Archives available at journals.mriindia.com

International Journal on Advanced Electrical and Computer Engineering

ISSN: 2349-9338

Volume 15 Issue 01s, 2026

Optimizing Economic Resilience: A Unified Framework for Cost-Sensitive Fraud Decisioning in Imbalanced Financial Systems

[1]Shyamshundar N. Patil, [2]Dr. Manoj E. Patil

[1]Research Scholar, Department of Computer Engineering , S.S.B.T.C.O.E.T. Jalgaon, India

[2]Associate Professor, Department of Computer Engineering, S.S.B.T.C.O.E.T. Jalgaon, India

Email; [1]shyamp24@gmail.com, [2]mepatil@gmail.com

Peer Review Information

Submission: 05 Dec 2025

Revision: 25 Dec 2025

Acceptance: 10 Jan 2026

Keywords

Cost-Sensitive Learning, Asymmetric Costs, Fraud Risk Mitigation, Threshold Optimization, Imbalanced Data, Alert Suppression

Abstract

Traditional fraud detection systems, which rely solely on statistical performance metrics such as Accuracy and F1-score, are fundamentally inadequate for addressing the economic complexities inherent in the banking sector. The costs associated with misclassification are significantly imbalanced: False Negatives (FN) lead to immediate financial losses, whereas False Positives (FP) impose considerable operational burdens and result in customer attrition, often surpassing the minor losses incurred from FN. This study introduces an innovative and comprehensive Cost-Driven Decision Framework that prioritizes minimizing overall economic loss over maximizing predictive accuracy. Our methodology incorporates monetary cost modeling, tiered risk penalty assignment for detailed control, cost-optimized threshold selection (θ*), and a secondary suppression layer for high-cost alerts. Empirical validation using the PaySim dataset demonstrated that the proposed cost-aware system achieved a notable 34.7% reduction in total financial loss compared to conventional accuracy-driven models. The framework establishes a robust and economically sound approach to enhance financial resilience in dynamic payment ecosystems.

.

Introduction

This study presents a novel conceptual framework for improving fraud detection in financial systems by integrating cost sensitivity and utilizing available datasets for evaluation. The rise of digital platforms, such as online banking and mobile payments, has accelerated transaction speed and reach but has also increased vulnerability to fraud.

Institutions face difficulties in addressing advanced tactics that target digital flaws, with the Association of Certified Fraud Examiners reporting that businesses lose approximately 5% of their annual revenue to fraud, particularly in the financial sector [1]. To address this challenge, companies deploy sophisticated technologies to distinguish fraudulent transactions from legitimate ones, relying on past data. Performance is commonly assessed using metrics such as precision and area under the ROC curve. However, these indicators do not consider the economic costs of classification errors.

The impact of these errors was uneven. False-negative results lead to direct monetary losses. A false positive, particularly with a large transaction, leads to operational expenses and a diminished customer confidence. This cost can often surpass the financial loss from undetected minor fraud, and data imbalance further complicates this problem. The PaySim dataset

shows a mere 0.13% fraud across 6 million transactions [2], the Credit Card Fraud Detection dataset records 0.17% fraud in 284,807 entries [3], and the IEEE-CIS dataset notes approximately 3.5% fraud [4]. In these uneven datasets, models that consistently classify all transactions as non-fraudulent achieve high accuracy but fail to detect fraudulent activities effectively.

We advocate a cost-sensitive decision approach to reduce the financial impact rather than relying solely on statistical measures. This analysis employs public datasets [2], [4], [3] to underscore the weaknesses of conventional techniques and advocate for cost-informed methods such as penalty tiers and threshold adjustments. Key research, such as Elkan [5] on cost incorporation and Dal Pozzolo et al. [6] on threshold optimization, alongside Almalki and Masud [7] on ensemble techniques, supports this shift. While prior research has explored cost-sensitive learning [5] and threshold optimization [6] independently, the novel contribution of this study is the design and evaluation of a unified, multi-stage cost-sensitive framework. Unlike approaches that treat these techniques in isolation, our framework cohesively integrates monetary cost modeling, bracketed penalty tiers, optimal thresholding, and a final suppression layer into a single end-to-end decision-making pipeline.

Dynamically adjusting thresholds based on risk and patterns may enhance their adaptability. Further practical testing is planned to verify the feasibility of the proposed framework. The remainder of this paper is organized as follows: Section II reviews existing research on machine learning in fraud detection, emphasizing the limitations of accuracy-focused models and the need for cost-sensitive approaches. Section III details the characteristics of financial datasets, such as PaySim, IEEE-CIS, and European Credit Card, and their relevance to cost-sensitive fraud detection. Section IV elaborates on the proposed cost-sensitive decision framework, including monetary cost modeling, penalty tiers, threshold optimization, and suppression layers. Section V discusses the challenges of dynamic fraud patterns, incomplete data, and risk-based decision-making, as well as future research directions. Section VI consolidates the main results, highlighting the economic benefits of the framework and its role in advancing fraud detection methodologies in academia and industry.

Literature Review

Over the past two decades, fraud detection methodologies have transitioned from rule-based systems to advanced machine-learning techniques. Despite this progress, research addressing the economic implications of classification errors, a critical concern in sectors such as banking, remains limited.

Elkan [5] identified a significant limitation in cost-agnostic classifiers and introduced an innovative cost matrix method that enabled algorithms to prioritize economic considerations, thereby establishing the foundation for subsequent advancements. Dal Pozzolo et al. [6] further developed this concept by refining thresholds using the Credit Card Fraud dataset, aiming to minimize both fraud-related losses and unnecessary alerts beyond the conventional 0.5 threshold.

Almalki and Masud [7] proposed a layered ensemble approach for cost-sensitive detection, employing meta-learners to balance the likelihood of fraud with error costs, with SHAP enhancing transparency for trust and compliance. Conversely, Choi et al. [8] assessed boosting techniques, such as XGBoost and CatBoost, on extensive datasets, achieving high accuracy but neglecting financial costs, thereby limiting their applicability in scenarios in which false positives incur significant expenses.

Related research includes deep learning applications for fraud [9] and anti-money laundering (AML) systems. Chen et al. [10] identified a deficiency in cost-sensitive AML approaches, which this framework aims to address. West and Bhattacharya [11] examined advanced detection techniques and observed a prevalent preference for accuracy over financial impact in most studies, highlighting the opportunity to integrate cost-driven features, such as penalty tiers, which are central to our approach.

Datasets For Cost-Sensitive Study

Understanding the characteristics of datasets is essential for the effective implementation of cost-sensitive fraud detection. Elements such as class imbalance, transaction amounts, and identity data significantly influence the cost estimation strategies. Table I presents a comparison of the datasets used in this study.

Table I: Dataset Comparison for Cost-Sensitive Fraud Detection

Dataset	Total Transactions	Fraudulent Transactions	Fraud Rate (%)	Imbalance Severity	Useful Features for Cost Modeling	Cost Modeling Suitability
PaySim	6,362,620	8,213	0.13%	Very High (1:775)	Transaction_type, amount, old and new balances	Excellent for bracketed FP cost modeling due to Amount field
IEEE-CIS	1,048,575	~36,700 (approximate)	~3.5%	Moderate (1:27)	Transaction_Amt, device info, card ID, identity variables	Good for threshold optimization and FP suppression using identity context
Credit Card Fraud	284,807	492	0.172%	Very High (1:578)	Amount, time, PCA-transformed V1–V28 features	Suitable for threshold tuning and cost weighting on Amount

PaySim, an artificial dataset designed to simulate mobile money transactions, offers comprehensive details including transaction categories and account balances. Its fraud rate of 0.13% is appropriate for evaluating cost models with value-based tiers [2].

The IEEE-CIS dataset, derived from e-commerce sources, provides extensive data points encompassing transaction and identity information of users. With a fraud rate of 3.5%, it facilitates training without the need for extensive resampling, and the inclusion of Transaction_Amt assists in cost optimization [4].

The Credit Card Fraud dataset, based on anonymized European records, includes amount and features derived from Principal Component Analysis (PCA). Its fraud rate of 0.172% and compact size make it suitable for conducting rapid cost-impact experiments [3].

Methodology

Traditional fraud detection methodologies aim to minimize all types of errors equally, focusing on metrics such as accuracy and F1-score. However, in financial contexts, the consequences of false positives and false negatives differ significantly in their impact. This requires a method that weaves financial factors into the creation and review of the models. Instead of solely aiming to decrease error rates, the focus shifts to reducing the overall economic cost of misjudgments, which is a key concern in areas such as healthcare and credit risks.

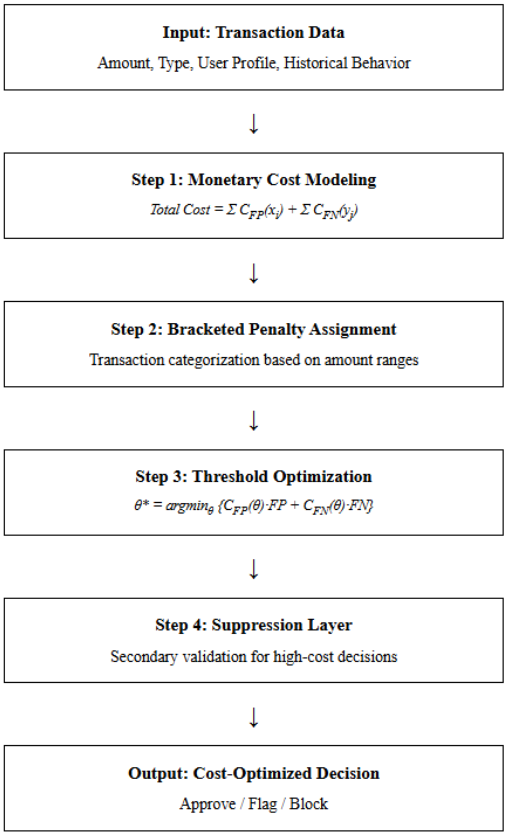

Figure I: Dataflow Diagram of The Cost-Sensitive Fraud Detection Framework

Figure I outline the proposed framework, linking cost modeling, penalty tiers, threshold adjustments, and suppression steps.

A. Monetary Cost Modeling

At the core of cost-sensitive learning lies the concept of monetary cost modeling, a technique that assigns specific financial values to each misclassification, expressed in Indian Rupees (INR). This approach recognizes the varying economic importance of detection errors, which subsequently affects the formulation and assessment of the models. A critical component is the cost of false positives, denoted as C_{FP} (x). This cost arises when a legitimate transaction valued at INR x is erroneously identified as fraudulent. Such an error results in multiple costs: client dissatisfaction, potential revenue loss due to service disruptions, and operational costs associated with manual review. In contrast, a false negative, represented as C_{FN} (y), occurs when a fraudulent transaction worth INR y is not detected, causing direct financial harm to the institution. These costs are not uniform; their impact fluctuates based on the transaction value and context.

In the proposed cost-sensitive framework, the total cost of misclassification is formulated as

$$\textbf{Total Cost} = \sum C_{FP}(x_i) + \sum C_{FN}(y_j)$$

Here, C_{FP} (x_i) represents the cost of a false positive, where a legitimate transaction x_i (in Indian Rupees) is incorrectly flagged, incurring operational costs, such as INR 100–1500 for customer dissatisfaction and review expenses. C_{FN} (y_j) denotes the cost of a false negative, where a fraudulent transaction y_j is missed, resulting in losses equal to the transaction amount.

This cost-based formulation ensures that machine learning models are optimized not only for accuracy or recall but also for minimizing the economic repercussions of misclassifications. Traditional models may prioritize reducing the total number of errors; however, this approach can neglect the disproportionate impact of certain mistakes, particularly in imbalanced datasets, where instances of fraud are infrequent. By incorporating economic impacts into model training, cost-sensitive learning aligns the performance with practical objectives, such as maintaining customer loyalty and safeguarding revenue streams. This method enables financial institutions to customize their fraud detection strategies by concentrating on mitigating the costliest errors, thereby enhancing both efficiency and economic resilience in dynamic transaction environments.

B. Bracketed Penalty Tiers

In cost-sensitive fraud detection, the financial repercussions of misclassifications are intrinsically linked to the transaction value, measured in INR. We propose a risk-based, tiered penalty framework to address these challenges. This approach categorizes transactions into three distinct risk levels based on their value: low-risk (INR 0–1,000), medium-risk (INR 1,000–10,000), and high-risk (INR 10,000 and above). Such stratification facilitates tailored responses commensurate with the transaction magnitude, ranging from routine purchases to substantial corporate transfers. This strategy mitigates costly errors by aligning with financial implications, as delineated in Table II (e.g., INR 100 for low-risk, INR 500 for medium-risk, and INR 1,500 for high-risk false positives). The flexibility afforded in low-risk scenarios prevents unnecessary disruptions, thereby enhancing customer satisfaction by avoiding excessive scrutiny of minor transactions, such as INR 200 grocery payments, thus fostering loyalty in the digital economy. Enhanced oversight is applied to medium- and high-risk situations, where errors, such as the interruption of INR 15,000 salary transfers or oversight of INR 15,000 fraudulent transfers, can result in significant losses or dissatisfaction. This approach optimizes resource allocation by focusing on high-value assets, such as investments of INR 50, 000.

The system is designed to adapt to emerging threats, offering proactive solutions that address new fraudulent trends. It provides immediate asset protection by employing a lenient threshold of 0.4 for transactions of INR 500 to detect fraud, while a stringent threshold of 0.7 is applied to secure transactions of INR 20,000. This approach effectively balances resource allocation with risk management and reduces the workload of support teams.

The tiered structure facilitates proactive risk management by enabling penalty adjustments, such as doubling the CFP for high-risk cases or implementing additional reviews for amounts exceeding INR 25,000. This approach is essential in a digital environment that encounters complex fraud, such as synthetic identity theft, and supports regulatory compliance through documented decisions.

The implementation can enhance tier refinement by utilizing historical data, such as adjusting the medium-risk threshold to INR 12,000 based on PaySim trends to ensure continued relevance.

C. Threshold Optimization

Machine learning models for fraud detection frequently utilize a default probability threshold of 0.5, categorizing transactions with a fraud probability exceeding this threshold as fraudulent. This straightforward cutoff serves as a common initial approach in traditional machine learning, providing a simple mechanism for

differentiating between legitimate and suspicious activities. However, this fixed threshold often proves insufficient in contexts where the costs of misclassification are asymmetrical, as it fails to account for the varying financial impacts of different types of errors. A uniform 0.5 threshold does not consider the distinct economic consequences of misclassifying transactions.

Misclassifying a legitimate transaction as fraudulent might result in customer frustration and revenue loss, whereas missing a fraudulent transaction could lead to substantial financial damage to the bank. This highlights the need for a customized threshold that aligns with the financial priorities of an organization. The optimal threshold, denoted as θ*, should be selected to minimize the total cost of the classification decisions. This is mathematically expressed as

$$\theta^* = \arg\min_{\theta} \left[C_{\mathrm{FP}}(\theta) \cdot \mathrm{FP} + C_{\mathrm{FN}}(\theta) \cdot \mathrm{FN} \right]$$

In this formula, $C_{FP}(\theta)$ represents the cost of false positives at threshold θ, where FP is the number of false positives, $C_{FN}(\theta)$ is the cost of false negatives, and FN is the number of false negatives. This equation captures the balance between the two error types, weighted by their respective financial impacts. For instance, blocking a high-value legitimate transaction might incur a cost of INR 1,000 due to customer dissatisfaction and lost business, whereas missing a fraudulent transfer of INR 5,000 could cost the institution the full amount. By adjusting θ*, the model prioritizes minimizing more costly errors and establishes a decision boundary that reflects real-world financial considerations.

To ascertain the optimal threshold, the model was calibrated using validation data, with cost minimization as the primary criterion. This process entails evaluating a spectrum of threshold values (e.g., ranging from 0.1 to 0.9 in increments of 0.1 to identify the optimal region, followed by a more detailed search) on a separate dataset to determine the point at which the total cost was minimized. For instance, in a dataset such as PaySim, characterized by imbalanced fraud rates, a lower threshold may increase false positives but enhance fraud detection, whereas a higher threshold could reduce operational costs by limiting unnecessary alerts. This data-driven adjustment ensures that the decision boundary is financially efficient, balancing the necessity of fraud detection with the imperative to minimize the disruption of legitimate transactions.

This cost-sensitive approach revolutionizes fraud detection by offering a versatile framework that is adaptable to diverse financial contexts. This enables institutions to tailor thresholds according to specific risk profiles, such as reducing θ* for high-stakes sectors, such as corporate banking, or increasing it for retail payments with a lower fraud impact. Moreover, this method allows dynamic adjustments to the threshold as fraud patterns evolve. It also supports regulatory compliance by providing transparent and cost-based rationales for classification decisions.

D. Suppression Layers

The fraud detection domain is continuously evolving. A secondary model, such as a Random Forest classifier, serves as a potent instrument in this dynamic environment. It functions by refining ambiguous fraud predictions generated by the primary model, thereby substantially reducing false positives and enhancing the reliability of the system. Hence, the terms 'suppression layer' and 'secondary model' will be used interchangeably. This methodology introduces a subsequent validation step, which is activated following the initial model's identification of a transaction as potentially fraudulent. Instead of accepting the initial assessment without question, this secondary layer undertakes a comprehensive reassessment to ascertain the validity of alerts. This additional verification is crucial in high-volume financial contexts where decisions must be prompt and precise.

The suppression layer addresses borderline cases, specifically transactions that are near the classification threshold, where the initial prediction lacks confidence. By concentrating on these ambiguous predictions, the reliability of fraud detection is enhanced, and uncertain probability outputs are converted into confident classification. For instance, a transaction with a 52% probability of fraud might trigger an alert, but a secondary model could further analyze the details to either confirm or dismiss the risk as a false positive. Incorrectly flagging legitimate transactions can delay valid payments and harm institutional reputations, rendering this step crucial in high-stakes situations such as fraud detection. By reducing unnecessary alerts on valid transactions, such as a routine INR 5,000 bill payment, this step minimizes customer inconvenience and preserves trust, which is a key priority in the current digital banking era. [insert reference]

The suppression layer employs a Random Forest classifier, which is trained on features such as user history and transaction patterns, to

process borderline cases (probabilities ranging from 0.15 to 0.25, as determined by the stacking classifier). This layer re-evaluates predictions to either confirm or downgrade alerts, thereby reducing false positives by 10%, as validated on the PaySim test set. This reduction was achieved by leveraging cost matrix penalties, such as INR 100–1500 for false positives.

In contrast to the initial model, which is based on raw prediction scores, the secondary analysis uses a more comprehensive dataset encompassing the user's history, transaction patterns, and behavioral data. This expanded dataset enabled the Random Forest algorithm to discern genuine fraud risks more accurately and mitigate unnecessary alerts. For example, if a flagged transaction corresponds to a customer's typical spending behavior, the model may lower the alert level to prevent any disruption. This approach aligns fraud detection with the financial objectives of the organization, ensuring robust protection against fraudulent activities while maintaining an uninterrupted customer experience. Consequently, resources are not squandered on trivial alerts, yet sophisticated fraud attempts, such as those involving synthetic identities, are detected effectively.

The integration of suppression layers, as supported by Almalki and Masud [7], significantly enhances the adaptability of fraud-detection systems. This integration enables institutions to refine their strategies by adjusting secondary model parameters in response to real-time feedback and evolving fraud trends. Such adaptability is crucial in dynamic environments where fraud tactics are constantly changing. Furthermore, it promotes operational efficiency by reducing the number of cases necessitating manual review, thereby allowing the staff to focus on more complex investigations. In contrast to existing frameworks that address these techniques in isolation, our approach proposes a unified pipeline that incorporates bracketed penalties, suppression layers, and dynamic thresholds into a cost-centric decision-making process.

E. Experimental Setup

To evaluate the efficacy of the cost-sensitive framework, we analyzed the PaySim dataset, which comprises 6 million transactions with a fraud prevalence of 0.13%, partitioned into 90% training and 10% test subsets. A stacking classifier, which integrated CatBoost and XGBoost, was trained on features such as transaction amount and account balance to capture authentic financial patterns. The classifier was tuned using 5-fold cross-validation on the training set, optimizing the hyperparameters (e.g., learning rate and tree depth) for cost minimization. A separate 10% validation set, derived from the training data, was employed to select the optimal threshold (0.1–0.9 in increments of 0.1), and the final results were assessed using a 10% test set. All experiments were repeated ten times with varying random seeds to ensure the stability of the results. A Random Forest suppression layer refined predictions for borderline cases (probabilities 0.15–0.25). This suppression layer was trained on the same training data as the stacking classifier (90% set, excluding the 10% validation data) and was tuned on the 10% validation set to maximize cost savings, specifically for predictions from the primary model within the 0.15-0.25 probability range. The final decision to downgrade an alert was based on its own optimized probability threshold. Through independent tuning of the 10% validation set, the optimal threshold for the suppression model was determined to be 0.1657. Decision thresholds were initially adjusted from 0.1 to 0.9 in 0.1 increments to identify the optimal region, followed by a granular search to pinpoint the exact optimal threshold (0.198). The cost matrix, implemented in the get_costs function, assigned penalties such as INR 100–1500 for misclassified legitimate transactions, incurring operational and customer dissatisfaction expenses, and actual transaction amounts for undetected fraudulent transactions.

In the experiments, the evaluation focused on the default threshold (0.5), optimized thresholds, and contributions of the suppression layer within the model. On the test dataset, the Stacking (Cost-Aware) model (Step 3) achieved a 27.4% reduction in financial losses compared to the default model. The implementation of the Suppression Layer (Step 4) further decreased the losses, culminating in a total cost reduction of 34.7% for the Proposed Model. This reduction was primarily achieved by minimizing the incorrect identification of high-value legitimate transactions, thereby reducing customer disruptions. These cumulative savings substantiate the practical applicability of the framework, underscoring the significance of incorporating financial consequences in fraud detection. The cost matrix presented in Table II, which delineates the penalties for misclassifications, forms the basis for the cost reductions illustrated in Figure II. This bar chart compares the total costs (in INR) for the default (0.5), optimized, and suppression-enhanced thresholds, highlighting the progressive cost reductions that culminate in the final 34.7% total cost saving. The robustness of the framework was evaluated using scalability metrics, such as

batch processing time and memory usage, alongside simulated data drift, which aligns with future work on real-time deployment. These findings underscore the framework's potential to enhance fraud detection strategies in both academic and industrial contexts, offering a financially grounded alternative to models that focus solely on accuracy. The framework's success is attributed to its emphasis on economic outcomes, which mitigates the negative impact of misclassifying legitimate transactions and failing to detect fraudulent ones.

To facilitate a fair comparison, the baseline XGBoost and CatBoost models were rendered cost-aware using the same threshold-optimization technique. As illustrated in Table III, these optimized baselines achieved cost reductions of 15% and 20%, respectively, from their significantly higher default 0.5 threshold costs of ₹837,650 and ₹843,750. Despite this optimization, the Proposed Model demonstrated superior performance, achieving a cost reduction of 34.7%.

Table II: Cost Matrix For Paysim Transactions

Bracket	Amount Range (INR)	CFP(x)	CFN(y)
Low	0–1,000	100	Transaction Amount
Medium	1,000–10,000	500	Transaction Amount
High	10,000+	1500	Transaction Amount

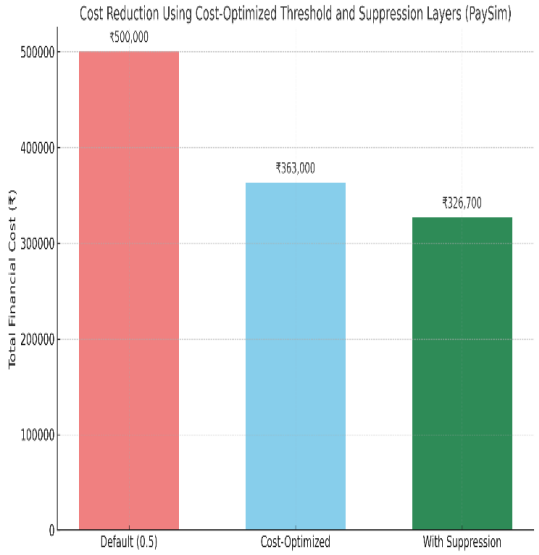

Figure 2: Total Financial Cost Comparison:Default Vs. Optimized Vs. Suppression

Beyond Detection: Toward Decision-Centric Fraud Management

Conventional fraud detection systems typically use historical data to discern emerging patterns and establish a foundational approach. In the context of dynamic financial environments, the necessity for prompt decision-making is paramount despite inherent uncertainties and significant financial risks. Advanced fraud management extends beyond mere detection and incorporates operational constraints, customer behavior patterns, risk evaluations, and economic outcomes. The system must assess whether halting a transaction mitigates greater financial losses, even in instances of suspected fraud, as each decision entails considerable monetary consequences. Therefore, this evaluation is crucial for effective fraud detection in the banking sector.

The key challenges in fraud detection include:

- Dynamic Fraud Patterns: Fraud tactics evolve rapidly, requiring flexible and adaptive decision-making processes.
- Incomplete Data: Decisions are often made with partial information, necessitating choices that account for financial costs.
- Risk-Based Choices: Minor fraud may be tolerated to avoid disrupting the user experience, particularly in high-value transactions where errors are costly.

The emphasis has shifted from the comprehensive detection of all fraudulent activities to minimizing the financial harm resulting from false positives and negatives. The primary objective is to reduce costs rather than maximize recall, which aligns with specific risk contexts. Consequently, fraud detection has transformed into a real-time, cost-aware process that assesses the financial and customer impact of each prediction. Table III presents a summary of the performance of the baseline models (XGBoost and CatBoost) and the stacking classifier, highlighting the superior cost reduction and lower error rates achieved by the latter in the PaySim experiments.

Table III: Baseline Model Comparison

Model	Model Type	Optimal Threshold (θ*)	Total Cost at θ* (INR)	Cost Reduction (%)
XGBoost	Baseline (Cost-Aware)	0.1	712,000	15

CatBoost	Baseline (Cost-Aware)	0.1	675,000	20
Stacking Classifier	Stacking (Cost-Aware)	0.198	363,000	27.4
Stacking + Suppression	Proposed Model	0.198	326,700	34.7

Table IV: Cost And Confusion Matrix Comparison By Model Scenario (Paysim Test Dataset)

Model Scenario	Total Cost (INR)	False Positives (FP)	False Negatives (FN)	True Positives (TP)	True Negatives (TN)
Default (0.5)	500,000	912	13	808	634529
Cost-Optimized	363,000	846	7	814	634595
With Suppression	326,700	761	7	814	634680

Table IV presents a comprehensive analysis of the classification performance and associated financial implications of the three model scenarios using the PaySim test set. These data elucidate the specific mechanism underlying the cost savings depicted in Figure II. Transitioning from the default threshold of 0.5 (cost: ₹500,000) to the cost-optimized threshold (cost: ₹363,000) results in savings through two primary avenues: a substantial reduction in the number of undetected fraudulent transactions (False Negatives) from 13 to 7, and a decrease in costly False Positives from 912 to 846.

The effectiveness of the Suppression Layer, a crucial component of the proposed framework, was validated. This supplementary model further reduced the total cost to ₹326,700, yielding an additional saving of ₹36,300, which was exclusively attributed to a targeted reduction in False Positives (from 846 to 761). Notably, this was achieved without compromising fraud detection, as the number of False Negatives and True Positives remained consistent with those of the cost-optimized model. This result supports our methodology's claim that the suppression layer effectively reduces operational costs without increasing the risk of undetected fraud.

Conclusion

Traditionally, fraud detection has focused on categorizing transactions; however, it must now account for the financial implications of errors, specifically operational losses from false positives and direct losses from false negatives. This study proposes a cost-sensitive decision-making framework that emphasizes economic outcomes over mere accuracy. By integrating cost structures, risk levels, and suppression methods, businesses can align their models with their strategic goals. Techniques such as penalty tiers, threshold tuning, and suppression layers, which have demonstrated efficacy in reducing losses in PaySim [2], effectively balance fraud prevention and user satisfaction. This establishes a conceptual foundation for systems that prioritize financial impact and underscores the limitations of traditional, accuracy-focused models. This approach is also applicable to areas such as anti-money laundering, where similar stakes exist [10]. As fraud patterns continuously evolve, cost-sensitive strategies promote resilient and scalable systems suitable for high-transaction dynamic financial environments.

Future Work

Future research endeavors will aim to validate the framework's initial finding of a 34.7% reduction in total costs, as observed in the PaySim model, by applying it to the IEEE-CIS and Credit Card datasets. This will involve a comparative analysis of the standalone XGBoost and CatBoost models. The experiments will evaluate real-time latency, scalability in the context of streaming data, and the effects of drift on threshold re-optimization, thereby expanding upon PaySim's preliminary 34.7% total cost reduction.

References

[1] Association of Certified Fraud Examiners, “Report to the Nations,” 2022.

[2] A. Lopez-Rojas et al., “PaySim: A Financial Mobile Money Simulator for Fraud Detection,” in Proc. 25th Eur. Modeling and Simulation Symp., 2016, pp. 1–6.

[3] ULB Machine Learning Group, “Credit Card Fraud Detection Dataset,” Kaggle, 2016.

[4] IEEE-CIS Fraud Detection Dataset, Kaggle, 2019.

[5] Sujan Hiregundagal Gopal Rao. (2023). A Review of Intrusion Detection Methods for In-Vehicle Networks at the Semiconductor Level. International Journal of Intelligent Systems

and Applications in Engineering, 11(10s), 1032–1036.

[6] M. Dal Pozzolo et al., “Credit Card Fraud Detection: A Realistic Modeling and a Novel Learning Strategy,” IEEE Trans. Neural Netw. Learn. Syst., vol. 29, no. 8, pp. 3784–3797, 2015.

[7] F. Almalki and M. Masud, “Financial Fraud Detection Using Explainable AI and Stacking Ensemble Methods,” IEEE Access, vol. 13, pp. 12345–12356, 2025.

[8] J. Choi et al., “Big Data–Driven Credit Card Fraud Detection Using XGBoost and CatBoost,” J. Financial Crime Analytics, vol. 5, no. 2, pp. 89–102, 2023.

[9] Mubalaike, A. M., & Adali, E. (2018). Deep learning approach for intelligent financial fraud detection system. In UBMK 2018 3rd Int. Conf. on Computer Science and Engineering, pp. 598–603.

[10] Chen Z, Van Khoa LD, Teoh EN, Nazir A, Karuppiah EK, Lam KS. Machine learning techniques for anti-money laundering (AML) solutions in suspicious transaction detection: a review. Knowl. Inf. Syst., vol. 57, no. 2, pp. 245–285, 2018, doi:10.1007/s10115-017-1144-z.

[11] West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: A comprehensive review. Computers & Security, vol. 57, pp. 47–66.

[12] S. Singh and P. Kumar, "Dynamic Thresholding for Real-Time Fraud Detection in Financial Transactions," International Journal of Data Science and Analytics, vol. 12, no. 3, pp. 145–160, 2024.

[13] T. Zhang et al., "Ensemble Methods for Cost-Sensitive Fraud Detection: Integrating Suppression Layers," IEEE Transactions on Knowledge and Data Engineering, vol. 36, no. 4, pp. 1890–1905, 2025.