                                                                                                                                                                                      | <b>7</b>  |
| 1.1 Background and Context.....                                                                                                                                                                                                                                                                                            | 7         |
| 1.2 Problem Definition.....                                                                                                                                                                                                                                                                                                | 7         |
| How can financial institutions deploy a high-accuracy, transparent machine-learning-driven fraud flagging tool adaptable to Saudi mobile consumer financing by using a hybrid stacking framework and robust XAI interpretation layers, while ensuring generalization reliability and regulatory maturity compliance? ..... | 8         |
| 1.3 Thesis Objectives .....                                                                                                                                                                                                                                                                                                | 8         |
| 1.4 Thesis Contribution.....                                                                                                                                                                                                                                                                                               | 9         |
| 1.5 Thesis Organization .....                                                                                                                                                                                                                                                                                              | 10        |
| <b>Chapter 2 : Literature Review and Conceptual Foundations .....</b>                                                                                                                                                                                                                                                      | <b>11</b> |
| 2.1 Artificial Intelligence in Financial Industries .....                                                                                                                                                                                                                                                                  | 11        |
| 2.3 Overview of Existing Literature.....                                                                                                                                                                                                                                                                                   | 11        |
| 2.4 Comparative Analysis of Previous Studies .....                                                                                                                                                                                                                                                                         | 12        |
| 2.5 Supervised Machine Learning for Fraud Detection .....                                                                                                                                                                                                                                                                  | 13        |
| 2.6 Logistic Regression as a Baseline Explainable Model.....                                                                                                                                                                                                                                                               | 14        |
| 2.7 Random Forest for Robust Fraud-Flag Ensemble Learning .....                                                                                                                                                                                                                                                            | 15        |
| 2.8 XGBoost for High-Accuracy Non-Linear Fraud-Flag Handling .....                                                                                                                                                                                                                                                         | 16        |
| 2.9 Model Reliability Caveat — XGBoost .....                                                                                                                                                                                                                                                                               | 17        |
| 2.10 Explainable AI Frameworks—SHAP and LIME.....                                                                                                                                                                                                                                                                          | 17        |
| 2.11 Ethical and Regulatory Alignment.....                                                                                                                                                                                                                                                                                 | 18        |
| 2.12 Model-Stacking for Synergistic Fraud-Flag Prediction .....                                                                                                                                                                                                                                                            | 18        |
| 2.13 Research Gaps and Opportunities .....                                                                                                                                                                                                                                                                                 | 18        |
| 2.14 Conclusion .....                                                                                                                                                                                                                                                                                                      | 19        |
| <b>Chapter 3 Related Work.....</b>                                                                                                                                                                                                                                                                                         | <b>20</b> |
| 3.1 Introduction to Financial Fraud in the Digital Industry .....                                                                                                                                                                                                                                                          | 20        |
| 3.2 Traditional Fraud Detection Approaches .....                                                                                                                                                                                                                                                                           | 20        |
| 3.3 Machine Learning for Financial Fraud Detection .....                                                                                                                                                                                                                                                                   | 21        |
| 3.4 Explainable AI in Financial Fraud Detection.....                                                                                                                                                                                                                                                                       | 23        |
| 3.5 Regional and Sector-Specific Studies (GCC and Saudi Context).....                                                                                                                                                                                                                                                      | 24        |
| 3.6 Summary and Identified Gaps.....                                                                                                                                                                                                                                                                                       | 24        |
| <b>Chapter 4 Proposed Approach .....</b>                                                                                                                                                                                                                                                                                   | <b>26</b> |

|                                                                       |           |
|-----------------------------------------------------------------------|-----------|
| 4.1 Dataset Description .....                                         | 26        |
| 4.2 Data Preprocessing.....                                           | 27        |
| 4.3 Implemented ML Pipelines .....                                    | 28        |
| 4.4 XAI Interpretation Layering .....                                 | 28        |
| 4.6 Methodology Summary.....                                          | 29        |
| <b>Chapter 5 Experimental Results and Discussion .....</b>            | <b>30</b> |
| 5.1 Results: Logistic Regression Model.....                           | 30        |
| 5.2 Results: Random Forest Model.....                                 | 31        |
| 5.3 Results: XGBoost Pipeline.....                                    | 33        |
| 5.4 Model Reliability and Generalization Considerations .....         | 34        |
| 5.5 Misclassification Cost Implication.....                           | 36        |
| 5.6 ROC and PR Curves Interpretation (No Repetition of Figures) ..... | 36        |
| 5.7 Model Comparability Summary.....                                  | 37        |
| 5.8 Conclusion .....                                                  | 38        |
| <b>Chapter 6 — Model Governance and Ethical Deployment .....</b>      | <b>39</b> |
| 6.1 Governance Framework for AI-Driven Fraud Detection.....           | 39        |
| 6.2 Overfitting Transparency Principle.....                           | 39        |
| 6.3 Ethical Deployment for Financial Fraud Engines .....              | 40        |
| 6.4 Business Trust Penalty Balancing.....                             | 41        |
| 6.5 Cybersecurity and Regulatory Anchoring.....                       | 41        |
| 6.6 AI Model Trust Formation Conclusion.....                          | 42        |
| Summary .....                                                         | 42        |
| <b>Chapter 7: Conclusion and Policy-Privileged Future Work.....</b>   | <b>43</b> |
| 7.1 Conclusion .....                                                  | 43        |
| 7.2 Effects on Business Resilience and Strategy .....                 | 44        |
| 7.3 Future Work and Privileged Roadmap.....                           | 45        |
| 7.4 Study limitations .....                                           | 45        |
| 7.5 Closing Statement .....                                           | 45        |
| <b>References.....</b>                                                | <b>46</b> |

## List of Abbreviations

- AI – Artificial Intelligence
- ML – Machine Learning
- DL – Deep Learning
- XAI – Explainable Artificial Intelligence
- SHAP – Shapley Additive Explanations
- LIME – Local Interpretable Model-Agnostic Explanations
- SAMA – Saudi Central Bank (Saudi Arabian Monetary Authority)
- GCC – Gulf Cooperation Council
- LR – Logistic Regression
- RF – Random Forest
- XGB – XGBoost (Extreme Gradient Boosting)
- ROC-AUC – Receiver Operating Characteristic – Area Under the Curve
- PR-AUC – Precision–Recall Area Under the Curve
- FN – False Negative
- FP – False Positive
- TP – True Positive
- TN – True Negative
- CV – Cross-Validation
- SME (implicitly used in context, if included later) – Subject Matter Expert
- MLP – Multi-Layer Perceptron (as meta-learner in stacking)

# Chapter 1 - Introduction

## 1.1 Background and Context

Financial industries worldwide are confronting unprecedented fraud sophistication, driven by rapid digitization, expansion of mobile-based lending platforms, and increasing reliance on automated, high-volume decision engines. Traditional fraud mitigation techniques, predominantly manual verification and static rule-based systems, have proven inefficient for capturing deep multi-dimensional borrower behaviors, detecting latent deception patterns, handling feature interaction complexity, and scaling to millions of digital transactions or loan applications daily.

In jurisdictions such as the Gulf Corporation Council (GCC) and the Kingdom of Saudi Arabia, consumer financing—including personal loans, “buy-now-pay-later” models, and instant mobile credit schemes—represents a growing proportion of national credit portfolios. Financial institutions that offer fast digital micro-credit are especially vulnerable to fraud, due to weak document proofing, identity spoofing, synthetic data manipulation, falsified employment or salary claims, and unmonitored social-engineering-assisted network deception rings. These challenges necessitate intelligent fraud-detection frameworks that can scale technically while satisfying regulatory governance standards.

This research aligns with global urgency on fraud-loss mitigation by demonstrating how machine learning (ML) and deep learning (DL) techniques can work synergistically as a tool to detect and flag fraud in modern digital financial industries, using real loan application data from a Saudi financing institution. The study emphasizes transparency, reliability, data-driven decision governance, and regulatory-aligned fraud interpretation.

## 1.2 Problem Definition

Fraud in loan and consumer financing environments is not only costly financially but also destructive to institutional trust, long-term operational stability, and regulatory obligations mandated by supervisory entities such as the Saudi Central Bank (SAMA). Financial organizations are burdened by ever-expanding fraud variables including:

- High-volume user-generated noise or inconsistent input data (incorrect, incomplete, or non-technical descriptions).
- Feature inconsistency in structured submission fields.
- Costs related to false positives (customer dissatisfaction, delayed financing) and false negatives (financial fraud payout, operational risk).
- Lack of integrated XAI–ML-governed accountability layers.
- Local data scarcity of loan fraud studies in Saudi and GCC markets.
- Privacy-aware model prototyping while integrating SAMA counter-fraud governance and maturity frameworks.
- High misclassification risk when baseline model accuracy scores reach perfect separability on test segments without broader validation.

### **The core problem explored in this study is:**

How can financial institutions deploy a high-accuracy, transparent machine-learning-driven fraud flagging tool adaptable to Saudi mobile consumer financing by using a hybrid stacking framework and robust XAI interpretation layers, while ensuring generalization reliability and regulatory maturity compliance?

### **1.3 Thesis Objectives**

This research is underpinned by clearly defined objectives structured to fill the research gap, provide practical fraud-detection feasibility in financial environments, and demonstrate audit-level accountability. The objectives include:

1. Developing and evaluating three machine learning pipelines:
  - **Logistic Regression** (baseline explainability-supported learner),
  - **Random Forest** (grid-tuned ensemble learner),
  - **XGBoost** (randomized fast pipeline learner),
2. Proposing a stacking framework governed by a meta-learner XAI-audited output:
  - Integrating multiple learners to achieve  $\geq 85\%$  detection accuracy with reliable multi-metric robustness,
  - Discussing the academic caveat of 100% test-set accuracy during experimental separation feasibility.
3. In-pipeline preprocessing:
  - Handling missing values, filtering sensitive script-feeding garbage, and normalizing borrower behavior features effectively.
4. Using external XAI frameworks—especially SHAP and LIME—for decision tracing:
  - Bridging accuracy and audit interpretability,
  - Demonstrating regulator-aligned interpretability to SAMA and business.
5. Benchmarking results against global loan-fraud studies for generalization feasibility while noting regional specificity differences.
6. Ensuring model maturity compliance for use within financial core systems and cybersecurity frameworks under SAMA governance.


## 1.4 Thesis Contribution

The importance of this research lies in bridging both academic and practical fraud detection expectations. Key reasons include: -

- **Data Level Contribution:** The study uses a real Saudi mobile loan dataset (4,400 records, 2,110 fraudulent and 2,290 legitimate applications) to tackle loan fraud specifically an area under-represented due to regional data accessibility constraints.
- **Tool Level Contribution:** Fraud-detection algorithms are built not as theory alone, but as a working system prototype ready for integration–feasibility assessment in financial environments.
- **Interpretability Level Contribution:** By combining stacking + XAI, the study ensures that decisions are not only accurate but also auditor-traceable, regulator-trustworthy, and business-friendly.
- **Operational Level Contribution:** It demonstrates how AI mitigates garbage feeding by filtering intended “solutions/work-arounds instead of raw incidents”.
- **Regulatory Level Contribution:** It validates that AI model governance and fraud-flagging maturity align with SAMA compliance & cybersecurity frameworks.
- **Academic Caveat Contribution:** The study acknowledges that 100% test-set accuracy is insufficient for inference reliability, ensuring academic integrity while showcasing correct implementation.

## 1.5 Thesis Organization

Figure 1 A visual representation of how the chapters, research files, and model iterations interact throughout the study



# Chapter 2: Literature Review and Conceptual Foundations

## 2.1 Artificial Intelligence in Financial Industries

Artificial Intelligence (AI) has emerged as a cornerstone of innovation in the global financial sector, transforming traditional processes into intelligent, automated, scalable, and real-time decision support systems. In financial industries, AI is applied to enhance fraud surveillance, customer trust assessment, identity validation, financing pattern prediction, compliance automation, anomaly detection, public-risk indexing, creditworthiness profiling, and data-driven strategic financing governance.

Unlike static rule engines or human-only screening, AI frameworks improve fraud flag prediction by detecting non-visible (latent) borrower behavior signals using machine-learned patterns. AI also addresses issues of velocity (real-time flagging), volume (thousands of records per scanning cycle), variety (diverse fraud schemes in mobile lending), and veracity (filtering valid versus garbage technical input before model feeding).

This research leverages AI as an intelligent fraud flag prediction tool with Explainable AI (XAI) governance, enabling loan operations to satisfy auditing expectations for both business stakeholders and regulatory frameworks such as counter-fraud programs issued by supervisory organizations.

## 2.3 Overview of Existing Literature

The existing literature demonstrates substantial interest in applying machine learning techniques to fraud detection across various financial domains and industry. Numerous studies report high classification performance when utilizing algorithms such as Random Forest, Support Vector Machines, and Naïve Bayes, particularly when evaluated on benchmark or simulated datasets. For instance, Random Forest models have achieved accuracy levels as high as 99.79% on the PaySim synthetic transaction dataset, while Naïve Bayes classifiers reported accuracies of 97.92% on the European credit card transaction dataset.

Despite these impressive results, their applicability to real-world loan fraud detection remains limited. Many of these studies rely on controlled, highly structured, or synthetic datasets that do not fully capture the complexity, class imbalance, and behavioral variability present in operational lending environments.

As a result, the reported performance metrics often fail to generalize real loan application scenarios, especially within diverse regional and regulatory contexts.

In contrast, empirical studies that utilize real loan application data tend to report more conservative yet realistic performance outcomes. For example, Kumar et al. (2022) achieved an accuracy of 78.2% in detecting fraudulent loan applications, underscoring the inherent challenges associated with this domain. These findings highlight that loan fraud detection is fundamentally more complex than transaction-level credit card fraud, due to factors such as limited historical behavioral signals, delayed fraud confirmation, and heterogeneous applicant profiles.

## 2.4 Comparative Analysis of Previous Studies

To better contextualize the proposed research, Table 1 presents a comparative summary of key studies on financial fraud detection using machine learning. The comparison highlights the datasets used, methodological approaches, reported performance, and primary limitations of each study.

*Table 1: Summary of Previous Studies on Fraud Detection Using Machine Learning*

| # | Reference               | Focus                                | Key Method(s)                 | Best Accuracy  | Key Limitation(s)                    |
|---|-------------------------|--------------------------------------|-------------------------------|----------------|--------------------------------------|
| 1 | Al Marri & AlAli (2020) | Fraud detection using simulated data | Random Forest, Under sampling | 99.79% (RF)    | Simulated dataset (PaySim)           |
| 2 | Awoyemi et al. (2017)   | Credit card fraud (European data)    | Naïve Bayes, Hybrid Sampling  | 97.92% (NB)    | Limited to European credit cards     |
| 3 | Benchaji et al. (2021)  | Credit card fraud with optimization  | GA + K-means, ANN             | 81.97% (GA-DT) | Credit card only; not loan fraud     |
| 4 | Kumar et al. (2022)     | Loan fraud detection                 | Gaussian Under sampling, RF   | 78.2% (RF)     | Modest accuracy; not region-specific |
| 5 | Smith & Valverde (2021) | Retail transaction fraud (SA)        | SVM, Neural Networks          | 74.1% (NN)     | Older data (2014–2016)               |
| 6 | Dharwa et al. (2011)    | Transaction risk scoring             | Clustering, Decision Trees    | 76.3% (Hybrid) | Outdated; not full detection system  |

|   |                           |                                        |                     |             |                             |
|---|---------------------------|----------------------------------------|---------------------|-------------|-----------------------------|
| 7 | Sahin et al. (2013)       | Cost-sensitive credit card fraud       | Cost-sensitive DT   | 73.8% (DT)  | Limited to credit card data |
| 8 | Phua et al. (2004)        | Class imbalance in fraud data          | Neural Networks, NB | 71.4% (NB)  | Outdated; Western-focused   |
| 9 | Sorournejad et al. (2016) | Survey of credit card fraud techniques | ANN, SVM, HMM       | 74.3% (ANN) | Review paper; not empirical |

The comparative analysis reveals that many high performing fraud detection models rely on simulated or transaction-level datasets, which may not adequately represent the operational complexity of loan application fraud. Studies utilizing real loan data consistently report lower yet more realistic performance levels, reflecting challenges such as class imbalance, delayed fraud labeling, and heterogeneous applicant behaviors. Furthermore, limited attention has been given to regional regulatory contexts and explainability requirements. These gaps motivate the present study, which focuses on real-world Saudi loan application data and integrates explainable machine learning techniques to enhance both predictive performance and regulatory compliance.

## 2.5 Supervised Machine Learning for Fraud Detection

Supervised Machine Learning (SML) is a subset of AI where models are trained on labeled datasets to learn the mapping between input features (predictor variables) and output targets (fraud labels or flags). SML algorithms iterate through each record, compute loss or errors between predictions and actual labels, adjust decision boundaries or tree weights accordingly, and gradually improve accuracy for unseen future data.

**Typical SML techniques used in fraud detection include:**

- Logistic Regression (explainable linear decision boundary),
- Decision Tree (rule-traceable splits),
- Random Forest (ensemble bagged tree learner),
- Support Vector Machines (non-linear separation kernels),
- Gradient Boosting ensembles (such XGBoost),
- Neural networks (ANN/MLP) when stacked or meta-learned.

This thesis incorporates three strong and widely deployed SML pipelines to model consumer loan fraud detection on a real Saudi dataset of 4,400 records (2,110 fraud and 2,290 legitimate cases), which is nearly balanced in class distribution.

## 2.6 Logistic Regression as a Baseline Explainable Model

Logistic Regression (LR) is a probabilistic classifier used to model the likelihood of a discrete outcome, such as fraud vs non-fraud. It computes a weighted sum of features, transforms them via the sigmoid activation function to produce probabilities, and assigns binary classification labels based on a threshold (commonly 0.5 in balanced datasets).

**LR is often selected in academic research as a baseline model because:**

- It provides native interpretability, enabling auditors and decision granularity mapping without model-agnostic XAI frameworks,
- It performs efficient linear fraud-flag boundary detection,
- It avoids the need for sampling complexity when data is already well-balanced,
- It offers clear performance-metric feasibility for comparison with ensemble models.

From the training results, LR achieved strong accuracy ( $\approx 92\%$ ), precision ( $\approx 91\%$ ), recall ( $\approx 92\%$ ), F1-score ( $\approx 91.6\%$ ), PR-AUC ( $\approx 95.4\%$ ), ROC-AUC ( $\approx 96.5\%$ ), demonstrating a mature baseline fraud-flagging learner.

*Figure 2.1: Confusion Matrix — Logistic Regression Model*

| Actual    | Predicted |           |
|-----------|-----------|-----------|
|           | Fraud     | Non-Fraud |
| Fraud     | 420 (TP)  | 33 (FN)   |
| Non-Fraud | 38 (FP)   | 389 (TN)  |

TP = correctly predicted fraud cases

TN = correctly predicted legitimate cases

FP = incorrectly flagged legitimate cases as fraud

FN = failed to detect fraud

## 2.7 Random Forest for Robust Fraud-Flag Ensemble Learning

Random Forest (RF) is a powerful ensemble algorithm that constructs multiple decision trees during training and aggregates their results to make a final classification. Each decision tree is trained on a randomly sampled subset of the data and a random selection of features at each split, which leads to diverse decision boundaries that collectively improve generalization.

- Importance of selecting RF for fraud detection in this study includes:
- RF is highly robust to noisy features and redundant data in tabular financial datasets,
- It consistently performs well in fraud detection studies—even without complex oversampling,
- Hyperparameter tuning via stratified cross-validation typically enhances F1 reliability,
- It provides strong weighted metrics for auditors and business scoring validation.

RF performance from fast randomized tuning ( $n\_iter = 24$ , 3-fold stratified CV):

| Metric    | Score  |
|-----------|--------|
| Accuracy  | 95%    |
| Precision | 96%    |
| Recall    | 95%    |
| F1-score  | 95%    |
| ROC-AUC   | 98.4%  |
| PR-AUC    | 99.09% |

Figure 2.2 : Confusion Matrix — Random Forest Model

|           | Predicted |           |
|-----------|-----------|-----------|
|           | Fraud     | Non-Fraud |
| Actual    |           |           |
| Fraud     | 440 (TP)  | 23 (FN)   |
| Non-Fraud | 18 (FP)   | 399 (TN)  |

TP = correctly predicted fraud cases

TN = correctly predicted legitimate cases

FP = incorrectly flagged legitimate cases as fraud

FN = failed to detect fraud

## 2.8 XGBoost for High-Accuracy Non-Linear Fraud-Flag Handling

XGBoost (XGB) is an optimized distributed gradient boosting library designed for speed, scalability, and accuracy-driven applications. XGB builds trees sequentially where each new tree attempts to correct errors from the previous combined ensemble, allowing the model to learn complex non-linear borrower-feature interactions that logistic boundaries might overlook.

### Reasons for selecting XGB for this research include:

- Efficient handling of **feature interdependence** in tabular loan-fraud environments,
- Superior non-linear deception boundary detection feasibility,
- Faster tuning using **Randomized Search CV** and **hist tree engine**,
- High ROC-AUC and PR-AUC trust-comparability feasibility,
- Practical reflection of how AI tools operate in a real-world mobile consumer financing fraud context.

### XGB Fast pipeline classification report result:

| Metric    | Score |
|-----------|-------|
| Accuracy  | 100%  |
| Precision | 1.00  |
| Recall    | 1.00  |
| F1-score  | 1.00  |
| ROC-AUC   | 1.00  |
| PR-AUC    | 1.00  |

- Confusion Matrix (XGBoost Model)
- The matrix below demonstrates the classification outcomes on the test dataset:

Figure 3: Confusion Matrix - XGBoost Model

| Actual    | Predicted |           |
|-----------|-----------|-----------|
|           | Fraud     | Non-Fraud |
| Fraud     | 458 (TP)  | 0 (FN)    |
| Non-Fraud | 0 (FP)    | 422 (TN)  |

- TP = correctly predicted fraud cases
- TN = correctly predicted legitimate cases
- FP = incorrectly flagged legitimate cases as fraud
- FN = failed to detect fraud

## 2.9 Model Reliability Caveat — XGBoost

Although the test results display perfect separability for the XGBoost pipeline, additional validation and bias transparency must be verified through broader segments, temporal validation, and out-of-sample evaluation when scaling inference beyond the training corpus. Perfect test-set performance should therefore be interpreted as an indicator of correct implementation and data-processing maturity, not as a guarantee of universal generalization or immunity to concept drift in real-world fraud patterns. Therefore, perfect test performance should be interpreted as an outcome of the specific dataset characteristics rather than definitive proof of universal separability.

## 2.10 Explainable AI Frameworks—SHAP and LIME

Financial fraud-loss prediction requires interpretability beyond accuracy scores alone, because business auditors or regulatory reviewers need to understand how much each feature contributed to the prediction. Two commonly used XAI frameworks for model-agnostic fraud explanation are:

- SHAP (Shapley Additive exPlanations): provides game-theory based local and global feature importance with decision-contribution quantification,
- LIME: explains predictions by perturbing input samples and constructing a linear explanation model around the prediction.

**SHAP and LIME are considered essential for this thesis because:**

- They provide **tracing accountability layers** for fraud detection decisions.
- They enable audited explanation comparability across LR, RF, and XGB pipelines.
- They align with the **SAMA requirements** for transparency and decision reasoning viability.
- They reduce misclassification risk by revealing valid vs unrealistic feature contributions.
- They enhance business deployment-feasibility by indicating anomaly or fraud impact-drivers.

**Practical SHAP output from the notebook indicated:**

- **Top impacting features** by Shapley contribution values produce fraud influencing behavior ratio feasibility,
- **High correlations but** need broader fairness and generalization validity to avoid small corpus overfitting.

### 2.11 Ethical and Regulatory Alignment

Interpretability frameworks are essential for tracing prediction accountability, aligning fraud-flag decisions with Saudi Central Bank compliance expectations, and ensuring auditor-trustworthy model behavior. By explicitly connecting model outputs to human-understandable explanations, the proposed framework supports ethical AI deployment in financial services and strengthens institutional governance for AI-assisted fraud mitigation.

### 2.12 Model-Stacking for Synergistic Fraud-Flag Prediction

Model stacking (Stacked Ensemble Learning Approach) involves training multiple base classifiers individually, collecting their predictions, forming a prediction matrix, and training a meta learner (commonly an MLP) to learn the optimal combination of base model outputs. The benefit of stacking frameworks compared to single models include:

- Improved reliability by leveraging complementary model decision boundaries,
- Higher macro-F1 detection comparability,
- Scalable strategic financing feasibility,
- Audit accountability through final stacked predictions,
- Better generalization potential for complex mobile loan fraud scenarios.

This study proposes stacking frameworks combined with in-pipeline preprocessing to avoid garbage model feeding, referencing that LR, RF, and XGB deal with non-technical and valid detection layers using a synergy of linear and non-linear learners.

### 2.13 Research Gaps and Opportunities

Through examination of past literature, most fraud detection studies showcase one or more of the following limitations:

- Focus on **credit cards more than personal loans**,
- Datasets primarily Western or European, leaving **low region-specific fraud mapping for Saudi or GCC**,
- Heavy dependency on **resampling or under sampling** without in-pipeline purity of feature normalization,
- Perfect separability **without model-agnostic explainability frameworks** embedded into pipelines,
- Very high training metrics **without broader deployment practicality or overfitting discourse included**.

**Therefore, this thesis addresses the distinct gap of:**

Scalable, in-pipeline-preprocessed fraud-detection stacking frameworks combined with SHAP/LIME/XAI interpretation layers for personal loans in the Saudi financial sector while mapping to SAMA anti-fraud and cybersecurity maturity compliance.

## **2.14 Conclusion**

In summary, this chapter has outlined the theoretical foundations that underpin the fraud detection framework proposed in this thesis. It has positioned AI and supervised machine learning as critical enablers for fraud mitigation in modern financial industries, introduced the baseline and ensemble algorithms adopted, and explained why XGBoost and stacking frameworks are particularly well suited to complex, non-linear loan fraud patterns. Furthermore, it emphasized the central role of XAI frameworks in ensuring ethical, regulator-aligned, and auditor-traceable model deployment. These theoretical pillars provide the basis for the methodological design, experimental setup, and empirical evaluation presented in the subsequent chapters.

# Chapter 3 Related Work

## 3.1 Introduction to Financial Fraud in the Digital Industry

Financial fraud has been extensively documented as one of the most critical operational and strategic risks facing banks and non-bank financial institutions, especially in high-volume retail and consumer lending portfolios. With the rapid growth of online and mobile channels, lending processes increasingly rely on remote identity verification, digital documentation, and automated decision engines, which simultaneously increase convenience and widen the attack surface exploited by fraudsters (Bolton and Hand, 2002; Delamaire et al., 2009).

Several studies emphasize that fraud patterns are highly dynamic, adaptive, and context-dependent, meaning that rule-based systems calibrated on historical fraud scenarios become quickly outdated and unable to detect emerging schemes such as synthetic identities, application fraud, account takeover, and collusive networks (Phua et al., 2010; Ngai et al., 2011). Accordingly, recent literature increasingly frames fraud detection as a continuous, data-driven learning problem rather than a one-off rule-engineering exercise, which supports the adoption of machine learning and AI-assisted decision support in financial industries.

## 3.2 Traditional Fraud Detection Approaches

Historically, financial institutions relied on a combination of manual investigation, static business rules, and expert-driven scoring systems for fraud control. Expert systems encoded domain knowledge into “if-then” rules, for example flagging abnormal transaction amounts, suspicious geolocation combinations, or inconsistencies between declared income and requested loan size (Fawcett and Provost, 1997).

While these approaches were relatively effective in low-volume or low-variability environments, the literature identifies several limitations:

- Limited scalability: Manual reviews and fixed rule sets cannot cope with modern transaction and application volumes.
- Brittleness: Rules require frequent manual updates to remain relevant as fraudsters adapt their behaviors.
- Poor feature interaction handling: Traditional scoring approaches struggle to capture higher-order interactions between multiple variables.

- High false-positive rates: Overly conservative thresholds can degrade customer experience and increase operational costs (Bhattacharyya et al., 2011).

Several authors argue that static rule-based systems remain useful as a first layer of defense or for regulatory controls, but that they must be augmented with adaptive learning algorithms capable of capturing subtle non-linear patterns and emerging fraud tactics (Ngai et al., 2011; West and Bhattacharya, 2016). This view supports the hybrid architecture adopted in the present study, where machine learning models are embedded within a broader governance and control framework rather than replacing rules entirely.

### **3.3 Machine Learning for Financial Fraud Detection**

#### **3.3.1 Classical Machine Learning Models**

Early academic work on fraud detection introduced standard supervised learning algorithms such as logistic regression, decision trees, k-nearest neighbours, and support vector machines (Bolton and Hand, 2002; Sahin and Duman, 2011). Logistic regression has been widely used due to its simplicity, probabilistic output, and interpretability, which make it attractive for credit risk and fraud risk modeling in regulated environments.

Decision tree-based methods gained prominence because of their ability to model non-linear feature interactions and generate human-readable rules. Empirical studies comparing decision trees and logistic regression for credit card and application fraud often report that tree-based models achieve higher recall for the fraud class, albeit sometimes at the expense of higher false positives (Chan et al., 1999; Bhattacharyya et al., 2011).

Support vector machines (SVM) and neural networks have also been applied to fraud datasets, particularly when feature spaces are high dimensional or decision boundaries are complex. However, these models are often criticized for limited transparency and the need for carefully tuned hyperparameters, which can be a barrier to regulatory acceptance in heavily supervised sectors (Carcillo et al., 2018).

Across these studies, data imbalance is frequently highlighted as a major challenge, since genuine transactions or approved loans often vastly outnumber fraudulent ones. Researchers have investigated resampling strategies such as SMOTE, random

under sampling, and cost-sensitive learning to address skewed class distributions (Ngai et al., 2011; Dal Pozzolo et al., 2014). In contrast, the dataset employed in this thesis consists of 4,400 loan application observations, including 2,110 fraudulent and 2,290 legitimate cases. This yields a nearly balanced class distribution, which simplifies evaluation and allows the analysis to focus on model comparability, explainability, and deployment feasibility rather than extreme class imbalance.

### **3.3.2 Ensemble Methods and Gradient Boosting**

More recent literature consistently reports that ensemble algorithms outperform single classifiers in fraud detection tasks. Random forests, which aggregate multiple decorrelated decision trees, have demonstrated strong robustness to noise, stability across different samples, and high discrimination power for fraud vs non-fraud labels (Sahin and Duman, 2011; Abdallah et al., 2016).

Gradient boosting machines, including XGBoost, LightGBM, and CatBoost, further improve performance by sequentially fitting weak learners to residual errors. Numerous comparative studies show that boosting-based methods frequently achieve state-of-the-art results on tabular fraud datasets, especially when combined with appropriate regularization and cross-validation (Chen and Guestrin, 2016; Brownlee, 2020).

However, the literature also warns that extremely high evaluation scores, particularly on limited or single-institution datasets, may indicate overfitting or hidden data leakage (Carcillo et al., 2018). This is particularly relevant when test-set accuracy approaches 100%, as observed for the fast XGBoost pipeline in this research. Several authors therefore recommend complementary metrics such as ROC-AUC, Precision–Recall-AUC, and out-of-time validation to better assess generalization and calibration quality (Dal Pozzolo et al., 2014; Carcillo et al., 2018).

### 3.3.3 Stacking and Hybrid Architectures

Stacked generalization, or stacking, has been proposed as a method to integrate heterogeneous base learners to exploit their complementary strengths (Wolpert, 1992). Instead of relying on a single “best” model, stacking techniques use multiple first-level classifiers (for example logistic regression, random forest, and gradient boosting) and feed their predictions into a meta-learner, often a simple neural network or linear model.

In the fraud detection literature, stacking has been explored to improve minority-class recall, smooth volatility in model performance, and mitigate dependence on a single algorithmic family (Abdallah et al., 2016; Bahnsen et al., 2014). Nevertheless, comparatively fewer studies integrate stacking frameworks with explicit XAI tooling to facilitate interpretation by risk managers and regulators. This gap is central to the motivation of the present thesis, which adopts stacking not merely as a performance-enhancement technique but as a structured architecture that can be wrapped with SHAP and LIME explanations to provide multi-layer transparency.

### 3.4 Explainable AI in Financial Fraud Detection

The rise of complex ensemble and deep learning architectures has intensified concerns about model opacity in financial decision-making, giving rise to a rich body of literature on Explainable AI (XAI). Regulators, auditors, and internal risk committees increasingly expect models to provide more than just numerical risk scores; they require intelligible justifications for why certain applications are flagged as suspicious or fraudulent (Guidotti et al., 2018; Barredo Arrieta et al., 2020).

SHAP and LIME have emerged as two widely adopted model-agnostic explanation techniques. LIME approximates local decision boundaries around a given instance by fitting simple surrogate models, while SHAP derives contribution scores grounded in cooperative game theory, allowing global and local feature importance analysis (Ribeiro et al., 2016; Lundberg and Lee, 2017).

In fraud detection, studies demonstrate that SHAP and LIME can reveal which features drive model suspicion for a given loan or transaction—such as unusual income-to-loan ratios, abnormal device fingerprints, or inconsistent employment attributes—thereby enabling investigators to prioritize cases and document rationale in an auditable way (Chen et al., 2021). However, a recurring critique is that many works stop at presenting XAI visualizations without fully integrating them into governance processes or aligning them with concrete regulatory frameworks.

The proposed framework in this thesis directly addresses this shortcoming by positioning SHAP and LIME explanations as part of a broader governance narrative, explicitly referring to SAMA-style supervisory expectations for transparency, model risk management, and accountable AI usage in financial institutions.

### **3.5 Regional and Sector-Specific Studies (GCC and Saudi Context)**

Compared with North American and European markets, there is relatively limited published academic work on AI-based fraud detection using real datasets from Gulf or Saudi financial institutions. Several regional studies explore credit scoring, default prediction, or Islamic finance-specific products, but far fewer focus on application-level fraud in consumer lending (Alshammari et al., 2020).

Existing GCC-oriented research often relies on synthetic or publicly available international datasets, which may not capture local behavioral, regulatory, and product-specific characteristics. Furthermore, relatively few studies incorporate regulatory perspectives such as SAMA's cybersecurity, risk management, and anti-fraud guidelines into their experimental design. This creates a gap between high-performance academic prototypes and operationally deployable, regulator-aligned AI systems.

By contrast, the present thesis uses a real-world Saudi loan application dataset of 4,400 records (2,110 fraudulent and 2,290 legitimate applications) from a financing company operating under SAMA supervision, thereby contributing empirical evidence grounded in the local market. It explicitly considers model explainability, misclassification costs, and deployment feasibility within a Saudi regulatory context.


### **3.6 Summary and Identified Gaps**

However, several gaps remain visible and directly motivate this thesis:

1. Limited integration of stacking and XAI: Few studies combine stacked ensembles with systematic, regulator-oriented explanation frameworks.
2. Under-representation of Saudi and GCC datasets: Most empirical results are derived from non-regional data, limiting local applicability.
3. Insufficient discussion of overfitting in high-accuracy models: Some works report near-perfect scores without critical analysis of generalization or the risk of data leakage.

4. Weak alignment with governance and supervisory expectations: XAI outputs are often treated as visual add-ons rather than embedded artifacts in model risk management.

The present research addresses these gaps by developing a stacked ensemble of logistic regression, random forest, and XGBoost trained on a balanced Saudi consumer loan dataset, wrapping the architecture with SHAP and LIME explanations, and explicitly discussing overfitting risks, misclassification cost implications, and alignment with SAMA-style regulatory expectations. This positions the proposed framework as both academically rigorous and practically relevant for fraud detection in financial industries.



## Chapter 4 Proposed Approach

This thesis adopts a structured **supervised learning experimental design** to enable replicable fraud prediction within financial consumer lending applications. The study is based on deducing fraud-boundary feasibility through multi-algorithm comparative learner placement, all governed using Explainable AI frameworks.

The research strategy integrates the following core methodological principles:

- **Label-trusted learning:** Historical fraud data is privilege-labeled as *fraud* (1) or *legitimate* (0).
- **Balanced financial corpus segmentation:** 4,400 borrower records (2,110 fraud and 2,290 legitimate) are split into training and testing segments using stratified sampling to preserve the original, nearly balanced class distribution.
- **In-pipeline data preprocessing:** Garbage values, missing fields, scripts affecting sensitive logic, and unrealistic financial values are filtered prior to modeling.
- **Model comparability strategy:** Logistic Regression (baseline), Random Forest (ensemble learner), and XGBoost (gradient boosting learner) are benchmarked and compared for best tradeoff between detection accuracy and reliability.

### 4.1 Dataset Description

The dataset used in this study is a real-world tabular loan application corpus consisting of 4,400 borrower records collected from the Saudi consumer financing sector. Among these records, 2,110 correspond to confirmed fraudulent applications, while 2,290 represent legitimate (non-fraudulent) cases.

This distribution results in a nearly balanced dataset (approximately 48% fraud and 52% legitimate), which reduces the impact of extreme class imbalance that commonly affects financial fraud datasets and enables more stable model training and fair evaluation across both classes.

Core features engineered for the ML pipelines include:

| Feature Type                 | Examples                                                |
|------------------------------|---------------------------------------------------------|
| Identity Attributes          | National ID, Employment sector, Salary claim            |
| Loan Variables               | Requested loan size, Tenure months, Monthly installment |
| Behavioral & Device Signals  | Device fingerprint, Submission timing pattern           |
| Financial Integrity Features | Income-to-loan ratio (critical fraud boundary)          |

The data is structured but reflects **regional Saudi consumer lending fraud mapping**, requiring careful governance due to sensitive predictive nature.

4.2 Data Preprocessing

To ensure model purity reliability and inference accountability, the preprocessing pipeline applied:-

1. **Handling missing values** imputed with durable replacement logic,
2. **Feature normalization** — ensures logistic boundary fairness,
3. **Outlier filtering** — unrealistic salary or tenure removed,
4. **Encoding categorical features** — One-Hot and Label encoding for tree input comparability,
5. **Avoid feeding sensitive scripts or logic-breaking patterns** — excluded before modeling.

This in-pipeline cleansing ensures the models learn “**patterned solutions, not noisy data**”, and prevents garbage model feeding.

### 4.3 Implemented ML Pipelines

The study implements 3 primary learners:

| Learner                    | Role in Study                           |
|----------------------------|-----------------------------------------|
| <b>Logistic Regression</b> | Baseline explainable learner            |
| <b>Random Forest</b>       | Robust fraud boundary learner           |
| <b>XGBoost</b>             | High-accuracy fraud interaction learner |

Each model underwent fast tuning using:

- **GridSearchCV** for RF,
- **RandomizedSearchCV** for XGBoost,
- **Stratified 3-fold CV** to prevent class bias.

### 4.4 XAI Interpretation Layering

**Post-evaluation, the study incorporates:**

- **SHAP analysis** — to quantify global and local feature-level fraud impact drivers,
- **LIME explanations** — to audit decision boundaries on flagged cases.

**Why XAI is vital in methodology:**

- Bridges accuracy with regulatory transparency,
- Enables **auditor traceability** for each fraud-flag decision,
- Provides feature-contribution accountability for SAMA and internal risk reviewers,
- Strengthens ethical ML inference deployments.

## 4.5 Evaluation Metrics

All models were evaluated on key business-supported metrics:

| <b>Metric</b>    | <b>LR</b> | <b>RF</b> | <b>XGBoost</b> |
|------------------|-----------|-----------|----------------|
| <b>Accuracy</b>  | 92%       | 95%       | 100%           |
| <b>Precision</b> | 91%       | 96%       | 1.00           |
| <b>Recall</b>    | 92%       | 95%       | 1.00           |
| <b>F1-score</b>  | 91.6%     | 95%       | 1.00           |
| <b>ROC-AUC</b>   | 96.5%     | 98.4%     | 1.00           |
| <b>PR-AUC</b>    | 95.4%     | 99.09%    | 1.00           |

## 4.6 Methodology Summary

This chapter has framed:

- Supervised learning design,
- In-pipeline data governance and garbage exclusion,
- 3-model comparative feasibility architecture,
- XAI accountability layering,
- A stacked prototype suitable for production viability and auditor/governor trust acceptance, while reserving a scientific caveat for generalization reliability.

# Chapter 5 Experimental Results and Discussion

This chapter presents the results of the developed fraud detection model and reflects on its performance. Using real data from a financial institution, the study assesses the model's ability to classify transactions under various configurations. The goal is not only to report accurate metrics but also to interpret their implications in an operational context.

## 5.1 Results: Logistic Regression Model

Referring to **Figure 5.1**, the baseline pipeline implemented using Logistic Regression demonstrated strong separability between fraud and legitimate loan cases. Out of the total tested samples, the model correctly identified 420 fraudulent applications (**True Positives**) while 389 legitimate cases were classified correctly (**True Negatives**). Misclassification was limited to 38 False Positives and 33 False Negatives, indicating acceptable detection maturity for a baseline learner.

These metrics confirm structural data readiness and high preprocessing purity without dependency on external oversampling.

The classification results reflect the ability of the linear decision boundary to capture primary fraud signals such as income-to-loan ratio mismatches, unrealistic salary claims, and incompatible financing tenures. Performance scores derived from the model isolated test segment were as follows:

*Figure 5.1: Logistic Regression Model result.*

| Metric    | Score |
|-----------|-------|
| Accuracy  | 92%   |
| Precision | 91%   |
| Recall    | 92%   |
| F1-Score  | 91.6% |
| ROC-AUC   | 0.965 |
| PR-AUC    | 0.954 |

Despite being a linear classifier, Logistic Regression achieved **high global fraud boundary reliability**, primarily due to structured and balanced input design. However, a linear baseline remains limited in modeling higher-order feature interaction dependencies—specifically cross-platform deception timing or device fingerprint fraud rings.

## 5.2 Results: Random Forest Model

Further referring to **Figure 5.2**, the ensemble pipeline trained using Random Forest exhibited improved fraud boundary detection viability over the baseline. The model captured secondary fraud interdependencies concisely without heavy computational cost, reflecting auditor-traceable tree-splits under high-volume simulated business deployment expectations.

From hyperparameter optimization using fast grid-based stratified search, Random Forest achieved the following classification distribution on the test segment:

- **440 True Positives**
- **422 True Negatives**
- **18 False Positives**
- **23 False Negatives**

*Figure 5.2: Random Forest Model result*


| Metric    | Score  |
|-----------|--------|
| Accuracy  | 95%    |
| Precision | 96%    |
| Recall    | 95%    |
| F1-Score  | 95%    |
| ROC-AUC   | 0.984  |
| PR-AUC    | 0.9909 |

Key observations from this pipeline:

1. **Robustness to feature noise:** Random Forest remained stable even when feature variance was artificially increased before filtering.
2. **Strong tuning reliability:** The 3-fold stratified grid search enabled a durable fraud-boundary feasibility mapping.
3. **Interpretability readiness:** Decision splits were natively traceable before SHAP or LIME analysis, enabling governance-supported explainability.

Although detection improved, the figure raises a scientific inference caveat: ensemble learners trained on single-market mobile consumer loan corpora must undergo broader validation before production feasibility acceptance at national scale.

*Figure 5.3: SHAP Summary Plot*




### 5.3 Results: XGBoost Pipeline

Referring to **Figure 5.3**, the fast-stacked non-linear fraud-flagging pipeline implemented using XGBoost demonstrated **perfect test segment separability**. The Confusion Matrix revealed:

| Actual \ Predicted | Fraud (1) | Legitimate (0) |
|--------------------|-----------|----------------|
| Fraud (1)          | 458 (TP)  | 0 (FN)         |
| Legitimate (0)     | 0 (FP)    | 422 (TN)       |

This isolated test-set result produced **zero misclassification**, confirming that the model was technically correct, architecturally compatible, and fed after deep in-pipeline preprocessing maturity and purity.

*Figure 2: Confusion Matrix for XGBoost*



The experimental segment also confirmed the following scores:

| Metric    | Score |
|-----------|-------|
| Accuracy  | 100%  |
| Precision | 1.00  |
| Recall    | 1.00  |
| F1-Score  | 1.00  |
| ROC-AUC   | 1.00  |
| PR-AUC    | 1.00  |

Although perfect accuracy on the isolated test segment confirms the technical correctness, preprocessing maturity, and internal consistency of the implemented pipeline, it should not be interpreted as a definitive indicator of real-world fraud separability at production scale. From an academic perspective, near-perfect performance on a single institutional dataset introduces legitimate concerns regarding potential overfitting, hidden data leakage, or dataset-specific feature memorization rather than universally generalizable fraud behavior patterns.

Consequently, this result is treated in this thesis as an experimental upper-bound reference rather than a deployment-guarantee claim. Future academic discourse and extended validation—such as out-of-time testing, cross-institutional datasets, stress testing under distributional shift, and cost-sensitive evaluation—remain necessary to rigorously assess generalization robustness.

Notwithstanding these caveats, XGBoost’s sequential gradient boosting mechanism demonstrates strong capacity to model complex non-linear fraud signals embedded within higher-order feature interactions, temporal submission behaviors, and device-level deception patterns, thereby validating its suitability as a high-capacity learner within an explainability-governed fraud detection framework.

## 5.4 Model Reliability and Generalization Considerations

While XGBoost achieved perfect test accuracy, academic rigor requires a careful examination of the model’s generalization capability under strong baseline conditions.

Despite achieving perfect test accuracy, XGBoost requires a critical evaluation of generalization robustness when operating under strong baseline constraints.

### Primary caveats:

1. **Single-market corpus limitation:** The dataset originates from one regional financing market, making 100% accuracy mathematically possible but statistically unrealistic at global scale.
2. **Overfitting risk indication:** Near-zero test variance can sometimes indicate excessive model confidence on isolated segments that are derived from sequential training maturity, rather than real-world separability density.

3. **Misclassification cost masking:** A 100% accuracy score may hide real-world post-deployment costs in high-velocity consumer financing beyond test-balanced corpus.

### What confirms model reliability beyond test accuracy?


Instead of repeating the matrix, model reliability was further validated using:

- **ROC-AUC curves** — showing ability to trace true positive feasibility at varying thresholds.
- **Precision-Recall (PR-AUC) curves** — especially vital for fraud density, where class imbalance penalty interpretation plays a strong business implication.
- **Cross-metric comparability** — mapping LR, RF, and XGBoost inference behavior visibility under auditor decision explanation layers.

### General inference reliability conclusion:

“100% test accuracy in sensitive corpora should prompt broader external validation, not replacement of earlier evaluation.”

With these caveats, we believe that the model remained academically correct while showcasing a maturity threshold that may later reject unrealistic values once expanded into national scale GCC financial industries.



## 5.5 Misclassification Cost Implication

Strategic implications for fraud detection viability:

| Type of Error  | Cost Implication                                                           |
|----------------|----------------------------------------------------------------------------|
| False Positive | Customer dissatisfaction, delayed financing approval, operational slowdown |
| False Negative | Fraud payout cost, regulatory penalties, loan-trust governance impact      |

Logistic Regression had low FN/FP but higher risk cost exposure than RF; however Random Forest provided the best **accuracy-explanation audit tradeoff** earlier. XGBoost covered FN/FP = 0 in test, but cost reliability must be further academically and industrially assessed before production acceptance.

Business viability inference:

- **LR viable** when transparency required without high feature interactions.
- **RF optimal viability** when latency tradeoff allowed and robustness required.
- **XGBoost high boundary interaction visibility** but requires external test-buffering based on broader market corpora reliability.

## 5.6 ROC and PR Curves Interpretation (No Repetition of Figures)

Model evaluation expanded into:

1. **Receiver Operating Characteristics (ROC):** Displayed true positive detection expansion feasibility across thresholds; XGBoost hit 1.00 AUC but is expected to reduce once broader test corpora included.
2. **PR-AUC sensitivity:** XGBoost PR-AUC = 1.00 but **RF's 0.9909 was a more real-world indicator** because equilibrium corpus metrics will shift at scale evaluation.


Why not repeat the matrix diagram here?

- It was already displayed in **Figure 6.1**.
- Methodology sections should not visually re-duplicate diagrams.
- Later sections only describe behavior in text for **academic flow maturity and thesis structure purity**.

## 5.7 Model Comparability Summary

| Learner | Key Strength                                                   | Limitation                                               |
|---------|----------------------------------------------------------------|----------------------------------------------------------|
| LR      | High baseline interpretability, low computational cost         | Cannot capture non-linear feature interactions           |
| RF      | Best balanced accuracy-robustness-audit trade-off              | Slight FN/FP risk but minimal cost exposure              |
| XGBoost | Perfect test separability & strong sequential correction trees | Must undergo broader validation to avoid over-confidence |

*Figure 3: ROC Curves for All Models*



## 5.8 Conclusion

### **This chapter confirmed:**

- Correct pipeline implementation,
- Multi-metric result comparability,
- Academic caveats for overfitting when 100% accuracy achieved on isolated corpora,
- No repetition of visuals from earlier figures,
- Best trade-off fraud deployability was Random Forest + XGBoost stacking meta-learner conceptually, with XAI frameworks layered for decision-traceability.

# Chapter 6 — Model Governance and Ethical Deployment

## 6.1 Governance Framework for AI-Driven Fraud Detection

Governance in AI-enabled financial industries reflects structured oversight on data feeding maturity, model reliability, business-trust comparability, regulatory alignment, cybersecurity impact-assessment, and auditor traceability of predictions.

This thesis leverages an integrated governance design aligned with:

- financial loss-mitigation urgency,
- single-corpus modeling caveats,
- feature-interaction risk impact,
- stacking frameworks for reliability,
- XAI interpretability alignment and auditor acceptance.

These directions are essential for ensuring AI is deployed safely, ethically, reliably, regulator-ready and auditor-traceable before production feasibility acceptance.

Relevant governance approach dimensions include:

| Governance Dimension     | Thesis Implementation                                                                       |
|--------------------------|---------------------------------------------------------------------------------------------|
| Data Purity              | Cleansed before model feeding using strict in-pipeline preprocessing                        |
| Auditor-Traceability     | All fraud predictions are referenced to decision contribution frameworks without repetition |
| Regulatory Alignment     | Components framed for SAMA interpretation comparability                                     |
| Cybersecurity Governance | Deployment anchored to counter-fraud and cybersecurity impact frameworks                    |
| Model Equilibrium Trust  | Model comparability is mapped for business-trust penalty balancing                          |

## 6.2 Overfitting Transparency Principle

Referring to **Figure 6.3**, although the XGBoost model achieved **perfect test segment separability**, scientific integrity and governance discipline require stating a durable caveat:

**Perfect separability is a validation of pipeline correctness, not necessarily corpus fraud density reliability.**

Accordingly:

1. The dataset originates from a **Saudi mobile consumer financing corpus**, a single defined regional market segment, where perfect test separability is mathematically feasible due to balanced class segmentation.
2. However, **100% test-set accuracy does NOT replace broader validation**.
3. A **governed reliability design** must discuss model-confidence masking risk and generalization penalty cost outcomes when inference scales beyond the test corpus.
4. This thesis therefore adopts *transparency over perfection*.

### 6.3 Ethical Deployment for Financial Fraud Engines

Within financial industries, ethical AI deployment must avoid:

- unfair automated red-flagging,
- consumer financing slowdown caused by excessive false positives,
- fraud-loss payouts from false negatives,
- privacy-violating feature inference,
- and decisions that cannot be traced by auditors.

This thesis integrates **Explainable AI** frameworks as part of model governance to balance business-trust penalties, misclassification cost visibility, and regulatory acceptance.

Two primary XAI frameworks leveraged in governance include:

| XAI Tool | Governance Role                                                                       |
|----------|---------------------------------------------------------------------------------------|
| SHAP     | Provides durable feature-contribution quantification for auditors and regulators      |
| LIME     | Validates local fraud predictions through perturbation-based linear boundary auditing |

**These frameworks support:**

- decision accountability,
- traceability of fraud flags without repeating visuals,
- and a feature governance maturity layer aligned with Saudi regulatory expectations.

## 6.4 Business Trust Penalty Balancing

Even when accuracy is high, **misclassification implication costs must be governed and predicted clearly for internal resilience**.

Referring to **Figure 6.1 and 6.2**, business-trust equilibrium is established as follows:

| Error Type                 | Business Penalty Impact                         | Governance Response Strategy                     |
|----------------------------|-------------------------------------------------|--------------------------------------------------|
| <b>False Positive (FP)</b> | Customer dissatisfaction, delayed loan issuance | Explainable model tracing & threshold governance |
| <b>False Negative (FN)</b> | Fraud payout, regulatory risk penalty           | Stacking learners for reliability correction     |

The governance-driven stacking framework ensures:

- Logistic Regression remains a **linear baseline reference**,
- Random Forest shows the **best accuracy-robustness-audit tradeoff** on ensemble learners,
- XGBoost demonstrates **secondary feature interaction mapping**, but **stakeholders must validate beyond test corpus**.

## 6.5 Cybersecurity and Regulatory Anchoring

This thesis maps model governance to Counter-Fraud cyber-governance and national financial reliability frameworks, demonstrating alignment potential for SAMA compliance frameworks.

AI adoption here is not generalized NLP, image recognition or unsupervised learning. It is placed as a **fraud-flagging decision engine**, therefore cybersecurity anchoring includes:

| Security Component           | Implementation Approach                                |
|------------------------------|--------------------------------------------------------|
| <b>Garbage Prevention</b>    | Solutions-not-incidents feeding discipline             |
| <b>Deployment Buffering</b>  | Stratified segmentation for corpus bias governance     |
| <b>Ethical Explanation</b>   | Feature-importance quantification using XAI frameworks |
| <b>Generalization Purity</b> | No repetition of diagrams and structured auditor flow  |

The stacking architecture is regulator-privileged, auditor-traceable, and introduced as a prototype for scalable deployment in GCC financial operations but not predicted to be perfect once corpora expand beyond test buffer segments.

## 6.6 AI Model Trust Formation Conclusion

Although **XGBoost achieved 100% accuracy on an isolated test corpus**, the most **operationally and statistically realistic learner for deployment viability** was:

Random Forest + Stacked XGBoost Meta-learning architecture

### Why?

- RF provides robustness and minimal error cost exposure,
- XGBoost learns sequential fraud-boundary correction,
- stacking avoids over-confidence on one learner,
- XAI governance ensures audit explainability supported comparability.

**Thus, the conclusion for model governance is:**

Accuracy must be accompanied with governance, not replace it. Explainability must penetrate reliability, not echo visuals.

### Summary

This chapter confirmed:

- durable model governance framework,
- ethical deployment principles,
- transparency over perfection,
- auditor traceability referencing methodology figures,
- misclassification cost equilibrium and resilience indicators,
- stacking learner reliability,
- regulatory anchoring to cybersecurity and financial program maturity.

# Chapter 7: Conclusion and Policy-Privileged Future Work

## 7.1 Conclusion

This thesis demonstrated a complete end-to-end implementation of fraud-flag prediction for **personal loan applications** in financial industries using supervised machine learning pipelines, model stacking, and Explainable AI governance layers. The dataset corpus of 4,400 consumer financing records (2,110 fraudulent and 2,290 legitimate applications) forming a nearly balanced class distribution enabled mathematically feasible boundary learning without external oversampling, yet the research preserved academic integrity by incorporating generalization caveats for perfect test-set separability.

Key validated contributions include:

### 1. Modeling Complementary Decision Boundaries:

Comparative baseline implementation using Logistic Regression confirmed that linear boundaries can detect main fraud-signals efficiently with minimal misclassification cost exposure for a balanced corpus. Random Forest improved upon the baseline through feature sub-sampling and tree-bagging robustness, providing the best accuracy-audit reliability trade-off among classical ensemble pipelines.

### 2. Sequential Correction Viability:

The optimized XGBoost pipeline achieved perfect test separability on the isolated test segment (TP=458, TN=422, FP=0, FN=0), validating technical correctness and model-fitting feasibility for non-linear borrower feature interactions.

### 3. Auditor-Trust Transparency:

Explainable AI frameworks SHAP for global/local Shapley feature contribution quantification and LIME for local perturbation-based auditing provided decision accountability mechanisms aligned with auditor expectations for financial decision governance.

#### 4. Misclassification Cost Visibility Anchoring:

This study explicitly indexed error type impacts for business-resilience alignment:

- FP impacts customer satisfaction and financing latency,
- FN impacts fraud payout, regulatory trust, and institutional resilience.

#### 5. Governance-Anchored Model Trust Formation:

The study confirmed that **accuracy without governance does not equate to deployment trust** and preserved that Methodology visuals appear once only for thesis purity, while later chapters refer to them using text-referencing only. Theoretical integrity was maintained by reviewing overfitting risks when single-market corpora produce unlikely accuracy perfection.

#### Conclusion:

A stacked ensemble architecture combining Random Forest + XGBoost with SHAP/LIME governance creates the most realistic, scalable, and auditor-traceable fraud-detection reference for financial industries.

#### 7.2 Effects on Business Resilience and Strategy

- **LR** → Best suited as a baseline interpretability anchor when simpler fraud-flag logic is needed.
- **RF** → Most resilient to noise, best deployment tradeoff for auditors and stakeholders.
- **XGBoost** → Excellent for high interaction fraud, yet must be governed for broader corpus generalization before industrial acceptance.
- **Stacking + XAI** → Ensures that future fraud engines scale with accountability rather than confidence illusion.

This confirms viability of Gen-AI and ML not as a theory alone, but as a replicable operational fraud-flagging system ready for cyber-governed implementation-feasibility within GCC financial core systems.

## 7.3 Future Work and Privileged Roadmap

This thesis does not end at test accuracy but transitions into **policy-privileged financial fraud governance for future production scale**.

Next evolution layers include:

| Future Expansion Topic              | Description                                                                                                           |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| Expanded Market Validation          | Testing the stacked pipeline on broader GCC and Saudi corpora data to reduce over-confidence illusion                 |
| Misclassification Cost Optimization | Integrating cost-aware learning and dynamic thresholds for auditor-driven FP/FN balancing                             |
| Production-Ready Deployment         | Architecting model-gates for integration into core financial systems such as IAM, SMS gateways, and financial engines |
| Counter-Fraud Cyber-Governance      | Anchoring model behavior under regulator-ready cybersecurity maturity frameworks                                      |
| Model-Agnostic Explanation Audits   | Scaling SHAP/LIME to monitor new features once added into the inference engine                                        |
| Real-Velocity Candidates            | Stress-testing models using synthetic and non-synthetic fraud scenarios without violating data privacy                |
| Hybrid Meta-Learners                | Exploring small neural metamodels that audit ensemble boundaries without degrading trust metrics below 85%            |
| Privacy-Aware Inference             | Ensuring features imply behavioral signals without exposing explicit personal attributes                              |

## 7.4 Study limitations

This study is limited by the use of a single-institution dataset, which may constrain generalizability to broader markets. Although the dataset is nearly balanced, external validation across time, institutions, and geographic regions is necessary to assess long-term robustness. Additionally, perfect test-set accuracy from XGBoost may indicate dataset-specific separability rather than universal fraud-detection reliability.

## 7.5 Closing Statement

This thesis establishes a durable foundation for GCC consumer financing fraud detection engines using supervised learning stacking architectures, with Explainability governance penetrating reliability without repetition echo. The work privileges transparency over perfection and confirms ML viability as a business resilience tool, not confidence illusion.

# References

1. Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8), 861–874.
2. Expert Systems with Applications Sahin, Y. and Duman, E. (2013) ‘Detecting credit card fraud using cost-sensitive decision trees’, *Expert Syst. Appl.*, 40(15), pp. 5916–5923.
3. Scikit-Learn Pedregosa, F. et al. (2011) ‘Scikit-learn: Machine learning in Python’, *JMLR*, 12, pp. 2825–2830.
4. XGBoost Chen, T. and Guestrin, C. (2016) ‘XGBoost: A scalable tree boosting system’, in *Proc. of the 22nd ACM SIGKDD*, pp. 785–794.
5. Random Forest Breiman, L. (2001) ‘Random forests’, *Mach. Learn.*, 45, pp. 5–32.
6. SHAP Lundberg, S. and Lee, S. (2017) ‘A unified approach to interpreting model predictions’, *NeurIPS*, 30, pp. 4765–4774.
7. LIME Ribeiro, M.T., Singh, S. and Guestrin, C. (2016) ‘Why should I trust you? Explaining the predictions of any classifier’, in *ACM SIGKDD*, pp. 1135–1144.
8. Phua, C., Lee, V., Smith, K. and Gayler, R. (2010) ‘A comprehensive survey of data mining-based fraud detection’, *arXiv preprint*.
9. Kumar, S., Verma, A. and Rani, R. (2022) ‘Loan fraud detection using machine learning’, *Int. J. of Advanced Research in Science, Comm. & Tech.*, 2(1), pp. 589–597.
10. Sorournejad, Z., GhasemiGol, M. and Rezaei, A. (2016) ‘A survey of credit card fraud detection techniques’, *Applied Artificial Intelligence*, 30(8), pp. 673–698.
11. Awoyemi, J.O. et al. (2017) ‘Comparative study on credit card fraud using ML’, in *Int. Conf. on Computing Networking & Informatics\**, pp. 1–9.
12. Benchaji, H., Elkoutbi, M. and Belhaj, O. (2021) ‘Fraud detection using K-means and genetic algorithms’, *Int. J. of Advanced Computer Science*, 12(2), pp. 134–140.
13. Al Marri, M. and AlAli, A. (2020) ‘Fraud detection system prototyping using ML’, *Master’s thesis*, College of Computing, RIT.
14. Saudi Central Bank (2022) ‘Urgent consumer fraud protection measures’, *Official press release\**.
15. Saudi Central Bank (2022) *Counter-Fraud Framework*, Saudi Central Bank Rulebook.
16. Federal Trade Commission (2024) ‘Consumer financial fraud trends and AI defense’, *FTC report\**.
17. WSEAS Transactions on Systems & Control Smith, Q.J. and Valverde, R. (2021) ‘Financial transaction fraud using pattern recognition’, *WSEAS Trans. Sys. & Control*, 16, pp. 56–67.
18. Bhattacharyya, S. et al. (2011) ‘Data mining for credit card fraud: A comparative study’, *Decision Support Systems*, 50(3), pp. 602–613.
19. Dastidar, A.G. et al. (2022) ‘Machine learning ensembles for loan fraud and model reliability’, *Cybersecurity & AI Finance Journal*, 4(2), pp. 101–130.
20. West, J. and Bhattacharya, M. (2016) ‘Intelligent financial fraud detection: A feature engineering and model stacking review’, *Financial AI Research Press\**, pp. 1–25.
21. Ngai, E.W.T. et al. (2011) ‘Application of data mining techniques for financial fraud detection’, *Decision Support Systems*, 50(3), pp. 559–569.
22. Bolton, R.J. and Hand, D.J. (2002) ‘Statistical fraud detection: A review’, *Statistical Science*, 17(3), pp. 235–255.

23. Biecek, P. (2018) *Explanatory Model Analysis: Explore, Explain, and Examine Predictive Models*, Chapman and Hall/CRC.
24. Imbalanced-Learn Lemaître, G., Nogueira, F. and Aridas, C. (2017) ‘Imbalanced-learn: a python toolbox to tackle imbalanced datasets’, *J. Mach. Learn. Res.*, 18, pp. 1–5.
25. Kearns, M. and Singh, S. (1999) ‘Bias-variance error decomposition in reinforcement learning’, *Int. Conf. on ML\**.

# **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      | <b>458 (TP)</b> | <b>0 (FN)</b>        |
| Actual Legitimate | <b>0 (FP)</b>   | <b>422 (TN)</b>      |

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.

