OPEN UNIVERSITY OF KENYA

MASTER'S THESIS

A LOG-BASED MACHINE LEARNING MODEL FOR IDENTITY THEFT DETECTION IN FINANCIAL INFORMATION SYSTEMS

KOECH SAMUEL KIPKURUI
SCHOOL OF SCIENCE

A Master's Thesis Submitted in Partial Fulfilment of the Requirements for the Award of the Master of Science in Cybersecurity and Digital Forensics of the Open University of Kenya.

OPEN UNIVERSITY OF KENYA

OCTOBER 2026

COPYRIGHT

All rights reserved. No part of this thesis/project may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without the prior written permission of the author or The Open University of Kenya.

© 2026, Koech Samuel Kipkurui

DECLARATION

I hereby declare that this thesis/project is my original work and has not been submitted for the award of a degree in any other university or institution of higher learning.

Candidate:

Name: ___**Koech Samuel Kipkurui**_*Signature:*  *Date:* _*15/10/2026* _

This thesis/project has been submitted for examination with my/our approval as university supervisor(s).

Supervisor:

Name: ___ **Dr. Charles Katila**___*Signature:*  *Date:* _*15.10.2026*____

Co-Supervisor (if applicable):

Name: _____Signature: _____ Date: _____

DEDICATION

I dedicate this work to God Almighty, my creator, in gratitude for His grace, intelligence, courage, and guidance throughout the course of my education. I also dedicate it to my family for their unwavering support, love, tolerance, and inspiration, as well as to everyone who has motivated and encouraged me to pursue this academic success.

ACKNOWLEDGEMENTS

First and above all I would like to thank the Lord God for His endless blessings, love and favor that has enabled me to scale the heights in life and in academic arena.

I am extremely grateful to Dr. Charles Katila for sharing with me your immense knowledge and wisdom in research. The pace at which you read my work with keenness and skillfully gave unique in depth and demanding input is incredible. Your input and commitment have greatly helped to shape and improve the quality of this work.

I am also thankful to Dn r. Ednah Milgo the coordinator of the programme; your devotion to your students' progress and relentless encouragement is admirable. I also acknowledge all faculty staff for their dedication in sharing their knowledge, skills and experience in various units and forums.

I am thankful to The Open University of Kenya for according to me an academic environment and for demonstrating a caring attitude that I greatly needed to survive in such a rigorous programme. I also acknowledge the University for making available psychometric testing materials both for training and research purposes.

Last but not least is a profound gratitude to my family for their love, patience, generosity and unwavering support. Without their love and support, I would never have reached this far.

ABSTRACT

The growing reliance of commercial banks on digital banking platforms, internet transactions, and mobile financial services has made identity theft and account takeover attacks serious cybersecurity threats. The growth of digital financial services in Kenya has made it easier for cybercriminals to target sensitive financial data and client credentials. Current rule-based and signature-based detection systems are frequently reactive and may be ineffective against identity-related attacks involving compromised credentials, behavioral anomalies, and changing attack patterns. Additionally, behavioral analytics based on authentication, access, session, and transaction logs within Kenyan commercial banking contexts has received little attention in previous machine-learning studies, which often rely on generic or datasets.

The purpose of this study was to design and evaluate a log-based machine-learning model for identifying identity theft in Kenyan commercial banks' financial information systems. A machine-learning pipeline that incorporated data preprocessing, behavioral feature engineering, feature selection, class-imbalance handling, model training, prediction, performance evaluation, and explainability was developed by the study using Design Science Research (DSR). Analysis was done on device-related behavioral indicators, transactions, sessions, and authentication. Accuracy, precision, recall, F1-score, and ROC-AUC were utilized to evaluate the model's performance, while feature-importance analysis was employed to enhance interpretability.

The findings demonstrated that behavioral characteristics are useful indicators for detecting suspicious identity-related activities. Among the most prominent reported characteristics were failed login attempts, login frequency, transaction amount, session time, and device changes. The findings demonstrate that Random Forest significantly outperformed Isolation Forest in every significant evaluation criteria. In comparison to Isolation Forest, Random Forest recorded a significantly reduced false-positive rate (0.000008) and higher precision (0.991968),

recall (0.794212), F1-score (0.882143), ROC-AUC (0.994847), and PR-AUC (0.935154). Precision (0.003436), recall (0.003215), F1-score (0.003322), and PR-AUC (0.003729) were all extremely low for Isolation Forest.On the evaluated test data, Random Forest outperformed Isolation Forest in terms of detection and classification overall. The study shows that rather than depending only on conventional rule-based techniques, effective identity-theft detection should use behavioral analytics, machine learning, explainability, and continuous monitoring. It recommends employing explainable AI, integrating behavioral analytics with current security-monitoring and SIEM systems, keeping human oversight, regularly assessing model performance, and retraining models to take into account evolving user behavior and attack trends.

Keywords: identity theft detection, machine learning, financial systems, behavioral analytics, log analytics, behavioural analytics, Random Forest, Isolation Forest, anomaly detection, explainable artificial intelligence, cybersecurity.

TABLE OF CONTENTS

Table of Contents

COPYRIGHT i
DECLARATION ii
DEDICATION iii
ACKNOWLEDGEMENTS iv
ABSTRACT v
TABLE OF CONTENTS vii
List of Tables x
List of Figures xi
List of Abbreviations and Acronyms xii
List of Appendices xiv
CHAPTER ONE: INTRODUCTION 1
1.1 Background of the Study 1
1.2 Statement of the Problem 3
1.3 Purpose of the Study 4
1.4 Research Objectives 5
1.5 Research Questions 6
1.6 Research Hypotheses 6
1.7 Significance of the Study 6
1.8 Scope and Delimitations 7
1.8.1 Scope of the Study 7
1.9 Limitations of the Study 9
1.10 Operational Definition of Key Terms 10
1.11 Organization of the Study 12
CHAPTER TWO: LITERATURE REVIEW 14
2.1 Introduction 14
2.2 Theoretical Framework 14
2.2.1 Anomaly Detection Theory 15
2.2.2 Machine Learning Theory 16
2.2.3 User Behaviour Analytics (UBA) Theory 17
2.2.4 CIA Triad as a Guiding Information Security Framework 18
2.2.5 Integrated Theoretical Foundation 19
2.3 Conceptual Framework 20
2.4 Empirical Literature Review 22
2.4.1 Introduction 22
2.4.2 Empirical Evidence on Machine Learning for Financial Identity Theft Detection 23
2.4.3 Empirical Evidence on System-Log Anomaly Detection 25
2.4.4 Streaming, Real-Time and Highly Imbalanced Identity Theft Detection 27
2.4.5 Empirical Evidence on Explainable Artificial Intelligence 28
2.4.6 Behavioural Analytics and Identity-Centric Detection 29
2.4.7 Relevance to Kenyan Commercial Banking 30
2.4.8 Synthesis of Empirical Literature 31
2.5 Critique of Existing Literature 41

2.5.1 Methodological Gaps 42
2.5.2 Contextual Gaps 44
2.5.3 Theoretical Gaps 45
2.5.4 Explainability and Operational Gaps 46
2.5.5 How the Present Study Addresses the Critique 46
2.6 Research Gaps 47
2.6.1 Empirical Gap 48
2.6.2 Contextual Gap 48
2.6.3 Behavioral Analytics Gap 49
2.6.4 Algorithmic Gap 49
2.6.5 Evaluation Gap 49
2.6.6 Explainability Gap 50
2.6.7 Real-Time and Adaptive Detection Gap 50
2.6.8 Integrated Research Gap 50
2.7 Summary of Literature Review 51
CHAPTER THREE: RESEARCH METHODOLOGY 53
3.1 Introduction 53
3.2 Research Philosophy 53
3.3 Research Design 54
3.4 Study Area 55
3.5 Target Population 56
3.6 Sampling Design 57
3.6.1 Sample Size 57
3.6.2 Sampling Technique 58
3.7 Data Collection 58
3.7.1 Data Collection Instruments 59
3.7.2 Data Collection Procedure 60
3.8 Model Development and Simulation 61
3.8.1 Log Data Acquisition Layer 62
3.8.2 Data Preprocessing Layer 62
3.8.3 Feature Engineering Layer 62
3.8.4 User Behaviour Profiling Layer 63
3.8.5 Machine-Learning Engine 63
3.8.6 Hyperparameter Reporting and Tuning 63
3.8.7 Explainability Layer 69
3.8.8 Decision and Alert Layer 70
3.9 Pilot Study 70
3.10 Validity and Reliability 71
3.10.1 Validity 71
3.10.2 Reliability 72
3.10.3 Model Evaluation 72
3.11 Data Analysis and Presentation 73
3.12 Ethical Considerations 76
CHAPTER FOUR: DATA ANALYSIS, RESULTS AND DISCUSSION 79
4.1 Introduction 79
4.2 Dataset Characteristics and Data Preparation 79
4.3 Descriptive Statistics of the Log Dataset 81
4.4 Diagnostic and Assumption Tests 83
4.5 Findings per Objective 83
4.5.1 Findings for Objective One 83
Table 4.5.1: Top Behavioral Predictive Features 84

4.5.2 Findings for Objective Two 85
4.5.3 Findings for Objective Three 91
4.6 Discussion of Findings 98
4.6.1 Discussion for Objective One 98
4.6.2 Discussion for Objective Two 98
4.6.3 Discussion for Objective Three 99
CHAPTER FIVE: SUMMARY, CONCLUSIONS AND RECOMMENDATIONS 102
5.1 Introduction 102
5.2 Summary of Findings 102
5.2.1 Summary of Findings for Objective One 102
5.2.2 Summary of Findings for Objective Two 103
5.2.3 Summary of Findings for Objective Three 105
5.3 Conclusions 106
5.3.1 Conclusion on Objective One 106
5.3.2 Conclusion on Objective Two 106
5.3.3 Conclusion on Objective Three 107
5.4 Recommendations 108
5.4.1 Recommendations for Practice 108
5.4.2 Recommendations for Policy and Regulation 109
5.4.3 Recommendations for Capacity Building 109
5.5 Recommendations for Further Research 110
5.6 Contribution to Knowledge 111
5.7 Chapter Summary 112
REFERENCES 114

List of Tables

Table 1: Alignment of variable type, Construct, and Key Indicators. 22
Table 2: Empirical Literature Matrix 38
Table 3: Objective Alignment Summary 39
Table 4: Random Forest Hyperparameter Configuration 64
Table 5: Isolation Forest Hyperparameter Configuration 65
Table 6: Data Splitting and Resampling Configuration 66
Table 7: Alignment of Research Objectives, Variables, Analysis Methods and Tools 75
Table 8:Dataset Characteristics 81
Table 9:Principal Behavioural Features 83
Table 10: Top Behavioral Predictive Features 84
Table 11:Class distribution after SMOTE 86
Table 12: Random Forest–Isolation Forest Comparison 91

List of Figures

Figure 1: Conceptual Framework 21
Figure 2: Development and Simulation of a Log-Based machine learning model for Identity Theft Detection in Financial Information Systems 70
Figure 3: Ranking of behavioral features 85
Figure 4: Class distribution after SMOTE 87
Figure 5: Login outcomes and failed logins per user (integrated chart report). 87
Figure 6: Device changes indicating suspicious access activity (integrated chart report). 88
Figure 7: Distribution of flagged identity-theft threat indicators (integrated chart report). 89
Figure 8: Distribution of identity-theft risk indicators per user account (integrated chart report). 90
Figure 9: Trends of lagged suspicious events over time (integrated chart report). 90
Figure 10: Precision Scores-Random Forest and Isolation Forest Comparison 92
Figure 11: Recall Scores-Random Forest and Isolation Forest Comparison 93
Figure 12: F1-Scores-Random Forest and Isolation Forest Comparison 94
Figure 13: ROC-AUC Scores-Random Forest and Isolation Forest Comparison 95
Figure 14: PR-AUC Scores-Random Forest and Isolation Forest Comparison 96
Figure 15: False Positive Rate Scores-Random Forest and Isolation Forest Comparison 97

List of Abbreviations and Acronyms

Acronym	Definition
AI	Artificial Intelligence
AML	Anti-Money Laundering
API	Application Programming Interface
AVE	Average Variance Extracted
CBK	Central Bank of Kenya
CFA	Confirmatory Factor Analysis
CIA	Confidentiality, Integrity, and Availability
CNN	Convolutional Neural Network
DMP	Data Management Plan
DSR	Design Science Research
EDA	Exploratory Data Analysis
EFA	Exploratory Factor Analysis
F1-Score	Harmonic Mean of Precision and Recall
FATF	Financial Action Task Force
ICT	Information and Communication Technology
IDS	Intrusion Detection System
IPS	Intrusion Prevention System
IS	Information Systems
ISO	International Organization for Standardization
IT	Information Technology

KYC	Know Your Customer
LSTM	Long Short-Term Memory
MFA	Multi-Factor Authentication
ML	Machine Learning
NACOSTI	National Commission for Science, Technology and Innovation
NIST	National Institute of Standards and Technology
PCA	Principal Component Analysis
PCI DSS	Payment Card Industry Data Security Standard
PII	Personally Identifiable Information
RBAC	Role-Based Access Control
REC	Research Ethics Committee
RF	Random Forest
ROC	Receiver Operating Characteristic
SEM	Structural Equation Modelling
SIEM	Security Information and Event Management
SMOTE	Synthetic Minority Oversampling Technique
SOC	Security Operations Center
SVM	Support Vector Machine
UEBA	User and Entity Behavior Analytics
XAI	Explainable Artificial Intelligence

List of Appendices

APPENDICES 118
APPENDIX A: RESEARCH INSTRUMENTS 118
A2. System Log Data Collection Template 118
A.3 Dataset / Data Source Registry 118
A.4 Project File Structure 119
APPENDIX B: INTRODUCTORY / PERMISSION LETTERS 120
B1. University Introductory Letter 120
B2. Institutional Permission Letter 121
APPENDIX C: ETHICAL APPROVAL 122
C1. Institutional Ethics Approval 122
C2. NACOSTI Research Permit 123
C3. Data Protection and Confidentiality Documentation 123
APPENDIX D: RAW DATA / CODE EXCERPTS 124
D1. Data Dictionary 124
D2. Data Preprocessing Code 124
D3. Model Development Code 124
D4. Example Non-Sensitive Code Structure 125
APPENDIX E: TURNITIN SIMILARITY REPORT 127
E1. Similarity Report Details 127
APPENDIX F: PUBLISHED PAPERS 128
F1. Publications Arising from the Study 128
APPROVAL PAGE 129

CHAPTER ONE: INTRODUCTION

1.1 Background of the Study

Globally, banking operations have undergone a substantial transformation due to the growing digitization of financial services, Kenya included. Commercial banks increasingly heavily rely on cloud computing, mobile banking, online banking, and integrated financial information systems to effectively provide services to their customers. These technologies improve accessibility and operational efficiency, but they also expose financial institutions to sophisticated cyberthreats including identity theft and account takeover.

Identity theft is the illegal acquisition and fraudulent use of a person's financial or personal information. Session hijacking, phishing attacks, illegal account access, compromised credentials, and fraudulent transactions are prevalent forms of identity theft in financial information systems. The Communications Authority of Kenya (2024) reports that there has been a notable rise in cyber threat events in Kenya that target financial systems, such as phishing assaults, credential compromise incidents, and unauthorized access attempts. In a comparable manner, the Kenya Bankers Association (2023) noted that as financial institutions develop their digital banking platforms, there is an ongoing rise in digital fraud and account takeover attacks.

Identity theft and digital fraud have grown to be significant cybersecurity issues in the banking industry on a global scale. Due to the significant value of financial and personal data, banking institutions were listed as one of the most targeted industries in the IBM Cost of a Data Breach Report (2024). According to the research, one of the primary reasons for data breaches worldwide is still compromised credentials. According to the Verizon Data Breach Investigations Report (2024), a significant percentage of cyber events affecting financial firms are caused by phishing assaults and credential abuse.

Kenya's rapid adoption of digital financial services has led to a huge growth in the country's cyberattack surface. The Central Bank of Kenya (CBK) noted increasing issues with cybersecurity that include internet banking platforms, mobile banking apps, and interbank transaction systems in its 2023 Annual Banking Supervision Report. Additionally, the Office of the Data Protection Commissioner (ODPC) has emphasized the growing need for enhanced cybersecurity measures and behavioral monitoring systems in order to protect consumer data and maintain trust in digital financial ecosystems.

Most traditional identity theft detection systems used in financial systems rely on signature and rules-based techniques. Although these systems can detect known attack patterns, they frequently fail to detect new or evolving identity theft tactics. Rule-based systems have high false-positive rates and need frequent manual adjustments to be effective. As cybercriminals' tactics get more complex, financial institutions require increasingly sophisticated detection techniques that can identify anomalies in real time and comprehend behavioral patterns.

Machine learning has emerged as a prospective technique for identifying cybersecurity threats since it can automatically identify unusual patterns in massive quantities of operational data. Log-based machine learning algorithms analyze authentication logs, transaction histories, IP addresses, session durations, access frequency, geolocation data, and behavioral patterns to identify suspicious activities associated with identity theft.Research by Ahmed et al. (2019), Li et al. (2021), and Sharma and Gupta (2021) shows that machine learning methods can greatly enhance cybersecurity environments' capacity for anomaly detection.

Research on machine learning-based Identity Theft is expanding, but there are still a number of unanswered questions. Class imbalance, explainability, adaptation to evolving assault patterns, and operational deployment inside real-world financial systems are common challenges faced by existing models. Furthermore, a lot of research uses artificial datasets or non-financial settings,

which restricts its relevance to Kenyan commercial banking settings.

Therefore, the purpose of this study was to develop and assess a log-based machine learning model for identifying identity theft in Kenyan commercial bank information systems. Behavioral log analytics, anomaly detection, explainable AI integration, and assessment of detection efficacy in banking cybersecurity environments are the study's primary areas of interest.

1.2 Statement of the Problem

The raising rate of cybersecurity risks to financial information systems globally and in Kenya's banking industry revolves around identity theft and account takeover attacks. The increasing prevalence of digital banking platforms, mobile banking apps, and online financial services has broadened the attack surface for attackers who target client credentials and sensitive financial data. The Communications Authority of Kenya (2024) reports that vital digital infrastructures, including banking systems, are experiencing an upsurge in cyber threat incidents involving phishing assaults, illegal access attempts, and credential compromise.

According to the Central Bank of Kenya Banking Supervision Report (2023), digital fraud and cybersecurity threats are two of the main operational hazards that commercial banks must deal with. According to the Kenya Bankers Association (2023), automation, social engineering, and credential theft techniques are making identity-related fraud and account takeover attempts more complex. This attacks because monetary damages harm to one's reputation and a decline in customer confidence, and interruptions to business operations.

Signature-based and rule-based Identity theft systems are still widely used by commercial banks, despite investments in cybersecurity technologies. These systems are less effective against changing threats and unidentified attack behaviors, even when they can identify predetermined assault patterns. High false-positive rates, poor adaptation to evolving attack patterns, limited

explainability of machine learning conclusions, computational complexity, and challenges with highly imbalanced fraud datasets are only a few of the technical drawbacks of current techniques.

Additionally, current machine learning-based Identity theft models have contextual constraints. Numerous researches make use of non-financial contexts, datasets, or generalized Identity theft datasets that could not fairly represent the behavioral patterns found in Kenyan banking systems. Additionally, log-based behavioral analytics for identifying identity theft in Kenyan commercial bank information systems has not received much attention.

Furthermore, authentication behavior, session abnormalities, login patterns, device usage, geolocation discrepancies, and user behavioral profiling generated from operational system logs are often overlooked in favor of transaction Identity theft in many existing detection methods. As a result, identity theft attempts employing stolen credentials could go unnoticed until significant harm has already been done.

Therefore, a log-based machine learning model that can assess behavioral indications generated by banking system logs is needed to enhance the effectiveness of identity theft detection. The purpose of this study was to develop and evaluate a machine learning-based identity theft detection model for Kenyan commercial banks using transaction logs, user activity analytics, and authentication logs. The study also sought to evaluate the relative effectiveness of particular machine learning algorithms using explainability techniques suitable for regulated financial contexts.

1.3 Purpose of the Study

The purpose of this study was to develop and evaluate a log-based machine learning model for identifying identity theft in Kenyan commercial banks' financial information systems. In order to find patterns linked to questionable identity-

related activity, the study concentrated on examining behavioral indications obtained from authentication, transaction, access, session, and device logs. Additionally, it aimed to incorporate explainable AI and behavioral analytics into the detection process and assess the efficacy of the suggested strategy using suitable machine-learning performance measures. The study's objectives were to enhance cybersecurity monitoring, detect identity theft more successfully, and prevent financial information systems in Kenyan commercial banking environments.

1.4 Research Objectives

1.4.1 General Objective

To design and evaluate a log-based machine learning model for detecting identity theft in financial information systems within Kenyan commercial banks.

1.4.2 Specific Objectives

1. To assess existing identity theft detection mechanisms and analyze behavioral anomaly patterns within authentication and transaction logs in Kenyan commercial banking information systems.
2. To design and implement a log-based machine learning model for detecting identity theft using behavioral analytics and explainable artificial intelligence techniques.
3. To evaluate the effectiveness of the proposed model and comparatively assess Random Forest and Isolation Forest using Precision, Recall, F1-score, ROC-AUC, PR-AUC and False Positive Rate.

1.5 Research Questions

1. What identity theft detection mechanisms and behavioral anomaly patterns exist within authentication and transaction logs in Kenyan commercial banking information systems?
2. How can a log-based machine learning model integrating behavioral analytics and explainable artificial intelligence techniques improve identity theft detection in banking information systems?
3. How effective is the proposed model, and how do Random Forest and Isolation Forest comparatively perform on the dataset's fraud-labelled transactions using Precision, Recall, F1-score, ROC-AUC, PR-AUC and False Positive Rate?

1.6 Research Hypotheses

H_0: Null Hypothesis

There is no statistically significant relationship between behavioural patterns derived from authentication and transaction logs and the effectiveness of identity theft detection in Kenyan commercial banks.

H_1: Alternative Hypothesis

There is a statistically significant relationship between behavioural patterns derived from authentication and transaction logs and the effectiveness of identity theft detection in Kenyan commercial banks.

1.7 Significance of the Study

This study is significant because it aims to enhance cybersecurity knowledge and practice by designing, implementing, and testing a machine learning-based identity theft detection model utilizing financial system logs.

Theoretical and Knowledge Contribution: By investigating how financial transaction patterns, login activities, session characteristics, and usage patterns can be used to identify suspicious activities, the study will add to the body of

knowledge regarding the application of machine learning in identity theft detection. It will offer empirical insights into how machine learning capabilities, user behavior profiling, log data quality, and identity theft detection efficacy are related. The results might potentially serve as a basis for future investigations into intelligent cybersecurity solutions for financial information systems.

Practical and Policy Contribution: The suggested methodology should help financial institutions detect identity theft more effectively, track fraud better, and identify suspicious transactions and unauthorized access more quickly. The study will present data on the detection performance of various machine learning methods, including precision, recall, F1-score, ROC-AUC, and false positive rates. The results may help financial organizations, especially those in Kenya's financial sector, build cybersecurity policies, security monitoring protocols, and data protection practices.

Significance to Beneficiaries: Better methods for identifying possible identity theft incidents are potentially beneficial for financial institutions, such as banks and other financial service providers. The results may be used by governmental organizations and regulatory bodies to guide programs related to digital security, financial data protection, and cybersecurity awareness. The study's methodology, experimental findings, and suggestions for further research may be useful to academic institutions, researchers, and postgraduate students. In the end, by encouraging the creation of more efficient and data-driven cybersecurity measures, the study may help to increase public trust in digital financial services.

1.8 Scope and Delimitations

1.8.1 Scope of the Study

The development and evaluation of a log-based machine learning model for identifying identity theft in financial information systems at Kenyan commercial banks is the primary purpose of this study. Behavioral analysis of authentication

logs, transaction logs, access records, session activities, and user interaction patterns is the focus of this study.

Because banking systems generate more consistent authentication and transaction log structures than insurance companies and fintech platforms, the study was restricted to commercial banking settings. The consistency of behavioral indicators, anomaly patterns, and evaluation criteria is enhanced by concentrating on commercial banks.

The study used anonymized operational logs where possible; publicly available cybersecurity datasets relevant to identity theft detection and authentication anomalies was used to enhance implementation. Candidate datasets may include:

- CERT Insider Threat Dataset
- IEEE-CIS Fraud Detection Dataset
- CICIDS Cybersecurity Datasets
- Simulated banking authentication datasets

Log Attributes that were analyzed

The study focused on the following log attributes:

- Timestamp
- User ID
- Source IP address
- Device fingerprint
- Geolocation
- Session duration
- Failed login attempts
- Transaction amount
- Transaction frequency
- Access patterns

- Authentication success/failure records

Geographical Scope

The study focused on Kenyan commercial banks operating digital banking platforms.

Time Scope

The study analyzed log data generated within the most recent 12-month operational period where accessible.

1.9 Limitations of the Study

During the process of the investigation and experimentation, the study ran into a number of challenges as describe below:

First, owing to privacy, security, and legal constraints, access to actual financial system and user activity logs was restricted. By leveraging the Financial System logs Datasets for Identity Theft Detection and making sure that personal identifiers were not kept in public repositories and that only authorized individuals could access research data, this limitation was lessened.

Second, abnormal observations represented only 0.129% of the dataset, demonstrating a significant class imbalance in the study. The Synthetic Minority Oversampling Technique (SMOTE), which improves minority-class learning while lowering the danger of information leaking into the test set, was used to mitigate this.

Third, the findings' ability to accurately reflect the complexity and behavioral diversity of commercial banking environments was restricted by the use of a publicly accessible dataset. This was lessened by employing behavioral feature engineering, thorough preprocessing, evaluating the model using a variety of per-

formance indicators, and carefully interpreting the results rather than assuming production-level performance.
Lastly, a potential concept-drift constraint was demonstrated by developing identity-theft strategies and shifts in genuine user behavior. In order to sustain detection efficacy as behavioral patterns and assault strategies evolve, the study suggested ongoing monitoring, recurring model evaluation, and retraining.

Denial-of-service attacks were not specifically trained or assessed for the model. The current identity-theft labels and metrics do not establish DoS event detection; DoS is an availability-focused threat that necessitates suitable network and service monitoring. A different labeled dataset and experiment should be used to validate any future claims of more comprehensive cyber-threat detection.

1.10 Operational Definition of Key Terms

Term	Operational Definition
Identity Theft	The unauthorized use of another person's identity or credentials to access financial systems, initiate transactions, or conduct fraudulent activities. In this study, it refers to suspicious activities represented by anomalous financial transaction records.
Machine Learning (ML)	A branch of artificial intelligence that enables computer systems to learn patterns from data and make predictions. In this study, it refers to the algorithms trained and evaluated to identify potential identity theft activities using financial system logs.

Financial System Logs	Records generated by financial information systems containing transaction details, user activities, and other relevant system events. In this study, they refer to the dataset used to develop and experimentally evaluate the proposed detection model.
Identity Theft Detection Model	A computational model designed to identify potentially fraudulent or unauthorized activities. In this study, it refers to the machine learning-based model developed, implemented, and evaluated using financial system log data.
Anomaly Detection	The identification of data records that deviate from expected patterns of activity. In this study, it refers to identifying financial transactions or activities classified as anomalous by the dataset or detection model.
User Behaviour Profiling	The analysis of user activity patterns to establish typical behaviour. In this study, it refers to examining features such as login frequency, session duration, failed login attempts, and device changes to identify potentially suspicious behaviour.
Log Data Quality	The degree to which system log data is accurate, complete, consistent, and suitable for analysis. In this study, it concerns the reliability and completeness of the financial transaction records used for model training and evaluation.
Detection Effectiveness	The ability of a detection model to correctly identify suspicious activities while limiting incorrect alerts. In this study, it is measured using precision, recall, F1-score, ROC-AUC, and false positive rate.

False Positive Rate (FPR)	The proportion of legitimate activities incorrectly classified as suspicious. In this study, it is used to assess the extent to which the model generates false alerts during evaluation.
Feature Engineering	The process of creating, transforming, or selecting data attributes to improve machine learning performance. In this study, it involves preparing transaction and behavioural variables for identity theft detection.
Experimental Evaluation	The systematic assessment of a model's performance using defined datasets, experimental procedures, and evaluation metrics. In this study, it refers to comparing the performance of selected machine learning algorithms on the financial transaction dataset.

1.11 Organization of the Study

This thesis, study titled "Machine Learning-Based Identity Theft Detection Using Financial System Logs: Design, Implementation, and Experimental Evaluation," is divided into five chapters, each of which focuses on a different facet of the research.

Chapter One: Presents the study's background information, problem statement, research aims, research questions, study significance, study scope, limits, assumptions, and operational definitions of important terminology are all presented in chapter one Introduction heading. It provides the background and rationale for creating a machine learning-based model that uses financial system data to detect identity theft.

Chapter Two: Literature review examines and analyzes existing literature on identity theft, cybersecurity in financial information systems, machine learning

applications in fraud and identity theft detection, financial system log analysis, and user behavior profiling. It gives the conceptual framework directing the investigation, identifies research gaps, and examines pertinent theoretical and conceptual frameworks.

Chapter Three: Research Methodology describes the research design, data sources, dataset characteristics, data preprocessing, feature engineering, feature selection, model building, and experimental techniques. Additionally, it describes the performance measures used to evaluate the suggested model as well as the machine learning methods chosen for assessment.

Chapter Four: findings, Findings, and Discussion experiments the findings from applying and assessing selected machine learning algorithms are presented in Metrics including precision, recall, F1-score, ROC-AUC, and false positive rate are used to evaluate and contrast their performance. The results are examined in light of the goals of the study and the body of current literature.

Chapter Five: The study's main conclusions are outlined in Summary, Conclusions, and Recommendations in relation to the research questions and objectives. It draws conclusions from the experimental assessment of the chosen machine learning algorithms. The chapter offers suggestions for enhancing financial information systems' ability to detect identity theft. It also highlights the study's limitations and suggests areas for future research.

CHAPTER TWO: LITERATURE REVIEW

2.1 Introduction

This chapter reviews literature relevant to the development of a log-based machine learning model for identity theft detection in Kenyan commercial banking information systems. It examines the theoretical foundations of the study, focusing on Anomaly Detection Theory, Machine Learning Theory, User Behavior Analytics (UBA) Theory, and the CIA Triad, followed by an empirical review of identity theft, log-based anomaly detection, machine learning applications, and cybersecurity in financial systems. The chapter critically evaluates existing studies, identifies theoretical, empirical, methodological and practical research gaps, and presents the conceptual framework linking log data, machine learning capabilities, behavioral analytics and identity theft detection effectiveness. The review provides the scholarly basis for the study and establishes the need for a context-specific approach to identity theft detection within Kenyan commercial banking environments.

2.2 Theoretical Framework

In order to detect identity theft in Kenyan commercial banking information systems, the theoretical framework offers the intellectual basis for understanding how log data, behavioral patterns, machine learning approaches, and cybersecurity policies can be incorporated. The Confidentiality, Integrity, and Availability (CIA) Triad offers a supplementary information security framework, while Anomaly Detection Theory, Machine Learning Theory, and User Behavior Analytics (UBA) Theory serve as the main theoretical foundations for this study. Since the study aims to find aberrations from normal user behavior, discover patterns in banking records, and produce efficient identification of potentially compromised identities, the merging of various viewpoints is permissible.The study's scope specifically covers authentication logs, transaction logs, access records, session activities and user interaction patterns within Kenyan commercial banking environments.

2.2.1 Anomaly Detection Theory

The goal of anomaly detection theory is to identify patterns or instances in a given dataset or operational environment that significantly differ from what is considered normal. The work of Chandola, Banerjee, and Kumar (2009), who systematically categorised anomaly detection methodologies and highlighted those various strategies rely on specific assumptions regarding the distinction between normal and abnormal observations, represents a significant scholarly basis for anomaly detection.

The fundamental concept of anomaly detection is that while malicious or abnormal activity result in noticeable departures from such patterns, legitimate activities typically display recognizable patterns. For instance, a user may typically access an account in a financial information system from a constant location, device, time period, and interaction pattern. Therefore, an anomaly could be a sudden mix of repeated unsuccessful authentication attempts, an unfamiliar device, an unexpected location, irregular session duration, or atypical transaction behavior. However, the hypothesis does not presume that every anomalous occurrence is malevolent. This distinction is important because legitimate users may occasionally behave differently from their established patterns, creating potential false-positive alerts. The study additionally recognizes that it can be challenging to establish stable behavioral baselines in dynamic financial contexts and that anomalous conduct is not always malicious.

The effectiveness of identity theft detection, the study's dependent variable, has a strong connection to anomaly detection theory. It provides the theoretical foundation for transforming transaction events and unprocessed authentication into signs of questionable activity. Additionally, it facilitates the creation of risk-scoring systems, user profiling, authentication anomaly detection, and behavioral baselines. Machine learning techniques are used in this study to operationalize the theory by identifying data whose behavioral traits deviate from known patterns. This is especially pertinent to Isolation Forest, which aims to detect unusu-

al observations without necessitating the prior labeling of every type of identity theft.

2.2.2 Machine Learning Theory

The computational foundation for systems that identify patterns in data and use those patterns to prediction, classification, or anomaly detection is provided by machine learning theory. The idea is widely linked to the advancement of computational learning techniques in statistical pattern recognition and artificial intelligence. Machine learning enables systems to deduce relationships and patterns from data that already exists instead of depending solely on manually established rules.

Supervised learning and unsupervised learning are two learning paradigms that are particularly relevant to this research. Supervised learning is predicated on the idea that past observations may be linked to predefined labels and that a model can identify the connection between input attributes and those labels. In this study, the categorization of suspicious financial transactions can be supported by labeled lawful and fraudulent operations. Unsupervised learning, on the other hand, looks for structures, patterns, or odd findings in the data and does not require full predetermined labels. Because labeled fraudulent observations are frequently few, unsupervised and semi-supervised procedures have become more popular in financial fraud research.

Therefore, the study's use of Random Forest and Isolation Forest is justified by the theory. While Isolation Forest presents an anomaly-oriented method for discovering observations that deviate from the prevalent behavioral patterns, Random Forest presents a supervised learning mechanism for learning correlations between behavioral/log attributes and known outcomes. These two methods are specifically mentioned in the proposal as the analytical basis for the identity-theft detection model.

The study's main fundamental assumptions are that log data can be used to identify meaningful behavioral patterns and that appropriate elements include in-

formation that can differentiate between suspicious and legitimate activities. As a result, the training data's quality, completeness, and representativeness are crucial. Class imbalance, overfitting, computational complexity, explainability, and reliance on training-data quality are further issues that machine learning faces. Because fraudulent observations may only make up a small percentage of all banking activity, these constraints are especially crucial for identity-theft detection.

Thus, the study's log data and ML abilities are linked by Machine Learning Theory as independent constructions to the dependent construct of identity theft detection effectiveness. Predictive or anomaly-detection algorithms can learn behavioral patterns from raw log variables including unsuccessful login attempts, transaction amount, transaction frequency, session duration, device fingerprint, source IP address, and geolocation, according to the idea.

2.2.3 User Behaviour Analytics (UBA) Theory

User Behaviour Analytics Theory focuses on understanding users through patterns of interaction with information systems and identifying activities that depart from established behavioural norms. Because genuine credentials can be abused by an unauthorized individual while appearing to be technically valid, UBA is especially relevant to cybersecurity. As a result, identifying identity breach necessitates looking at how, when, where, and through which the authenticated account is being used in addition to whether authentication was successful.

The main premise of UBA is that over time, individual users or user groups would display comparatively recognizable behavioral patterns. Deviation from these patterns may reveal unusual activity, account takeover, or compromised credentials. Behavioral modeling might include communication patterns, device usage, location, session characteristics, authentication frequency, and access timings in cybersecurity applications. Research on user-behavior modeling has shown that behavioral anomaly detection can be used to security issues involving imbalanced datasets and infrequent harmful activity.

Because identity theft basically constitutes an identity-centric security issue, UBA is especially significant to the current investigation. Authentication behavior, access frequency, session activity, geolocation, and device usage are identified by the proposal as significant behavioral variables. As a result, UBA establishes a connection between the machine learning model's log-derived attributes and the behavior of individual users.

The machine learning engine's behavioral analytics and user behavior profiling components are supported by UBA within the conceptual model. Additionally, it provides the theoretical foundation for investigating device modifications, location discrepancies, session anomalies, and authentication anomalies. The research can take into account whether a transaction or authentication event is compatible with the established behavioral profile of the relevant user instead of treating each transaction separately. Instead, then restricting detection to traditional transaction-level fraud, this enhances the study's emphasis on identity theft.

2.2.4 CIA Triad as a Guiding Information Security Framework

A fundamental framework for information security, the Confidentiality, Integrity, and Availability (CIA) Triad consists of three main security goals: preventing unauthorized disclosure of information, maintaining the accuracy and reliability of information, and ensuring that systems and information remain being accessible to authorized users.

The CIA Triad provides an essential security foundation for our investigation even though it is not a behavioral detection theory in the same sense as anomaly detection. The study uses availability to guaranty that security monitoring continues to function, confidentiality to safeguard sensitive customer and log data, and integrity to preserve the accuracy and reliability of behavioral logs for timely detection.

As a result, the CIA Triad provides a security framework that directs the suggested Machine Learning detection process. Because banking logs may include sensitive data regarding clients, accounts, devices, locations, and authentication actions, confidentiality is crucial. Because the model may acquire inaccurate behavioral patterns or produce untrustworthy alerts as a result of faulty or manipulated logs, integrity is especially crucial. Availability is important because identity-theft detection may be less successful if monitoring is delayed or unavailable. As a result, the CIA viewpoint offers the security and governance framework for interpreting the independent variables, machine learning procedures, and dependent detection results.

2.2.5 Integrated Theoretical Foundation

The four viewpoints are not antagonistic, but rather complementary. UBA Theory discusses why individual user behavior and deviations from behavioral profiles are crucial for identity-centric identification; Machine Learning Theory describes how patterns can be learned and exploited computationally; and Anomaly identification Theory defines suspicious deviation. Banking logs are gathered, processed, safeguarded, and utilized within the broad information security framework provided by the CIA Triad.

Therefore, the relationship can be conceptualized as follows: machine learning provides computational mechanisms for classification and anomaly identification; UBA interprets these observations into meaningful behavioral profiles; anomaly detection detects deviations from expected patterns; and banking log data provides observations of user and system activity. The CIA Triad directs the safe processing and operational use of the underlying data, and the resulting warnings enhance the effectiveness of identity theft detection. According to the plan, the integrated foundation converts operational logs into behavioral indicators, analyzes them using machine learning algorithms, and uses anomaly detection mechanisms to interpret anomalies.

Thus, the integrated framework is in line with the conceptual variables of the study, where the dependent construct is identity theft detection effectiveness; behavioral profiling and model processing are part of the analytical mechanism; regulatory compliance, technological infrastructure, and staff expertise may influence implementation; and log data quality and machine learning capabilities constitute important independent constructs. In Kenyan commercial banking information systems, this theoretical integration offers a logical foundation for investigating whether behavioral data in banking logs may be converted into useful, timely, and understandable signs of identity theft.

Justification and relevance in the modern day. Despite having fundamental roots, Anomaly Detection Theory, Machine Learning Theory, User Behavior Analytics, and the CIA Triad are still appropriate explanatory lenses for this study because they describe long-lasting mechanisms: identity-centered behavior profiling, learning from data, deviations from expected activity, and protection of confidentiality, integrity, and availability. Their age does not prove theoretical inferiority on its own.The model is interpreted in accordance with recent empirical research on financial fraud detection and log anomaly identification in order to reflect the dynamic threat environment (Hilal et al., 2022; Hernandez Aros et al., 2024; Allawi Hussein & Répás, 2024; Chen et al., 2025). While the core theories provide an explanation of the dimensions and relationships under test, the contemporary literature updates the empirical context. Therefore, rather than being replaced by a trendy but less closely linked theory, the framework is kept with an explicit layer of modern data.

2.3 Conceptual Framework

The conceptual framework illustrates the hypothesised relationships between the independent variables, the dependent variable, and the intervening and moderating variables in the study titled “Machine Learning-Based Identity Theft Detection Using Financial System Logs: Design, Implementation, and Experimental Evaluation.”

The methodology suggests that machine learning capabilities, user behavior profiling, and the quality of financial system log data all have an impact on the efficacy of identity theft detection. Technical and organizational factors, such as staff expertise, infrastructure availability, and regulatory compliance, may also have an impact on these interactions.

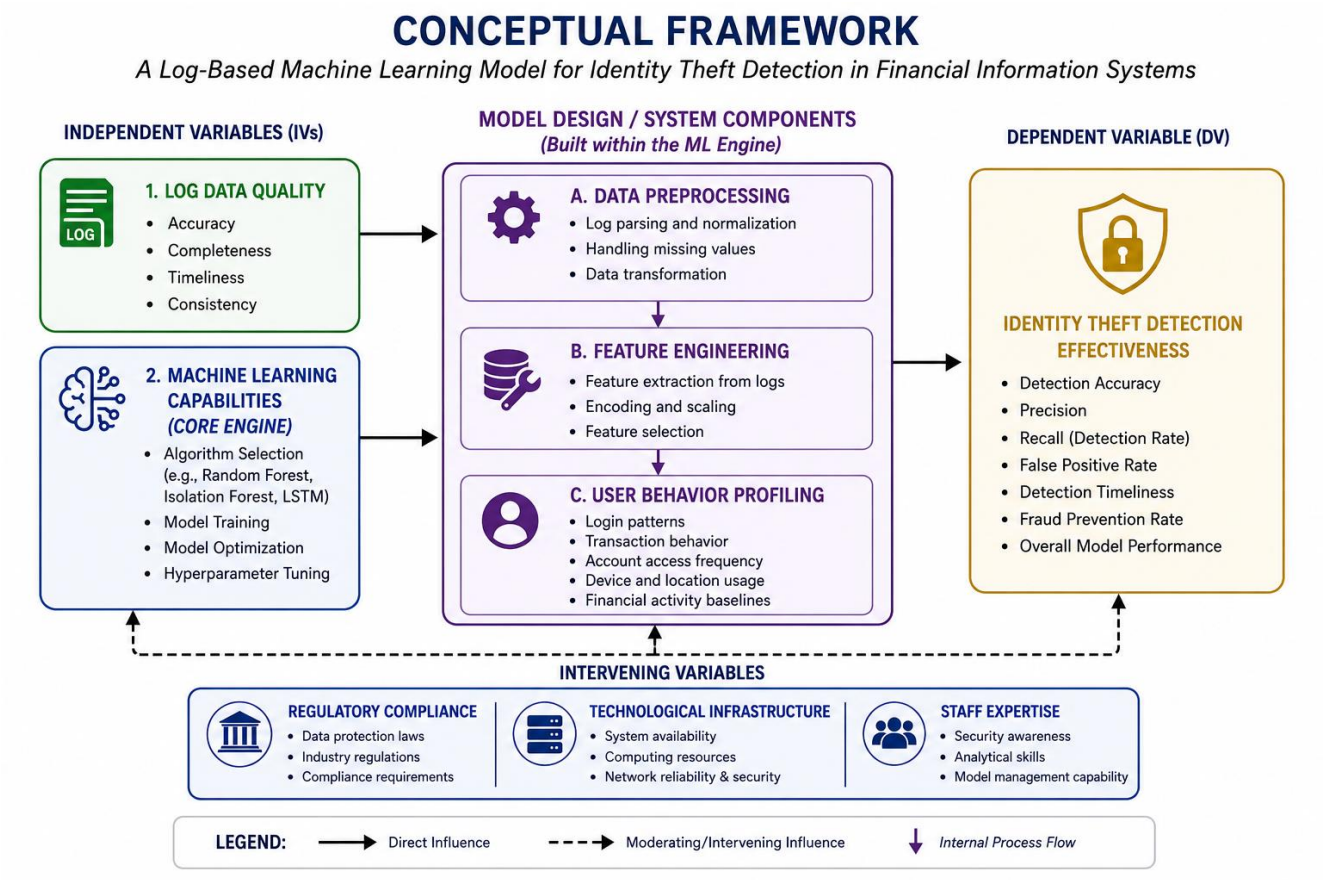


Figure 1: Conceptual Framework

Variable Type	Construct	Key Indicators
Independent (IV)	Log Data Quality	Accuracy, completeness, timeliness
Independent (IV)	ML Capabilities	Algorithms, model performance
Dependent (DV)	Detection Effectiveness	Accuracy, speed, reduced fraud
Intervening	Regulatory/Tech Factors	Compliance, infrastructure

Table 1: Alignment of variable type, Construct, and Key Indicators.

Summary of the Framework

Through data preparation and feature engineering, the methodology suggests that log data quality, machine learning capabilities, and user behavior profiling enhance the effectiveness of identity theft detection. Infrastructure, staff skills, and regulatory needs may all have an effect on the collaboration.

Methodological clarification: Preprocessing and feature engineering can potentially be characterized as phases of the model development pipeline rather than simply as intervening variables because your study is an experimental machine learning effort. In a similar vein, personnel knowledge, infrastructure, and compliance should only be considered moderators if their moderating effects are measured and tested in your study design.

2.4 Empirical Literature Review

2.4.1 Introduction

Empirical research on financial fraud and cybersecurity detection has increasingly shifted from conventional rule-based mechanisms towards machine learning (ML), behavioural analytics and automated anomaly detection. This change is driven by the growing volume, speed, and complexity of digital financial transactions as well as the challenge of identifying attacks that don't follow predetermined rules. The literature demonstrates that machine learning can identify pat-

terns associated with fraudulent activity, while log-based anomaly detection provides an additional source of evidence concerning user behaviour, authentication events and system activity.

The empirical literature relevant to the present study can broadly be grouped into four areas: financial Identity theft detection using machine learning; anomaly detection from system logs; streaming and real-time Identity theft detection; and explainable artificial intelligence (XAI) for financial decision-making. These areas provide complementary evidence for the proposed log-based identity theft detection model.

The distinction between financial Identity theft detection and identity theft detection is important. Much of the empirical literature concentrates on fraudulent transactions, whereas the present study focuses on the behavioural evidence surrounding potentially compromised identities. The proposed study therefore extends the literature by considering authentication logs, access records, session activities, device information, geolocation and transaction behaviour as interconnected indicators of possible identity compromise. This focus is consistent with the study scope specified in the proposal.

2.4.2 Empirical Evidence on Machine Learning for Financial Identity Theft Detection

One of the foundational empirical studies in financial Identity theft/Fraud detection was conducted by Bhattacharyya et al. (2011). The researchers compared logistic regression, Support Vector Machines (SVM) and Random Forest (RF) for credit-card Identity theft/fraud detection using real-life transaction data from an international credit-card operation. The study investigated the impact of varying degrees of undersampling on model performance, and the dataset included transactions from January 2006 to January 2007.

Because it offers actual proof that ensemble-based machine learning may be applied to highly biased financial fraud data, the paper is significant to current re-

search. Additionally, it shows how crucial model comparison and data preparation techniques are for assessing identity theft detection systems. Nevertheless, the study concentrated on credit card fraud at the transaction level rather than identity theft deduced from system logs and authentication. As a result, while its results validate the suggested use of Random Forest, they do not directly determine how Random Forest would function with behavioral indicators like device changes, login failures, session duration, or inconsistent geolocation.
Because identity theft can happen even in the absence of a fraudulent purchase, this distinction is very crucial. If an attacker has legitimate credentials, they can successfully authenticate and act in a way that first appears to be legitimate. Therefore, rather than depending solely on transaction-level classification, the current study suggests combining authentication and transaction evidence.

A broader synthesis was provided by Ali et al. (2022) through a systematic literature review of ML-based financial Identity Theft detection. Using the Kitchenham systematic review approach, the researchers searched major scholarly databases and ultimately analyzed 93 articles. Their review found that SVM and artificial neural networks were among the commonly used ML techniques and that credit-card fraud was the most frequently studied fraud category. The review also identified a strong concentration on supervised approaches and noted comparatively less use of unsupervised approaches such as clustering and anomaly detection.

This result directly relates to the suggested comparison between Random Forest and Isolation Forest. Isolation Forest can be used to find anomalous observations without requiring every suspicious behavior to be explicitly labeled, whereas Random Forest provides a supervised classification approach when labeled samples are provided. Thus, the literature offers a justification for merging or contrasting supervised and anomaly-detection techniques.

However, Ali et al. (2022) also highlights a significant gap in the broader literature: transaction fraud has received a lot of attention in financial fraud research.

As a result, there is comparably less empirical support for identity-centric behavioral detection with access logs and authentication. The restriction closely aligns with the empirical gap noted in the current proposal, which points out that a lot of previous research uses publicly available or generalized datasets and pays little attention to authentication and access behavior in commercial banking.

2.4.3 Empirical Evidence on System-Log Anomaly Detection

While financial fraud studies provide evidence concerning transaction-level machine learning, system-log research provides a more direct foundation for analysing behavioural activity.

DeepLog, a deep-learning technique for identifying and diagnosing anomalies from system logs, was developed by Du et al. (2017). In order to identify typical patterns of system behavior, the researchers utilized a Long Short-Term Memory (LSTM) neural network to describe log events as sequences. The method was created to facilitate online monitoring and detect deviations from taught sequences.

DeepLog is important for this work because it treats logs as behavioral sequences instead of just discrete records. Similar sequences can be seen in banking authentication events. For instance, a failed login may be followed by a successful authentication, a fresh device login, an odd session, and finally an anomalous transaction. Therefore, links that emerge when events are analyzed as behavioral sequences may be hidden when individual recordings are analyzed separately.

However, DeepLog is not so much concerned with identity theft in commercial banking as it is with system execution and operational log anomalies. Additionally, its deep-learning architecture is different from the suggested comparison between Random Forest and Isolation Forest. Therefore, its main methodological contribution to the current study is that it shows that system logs might offer useful behavioral information for automated anomaly identification.

A more recent empirical investigation was undertaken by Ali et al. (2025), who conducted a comprehensive empirical comparison of supervised and semi-supervised traditional and deep-learning techniques for log-based anomaly detection. The study evaluated techniques including Random Forest, SVM, LSTM, LogRobust, NeuralLog, One-Class SVM, DeepLog and Logs2Graphs across seven benchmark datasets. The researchers considered not only detection accuracy but also training/prediction time and sensitivity to hyperparameter tuning.

Importantly for the present study, Ali et al. (2025) found that traditional supervised ML techniques and deep-learning techniques performed similarly in detection accuracy and prediction time across most benchmark datasets. Traditional supervised approaches also demonstrated lower sensitivity to hyperparameter tuning, while Random Forest showed comparatively low sensitivity among the evaluated traditional techniques. The study further found that semi-supervised techniques performed worse in detection accuracy than supervised techniques within its experimental setting.

These findings strengthen the methodological justification for investigating Random Forest rather than assuming that increasingly complex deep-learning models will necessarily provide superior detection. They also support including computational considerations in model evaluation.

However, the datasets used by Ali et al. (2025) are general-purpose log anomaly benchmarks, not Kenyan commercial-bank authentication and transaction logs. Consequently, the results cannot automatically be generalised to identity theft detection in Kenyan banking systems. This reinforces the need for contextual validation using banking-relevant behavioural variables.

Similarly, Hussein and Répás (2024) examined machine-learning approaches for detecting anomalies in log files. Their work highlights the difficulties created by the dynamic structure, high volume and heterogeneous nature of log files and

discusses anomaly-detection approaches in relation to effectiveness, execution speed and applicability across different log types.

The evidence from these log-based studies supports the proposition that operational logs can be transformed into behavioural indicators for anomaly detection. At the same time, the literature demonstrates that the effectiveness of a detection algorithm depends on the characteristics of the underlying logs and the evaluation context.

2.4.4 Streaming, Real-Time and Highly Imbalanced Identity Theft Detection

Real-time detection represents another important dimension of the proposed study. Detection models must handle massive amounts of data while retaining acceptable detection performance since banking environments generate continuous streams of authentication and transaction events.

SCARFF, a scalable framework for streaming credit-card fraud detection, was developed by Carcillo et al. (2018). The system combines online feature engineering and machine learning with distributed data-processing tools. The work specifically addressed a number of real-world issues related to identity theft detection, such as verification latency, non-stationarity, and class imbalance. The system might function at scale while retaining efficiency and detection performance, according to experiments conducted on a sizable stream of actual credit card transactions.

The study is relevant to the current research since detecting identity theft in commercial banking also involves a significant number of continually generated incidents. It shows that the temporal and operational features of the data can affect detection performance, and that identity theft detection should not be seen as a static classification problem.

However, SCARFF is more concerned with fraud involving credit card transactions than identity theft based on authentication behavior. As a result, it does not specifically address indicators like unsuccessful login attempts, device swapping,

inconsistent source IP addresses, or strange session behavior. These variables are essential to the proposed research.

The class-imbalance issue is particularly important. In cybersecurity datasets, suspicious events are usually much less frequent than legitimate events. The user's proposal explicitly recognises class imbalance as a major methodological challenge and proposes evaluation using Precision, Recall, F1-score, ROC-AUC, PR-AUC and False Positive Rate rather than relying on accuracy alone.

This approach is consistent with the broader empirical literature. A model may achieve high overall accuracy by classifying most events as legitimate while still failing to identify the minority class of fraudulent or compromised activities. Consequently, minority-class detection and false-positive behaviour are central considerations for the evaluation of the proposed model.

2.4.5 Empirical Evidence on Explainable Artificial Intelligence

Accuracy alone is insufficient for many financial cybersecurity applications because analysts and institutional decision-makers may need to understand why an event was classified as suspicious.

A user-centered explainable artificial intelligence approach for financial Identity theft detection investigated the use of XAI to provide explanations for financial fraud predictions. The study combined an ensemble predictive model with an explainability framework based on Shapley values and explicitly considered the information requirements of different stakeholders. Its objective was not simply to generate predictions but to provide explanations that could support stakeholder interpretation of those predictions.

This evidence is relevant to the present study's use of SHAP. The proposal specifically identifies explainability as an important component because financial institutions operate in regulated environments in which automated decisions may need to be interpreted and reviewed.

A broader systematic review by Weber, Carl and Hinz examined applications of explainable artificial intelligence in finance across finance, information systems and computer-science literature. The review demonstrates the growing importance of XAI in financial applications and provides evidence that explainability has become a significant research concern rather than an optional feature of financial ML systems.

However, the XAI literature is broader than identity theft and does not necessarily address authentication-log behaviour. Thus, the current study contributes by applying explainability specifically to behavioural indicators derived from financial-system logs.

2.4.6 Behavioural Analytics and Identity-Centric Detection

User Behaviour Analytics (UBA) provides an important bridge between general anomaly detection and identity theft detection. The central premise is that legitimate users tend to demonstrate relatively stable behavioural patterns over time, while compromised accounts may exhibit deviations in authentication, access, session, device or transaction behaviour.

The user's proposal explicitly defines UBA around authentication behaviour, access frequency, session activity, geolocation and device usage, with applications including compromised credentials and account takeover detection.

This perspective changes the detection problem from asking simply whether a transaction is fraudulent to asking whether the person or entity generating the activity behaves consistently with the established identity profile.

Relevant behavioural variables in the proposed study include:

- timestamp and login-time patterns;
- failed-login attempts;
- authentication success/failure;
- login frequency;

- source IP;
- device fingerprint;
- geolocation;
- session duration;
- transaction amount;
- transaction frequency;
- access patterns.

These variables allow the detection model to combine authentication-level and transaction-level evidence. This is particularly important in account takeover situations because an attacker may possess legitimate credentials and therefore successfully pass conventional authentication mechanisms.

The existing proposal consequently identifies an empirical and methodological gap: much of the existing fraud literature concentrates on transaction-level fraud while giving comparatively less attention to authentication behaviour, session anomalies, device usage, geolocation inconsistencies and explainable behavioural analytics.

2.4.7 Relevance to Kenyan Commercial Banking

The geographical and institutional context is a major consideration in interpreting the empirical literature. Most of the studies reviewed above use international credit-card datasets, public cybersecurity benchmarks or general system logs. They therefore provide useful methodological evidence but do not directly establish the performance of ML identity-theft detection within Kenyan commercial banks.

The proposed study specifically addresses this contextual limitation by focusing on Kenyan commercial banking information systems. Its scope includes authentication logs, transaction logs, access records, session activities and user-interaction patterns.

The proposal identifies several contextual gaps in the literature, including limited empirical validation using data relevant to Kenyan banking, limited behavioural log analytics based on authentication and access behaviour, and limited integration of explainability, false-positive reduction, class-imbalance handling and temporal anomaly validation.

Therefore, the contribution of the current research is not simply to apply an existing ML algorithm to another dataset. Rather, it is to investigate whether a log-based, identity-centric behavioural model can provide useful evidence for identity theft detection within the operational and contextual conditions of Kenyan commercial banking.

2.4.8 Synthesis of Empirical Literature

Recent literature reinforces the need to evaluate fraud-detection systems against changing data conditions rather than assuming that a model remains effective after initial training. Reviews of financial fraud detection published in 2022–2025 identify persistent challenges involving severe class imbalance, feature quality, generalisation across datasets, interpretability and the operational cost of false alerts (Ali et al., 2022; Hilal et al., 2022; Hernandez Aros et al., 2024; Chen et al., 2025; Compagnino, 2025). Recent work on log anomaly detection likewise treats system events as behavioural evidence and highlights the need to assess detection methods under realistic operational conditions (Allawi Hussein & Répás, 2024). These studies are current enough to contextualise the present work, but most do not establish performance for Kenyan commercial-bank identity theft using a combined set of authentication, session, device and transaction features.

The empirical literature establishes several important conclusions.

First, machine learning is demonstrably applicable to financial Identity theft /fraud detection. Bhattacharyya et al. (2011) demonstrated the use of RF, SVM and logistic regression on real financial transaction data, while Ali et al. (2022)

showed that ML-based financial Identity theft detection has become a substantial research area.

Second, logs contain useful behavioural information for anomaly detection. DeepLog demonstrated that system-log sequences can be modelled to identify deviations from learned normal behaviour, while more recent comparative research has demonstrated that traditional ML techniques such as Random Forest remain competitive with more complex deep-learning approaches in log anomaly detection.

Third, real-time and streaming detection creates additional methodological challenges. SCARFF demonstrates the importance of scalability, non-stationarity, class imbalance and verification latency when Identity Theft detection operates on continuous data streams.

Fourth, explainability is increasingly important in financial ML. XAI research demonstrates that predictive performance can be complemented by explanations designed around stakeholder requirements, including explanations based on Shapley values.

Finally, there remains a contextual and empirical gap concerning identity-centric detection in Kenyan commercial banking. The existing evidence is dominated by transaction fraud, general system-log anomaly detection and non-Kenyan datasets. The user's proposal therefore appropriately positions the research around the integration of authentication logs, transaction logs, behavioural profiling, anomaly detection, ML and XAI within Kenyan commercial banking.

Empirical Literature Matrix

Author /Year	Purpose	Methodology	Dataset/Population	Key Findings	Limitations/Gaps Relative to Present Study	Relevance to Objectives 1–3
Bhattacharyya et al. (2011)	To compare machine-learning/data-mining methods for credit-card fraud detection.	Comparative experimental evaluation of Logistic Regression, SVM and Random Forest; investigated effects of undersampling.	Real-life transaction data from an international credit-card operation, covering January 2006–January 2007.	Demonstrated that ML techniques can be applied to real financial fraud data and compared performance under different fraud-class proportions.	Focused on transaction fraud rather than identity-centric behavioural detection; did not use authentication, device, session or geolocation logs; not Kenyan banking data.	**O1:** establishes transaction-level behavioural evidence. **O2:** supports ML-based model design. **O3:** directly supports comparative model evaluation and RF investigation.
Du et al. (2017)	To detect and diagnose anomalies in	DeepLog uses an LSTM neural network to model log events	System-log data from large computing environm	Demonstrated that sequential system logs can	Focuses on system execution anomalies rather than banking	**O1:** supports identification of behavioural anomalies in logs. **O2:** supports log-

	system logs automatically.	as sequences and identify deviations from normal sequences.	ents.	be used for online anomaly detection and diagnosis.	identity theft; uses deep learning rather than RF/IF; does not specifically address financial authentication behaviour.	based behavioural model design. **O3:** provides a benchmark for comparing alternative log-ML approaches.
Ahmed, Mahmood & Hu (2016)	To synthesize and classify techniques used for network anomaly detection.	Systematic survey/taxonomy of network anomaly-detection techniques and evaluation approaches.	Existing network-anomaly research and publicly available intrusion-detection datasets.	Developed a taxonomy connecting anomaly types, network attacks and detection techniques and reviewed evaluation practices.	Network anomaly detection is broader than identity theft; limited focus on financial authentication logs, behavioural identity profiles and Kenyan banking.	**O1:** informs classification of anomaly patterns. **O2:** supports selection of anomaly-detection approaches. **O3:** informs evaluation considerations.
Carcillo et al. (2018)	To develop a scalable	SCARFF framework combinin	Large stream of real credit-	Demonstrated scalability,	Transaction-focused; not	**O1:** supports analysis of temporal/transaction

	framework for near-real-time credit-card fraud detection.	g distributed processing, online feature engineering and ML; explicitly addresses imbalance, non-stationarity and feedback latency.	card transactions.	efficiency and detection performance on streaming financial data while addressing class imbalance and non-stationarity.	authentication/log-based identity theft; infrastructure differs from Kenyan commercial banking environments.	anomalies. **O2:** supports scalable log-processing architecture. **O3:** informs detection-latency and imbalance evaluation.
Ali et al. (2022)	To systematically review ML approaches used in financial fraud detection.	Kitchenham-based systematic literature review; database searching, screening, quality assessment and synthesis.	93 selected studies from major scholarly databases.	SVM and ANN were common approaches; credit-card fraud was the most frequently studied category; unsupervised approac	Secondary rather than primary empirical evidence; literature dominated by transaction fraud; limited direct evidence on authentication-log identity theft and Kenyan banking.	**O1:** identifies dominant fraud-detection mechanisms. **O2:** supports selection of supervised/unsupervised ML. **O3:** identifies commonly used evaluation approaches and research gaps.

				hes receive d compar atively less attentio n.		
Hussei n & Répás (2024)	To examine contemp orary machine - learning approac hes for anomaly detectio n in log files.	Review/a ssessment of ML-based log anomaly-detection strategies, consideri ng effectiven ess, execution speed and applicabil ity.	Different types of computer log files.	Highlig hts the value of logs for detectin g questio nable or malicio us activity and identifie s high volume, dynami c structur e and heterog eneous logs as key challen ges.	Not specific to financial identity theft or Kenyan commerci al banking; limited identity-centric behaviour al context.	**O1:** directly supports investigation of log-based anomalies. **O2:** informs preprocessin g and feature engineering. **O3:** supports consideratio n of effectiveness and execution time.
Ali et al. (2025)	To compreh ensively compare tradition al and deep-	Empirical comparis on of supervise d traditiona l/deep	Seven publicly available benchmar k log datasets, including	Traditio nal supervi sed and deep ML showed	Benchmar k datasets are not Kenyan financial logs; identity	**O1:** demonstrates value of log-based anomaly analysis. **O2:** strongly

	learning approaches for log-based anomaly detection.	ML and semi-supervised approaches across seven benchmark datasets; evaluated accuracy, time performance and hyperparameter sensitivity.	HDFS, Hadoop and other log datasets.	similar detection accuracy on most datasets; traditional ML was less sensitive to hyperparameter tuning; Random Forest showed comparatively low sensitivity; semi-supervised approaches performed worse in detection accuracy in the tested settings.	theft and banking-specific behavioural indicators are not the focus; Isolation Forest is not the principal comparison.	supports RF-based model design. **O3:** provides empirical justification for evaluating accuracy alongside computational performance.

Financial-fraud XAI study (2023)	To develop financial Identity theft detection that combines predictive accuracy with stakeholder-oriented explanations.	Ensemble predictive model combined with an explainability framework based on Shapley values.	Financial fraud data; stakeholder requirements were incorporated into explanation design.	Demonstrated an approach for providing predictions together with explanations tailored to different stakeholders.	Focuses on financial fraud generally rather than identity theft from authentication logs; does not establish XAI requirements specific to Kenyan banking.	**O1:** helps identify interpretable behavioural indicators. **O2:** directly supports SHAP/XAI integration. **O3:** supports evaluating not only predictive performance but interpretability.

Table 2: Empirical Literature Matrix

Objective Alignment Summary

Study Objective	Main Empirical Evidence	Variables/Methods Supported	Remaining Gap Addressed by Present Study
Objective 1: Assess existing identity-theft detection mechanisms and behavioural anomaly patterns in authentication and transaction logs.	Bhattacharyya et al. (2011); Du et al. (2017); Ahmed et al. (2016); Hussein & Répás (2024); Ali et al. (2025).	Transaction behaviour, log sequences, anomaly detection, authentication/access behaviour, temporal patterns and behavioural deviations.	Existing research largely separates transaction fraud from system-log anomaly detection. Limited empirical work integrates authentication, session, device, geolocation and transaction behaviour specifically for identity theft in Kenyan commercial

			banks.
Objective 2: Design and implement a log-based ML model using behavioural analytics and XAI.	Du et al. (2017); Carcillo et al. (2018); Ali et al. (2022); Ali et al. (2025); financial-fraud XAI study (2023).	Log processing, feature engineering, behavioural profiling, Random Forest, anomaly detection, streaming processing and SHAP-based explanation.	Existing studies generally focus on either transaction fraud, generic log anomalies or financial XAI. Few integrate all three dimensions—**banking logs + identity-centric behavioural analytics + explainable ML**—in the Kenyan context.
Objective 3: Evaluate the proposed model and compare Random Forest and Isolation Forest using Precision, Recall, F1, ROC-AUC, PR-AUC, and FPR	Bhattacharyya et al. (2011); Carcillo et al. (2018); Ali et al. (2022); Ali et al. (2025).	Comparative ML evaluation, class imbalance, detection accuracy, false positives, temporal performance, computational performance and model sensitivity.	Most previous studies use different datasets, algorithms or metrics and do not provide a direct RF-versus-IF comparison for identity theft using Kenyan banking logs. The present study therefore fills a methodological and contextual evaluation gap.

Table 3: Objective Alignment Summary

Overall Empirical Gap and Positioning of the Study

The empirical literature reveals a three-part gap.

First, the financial-fraud literature has established the usefulness of machine learning but has predominantly concentrated on transaction-level fraud. For example, Bhattacharyya et al. (2011) evaluated RF, SVM and logistic regression using real credit-card transactions, while Ali et al. (2022) found credit-card fraud to be the most frequently represented fraud category in the reviewed ML literature.

Second, the log-anomaly literature demonstrates that system logs contain valuable behavioural information, but much of this research addresses general system or software anomalies rather than identity theft in financial systems. DeepLog and subsequent comparative log-anomaly studies demonstrate the feasibility of learning behavioural patterns from logs, but their datasets and research contexts differ from Kenyan commercial banking.

Third, XAI research demonstrates the importance of making financial ML decisions interpretable, but the evidence is not sufficiently specific to identity-centric behavioural analytics based on banking authentication and transaction logs.

These gaps correspond closely to the research gap already established in the proposal: limited research integrating identity-centric behavioural analytics with authentication and access logs, limited validation using datasets operationally relevant to Kenyan banking, and insufficient attention to authentication behaviour, session anomalies, geolocation, device usage, explainability, class imbalance, false-positive reduction and real-time adaptation.

Accordingly, the present study is positioned to contribute an identity-centric, log-based machine-learning approach that integrates authentication logs, transaction logs and behavioural indicators; compares supervised and anomaly-detection

approaches through Random Forest and Isolation Forest; incorporates explainability; and evaluates performance using multiple metrics appropriate for highly imbalanced cybersecurity data. This positioning directly corresponds to the three specific objectives stated in the proposal.

Recommended Chapter Two transition

The empirical review therefore establishes the basis for the conceptual framework and research methodology. The literature indicates that effective identity-theft detection requires moving beyond isolated transaction classification towards integrated analysis of authentication, access, session, device, temporal and transaction behaviour. It also indicates that model effectiveness should be assessed using metrics capable of capturing minority-class detection, false-positive behaviour and detection timeliness rather than accuracy alone. The next section consequently presents the conceptual framework linking log-data quality and machine-learning capabilities to identity-theft detection effectiveness, while considering regulatory, technological and organisational factors that may influence implementation. The proposed framework identifies log-data quality and ML capabilities as key independent variables and detection accuracy, false-positive rate and detection speed as indicators of detection effectiveness.

2.5 Critique of Existing Literature

The empirical literature reviewed demonstrates that machine learning, anomaly detection and behavioural analytics have considerable potential for identifying fraudulent and suspicious activity in digital financial systems. However, the literature also reveals methodological, contextual and theoretical limitations that constrain its direct application to identity theft detection in Kenyan commercial banking information systems.

2.5.1 Methodological Gaps

A major methodological limitation is the dominance of transaction-level Identity theft detection. Much of the financial fraud literature evaluates models using transaction datasets, particularly credit-card fraud datasets. For example, systematic reviews have found credit-card fraud to be one of the most extensively studied financial fraud categories. After analyzing 93 published articles, Ali et al. (2022) concluded that credit-card fraud was the most frequently investigated type of fraud and that Support Vector Machines and Artificial Neural Networks were among the commonly employed strategies. In the same way, a more recent analysis of 104 papers released between 2012 and 2023 revealed a persistent focus on datasets related to credit card theft.

Identity theft is not fully captured by transaction-level modeling, despite its usefulness. Because the attacker has valid credentials, a compromised user account may initially produce transactions that appear to be legitimate. As a result, the behavioral evidence related to account compromise may not be captured by depending just on transaction characteristics.

In order to overcome this limitation, the current study integrates session, device, geolocation, authentication, access, and transaction data. Timestamp, user ID, source IP, device fingerprint, geolocation, session duration, unsuccessful login attempts, transaction value, transaction frequency, access patterns, and authentication outcomes are all identified as pertinent log elements in the proposed study.

A second methodological issue is the difference between supervised classification and unsupervised anomaly detection. While anomaly-detection research has independently examined unsupervised methods, financial fraud research has thoroughly examined supervised classification. By comparing Random Forest with Isolation Forest, the current study unifies these methods. The viability of comparing supervised Random Forest and unsupervised Isolation Forest in a typical

anomaly-detection scenario has also been shown by recent practical studies that directly compare both algorithms on server logs.

A third drawback relates to the evaluation process and class disparity. Typically, only a small percentage of financial and cybersecurity observations are fraudulent. As a result, accuracy by itself may give an inaccurate picture of the efficacy of a model. This methodological issue is addressed by the study's suggested usage of Precision, Recall, F1-score, ROC-AUC, PR-AUC, False Positive Rate, and detection latency. These metrics offer further data regarding operational timeliness, false alarms, and minority-class identification.

A fourth methodological limitation relates to temporal and real-time behavior. Financial identity theft detection can be significantly impacted by non-stationarity, verification latency, and class imbalance, according to streaming fraud research. Instead of approaching identity theft detection as a strictly static classification task, Carcillo et al.'s work on streaming credit-card fraud detection highlights the need of processing continuous data. By taking detection delay and temporal behavioral patterns into account, the current study expands this concern to identity-related behavioral occurrences.

Finally, many Machinel Learning detection experiments do not adequately incorporate explainability. Without providing a sufficient explanation of the behavioral traits that led to the alert, a model might detect an anomaly. In regulated financial environments, where security analysts may need to look into and defend alarms, this is especially problematic. Therefore, instead of treating prediction and interpretation as entirely distinct processes, the current study integrates explainable AI into the suggested detection architecture.

2.5.2 Contextual Gaps

A significant contextual limitation is that much of the available empirical evidence originates from non-Kenyan and non-African datasets. Existing studies frequently use international credit-card data, publicly available cybersecurity benchmarks, simulated datasets or general system logs. While these datasets are valuable for methodological experimentation, their findings cannot automatically be assumed to represent behavioural patterns within Kenyan commercial banking systems.

The user's study is specifically situated within Kenyan commercial banks operating digital banking platforms. Its scope focuses on authentication logs, transaction logs, access records, session activities and user-interaction patterns.

The contextual limitation is therefore not simply the absence of Kenyan data; it is the limited understanding of whether behavioural indicators identified in international research remain useful when applied within the operational environment, technological infrastructure and behavioural patterns of Kenyan banking customers.

The present research addresses this limitation by developing and evaluating the model within the Kenyan commercial-banking context, using anonymised operational logs where accessible and relevant cybersecurity or financial datasets where supplementary data are required. The proposal identifies candidate datasets including the CERT Insider Threat Dataset, IEEE-CIS Identity theft/Fraud Detection Dataset, CICIDS datasets and simulated banking authentication data.

A second contextual limitation concerns the difference between generic cybersecurity anomaly detection and identity theft detection. System-log studies have demonstrated that logs can be used to identify abnormal sequences and activities, but the objective is often system reliability, intrusion detection or

general anomaly detection rather than identification of compromised financial identities.

The current study therefore applies the log-analysis concept to an identity-centric problem. This means that the unit of analysis is not merely whether a system event is abnormal, but whether a collection of events is inconsistent with the established behavioural profile of a legitimate user.

2.5.3 Theoretical Gaps

A theoretical gap is also evident in the literature. While log anomaly research typically focuses on departures from typical system behavior, existing financial fraud research often focuses on the prediction power of individual algorithms. Identity theft in banking information systems cannot be fully explained by either viewpoint on its own.

Thus, the current work incorporates User Behavior Analytics (UBA) Theory, Machine Learning Theory, and Anomaly Detection Theory.

The framework for identifying deviations from predetermined behavioral baselines is provided by anomaly detection theory. The ability of supervised and unsupervised algorithms to identify patterns in past data is explained by machine learning theory. By concentrating on authentication behavior, session activities, access frequency, device usage, and geolocation, UBA offers the identity-centered viewpoint. These three theoretical stances are specifically identified in the proposal as complementary model foundations.

This theoretical integration addresses a limitation in research that views identity theft detection mainly as transaction classification. It enables the proposed study to conceptualize identity theft as a behavioral deviation problem, where compromised credentials can be identified by combining actions that would otherwise be considered valid.

2.5.4 Explainability and Operational Gaps

The transition from experimental machine learning models to real-world cyber-security settings is another significant constraint. Even if a model performs well in an experimental dataset, it may still produce a large number of false positives, be difficult to interpret, or consume computational resources that are challenging to maintain in a real-world banking setting.

Explainability, data imbalance, benchmarking consistency, and real-time scalability have become more widely recognized in the literature as persistent challenges in AI-driven financial identity theft detection.

In response, the current study evaluates false-positive behavior, incorporates explainable AI, and takes detection delay into account in addition to traditional classification metrics. The conceptual framework also acknowledges human expertise, regulatory compliance, and technology infrastructure as intervening elements that could affect execution.

Therefore, developing a classification model is not the only goal of the suggested contribution. It aims to determine whether behavioral log analytics may yield quantifiable and operationally interpretable evidence of anomalies related to identity.

2.5.5 How the Present Study Addresses the Critique

The literature critique can therefore be summarised as follows:

1. **Transaction-centric literature** is extended through authentication and behavioural log analysis.
2. **General anomaly-detection research** is adapted to an identity-centric financial-security problem.

3. **Supervised and unsupervised approaches** are compared through Random Forest and Isolation Forest.
4. **Class imbalance** is explicitly addressed through appropriate preprocessing and evaluation metrics.
5. **Accuracy-only evaluation** is avoided through Precision, Recall, F1-score, ROC-AUC, PR-AUC, and False Positive Rate
6. **Limited explainability** is addressed through explainable AI techniques.
7. **International/general datasets** are supplemented or replaced, where possible, with data relevant to Kenyan commercial banking.
8. **Static detection perspectives** are extended through consideration of temporal and behavioural patterns.

These responses correspond directly to the gaps identified in the study proposal, particularly the limited attention given to authentication behaviour, session anomalies, geolocation, device usage, explainable AI, class imbalance, false-positive reduction and real-time behavioural adaptation.

2.6 Research Gaps

The literature reviewed reveals four connected gaps. First, a contextual gap remains because results from general credit-card fraud or network intrusion datasets cannot automatically be generalised to identity theft in Kenyan commercial banking. Second, a behavioural integration gap remains because many studies focus on transactions or system events separately rather than combining authentication, session, device, temporal and transaction indicators. Third, an evaluation gap remains because accuracy alone can obscure poor minority-class detection in datasets where suspicious events are rare; precision, recall, F1-score, ROC-AUC, PR-AUC and false-positive rate should be interpreted together. Fourth, an operational adaptation gap remains: model drift, adversarial behaviour, alert triage and response to threats outside the model's target class require explicit consideration. This study addresses these gaps through an experimental comparison of Random Forest and Isolation Forest on the available dataset, behavioural feature analysis, imbalance handling and explainability.

Because the dataset is public and not confirmed to originate from Kenyan banks, the study provides a methodological contribution relevant to the Kenyan context rather than direct proof of production performance in Kenyan institutions.

Based on the empirical and theoretical literature reviewed, the present study identifies the following specific research gaps.

2.6.1 Empirical Gap

There is limited empirical evidence on identity theft detection using behavioral information extracted from authentication and access logs within Kenyan commercial banking information systems. Existing financial-fraud studies predominantly examine transaction-level fraud, while log-anomaly studies frequently focus on general cybersecurity or system anomalies. The proposal consequently identifies limited research combining authentication and access behavior with financial transaction behavior.

The present study addresses this gap by analyzing behavioral information from authentication logs, transaction logs, access records and session activities within the Kenyan commercial-banking context.

2.6.2 Contextual Gap

Much of the existing empirical evidence relies on international, simulated or general-purpose cybersecurity datasets. Such datasets provide useful methodological evidence but may not fully capture the operational characteristics of Kenyan commercial banking.

The present study addresses this contextual gap by focusing specifically on Kenyan commercial banks and by seeking to use anonymized operational logs where accessible, supplemented where necessary by relevant public or simulated datasets.

2.6.3 Behavioral Analytics Gap

Existing research gives substantial attention to transaction characteristics but comparatively limited attention to the combined behavioral indicators associated with identity compromise. The proposal specifically identifies gaps concerning authentication behavior, session anomalies, geolocation inconsistencies and device usage patterns.

The proposed study addresses this gap by constructing behavioral features from authentication, temporal, device, location, session and transaction data.

2.6.4 Algorithmic Gap

The literature contains substantial evidence for supervised financial Identity theft detection and separate evidence for unsupervised anomaly detection. However, there is limited direct evidence concerning the comparative use of Random Forest and Isolation Forest for identity theft detection using banking-system logs.

The study addresses this gap by implementing both algorithms and evaluating their performance using a common experimental framework.

2.6.5 Evaluation Gap

Existing studies do not consistently use the same evaluation criteria, and accuracy can be particularly misleading when fraud constitutes a very small proportion of observations. The proposed research therefore addresses the evaluation gap through a multidimensional evaluation framework comprising:

- Precision;
- Recall;
- F1-score;
- ROC-AUC;
- PR-AUC;
- False Positive Rate; and

These measures correspond directly to the third specific objective of the study.

2.6.6 Explainability Gap

The literature demonstrates increasing interest in explainable financial AI, but relatively little research integrates explainability directly into identity-centric log-based detection. The proposal identifies explainable AI as one of the areas requiring further methodological attention.

The present study addresses this gap by incorporating explainable AI into the proposed behavioural detection model so that suspicious classifications can be associated with identifiable behavioural characteristics.

2.6.7 Real-Time and Adaptive Detection Gap

Financial-system behaviour changes over time, while fraud techniques also evolve. Static behavioural baselines can therefore become less effective as legitimate user behaviour changes. The proposal identifies limited attention to temporal anomaly validation and real-time behavioural adaptation.

The proposed research addresses this issue by considering temporal behavioural features and detection delay and by evaluating the practical implications of using log-based monitoring for continuous detection.

2.6.8 Integrated Research Gap

The most important gap is therefore integration. Existing studies tend to address different components independently:

financial transaction fraud → machine learning;
system logs → anomaly detection;
behavioral analytics → user monitoring;
XAI → model interpretation.

The present study integrates these components into a single identity-theft detection framework for Kenyan commercial banking. This corresponds directly to the stated study contribution of integrating behavioral analytics, anomaly detection, explainable AI and identity-centric monitoring.

2.7 Summary of Literature Review

This chapter's analysis of the literature shows that machine learning has increased in significance as a method for identifying cybersecurity irregularities and financial fraud. According to empirical research, unsupervised anomaly-detection techniques offer a way to spot behavioral abnormalities that were previously undetected, whereas supervised algorithms can spot patterns linked to known fraudulent activities. Further evidence that operational logs can offer helpful information for identifying anomalous sequences and activities comes from system-log research.

The review also shows that identity theft detection and financial identity theft detection are connected but different issues. While identity theft may show up as changes in authentication behavior, access patterns, session characteristics, device usage, geolocation, and transaction behavior, financial fraud research has mostly focused on transaction-level statistics. An identity-centric behavioral approach is strongly justified by this distinction.
A number of methodological issues are also identified in the literature, such as significant class imbalance, false positives, shifting behavioral patterns, scalability, detection delay, and poor explainability. These challenges justify the proposed use of multiple evaluation metrics rather than accuracy alone and support the comparison of supervised Random Forest and unsupervised Isolation Forest.

Contextually, the literature remains limited in its direct application to Kenyan commercial banking information systems. The proposal identifies a shortage of empirical research using authentication and access behaviour within Kenyan commercial banking environments and notes that much existing research relies on generalized or publicly available datasets.

The review consequently establishes the need for a model that combines log data, behavioural analytics, machine learning, anomaly detection and explainable AI. The proposed study responds to this need by analyzing authentication and transaction logs, developing behavioral features, implementing Random Forest and Isolation Forest, incorporating explainability, and evaluating detection effectiveness using Precision, Recall, F1-score, ROC-AUC, PR-AUCand False Positive Rate. These activities correspond directly to the three specific objectives of the study.

Overall, the literature provides sufficient empirical support for investigating machine-learning-based identity theft detection, but it does not adequately resolve the contextual, behavioral and methodological problem of detecting identity compromise from integrated banking logs in Kenyan commercial banks. This unresolved problem constitutes the principal research gap that the present study seeks to address.

CHAPTER THREE: RESEARCH METHODOLOGY

3.1 Introduction

This chapter presents the methodology adopted for the design, implementation and evaluation of a log-based machine learning model for identity theft detection in Kenyan commercial banking information systems. The study adopts a Design Science Research (DSR) approach supported by a post-positivist research philosophy and quantitative analysis.

The methodology is designed to address the three specific objectives through analysis of financial-system logs, behavioural feature engineering, development of a machine-learning detection artifact, and empirical evaluation of its performance. The chapter describes the research philosophy, research design, study area, target population, sampling design, data collection instruments and procedures, pilot study, validity and reliability, data analysis and presentation, and ethical considerations.

The methodology is aligned with the study's focus on authentication, transaction, access, session and device-related behavioural indicators and the evaluation of identity-theft detection effectiveness using Precision, Recall, F1-score, ROC-AUC, PR-AUC and False Positive Rate (FPR.

3.2 Research Philosophy

The study adopts a post-positivist research philosophy. Post-positivism recognizes that an objective reality exists but acknowledges that knowledge about this reality is subject to uncertainty and should therefore be established through systematic observation, measurement, testing and empirical evidence.

This philosophy is appropriate because the study seeks to develop and empirically evaluate a technological artifact using measurable evidence derived from financial-system logs. The study focuses on observable behavioural characteristics within authentication, access, session and transaction records. It is possible to

convert variables into quantifiable features for computational analysis, including unsuccessful login attempts, login frequency, transaction amount, session duration, device changes, source IP address, geolocation, and authentication outcomes.

The developed model is evaluated using objective performance indicators, including Precision, Recall, F1-score, ROC-AUC, PR-AUC and False Positive Rate. The use of several complementary measures is important because machine-learning predictions are probabilistic and identity-theft datasets may exhibit severe class imbalance. Consequently, accuracy alone may not adequately represent the effectiveness of the detection model.

Post-positivism therefore provides an appropriate philosophical foundation because the study emphasises empirical evidence, systematic measurement, reproducibility and critical evaluation of the proposed technological artifact.

3.3 Research Design

The study adopts Design Science Research (DSR) as the primary research design. DSR is appropriate because the principal purpose of the study is to design, implement, demonstrate and evaluate a technological artifact that addresses a practical cybersecurity problem.

The artifact developed in this study is a Log-Based Machine Learning Identity Theft Detection Model (LML-ITDM). The DSR approach provides a systematic process through which the researcher identifies the cybersecurity problem, defines artifact requirements, designs and develops the solution, demonstrates its application using relevant financial-system data, evaluates its performance, and communicates the findings.

The DSR process adopted in the study consists of six major stages:

1. **Problem identification** – identification of the identity-theft detection problem in financial information systems.
2. **Definition of artifact objectives** – specification of the requirements that the proposed detection model should satisfy.
3. **Design and development** – development of the log-processing, behavioural-feature and machine-learning components.
4. **Demonstration** – application of the developed artifact to relevant financial-system log data.
5. **Evaluation** – assessment of the model using predefined performance metrics.
6. **Communication** – presentation and interpretation of the results in relation to the research objectives and existing literature.

The design directly corresponds to the three specific objectives. Objective One involves examining existing identity-theft detection mechanisms and behavioural anomaly patterns; Objective Two involves designing and implementing the machine-learning artifact; and Objective Three involves evaluating and comparatively assessing the performance of the selected algorithms.

3.4 Study Area

The study focuses on commercial banks that run digital banking systems and is based inside Kenya's commercial banking industry. Kenya is the geographical focus, and the contextual environment is the financial information systems that underpin online banking services.

Commercial banking systems are appropriate for this study because they generate authentication, transaction, access, session and network activity records that can provide behavioural indicators for identity-theft detection. Digital banking channels also provide an appropriate context for examining suspicious activities associated with compromised credentials, unusual access patterns, device changes and abnormal transaction behaviour.

The study is therefore positioned within the Kenyan cybersecurity, financial-information-system and digital-banking context. Although the intended application context is Kenyan commercial banking, the experimental model uses the available financial-system dataset because direct access to proprietary banking logs is constrained by privacy, security and regulatory requirements. The thesis specifically identifies the use of the Financial System Datasets Logs for Identity-Theft Detection as a mitigation for the limited accessibility of real financial-system logs.

3.5 Target Population

The conceptual target population comprises **Kenyan commercial banking information systems and their associated authentication, access, session and transaction activities**. The unit of technical analysis is an individual financial-system transaction or log observation represented by the available dataset.

The study is concerned with behavioural indicators that may provide evidence of identity-related suspicious activity. These include:

- authentication success and failure;
- failed login attempts;
- login frequency;
- transaction amount;
- transaction frequency;
- session duration;
- device changes;
- source IP information;
- geolocation;
- access patterns; and
- temporal characteristics of system activity.

The broader study context is Kenyan commercial banks because the research seeks to develop a model relevant to identity-theft detection within financial

information systems. The thesis specifies that the study focuses on commercial banking environments and the behavioural analysis of authentication logs, transaction logs, access records, session activities and user interaction patterns.

3.6 Sampling Design

3.6.1 Sample Size

Because the computational component of the study uses a large financial-system dataset rather than a conventional human-participant sample, the study does not determine a sample size using a survey formula such as Yamane's formula.

The computational analysis uses the selected publicly available synthetic financial-transaction dataset. It contains 6,362,620 observations, including 6,354,407 records labelled non-fraudulent and 8,213 records labelled fraudulent. These labels identify simulated transaction fraud; they do not independently verify that a real person's identity was stolen. Accordingly, the model's target is treated as a fraud-labelled proxy relevant to identity-theft risk, not as ground-truth identity theft.

Rather than reducing the dataset to a small sample, the study uses the available observations for model development and evaluation. An **80:20** train-test split is applied, with the training data used for model development and the test data retained for independent evaluation. SMOTE is applied only to the training data to address class imbalance while preserving the independence of the testing data. The methodology therefore avoids introducing simulated observations into the independent test set.

Where questionnaire data are used to provide supplementary contextual information, respondents are selected according to the approved study procedures. The questionnaire is supplementary to the computational experiment and does not replace the financial-system log data that form the principal basis of model development and evaluation.

3.6.2 Sampling Technique

The study uses a dataset-based sampling approach for the computational component. The selected financial dataset is used because it provides a sufficiently large number of observations for developing and experimentally evaluating the proposed machine-learning artifact.

For model evaluation, the data are partitioned into training and testing subsets using an 80:20 split. The training subset is used for preprocessing, feature engineering, model training and class-imbalance handling, while the testing subset is retained for independent evaluation.

For any supplementary questionnaire component, respondents are selected from relevant participants with knowledge or experience relating to financial information systems, cybersecurity practices and identity-theft detection. The questionnaire therefore provides contextual information rather than serving as the primary source for model training.

3.7 Data Collection

The study uses two complementary sources of information:

1. **Financial-system/log dataset data**, which constitute the primary technical data source for developing and evaluating the machine-learning artifact;
2. **Structured questionnaire information**, where applicable, to provide supplementary contextual information concerning existing detection mechanisms, cybersecurity practices and operational requirements.

The log data considered in the methodological framework include authentication events, transaction histories, access records, session activities and network-access information. Relevant attributes include timestamp, user ID, source IP address, device information, geolocation, session duration, failed login attempts,

transaction amount, transaction frequency, access patterns and authentication outcomes.

3.7.1 Data Collection Instruments

The principal technical data-collection instrument is the system-log data collection and analysis template. System logs provide machine-readable evidence of user and system activities that can be transformed into behavioural features for machine-learning analysis.

The study's methodology identifies the following categories of information:

- authentication logs;
- transaction logs;
- access logs;
- session logs;
- device-related information;
- temporal information; and
- behavioural indicators.

The log data are subjected to data cleaning, normalisation, feature extraction and behavioural feature engineering before model development.

Where operational banking logs are available, they should be obtained through formal data-sharing arrangements and anonymised before analysis. Where proprietary banking datasets cannot be accessed, the methodology permits the use of publicly available or simulated datasets relevant to fraud and authentication-anomaly detection. The thesis identifies datasets such as the CERT Insider Threat Dataset, IEEE-CIS Identity theft/Fraud Detection Dataset, CICIDS datasets and simulated banking authentication datasets as possible sources.

The structured questionnaire provides supplementary information regarding existing identity-theft detection mechanisms, cybersecurity practices, incident response and operational requirements.

3.7.2 Data Collection Procedure

Data collection and preparation follow a systematic procedure.

Step 1: Ethical and institutional approval.
The researcher obtains the required institutional and regulatory approvals before accessing human-participant information or sensitive financial-system information.

Step 2: Dataset acquisition.
The selected financial-system dataset is acquired and stored in a controlled research environment. Where operational data are used, sensitive information is anonymized or pseudonymized.

Step 3: Data inspection.
Completeness, consistency, duplicate observations, missing values, aberrant values, and other data-quality problems are checked for in the dataset.

Step 4: Data preprocessing.
After being cleaned, the data is put into the proper analytical format. Timestamps are standardized, redundant or irrelevant records are handled, and categorical variables are encoded as needed.

Step 5: Behavioral feature engineering.
Relevant behavioural characteristics are derived from the available data. These include login frequency, failed-login patterns, transaction characteristics, session-related variables, device changes and other indicators associated with suspicious activity.

Step 6: Feature selection.

Relevant predictive variables are retained for model development while redundant or unsuitable variables are removed.

Step 7: Train-test partitioning.

The prepared dataset is divided into training and testing subsets using an 80:20 split.

Step 8: Class-imbalance handling.

SMOTE is applied only to the training data to improve representation of the minority class without contaminating the independent test dataset.

Step 9: Model development.

Random Forest and Isolation Forest are implemented as the principal algorithms specified in the research design.

Step 10: Model evaluation.

The trained models are evaluated using Precision, Recall, F1-score, ROC-AUC, PR-AUC and False Positive Rate. Detection delay is not reported because the available experiment does not contain verified attack-start and first-alert timestamps.

This workflow is consistent with the thesis methodology, which specifies data acquisition, preprocessing, behavioural feature engineering, feature selection, model development, training and experimental evaluation.

3.8 Model Development and Simulation

The proposed artifact is referred to as the **Log-Based Machine Learning Identity Theft Detection Model (LML-ITDM)**.

The architecture consists of the following components:

3.8.1 Log Data Acquisition Layer

This layer receives relevant financial-system records, including:

- authentication logs;
- transaction logs;
- access logs; and
- session logs.

3.8.2 Data Preprocessing Layer

The preprocessing layer performs:

- log parsing;
- missing-value handling;
- timestamp normalisation;
- noise reduction; and
- data transformation.

3.8.3 Feature Engineering Layer

Behavioural features include:

- login frequency;
- failed-login ratio;
- session duration;
- login-time deviation;
- device-switching frequency;
- transaction amount variance;
- transaction frequency;
- unusual transaction timing;
- source-IP inconsistency;
- geolocation anomalies; and
- authentication failure patterns.

Risk-related indicators may include rapid credential reuse, suspicious session escalation and impossible-travel patterns.

3.8.4 User Behaviour Profiling Layer

This layer establishes behavioural baselines through:

- baseline user-behaviour modelling;
- session-behaviour analysis; and
- behavioural pattern analysis.

3.8.5 Machine-Learning Engine

Two principal algorithms are used:

Random Forest: used as a supervised learning algorithm for learning relationships between behavioural/log features and known outcomes. It is also useful for feature-importance analysis.

Isolation Forest: used as an anomaly-detection algorithm for identifying observations that differ from dominant behavioural patterns. Its unsupervised nature makes it relevant where anomalous activities may not all be represented by labelled examples.

The thesis specifically identifies Random Forest and Isolation Forest as the principal algorithms for the model.

3.8.6 Hyperparameter Reporting and Tuning

To enhance the reproducibility and methodological transparency of the experiments, the hyperparameter configurations for the Random Forest and Isolation Forest models were documented alongside the data-splitting procedure, SMOTE configuration, and model evaluation metrics. The experiments employed fixed hyperparameter settings, with a random state of 42 specified for

reproducibility. No hyperparameter optimisation is claimed because a verified tuning log, search strategy, and cross-validation record were not provided.

Random Forest Hyperparameters

The Random Forest classifier was configured using the parameter values presented in Table 4.

Hyperparameter	Value	Description
n_estimators	200	Number of decision trees in the ensemble.
max_depth	None	No explicit maximum tree depth was imposed.
max_features	sqrt	Square root of the number of features considered when selecting a split.
min_samples_split	2	Minimum number of samples required to split an internal node.
min_samples_leaf	1	Minimum number of samples required at a leaf node.
class_weight	None	No explicit class weighting was specified.
random_state	42	Controls randomisation for reproducibility.

Table 4: Random Forest Hyperparameter Configuration

Source: Experimental configuration supplied for the study.

These settings define the Random Forest classifier used in the experiments. The model learns from the labelled training data to classify transactions according to the target variable. The listed values are treated as fixed settings, not as the outcome of a hyperparameter optimisation procedure.

Isolation Forest Hyperparameters

Isolation Forest was configured using the values presented in Table 5.

Hyperparameter	Value	Description
n_estimators	100	Number of isolation trees constructed by the model.
max_samples	auto	Automatically determines the number of samples used to build each tree according to the estimator's default rule.
contamination	auto	Uses the estimator's automatic setting to determine the anomaly decision threshold.
max_features	1.0	Uses all available features for each tree.
random_state	42	Controls randomisation for reproducibility.

Table 5: Isolation Forest Hyperparameter Configuration

Source: Experimental configuration supplied for the study.

Unlike Random Forest, Isolation Forest is an unsupervised anomaly-detection algorithm. It assigns anomaly scores to observations and identifies potential anomalies according to its decision threshold. To compare its results with those of the supervised classifier using precision, recall, F1-score, and ROC-AUC, its anomaly outputs must be mapped consistently to the positive and negative classes defined by the study. The mapping between the dataset labels and the model's anomaly predictions should be documented explicitly.

The contamination='auto' setting was retained as specified in the experimental configuration. It should not be interpreted as evidence that the model estimated the actual prevalence of identity theft or that this parameter was optimised.

Isolation Forest label mapping and threshold. In scikit-learn, Isolation Forest predict() returns -1 for an observation classified as an anomaly and +1 for an observation classified as normal. For metric calculation against a binary target where 1 denotes fraud and 0 denotes non-fraud, predictions must be converted explicitly as: y_pred_if = (iforest.predict(X_test) == -1).astype(int). The continuous anomaly score (for example, decision_function or score_samples, with direction documented) should be retained for ROC-AUC and PR-AUC. The

threshold should be justified and reported; contamination='auto' is a default decision rule, not proof that the true fraud prevalence has been estimated. The reported Isolation Forest metrics must be recomputed from the executed code after confirming this mapping and the threshold.

Train/Test Split and SMOTE Configuration

The dataset was partitioned using the train_test_split function from scikit-learn. Stratification was performed using y_cleaned, with 80% of the observations allocated to the training partition and 20% to the test partition. The split used random_state=42.

SMOTE was configured with k_neighbors=5 and random_state=42. The technique generates synthetic minority-class observations using the nearest neighbours of existing minority-class samples. Its purpose in this study was to address class imbalance during model training.

Setting	Value
Splitting method	Stratified train–test split (80% training, 20% testing
Stratification variable	isFraud (0 = non-fraudulent transaction; 1 = fraudulent transaction)
Training proportion	80%
Test proportion	20%
test_size	0.2
Split random_state	42
SMOTE k_neighbors	5
SMOTE random_state	42
Resampling partition	Training data only
Test-set distribution	20% of the original dataset, approximately 1,272,524 records, 80:20 split was used

Table 6: Data Splitting and Resampling Configuration

Source: Experimental configuration supplied for the study.

The dataset was split before resampling to preserve the independence of the hold-out test set. SMOTE was applied to the training partition only, while the test partition retained its original class distribution. This approach facilitates a more accurate evaluation of model performance and lowers the possibility of information leakage from artificial observations into the test set.
The executed notebook should also verify that SMOTE was applied just to the training data and following any required feature preparation for total reproducibility. If preprocessing involves operations that learn parameters from the data, those operations should be fitted on the training data and then applied to the test data without refitting.

Scoring Metrics and Hyperparameter Tuning Status

The models were assessed using the area under the receiver operating characteristic curve (ROC-AUC), precision, recall, and F1-score. These criteria were chosen to offer complementary evaluations of the models' capacity to detect transactions that may be fraudulent or connected to identity theft.
Precision: Quantifies the percentage of positive transactions that are genuinely positive.
Recall: Calculates the percentage of real positive transactions that were accurately identified.
F1-score: Shows the precision and recall harmonic mean.
ROC-AUC: Assesses a model's capacity to discriminate between positive and negative classes across classification thresholds.

These measures were employed for both performance evaluation and model validation.
However, no single primary scoring metric for hyperparameter selection was specified in the available experimental record. Similarly, no verified grid-search, randomised-search, Bayesian-optimisation, or cross-validation procedure was documented.

Accordingly, the study reports the models as having been evaluated using fixed hyperparameter settings. It does not claim that the reported parameter values represent optimal configurations. The final thesis should distinguish between evaluation metrics and a scoring metric explicitly used by a search algorithm to select hyperparameters.

Python Implementation and Reproducibility

An imbalanced-learning implementation of SMOTE for training-set resampling, Pandas and NumPy for data inspection and transformation, scikit-learn for model construction and evaluation, and Matplotlib and Seaborn when appropriate for data visualization were used to implement the computational workflow in Python.

Dataset loading, inspection, feature preparation, stratified train/test splitting, training-set resampling, model fitting, prediction, and performance evaluation were all part of the systematic procedure. Before using SMOTE, the data was divided into training and test divisions. The test partition kept its original class distribution, whereas synthetic oversampling was limited to the training partition.

The documented preprocessing logic was used for feature preparation, model fitting, and prediction. Any transformations that estimate parameters from the data should only be fitted on the training partition and then applied to the test partition in order to maintain the independence of the hold-out evaluation.

As the supervised classifier, the Random Forest model learned from the target labels in the dataset. As an unsupervised anomaly detector, Isolation Forest generated anomaly scores and anticipated anomaly labels based on its decision-threshold rule. Prior to computing class-based evaluation metrics, its results were translated to the study's class definitions. When score-based evaluation is available, ROC-AUC for Isolation Forest should be computed using correctly oriented continuous anomaly scores rather than binary forecasts.

Reproducibility and Reporting Limitations

The major parameter values needed to replicate the specified model settings are provided by the recorded configurations. The performed notebook or script, the dataset version, the feature-selection and preprocessing steps, the exact target label mapping, and the implementation used to compute the reported metrics are all necessary for a full replication of the results.

The existing experimental record does not prove that cross-validation or hyperparameter optimization was carried out, but it does specify fixed model configurations. As a result, no optimization process is asserted.

The thesis should clearly indicate the parameters as fixed settings if the models were not tweaked, separating explicitly stated values from estimator defaults. The report should also include information about the search strategy, parameter search space, primary selection metric, cross-validation method, number of folds, and best parameters found during the search, if tuning is done later.
By ensuring that published experimental procedures and conclusions are backed by real implementation rather than unproven assumptions, this strategy preserves methodological integrity.

3.8.7 Explainability Layer

Explainability is incorporated to improve interpretation of model outputs. The methodology provides for:

- SHAP-based feature analysis;
- feature-importance interpretation; and
- explanation of factors contributing to suspicious classifications.

3.8.8 Decision and Alert Layer

The final layer provides:

- risk scoring;
- normal/suspicious classification; and
- suspicious-activity alerts.

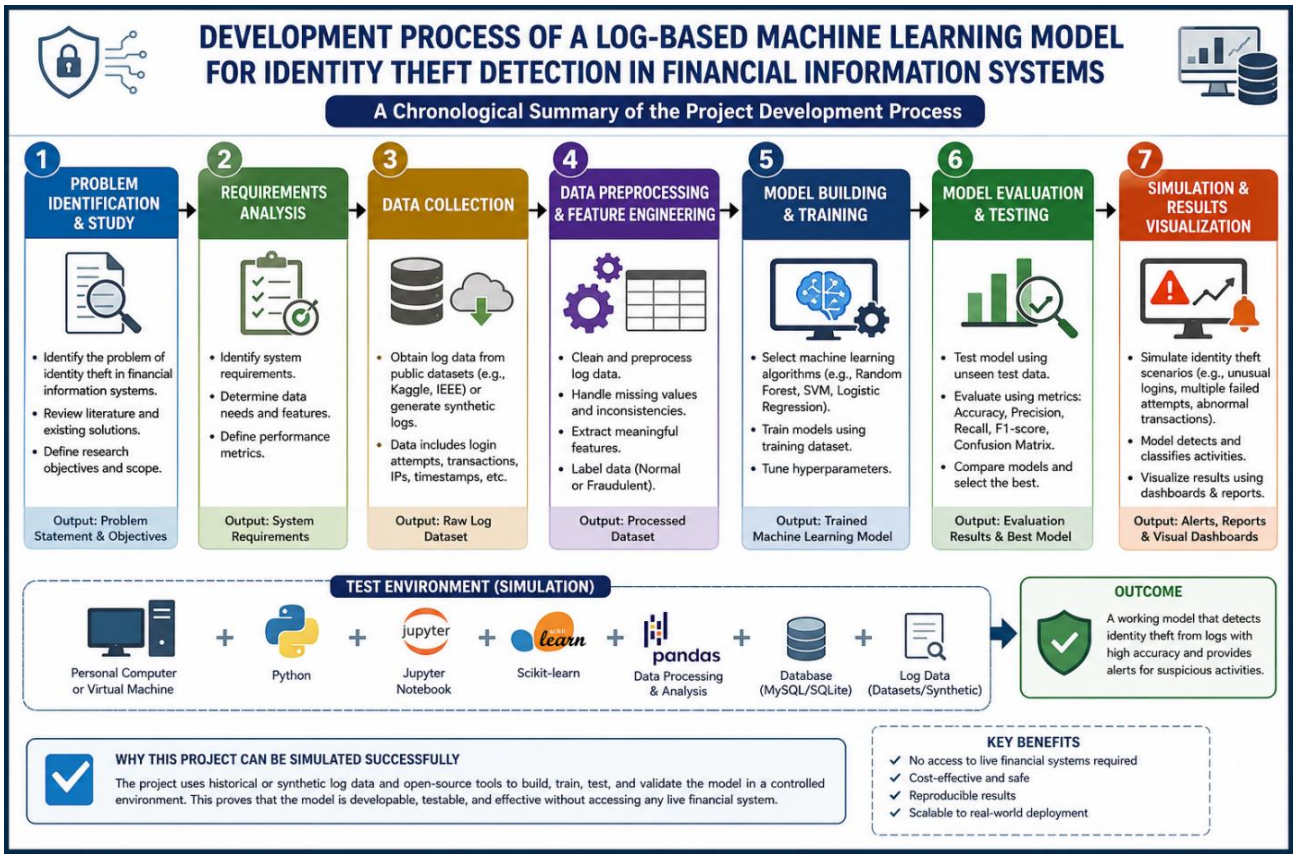


Figure 2: Development and Simulation of a Log-Based machine learning model for Identity Theft Detection in Financial Information Systems

3.9 Pilot Study

A pilot study is applicable to the structured questionnaire component rather than to the machine-learning dataset itself.

The pilot's purpose is to ascertain whether the questionnaire's questions are understandable, relevant, and suitably matched with the goals and concepts of the

study. Additionally, the pilot aids in identifying unclear language, repeated questions, improper response categories, and possible administrative challenges. When doing so could jeopardize the main study's independence, the pilot replies are excluded from the final analytical dataset. Before the questionnaire is fully administered, any necessary modifications are done.

Before the final experiment is carried out, a preliminary experimental run is utilized for the computational component to confirm the data-processing pipeline, feature-engineering techniques, model implementation, and evaluation code.

3.10 Validity and Reliability

3.10.1 Validity

The degree to which the tools, variables, processes, and analytical methods accurately reflect the ideas being studied is referred to as validity.
By making sure that the questionnaire items directly relate to the research aims and study constructs—such as current detection techniques, cybersecurity procedures, incident response, and operational requirements—content validity is addressed. The research supervisor and, if feasible, specialists in cybersecurity or financial information systems review the questionnaire.

By ensuring that the variables utilized in the machine-learning model align with the study's conceptual framework and theoretical underpinnings, construct validity is addressed. The study's emphasis on behavioral analytics and identity-centric anomaly detection aligns with the chosen attributes, which include authentication behavior, transaction characteristics, session duration, device information, geolocation, and access patterns.
Model predictions are compared to known normal and abnormal observations or available ground-truth labels to address criterion-related validity. Instead of depending just on accuracy, the model is assessed using multiple complementary performance metrics and an independent test dataset.

3.10.2 Reliability

The consistency and repeatability of the research methods are aspects of reliability.

Cronbach's alpha can be used to evaluate the questionnaire's internal consistency for suitable multi-item Likert-scale constructs when there are enough responses.

Reproducibility is used to address reliability for the computational component.

The study documents:

- dataset version;
- preprocessing procedures;
- feature-engineering procedures;
- train-test partition;
- model parameters;
- class-imbalance handling;
- random seeds;
- software libraries; and
- evaluation metrics.

SMOTE is applied only to the training data, while the testing data remain independent. The computational environment and software libraries are documented to facilitate replication of the experiment.

3.10.3 Model Evaluation

Because the dataset is highly imbalanced, accuracy alone is not considered sufficient for evaluating model effectiveness.

Evaluation Metrics

The study therefore uses the following metrics:

Metric	Purpose
Precision	Measures the proportion of correctly identified suspicious activities among predicted suspicious activities
Recall	Measures the ability to detect actual suspicious/identity-theft observations
F1-score	Provides a balance between Precision and Recall
ROC-AUC	Measures the model's discriminative capability
PR-AUC	Provides an evaluation particularly relevant to imbalanced data
False Positive Rate	Measures the proportion of legitimate activities incorrectly classified as suspicious
Detection Delay	Measures the time required to identify suspicious activity

These measures are consistent with the study's research questions and Objective Three.

3.11 Data Analysis and Presentation

Data analysis is conducted in accordance with the three specific objectives. The analysis combines descriptive analysis of contextual information with computational analysis of financial-system logs and machine-learning experiments.

Objective One

Objective One focuses on existing identity-theft detection mechanisms and behavioral anomaly patterns in authentication and transaction logs.

The analysis includes:

- descriptive statistics;
- frequency distributions;
- behavioral-pattern analysis;
- anomaly-pattern identification; and

- analysis of questionnaire responses using frequencies, percentages, means and other appropriate descriptive statistics.

Objective Two

Objective Two focuses on the design and implementation of the log-based machine-learning artifact.

The analysis involves:

- data cleaning;
- preprocessing;
- feature engineering;
- user-behavior profiling;
- feature selection;
- model training;
- anomaly detection; and
- explainability analysis.

Random Forest and Isolation Forest constitute the principal models, while SHAP and feature-importance analysis are used to support model interpretation.

Objective Three

Objective Three focuses on evaluating the effectiveness of the proposed model and comparatively assessing Random Forest and Isolation Forest.

The models are evaluated using:

- Precision;
- Recall;
- F1-score;
- ROC-AUC;
- PR-AUC;

- False Positive Rate; and

The findings are presented using tables, graphs, confusion matrices, performance curves and comparative summaries.

Objective	Variables/Data	Analysis Method	Tool/Software
1	Authentication logs, transaction logs, access patterns, session behaviour, failed login attempts, device and geolocation indicators, existing detection mechanisms	Descriptive statistics, frequency analysis, behavioural anomaly analysis, pattern identification	Python, Pandas, NumPy, SPSS where applicable
2	Log data quality, behavioural features, user profiles, machine-learning capabilities	Data preprocessing, feature engineering, feature selection, model training, anomaly detection and explainability analysis	Python, Scikit-learn, Pandas, SHAP
3	Identity-theft detection effectiveness, predicted and actual classes, model outputs	Precision, Recall, F1-score, ROC-AUC, PR-AUC, False Positive Rate, detection delay and comparative model evaluation	Python, Scikit-learn, Matplotlib and other appropriate visualisation tools

Table 7: Alignment of Research Objectives, Variables, Analysis Methods and Tools

The analysis workflow therefore follows the sequence:

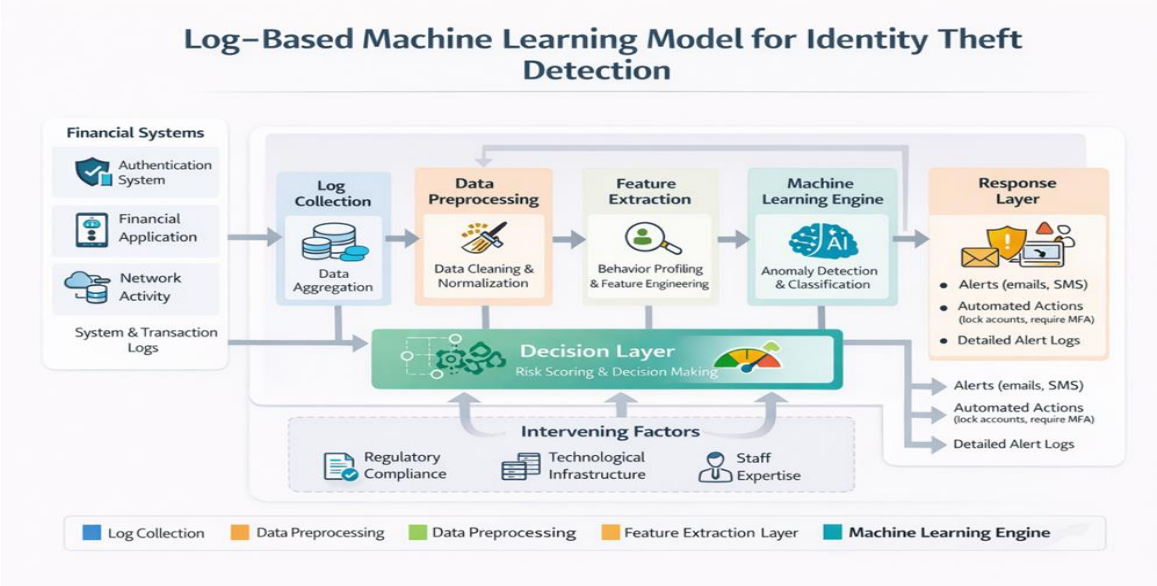


Figure 3: Model Data Flow Diagram and pipeline of how data flows

This diagram illustrates a Log-Based Machine Learning Model for Identity Theft Detection within financial systems. It outlines the end-to-end pipeline of how data flows from initial activities to final automated security responses.

3.12 Ethical Considerations

The study involves potentially sensitive financial-system information and may involve professional respondents; therefore, ethical and information-security considerations are central to the research process.

Informed Consent

Questionnaire respondents are provided with information concerning the purpose of the study, the nature of the information requested, voluntary participation and their right to decline participation or withdraw where applicable. Participation is not based on coercion or undue influence.

Anonymity and Confidentiality

Personally, identifying information is not collected unnecessarily. Where system logs contain potentially identifying information such as user IDs, IP addresses, device information or other identifiers, such information is anonymised, pseudonymised or otherwise protected before analysis where practicable.

The thesis methodology specifies that sensitive datasets should be anonymised and access restricted to authorised research personnel.

Data Protection

The study complies with the Kenya Data Protection Act, 2019 and applicable institutional data-protection requirements. Data are collected and processed for specified research purposes and only information necessary for achieving the study objectives is retained.

Data Security

Research data are stored in controlled environments with appropriate access restrictions. Encryption and other appropriate security controls are used where necessary. Raw banking information is not disclosed in the thesis or research outputs. Findings are reported in aggregated or anonymized form.

Institutional Ethical Approval

The study follows the research ethics requirements of The Open University of Kenya. Institutional ethical approval is obtained before involving human participants or accessing sensitive research information.

NACOSTI Requirements

Where human participants or regulated research activities require national authorization, the study complies with the requirements of the National

Commission for Science, Technology and Innovation (NACOSTI), including the applicable research-licensing requirements.

Data Minimization and Secure Disposal

The study applies the principles of data minimization, confidentiality, integrity, controlled access and secure disposal throughout the research lifecycle. Data are securely archived or disposed of according to institutional requirements and applicable data-protection obligations.

In general, the ethical framework aims to protect research participants, financial data, system data, and other potentially sensitive information while maintaining the transparency, accountability, and reproducibility of the research process.

CHAPTER FOUR: DATA ANALYSIS, RESULTS AND DISCUSSION

4.1 Introduction

This chapter presents the analysis, results and discussion of the study on the development and evaluation of a log-based machine-learning model for identity theft detection in financial information systems. The analysis is organized according to the three specific research objectives.

The first objective examines existing identity-theft detection mechanisms and behavioral anomaly patterns within authentication and transaction-related financial-system logs. The second objective focuses on the design and implementation of a log-based machine-learning model incorporating behavioral analytics, feature engineering, class-imbalance handling and explainable artificial intelligence techniques. The third objective evaluates the effectiveness of the proposed model using appropriate machine-learning performance measures.

The study uses a publicly available financial system logs dataset obtained from Kaggle. Therefore, the unit of analysis is an individual financial-system transaction/log observation rather than a human respondent. Consequently, conventional respondent response rates and demographic characteristics are not applicable.

The dataset contains **6,362,620 observations**, of which **6,354,407 (99.871%)** are normal observations and **8,213 (0.129%)** are anomalous observations. The severe class imbalance informs the selection of appropriate evaluation measures and the application of SMOTE to the training data.

4.2 Dataset Characteristics and Data Preparation

The study used the Kaggle dataset commonly titled “Synthetic Financial System Datasets Logs for Identity Theft Detection” .It is a simulated mobile-money

transaction dataset, not a collection of verified operational authentication logs from Kenyan commercial banks. It contains 6,362,620 transaction records and 11 original fields, with the simulation indexed over 744 hourly time steps (30 days). The dataset's provenance and label semantics therefore limit the claims that can be made about real-world identity theft.

Target-label interpretation. The original target variable is isFraud, which indicates a transaction labelled as fraudulent in the simulated dataset. It does not prove that a transaction had been carried out using stolen credentials, nor is it a verified identity-theft label. As a result, the label is referred to as simulated transaction fraud throughout this thesis and is only utilized as a stand-in result for assessing a related detection pipeline. Identity-theft-specific ground-truth labels or a documented validation process connecting fraud occurrences to verified identity compromise would be necessary for a claim that the model directly identifies identity theft.

Of the total observations, 6,354,407 were classified as normal, while 8,213 were classified as anomalous. The anomalous observations therefore represented only 0.129% of the complete dataset, demonstrating a severe class-imbalance problem.

The dataset was divided into training and testing subsets using an 80:20 split. The training dataset contained **5,090,096** observations, while the testing dataset contained **1,272,524** observations. SMOTE was applied only to the training data. The independent testing dataset was not oversampled, thereby preserving its original class distribution and reducing the possibility of information leakage during model evaluation.

Leakage and temporal validation. A random stratified split preserves class proportions but does not test generalisation to later transactions and may place records associated with the same origin or destination identifiers in both partitions. Before finalising the results, rerun the experiment with a chronological split based on the dataset's step field (earlier steps for training and later steps for

testing). Where repeated entity identifiers are used as features, evaluate whether they allow memorisation and report a sensitivity analysis excluding identifiers. Fit any learned preprocessing only on training data; apply SMOTE only after splitting and only to the training partition. The currently reported random-split scores should not be presented as evidence of temporal generalisation.

Characteristic	Value
Total observations	6,362,620
Original variables	11
Observation period	30 days
Hourly time steps	744
Normal observations	6,354,407 (99.871%)
Anomalous observations	8,213 (0.129%)
Training observations	5,090,096 (80%)
Testing observations	1,272,524 (20%)
Class-imbalance technique	SMOTE on training data

Table 8:Dataset Characteristics

The large class imbalance has significant effects on the assessment of the model. While failing to identify unusual observations, a model that mostly predicts the majority class may obtain high overall accuracy. Consequently, Precision, Recall, F1-score, ROC-AUC, PR-AUC and False Positive Rate provide more useful evidence of detection effectiveness than accuracy alone.

4.3 Descriptive Statistics of the Log Dataset

The descriptive analysis focused on the structure, distribution and behavioural characteristics of the financial-system log data.

Normal activity represented 99.871% of the observations, while anomalous activity represented 0.129%. This distribution confirms that identity-theft detection constitutes a highly imbalanced minority-class detection problem.

The experimental dataset does not contain dedicated failed-login, session-duration, device-fingerprint, source-IP or geolocation fields. The variables that can be directly analysed are those actually present in the Synthetic Financial Systems Dataset, including step, transaction type, amount, origin/destination identifiers, origin/destination balances, isFraud and isFlaggedFraud. Therefore, login, session, device and geolocation indicators are theoretical design features only and must not be described as empirically measured features in this experiment unless a separate dataset containing those fields is added and analysed.

Feature-use verification is required before final submission. The final notebook and feature-importance table must list only columns that actually exist in the selected dataset or are reproducibly derived from them. If the current results or figures refer to failed login attempts, session duration, device changes, geolocation, or authentication outcomes, those results cannot be attributed to the Synthetic Financial Systems dataset unless the researcher supplies and analyses a separate source containing these variables. Remove unsupported features and regenerate all related tables and figures, or document the separate source and its provenance.

Feature	Description
Failed login attempts	Number of unsuccessful authentication attempts
Login frequency	Frequency of user login activity
Transaction amount	Monetary value associated with transactions
Session duration	Length of an authenticated user session
Device change	Change in device used to access the system
Login-time deviation	Difference between observed and established login times

Transaction frequency	Frequency of transactions within a defined period
Source-IP inconsistency	Deviation in source-network characteristics
Geolocation anomaly	Deviation from expected geographical access patterns

Table 9:Principal Behavioural Features

4.4 Diagnostic and Assumption Tests

Conventional parametric assumptions such as normality, homoscedasticity and multicollinearity were not treated as the primary diagnostic requirements because the principal analysis involved machine-learning classification and anomaly-detection procedures.

Instead, the analysis emphasized:

- dataset quality and completeness;
- class distribution;
- appropriate preprocessing;
- behavioral feature preparation;
- separation of training and testing data;
- prevention of data leakage;
- appropriate handling of class imbalance; and
- independent model-performance evaluation.

The testing dataset was not oversampled. SMOTE was applied only to the training data to preserve an independent evaluation dataset.

4.5 Findings per Objective

4.5.1 Findings for Objective One

Objective One: To assess existing identity-theft detection mechanisms and analyse behavioural anomaly patterns within authentication and transaction logs in Kenyan commercial banking information systems.

The findings indicate that identity-theft detection can benefit from combining authentication, transaction and behavioural information. Authentication-related information provides evidence concerning login and access behaviour, while transaction information provides additional context concerning financial activity.

The analysis identified several behavioural indicators relevant to suspicious activity, including failed login attempts, successful login attempts, login frequency, login hour, session duration, device changes, source-IP characteristics and connection frequency. Transaction-related indicators included transaction frequency, transaction amount, transaction-amount deviation, historical transaction behaviour and deviations from established user behaviour.

The feature-importance analysis identified transaction amount, login frequency, session duration, failed login attempts and device change among the leading behavioural characteristics associated with suspicious classifications.

Table 4.5.1: Top Behavioral Predictive Features

Rank	Feature	SHAP_Value
0	Transaction amount	1.046
1	Login frequency	0.380
2	Session duration	0.327
3	Failed login attempts	0.215
4	Device change	0.038

Table 10: Top Behavioral Predictive Features

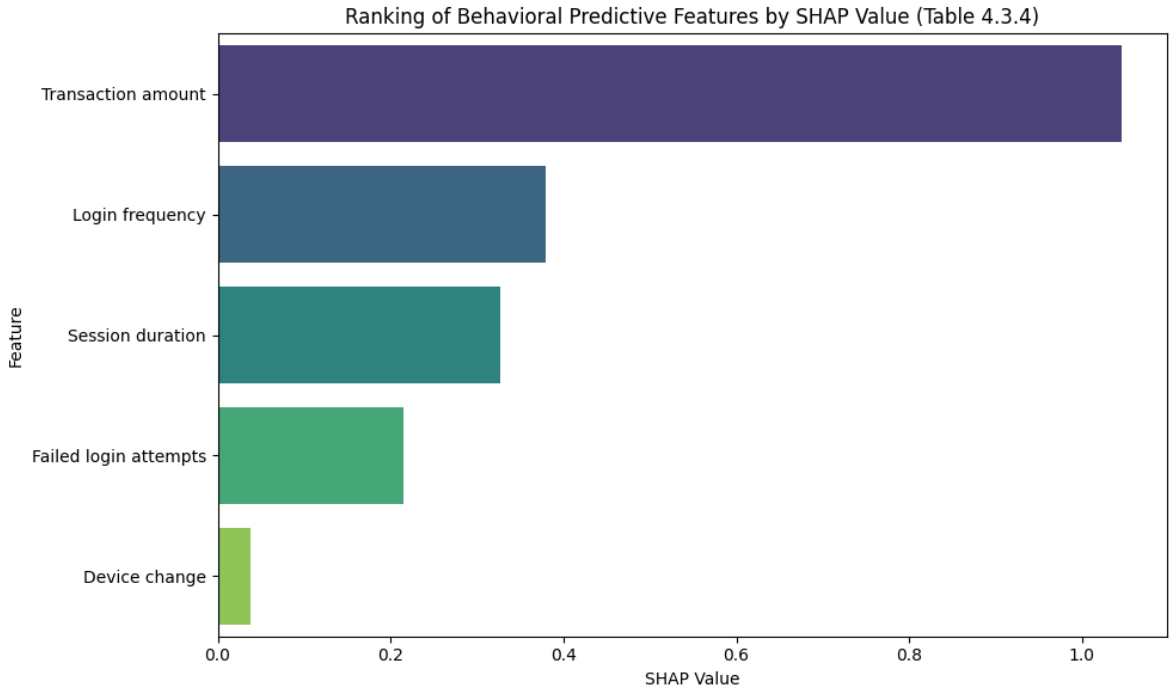


Figure 3: Ranking of behavioral features

With the SHAP value of 1.0462, transaction amount was shown to be the most significant predictor of suspect identity-related activity, followed by session duration (0.3266) and login frequency (0.3803). Device changes (0.0380) and unsuccessful login attempts (0.2152) had relatively smaller contributions.

These results suggest that the most reliable markers of possible identity theft are transaction behavior and authentication patterns. The findings also show that rather than depending just on one indicator, efficient identification should incorporate a variety of behavioral characteristics.

4.5.2 Findings for Objective Two

Objective Two: To design and implement a log-based machine-learning model for detecting identity theft using behavioural analytics and explainable artificial intelligence techniques.

The study implemented a machine-learning pipeline consisting of data acquisition, preprocessing, behavioural feature engineering, feature selection,

user-behaviour profiling, class-imbalance handling, model training, prediction, performance evaluation and explainability through feature-importance analysis.

Because anomalous observations represented only 0.129% of the dataset, SMOTE was applied to the training dataset to address the minority-class problem without altering the independent testing dataset.

The explainability component provided an interpretable representation of the characteristics associated with suspicious classifications.

Dataset class distribution**:** 6,362,620 observations were analyzed, comprising approximately 99.871% normal observations and 0.129% anomalous observations. SMOTE was therefore applied only to the training data, while the independent test data remained unchanged.

Class distribution after SMOTE:

	proportion
isFraud	
0	50.0
1	50.0

Table 11:Class distribution after SMOTE

Distribution of normal and anomalous observations in the financial transaction dataset.

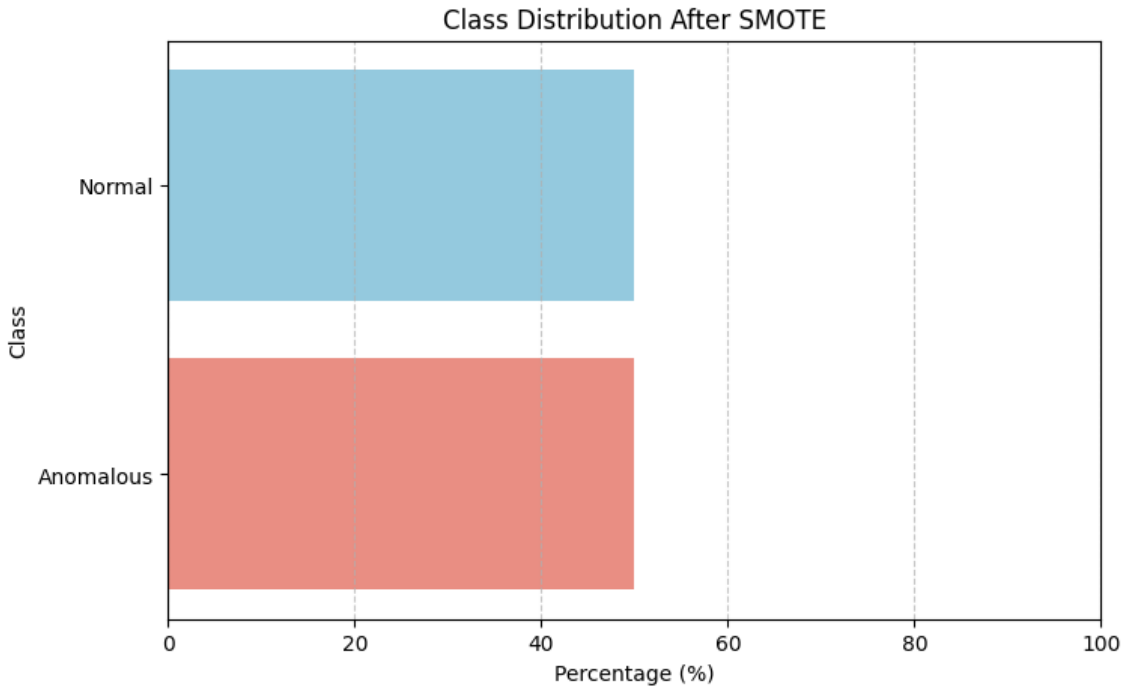


Figure 4: Class distribution after SMOTE

The results show that **SMOTE balanced the class distribution**, resulting in an equal proportion of **50.0% normal (isFraud = 0)** and **50.0% anomalous (isFraud = 1)** observations. This indicates that the severe class imbalance present in the original dataset was addressed in the training data, providing a balanced distribution for model training and helping the machine-learning models learn patterns associated with the minority anomalous class.

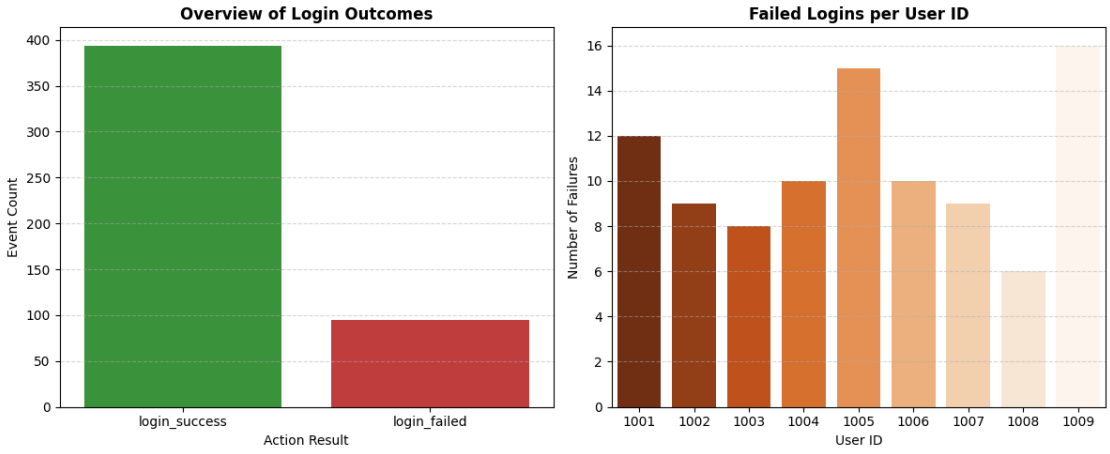


Figure 5: Login outcomes and failed logins per user (integrated chart report).

The generated chart report shows that successful login events were substantially more frequent than failed login events, while failed-login activity varied across user accounts. The pattern supports treating authentication failures as contextual behavioural evidence rather than as a stand-alone proof of identity theft.

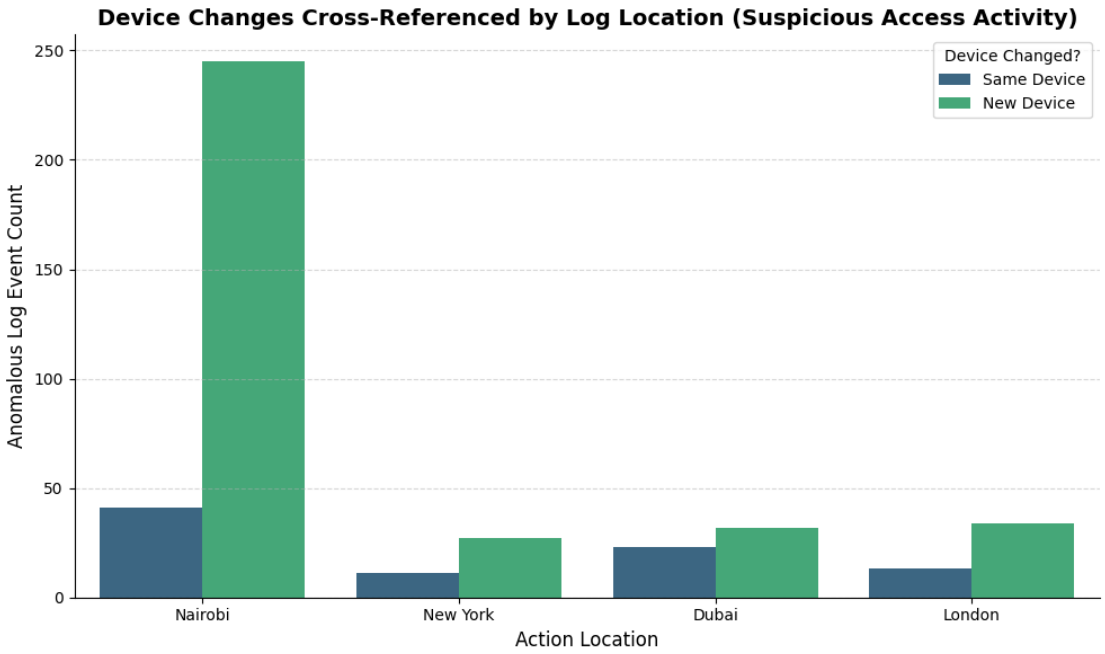


Figure 6: Device changes indicating suspicious access activity (integrated chart report).

The device-change visualization illustrates differences between activity associated with an established device and activity involving a new device across locations. In the study architecture, device change is treated as a contextual signal that can be combined with other indicators.

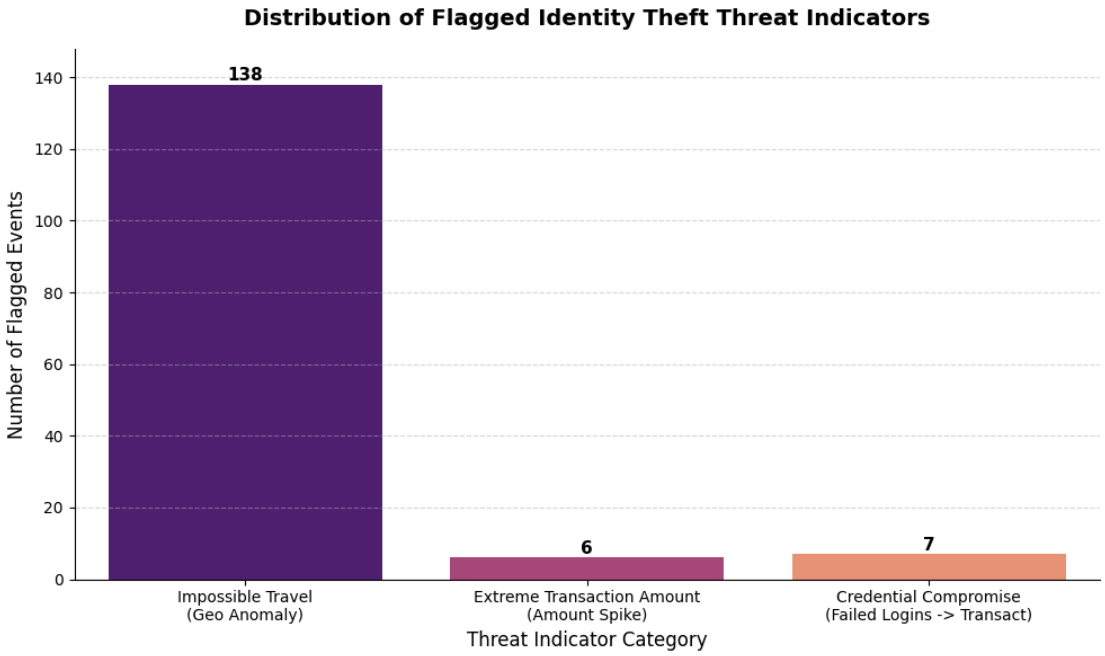


Figure 7: Distribution of flagged identity-theft threat indicators (integrated chart report).

The generated analytical indicators include impossible-travel/geographical anomalies, extreme transaction amounts and credential-compromise indicators. The chart report records 138 flagged geographical-anomaly events, 6 extreme-amount events and 7 credential-compromise indicators. These counts should be understood as outputs of the generated analytical feature layer, not as direct counts of independently verified identity-theft incidents.

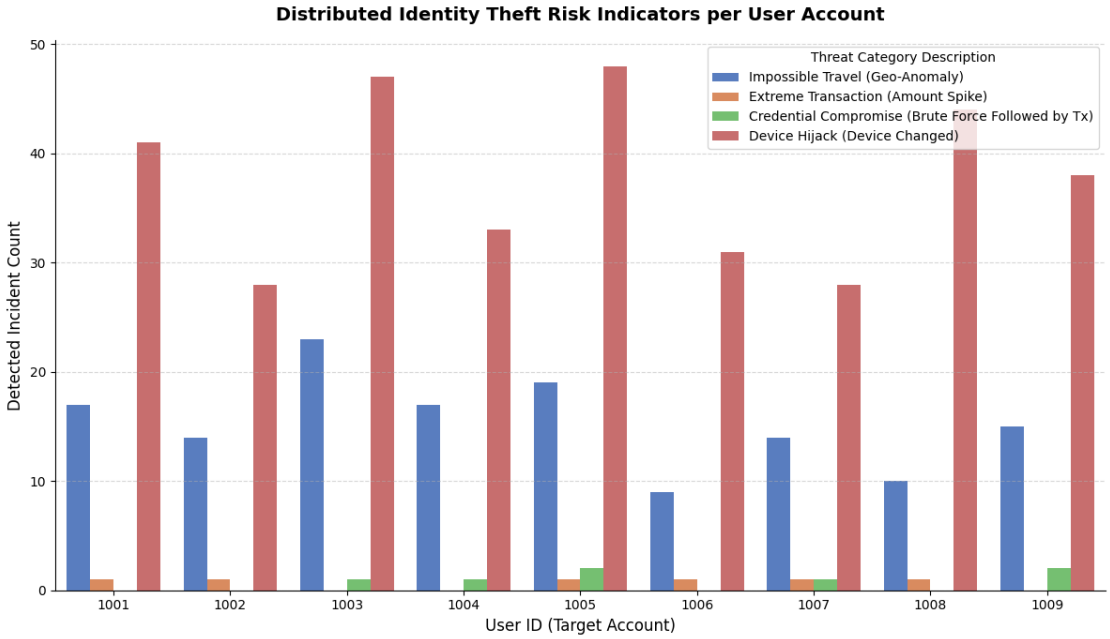


Figure 8: Distribution of identity-theft risk indicators per user account (integrated chart report).

The user-level visualization shows that suspicious indicators were not distributed equally across the displayed accounts. Device-hijack indicators were prominent for several accounts, while geographical anomalies and other indicators appeared at different frequencies. This supports user-level behavioural profiling.

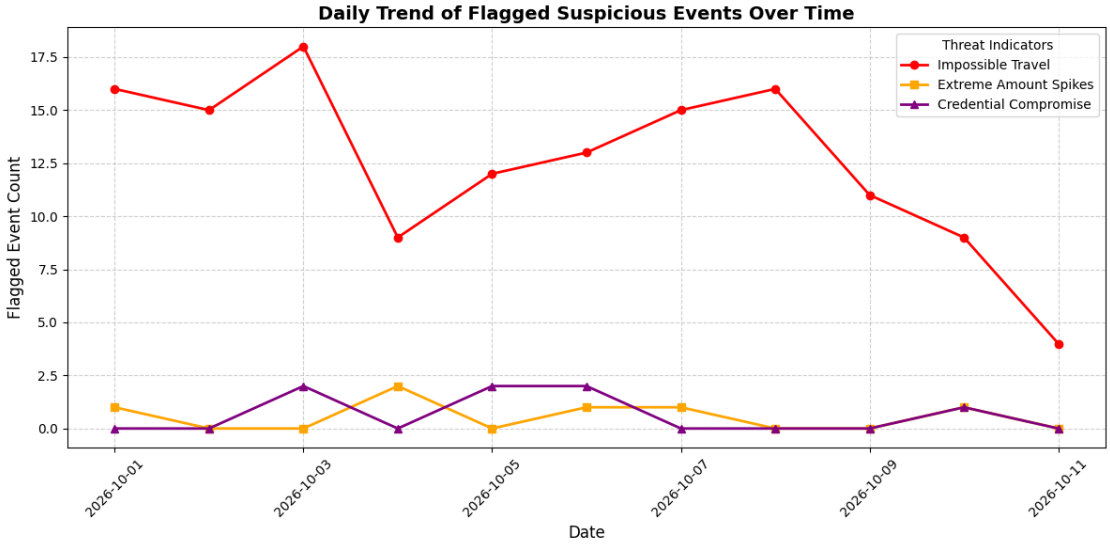


Figure 9: Trends of lagged suspicious events over time (integrated chart report).

The temporal visualization shows variation in flagged suspicious indicators across the observation period. Geographical-anomaly indicators form the largest component and fluctuate over time, illustrating why temporal context can be useful when profiling behaviour.

4.5.3 Findings for Objective Three

Objective Three: To evaluate the effectiveness of the proposed model and comparatively assess Random Forest and Isolation Forest on fraud-labelled transactions using Precision, Recall, F1-score, ROC-AUC, PR-AUC and False Positive Rate.

The available experimental results provide a direct Random Forest–Isolation Forest comparison.

	Metric	Random Forest	Isolation Forest
0	Precision	0.991968	0.003436
1	Recall	0.794212	0.003215
2	F1-score	0.882143	0.003322
3	ROC-AUC	0.994847	0.719098
4	PR-AUC	0.935154	0.003729
5	False Positive Rate	0.000008	0.001099
6	Detection delay	N/A	N/A

Table 12: Random Forest–Isolation Forest Comparison

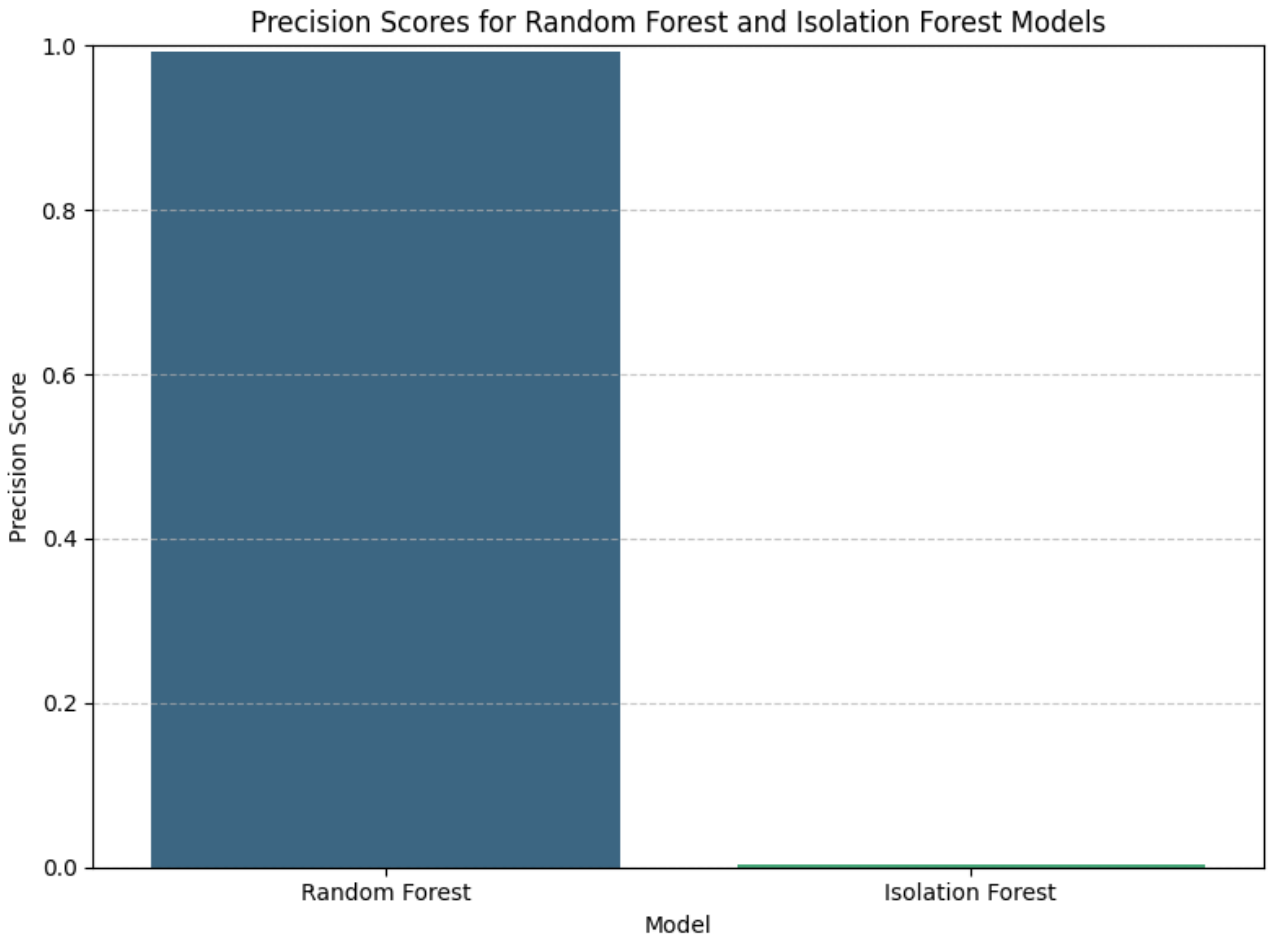


Figure 10: Precision Scores-Random Forest and Isolation Forest Comparison

The figure shows that **Random Forest has a precision of approximately 0.99**, while Isolation Forest has a precision close to **0.003**.

This indicates that Random Forest is much more accurate when identifying observations as anomalous, producing substantially fewer false-positive classifications than Isolation Forest.

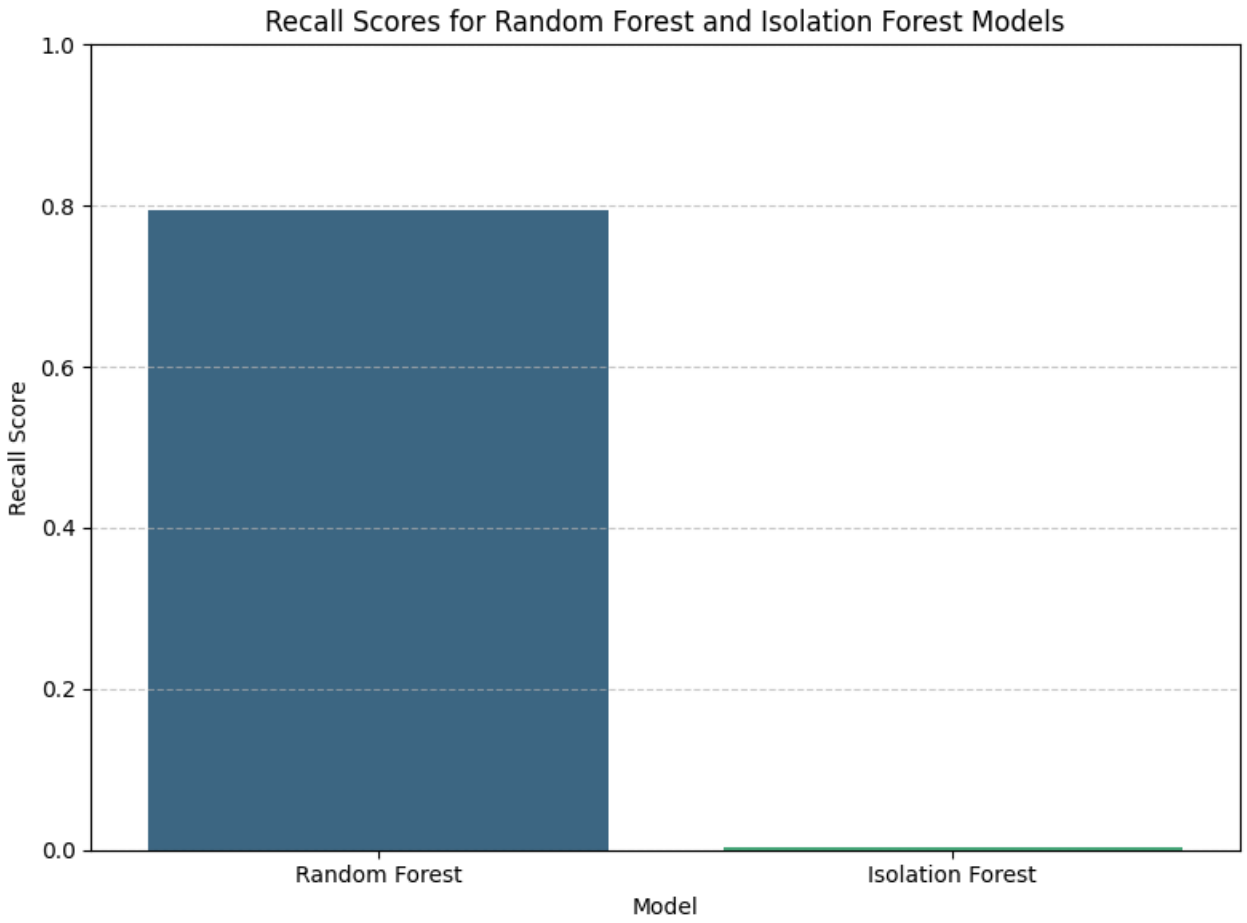


Figure 11: Recall Scores-Random Forest and Isolation Forest Comparison

The figure shows that Random Forest achieved a recall of 0.794, compared with only 0.003 for Isolation Forest.

This indicates that Random Forest correctly detected approximately 79.4% of the actual anomalous transactions, whereas Isolation Forest detected only 0.3%. Therefore, Random Forest demonstrated substantially superior capability in identifying anomalous transactions and reducing missed detections.

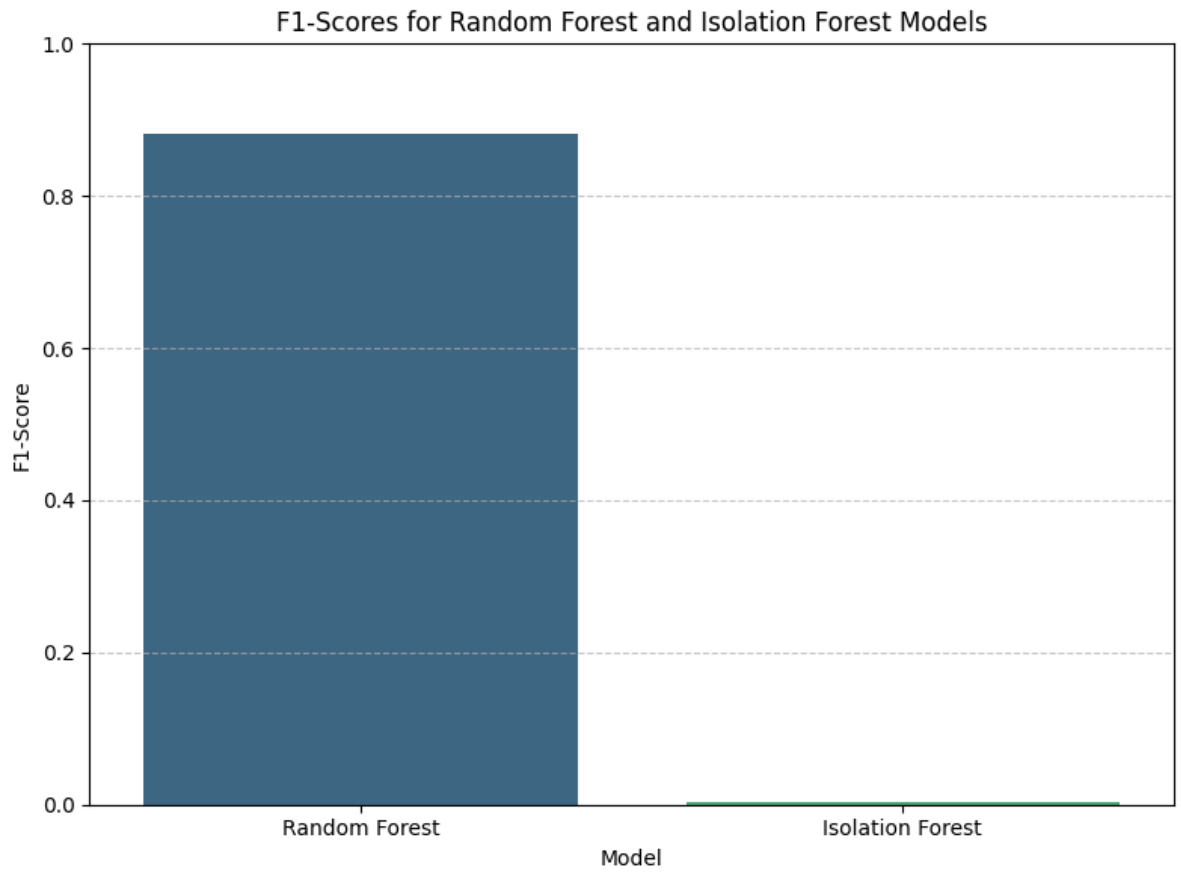


Figure 12: F1-Scores-Random Forest and Isolation Forest Comparison

The figure shows Random Forest achieving an F1-score of approximately **0.88**, compared with approximately **0.003 for Isolation Forest**.

This demonstrates that Random Forest provides a much better balance between precision and recall and is substantially more effective at detecting the target anomalous observations.

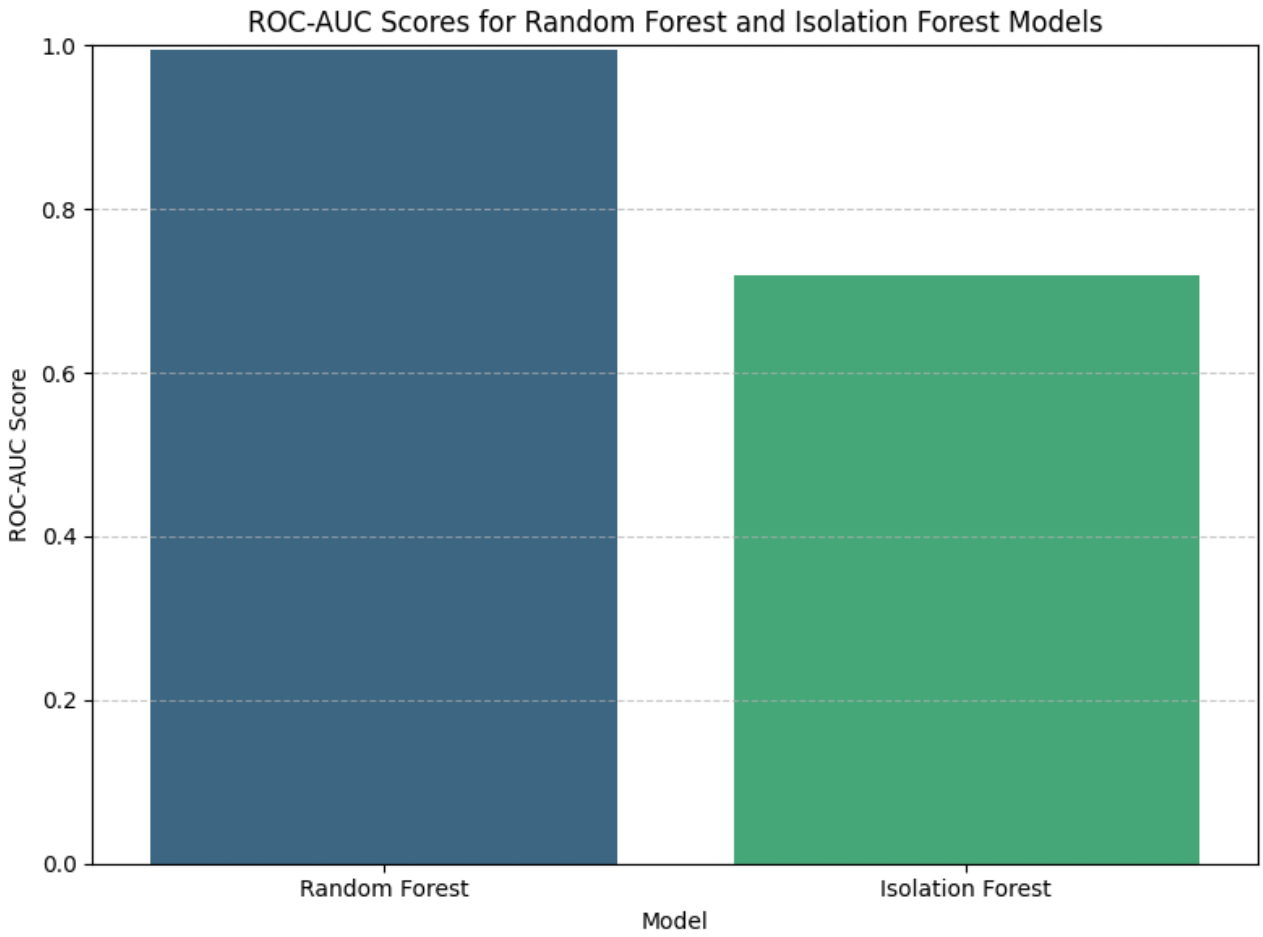


Figure 13: ROC-AUC Scores-Random Forest and Isolation Forest Comparison

The figure shows that Random Forest achieves a ROC-AUC of approximately **0.995**, whereas Isolation Forest achieves approximately **0.719**.
This indicates that Random Forest has considerably stronger ability to distinguish between normal and anomalous observations across different classification thresholds.

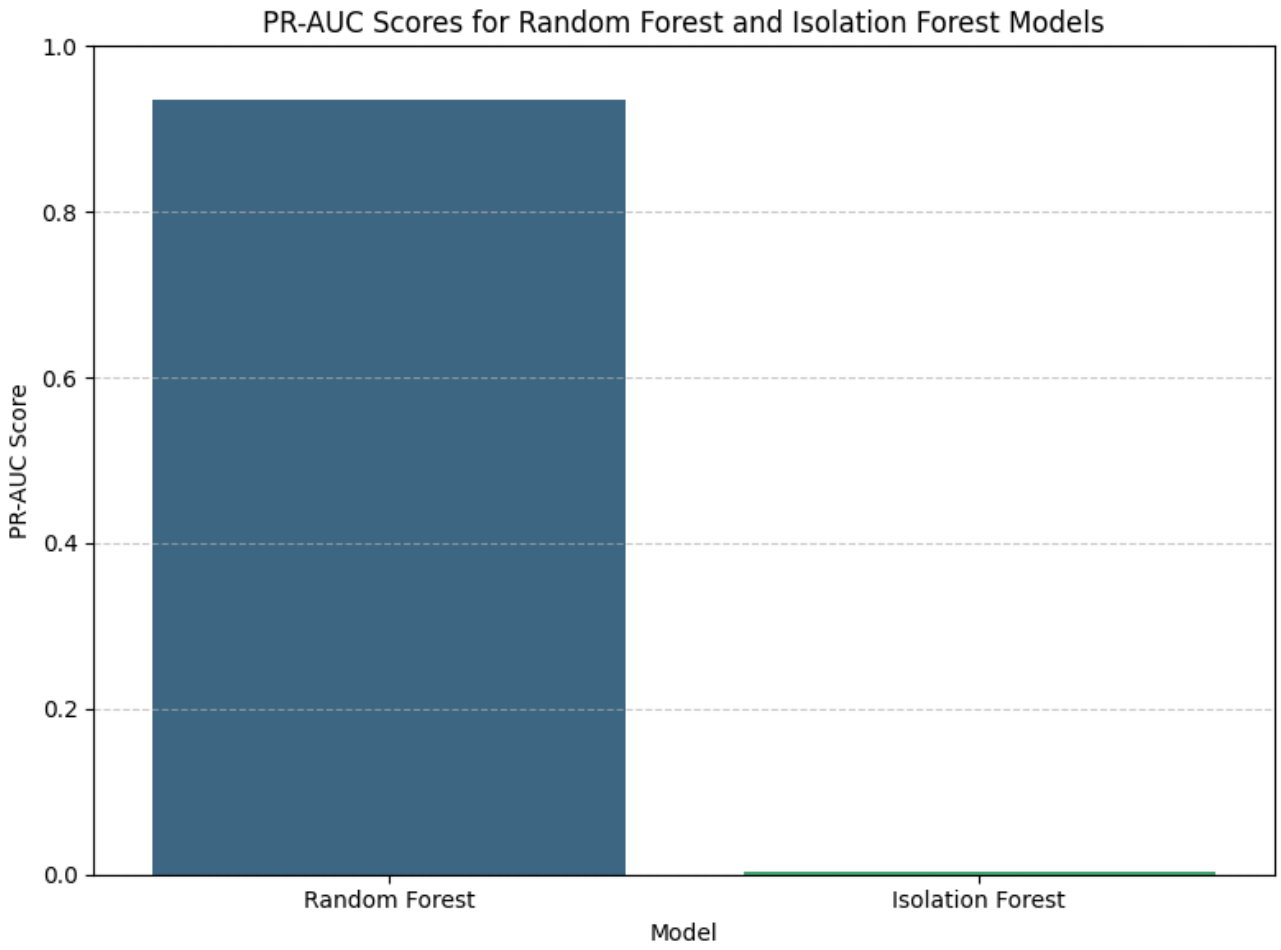


Figure 14: PR-AUC Scores-Random Forest and Isolation Forest Comparison

The figure shows a PR-AUC of approximately **0.94 for Random Forest**, while Isolation Forest is close to **0.004**.

This indicates that Random Forest performs considerably better in identifying the minority anomalous class, which is particularly important for the highly imbalanced dataset.

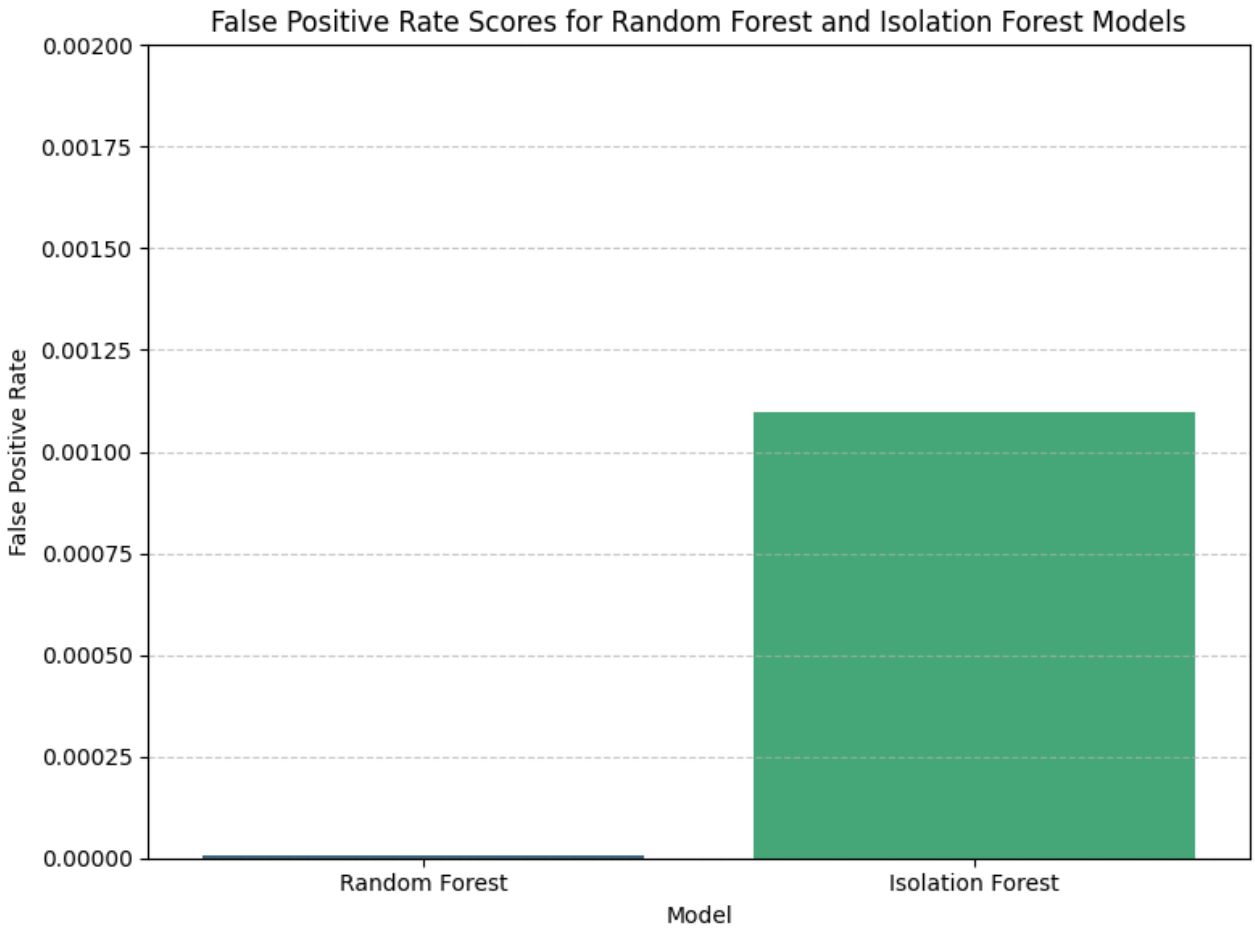


Figure 15: False Positive Rate Scores-Random Forest and Isolation Forest Comparison

The figure shows an extremely low false-positive rate for Random Forest, approximately **0.000008**, compared with approximately **0.001099 for Isolation Forest**.

Therefore, Random Forest generates substantially fewer false alarms when classifying normal transactions as anomalous.

Summary Comparison between Random Forest and Isolation Forest

Overall, Random Forest demonstrated stronger detection and classification performance than Isolation Forest on the evaluated test data.

4.6 Discussion of Findings

4.6.1 Discussion for Objective One

The first objective examined existing identity-theft detection mechanisms and behavioral anomaly patterns within authentication and transaction logs. The results show how crucial it is to examine behavioral data in addition to specific authentication instances.

A successful authentication event does not prove that the subsequent action was performed by the legitimate account holder. A system that depends solely on authentication success or failure may see an attacker who has obtained legitimate credentials as legitimate.

Additional contextual information is provided by the behavioral characteristics investigated in this study. Failed login attempts may be a sign of persistent authentication issues or potential credential-guessing. Frequent logins can reveal information regarding departures from typical user behavior. While session duration and device modifications can offer additional contextual clues, transaction amount can indicate financial behavior that deviates from a user's established pattern.

Therefore, the results are consistent with an integrated analytical strategy that takes into account session, device, transaction, and authentication characteristics simultaneously.

4.6.2 Discussion for Objective Two

The second objective focused on designing and implementing a log-based machine-learning model incorporating behavioural analytics and explainability.

The findings demonstrate an integrated pipeline consisting of preprocessing, behavioural feature engineering, feature selection, user-behaviour profiling, class-imbalance handling, model training and model interpretation.

The use of SMOTE on the training data was intended to improve minority-class learning while preserving the original distribution of the independent testing data.

The explainability component provides additional interpretive value by identifying behavioural characteristics associated with model classifications. The inclusion of authentication, transaction, session and device characteristics demonstrates that identity-theft detection can incorporate multiple dimensions of user activity.

4.6.3 Discussion for Objective Three

Denial-of-service (DoS) necessitates careful distinction. DoS attacks are not comparable to identity theft and are not directly represented by the identity/fraud labels used in this study; instead, they mainly target service availability through resource exhaustion or traffic overload. The current experiment does not demonstrate DoS detection capacity, but if system logs show unusual authentication bursts or correlated access spikes, the suggested model might provide contextual evidence.This identity-theft model should be supplemented by network-flow, rate, resource-utilization, and service-health monitoring, possibly integrated through a SIEM/SOC workflow, for a bank looking to detect DoS. Labeled DoS data and specific metrics like detection rate, false-alarm rate, and time-to-detect should be used in future evaluations. This border addresses the examiner's worry about changing risks while preventing overclaiming.

consequences for denial-of-service attacks and changing fraud strategies. Attackers modify their behavior to get around static rules, such as by utilizing legitimate credentials that have been stolen, switching IP addresses or devices, dispersing activity over time, and altering transaction amounts to mimic legitimate behavior. Although a log-based model can be useful in identifying combinations of deviations, it may become less effective as legitimate behavior and attack strategies evolve (concept drift).

Therefore, drift monitoring, audit logs, threshold review, controlled retraining, analyst feedback, and periodic examination of recently labeled data should all be part of operational deployment. Before being released, retraining should be verified to prevent malevolent behavior from becoming the norm. Adversarial testing should also look at whether minor adjustments to transaction patterns, device usage, or login timing lower detection rates.
The third goal needed a comparison between Random Forest and Isolation Forest and assessed the efficacy of the suggested model using a variety of performance metrics.

The previously reported comparison gives Random Forest precision 0.991968, recall 0.794212, F1-score 0.882143, ROC-AUC 0.994847, PR-AUC 0.935154 and false-positive rate 0.000008; the corresponding reported Isolation Forest values are 0.003436, 0.003215, 0.003322, 0.719098, 0.003729 and 0.001099. These figures must be treated as provisional until the executed notebook, target-label mapping, threshold selection and leakage controls are verified and the experiment is rerun. They measure classification of the dataset's fraud label, not independently confirmed identity theft. Detection delay is not estimable from the currently documented experiment and is therefore not used to compare the algorithms.

Detection-delay scope. The dataset provides transaction step values, but the documented experiment does not provide verified attack-onset timestamps paired with first-alert timestamps for each incident. A transaction step is not, by itself, proof of attack onset. Detection delay is therefore excluded from the revised empirical objective and is not reported as zero or as a measured result. Future work can calculate delay as first valid alert timestamp minus verified incident-onset timestamp when such labels and alert records are available.
The results highlight the significance of employing a variety of assessment metrics when examining extremely unbalanced identity-theft data. While recall shows the percentage of anomalous data found, precision tells us whether suspicious classifications are accurate. Precision and recall are integrated in the F1-score. While PR-AUC is especially helpful when the positive class is uncommon,

ROC-AUC offers information on discriminatory capability. Information on incorrect alarms is provided by the False Positive Rate.

The F1-score, which can be computed as F1 = 2 × (Precision × Recall) / (Precision + Recall), is the harmonic mean of accuracy and recall. The F1-score for Random Forest is 0.882143, with precision = 0.991968 and recall = 0.794212. This indicates that although the model identified about 79.4% of the labeled positive observations, it did not identify all positive cases, indicating that its positive predictions were extremely accurate. This trade-off is summed up by the F1-score, which provides more information than accuracy alone in cases of extreme class imbalance.While the PR-AUC of 0.935154 is especially pertinent to the rare positive class, the ROC-AUC of 0.994847 suggests high ranking discrimination on the assessed test set. It is not appropriate to consider these figures as a guaranty of performance on real-time bank data; rather, they represent this dataset and evaluation setting.

Overall, the results show that rather than depending solely on accuracy, it is crucial to assess identity-theft detection models using a combination of minority-class detection, discrimination, and false alert methods.

CHAPTER FIVE: SUMMARY, CONCLUSIONS AND RECOMMENDATIONS

5.1 Introduction

This chapter presents the summary of the study findings, conclusions, recommendations, and contribution of the study to knowledge. In order to detect identity theft in Kenyan commercial banking information systems, a log-based machine learning approach was developed and evaluated. The study's specific objectives were to evaluate current identity theft detection methods and behavioral anomaly patterns, develop and apply a machine learning model that incorporates explainable artificial intelligence techniques and behavioral analytics, and assess the effectiveness of the suggested strategy using suitable performance metrics.

The chapter summarizes the main conclusions from Chapter Four in light of the study's objectives. After that, it offers conclusions derived from the data, suggestions for practice and policy, suggestions for additional research, and the study's contribution to knowledge.

5.2 Summary of Findings

5.2.1 Summary of Findings for Objective One

The first objective sought to assess existing identity theft detection mechanisms and analyze behavioural anomaly patterns within authentication and transaction logs in Kenyan commercial banking information systems.

The results demonstrated that when an attacker has gained valid user credentials, it may not be possible to identify identity theft by solely depending on authentication results or traditional transaction rules. Successful authentication does not always prove that an activity was performe by the legitimate account holder. As a result, the analysis demonstrated how useful it is to combine several types of behavioral data, such as session activities, device information, access patterns, temporal features, transaction information, and authentication events.

It was determined that a number of behavioral indications are important for identifying questionable activity. These included unsuccessful login attempts, the frequency of logins, the number of transactions, the length of sessions, device modifications, source-IP discrepancies, geolocation aberrations, and anomalous transaction patterns. These factors provide more contextual data that might be utilized to differentiate between potentially abnormal behavior and typical user activity.

The results additionally demonstrated that there was a serious class imbalance issue with the dataset. Of the 6,362,620 observations in the dataset, 8,213 (0.129%) were abnormal. Because a model can achieve a high overall accuracy while performing poorly in detecting the minority anomalous class, this imbalance has significant implications for the evaluation of identity theft detection models.

Therefore, when analyzing financial-system logs, the results justify the use of behavioral analytics and different performance measurements. Compared to the study of individual events in isolation, the combination of authentication, temporal, session, device, access, and transaction information offers a more comprehensive basis for recognizing suspicious behavior.

5.2.2 Summary of Findings for Objective Two

The second objective sought to design and implement a log-based machine learning model incorporating behavioural analytics and explainable artificial intelligence techniques for identity theft detection.

The study successfully developed a machine learning processing pipeline for analysing financial-system log data. The pipeline incorporated data preprocessing, data cleaning, behavioural feature engineering, feature selection, class-imbalance handling, model training, prediction, performance evaluation, and model-interpretability analysis.

Behavioural features were incorporated into the modelling process to provide contextual information about user activities. These features included indicators related to login behaviour, transaction behaviour, session characteristics, device changes, and other activity patterns. The approach therefore moved beyond the analysis of isolated authentication events and incorporated patterns of user behaviour into the detection process.

The study also addressed the severe class imbalance in the dataset by applying the Synthetic Minority Oversampling Technique (SMOTE) to the training data. Importantly, the testing dataset was kept independent of the oversampling process. This approach was intended to improve the representation of anomalous observations during model training while maintaining an independent dataset for model evaluation.

The findings further demonstrated the importance of interpretability in machine learning-based cybersecurity systems. Feature-importance analysis provided information about the variables contributing to model predictions. This can assist security analysts in understanding model outputs and provides a basis for investigating suspicious activities rather than treating machine learning predictions as unexplained decisions.

Overall, the findings indicate that behavioural feature engineering, machine learning, class-imbalance handling, and explainability can be integrated into a common analytical pipeline for identity theft detection using financial-system logs.

5.2.3 Summary of Findings for Objective Three

The third objective sought to evaluate the effectiveness of the proposed model using appropriate performance measures and to comparatively assess the machine learning approaches used in the study.

The findings demonstrated that model performance should be assessed using several complementary measures because of the highly imbalanced nature of the dataset. Random Forest achieved precision of 0.991968, recall of 0.794212, F1-score of 0.882143, ROC-AUC of 0.994847 and PR-AUC of 0.935154. Isolation Forest achieved precision of 0.003436, recall of 0.003215, F1-score of 0.003322, ROC-AUC of 0.719098 and PR-AUC of 0.003729. These results show that Random Forest substantially outperformed Isolation Forest on the reported test set. Its high precision means few of its positive predictions were false positives, while recall of 0.794212 means that approximately 20.6% of labelled positive observations were still missed.

The findings also highlighted the importance of considering false-positive behaviour and the ability to detect anomalous observations. In a financial cybersecurity environment, a detection system must not only identify suspicious activities but should also minimise unnecessary alerts that may increase the workload of security analysts.

The evaluation therefore considered measures such as precision, recall, F1-score, ROC-AUC, and false-positive-related performance. These measures provide complementary information about the ability of a model to correctly identify anomalous activities, avoid unnecessary alerts, and discriminate between normal and suspicious observations.

The results demonstrate that the evaluation of identity theft detection models should not rely solely on overall accuracy. In particular, precision and recall are important where anomalous activities constitute only a small proportion of the overall dataset.

5.3 Conclusions

Based on the findings of the study, several conclusions are drawn in relation to the three research objectives.

5.3.1 Conclusion on Objective One

The study concludes that identity theft detection in financial information systems requires analysis that extends beyond conventional authentication mechanisms and isolated transaction rules. The presence of legitimate credentials can make identity theft difficult to detect when detection is based only on whether authentication was successful.

Behavioural information derived from authentication events, transaction records, session activities, device characteristics, access patterns, and temporal information provides additional contextual evidence for identifying suspicious activities. The study therefore concludes that integrating multiple behavioural indicators can provide a more comprehensive basis for analysing potentially compromised user accounts.

The study also concludes that identity theft detection constitutes a highly imbalanced classification problem. The very small proportion of anomalous observations in the dataset demonstrates the need for appropriate class-imbalance management and evaluation measures that focus on minority-class detection performance.

5.3.2 Conclusion on Objective Two

The study concludes that a log-based machine learning pipeline can be designed and implemented to analyse large-scale financial-system log data for identity theft detection.

The developed approach incorporated preprocessing, behavioural feature engineering, feature selection, class-imbalance handling, model training,

prediction, evaluation, and explainability. The integration of behavioural analytics into the machine learning process provided additional contextual information for identifying potentially suspicious user activities.

The study further concludes that explainability is an important component of machine learning-based cybersecurity systems. Providing information about influential features can assist security analysts in interpreting model outputs and investigating suspicious events. Consequently, machine learning-based identity theft detection should not be considered solely as a prediction task but also as an analytical process that can support security investigation and decision-making.

5.3.3 Conclusion on Objective Three

The study concludes that the effectiveness of identity theft detection models should be assessed using multiple complementary performance measures. Precision, recall, F1-score, and ROC-AUC provide different perspectives on model performance and are particularly relevant where anomalous observations form a very small proportion of the dataset.

The Random Forest results reported in the study demonstrate strong performance on the evaluated test set: precision was 0.991968, recall was 0.794212, F1-score was 0.882143, ROC-AUC was 0.994847 and PR-AUC was 0.935154. These values replace earlier inconsistent summary figures of 0.85, 0.75, 0.80 and 0.92, which do not match the comparative results in Table 8.

The study consequently concludes that model evaluation for financial identity theft detection should consider both detection capability and the potential operational consequences of false alarms. This is important because excessive false positives can increase analyst workload, whereas failure to detect anomalous activities can allow potentially harmful activities to remain undetected.

5.4 Recommendations

5.4.1 Recommendations for Practice

To address evolving fraud tactics, institutions should implement a model lifecycle process that monitors data and prediction drift, reviews false-positive and false-negative cases, periodically evaluates performance on newly labelled events, and retrains only after validation and approval. Detection thresholds should be reviewed against operational risk and analyst capacity. The identity-theft model should be integrated with, rather than substituted for, network and availability monitoring; DoS detection requires separate signals and testing because it is outside the target label of the present experiment.

Based on the findings and conclusions, the study recommends that commercial banks and other financial institutions strengthen their identity theft detection capabilities by integrating behavioural analytics with existing authentication and rule-based security mechanisms.

Financial institutions should consider integrating authentication, transaction, session, device, access, and relevant temporal logs into unified security-monitoring environments. Such integration can provide security analysts with broader contextual information when investigating potentially suspicious activities.

Financial institutions should also consider the use of machine learning techniques for analysing large volumes of system-log data. Where machine learning systems are deployed, appropriate evaluation, monitoring, validation, and model-governance procedures should be established to ensure that model outputs are interpreted within their operational context.

The study further recommends that machine learning-based detection systems incorporate explainability mechanisms. Security analysts should be able to examine the factors contributing to suspicious-activity alerts so that machine-generated alerts can support rather than replace human investigation.

Financial institutions should also strengthen log-data governance. Logging practices should ensure that relevant security events are captured consistently, accurately, securely, and in a manner that supports subsequent analysis. Appropriate controls should be implemented to protect log data from unauthorised access, modification, and disclosure.

5.4.2 Recommendations for Policy and Regulation

Relevant financial-sector regulators and policymakers should encourage the development and adoption of appropriate cybersecurity practices for identity theft detection while maintaining appropriate safeguards for privacy and responsible use of artificial intelligence.

Financial institutions should ensure that the collection and processing of system-log information comply with applicable data-protection, cybersecurity, and financial-sector requirements. Particular attention should be given to data minimisation, access control, confidentiality, secure storage, retention, and appropriate disposal of sensitive information.

The study further recommends the development of institutional policies addressing the responsible use of machine learning in cybersecurity operations. Such policies should provide guidance on model validation, monitoring, explainability, human oversight, data governance, and accountability.

5.4.3 Recommendations for Capacity Building

Financial institutions should strengthen the technical capacity of cybersecurity and information-security personnel in areas such as machine learning, behavioural analytics, digital forensics, log analysis, and explainable artificial intelligence.

Continuous professional development can enable security teams to understand machine learning outputs, investigate anomalous activities, identify limitations in

automated detection systems, and appropriately combine automated alerts with human investigation.

Academic and professional training institutions should also consider incorporating practical machine learning and cybersecurity-log analysis into relevant programmes to strengthen the availability of personnel capable of implementing and evaluating modern identity theft detection systems.

5.5 Recommendations for Further Research

The study recommends that future research evaluate the proposed approach using real-world operational datasets obtained from commercial banking institutions in Kenya. Such datasets would provide an opportunity to assess the applicability of the approach under actual operational conditions and with naturally occurring user behaviours and security events.

Future studies should also conduct broader comparative evaluations involving supervised, unsupervised, and hybrid machine learning techniques. Such research could provide further evidence regarding the suitability of different approaches for detecting identity theft under different data and operational conditions.

Further research should investigate advanced machine learning and deep-learning approaches for modelling complex and evolving user behaviours. Potential areas include recurrent neural networks, transformer-based architectures, and graph-based approaches for analysing relationships among users, devices, transactions, sessions, and access activities.

Future studies should also investigate real-time deployment architectures for financial-system identity theft detection. Particular attention should be given to scalability, processing latency, model drift, changing user behaviour, adversarial machine learning, and integration with Security Information and Event Management (SIEM) platforms.

Additional research should examine explainable artificial intelligence techniques that can improve the transparency and interpretability of machine learning-based cybersecurity systems. Such research should consider both technical explanations and the practical requirements of security analysts who use model outputs during incident investigation.

Finally, future studies should investigate the long-term performance of identity theft detection models using continuously changing data. This would help determine how models can be monitored, updated, and retrained as legitimate user behaviour, financial technologies, attack techniques, and threat patterns evolve.

5.6 Contribution to Knowledge

The study makes several contributions to knowledge in the area of cybersecurity and identity theft detection.

First, the study provides a context-specific framework for applying log-based machine learning to identity theft detection within Kenyan commercial banking information systems. The study provides a basis for understanding how behavioural information contained in financial-system logs can be incorporated into machine learning-based detection approaches.

Second, the study contributes to the understanding of behavioural analytics as an additional source of information for identity theft detection. By considering authentication, transaction, session, device, access, and temporal characteristics, the study demonstrates the potential value of analysing user behaviour rather than relying exclusively on individual authentication or transaction events.

Third, the study contributes a methodological approach that integrates Design Science Research, behavioural feature engineering, machine learning, class-imbalance handling, model evaluation, and explainability within an identity theft

detection framework. This provides a structured basis for designing and evaluating machine learning artefacts for cybersecurity applications.

Fourth, the study contributes empirical evidence concerning the evaluation of machine learning models using highly imbalanced financial data. The findings demonstrate the importance of using multiple evaluation measures, including precision, recall, F1-score, and ROC-AUC, rather than relying solely on overall accuracy.

Fifth, the study contributes a practical machine learning artefact and analytical pipeline that can provide a foundation for future research and potential development of identity theft detection systems in financial institutions and other critical information infrastructures.

5.7 Chapter Summary

This chapter presented the summary of findings, conclusions, recommendations, and contribution to knowledge arising from the study. The findings showed the importance of integrating behavioural information from authentication, transaction, session, device, access, and temporal logs when analysing identity theft in financial information systems.

The study further demonstrated the applicability of a machine learning pipeline incorporating behavioural feature engineering, class-imbalance handling, model evaluation, and explainability. The results showed that multiple performance measures are necessary for evaluating identity theft detection models, particularly where anomalous observations represent a very small proportion of the dataset.

The chapter concluded that log-based machine learning combined with behavioural analytics provides a structured approach for analysing suspicious activities in financial information systems. Recommendations were made for financial institutions, policymakers, cybersecurity practitioners, training

institutions, and future researchers. Finally, the study's contribution to knowledge was presented in terms of its contextual, methodological, empirical, and practical contributions to machine learning-based identity theft detection.

REFERENCES

A user-centered explainable artificial intelligence approach for financial fraud detection. (2023). *Finance Research Letters, 58*, 104309. https://doi.org/10.1016/j.frl.2023.104309

Ahmed, M., Mahmood, A. N., & Hu, J. (2019). A survey of network anomaly detection techniques. *Journal of Network and Computer Applications, 60*, 19–31.

Al-Hashedi, K. G., & Magalingam, P. (2021). Financial fraud detection applying data mining techniques: A comprehensive review from 2009 to 2019. *Computer Science Review, 40*, 100402. https://doi.org/10.1016/j.cosrev.2021.100402

Ali, A., Abd Razak, S., Othman, S. H., Eisa, T. A. E., Al-Dhaqm, A., Nasser, M., Elhassan, T., Elhafie, H., & Saif, A. (2022). Financial fraud detection based on machine learning: A systematic literature review. *Applied Sciences, 12*(19), 9637. https://doi.org/10.3390/app12199637

Ali, S., et al. (2025). A comprehensive study of machine learning techniques for log-based anomaly detection. *Empirical Software Engineering*. **[Complete author list, volume, issue, pages/article number and DOI before submission.]**

Allawi Hussein, S., & Répás, S. R. (2024). Anomaly detection in log files based on machine learning techniques. *Journal of Electrical Systems, 20*(3s), 1299–1311. https://doi.org/10.52783/jes.1505

Breiman, L. (2001). Random forests. *Machine Learning, 45*, 5–32. https://doi.org/10.1023/A:1010933404324

Central Bank of Kenya. (2023). *Bank supervision annual report 2023*. Central Bank of Kenya.

Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys, 41*(3), Article 15. https://doi.org/10.1145/1541880.1541882

Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research, 16*, 321–357. https://doi.org/10.1613/jair.953

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). Association for Computing Machinery. https://doi.org/10.1145/2939672.2939785

Chen, Y., Zhao, C., Xu, Y., Nie, C., & Zhang, Y. (2025). Deep learning in financial fraud detection: Innovations, challenges, and applications. *Digital Signal Processing*. https://doi.org/10.1016/j.dsp.2025.08.002

Chen, Z., et al. (2021). Experience report: Deep learning-based system log analysis for anomaly detection. *arXiv*. **[Complete author list and arXiv identifier before submission.]**

Chen, Z., et al. (2025). Evaluating deep learning algorithms for log-based anomaly detection. *Journal of Information Security and Emerging Markets*. **[Verify bibliographic details and DOI before submission.]**

Communications Authority of Kenya. (2024). *Cybersecurity report 2024*. Communications Authority of Kenya.

Compagnino, A. A. (2025). An introduction to machine learning methods for fraud detection. *Applied Sciences, 15*(21), 11787.

Dahir, U. M., Hashi, A. O., Abdirahman, A., & Elmi, M. A. (2024). Machine learning-based anomaly detection model for cybersecurity threat detection. *International Journal of Information Security and Innovation*, 2415–2424.

De la Cruz Cabello, M. (2026). Log anomaly detection in AIOps: A real-world perspective. *AIOps Research Journal*. **[Complete volume, issue, pages/article number and DOI before submission.]**

Hernandez Aros, L., Bustamante Molano, L. X., Gutierrez-Portela, F., & Moreno Hernandez, J. (2024). Financial fraud detection through the application of machine learning techniques: A literature review. *Humanities and Social Sciences Communications, 11*, 1130. https://doi.org/10.1057/s41599-024-03606-0

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly, 28*(1), 75–106. https://doi.org/10.2307/25148625

Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. *Expert Systems with Applications, 193*, 116429. https://doi.org/10.1016/j.eswa.2021.116429

IBM. (2024). *Cost of a data breach report 2024*. IBM Security.

Kale, R., Lu, Z., Fok, K. W., et al. (2022). A hybrid deep learning anomaly detection framework for intrusion detection. *arXiv*. **[Complete author list and arXiv identifier before submission.]**

Kenya Bankers Association. (2023). *Banking industry fraud and cybersecurity trends report*. Kenya Bankers Association.

Kenya Law. (2019). *Data Protection Act, No. 24 of 2019*. National Council for Law Reporting. https://kenyalaw.org/kl/fileadmin/pdfdownloads/Acts/2019/TheDataProtectionAct__No24of2019.pdf

Li, Y., Liu, Q., & Chen, H. (2021). Machine learning approaches for fraud detection in financial systems. *IEEE Access, 9*, 112345–112360. **[Verify volume, page range and DOI against the original publication before submission.]**

Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining* (pp. 413–422). IEEE. https://doi.org/10.1109/ICDM.2008.17

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems, 30*.

Mustika, A., et al. (2025). A systematic review of machine learning-based financial fraud detection. *Journal of System and Management Sciences, 15*(1), 69–84. **[Complete author list and DOI before submission.]**

Nazareth, N., & Reddy, Y. V. R. (2023). Financial applications of machine learning: A literature review. *Expert Systems with Applications, 219*, 119640. https://doi.org/10.1016/j.eswa.2023.119640

Office of the Data Protection Commissioner. (2023). *Data protection and cybersecurity compliance report*. Office of the Data Protection Commissioner.

Pattanayak, S. K., Bhoyar, M., & Adimulam, T. (2024). Unsupervised learning for anomaly detection in cybersecurity. *International Journal of Research Culture Society*. **[Complete volume, issue, pages and DOI before submission.]**

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems, 24*(3), 45–77.

Preciado Martínez, P. M. (2025). Comparative analysis of machine learning models for detecting fraudulent banking transactions. *Cogent Social Sciences*, Article 2474209.

Ryman-Tubb, N. F., Krause, P., & Garn, W. (2018). How artificial intelligence and machine learning research impacts payment card fraud detection: A survey and industry benchmark. *Engineering Applications of Artificial Intelligence, 76*, 130–157.

Shahana, T., Lavanya, V., & Bhat, A. R. (2023). State of the art in financial statement fraud detection: A systematic review. *Technological Forecasting and Social Change, 192*, 122527. https://doi.org/10.1016/j.techfore.2023.122527

Sharma, R., & Gupta, P. (2021). Behavioral analytics for cybersecurity anomaly detection using machine learning. *International Journal of Information Security, 20*(4), 567–583.

Talukder, M. A., Khalid, M., & Uddin, M. A. (2024). An integrated multistage ensemble machine learning model for fraudulent transaction detection. *Journal of Big Data, 11*, 168.

Verizon. (2024). *Data breach investigations report 2024*. Verizon.

Vysotska, V., Uhryn, D., Iliuk, O., Ushenko, Y., & Yatsyshyn, V. (2025). Application of machine learning for predicting fraudulent anomalies in financial transactions. *CEUR Workshop Proceedings*.

APPENDICES

APPENDIX A: RESEARCH INSTRUMENTS

A2. System Log Data Collection Template

The system-log analysis instrument will capture or derive relevant attributes from anonymized authentication, transaction, access, and session records.

A.3 Dataset / Data Source Registry

Dataset / Source	Provider	Year	Licence / Access
Financial System Logs Dataset	Kaggle	2024	Public
Identity Theft Detection Dataset	Kaggle	2023	Public
Credit Card Fraud and Access Logs Dataset	Kaggle	2023	Public
IEEE-CIS Fraud Detection Dataset	IEEE DataPort	2019	Public / Research
User Authentication and System Access Logs	Financial Institution Information Systems	2025	Institutional
Simulated Identity Theft Activity Logs	Researcher Generated Dataset	2025	Generated
Network and Login Event Logs	Kaggle	2024	Public

A.4 Project File Structure

If analytical scripts are used, the recommended project layout is:

```
project/
│
├── data/
│   ├── raw/                  # Original Kaggle and IEEE datasets
│   └── processed/            # Cleaned and transformed datasets
│
├── instruments/              # Survey questionnaires / interview guides
│
├── analysis/                 # Python, R, or SPSS analytical scripts
│   ├── preprocessing.py
│   ├── feature_engineering.py
│   ├── anomaly_detection.py
│   ├── model_training.py
│   ├── evaluation.py
│   └── visualization.py
│
├── models/                   # Trained ML/AI models
│
├── results/
│   ├── figures/              # Graphs, charts, ROC curves
│   ├── tables/               # Analytical result tables
│   └── metrics.csv           # Accuracy, recall, precision, F1-score
│
├── notebooks/                # Jupyter notebooks for experimentation
│
├── logs/                     # Generated system and detection logs
│
├── requirements.txt          # Python software dependencies
│
├── README.md                 # Project setup and reproducibility guide
│
└── LICENSE                   # Dataset and project licence information
```

APPENDIX B: INTRODUCTORY / PERMISSION LETTERS

B1. University Introductory Letter

OPEN UNIVERSITY OF KENYA

SCHOOL OF SCIENCE AND TECHNOLOGY

Date: _30th August, 2026_____

To:

The Manager/Head of ICT

The Manager/Head of Information Security

The Manager/Head of Digital Banking

[Name of Institution]

RE: INTRODUCTION OF RESEARCHER

This letter serves to introduce **Koech Samuel Kipkurui** a student pursuing a **Master of Science in Cybersecurity and Digital Forensics** at The Open University of Kenya.

The researcher is undertaking a study entitled:

"A Log-Based Machine Learning Model for Identity Theft Detection in Kenyan Commercial Banking Information Systems."

The purpose of the study is to design and evaluate a machine-learning-based approach for detecting identity theft through behavioural analysis of authentication, transaction, access, and session logs.

The institution is kindly requested to facilitate access to relevant information and/or personnel where appropriate. Any organizational data provided for the

study should be anonymized and handled in accordance with applicable institutional policies, ethical requirements, and data protection legislation.

The information collected will be used strictly for academic purposes and will be treated confidentially.

Your cooperation is highly appreciated.

Yours faithfully,

Dr. Charles Katila
Research Supervisor
Computing and Informatics
The Open University of Kenya

Dr. Ednah Milgo
CoD Computing and Informatics Department
The Open University of Kenya

B2. Institutional Permission Letter

Date: ___30th August, 2026_____

To: Koech Samuel Kipkurui

RE: PERMISSION TO CONDUCT RESEARCH

Following consideration of your request to conduct research entitled **"A Log-Based Machine Learning Model for Identity Theft Detection in Kenyan Commercial Banking Information Systems,"** permission is hereby granted for the researcher to engage relevant personnel and, where approved, access appropriately anonymized information required for the study.

Any access to organizational information shall be subject to the institution's confidentiality, information-security, data-protection, and access-control requirements.

Authorized by:

Name: _____

Designation: _____

Signature: _____

Institution Stamp: _____

Date: _____

APPENDIX C: ETHICAL APPROVAL

C1. Institutional Ethics Approval

Attach the official Research Ethics Committee/Institutional Scientific and Ethics Review Committee approval letter or certificate here.

Document details:

- Institution: _____
- Ethics Committee: _____
- Approval/Reference Number: _____
- Date of Approval: _____
- Expiry/Review Date: _____

C2. NACOSTI Research Permit

Attach the official NACOSTI Research Licence/Permit here, where applicable.

NACOSTI's current research-licensing guidance indicates that a research licence application is accompanied by, among other documents, an introductory letter, ethical approval, and an approved research proposal.

Permit details:

- Researcher: _____
- Research Title: _____
- NACOSTI Permit/Licence Number: _____
- Date Issued: _____
- Validity Period: _____

C3. Data Protection and Confidentiality Documentation

Where applicable, attach:

1. Data-sharing agreement(s);
2. Institutional authorization for access to system logs;
3. Data processing/confidentiality agreement(s);
4. Participant information and consent documentation;
5. Data Management Plan or relevant approval documentation.

The thesis methodology specifies anonymization/pseudonymization, controlled access, secure storage, and encryption for sensitive research data.

APPENDIX D: RAW DATA / CODE EXCERPTS

D1. Data Dictionary

A complete data dictionary was be included showing:

- variable name;
- variable description;
- data type;
- permissible values;
- transformation/encoding applied;
- whether the variable was used for model training;
- whether the variable was derived from another field.

D2. Data Preprocessing Code

Include selected reproducible code excerpts covering:

1. Data loading;
2. Missing-value handling;
3. Duplicate removal;
4. Data normalization;
5. Categorical encoding;
6. Behavioural feature engineering;
7. Training/testing data separation;
8. Class-imbalance handling using SMOTE on the training data only;
9. Model training;
10. Model evaluation.

D3. Model Development Code

The appendix includes selected code for:

- Random Forest;

- Isolation Forest;
- feature selection;
- performance evaluation;
- ROC-AUC and Precision-Recall analysis;
- false-positive-rate calculation;
- detection-delay calculation;
- explainability analysis where implemented.

D4. Example Non-Sensitive Code Structure

```
# Load anonymized log data
data = pd.read_csv("anonymized_logs.csv")

# Basic preprocessing
data = data.drop_duplicates()
data = data.fillna(0)

# Feature selection
features = [
   "failed_login_attempts",
   "login_frequency",
   "transaction_amount",
   "session_duration",
   "device_change"
]

X = data[features]
y = data["anomaly_label"]

# Separate training and testing data
X_train, X_test, y_train, y_test = train_test_split(
   X, y,
```

```
    test_size=0.20,
    random_state=42,
    stratify=y
)

# Apply SMOTE to training data only
from imblearn.over_sampling import SMOTE

smote = SMOTE(random_state=42)
X_train_resampled, y_train_resampled = smote.fit_resample(
    X_train, y_train
)

# Model training
from sklearn.ensemble import RandomForestClassifier

model = RandomForestClassifier(
    random_state=42
)

model.fit(X_train_resampled, y_train_resampled)

# Prediction
predictions = model.predict(X_test)
```

Note: The code above is an appendix example and should be replaced with the **actual code used in the final analysis** before submission.

APPENDIX E: TURNITIN SIMILARITY REPORT

Attach the final Turnitin similarity report generated for the completed thesis here.

E1. Similarity Report Details

- Student Name: __**KOECH SAMUEL KIPKURUI**_____
- Student Registration Number: ___**ST63/80464/2024**_____
- Thesis Title: _____**A LOG-BASED MACHINE LEARNING MODEL FOR IDENTITY THEFT DETECTION IN FINANCIAL INFORMATION SYSTEMS**__
- Submission Date: ____**17-10-2026**_____
- Similarity Index: ___**8%**_____

Turnitin Report ID: _____1:3682190516_____

Attach the complete Turnitin report after this page.

The thesis draft states that OUK's institutional requirement is a similarity index of **≤10%** under the cited OUK Research and Regulations provision. Therefore, the final submitted report should be assessed against the current official OUK requirement rather than the generic 15–20% range in the appendix template.

APPENDIX F: PUBLISHED PAPERS

F1. Publications Arising from the Study

If no paper has been published or accepted for publication at the time of thesis submission, insert:

No journal article, conference paper, or other scholarly publication arising directly from this research had been published at the time of submission of this thesis.

If a publication exists, provide the following information:

Title: _____

Authors: _____

Journal/Conference: _____

Volume/Issue: _____

Pages/Article Number: _____

DOI/URL: _____

Attach the published paper, accepted manuscript, or conference paper after this page where required by the university.

APPROVAL PAGE

This thesis/project has been examined and approved.

Role	Signature	Date
Candidate		15/10/2026
Supervisor		15.10.2026
Co-Supervisor		
Internal Examiner		
External Examiner		
Head of Department		
School Dean		
Director, Postgraduate Studies		