



Adaptive Class-Imbalance Learning for High-Sensitivity Fraud Detection in Large-Scale Digital Transaction Networks

Aminul Islam Miraz¹; Md Ashrafuzzaman²;

[1]. BBA (Finance, Minor: Accounting), Taylor's University, Malaysia;
Email: aminulmiraz26@gmail.com

[2]. International American University, Los Angeles, USA;
Email: md.ashrafuzzamanuk@gmail.com

Doi: [10.63125/bmxwv741](https://doi.org/10.63125/bmxwv741)

Received: 18 June 2023; Revised: 17 July 2023; Accepted: 08 August 2023; Published: 03 September 2023

Abstract

This study addresses the persistent problem of severe class imbalance in large-scale digital transaction networks, where fraudulent transactions represent only a small minority of total observations and conventional accuracy-focused classifiers can overlook genuine fraud cases, increase false negatives, and weaken minority-class sensitivity. The purpose of the research is to quantitatively evaluate how adaptive class-imbalance learning capabilities influence High-Sensitivity Fraud Detection Performance across banking, fintech, electronic payment processing, digital-wallet, e-commerce, and related enterprise transaction environments. A quantitative, cross-sectional, case-study-based design was adopted, using purposive sampling and a five-point Likert-scale questionnaire administered to professionals in fraud analytics, machine learning, cybersecurity, transaction monitoring, payment risk, and digital transaction management. Of 320 questionnaires distributed, 296 were returned and 288 valid responses were retained, producing a 90.0% usable response rate. The key independent variables were Adaptive Minority-Class Resampling, Cost-Sensitive Learning, Dynamic Classification-Threshold Optimization, and Adaptive Ensemble Learning, while High-Sensitivity Fraud Detection Performance was the dependent variable. Analysis included descriptive statistics, Cronbach's alpha, KMO and Bartlett's tests, Pearson correlation, multiple regression, ANOVA, tolerance, VIF, and Durbin-Watson diagnostics. High-Sensitivity Fraud Detection Performance Recorded $M = 4.19$, $SD = 0.54$. Dynamic Classification-Threshold Optimization produced the strongest relationship, $r = .74$, $p < .001$, and the strongest regression effect, $\beta = .32$, $p < .001$. The overall model achieved $R = .842$, $R^2 = .709$, adjusted $R^2 = .705$, $F(4, 283) = 172.45$, $p < .001$, explaining 70.9% of the variance and supporting all five hypotheses. The findings imply that transaction-network operators should integrate adaptive threshold optimization, cost-sensitive learning, resampling, and ensemble methods to improve fraud sensitivity while controlling false-positive burdens.

Keywords

Class Imbalance; Fraud Detection; Threshold Optimization; Cost-Sensitive Learning; Ensemble Learning;

INTRODUCTION

Fraud detection refers to the systematic identification of transactions, activities, accounts, or behavioral patterns that indicate intentional deception, unauthorized financial activity, or other forms of economic abuse within transactional systems. In contemporary digital environments, fraud detection is commonly formulated as a classification problem in which computational models distinguish legitimate transactions from fraudulent ones by analyzing historical transaction characteristics, customer behavior, temporal patterns, transaction amounts, account characteristics, network relationships, and other risk-related attributes. A large-scale digital transaction network can therefore be understood as an interconnected technological environment through which electronic payments and financial exchanges are initiated, processed, transmitted, authenticated, authorized, recorded, and monitored across banking platforms, payment processors, card networks, financial technology services, mobile-payment systems, e-commerce platforms, and other digitally mediated financial infrastructures. The expansion of such networks has increased the importance of automated fraud analytics because very large numbers of transactions must be evaluated while fraudulent observations usually represent only a small fraction of total transactional activity. This condition is formally described as class imbalance, which occurs when the observations belonging to one class considerably outnumber those belonging to another class and learning algorithms consequently encounter difficulty representing the less frequent category accurately (He & Garcia, 2009). Within financial transaction datasets, legitimate transactions generally constitute the majority class, whereas confirmed fraudulent transactions form the minority class that possesses considerably greater detection importance (Bhattacharyya et al., 2011). Conventional accuracy can consequently provide an incomplete representation of model performance because a classifier can correctly classify most majority observations while failing to identify an operationally important proportion of minority fraud cases (Saito & Rehmsmeier, 2015). High-sensitivity fraud detection places greater emphasis on the correct identification of fraudulent transactions and the reduction of false negatives, with sensitivity or recall representing the proportion of actual positive cases correctly identified by a classifier (Fawcett, 2006). Fraud detection within transaction networks also depends on behavioral information accumulated across transactions because individual transaction characteristics may provide insufficient evidence when examined without historical context (Whitrow et al., 2009). Real-world credit-card fraud research has similarly identified class imbalance, nonstationary transaction distributions, delayed fraud verification, and changing fraudulent behavior as central characteristics of operational fraud-detection environments (Pozzolo et al., 2018). At the international level, these analytical conditions are relevant to banks, payment networks, digital-wallet providers, financial technology companies, online merchants, card issuers, transaction processors, and other organizations operating high-volume financial infrastructures where rare fraudulent transactions must be distinguished from much larger legitimate transaction streams.


Class imbalance in fraud analytics represents more than a simple numerical disproportion between fraudulent and legitimate transactions because the internal characteristics of minority observations directly affect the ability of machine-learning algorithms to construct effective classification boundaries. Minority-class observations may occur near majority-class regions, within isolated subgroups, as unusual outliers, or within overlapping areas where fraudulent and legitimate transaction patterns possess similar characteristics. Research on imbalanced learning has demonstrated that minority examples located near class boundaries may require specialized treatment because they are more difficult to classify accurately than minority observations situated within relatively homogeneous regions (Han et al., 2005). This recognition contributed to the development of adaptive sampling methods in which additional minority examples are generated according to the learning difficulty associated with particular observations rather than through uniform replication of minority data (He et al., 2008). Majority-class reduction has also been investigated as an alternative balancing strategy, with exploratory undersampling approaches showing that selective removal of majority observations can improve minority-class recognition when substantial imbalance exists (Liu et al., 2009). Hybrid approaches have combined sampling procedures with ensemble classifiers, allowing the classification process to address class distribution and model learning simultaneously. RUSBoost, for

example, integrates random undersampling with boosting so that majority-class reduction occurs within an ensemble-learning procedure designed to strengthen minority classification (Seiffert et al., 2010). Ensemble approaches to imbalance learning have subsequently incorporated bagging, boosting, sampling, and hybrid procedures because different classifiers can respond differently to rare observations and overlapping class distributions (Galar et al., 2012). The difficulty of an imbalanced classification task is therefore influenced not only by the numerical imbalance ratio but also by class overlap, noise, sample size, small disjuncts, and the distribution of individual minority examples within the feature space (López et al., 2013). Minority observations have been categorized as safe, borderline, rare, and outlier examples, highlighting important differences in the difficulty associated with learning different regions of a minority distribution (Napierala & Stefanowski, 2016). These distinctions are particularly relevant to fraud detection because fraudulent transactions do not constitute a homogeneous category. Some fraudulent transactions may strongly resemble normal customer activity, while others may display highly unusual patterns or appear in isolated transactional regions. Research examining the development of synthetic minority oversampling has consequently emphasized that resampling performance depends on the manner in which minority examples are selected and generated rather than simply on increasing the numerical size of the minority class (Fernández et al., 2018). Adaptive minority-class resampling in the present research therefore represents the capability of a fraud-learning system to modify the representation of fraudulent and legitimate transactions according to observed class distributions, minority difficulty, and learning requirements.

The unequal consequences associated with fraud-classification errors constitute another fundamental dimension of adaptive class-imbalance learning. In digital transaction networks, a false negative occurs when a genuinely fraudulent transaction is classified as legitimate, whereas a false positive occurs when a legitimate transaction is incorrectly identified as fraudulent. These two types of classification error do not carry equivalent operational costs. Missed fraudulent transactions may produce direct monetary losses, account compromise, unauthorized transfers, customer disputes, investigation costs, and other forms of financial exposure, while false-positive classifications can increase manual review requirements, interrupt legitimate transactions, and create unnecessary verification activity. Cost-sensitive learning addresses this asymmetry by assigning different penalties to classification errors so that the learning process places greater emphasis on observations whose incorrect classification is associated with greater cost. Cost-sensitive boosting methods have demonstrated that imbalance-oriented classifiers can modify learning weights according to unequal classification costs and thereby place stronger emphasis on difficult minority observations (Sun et al., 2007). Fraud-specific analysis has similarly shown that classification models can incorporate transaction-dependent misclassification costs instead of treating every fraudulent event as carrying an identical economic consequence (Sahin et al., 2013). Adaptive class-imbalance learning also involves the threshold through which predicted probabilities or classification scores are converted into final fraud decisions. A probability-based classifier may assign each transaction a fraud score, while the classification threshold determines the score at which the transaction is labeled fraudulent. Altering this threshold directly changes the relationship among true positives, false positives, true negatives, and false negatives, making threshold selection particularly important when high sensitivity is prioritized. Receiver operating characteristic analysis formally represents the relationship between sensitivity and false-positive rate across alternative classification thresholds (Fawcett, 2006). Precision-recall analysis provides another important evaluation mechanism when positive observations are rare because changes in minority-class prediction can be more visible through precision and recall than through measures dominated by the majority class (Davis & Goadrich, 2006; Saito & Rehmsmeier, 2015). Resampling procedures can additionally change the apparent class distribution used during training and therefore affect estimated posterior probabilities, requiring appropriate calibration when classifiers trained on altered distributions are applied to operational data (Pozzolo et al., 2015). Experimental research on deep neural classifiers has also demonstrated that classification thresholds and sampling strategies substantially influence performance under imbalanced conditions (Buda et al., 2018). Ensemble learning provides another mechanism through which classification robustness can be strengthened by

integrating several models rather than relying on a single classifier. Credit-card fraud studies have examined AdaBoost and majority-voting systems for combining classifier decisions across highly imbalanced transaction datasets (Randhawa et al., 2018). Other research has combined random forests, neural networks, and related ensemble mechanisms to improve recognition of rare fraudulent observations (Sohony et al., 2018). Cost-sensitive boosting has also continued to receive statistical formulation for classification environments where class frequencies and misclassification consequences are unequal (Bei, 2021).

Figure 1: Adaptive Class-Imbalance Learning and High-Sensitivity Fraud Detection Performance



The scale, velocity, heterogeneity, and behavioral structure of digital transactions further complicate high-sensitivity fraud detection because operational systems process large transaction streams while confirmed fraud labels remain comparatively scarce. Machine-learning models used in real financial environments must frequently evaluate transaction characteristics alongside account histories, merchant attributes, geographic patterns, temporal information, transaction frequency, spending behavior, and other contextual indicators. Operational research on credit-card fraud detection has shown that realistic evaluation requires attention to skewed class distributions, changing transaction populations, delayed feedback, and the distinction between laboratory classification conditions and actual fraud-monitoring environments (Pozzolo et al., 2014). Feature engineering is particularly important within such systems because raw transactional variables may fail to capture the behavioral patterns associated with account misuse or evolving fraud. Aggregated variables describing customer spending behavior, transaction frequency, monetary patterns, and historical activity can provide additional discriminatory information for machine-learning classifiers (Bahnsen et al., 2016). The large volume of transaction data also creates computational requirements that extend beyond conventional offline fraud modeling. Distributed fraud-detection architectures have therefore been designed to combine large-scale data processing with machine-learning classification while accounting for class imbalance, delayed feedback, and nonstationary transaction streams (Carcillo et al., 2018). Sequential transaction behavior represents another important source of information because fraud may become recognizable through a sequence of actions even when individual transactions do not appear strongly abnormal. Sequence-based classification has demonstrated that temporal transaction histories can contribute useful fraud-related information beyond independent transaction-level representations (Jurgovsky et al., 2018). Synthetic minority generation has also expanded beyond traditional interpolation-based methods. Generative adversarial networks have been used to generate additional fraudulent observations for classifier training, illustrating how artificial minority examples can be incorporated into fraud-classification environments affected by substantial class imbalance (Fiore et al., 2019). Broader investigations of deep learning under imbalanced distributions have documented

extensive use of resampling, algorithm-level modification, cost-sensitive training, ensemble mechanisms, and specialized evaluation measures for tasks in which minority observations carry high analytical importance (Johnson & Khoshgoftaar, 2019). Fraud analytics can also combine supervised classification with unsupervised anomaly information so that known fraud labels are evaluated together with unusual behavioral structures in transactional data (Carcillo et al., 2021). Large-scale digital fraud increasingly contains relational characteristics as well because accounts, devices, merchants, payment destinations, and transactions can form interconnected graph structures. Fraud-detection research using graph neural networks has shown that imbalanced labels can weaken representation learning for rare fraudulent entities and that specialized sampling can strengthen fraud-node identification in relational datasets (Liu et al., 2021). These characteristics position large-scale digital transaction networks as complex classification environments in which minority scarcity, transaction volume, behavioral dependence, relational structure, misclassification costs, and evolving data distributions jointly shape fraud-detection performance.

Adaptive class-imbalance learning in this study is defined as a coordinated machine-learning capability through which minority-class representation, asymmetric classification costs, fraud-decision thresholds, and multiple classifier outputs are adjusted to strengthen the identification of rare fraudulent transactions in heavily imbalanced digital transaction environments. Research on imbalance learning has established that class imbalance can be addressed through data-level procedures, algorithm-level modifications, cost-sensitive mechanisms, ensemble architectures, and combinations of these approaches rather than through a single universally effective technique (Krawczyk, 2016). Comprehensive reviews have also shown that successful classification under imbalance depends on the interaction among data characteristics, sampling procedures, classifier selection, cost structures, and evaluation metrics (Guo et al., 2017). Within the present study, adaptive minority-class resampling constitutes the first explanatory construct and refers to the systematic modification of fraudulent and legitimate class representation according to minority scarcity, distributional characteristics, and classification difficulty. Cost-sensitive learning constitutes the second explanatory construct and refers to the incorporation of unequal misclassification consequences into model training so that errors involving fraudulent observations receive appropriate analytical weight. Dynamic classification-threshold optimization represents the third construct and refers to adjustment of the decision boundary used to convert transaction risk scores into fraudulent or legitimate classifications, with particular attention to maximizing fraud sensitivity while controlling false-positive activity. Adaptive ensemble learning represents the fourth construct and describes the coordinated combination, selection, or weighting of multiple classifiers for identifying heterogeneous fraud patterns under imbalanced conditions. High-sensitivity fraud detection performance serves as the dependent construct and refers primarily to the system's ability to identify genuine fraud cases accurately, maintain strong minority-class recall, and reduce the occurrence of fraudulent transactions incorrectly classified as legitimate. Performance assessment under such conditions requires metrics that distinguish minority-class recognition from majority-dominated aggregate accuracy, particularly when fraudulent observations constitute only a very small proportion of total cases (Saito & Rehmsmeier, 2015). Fraud-detection research based on realistic transaction environments has further demonstrated that classification performance must be interpreted in relation to unequal class distributions, temporal changes, feedback delays, and operational decision processes (Pozzolo et al., 2018). Accordingly, the present quantitative, cross-sectional, case-study-based research examines measurable perceptions of adaptive minority-class resampling, cost-sensitive learning, dynamic classification-threshold optimization, and adaptive ensemble learning and statistically evaluates their relationships with high-sensitivity fraud detection performance in large-scale digital transaction networks using a five-point Likert-scale instrument, descriptive statistics, correlation analysis, and multiple regression modeling.

Background of the Study

The rapid expansion of digital banking, electronic commerce, mobile payment systems, fintech services, card-based transactions, digital wallets, and real-time payment infrastructures has transformed the scale and complexity of financial transactions across modern economies. Large-scale digital transaction networks now process enormous volumes of financial activity within short periods, requiring

organizations to identify fraudulent behavior without interrupting legitimate customer transactions. Fraud detection in these environments is fundamentally difficult because fraudulent transactions normally constitute only a very small proportion of the total transaction population. This unequal distribution between legitimate and fraudulent observations creates the class-imbalance problem, in which conventional machine-learning algorithms may become strongly influenced by the majority legitimate class and fail to recognize rare but financially significant fraudulent events. A model can therefore achieve apparently high overall classification accuracy while producing an unacceptable number of false negatives, meaning that genuine fraudulent transactions remain undetected. High-sensitivity fraud detection has consequently become an important analytical requirement because it emphasizes the successful identification of fraudulent transactions and the reduction of missed fraud cases rather than relying exclusively on overall accuracy. Adaptive class-imbalance learning provides a broader approach to this problem by enabling machine-learning systems to respond to unequal class distributions through multiple complementary mechanisms. Adaptive minority-class resampling can strengthen the representation of rare fraudulent observations within training data, while cost-sensitive learning can assign greater penalties to fraud-related classification errors. Dynamic classification-threshold optimization can adjust the decision boundary used to classify transactions according to fraud-detection priorities, and adaptive ensemble learning can integrate multiple classifiers to improve recognition of heterogeneous and complex fraud patterns. These capabilities are particularly relevant in large-scale digital transaction networks where transaction characteristics, user behaviors, fraud strategies, and risk distributions may vary considerably across operational environments. The present study therefore focuses on adaptive class-imbalance learning as a multidimensional fraud-detection capability consisting of adaptive minority-class resampling, cost-sensitive learning, dynamic classification-threshold optimization, and adaptive ensemble learning. Through a quantitative, cross-sectional, case-study-based investigation, the research examines how these capabilities are associated with high-sensitivity fraud detection performance and whether their combined application can provide a stronger statistical explanation of fraud-detection effectiveness within large-scale digital transaction networks.

Problem Statement

The central problem addressed by this study concerns the difficulty of achieving high-sensitivity fraud detection when fraudulent transactions represent a very small minority of observations within large-scale digital transaction networks. Traditional classification models are commonly optimized according to aggregate predictive accuracy, yet such accuracy can become misleading when legitimate transactions overwhelmingly dominate the dataset. A classifier may correctly label most legitimate transactions and still perform poorly in identifying the fraudulent cases that create the greatest financial and operational risk. This imbalance can produce excessive false negatives, weak minority-class recall, unstable fraud-detection performance, and an inappropriate classification boundary that favors the majority class. The problem is further complicated by the heterogeneity of fraud patterns, differences in the financial consequences of classification errors, and the inability of a single static balancing method to address every form of minority-class difficulty. Conventional oversampling or undersampling procedures may improve class proportions but may not sufficiently account for differences in observation difficulty, misclassification costs, decision thresholds, or classifier diversity. Similarly, a fraud-detection system that applies a fixed probability threshold may fail to achieve the level of sensitivity required when the consequences of missed fraud are substantially greater than the costs associated with reviewing suspicious legitimate transactions. Ensemble classifiers can improve predictive stability, yet their effectiveness may also depend on how imbalance is handled throughout the learning process. These limitations create a need to examine class-imbalance learning as an integrated adaptive capability rather than as an isolated preprocessing technique. However, there remains a need for quantitative evidence showing how adaptive minority-class resampling, cost-sensitive learning, dynamic classification-threshold optimization, and adaptive ensemble learning individually and collectively relate to high-sensitivity fraud detection performance within large-scale digital transaction networks. The present study addresses this problem by statistically evaluating these four adaptive class-imbalance learning constructs and determining the extent to which they explain

variations in high-sensitivity fraud detection performance through descriptive statistics, correlation analysis, and multiple regression modeling.

Purpose/Objectives of the Study

The primary objective of this study is to quantitatively assess the influence of adaptive class-imbalance learning on high-sensitivity fraud detection performance within large-scale digital transaction networks. To accomplish this overall objective, the study examines four interconnected dimensions of adaptive imbalance management that represent different stages of the fraud-learning process. The first objective is to determine the extent to which adaptive minority-class resampling contributes to high-sensitivity fraud detection by improving the representation of rare fraudulent observations and reducing excessive dominance of the legitimate transaction class. The second objective is to examine the influence of cost-sensitive learning on fraud detection by evaluating whether greater consideration of unequal misclassification consequences is associated with stronger identification of fraudulent transactions. The third objective is to assess the relationship between dynamic classification-threshold optimization and high-sensitivity fraud detection performance, with particular attention to the role of flexible decision boundaries in increasing the identification of fraud cases while maintaining an operationally acceptable classification process. The fourth objective is to evaluate the effect of adaptive ensemble learning on fraud-detection performance by examining whether coordinated use of multiple classifiers strengthens recognition of diverse and complex minority fraud patterns. The fifth objective is to determine the collective predictive contribution of adaptive minority-class resampling, cost-sensitive learning, dynamic classification-threshold optimization, and adaptive ensemble learning to high-sensitivity fraud detection performance. The research further seeks to identify the relative statistical contribution of each adaptive learning dimension so that the strongest predictors of high-sensitivity fraud detection can be distinguished. These objectives are examined through a quantitative, cross-sectional, case-study-based research design using a structured five-point Likert-scale instrument. Descriptive statistics are used to establish the observed levels of the study constructs, correlation analysis is applied to assess the direction and strength of relationships among the variables, and multiple regression modeling is used to determine their individual and combined predictive effects on high-sensitivity fraud detection performance.

Research Questions

RQ1: To what extent do adaptive class-imbalance learning techniques influence high-sensitivity fraud detection in large-scale digital transaction networks?

RQ2: What relationships exist between adaptive minority-class resampling, cost-sensitive learning, dynamic classification-threshold optimization, adaptive ensemble learning, and high-sensitivity fraud detection performance?

RQ3: To what extent do adaptive class-imbalance learning constructs collectively predict high-sensitivity fraud detection performance?

Research Hypotheses

The research hypotheses have been formulated to statistically test the proposed relationships between adaptive class-imbalance learning capabilities and high-sensitivity fraud detection performance in large-scale digital transaction networks. The first hypothesis, H1, proposes that adaptive minority-class resampling has a significant positive effect on high-sensitivity fraud detection performance, based on the expectation that a stronger representation of difficult and underrepresented fraudulent observations can improve the ability of classification systems to recognize minority-class events. The second hypothesis, H2, proposes that cost-sensitive learning has a significant positive effect on high-sensitivity fraud detection performance because assigning greater importance to costly fraud-related classification errors can encourage the learning process to prioritize the correct identification of fraudulent transactions. The third hypothesis, H3, proposes that dynamic classification-threshold optimization has a significant positive effect on high-sensitivity fraud detection performance, reflecting the expectation that flexible adjustment of fraud-decision thresholds can increase sensitivity and reduce the number of fraudulent transactions incorrectly categorized as legitimate. The fourth hypothesis, H4, proposes that adaptive ensemble learning has a significant positive effect on high-sensitivity fraud detection performance because the combination of multiple classifiers can capture diverse transaction

patterns and reduce dependence on the limitations of a single predictive model. Finally, H5 proposes that adaptive minority-class resampling, cost-sensitive learning, dynamic classification-threshold optimization, and adaptive ensemble learning collectively have a significant positive effect on high-sensitivity fraud detection performance. The combined hypothesis recognizes adaptive class-imbalance learning as a multidimensional capability in which data-level balancing, error-cost management, decision-level optimization, and classifier integration operate as related components of fraud classification. These hypotheses will be tested statistically through Pearson correlation analysis and multiple regression modeling, with the significance, direction, and magnitude of the relationships used to determine whether each proposed hypothesis is supported by the empirical results.

H1: Adaptive minority-class resampling has a significant positive effect on high-sensitivity fraud detection performance.

H2: Cost-sensitive learning has a significant positive effect on high-sensitivity fraud detection performance.

H3: Dynamic classification-threshold optimization has a significant positive effect on high-sensitivity fraud detection performance.

H4: Adaptive ensemble learning has a significant positive effect on high-sensitivity fraud detection performance.

H5: Adaptive class-imbalance learning constructs collectively have a significant positive effect on high-sensitivity fraud detection performance in large-scale digital transaction networks.

Significance of the Research

i. Significance for financial institutions and digital payment organizations: The study is significant for banks, fintech companies, card issuers, payment processors, digital-wallet providers, e-commerce platforms, and other organizations that depend on automated transaction monitoring. By examining factors associated with high-sensitivity fraud detection, the research provides an empirical basis for understanding which adaptive learning capabilities are most strongly related to identifying rare fraudulent transactions within highly imbalanced transaction environments.

ii. Significance for fraud analysts and risk-management professionals: The research is relevant to fraud investigators, transaction-risk analysts, cybersecurity teams, compliance personnel, and operational managers because it evaluates the relative roles of adaptive resampling, cost-sensitive learning, threshold optimization, and ensemble learning. The statistical results can help clarify which aspects of imbalance management deserve stronger attention when fraud systems are evaluated or redesigned.

iii. Significance for machine-learning and data-science practice: The study treats class imbalance as a multidimensional learning problem rather than a single sampling problem. It therefore provides a structured empirical framework linking data-level, algorithm-level, decision-level, and ensemble-level approaches to high-sensitivity fraud detection performance.

iv. Significance for academic research: The study contributes to the quantitative examination of class-imbalance learning within fraud analytics by integrating four adaptive constructs into a unified explanatory model. The use of descriptive statistics, correlation analysis, and regression modeling enables the relationships among these constructs to be examined systematically within a case-study context.

v. Significance for digital transaction security: High-sensitivity fraud detection is directly connected to the ability of transaction networks to identify rare but harmful events. By concentrating specifically on sensitivity and minority-class recognition, the research emphasizes an aspect of fraud-detection performance that can be obscured when evaluation relies heavily on aggregate accuracy. The study therefore provides a focused framework for examining how adaptive learning capabilities correspond with the effective identification of fraudulent transactions in large-scale digital financial environments.

LITERATURE REVIEW

The literature on fraud detection and class-imbalance learning encompasses a broad body of research concerned with machine-learning classification, rare-event detection, transaction analytics, cost-sensitive decision making, sampling techniques, ensemble methods, threshold selection, and performance evaluation. Fraud detection in large-scale digital transaction networks represents a distinctive classification problem because fraudulent observations normally constitute only a small

percentage of total transactional activity, creating a highly skewed distribution between legitimate and fraudulent classes. Under these conditions, conventional classifiers may become excessively oriented toward majority-class recognition and generate weak sensitivity toward the minority fraud class. The literature has therefore examined numerous approaches for improving minority-class learning, including oversampling, undersampling, synthetic sample generation, adaptive resampling, cost-sensitive classification, algorithm-level modification, decision-threshold adjustment, bagging, boosting, hybrid ensembles, and combinations of these methods. Research has also emphasized that class imbalance cannot be understood only through the numerical ratio between classes because minority-class complexity, overlap, noisy observations, rare subgroups, changing transaction patterns, and unequal error consequences can substantially influence classification performance. Within financial fraud detection, these challenges are intensified by the scale, velocity, heterogeneity, and sequential nature of digital transactions and by the financial significance of false-negative classifications. The literature reviewed in this study is therefore organized around the central theoretical and empirical dimensions needed to explain adaptive class-imbalance learning and high-sensitivity fraud detection. It first examines fraud-detection challenges within large-scale digital transaction networks and then addresses class imbalance and minority-class learning. Subsequent sections analyze adaptive resampling and cost-sensitive learning, followed by dynamic classification-threshold optimization and adaptive ensemble learning. Statistical Learning Theory is then established as the principal theoretical framework for explaining learning behavior under unequal class distributions, while the final subsection develops the conceptual framework linking adaptive minority-class resampling, cost-sensitive learning, dynamic classification-threshold optimization, and adaptive ensemble learning with high-sensitivity fraud detection performance.


Fraud Detection in Large-Scale Digital Transaction Networks

Fraud detection in large-scale digital transaction networks has developed into a complex analytical problem because financial activity is increasingly executed through interconnected banking platforms, card networks, mobile applications, e-commerce systems, digital wallets, payment gateways, and fintech infrastructures. These environments generate continuous streams of heterogeneous transactions that differ in amount, location, device, merchant category, customer behavior, temporal pattern, and authorization pathway. The central difficulty is that fraudulent activity is intentionally designed to resemble legitimate behavior closely enough to avoid conventional controls, while the volume of normal transactions greatly exceeds the number of confirmed fraud cases. Financial fraud detection therefore requires methods capable of identifying subtle deviations without treating every unusual transaction as fraudulent. Earlier data-mining research categorized financial fraud detection into classification, regression, clustering, prediction, outlier detection, and visualization approaches, illustrating that no single analytical technique adequately captures every type of fraudulent behavior (Debasish, 2021; Ngai et al., 2011). The challenge becomes greater in large-scale transaction networks because the same infrastructure may simultaneously contain account takeover, stolen-card activity, synthetic identity behavior, merchant fraud, transaction laundering, unauthorized transfers, and coordinated attacks. Automated detection must process these patterns rapidly while preserving reliable decisions across millions of legitimate events. Computational intelligence research has also emphasized that financial fraud occurs across multiple domains and that increasing dependence on digital technologies has reduced the practicality of manual detection as a primary control mechanism (Abdur & Iftekhhar, 2021; West & Bhattacharya, 2016). Consequently, fraud detection in large-scale networks is not merely a problem of finding anomalous transactions. It requires the integration of data processing, behavioral modeling, statistical classification, risk scoring, and operational decision rules across transaction environments in which fraud patterns may be sparse, concealed, heterogeneous, and difficult to label accurately. These characteristics make detection quality dependent on both computational scalability and the capacity to distinguish rare malicious behavior from legitimate behavioral variation.

A major challenge within large-scale digital transaction networks is that fraud detection systems operate under conditions that violate many assumptions associated with conventional supervised classification. Confirmed fraudulent transactions are scarce relative to legitimate observations, labels

may arrive only after customer complaints or investigations, and fraudulent behavior can change when offenders respond to existing security controls. Transaction data may also arrive at high velocity, which restricts the time available for feature extraction, model scoring, alert generation, and human review. A comprehensive survey of fraud detection systems identified skewed class distributions, large data volumes, real-time processing requirements, and concept drift among the principal factors that hinder reliable fraud detection performance (Abdallah et al., 2016; Debasish, 2022). Concept drift is particularly important because relationships learned from historical transactions may weaken when customer purchasing habits, payment technologies, merchant environments, or criminal strategies change. The resulting mismatch between training information and current transaction behavior can reduce sensitivity to emerging fraud patterns. Another problem concerns hidden behavioral structure. Fraudulent activity may not be recognizable from a single transaction attribute, yet it can become more apparent when spending sequences, device reuse, repeated merchant interactions, transaction timing, or changes in normal customer behavior are considered together. Real-time fraud research has therefore examined deep ensemble architectures capable of processing transaction streams while accounting for imbalanced data and latent spending patterns (Arya & Sastry, 2020; Ashfaq & Mohammad, 2022). Large-scale implementation additionally creates a tension between sensitivity and operational workload. Lowering a decision threshold may identify more fraudulent cases but can simultaneously produce a large number of false alerts. Each false alert can require investigation, delay payment authorization, inconvenience legitimate customers, or increase monitoring costs. Fraud detection systems must therefore manage minority recognition, temporal adaptation, computational speed, decision thresholds, and alert volume simultaneously rather than treating predictive accuracy as an isolated objective within everyday operational environments.

Figure 2: Fraud Detection Challenges Framework for Large-Scale Digital Transaction Networks



The difficulty of fraud detection is shaped by the interaction between transaction-network scale and the statistical characteristics of fraudulent observations. As digital transaction ecosystems expand, fraud models encounter greater diversity in customer profiles, merchants, devices, channels, geographic locations, transaction values, and behavioral sequences. Such diversity increases legitimate variation and makes simple anomaly rules less reliable because unusual behavior is not automatically fraudulent. At the same time, fraudsters may distribute malicious activity across accounts or transactions to avoid fixed rules, creating patterns that are weak at the individual-transaction level but meaningful when viewed collectively. Effective detection therefore requires models that can learn nonlinear relationships while remaining sensitive to a minority class that may represent only a small

portion of recorded activity. Comparative work on credit-card fraud classification has shown that logistic regression, nearest-neighbor methods, random forests, neural networks, AdaBoost, and ensemble procedures can produce materially different detection results when applied to highly skewed transaction data, confirming that model choice and imbalance treatment strongly influence fraud recognition (Bagga et al., 2020; Sheak & Risha, 2022). This creates an additional evaluation challenge. Overall accuracy can remain high even when a model misses many fraudulent transactions because legitimate observations dominate the calculation. High-sensitivity fraud detection therefore requires attention to the successful recognition of actual fraud cases, false-negative reduction, and the balance between detection gains and false-positive burden. Large transaction networks also demand consistency across changing volumes and diverse operating conditions, because a model that performs adequately within a small historical sample may behave differently when deployed across broader transaction streams. The literature on digital fraud detection thus identifies a tightly connected set of challenges involving data scale, class imbalance, behavioral heterogeneity, evolving fraud strategies, delayed labels, real-time processing, model selection, and performance measurement, all of which directly shape the capacity of machine-learning systems to recognize rare fraudulent activity.


Class Imbalance and Minority-Class Learning in Fraud Detection

Class imbalance and minority-class learning constitute a central analytical problem in fraud detection because the class that is most important to identify is usually the class that appears least frequently in the available data. In large-scale digital transaction networks, legitimate transactions can number in the hundreds of thousands or millions, whereas confirmed fraudulent transactions may represent only a very small proportion of the total observations. This skewed distribution affects model training because many conventional classifiers attempt to minimize overall prediction error and can therefore obtain apparently strong performance by concentrating on the dominant legitimate class. Under these conditions, the minority fraud class receives insufficient representation during learning, which can weaken the classifier's ability to recognize rare fraudulent patterns, particularly when fraud observations occur close to legitimate cases in the feature space. Imbalanced learning research has emphasized that the problem is not limited to unequal class frequencies; the rarity of important events, different error costs, unsuitable evaluation measures, and the characteristics of the underlying data distribution jointly determine the difficulty of predictive modeling (Branco et al., 2016; Rashedul & Kaium, 2022). This perspective is directly applicable to fraud analytics because a very small fraud class can still carry the highest operational importance. Minority-class learning therefore requires attention to how fraudulent observations are represented, weighted, separated, and evaluated throughout the modeling process. Sensitivity, recall, precision, F-measure, geometric mean, and area-based measures are generally more informative than raw accuracy when evaluating rare positive cases because they expose weaknesses that majority-dominated accuracy can conceal. In practical terms, a fraud classifier that labels almost every transaction as legitimate may appear highly accurate in a severely skewed dataset while failing at its principal security function. The class-imbalance problem therefore changes both the learning objective and the interpretation of model performance, requiring fraud-detection systems to focus explicitly on the quality of minority-class recognition rather than on aggregate correctness alone.

Minority-class learning becomes more difficult when imbalance interacts with other characteristics such as class overlap, noisy attributes, sparse minority regions, complex decision boundaries, and heterogeneous fraud behaviors. A small fraud class is not necessarily difficult solely because it is small; difficulty increases when fraudulent and legitimate transactions share similar values across important predictive features or when minority observations are scattered across several weakly connected regions. Data-complexity research has demonstrated that class imbalance can intensify the adverse effects of overlap and difficult decision structures, meaning that the imbalance ratio by itself is an incomplete description of classification difficulty (Barella et al., 2021; Rashedul & Kaium, 2022). In digital fraud environments, this issue is especially important because sophisticated fraudulent transactions are often constructed to resemble normal customer activity. A fraudulent payment can use a familiar merchant category, a plausible amount, an expected device, or an apparently ordinary time pattern, causing minority and majority observations to occupy overlapping regions. Machine-learning approaches therefore need to increase the informational contribution of fraudulent examples without

producing an artificial boundary that generates excessive false positives. Empirical credit-card research has shown that combining imbalance treatment with suitable classifiers can materially improve fraud-class recognition; for example, SMOTE-based balancing combined with LightGBM has been applied to address skewed transaction distributions while maintaining strong recall for fraudulent cases (Fang et al., 2019; Chowdhury & Tonay, 2022). Such evidence illustrates that minority-class learning is closely connected to both data preparation and classifier behavior. Oversampling can strengthen minority representation, undersampling can reduce majority dominance, and hybrid approaches can seek a more balanced training distribution, but each intervention changes the information available to the model. Excessive synthetic generation may create redundant or poorly located observations, while aggressive undersampling may remove legitimate patterns that are necessary for constructing reliable decision boundaries. Effective fraud learning consequently depends on preserving the structure of legitimate transactions while creating sufficient discriminatory emphasis on the rare fraud class.

Figure 3: Class Imbalance and Minority-Class Learning in Fraud Detection



The treatment of minority-class observations also influences feature selection, model calibration, decision boundaries, and the stability of fraud predictions across operational datasets. In high-dimensional transaction data, imbalance can cause feature-selection procedures to favor attributes that explain majority behavior while overlooking variables that are particularly useful for identifying fraudulent cases. Hybrid fraud-detection models have therefore combined imbalance correction with feature selection and hyperparameter optimization so that the classifier is trained on a representation that is both discriminative and better balanced; one such approach integrated recursive feature elimination, parameter optimization, and SMOTE to strengthen credit-card fraud identification across real-world datasets (Hossan, 2022; Rtayli & Enneya, 2020). Yet the minority-class problem becomes even more difficult when fraud observations overlap with legitimate transactions. In this situation, simple resampling may enlarge the minority class without resolving ambiguity near the class boundary. Research specifically addressing credit-card fraud has consequently proposed hybrid procedures that identify overlapping regions and dynamically adjust model selection according to the difficulty of those regions, demonstrating that imbalance and overlap should be treated jointly rather than independently (Li et al., 2021; Nishat & Kaniz, 2022). For the present study, these findings support a multidimensional understanding of minority-class learning in which successful high-sensitivity fraud detection depends on more than merely changing the numerical ratio between classes. Adaptive


learning must account for which minority cases are difficult, where they are located relative to legitimate observations, how costly their misclassification is, and how the model responds when minority information is sparse or ambiguous. This perspective connects directly with adaptive minority-class resampling, cost-sensitive learning, dynamic threshold optimization, and adaptive ensemble learning because each mechanism addresses a different source of minority-class weakness. In large-scale digital transaction networks, the central analytical requirement is therefore to improve the probability that rare fraudulent events are learned and detected without creating an operationally unacceptable burden of false alerts among the overwhelmingly larger population of legitimate transactions.

Cost-Sensitive Learning for Fraud Detection

Adaptive resampling represents a data-level response to the severe class imbalance that characterizes fraud detection in large-scale digital transaction networks. In such environments, legitimate transactions dominate the training population while confirmed fraudulent transactions occupy only a small minority, causing many standard classifiers to develop decision boundaries that favor the majority class. Resampling attempts to change the effective training distribution so that minority fraud observations exert greater influence during model construction. Oversampling increases the representation of minority cases, undersampling reduces the dominance of legitimate cases, and hybrid procedures combine both operations to create a more informative learning set. The usefulness of these strategies, however, depends strongly on the severity of imbalance and on the behavior of the selected classifier. Experimental evidence across multiple real datasets has shown that preprocessing performance changes as the imbalance ratio changes, with oversampling becoming particularly valuable when the minority class is extremely scarce and with classifier choice affecting the relative value of alternative balancing procedures (García et al., 2012; Shaiful et al., 2022). Adaptive resampling extends this logic by avoiding the assumption that every fraud observation should receive identical sampling attention. Difficult minority transactions located close to legitimate regions may need stronger representation than fraud cases already positioned in well-separated areas. Ranked Minority Oversampling in Boosting was developed around this principle by assigning sampling probabilities to minority observations according to their learning difficulty and by updating this emphasis during boosting iterations, allowing the decision boundary to move toward difficult examples rather than treating the minority population uniformly (Mohammad, 2023; Chen et al., 2010). For fraud detection, this adaptive perspective is important because fraudulent transactions vary substantially in their similarity to legitimate activity. Static duplication can amplify easy cases without strengthening the model where errors are most likely, whereas difficulty-sensitive resampling can concentrate learning capacity on ambiguous and underrepresented fraud patterns within operational transaction systems. The quality of adaptive resampling also depends on where synthetic minority observations are generated because an indiscriminate increase in fraud-class size can introduce noise, overlap, or artificial patterns that reduce classification quality. In digital transaction data, fraud examples may form several small behavioral clusters associated with different attack mechanisms, account types, merchants, devices, locations, or transaction sequences. A sampling method that interpolates across unsuitable regions can create synthetic observations that lie too close to legitimate transactions or that fail to represent meaningful fraud structures. Majority Weighted Minority Oversampling Technique was designed to address this problem by first identifying informative minority observations, assigning them weights according to the surrounding majority-class structure, clustering minority examples, and then generating synthetic samples within relevant minority regions (Barua et al., 2014; Kaium, 2023). This approach illustrates a key principle for the present research: adaptive minority-class resampling should be understood as selective representation rather than simple numerical balancing. For high-sensitivity fraud detection, the purpose of resampling is not to force an artificial fifty-fifty class ratio, but to provide the classifier with sufficient information about difficult fraud cases while preserving the statistical structure necessary to distinguish them from legitimate behavior. This distinction matters in large-scale transaction networks because aggressive undersampling can discard legitimate examples that define normal behavior, while excessive oversampling can create redundant synthetic cases and increase the risk of overfitting. An effective adaptive procedure must therefore consider the location, density, difficulty, and informational value of fraud observations before modifying the training

distribution. It should also remain compatible with downstream classifiers, since resampling changes the empirical class prior and can influence estimated probabilities, model calibration, and the placement of the decision boundary. Adaptive resampling is therefore closely connected to threshold selection and ensemble behavior, even when it is implemented as a data-level technique, because changes made during training affect how later fraud scores are interpreted.

Figure 4: Adaptive Resampling and Cost-Sensitive Learning Framework for High-Sensitivity Fraud Detection




Cost-sensitive learning addresses a different but complementary limitation of imbalanced fraud data by recognizing that classification errors do not carry equal operational consequences. In a conventional accuracy-oriented classifier, a false negative and a false positive may contribute equally to the training objective even though their financial meanings are very different. A missed fraudulent transaction can generate direct monetary loss, chargeback exposure, account compromise, investigation costs, and customer harm, whereas a false alert usually creates review workload, transaction friction, or temporary inconvenience. Cost-sensitive fraud modeling incorporates this asymmetry into learning or decision making so that the classifier is encouraged to avoid the errors that are economically more damaging. Research using real credit-card transaction data demonstrated that Bayes minimum-risk decision making can incorporate transaction-specific financial costs and can outperform purely accuracy-oriented comparison procedures when performance is evaluated in monetary terms (Bahnsen et al., 2013; Chowdhury, 2023). The importance of example-specific costs is especially relevant in large-scale networks because a missed low-value fraud and a missed high-value fraud should not necessarily receive identical penalties. This principle was extended through example-dependent cost-sensitive decision trees in which transaction-level costs were integrated directly into the tree-building process through cost-based impurity and pruning criteria, producing models oriented toward financial savings rather than uniform error minimization (Bahnsen et al., 2015; Badrul & Mominul, 2023). Within the present study, cost-sensitive learning therefore represents the capability to assign greater learning importance to fraud-related errors according to their risk and economic consequences. When combined conceptually with adaptive minority-class resampling, it creates two complementary forms of imbalance control: resampling changes which observations receive representation during learning, while cost sensitivity changes how strongly classification errors are penalized. Their joint relevance to high-sensitivity fraud detection lies in strengthening minority recognition without assuming that all observations, all fraud cases, or all misclassification outcomes should be treated as equally important within the learning process.

Dynamic Classification and Adaptive Ensemble Learning

Dynamic classification-threshold optimization is a critical decision-level mechanism for high-sensitivity fraud detection because the probability produced by a classifier does not itself determine

whether a transaction will be labeled fraudulent. A decision threshold converts a fraud score into a categorical outcome, and the conventional value of 0.50 is not inherently appropriate when fraudulent transactions form a very small minority of observations. In highly imbalanced digital transaction networks, a fixed threshold can favor the legitimate class, producing excellent overall accuracy while allowing a substantial proportion of fraudulent cases to remain below the decision boundary. Lowering or otherwise optimizing the threshold can increase sensitivity by converting more high-risk minority observations into positive fraud predictions, although this adjustment can simultaneously increase false-positive alerts. Threshold optimization must therefore be treated as a controlled trade-off among fraud recall, false-positive burden, precision, operational review capacity, and the cost of missed fraud. Automated threshold-selection procedures developed for imbalanced classification have demonstrated that modifying the decision threshold can improve minority-class performance without requiring the underlying model to be retrained or the training dataset to be resampled (Esposito et al., 2021; Zahidul, 2023). This principle is especially relevant to large-scale global transaction networks because fraud prevalence, transaction value, customer behavior, and risk tolerance can differ rapidly across channels and operational business periods. A threshold that performs adequately for one transaction population may not remain suitable for another population with a different class distribution or cost structure. Dynamic classification-threshold optimization can therefore be understood as an adaptive decision process that continually aligns the classification boundary with the performance objective of the fraud-detection system. In the context of the present study, the principal objective is high sensitivity, meaning that threshold selection should place explicit emphasis on correctly identifying genuine fraudulent transactions while controlling, rather than ignoring, the operational consequences associated with false alerts.

Figure 5: Dynamic Classification-Threshold Optimization and Adaptive Ensemble Learning Framework for High-Sensitivity Fraud Detection




Adaptive ensemble learning addresses another limitation of fraud classification by reducing dependence on the strengths and weaknesses of a single predictive model. Individual classifiers can respond differently to skewed transaction data: one model may classify legitimate observations with high stability, another may be more sensitive to minority fraud patterns, and a third may capture nonlinear or sequential relationships that are not represented effectively by simpler algorithms. Ensemble learning combines such models through voting, averaging, stacking, or boosting procedures so that their complementary predictive information can be used within one architecture. In operational credit-card environments, hybrid ensemble and deep-learning systems have been evaluated through champion-challenger frameworks that compare competing models with realistic measures before and after deployment, demonstrating the importance of selecting models according to operational fraud-detection behavior rather than a single global accuracy score (Kim et al., 2019). Deep sequential ensemble approaches have further treated fraud detection as a sequence-classification problem and combined recurrent neural-network base classifiers through learned voting mechanisms, allowing temporal transaction behavior to contribute to the ensemble decision (Forough & Momtazi, 2021). These approaches are relevant to large-scale digital transaction networks because fraudulent behavior is heterogeneous and may emerge through different combinations of transaction amount, timing, device, merchant, location, account history, and sequential activity. Adaptive ensemble learning extends conventional ensemble construction by allowing classifier selection, weighting, or combination to reflect changes in performance and data characteristics. Adaptation can improve robustness when one model becomes less effective for a particular fraud pattern or transaction segment. For high-sensitivity detection, the value of an ensemble lies not simply in combining many models, but in coordinating classifiers so that minority-class recognition is strengthened without allowing the majority class to dominate the final decision. This makes ensemble diversity, base-model specialization, aggregation strategy, and minority-class performance central considerations when constructing adaptive fraud-detection systems.

A heterogeneous ensemble can combine different families of classifiers, increasing the probability that distinct structures within fraud data are captured. Research on credit-card fraud detection has shown that distribution-aware heterogeneous ensembles can preserve informative characteristics of majority transactions while improving recall for fraudulent cases, while limiting excessive false positives among legitimate transactions (Xie et al., 2021). Ensemble construction can also incorporate cost sensitivity at the meta-learning stage. A cost-sensitive meta-learning ensemble has been proposed in which base classifiers are trained conventionally while the higher-level meta-classifier incorporates asymmetric costs, producing an ensemble that remains sensitive to fraud under changing fraud proportions without forcing identical cost adjustments on every component model (Olowookere & Adewale, 2020). Dynamic threshold optimization and ensemble learning should be viewed as complementary rather than isolated techniques. The ensemble determines how different predictive signals are combined, while the threshold determines the operating point at which the combined risk score becomes an actionable fraud classification. In a high-sensitivity environment, the optimal configuration may therefore involve stronger weights for classifiers that identify difficult minority patterns, a decision threshold calibrated toward fraud recall, and controls that prevent the resulting false-positive volume from becoming operationally unsustainable. The interaction is especially important in large-scale networks, where millions of legitimate transactions can convert even a small increase in false-positive rate into a substantial number of alerts. Accordingly, adaptive fraud detection requires simultaneous attention to minority recall, majority-class preservation, model diversity, score calibration, decision boundaries, and misclassification costs. Within the present research, dynamic classification-threshold optimization and adaptive ensemble learning represent two distinct explanatory constructs whose coordinated operation provides a decision-level and model-level basis for achieving high-sensitivity fraud detection performance under severe class imbalance. This balance is essential because fraud systems must improve minority detection while retaining sufficient discrimination among the far larger legitimate transaction population.

Statistical Learning Theory as the Theoretical Framework

Statistical Learning Theory provides an appropriate theoretical foundation for examining adaptive class-imbalance learning because the theory explains how predictive functions are learned from finite observations, how empirical error relates to expected error, and how model complexity influences generalization to unseen cases. In fraud detection, these principles are especially important because the available learning distribution is highly asymmetric: legitimate transactions dominate the observed population, while fraudulent transactions form a sparse but operationally critical minority. Statistical Learning Theory treats learning as the selection of an appropriate predictive function from a hypothesis space that minimizes expected predictive risk while controlling excessive model complexity, thereby providing a formal basis for interpreting how fraud classifiers construct decision boundaries from observed transactional information. Modern classification theory has emphasized empirical risk minimization, model complexity, concentration behavior, generalization, and model selection as fundamental components determining whether a classifier developed from finite observations can perform reliably when exposed to new cases (Boucheron et al., 2005).

Figure 6: Statistical Learning Theory Framework for Adaptive Class-Imbalance Learning and High-Sensitivity Fraud Detection



The importance of these principles is heightened under class imbalance because an ordinary empirical loss function can become dominated by the much larger legitimate class. A learning algorithm can consequently minimize aggregate classification error while maintaining weak sensitivity toward fraudulent observations. Classification theory has further established that many machine-learning methods, including margin-based and boosting approaches, can be interpreted through optimization of surrogate loss functions and that the relationship between surrogate risk and actual classification risk is fundamental to reliable predictive learning (Bartlett et al., 2006). For the present research, this theoretical learning process can be expressed in a weighted and regularized empirical-risk form. Here, the optimal predictive function is selected from a hypothesis space by minimizing weighted classification loss while regularization controls excessive model complexity. The weighting component

is particularly relevant to fraud detection because it allows rare fraudulent observations or costly fraud-related errors to influence the learning objective more strongly than they would under an unweighted majority-dominated loss structure.

$$f^* = \arg \min_{f \in H} \left[\left(\frac{1}{n} \sum_{i=1}^n w_i L(y_i, f(x_i)) \right) + \lambda \Omega(f) \right]$$

Within this theoretical structure, the four explanatory constructs of the present study can be interpreted as different mechanisms for modifying the effective learning problem faced by a fraud-detection system. Adaptive Minority-Class Resampling changes the empirical training distribution by increasing, decreasing, or selectively modifying the representation of observations so that rare fraudulent transactions contribute sufficient information to classifier construction. Cost-Sensitive Learning modifies the loss function by assigning unequal penalties to different classification errors, meaning that a fraudulent transaction incorrectly classified as legitimate can exert greater influence on learning than an error associated with a legitimate observation. Dynamic Classification-Threshold Optimization modifies the decision rule through which a probability or risk score becomes an operational fraud classification, while Adaptive Ensemble Learning broadens the effective hypothesis space by integrating complementary information generated by multiple predictive models. Statistical Learning Theory therefore offers a unified explanation of how these mechanisms can influence high-sensitivity fraud detection even when the original transaction population remains severely imbalanced. Theoretical investigation of cost-sensitive neural classification has demonstrated that equal-cost learning under highly imbalanced data can generate decision boundaries biased toward the majority class, while the incorporation of asymmetric costs can modify the optimization objective and strengthen discrimination of the minority category (Castro & Braga, 2013). This principle is directly relevant to fraud detection because the operational importance of fraudulent transactions is substantially greater than their numerical frequency within most transaction datasets. A related statistical interpretation is provided by Neyman-Pearson classification, where one type of classification error is minimized subject to a predefined constraint on the other type of error, allowing learning objectives to reflect asymmetric operational priorities rather than assuming that both errors deserve identical treatment (Tong et al., 2016). High-sensitivity fraud detection corresponds strongly with this asymmetric logic because reducing false negatives is a principal requirement while false-positive activity must remain within an acceptable operational range. The theoretical framework therefore does not assume that effective fraud learning requires equal class frequencies, equal error costs, or a universally fixed classification threshold.

The theoretical framework is translated into the empirical design of this research by modeling High-Sensitivity Fraud Detection Performance as an outcome of four measurable adaptive class-imbalance learning capabilities. This translation is important because the present quantitative, cross-sectional, case-study-based study does not directly train a transaction-level classifier; rather, it measures the extent to which Adaptive Minority-Class Resampling, Cost-Sensitive Learning, Dynamic Classification-Threshold Optimization, and Adaptive Ensemble Learning are applied or supported within the selected digital transaction environment and statistically examines their association with fraud-detection performance. Contemporary research on imbalanced classification distinguishes conventional overall-error minimization from cost-sensitive and Neyman-Pearson learning paradigms and demonstrates that resampling techniques, base classifiers, imbalance ratios, and evaluation objectives can interact in determining minority-class performance (Feng et al., 2021). This distinction supports the present study's emphasis on high sensitivity because an aggregate accuracy objective is theoretically different from an objective that explicitly prioritizes the recognition of fraudulent minority observations. The principal empirical formula selected for application throughout the methodology, results, hypothesis testing, and discussion models High-Sensitivity Fraud Detection Performance as a function of Adaptive Minority-Class Resampling, Cost-Sensitive Learning, Dynamic Classification-Threshold Optimization, and Adaptive Ensemble Learning. This is the most appropriate primary formula for the whole empirical study because it directly operationalizes the theoretical proposition that adaptive interventions at the data, loss, decision, and ensemble levels explain variations in fraud-detection sensitivity.

$$HSFDP_i = \beta_0 + \beta_1 AMCR_i + \beta_2 CSL_i + \beta_3 DCTO_i + \beta_4 AEL_i + \varepsilon_i$$


Conceptual Framework

The conceptual framework of the present study establishes High-Sensitivity Fraud Detection Performance as the dependent construct and Adaptive Minority-Class Resampling, Cost-Sensitive Learning, Dynamic Classification-Threshold Optimization, and Adaptive Ensemble Learning as four independent constructs that represent complementary dimensions of adaptive class-imbalance management. The framework assumes that fraud-detection performance within large-scale digital transaction networks cannot be adequately explained by a single balancing mechanism because class imbalance affects multiple stages of the machine-learning process. Adaptive Minority-Class Resampling operates primarily at the data level by modifying the representation of fraudulent and legitimate transactions so that minority observations receive adequate exposure during model learning. This dimension becomes especially important in large-scale environments because the combination of enormous datasets and severe minority scarcity can create both computational and predictive difficulties. Research on imbalanced big-data classification has demonstrated that conventional learning procedures can lose effectiveness when massive datasets contain severely underrepresented classes, requiring scalable strategies that preserve minority information while remaining computationally manageable (Rio et al., 2014). Adaptive sampling research has additionally shown that ensemble-oriented sampling procedures can use information about class distribution to reduce learning bias and adapt the amount or type of sampling according to the underlying imbalance characteristics (Tang & He, 2017). Within this conceptual framework, Adaptive Minority-Class Resampling is therefore expected to increase the informational contribution of fraudulent observations without requiring an artificial assumption that both classes must contain equal numbers of cases. The first proposed relationship is AMCR to HSFDP, where stronger adaptive resampling capability is expected to correspond with stronger High-Sensitivity Fraud Detection Performance. This relationship reflects the principle that minority-class examples must be represented sufficiently for a classifier to learn fraud-related decision regions. At the same time, the framework recognizes that resampling alone cannot determine fraud sensitivity because the consequences of classification errors, the operating threshold used for final decisions, and the manner in which competing classifiers are combined also influence the probability that actual fraudulent transactions will be successfully identified.

The second and third relationships in the conceptual framework concern Cost-Sensitive Learning and Dynamic Classification-Threshold Optimization, both of which address the asymmetry between fraudulent and legitimate classification outcomes but operate at different points in the predictive process. Cost-Sensitive Learning modifies the learning objective by assigning greater importance to particular errors, especially false negatives in which actual fraudulent transactions are classified as legitimate. This treatment is consistent with the economic and security characteristics of fraud detection because the consequences of missing a genuine fraudulent transaction can differ substantially from those associated with incorrectly flagging a legitimate transaction. Cost-sensitive margin-based research has demonstrated that increasing minority-class penalties can alter classification boundaries and progressively improve minority detection rates under imbalanced training distributions (Cheng et al., 2016). The proposed relationship is therefore CSL to HSFDP, indicating that stronger incorporation of asymmetric misclassification costs is expected to correspond with greater fraud-detection sensitivity. Dynamic Classification-Threshold Optimization operates after a model has generated continuous probabilities or risk scores and determines the point at which an observation is converted into a fraud classification. This distinction is important because a model can produce informative fraud probabilities while still achieving weak sensitivity if its decision threshold is poorly aligned with the minority-detection objective. Threshold-moving research combined with bagging ensembles has demonstrated that posterior decision thresholds can be adjusted after model estimation to accommodate imbalanced classification objectives while preserving the natural class distribution during training (Collell et al., 2018). The third conceptual relationship is therefore DCTO to HSFDP. Within the present study, dynamic optimization does not mean that the threshold should simply be lowered as far as possible. Rather, it refers to systematic adjustment of the operating point so that stronger fraud sensitivity is achieved while the volume of false-positive alerts remains manageable.

Cost-Sensitive Learning and Dynamic Classification-Threshold Optimization are therefore conceptually distinct: the first changes how classification errors influence model learning, while the second changes how model outputs are transformed into final transaction decisions.

Figure 7: Conceptual Framework for Adaptive Class-Imbalance Learning and High-Sensitivity Fraud Detection




Adaptive Ensemble Learning constitutes the fourth explanatory construct and completes the conceptual framework by representing model-level integration across multiple classifiers. Fraudulent transactions are heterogeneous, and different machine-learning algorithms can recognize different portions of the fraud distribution according to their assumptions, feature representations, decision structures, and sensitivity to minority observations. Adaptive ensembles seek to exploit this diversity by combining multiple base learners and modifying their weights, membership, or influence according to predictive competence. The importance of strong machine-learning architectures in fraud classification has been demonstrated through optimized gradient-boosting approaches developed specifically for credit-card fraud detection, where model optimization can improve discrimination within highly skewed financial datasets (Taha & Malebary, 2020). The proposed fourth relationship is AEL to HSFDP, indicating that greater adaptive ensemble capability is expected to correspond with stronger fraud-detection sensitivity. The four relationships are integrated into one empirical conceptual model in which High-Sensitivity Fraud Detection Performance is predicted by Adaptive Minority-Class Resampling, Cost-Sensitive Learning, Dynamic Classification-Threshold Optimization, and Adaptive Ensemble Learning. In that model, the intercept represents baseline performance, the four regression coefficients represent the individual effects associated with the explanatory constructs, and the residual represents unexplained variation. This formula will serve as the principal empirical model throughout the study because it corresponds directly with H1 through H4, while the overall model significance will provide the statistical basis for H5 concerning the collective effect of adaptive class-imbalance learning. The conceptual framework therefore links data-level adaptation, learning-cost adaptation, decision-level adaptation, and classifier-level adaptation with one measurable outcome, enabling descriptive statistics to establish construct levels, Pearson correlations to examine bivariate associations, and multiple regression modeling to determine the relative and combined contributions of the four predictors to High-Sensitivity Fraud Detection Performance in large-scale digital transaction networks.

$$HSFDP_i = \beta_0 + \beta_1 AMCR_i + \beta_2 CSL_i + \beta_3 DCTO_i + \beta_4 AEL_i + \varepsilon_i$$

METHOD

This research has adopted a quantitative, cross-sectional, case-study-based design to examine the relationship between Adaptive Class-Imbalance Learning and High-Sensitivity Fraud Detection Performance in large-scale digital transaction networks. The research design has been structured to generate measurable evidence regarding Adaptive Minority-Class Resampling, Cost-Sensitive Learning, Dynamic Classification-Threshold Optimization, Adaptive Ensemble Learning, and High-Sensitivity Fraud Detection Performance.

Figure 8: Research Methodology




The case-study context has focused on organizations and professional environments associated with large-scale digital transactions, including banking, fintech, electronic payment processing, digital wallets, e-commerce payment systems, and related transaction-security operations. The target population has consisted of professionals directly involved in fraud analytics, machine learning, transaction monitoring, cybersecurity, data science, payment risk, financial technology, and digital transaction management, while the individual professional respondent has served as the primary unit of analysis. A purposive sampling strategy has been used to ensure that participants have possessed relevant professional knowledge and practical exposure to fraud-detection systems and transaction-risk management. Data have been collected through a structured questionnaire developed specifically for this study. The instrument has used a five-point Likert scale ranging from 1 = Strongly Disagree to 5 = Strongly Agree, and measurement items have been organized according to the five principal research constructs. A pilot test has been conducted before the main data collection to assess questionnaire clarity, wording, structure, and internal consistency, and necessary refinements have been made based on pilot responses. Content and construct validity have been evaluated to ensure that the questionnaire items have adequately represented the intended variables. Reliability has been assessed through Cronbach's alpha, while KMO and Bartlett's Test of Sphericity have been applied where appropriate to evaluate sampling adequacy and construct suitability. The collected data have been coded, cleaned, and initially organized using Microsoft Excel, while IBM SPSS Statistics has been used to compute descriptive statistics, reliability indices, Pearson correlation analysis, multiple regression modeling, ANOVA, and regression diagnostics including tolerance, VIF, and the Durbin-Watson statistic. EndNote has been used to organize scholarly sources and manage APA 7th edition citations and references, and Microsoft Word has been used for manuscript preparation, table presentation, equation placement, and final journal-style formatting.

FINDINGS/RESULTS

The findings of the study have provided quantitative evidence supporting the proposed relationships between Adaptive Class-Imbalance Learning and High-Sensitivity Fraud Detection Performance in large-scale digital transaction networks. A total of 320 questionnaires have been distributed among

professionals working in fraud analytics, financial technology, transaction monitoring, cybersecurity, machine learning, payment processing, and digital risk management, of which 296 questionnaires have been returned and 288 have been considered complete and usable, producing a usable response rate of 90.0%. The descriptive results based on the five-point Likert scale have shown consistently high levels of agreement across all research constructs.

Figure 9: Findings of The Study



Dynamic Classification-Threshold Optimization has recorded the highest mean among the independent variables, $M = 4.17$, $SD = 0.55$, followed by Cost-Sensitive Learning, $M = 4.13$, $SD = 0.58$, Adaptive Minority-Class Resampling, $M = 4.08$, $SD = 0.61$, and Adaptive Ensemble Learning, $M = 4.05$, $SD = 0.63$. High-Sensitivity Fraud Detection Performance has recorded an overall mean of $M = 4.19$, $SD = 0.54$, indicating that respondents have generally agreed that adaptive imbalance-management capabilities have contributed to stronger fraud identification and reduced false-negative outcomes. Reliability analysis has demonstrated satisfactory internal consistency, with Cronbach's alpha coefficients ranging from .84 to .91 across the individual constructs and an overall reliability coefficient of $\alpha = .93$. The Kaiser-Meyer-Olkin value has been .91, while Bartlett's Test of Sphericity has been statistically significant, $p < .001$, confirming the adequacy of the measurement structure. Pearson correlation analysis has revealed significant positive relationships between High-Sensitivity Fraud Detection Performance and Adaptive Minority-Class Resampling, $r = .62$, $p < .001$; Cost-Sensitive Learning, $r = .68$, $p < .001$; Dynamic Classification-Threshold Optimization, $r = .74$, $p < .001$; and Adaptive Ensemble Learning, $r = .66$, $p < .001$. The multiple regression model has further demonstrated a strong collective predictive effect, $R = .842$, $R^2 = .709$, adjusted $R^2 = .705$, $F(4, 283) = 172.45$, $p < .001$, indicating that approximately 70.9% of the variance in High-Sensitivity Fraud Detection Performance has been explained jointly by the four adaptive learning constructs. Dynamic Classification-Threshold Optimization has emerged as the strongest predictor, $\beta = .32$, $p < .001$, followed by Cost-Sensitive Learning, $\beta = .27$, $p < .001$, Adaptive Ensemble Learning, $\beta = .23$, $p < .001$, and Adaptive Minority-Class Resampling, $\beta = .19$, $p = .002$. Multicollinearity has not presented a concern, with VIF values ranging from 1.42 to 1.88, while the Durbin-Watson statistic of 1.96 has indicated acceptable residual independence. These results have supported H1, H2, H3, H4, and H5 and have simultaneously fulfilled the research objectives by demonstrating that each adaptive class-imbalance learning capability has had a significant positive relationship with fraud-detection sensitivity and that their combined effect has substantially explained variation in high-sensitivity fraud detection performance.

Questionnaire Distribution and Response Rate

Table 1 has presented the questionnaire distribution and response pattern obtained for the quantitative investigation of Adaptive Class-Imbalance Learning and High-Sensitivity Fraud Detection Performance. A total of 320 questionnaires have been distributed among professionals whose

responsibilities have involved fraud analytics, transaction monitoring, machine learning, cybersecurity, digital payments, financial technology, data analytics, or related digital transaction activities. Of these questionnaires, 296 have been returned, representing a gross response rate of 92.5%. Following data screening, eight questionnaires have been identified as incomplete or unsuitable for statistical analysis because important measurement items have not been sufficiently completed. Consequently, 288 questionnaires have remained valid and usable, producing a final usable response rate of 90.0%. This response level has provided a sufficiently substantial quantitative basis for calculating descriptive statistics, reliability coefficients, correlations, and multiple regression estimates across the five study constructs. The final sample has also allowed the study to examine variation in professional perceptions of Adaptive Minority-Class Resampling, Cost-Sensitive Learning, Dynamic Classification-Threshold Optimization, Adaptive Ensemble Learning, and High-Sensitivity Fraud Detection Performance. From the perspective of Statistical Learning Theory, adequate observational coverage has been important because empirical relationships have depended on information derived from the observed sample. The usable sample has supplied sufficient variation across respondents for evaluating whether the four adaptive learning capabilities have been systematically associated with the dependent construct. The high response rate has also reduced the likelihood that the statistical findings have been excessively influenced by nonresponse within the selected case-study environment. The resulting $N = 288$ has therefore been retained consistently across the major construct analyses and regression model. This sample has formed the empirical foundation for testing H1 through H5 and for examining the research objectives concerning individual and collective contributions of adaptive class-imbalance learning mechanisms to high-sensitivity fraud detection.

Table 1: Questionnaire Distribution and Response Rate

Questionnaire Status	Frequency	Percentage
Questionnaires distributed	320	100.0%
Questionnaires returned	296	92.5%
Incomplete/unusable questionnaires	8	2.5%
Valid and usable questionnaires	288	90.0%
Non-returned questionnaires	24	7.5%
Final sample used for analysis	288	90.0%

Respondent and Digital Transaction Network Characteristics

Table 2 has demonstrated that the final sample has represented a diverse group of professionals and digital transaction environments relevant to the objectives of the study. Fraud and transaction-risk analysts have constituted the largest professional group, accounting for 25.0% of respondents, while data scientists and machine-learning specialists have represented 20.1%. Cybersecurity specialists have accounted for 17.0%, payment and fintech professionals have represented 15.6%, IT and data managers have accounted for 12.5%, and compliance or financial-risk professionals have represented 9.7%. This distribution has ensured that the study has incorporated perspectives from both technical model-development functions and operational fraud-management functions. In terms of professional experience, 32.6% of respondents have reported three to five years of relevant experience, while 27.1% have reported six to ten years. A further 18.8% have possessed more than ten years of experience, indicating that a substantial proportion of the sample has had considerable exposure to digital transaction risk and fraud-detection activities. Regarding organizational environments, banking and card-based transaction systems have accounted for 29.9% of respondents, followed by fintech platforms at 20.1%, payment-processing networks at 18.8%, e-commerce payment systems at 17.0%, and digital-wallet or mobile-payment environments at 14.2%. These distributions have supported the case-study objective by representing several forms of large-scale digital transaction activity rather than a single narrow payment channel. Statistical Learning Theory has emphasized that learning behavior depends

partly on the characteristics of observed data environments, and the diversity represented in Table 2 has therefore strengthened the empirical context for evaluating adaptive learning capabilities. The sample composition has provided relevant professional knowledge concerning class imbalance, asymmetric fraud costs, transaction-risk thresholds, ensemble modeling, and high-sensitivity detection, thereby supporting the appropriateness of the respondents as the unit of analysis.

Table 2: Respondent and Digital Transaction Network Characteristics

Characteristic	Category	Frequency	Percentage
Professional role	Fraud/transaction risk analyst	72	25.0%
	Data scientist/ML specialist	58	20.1%
	Cybersecurity specialist	49	17.0%
	Payment/fintech professional	45	15.6%
	IT/data manager	36	12.5%
	Compliance/financial risk professional	28	9.7%
Professional experience	Less than 3 years	62	21.5%
	3–5 years	94	32.6%
	6–10 years	78	27.1%
	More than 10 years	54	18.8%
Transaction environment	Banking/card transactions	86	29.9%
	Fintech platforms	58	20.1%
	Payment processing networks	54	18.8%
	E-commerce payment systems	49	17.0%
	Digital-wallet/mobile-payment systems	41	14.2%
Total		288	100.0%

Descriptive Statistics of Adaptive Class-Imbalance Learning Constructs

Table 3: Descriptive Statistics of Adaptive Class-Imbalance Learning Constructs

Construct	N	Mean	SD	Likert Interpretation	Rank
Adaptive Minority-Class Resampling (AMCR)	288	4.08	0.61	Agree	3
Cost-Sensitive Learning (CSL)	288	4.13	0.58	Agree	2
Dynamic Classification-Threshold Optimization (DCTO)	288	4.17	0.55	Agree	1
Adaptive Ensemble Learning (AEL)	288	4.05	0.63	Agree	4
Overall Adaptive Class-Imbalance Learning	288	4.11	0.59	Agree	—

Table 3 has summarized respondents' assessments of the four independent constructs measured through the five-point Likert scale. Dynamic Classification-Threshold Optimization has recorded the highest mean score, $M = 4.17$, $SD = 0.55$, indicating a strong level of agreement that flexible fraud-decision thresholds have supported the identification of fraudulent transactions within highly

imbalanced transaction environments. Cost-Sensitive Learning has produced the second-highest mean, $M = 4.13$, $SD = 0.58$, demonstrating that respondents have recognized the importance of assigning greater analytical weight to costly fraud-related errors, particularly false negatives. Adaptive Minority-Class Resampling has achieved $M = 4.08$, $SD = 0.61$, showing agreement that modifying the representation of minority fraud observations has contributed meaningfully to imbalance management. Adaptive Ensemble Learning has recorded $M = 4.05$, $SD = 0.63$, which has also fallen within the Agree category and has indicated that combining or adapting multiple classifiers has been regarded as a relevant fraud-detection capability.

The overall composite mean of 4.11 has shown that adaptive class-imbalance learning has been present at a comparatively high perceived level across the investigated environments. The relatively moderate standard deviations, ranging from 0.55 to 0.63, have further indicated that respondents' evaluations have not been excessively dispersed. These descriptive findings have addressed the study's objective of assessing the observed level of each adaptive learning capability before testing statistical relationships with fraud-detection performance. Statistical Learning Theory has provided a direct explanation for these patterns because AMCR has altered the effective empirical distribution, CSL has modified the loss associated with classification errors, DCTO has adjusted the decision boundary, and AEL has expanded the effective predictive structure through classifier combination. The ranking has also suggested that decision-level adaptation has received the strongest assessment among the four mechanisms, which has been consistent with the later regression finding that DCTO has become the strongest individual predictor of High-Sensitivity Fraud Detection Performance.

Descriptive Statistics of High-Sensitivity Fraud Detection Performance

Table 4: Descriptive Statistics of High-Sensitivity Fraud Detection Performance

High-Sensitivity Fraud Detection Indicator	N	Mean	SD	Likert Interpretation
Improved identification of genuine fraudulent transactions	288	4.23	0.52	Strongly Agree
Reduction of fraudulent transactions classified as legitimate	288	4.21	0.53	Strongly Agree
Improved minority-class recall/sensitivity	288	4.18	0.55	Agree
Stronger recognition of difficult fraud patterns	288	4.16	0.58	Agree
Effective fraud sensitivity with manageable false alerts	288	4.17	0.57	Agree
Overall HSFPD	288	4.19	0.54	Agree

Table 4 has presented the descriptive results for High-Sensitivity Fraud Detection Performance, which has served as the dependent variable of the research. The indicator concerning improved identification of genuine fraudulent transactions has recorded the highest individual mean, $M = 4.23$, $SD = 0.52$, placing it within the Strongly Agree range. The reduction of fraudulent transactions incorrectly classified as legitimate has followed closely with $M = 4.21$, $SD = 0.53$, indicating particularly strong respondent agreement regarding false-negative reduction. Improved minority-class recall or sensitivity has obtained $M = 4.18$, $SD = 0.55$, while stronger recognition of difficult fraud patterns has recorded $M = 4.16$, $SD = 0.58$.

The indicator representing effective fraud sensitivity combined with manageable false-alert levels has produced $M = 4.17$, $SD = 0.57$. Collectively, these indicators have generated an overall High-Sensitivity Fraud Detection Performance mean of 4.19 with $SD = 0.54$. This result has indicated that respondents have generally perceived strong fraud-detection performance within environments characterized by adaptive class-imbalance learning capabilities. More importantly, the highest-rated indicators have related directly to genuine fraud recognition and false-negative reduction, which have represented the central meaning of sensitivity within the present study. The findings have therefore addressed the

overall research objective of evaluating the outcome variable before examining its statistical dependence on AMCR, CSL, DCTO, and AEL. Statistical Learning Theory has supported this interpretation because classification quality under asymmetric risk has not been adequately represented by overall accuracy alone.

Instead, the learned decision function has needed to provide meaningful discrimination for the minority class, particularly when fraudulent cases have carried greater operational importance. The high HSFD mean has therefore provided descriptive evidence that the investigated adaptive learning environment has been associated with strong minority-focused detection capability. This result has subsequently been tested more rigorously through correlation and multiple regression rather than being treated as sufficient evidence on its own.

Reliability and Construct Validity Analysis

Table 5: Reliability and Construct Validity Analysis

Construct/Test	Items	Cronbach's α / Statistic	Threshold	Result
Adaptive Minority-Class Resampling	4	.86	$\geq .70$	Reliable
Cost-Sensitive Learning	4	.88	$\geq .70$	Reliable
Dynamic Classification-Threshold Optimization	4	.91	$\geq .70$	Highly reliable
Adaptive Ensemble Learning	4	.84	$\geq .70$	Reliable
High-Sensitivity Fraud Detection Performance	4	.90	$\geq .70$	Highly reliable
Overall instrument	20	.93	$\geq .70$	Highly reliable
Kaiser-Meyer-Olkin measure	—	.91	$\geq .60$	Excellent adequacy
Bartlett's Test of Sphericity	—	$\chi^2 = 3,126.84, df = 190, p < .001$	$p < .05$	Significant

Table 5 has demonstrated that the research instrument has achieved satisfactory to excellent reliability and construct adequacy before the hypothesized relationships have been statistically evaluated. Adaptive Minority-Class Resampling has produced a Cronbach's alpha coefficient of .86, Cost-Sensitive Learning has recorded $\alpha = .88$, Dynamic Classification-Threshold Optimization has achieved $\alpha = .91$, Adaptive Ensemble Learning has produced $\alpha = .84$, and High-Sensitivity Fraud Detection Performance has recorded $\alpha = .90$. All coefficients have exceeded the conventional .70 reliability criterion, while the overall 20-item instrument has achieved $\alpha = .93$. These results have indicated that the measurement items within each construct have demonstrated strong internal consistency and have measured their intended dimensions in a sufficiently coherent manner.

Dynamic Classification-Threshold Optimization has recorded the highest construct-specific reliability, which has provided additional measurement support for the variable that has subsequently emerged as the strongest regression predictor. Construct suitability has also been assessed using the Kaiser-Meyer-Olkin measure and Bartlett's Test of Sphericity. The KMO value of .91 has indicated excellent sampling adequacy, while Bartlett's test has been statistically significant, $\chi^2 = 3,126.84, df = 190, p < .001$, demonstrating that the inter-item correlation structure has been sufficiently strong for construct-oriented analysis. These findings have strengthened confidence that subsequent correlation and regression results have not been based on highly inconsistent measurement scales. From the perspective of Statistical Learning Theory, reliable measurement has been particularly important because the empirical model has attempted to estimate systematic relationships among theoretically defined learning capabilities. Excessive measurement error would have weakened the observed relationships and obscured the theoretical effects associated with data-level, cost-level, decision-level, and ensemble-level adaptation. The reliability and validity evidence has therefore supported the

measurement stage of the study and has allowed H1 through H5 to be evaluated using composite construct scores with an acceptable level of statistical consistency.

Correlation Analysis

Table 6: Pearson Correlation Matrix of Study Constructs

Variable	AMCR	CSL	DCTO	AEL	HSFDP
AMCR	1.00				
CSL	.49***	1.00			
DCTO	.52***	.58***	1.00		
AEL	.47***	.54***	.61***	1.00	
HSFDP	.62***	.68***	.74***	.66***	1.00

Note. $N = 288$. *** $p < .001$. AMCR = Adaptive Minority-Class Resampling; CSL = Cost-Sensitive Learning; DCTO = Dynamic Classification-Threshold Optimization; AEL = Adaptive Ensemble Learning; HSFDP = High-Sensitivity Fraud Detection Performance.

Table 6 has shown that all four adaptive class-imbalance learning constructs have been positively and statistically significantly correlated with High-Sensitivity Fraud Detection Performance. Adaptive Minority-Class Resampling has demonstrated a positive correlation with HSFDP, $r = .62$, $p < .001$, indicating that stronger adaptive representation of minority fraud observations has been associated with stronger fraud-detection sensitivity. Cost-Sensitive Learning has produced a larger positive relationship, $r = .68$, $p < .001$, suggesting that greater consideration of asymmetric fraud-related misclassification costs has been associated with higher detection performance. Dynamic Classification-Threshold Optimization has recorded the strongest bivariate relationship with HSFDP, $r = .74$, $p < .001$. This result has indicated that adaptive management of the fraud-decision threshold has been especially strongly associated with the ability to identify genuine fraudulent transactions and reduce false negatives. Adaptive Ensemble Learning has also demonstrated a substantial positive correlation, $r = .66$, $p < .001$. Therefore, the preliminary relational evidence has been consistent with H1, H2, H3, and H4. The intercorrelations among the independent variables have ranged from .47 to .61, showing that the constructs have been related but have not represented identical dimensions of adaptive learning. This pattern has supported their treatment as complementary predictors within the subsequent regression model. Statistical Learning Theory has provided a coherent explanation for the correlation structure because the four predictors have represented modifications to different components of the classification process. AMCR has influenced the empirical sample distribution, CSL has influenced risk weighting, DCTO has influenced the final classification boundary, and AEL has influenced the composition of the predictive function. Their positive associations with HSFDP have therefore been theoretically consistent with an adaptive learning system that has reduced majority-class dominance and strengthened minority-event recognition. Correlation alone has not established the unique effect of each predictor, so multiple regression analysis has subsequently been used to determine whether each construct has retained significant explanatory power after the other adaptive capabilities have been statistically controlled.

Multiple Regression Analysis of High-Sensitivity Fraud Detection Performance

Table 7 has provided the principal inferential evidence for determining the individual and collective effects of adaptive class-imbalance learning on High-Sensitivity Fraud Detection Performance. The overall model has produced $R = .842$, demonstrating a strong relationship between the combined predictor set and the dependent variable. The coefficient of determination has been $R^2 = .709$, while adjusted R^2 has been .705, indicating that approximately 70.9% of the observed variance in HSFDP has been explained collectively by Adaptive Minority-Class Resampling, Cost-Sensitive Learning, Dynamic Classification-Threshold Optimization, and Adaptive Ensemble Learning. The ANOVA result has confirmed that the overall regression model has been statistically significant, $F(4, 283) = 172.45$, $p < .001$. Therefore, H5 has received strong statistical support. At the individual predictor level,

Dynamic Classification-Threshold Optimization has produced the largest standardized effect, $\beta = .32$, $p < .001$, and has therefore emerged as the strongest predictor of HSFDP. Cost-Sensitive Learning has followed with $\beta = .27$, $p < .001$, Adaptive Ensemble Learning has produced $\beta = .23$, $p < .001$, and Adaptive Minority-Class Resampling has produced $\beta = .19$, $p = .002$. All four coefficients have remained positive and statistically significant after the effects of the other predictors have been controlled, providing regression-based support for H1 through H4. The VIF values have ranged from 1.42 to 1.88, while tolerance values have ranged from .53 to .70, indicating that multicollinearity has not threatened interpretation of the individual coefficients. The Durbin-Watson statistic of 1.96 has further indicated acceptable residual independence. These findings have strongly reflected Statistical Learning Theory because adaptations to the empirical distribution, weighted loss, decision boundary, and hypothesis combination have each explained unique variation in fraud sensitivity. The empirical model has therefore received quantitative support from the observed results.

Table 7: Multiple Regression Analysis Predicting High-Sensitivity Fraud Detection Performance
Panel A. Model Summary

R	R ²	Adjusted R ²	Std. Error of Estimate	Durbin-Watson
.842	.709	.705	.291	1.96

Panel B. ANOVA

Source	Sum of Squares	df	Mean Square	F	p
Regression	59.34	4	14.84	172.45	< .001
Residual	24.35	283	0.086		
Total	83.69	287			

Panel C. Regression Coefficients

Predictor	B	SE	Standardized β	t	p	Tolerance	VIF
Constant	0.62	0.16	—	3.88	< .001	—	—
AMCR	0.17	0.05	.19	3.40	.002	.70	1.42
CSL	0.25	0.05	.27	5.00	< .001	.62	1.61
DCTO	0.31	0.05	.32	6.20	< .001	.53	1.88
AEL	0.20	0.04	.23	5.00	< .001	.58	1.72

$$\text{HSFDP}_i = \beta_0 + \beta_1 \text{AMCR}_i + \beta_2 \text{CSL}_i + \beta_3 \text{DCTO}_i + \beta_4 \text{AEL}_i + \varepsilon_i$$

Hypothesis Testing and Research Objective Validation

Table 8 has consolidated the descriptive, correlational, and regression evidence and has demonstrated that all five hypotheses have been supported. H1 has proposed that Adaptive Minority-Class Resampling has a significant positive effect on High-Sensitivity Fraud Detection Performance. This proposition has been supported by the positive correlation of $r = .62$, $p < .001$ and the significant regression coefficient of $\beta = .19$, $p = .002$. Accordingly, the first research objective concerning the contribution of adaptive minority representation has been achieved. H2 has been supported through $r = .68$, $p < .001$ and $\beta = .27$, $p < .001$, confirming that Cost-Sensitive Learning has contributed significantly to HSFDP and validating the second objective. H3 has received the strongest individual empirical support, with DCTO producing $r = .74$, $p < .001$ and $\beta = .32$, $p < .001$. The third objective has therefore been achieved, and threshold optimization has emerged as the strongest individual predictor. H4 has also been supported because Adaptive Ensemble Learning has demonstrated $r = .66$, $p < .001$ and $\beta = .23$, $p < .001$, validating the fourth objective concerning classifier integration. Finally, H5 has proposed a collective positive effect of the four adaptive class-imbalance learning constructs. The regression model has produced $R^2 = .709$ and $F(4, 283) = 172.45$, $p < .001$, demonstrating that the four predictors

have jointly explained 70.9% of the variance in HSFDP. The fifth research objective has therefore been achieved. These findings have also aligned closely with Statistical Learning Theory. Data-level adaptation through AMCR has strengthened minority representation, error-level adaptation through CSL has altered the importance of classification losses, decision-level adaptation through DCTO has optimized the operating boundary, and model-level adaptation through AEL has combined complementary predictive capabilities. The complete pattern of supported hypotheses has therefore demonstrated that high-sensitivity fraud detection has been associated not with one isolated imbalance correction method but with a multidimensional adaptive learning structure. Among these dimensions, DCTO has exerted the greatest standardized effect, while all four constructs have retained statistically significant independent contributions to the dependent variable.

Table 8: Summary of Hypothesis Testing and Research Objective Validation

Hypothesis	Proposed Relationship	Key Statistical Evidence	Decision	Objective Validation
H1	AMCR → HSFDP	$r = .62^{***}$; $\beta = .19$, $p = .002$	Supported	Objective 1 achieved
H2	CSL → HSFDP	$r = .68^{***}$; $\beta = .27$, $p < .001$	Supported	Objective 2 achieved
H3	DCTO → HSFDP	$r = .74^{***}$; $\beta = .32$, $p < .001$	Supported	Objective 3 achieved
H4	AEL → HSFDP	$r = .66^{***}$; $\beta = .23$, $p < .001$	Supported	Objective 4 achieved
H5	AMCR + CSL + DCTO + AEL → HSFDP	$R = .842$; $R^2 = .709$; Adj. $R^2 = .705$; $F(4,283) = 172.45$, $p < .001$	Supported	Objective 5 achieved

Note. $^{***}p < .001$ for Pearson correlations.

DISCUSSION

The findings have demonstrated that adaptive class-imbalance learning has provided a strong empirical explanation of High-Sensitivity Fraud Detection Performance within large-scale digital transaction networks. All four explanatory constructs have received high Likert-scale evaluations, with Dynamic Classification-Threshold Optimization recording the highest mean, $M = 4.17$, $SD = 0.55$, followed by Cost-Sensitive Learning, $M = 4.13$, $SD = 0.58$, Adaptive Minority-Class Resampling, $M = 4.08$, $SD = 0.61$, and Adaptive Ensemble Learning, $M = 4.05$, $SD = 0.63$. High-Sensitivity Fraud Detection Performance has recorded an overall mean of $M = 4.19$, $SD = 0.54$, suggesting that the investigated transaction environments have exhibited comparatively strong capability for identifying fraudulent minority events. The correlation results have reinforced this pattern because each adaptive learning construct has been significantly and positively associated with HSFDP. DCTO has shown the strongest relationship, $r = .74$, followed by CSL, $r = .68$, AEL, $r = .66$, and AMCR, $r = .62$, with all relationships statistically significant at $p < .001$. These findings correspond with the broader imbalance-learning literature, which has established that high aggregate accuracy can conceal weak minority-class recognition and that evaluation must therefore prioritize metrics that directly capture positive-class identification when the event of interest is rare (He & Garcia, 2009; Saito & Rehmsmeier, 2015). The results also support research indicating that fraud-detection systems must simultaneously manage severe imbalance, changing transactional behavior, feedback limitations, and minority-class recognition rather than treating imbalance as a simple numerical property of the dataset (Pozzolo et al., 2018). The high mean for HSFDP has therefore been meaningful because it has reflected respondents' assessments of genuine fraud identification, false-negative reduction, minority-class recall, difficult fraud-pattern recognition, and manageable false-alert performance rather than undifferentiated predictive accuracy. The observed relationships have further suggested that high fraud sensitivity has

not depended on a single technical intervention. Instead, it has been associated with a coordinated set of adaptive mechanisms operating at different stages of the classification process. This interpretation is consistent with ensemble and imbalance-learning research showing that sampling, learning objectives, model combination, and performance criteria interact in determining minority-class detection quality (Galar et al., 2012; Guo et al., 2017). The empirical evidence has therefore validated the study's central proposition that adaptive class-imbalance learning should be understood as a multidimensional capability rather than as one preprocessing technique applied before conventional classification.

The strongest finding has concerned Dynamic Classification-Threshold Optimization, which has produced both the highest independent-variable mean and the largest correlation and regression effect. DCTO has recorded $M = 4.17$, $r = .74$, and $\beta = .32$, $p < .001$, establishing it as the strongest individual predictor of High-Sensitivity Fraud Detection Performance after the effects of the other adaptive learning constructs have been controlled. This result has provided particularly strong support for H3 and has highlighted the importance of the final decision boundary through which fraud probabilities or risk scores are converted into operational classifications. The finding aligns with prior research demonstrating that the default classification threshold is often unsuitable for highly imbalanced problems because it can favor the majority class even when the underlying classifier contains useful information about minority cases. Threshold-moving methods have shown that altering the operating point can improve minority-class recognition without requiring complete retraining of the predictive model (Collell et al., 2018). More recent threshold-optimization research has similarly shown that systematic adjustment of decision thresholds can improve classification behavior in imbalanced settings by shifting attention toward the class of greater analytical importance (Esposito et al., 2021). The present study has extended this logic into the context of large-scale digital transaction networks by showing that professionals have associated dynamic threshold management with improved fraud sensitivity and reduced false-negative outcomes. Cost-Sensitive Learning has emerged as the second-strongest predictor, $\beta = .27$, $p < .001$, and has shown a substantial correlation with HSFD, $r = .68$. This finding has supported H2 and corresponds closely with fraud-specific research showing that errors should not necessarily receive equal penalties because the financial consequences of missing a fraudulent transaction can differ considerably from those of temporarily flagging a legitimate transaction (Bahnsen et al., 2015; Bahnsen et al., 2013). The joint importance of DCTO and CSL has indicated that high-sensitivity fraud detection depends strongly on two forms of asymmetry: asymmetric learning costs and asymmetric decision priorities. Cost-sensitive learning has influenced how strongly the model responds to particular classification errors during training, whereas threshold optimization has influenced how estimated fraud risk is translated into final operational action. Their independent statistical significance has shown that these mechanisms have not been interchangeable. Instead, the results have suggested that large-scale fraud-detection environments can benefit when error costs and decision thresholds are managed as separate but coordinated components of the adaptive classification process.


Adaptive Minority-Class Resampling and Adaptive Ensemble Learning have also demonstrated significant positive effects and have provided additional evidence that high-sensitivity fraud detection requires both data-level and model-level adaptation. AMCR has produced $M = 4.08$, $r = .62$, and $\beta = .19$, $p = .002$, supporting H1 and indicating that stronger representation of rare fraudulent observations has been significantly associated with improved HSFD. Although its standardized regression effect has been smaller than those of DCTO, CSL, and AEL, the positive and statistically significant coefficient has confirmed that data-level imbalance management remains an important foundation for learning from rare fraud examples. This result corresponds with research showing that adaptive oversampling can improve representation of difficult minority examples by concentrating synthetic generation where learning is more challenging rather than applying uniform duplication across the minority class (Barua et al., 2014; Chen et al., 2010). It is also consistent with evidence that resampling effectiveness depends on both imbalance severity and the characteristics of the underlying classifier, meaning that no universal sampling ratio can be assumed to be optimal across all datasets (García et al., 2012). Adaptive Ensemble Learning has produced a stronger standardized effect, $\beta = .23$, $p < .001$, along with $r = .66$ and $M = 4.05$, providing support for H4. This finding agrees with prior fraud research demonstrating

that ensemble architectures can integrate complementary model strengths and improve recognition of complex or heterogeneous fraudulent behavior (Forough & Momtazi, 2021; Kim et al., 2019). Heterogeneous ensemble research has further shown that preserving majority-class information while strengthening minority recall can be particularly useful in credit-card fraud classification, where legitimate observations greatly outnumber fraudulent ones (Xie et al., 2021). The present results have therefore suggested that ensemble learning contributes not merely by increasing model complexity but by allowing different classifiers to capture different portions of the fraud distribution. From a practical perspective, this combination has been especially relevant in large transaction networks where fraudulent behavior may vary across merchants, devices, payment channels, customer segments, locations, and transaction sequences. The smaller coefficient for AMCR relative to DCTO has also been informative. It has suggested that changing the training distribution alone may be insufficient to achieve the strongest gains in fraud sensitivity unless the resulting model scores are subsequently managed through cost-aware learning, optimized thresholds, and adaptive classifier combinations. The empirical pattern has therefore supported a layered fraud-detection architecture in which sampling improves learnability, ensembles strengthen representation, cost sensitivity modifies error priorities, and threshold optimization translates learned risk into operational fraud decisions.

The findings have also produced important practical and theoretical implications while requiring careful interpretation of the study's methodological limitations. Practically, the regression model has achieved $R = .842$, $R^2 = .709$, adjusted $R^2 = .705$, and $F(4, 283) = 172.45$, $p < .001$, meaning that the four adaptive class-imbalance learning constructs have jointly explained 70.9% of the variance in HSFDP. This substantial explanatory power has indicated that banks, fintech companies, payment processors, digital-wallet providers, e-commerce platforms, and transaction-security teams should avoid evaluating fraud models solely according to global accuracy or a single balancing technique. Instead, fraud-detection governance can be organized around coordinated monitoring of minority representation, misclassification costs, threshold performance, and ensemble behavior. Real-time fraud research has shown that scalable detection systems must combine machine-learning capabilities with transaction-stream processing and mechanisms capable of responding to nonstationarity and feedback limitations (Carcillo et al., 2018). The present findings have supported this operational perspective by showing that adaptive mechanisms at multiple stages have remained significant even when estimated simultaneously. Theoretically, the results have supported Statistical Learning Theory by demonstrating that modifications to empirical data representation, loss weighting, decision boundaries, and the effective hypothesis structure have each been associated with variations in classification performance. The significant coefficients have been consistent with the theoretical proposition that minimizing unweighted empirical error is insufficient when the minority class is rare and carries greater operational importance. Cost-sensitive classification theory and asymmetric-error learning similarly recognize that the optimal classifier depends on the objective being minimized rather than on raw accuracy alone (Castro & Braga, 2013; Feng et al., 2021). Several limitations must nevertheless be revisited. The cross-sectional design has captured professional assessments at one period rather than observing changes in model performance longitudinally. The use of Likert-scale responses has measured organizational and professional perceptions of adaptive capabilities rather than directly estimating transaction-level sensitivity, precision, AUC-PR, or monetary savings from raw fraud datasets. The purposive case-study sample has been appropriate for accessing knowledgeable respondents but has limited unrestricted generalization to all digital transaction systems. Self-reported assessments may also contain common-method or perceptual bias, and the high R^2 should therefore be interpreted as explanatory strength within the specified survey model rather than as direct proof that the same proportion of variance would occur in a live transaction-level classifier. These limitations have not invalidated the supported hypotheses, but they have defined the boundary within which the findings should be interpreted and have created a clear basis for more technically intensive validation. Future research should build on these findings by developing and empirically testing an Adaptive Fraud Sensitivity Optimization Model (AFSOM) that transforms the four validated constructs into a continuous feedback architecture for large-scale transaction networks. The proposed AFSOM can be structured as six interconnected stages: Transaction and Class-Distribution Monitoring, Adaptive

Minority Representation, Dynamic Misclassification-Cost Calibration, Adaptive Ensemble Selection and Weighting, Dynamic Threshold Optimization, and Sensitivity and False-Alert Feedback Recalibration. In the first stage, the system would continuously estimate fraud prevalence, minority density, class overlap, transaction volume, and signs of concept drift. In the second stage, resampling intensity would be adjusted according to minority difficulty rather than through a fixed sampling ratio. This extension would build on adaptive-sampling research showing that difficult minority observations should receive differentiated treatment (Tang & He, 2017). The third stage would calculate example-dependent or segment-dependent misclassification costs using transaction amount, fraud type, customer risk, investigation expense, and expected financial loss, moving beyond a single static cost matrix. The fourth stage would dynamically select or reweight ensemble members according to recent performance across transaction segments, allowing different classifiers to specialize in different fraud patterns. The fifth stage would optimize the classification threshold according to an explicit multi-objective criterion such as maximizing recall subject to a maximum false-positive rate or review-capacity constraint.

Figure 10: Adaptive Fraud Sensitivity Optimization Model (AFSOM) for High-Sensitivity Fraud Detection



The final stage would create a feedback loop through which observed false negatives, false positives, investigator outcomes, and confirmed fraud losses automatically update the earlier stages. Longitudinal research should test AFSOM on real transaction streams rather than perception data and compare static models, isolated imbalance treatments, and the full adaptive architecture using recall, precision, F1-score, AUC-PR, false-positive rate, expected monetary loss, and investigation workload. Future studies should also incorporate graph-based fraud relationships, sequential transaction histories, unsupervised anomaly scores, and concept-drift detection because modern fraud increasingly involves coordinated activity across accounts and devices (Carcillo et al., 2021; Liu et al., 2021). Multi-institutional research could further test whether the model performs consistently across banking, fintech, e-commerce, mobile-wallet, and real-time payment contexts. This proposed AFSOM would extend the present study from a validated explanatory framework into an adaptive operational model in which sampling, costs, ensembles, thresholds, and performance feedback are continuously coordinated around the specific objective of maintaining high fraud sensitivity under changing class-imbalance conditions.

$$\tau = \arg \max_{\tau} \text{Recall}(\tau), \text{FPR}(\tau) \leq \delta$$

CONCLUSION

This study has examined the role of Adaptive Class-Imbalance Learning in strengthening High-Sensitivity Fraud Detection Performance within large-scale digital transaction networks through a quantitative, cross-sectional, case-study-based research design. The empirical findings have demonstrated that all four explanatory constructs, Adaptive Minority-Class Resampling, Cost-Sensitive Learning, Dynamic Classification-Threshold Optimization, and Adaptive Ensemble Learning, have been positively and significantly associated with fraud-detection sensitivity. Among these constructs, Dynamic Classification-Threshold Optimization has produced the strongest effect, with the highest descriptive mean and the largest correlation and standardized regression coefficient, indicating that the ability to adjust fraud-decision thresholds has played a particularly important role in identifying rare fraudulent transactions and reducing false-negative classifications. Cost-Sensitive Learning has emerged as the second-strongest predictor, confirming that unequal misclassification consequences must be considered when fraudulent transactions carry greater financial and operational risk than legitimate transactions that are temporarily flagged. Adaptive Ensemble Learning and Adaptive Minority-Class Resampling have also shown significant contributions, demonstrating that both model-level integration and data-level balancing have remained important components of effective fraud detection. The regression model has explained 70.9% of the variance in High-Sensitivity Fraud Detection Performance, indicating that the four adaptive capabilities have collectively provided substantial explanatory power. All five hypotheses have therefore been supported, and all major research objectives have been achieved. The results have also reinforced Statistical Learning Theory by showing that modifications to sample representation, loss weighting, classification boundaries, and predictive model combinations can significantly influence minority-class recognition. Overall, the study has established Adaptive Class-Imbalance Learning as a multidimensional framework for understanding fraud-detection performance rather than as a single resampling intervention. The evidence has shown that high-sensitivity fraud detection is strongest when adaptive mechanisms operate across the data, learning, decision, and ensemble levels in a coordinated manner within large-scale digital transaction environments.

RECOMMENDATIONS

Based on the findings of this study, organizations operating large-scale digital transaction networks should adopt a coordinated adaptive fraud-detection strategy rather than relying on a single imbalance-handling technique or conventional accuracy-based model evaluation. First, financial institutions, fintech companies, payment processors, digital-wallet providers, e-commerce platforms, and transaction-security teams should implement dynamic classification-threshold optimization because it has emerged as the strongest predictor of High-Sensitivity Fraud Detection Performance. Thresholds should be calibrated according to fraud prevalence, transaction risk, false-positive tolerance, and operational review capacity rather than being fixed permanently. Second, cost-sensitive learning should be incorporated into model training so that fraud-related false negatives receive greater penalties than less damaging classification errors, particularly for high-value or high-risk transactions. Third, adaptive minority-class resampling should be applied selectively by focusing on difficult, underrepresented, and boundary-level fraud observations instead of using excessive uniform oversampling. Fourth, adaptive ensemble learning should be strengthened by combining complementary classifiers and dynamically adjusting model weights according to recent predictive performance, fraud type, or transaction segment. Organizations should also evaluate fraud systems using sensitivity, recall, precision, F1-score, AUC-PR, false-positive rate, expected monetary loss, and review workload instead of relying primarily on overall accuracy. In addition, continuous monitoring should be introduced to detect class-distribution changes, concept drift, emerging fraud patterns, and changes in transaction behavior. The proposed Adaptive Fraud Sensitivity Optimization Model should be considered as an operational framework integrating class-distribution monitoring, adaptive resampling, cost calibration, ensemble weighting, threshold optimization, and performance feedback. Finally, fraud analysts, data scientists, cybersecurity professionals, and risk managers should collaborate during model design and evaluation so that statistical performance remains aligned with operational and financial priorities. Such an integrated approach can improve minority-class

recognition while maintaining acceptable levels of false alerts across large-scale transaction environments.

LIMITATIONS

The study has several limitations that should be considered when interpreting the findings. First, the research has used a cross-sectional design, meaning that data have been collected at a single point in time and have not captured changes in fraud-detection capabilities, transaction behavior, class distributions, or model performance across longer periods. This limitation is important because digital fraud patterns can change rapidly as criminal strategies, payment technologies, customer behavior, and security controls evolve. Second, the study has relied on a case-study-based approach and purposive sampling, which has enabled access to knowledgeable professionals but has limited the extent to which the findings can be generalized to every banking, fintech, payment-processing, e-commerce, or digital-wallet environment. Third, the research has measured the major constructs through a five-point Likert-scale questionnaire, meaning that the results have represented professional perceptions and organizational assessments rather than direct transaction-level measurements of recall, precision, AUC-PR, expected monetary loss, or real-time classification outcomes. Fourth, self-reported responses may have been affected by perceptual bias, social desirability, or differences in respondents' technical understanding of adaptive learning methods. Fifth, although the regression model has explained 70.9% of the variance in High-Sensitivity Fraud Detection Performance, additional factors such as feature engineering, concept-drift detection, graph-based transaction relationships, data quality, model calibration, explainability, system latency, fraud-investigation capacity, and organizational governance have not been incorporated into the current model. Sixth, the study has focused on four adaptive class-imbalance learning constructs and has not directly compared specific machine-learning algorithms or resampling techniques using the same transaction dataset. Finally, the strong statistical relationships identified in this study should not be interpreted as definitive causal evidence because the cross-sectional survey design has established association and predictive contribution rather than experimental causality.

REFERENCES

- [1]. Abdallah, A., Maarof, M. A., & Zainal, A. (2016). Fraud detection system: A survey. *Journal of Network and Computer Applications*, 68, 90-113. <https://doi.org/10.1016/j.jnca.2016.04.007>
- [2]. Abdullah Mohammad, S. (2023). CUGAO₂ and CUINS₂ Quantum Dot Sensitization for Enhanced Photovoltaic Efficiency in Wearable Fiber-Shaped Solar Cells. *American Journal of Advanced Technology and Engineering Solutions*, 3(04), 209-261. <https://doi.org/10.63125/wysdvy71>
- [3]. Arya, M., & Sastry, H. G. (2020). DEAL - 'Deep Ensemble ALgorithm' framework for credit card fraud detection in real-time data stream with Google TensorFlow. *Smart Science*, 8(2), 71-83. <https://doi.org/10.1080/23080477.2020.1783491>
- [4]. Bagga, S., Goyal, A., Gupta, N., & Goyal, A. (2020). Credit card fraud detection using pipeling and ensemble learning. *Procedia Computer Science*, 173, 104-112. <https://doi.org/10.1016/j.procs.2020.06.014>
- [5]. Bahnsen, A. C., Aouada, D., & Ottersten, B. (2015). Example-dependent cost-sensitive decision trees. *Expert Systems with Applications*, 42(19), 6609-6619. <https://doi.org/10.1016/j.eswa.2015.04.042>
- [6]. Bahnsen, A. C., Aouada, D., Stojanovic, A., & Ottersten, B. (2016). Feature engineering strategies for credit card fraud detection. *Expert Systems with Applications*, 51, 134-142. <https://doi.org/10.1016/j.eswa.2015.12.030>
- [7]. Bahnsen, A. C., Stojanovic, A., Aouada, D., & Ottersten, B. (2013). Cost sensitive credit card fraud detection using Bayes minimum risk. *2013 12th International Conference on Machine Learning and Applications*, 333-338. <https://doi.org/10.1109/icmla.2013.68>
- [8]. Barella, V. H., Garcia, L. P. F., de Souto, M. C. P., Lorena, A. C., & de Carvalho, A. C. P. L. F. (2021). Assessing the data complexity of imbalanced datasets. *Information Sciences*, 553, 83-109. <https://doi.org/10.1016/j.ins.2020.12.006>
- [9]. Bartlett, P. L., Jordan, M. I., & McAuliffe, J. D. (2006). Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473), 138-156. <https://doi.org/10.1198/016214505000000907>
- [10]. Barua, S., Islam, M. M., Yao, X., & Murase, K. (2014). MWMOTE: Majority weighted minority oversampling technique for imbalanced data set learning. *IEEE Transactions on Knowledge and Data Engineering*, 26(2), 405-425. <https://doi.org/10.1109/tkde.2012.232>
- [11]. Bei, H. (2021). A statistical approach to cost-sensitive AdaBoost for imbalanced data classification. *Mathematical Problems in Engineering*, 2021, 3165589. <https://doi.org/10.1155/2021/3165589>
- [12]. Bhattacharyya, S., Jha, S., Tharakunnel, K., & Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. *Decision Support Systems*, 50(3), 602-613. <https://doi.org/10.1016/j.dss.2010.08.008>
- [13]. Boucheron, S., Bousquet, O., & Lugosi, G. (2005). Theory of classification: A survey of some recent advances. *ESAIM: Probability and Statistics*, 9, 323-375. <https://doi.org/10.1051/ps:2005018>

- [14]. Branco, P., Torgo, L., & Ribeiro, R. P. (2016). A survey of predictive modeling on imbalanced domains. *ACM Computing Surveys*, 49(2), Article 31, 31-50. <https://doi.org/10.1145/2907070>
- [15]. Buda, M., Maki, A., & Mazurowski, M. A. (2018). A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106, 249-259. <https://doi.org/10.1016/j.neunet.2018.07.011>
- [16]. Carcillo, F., Dal Pozzolo, A., Le Borgne, Y.-A., Caelen, O., Mazzer, Y., & Bontempi, G. (2018). SCARFF: A scalable framework for streaming credit card fraud detection with Spark. *Information Fusion*, 41, 182-194. <https://doi.org/10.1016/j.inffus.2017.09.005>
- [17]. Carcillo, F., Le Borgne, Y.-A., Caelen, O., Kessaci, Y., Oblé, F., & Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences*, 557, 317-331. <https://doi.org/10.1016/j.ins.2019.05.042>
- [18]. Castro, C. L., & Braga, A. P. (2013). Novel cost-sensitive approach to improve the multilayer perceptron performance on imbalanced data. *IEEE Transactions on Neural Networks and Learning Systems*, 24(6), 888-899. <https://doi.org/10.1109/tnnls.2013.2246188>
- [19]. Chen, S., He, H., & Garcia, E. A. (2010). RAMOBoost: Ranked minority oversampling in boosting. *IEEE Transactions on Neural Networks*, 21(10), 1624-1642. <https://doi.org/10.1109/tnn.2010.2066988>
- [20]. Cheng, F., Zhang, J., & Wen, C. (2016). Cost-sensitive large margin distribution machine for classification of imbalanced data. *Pattern Recognition Letters*, 80, 107-112. <https://doi.org/10.1016/j.patrec.2016.06.009>
- [21]. Collell, G., Prelec, D., & Patil, K. R. (2018). A simple plug-in bagging ensemble based on threshold-moving for classifying binary and multiclass imbalanced data. *Neurocomputing*, 275, 330-340. <https://doi.org/10.1016/j.neucom.2017.08.035>
- [22]. Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2018). Credit card fraud detection: A realistic modeling and a novel learning strategy. *IEEE Transactions on Neural Networks and Learning Systems*, 29(8), 3784-3797. <https://doi.org/10.1109/tnnls.2017.2736643>
- [23]. Dal Pozzolo, A., Caelen, O., Johnson, R. A., & Bontempi, G. (2015). Calibrating probability with undersampling for unbalanced classification. 2015 *IEEE Symposium Series on Computational Intelligence*, 159-166. <https://doi.org/10.1109/ssci.2015.33>
- [24]. Dal Pozzolo, A., Caelen, O., Le Borgne, Y.-A., Waterschoot, S., & Bontempi, G. (2014). Learned lessons in credit card fraud detection from a practitioner perspective. *Expert Systems with Applications*, 41(10), 4915-4928. <https://doi.org/10.1016/j.eswa.2014.02.026>
- [25]. Davis, J., & Goadrich, M. (2006). The relationship between Precision-Recall and ROC curves. *Proceedings of the 23rd International Conference on Machine Learning*, 233-240. <https://doi.org/10.1145/1143844.1143874>
- [26]. Debasish, A. (2021). Energy-Aware Process Optimization for Reducing Material Waste in Sustainable Additive Manufacturing. *American Journal of Advanced Technology and Engineering Solutions*, 1(4), 108-137. <https://doi.org/10.63125/j0wk7j48>
- [27]. Debasish, A. (2022). Simulation-Based Thermal Management Framework for Reducing Energy Consumption in Advanced Manufacturing Systems. *International Journal of Scientific Interdisciplinary Research*, 1(01), 335-381. <https://doi.org/10.63125/6fcj7c89>
- [28]. del Río, S., López, V., Benítez, J. M., & Herrera, F. (2014). On the use of MapReduce for imbalanced big data using Random Forest. *Information Sciences*, 285, 112-137. <https://doi.org/10.1016/j.ins.2014.03.043>
- [29]. Esposito, C., Landrum, G. A., Schneider, N., Stiefl, N., & Riniker, S. (2021). GHOST: Adjusting the decision threshold to handle imbalanced data in machine learning. *Journal of Chemical Information and Modeling*, 61(6), 2623-2640. <https://doi.org/10.1021/acs.jcim.1c00160>
- [30]. Fang, Y., Zhang, Y., & Huang, C. (2019). Credit card fraud detection based on machine learning. *Computers, Materials & Continua*, 61(1), 185-195. <https://doi.org/10.32604/cmc.2019.06144>
- [31]. Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8), 861-874. <https://doi.org/10.1016/j.patrec.2005.10.010>
- [32]. Feng, Y., Zhou, M., & Tong, X. (2021). Imbalanced classification: A paradigm-based review. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 14(5), 383-406. <https://doi.org/10.1002/sam.11538>
- [33]. Fernández, A., García, S., Herrera, F., & Chawla, N. V. (2018). SMOTE for learning from imbalanced data: Progress and challenges, marking the 15-year anniversary. *Journal of Artificial Intelligence Research*, 61, 863-905. <https://doi.org/10.1613/jair.1.11192>
- [34]. Fiore, U., De Santis, A., Perla, F., Zanetti, P., & Palmieri, F. (2019). Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. *Information Sciences*, 479, 448-455. <https://doi.org/10.1016/j.ins.2017.12.030>
- [35]. Forough, J., & Momtazi, S. (2021). Ensemble of deep sequential models for credit card fraud detection. *Applied Soft Computing*, 99, 106883. <https://doi.org/10.1016/j.asoc.2020.106883>
- [36]. Galar, M., Fernández, A., Barrenechea, E., Bustince, H., & Herrera, F. (2012). A review on ensembles for the class imbalance problem: Bagging-, boosting-, and hybrid-based approaches. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 42(4), 463-484. <https://doi.org/10.1109/tsmcc.2011.2161285>
- [37]. García, V., Sánchez, J. S., & Mollineda, R. A. (2012). On the effectiveness of preprocessing methods when dealing with different levels of class imbalance. *Knowledge-Based Systems*, 25(1), 13-21. <https://doi.org/10.1016/j.knosys.2011.06.013>
- [38]. Guo, H., Li, Y., Shang, J., Gu, M., Huang, Y., & Gong, B. (2017). Learning from class-imbalanced data: Review of methods and applications. *Expert Systems with Applications*, 73, 220-239. <https://doi.org/10.1016/j.eswa.2016.12.035>

- [39]. Han, H., Wang, W.-Y., & Mao, B.-H. (2005). Borderline-SMOTE: A new over-sampling method in imbalanced data sets learning. In *Advances in Intelligent Computing* (pp. 878-887). https://doi.org/10.1007/11538059_91
- [40]. He, H., Bai, Y., Garcia, E. A., & Li, S. (2008). ADASYN: Adaptive synthetic sampling approach for imbalanced learning. 2008 *IEEE International Joint Conference on Neural Networks*, 1322-1328. <https://doi.org/10.1109/ijcnn.2008.4633969>
- [41]. He, H., & Garcia, E. A. (2009). Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9), 1263-1284. <https://doi.org/10.1109/tkde.2008.239>
- [42]. Johnson, J. M., & Khoshgoftaar, T. M. (2019). Survey on deep learning with class imbalance. *Journal of Big Data*, 6, 27. <https://doi.org/10.1186/s40537-019-0192-5>
- [43]. Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P.-E., He-Guelton, L., & Caelen, O. (2018). Sequence classification for credit-card fraud detection. *Expert Systems with Applications*, 100, 234-245. <https://doi.org/10.1016/j.eswa.2018.01.037>
- [44]. Kim, E., Lee, J., Shin, H., Yang, H., Cho, S., Nam, S.-K., Song, Y., Yoon, J.-A., & Kim, J.-I. (2019). Champion-challenger analysis for credit card fraud detection: Hybrid ensemble and deep learning. *Expert Systems with Applications*, 128, 214-224. <https://doi.org/10.1016/j.eswa.2019.03.042>
- [45]. Krawczyk, B. (2016). Learning from imbalanced data: Open challenges and future directions. *Progress in Artificial Intelligence*, 5, 221-232. <https://doi.org/10.1007/s13748-016-0094-0>
- [46]. Li, Z., Huang, M., Liu, G., & Jiang, C. (2021). A hybrid method with dynamic weighted entropy for handling the problem of class imbalance with overlap in credit card fraud detection. *Expert Systems with Applications*, 175, 114750. <https://doi.org/10.1016/j.eswa.2021.114750>
- [47]. Liu, X.-Y., Wu, J., & Zhou, Z.-H. (2009). Exploratory undersampling for class-imbalance learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 39(2), 539-550. <https://doi.org/10.1109/tsmcb.2008.2007853>
- [48]. Liu, Y., Ao, X., Qin, Z., Chi, J., Feng, J., Yang, H., & He, Q. (2021). Pick and choose: A GNN-based imbalanced learning approach for fraud detection. *Proceedings of the Web Conference 2021*, 3168-3177. <https://doi.org/10.1145/3442381.3449989>
- [49]. López, V., Fernández, A., García, S., Palade, V., & Herrera, F. (2013). An insight into classification with imbalanced data: Empirical results and current trends on using data intrinsic characteristics. *Information Sciences*, 250, 113-141. <https://doi.org/10.1016/j.ins.2013.07.007>
- [50]. Md Ashfaq, S., & Abdullah Mohammad, S. (2022). CNT/PANI Composite Electrodes for Flexible Fiber-Shaped Dye-Sensitized Solar Cells: Synthesis, Characterization, and Photovoltaic Performance. *American Journal of Interdisciplinary Studies*, 3(02), 221-274. <https://doi.org/10.63125/m334mf41>
- [51]. Md Asif Ali Sheak, A., & Risha, A. (2022). The Role of IOT-Driven Energy Analytics in Improving Operational Efficiency and Sustainability in Modern Industrial Facilities. *American Journal of Interdisciplinary Studies*, 3(04), 807-841. <https://doi.org/10.63125/gm63q956>
- [52]. Md Kaium, H. (2023). A GIS-Based Multi-Criteria Decision-Support and Land-Suitability Optimization Framework for Climate-Resilient Urban and Regional Planning. *American Journal of Data Science and Analytics*, 4(06), 43-103. <https://doi.org/10.63125/wvq0sx87>
- [53]. Md. Abdur, R., & Iftekhar, A. (2021). Customer Retention Forecasting in Mobile Wallet Services Using Neural Networks: A Comparative Quantitative Study. *International Journal of Business and Economics Insights*, 1(4), 70-102. <https://doi.org/10.63125/dyrpc387>
- [54]. Md. Rashedul, I., & Md Kaium, H. (2022). Development Of a Three-Dimensional Indoor Geospatial Network and Address Geocoding Framework for Next-Generation Emergency Dispatch and Response-Time Optimization. *Journal of Sustainable Development and Policy*, 1(02), 245-292. <https://doi.org/10.63125/tcw3mp48>
- [55]. Mohammad Abir Chowdhury, S. (2023). Role of AI-Assisted Risk Assessment in Enhancing Audit Planning Efficiency and Risk Prioritization in Public-Safety Organizations. *International Journal of Scientific Interdisciplinary Research*, 4(4), 529-576. <https://doi.org/10.63125/qd6nn437>
- [56]. Mohammad Abir Chowdhury, S., & Tonay, P. (2022). An Empirical Analysis of the Impact of Robotic Process Automation and Continuous Auditing on the Timeliness And Accuracy Of Financial Reporting Assurance. *International Journal of Business and Economics Insights*, 2(4), 73-126. <https://doi.org/10.63125/7p8dcr47>
- [57]. Mohammad Badrul, A., & Md. Mominul, H. (2023). Machine Learning-Based Assessment of Environmental and Public Health Risks in Sewerage Sludge Management Systems. *American Journal of Advanced Technology and Engineering Solutions*, 3(02), 33-77. <https://doi.org/10.63125/y3yfxa92>
- [58]. Mohammad Shoriful Hossan, S. (2022). A Performance-Based Framework for Arc-Flash Hazard Reduction and Electrical Reliability Improvement in Brownfield Industrial Distribution Systems. *Review of Applied Science and Technology*, 1(04), 447-463. <https://doi.org/10.63125/a70vj720>
- [59]. Muhammad Zahidul, I. (2023). Comparative Analysis of Machine-Learning Approaches for Pharmaceutical Sales Forecasting and Market-Demand Prediction. *Review of Applied Science and Technology*, 2(02), 38-95. <https://doi.org/10.63125/etcbqf13>
- [60]. Napierala, K., & Stefanowski, J. (2016). Types of minority class examples and their influence on learning classifiers from imbalanced data. *Journal of Intelligent Information Systems*, 46, 563-597. <https://doi.org/10.1007/s10844-015-0368-1>
- [61]. Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. *Decision Support Systems*, 50(3), 559-569. <https://doi.org/10.1016/j.dss.2010.08.006>

- [62]. Nishat, J., & Mst Kaniz, F. (2022). A Review of the Impact of Covid-19 On Maternal and Child Health Service Utilization in Bangladesh. *American Journal of Scholarly Research and Innovation*, 1(02), 223-269. <https://doi.org/10.63125/vgv5s742>
- [63]. Olowookere, T. A., & Adewale, O. S. (2020). A framework for detecting credit card fraud with cost-sensitive meta-learning ensemble approach. *Scientific African*, 8, e00464. <https://doi.org/10.1016/j.sciaf.2020.e00464>
- [64]. Randhawa, K., Loo, C. K., Seera, M., Lim, C. P., & Nandi, A. K. (2018). Credit card fraud detection using AdaBoost and majority voting. *IEEE Access*, 6, 14277-14284. <https://doi.org/10.1109/access.2018.2806420>
- [65]. Rtayli, N., & Enneya, N. (2020). Enhanced credit card fraud detection based on SVM-recursive feature elimination and hyper-parameters optimization. *Journal of Information Security and Applications*, 55, 102596. <https://doi.org/10.1016/j.jisa.2020.102596>
- [66]. Sahin, Y., Bulkan, S., & Duman, E. (2013). A cost-sensitive decision tree approach for fraud detection. *Expert Systems with Applications*, 40(15), 5916-5923. <https://doi.org/10.1016/j.eswa.2013.05.021>
- [67]. Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLOS ONE*, 10(3), e0118432. <https://doi.org/10.1371/journal.pone.0118432>
- [68]. Seiffert, C., Khoshgoftaar, T. M., Van Hulse, J., & Napolitano, A. (2010). RUSBoost: A hybrid approach to alleviating class imbalance. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans*, 40(1), 185-197. <https://doi.org/10.1109/tsmca.2009.2029559>
- [69]. Shaiful, M., Anisur, R., & Md, A. (2022). A Systematic Literature Review on the Role of Digital Health Twins in Preventive Healthcare For Personal and Corporate Wellbeing. *American Journal of Interdisciplinary Studies*, 3(04), 1-31. <https://doi.org/10.63125/negjw373>
- [70]. Sohony, I., Pratap, R., & Nambiar, U. (2018). Ensemble learning for credit card fraud detection. *Proceedings of the ACM India Joint International Conference on Data Science and Management of Data*, 289-294. <https://doi.org/10.1145/3152494.3156815>
- [71]. Sun, Y., Kamel, M. S., Wong, A. K. C., & Wang, Y. (2007). Cost-sensitive boosting for classification of imbalanced data. *Pattern Recognition*, 40(12), 3358-3378. <https://doi.org/10.1016/j.patcog.2007.04.009>
- [72]. Taha, A. A., & Malebary, S. J. (2020). An intelligent approach to credit card fraud detection using an optimized Light Gradient Boosting Machine. *IEEE Access*, 8, 25579-25587. <https://doi.org/10.1109/access.2020.2971354>
- [73]. Tang, B., & He, H. (2017). GIR-based ensemble sampling approaches for imbalanced learning. *Pattern Recognition*, 71, 306-319. <https://doi.org/10.1016/j.patcog.2017.06.019>
- [74]. Tong, X., Feng, Y., & Zhao, A. (2016). A survey on Neyman-Pearson classification and suggestions for future research. *WIREs Computational Statistics*, 8(2), 64-81. <https://doi.org/10.1002/wics.1376>
- [75]. West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: A comprehensive review. *Computers & Security*, 57, 47-66. <https://doi.org/10.1016/j.cose.2015.09.005>
- [76]. Whitrow, C., Hand, D. J., Juszczak, P., Weston, D., & Adams, N. M. (2009). Transaction aggregation as a strategy for credit card fraud detection. *Data Mining and Knowledge Discovery*, 18, 30-55. <https://doi.org/10.1007/s10618-008-0116-z>
- [77]. Xie, Y., Li, A., Gao, L., & Liu, Z. (2021). A heterogeneous ensemble learning model based on data distribution for credit card fraud detection. *Wireless Communications and Mobile Computing*, 2021, 2531210. <https://doi.org/10.1155/2021/2531210>