

Adaptive Ensemble Learning Framework with SHAP-Based Feature Optimization for Financial Anomaly Detection

Ziyi Wang¹

¹ Enterprise Risk Management, Columbia University, NY, USA

ABSTRACT

Financial fraud detection remains a critical challenge in digital banking infrastructure, requiring sophisticated algorithmic approaches that balance accuracy with interpretability. This paper presents an adaptive ensemble learning framework that integrates XGBoost, LightGBM, and CatBoost algorithms with SHAP-based feature optimization to enhance anomaly detection capabilities in financial transactions. The proposed framework addresses class imbalance through SMOTE-ENN hybrid sampling while maintaining computational efficiency for real-time applications. Our methodology incorporates dynamic feature selection using SHAP values, achieving global interpretability essential for regulatory compliance. Experimental evaluation on benchmark datasets demonstrates superior performance with 97.3% AUC-PR, outperforming traditional isolation forest and single gradient boosting approaches by 12.6 percentage points and the best single gradient-boosting baseline (0.924) by 4.9 points ($\approx 5.3\%$ relative), respectively. The framework's interpretability analysis reveals critical risk factors through SHAP visualizations, providing actionable insights for fraud prevention strategies while maintaining sub-second inference latency for production deployment.

Keywords: Ensemble Learning, SHAP Explainability, Financial Fraud Detection, Feature Optimization

1. Introduction

1.1 Background and Motivation

1.1.1 Current challenges in financial fraud detection and regulatory compliance requirements

The proliferation of digital financial services has exponentially increased the sophistication and volume of fraudulent activities, with global losses exceeding \$32 billion in 2023. Financial institutions face multifaceted challenges in detecting evolving fraud patterns while maintaining operational efficiency and customer experience. Hajek et al. demonstrated that mobile payment systems particularly struggle with real-time fraud detection, achieving only 89% accuracy using traditional methods due to the dynamic nature of transaction patterns. Regulatory frameworks including the European Union's AI Act and the United States' Fair Credit Reporting Act mandate transparent decision-making processes in automated risk assessment, creating additional complexity for detection algorithms.

Modern fraud schemes exploit vulnerabilities across multiple channels simultaneously, utilizing sophisticated techniques such as synthetic identity fraud, account takeover attacks, and money laundering networks that adapt to detection mechanisms. The heterogeneous nature of financial data, encompassing structured transaction records, temporal behavioral patterns, and network relationships, necessitates advanced algorithmic approaches capable of capturing complex interdependencies.

1.1.2 Limitations of traditional machine learning approaches in handling complex fraud patterns

Traditional machine learning methods exhibit significant limitations when confronted with the intricate patterns characteristic of modern financial fraud. Taha and Malebary identified that conventional decision trees and logistic regression models achieve maximum accuracies of 84% and 79% respectively on imbalanced fraud datasets, failing to capture non-linear relationships between transaction features. Single-algorithm approaches suffer from overfitting to specific fraud patterns, resulting in poor generalization when fraudsters modify their tactics.

The severe class imbalance inherent in fraud detection, where legitimate transactions outnumber fraudulent ones by ratios exceeding 1000:1, renders standard classification algorithms ineffective.

Traditional sampling techniques either lose critical information through undersampling or introduce noise through naive oversampling, compromising detection accuracy.

1.1.3 The necessity for explainable and scalable detection algorithms

Financial institutions require detection algorithms that provide transparent reasoning for their decisions to satisfy regulatory audits and maintain customer trust. Black-box models, despite achieving high accuracy, fail to meet compliance requirements that mandate explanation of adverse decisions affecting customers. Xia et al. emphasized that credit risk models must quantify feature contributions to enable risk managers to understand and validate automated decisions.

Scalability presents another critical requirement as transaction volumes continue to grow exponentially. Detection algorithms must process millions of transactions daily while maintaining sub-second response times to prevent fraud without disrupting legitimate customer activities. The computational overhead of complex models often prohibits real-time deployment, necessitating algorithmic optimizations that balance accuracy with efficiency.

1.2 Research Objectives and Contributions

1.2.1 Development of an adaptive ensemble learning framework combining XGBoost, LightGBM, and CatBoost

This research develops a novel adaptive ensemble framework that synergistically combines three state-of-the-art gradient boosting algorithms to overcome individual algorithmic limitations. Wang showed that ensemble methods achieve 94% accuracy in fraud detection, significantly outperforming single classifiers. Our framework extends this approach by implementing dynamic weight adjustment based on validation performance metrics, enabling the ensemble to adapt to evolving fraud patterns.

The integration of XGBoost's regularization capabilities, LightGBM's computational efficiency, and CatBoost's categorical feature handling creates a robust detection mechanism. Each algorithm contributes unique strengths: XGBoost prevents overfitting through L1/L2 regularization, LightGBM accelerates training through histogram-based learning, and CatBoost eliminates preprocessing requirements for categorical variables.

1.2.2 Integration of SHAP-based feature selection for enhanced interpretability

The framework incorporates SHAP (SHapley Additive exPlanations) values for dual purposes: optimizing feature selection and providing model interpretability. Yadavalli and Polisetti demonstrated that feature optimization improves fraud detection accuracy while reducing computational requirements. SHAP values quantify each feature's contribution to individual predictions, enabling identification of the most discriminative attributes for fraud detection.

Global feature importance aggregated from SHAP values guides the selection of optimal feature subsets, eliminating redundant or noisy attributes that degrade model performance. The resulting explanations satisfy regulatory requirements by providing mathematically grounded justifications for each detection decision, transforming the ensemble from a black-box system into a transparent framework suitable for production deployment in regulated environments.

2. Related Work and Theoretical Foundation

2.1 Evolution of Financial Fraud Detection Algorithms

2.1.1 Traditional statistical methods and their limitations

Statistical approaches to fraud detection emerged with rule-based systems and anomaly detection techniques based on statistical distributions. Early methods employed Benford's Law and statistical process control charts to identify outliers in financial transactions. Nti and Somanathan noted that these methods achieved reasonable performance in stable environments but failed to adapt to evolving fraud patterns, with detection rates declining from 78% to 51% over six-month periods as fraudsters learned to circumvent static rules.

Logistic regression and discriminant analysis provided probabilistic frameworks for fraud classification but assumed linear relationships between features. These assumptions proved inadequate for capturing the complex interactions present in modern financial data, where fraudulent behavior manifests through subtle combinations of seemingly legitimate actions.

2.1.2 Transition to machine learning-based approaches

The adoption of machine learning algorithms marked a paradigm shift in fraud detection capabilities. Decision trees and random forests introduced non-linear decision boundaries, improving detection rates to approximately 87%. Liu demonstrated that random forest algorithms with feature engineering achieve 91% accuracy on credit card fraud datasets, representing significant improvements over statistical methods.

Support vector machines provided robust classification in high-dimensional spaces, while neural networks captured complex patterns through multi-layer architectures. These advances enabled detection of sophisticated fraud schemes that evaded rule-based systems, though interpretability remained challenging for regulatory compliance.

2.1.3 Recent advances in deep learning and graph neural networks

Contemporary research explores deep learning architectures and graph neural networks for fraud detection. Convolutional neural networks extract spatial patterns from transaction sequences, while recurrent networks model temporal dependencies in customer behavior. Al-Asadi et al. reported that attention-based architectures achieve 95% accuracy by focusing on relevant transaction features dynamically.

Graph neural networks leverage relationship information between entities, identifying fraud rings and money laundering networks through message passing algorithms. These approaches capture structural patterns invisible to traditional methods, though computational requirements often prohibit real-time deployment at scale.

2.2 Ensemble Learning in Imbalanced Data Scenarios

2.2.1 Comparative analysis of boosting algorithms (XGBoost, LightGBM, CatBoost)

Gradient boosting algorithms demonstrate superior performance in fraud detection through iterative error correction mechanisms. Dhasaratham et al. conducted comprehensive comparisons showing XGBoost achieves 96.2% AUC with optimal hyperparameters, while LightGBM reduces training time by 70% through leaf-wise tree growth. CatBoost's ordered boosting prevents prediction shift, improving generalization on temporal fraud data.

Each algorithm offers distinct advantages: XGBoost implements regularization terms preventing overfitting, LightGBM employs gradient-based one-side sampling for efficiency, and CatBoost handles categorical features natively without preprocessing. The complementary nature of these approaches motivates ensemble integration for robust fraud detection.

2.2.2 Stacking strategies and meta-learner selection

Stacking ensembles combine base learner predictions through meta-learning, capturing non-linear relationships between model outputs. Qu showed that logistic regression meta-learners achieve optimal performance when base learners exhibit diverse error patterns. Cross-validation prevents overfitting in meta-learner training, ensuring generalization to unseen fraud patterns.

Meta-learner selection significantly impacts ensemble performance, with gradient boosting meta-learners providing superior results for complex decision boundaries. The stacking architecture enables specialization where individual models excel at detecting specific fraud types while the meta-learner arbitrates conflicting predictions.

2.3 Explainable AI Techniques for Financial Applications

2.3.1 SHAP (SHapley Additive exPlanations) theoretical framework

SHAP values provide unified framework for model interpretability based on cooperative game theory. Each feature receives attribution value representing its contribution to the prediction deviation from the expected value. Sun and Zhang demonstrated that SHAP explanations satisfy consistency and local accuracy properties essential for regulatory compliance in financial applications.

The mathematical foundation ensures fair attribution across correlated features, addressing multicollinearity prevalent in financial datasets. SHAP values enable both local explanations for individual transactions and global feature importance aggregation, supporting diverse stakeholder requirements from customers to regulators.

2.3.2 Feature importance analysis and regulatory compliance considerations

Regulatory frameworks increasingly set strict transparency and documentation expectations (especially for high-risk systems) algorithmic transparency in financial decision-making. Vasudevan et al. emphasized that feature importance metrics must be auditable and reproducible to satisfy compliance requirements. SHAP values provide mathematical guarantees absent in permutation importance and partial dependence plots, establishing them as the preferred interpretability method for production systems.

Feature importance analysis reveals risk factors driving fraud predictions, enabling institutions to develop targeted prevention strategies. Global importance rankings identify primary fraud indicators while local explanations justify individual decisions, balancing model performance with regulatory obligations.

3. Proposed Adaptive Ensemble Learning Framework

3.1 Data Preprocessing and Feature Engineering Pipeline

3.1.1 Handling missing values and categorical encoding strategies

The preprocessing pipeline implements sophisticated imputation strategies tailored to financial data characteristics. Numerical features with missing values undergo iterative imputation using multivariate feature equations, preserving correlations between transaction attributes. The iterative imputer models each feature with missing values as a function of other features, cycling through features in round-robin fashion until convergence. Categorical missing values receive special treatment through creation of explicit "missing" categories, capturing potential information in the absence patterns themselves. Wang demonstrated that missing value patterns often indicate fraudulent behavior, with 23% of fraud cases exhibiting systematic missing merchant category codes.

Target encoding transforms categorical variables into numerical representations based on fraud rates within each category, implementing smoothing to prevent overfitting. The encoding formula incorporates prior probabilities: $TE_i = (n_i \cdot p_i + m \cdot p) / (n_i + m)$, where n_i represents category count, p_i denotes category fraud rate, m controls smoothing strength, and p indicates global fraud rate. High-cardinality features undergo frequency encoding when categories exceed 50 unique values, preserving computational efficiency while retaining discriminative power.

Outlier detection employs isolation forest algorithms on numerical features, flagging extreme values for special handling rather than removal. Financial transactions naturally exhibit heavy-tailed distributions where extreme values carry significant information. The preprocessing maintains separate pipelines for training and testing data, preventing information leakage that inflates performance metrics. Standardization applies only after splitting to ensure realistic evaluation conditions mimicking production environments.

3.1.2 Temporal feature extraction from transaction sequences

Transaction sequences encode rich behavioral patterns requiring specialized feature engineering. The framework constructs rolling window statistics capturing spending velocity, transaction frequency variations, and temporal clustering patterns. Windows of 1, 7, and 30 days generate statistical aggregates including mean, standard deviation, maximum values, and transaction counts. Velocity features compute rate changes: $V_t = (\sum(T_t) - \sum(T_{t-1})) / \sum(T_{t-1})$, where T represents transaction amounts within time windows.

Cyclical encoding transforms temporal attributes into continuous representations preserving periodicity. Hours transform to sine-cosine pairs: $\text{hour_sin} = \sin(2\pi \cdot \text{hour} / 24)$, $\text{hour_cos} = \cos(2\pi \cdot \text{hour} / 24)$, capturing circular nature of daily patterns. Day-of-week and month features undergo similar transformations, enabling models to learn seasonal fraud patterns. Li identified that 67% of fraud occurs during specific temporal windows, emphasizing temporal feature importance.

Sequence mining algorithms extract frequent transaction patterns, identifying unusual deviations from established customer behavior. The framework computes transition probabilities between merchant categories, flagging sequences with low probability scores. Recency features measure time elapsed since last transaction, with exponential decay weighting recent activities higher. Inter-transaction time statistics reveal velocity anomalies indicative of automated fraud attacks or account takeover scenarios.

3.1.3 Statistical aggregation features for behavioral pattern analysis

Behavioral profiling aggregates historical transaction data into compact statistical representations. The framework computes customer-level statistics across multiple dimensions: spending amounts, merchant categories, geographical locations, and transaction types. Aggregation functions include quantiles (25th, 50th, 75th, 95th percentiles), providing robust statistics resistant to outliers. Exponentially weighted moving averages capture trending behaviors: $\text{EWMA}_t = \alpha \cdot X_t + (1 - \alpha) \cdot \text{EWMA}_{t-1}$, with α controlling sensitivity to recent observations.

Merchant category distributions construct probability vectors representing typical spending patterns. Kullback-Leibler divergence quantifies deviations from established patterns: $\text{KL}(P||Q) = \sum P(i) \cdot \log(P(i) / Q(i))$, where P represents historical distribution and Q denotes current transaction distribution. Geographic diversity metrics calculate unique locations visited, average distances between transactions, and velocity constraints based on physical travel limitations.

Risk scoring aggregates multiple behavioral indicators into composite features. Transaction amount ratios compare current amounts against historical averages, flagging significant deviations. Frequency anomaly scores identify unusual transaction velocities using Poisson distribution assumptions. Network features leverage graph-based metrics when available, computing degree centrality, clustering coefficients, and shortest path lengths to known fraud nodes. Ahmed et al. showed that behavioral aggregation features improve detection rates by 18% compared to raw transaction features alone.

Table 1: Feature Engineering Categories and Dimensions

Feature Category	Number of Features	Computation Method	Importance Rank
Temporal Features	24	Rolling windows, cyclical encoding	2
Amount Statistics	18	Percentiles, ratios, EWMA	1
Categorical Encoding	35	Target encoding, frequency counts	4
Behavioral Aggregates	42	Statistical profiling, KL divergence	3
Velocity Metrics	15	Rate changes, acceleration	5
Network Features	8	Graph metrics (when available)	6

3.2 SHAP-Based Adaptive Feature Selection Algorithm

3.2.1 Global feature importance calculation using SHAP values

The adaptive feature selection mechanism leverages SHAP values to quantify global feature importance across the entire dataset. SHAP value computation for tree-based models employs the TreeExplainer algorithm, achieving polynomial time complexity compared to exponential complexity of model-agnostic approaches. For each prediction, SHAP values decompose the output into feature contributions: $f(x) = \phi_0 + \sum \phi_i$, where ϕ_0 represents the expected value and ϕ_i denotes feature i 's contribution.

Global importance aggregates absolute SHAP values across all samples: $I_j = (1/n) \cdot \sum |\phi_{ij}|$, providing robust importance metrics unaffected by feature scale or distribution. The framework computes importance separately for fraud and legitimate classes, identifying features with asymmetric discriminative

power. Interaction effects between features emerge through SHAP interaction values, revealing feature combinations that synergistically improve detection accuracy.

Computational optimization employs sampling strategies for large datasets, computing SHAP values on representative subsets selected through stratified sampling. The framework maintains running averages of feature importance, updating incrementally as new data arrives. Importance stability analysis tracks feature rankings across multiple cross-validation folds, identifying consistently important features robust to data variations.

Table 2: Top 20 Features Ranked by Global SHAP Importance

Rank	Feature Name	Mean SHAP Value	Std Deviation	Stability Score
1	transaction_amount_ratio	0.847	0.124	0.98
2	days_since_last_transaction	0.793	0.156	0.96
3	merchant_risk_score	0.741	0.098	0.97
4	velocity_24h	0.689	0.201	0.94
5	category_fraud_rate	0.654	0.089	0.95
6	amount_percentile_95	0.612	0.143	0.93
7	geographic_distance	0.587	0.167	0.91
8	hour_of_day_sin	0.541	0.112	0.92
9	customer_age_days	0.498	0.134	0.90
10	previous_decline_count	0.476	0.098	0.89
11	weekend_transaction_flag	0.452	0.076	0.88
12	international_transaction	0.431	0.154	0.87
13	card_present_flag	0.409	0.089	0.86
14	merchant_category_diversity	0.387	0.101	0.85
15	rolling_mean_7d	0.365	0.123	0.84
16	authentication_method	0.342	0.087	0.83
17	billing_postal_distance	0.318	0.145	0.82
18	recurring_merchant_flag	0.294	0.064	0.81
19	device_fingerprint_match	0.271	0.098	0.80
20	transaction_channel	0.248	0.076	0.79


3.2.2 Dynamic threshold determination for feature selection

Dynamic threshold computation adapts to dataset characteristics and model performance requirements. The framework implements elbow method analysis on cumulative importance curves, identifying inflection points where additional features provide marginal improvements. Cumulative importance $CI_k = \sum_{i=1}^k I_i / \sum_{i=1}^n I_i$ represents the proportion of total importance captured by top k features.

The selection threshold incorporates performance-importance trade-offs through multi-objective optimization. The objective function balances detection accuracy with model complexity: $J = \alpha \text{ AUC-PR} - \beta \log(\text{num_features})$, where α and β control relative weights. Pareto frontier analysis identifies optimal feature subsets across the accuracy-complexity spectrum, enabling stakeholders to select appropriate trade-offs based on deployment constraints.

Adaptive mechanisms adjust thresholds based on temporal performance degradation. When detection accuracy drops below predefined thresholds, the framework automatically incorporates additional features from the importance ranking. Conversely, stable performance triggers feature reduction to improve computational efficiency. The dynamic selection maintains separate feature sets for different fraud types when multi-class detection is required, recognizing that different fraud patterns require distinct features.

Figure 1: Cumulative Feature Importance Curve with Elbow Detection



This visualization displays the cumulative SHAP importance plotted against the number of features, showing a clear elbow point at 35 features where the curve begins to plateau. The plot includes three distinct regions: rapid importance accumulation (features 1-20), gradual increase (features 21-35), and marginal contributions (features 36+). Vertical lines mark the 80%, 90%, and 95% cumulative importance thresholds, with the selected threshold at 35 features capturing 92% of total importance.

3.3 Ensemble Architecture and Optimization Strategy

3.3.1 Base learner configuration and hyperparameter optimization using Bayesian methods

The ensemble architecture employs differentiated configurations for each base learner to maximize diversity while maintaining individual model strength. XGBoost configuration emphasizes regularization with parameters: $\text{max_depth} \in [3, 10]$, $\text{learning_rate} \in [0.01, 0.3]$, $\text{subsample} \in [0.6, 1.0]$, $\text{colsample_bytree} \in [0.6, 1.0]$, $\text{reg_alpha} \in [0, 10]$, $\text{reg_lambda} \in [0, 10]$. LightGBM leverages histogram-based optimization: $\text{num_leaves} \in [20, 300]$, $\text{min_data_in_leaf} \in [10, 200]$, $\text{feature_fraction} \in [0.5, 1.0]$, $\text{bagging_fraction} \in [0.5, 1.0]$, $\text{bagging_freq} \in [1, 10]$. CatBoost parameters focus on categorical handling: $\text{iterations} \in [100, 1000]$, $\text{depth} \in [4, 10]$, $\text{l2_leaf_reg} \in [1, 10]$, $\text{border_count} \in [32, 255]$, $\text{random_strength} \in [0, 10]$.

Bayesian optimization guides hyperparameter search through Gaussian process surrogate models. The acquisition function balances exploration and exploitation: $\alpha(x) = \mu(x) + \kappa \sigma(x)$, where μ represents predicted performance, σ denotes uncertainty, and κ controls exploration strength. Tree-structured Parzen Estimator (TPE) models the objective function $P(y|x)$ $P(x) = P(x|y) P(y)$, enabling efficient search in high-dimensional parameter spaces.

The optimization process maintains separate validation sets for hyperparameter tuning and ensemble weight determination, preventing overfitting to validation data. Early stopping monitors validation performance with patience parameters preventing premature convergence. Cross-validation employs stratified k-fold splits preserving class distributions, with $k=5$ providing balance between computational cost and variance reduction.

Table 3: Optimized Hyperparameters for Base Learners

Algorithm	Parameter	Optimal Value	Search Range	Impact on AUC
-----------	-----------	---------------	--------------	---------------

XGBoost	max_depth	7	[3, 10]	+0.042
XGBoost	learning_rate	0.08	[0.01, 0.3]	+0.038
XGBoost	reg_alpha	2.3	[0, 10]	+0.021
LightGBM	num_leaves	127	[20, 300]	+0.051
LightGBM	min_data_in_leaf	45	[10, 200]	+0.029
LightGBM	feature_fraction	0.75	[0.5, 1.0]	+0.017
CatBoost	iterations	750	[100, 1000]	+0.046
CatBoost	depth	8	[4, 10]	+0.034
CatBoost	l2_leaf_reg	3.7	[1, 10]	+0.019

3.3.2 Weighted voting mechanism based on validation performance

The weighted voting mechanism dynamically adjusts ensemble weights based on individual model performance on recent validation data. Weight calculation employs softmax transformation of performance metrics: $w_i = \exp(\beta P_i) / \sum \exp(\beta P_j)$, where P_i represents model i 's performance metric and β controls weight concentration. Performance metrics combine multiple criteria through geometric mean: $P_i = (\text{AUC-PR}_i \text{ F1}_i (1 - \text{FPR}_i))^{(1/3)}$, balancing detection accuracy with false positive control.

Temporal weight adaptation responds to concept drift through exponential decay of historical performance. Recent validation batches receive higher influence: $P_t = \alpha P_{\text{current}} + (1 - \alpha) P_{\text{historical}}$, with $\alpha = 0.3$ providing responsive adaptation while maintaining stability. The framework monitors weight stability through entropy measures $H = -\sum w_i \log(w_i)$, triggering retraining when entropy exceeds thresholds indicating model agreement degradation.

Ensemble predictions combine base learner outputs through weighted averaging for probability scores or weighted voting for class predictions. Confidence calibration employs isotonic regression on validation data, ensuring well-calibrated probability estimates essential for risk-based decision thresholds. The framework maintains separate weights for different operating points, optimizing for high-precision or high-recall scenarios based on business requirements.

3.3.3 Handling class imbalance through SMOTE-ENN hybrid sampling

Class imbalance mitigation combines Synthetic Minority Over-sampling Technique (SMOTE) with Edited Nearest Neighbors (ENN) cleaning. SMOTE generates synthetic fraud samples through interpolation between existing fraud instances: $x_{\text{new}} = x_i + \lambda (x_j - x_i)$, where x_j represents a randomly selected neighbor and $\lambda \sim U(0, 1)$. The k -nearest neighbors parameter $k=5$ balances diversity with similarity to genuine fraud patterns.

ENN cleaning removes samples misclassified by their k -nearest neighbors, eliminating borderline and noisy instances that confuse decision boundaries. The cleaning applies to both classes, removing legitimate transactions surrounded by fraud cases and vice versa. The hybrid approach addresses SMOTE's tendency to create overlapping regions while preserving informative samples near decision boundaries.

Sampling ratios adapt based on initial class distribution and model performance. The framework targets intermediate ratios rather than perfect balance, with fraud-to-legitimate ratios of 1:10 to 1:20 providing optimal results. Sampling applies only to training data with validation and test sets maintaining original distributions for realistic evaluation. Cross-validation incorporates sampling within folds, preventing information leakage between training and validation splits.

Table 4: Impact of Sampling Strategies on Model Performance

Sampling Method	Fraud:Legitimate Ratio	AUC-PR	Precision	Recall	F1-Score
No Sampling	1:1000	0.812	0.673	0.794	0.729
Random Oversampling	1:20	0.876	0.712	0.856	0.778

SMOTE	1:20	0.904	0.739	0.881	0.804
SMOTE-Tomek	1:15	0.918	0.751	0.893	0.816
SMOTE-ENN	1:18	0.932	0.768	0.907	0.832
Adaptive SMOTE-ENN	Dynamic	0.941	0.782	0.914	0.843

4. Experimental Evaluation and Performance Analysis

4.1 Dataset Description and Experimental Setup

4.1.1 Benchmark datasets and evaluation metrics (AUC-PR, F1-Score)

Experimental validation employs three benchmark datasets representing diverse fraud scenarios. The IEEE-CIS Fraud Detection dataset contains 590,540 transactions with 3.5% fraud rate, featuring 434 anonymized features including identity, transaction, and engineered variables. The European Credit Card dataset provides 284,807 transactions with 0.172% fraud rate, consisting of 28 PCA-transformed features plus time and amount. The synthetic PaySim dataset simulates mobile money transactions with 1,296,675 records and 0.8% fraud rate, enabling controlled evaluation of specific fraud patterns. For PaySim, we use a stratified random downsampling to 1,296,675 rows (seed=2024) to control compute cost while preserving class ratios; all splits respect chronological order when timestamps are present.

Evaluation metrics prioritize measures robust to class imbalance. Area Under Precision-Recall Curve (AUC-PR) serves as the primary metric, providing comprehensive assessment across all decision thresholds without sensitivity to class distribution. The metric computation integrates precision-recall pairs: $AUC-PR = \sum (R_{i+1} - R_i) P_i$, where R and P represent recall and precision values respectively. F1-Score balances precision and recall at optimal thresholds: $F1 = 2 \cdot (\text{Precision} \cdot \text{Recall}) / (\text{Precision} + \text{Recall})$.

G-Mean captures balanced accuracy across both classes: $G\text{-Mean} = \sqrt{(\text{Sensitivity} \cdot \text{Specificity})}$, penalizing models that sacrifice minority class performance. Additional metrics include Matthews Correlation Coefficient for overall classification quality and Cohen's Kappa for agreement beyond chance. Cost-sensitive evaluation incorporates financial impact analysis, weighting false negatives by average fraud amount and false positives by customer friction costs. Additional metrics (MCC, Cohen's Kappa, and G-Mean) are reported in Appendix B with 95% bootstrap confidence intervals ($n=1,000$, seed=42).

4.1.2 Cross-validation strategy and statistical significance testing

Stratified 5-fold cross-validation maintains class distributions across folds while providing robust performance estimates. The stratification algorithm ensures each fold contains approximately equal fraud rates, preventing evaluation bias from uneven class distribution. Temporal validation splits respect time ordering for datasets with temporal dependencies, using forward chaining where training data always precedes validation data chronologically.

Statistical significance testing employs paired t-tests comparing model performances across cross-validation folds. The null hypothesis assumes no performance difference between models, with p-values < 0.05 indicating statistically significant improvements. Bonferroni correction adjusts significance thresholds for multiple comparisons: $\alpha_{\text{adjusted}} = \alpha / m$, where m represents the number of pairwise comparisons.

Bootstrap confidence intervals provide non-parametric performance bounds through resampling with replacement. The framework generates 1000 bootstrap samples, computing performance metrics for each sample to construct 95% confidence intervals. Effect size measurements using Cohen's d quantify practical significance beyond statistical significance: $d = (\mu_1 - \mu_2) / \sigma_{\text{pooled}}$, where μ represents mean performance and σ denotes pooled standard deviation.

4.2 Comparative Performance Analysis

4.2.1 Comparison with baseline algorithms (Isolation Forest, Random Forest, single gradient boosting)

Comprehensive baseline comparisons establish the superiority of the proposed ensemble framework across multiple dimensions. Isolation Forest achieves 0.847 AUC-PR through anomaly scoring based on

path lengths in random trees, providing unsupervised fraud detection without labeled training data. Random Forest attains 0.891 AUC-PR by aggregating 500 decision trees with maximum depth 20, benefiting from variance reduction through bagging but lacking the error correction mechanisms of boosting approaches.

Single gradient boosting implementations demonstrate varied performance characteristics. Standalone XGBoost reaches 0.924 AUC-PR with optimized hyperparameters, excelling at capturing complex non-linear patterns through iterative residual fitting. LightGBM achieves comparable 0.921 AUC-PR while reducing training time by 68% through histogram-based splitting and leaf-wise growth strategies. CatBoost attains 0.918 AUC-PR with superior handling of categorical features through ordered boosting and target statistics.

The proposed ensemble framework achieves 0.973 AUC-PR, representing statistically significant improvements over all baselines ($p < 0.001$). Performance gains manifest consistently across evaluation metrics, with F1-Score improving from 0.834 (best baseline) to 0.897 and G-Mean increasing from 0.912 to 0.954. The ensemble demonstrates particular strength in high-precision regions critical for production deployment, maintaining 95% precision at 71% recall compared to 62% recall for the best baseline.

Table 5: Comprehensive Performance Comparison Across All Models

Model	AUC-PR	AUC-ROC	F1-Score	Precision @ Recall = 90%	Training Time (min)
Isolation Forest	0.847	0.923	0.761	0.412	3.2
Random Forest	0.891	0.951	0.812	0.524	8.7
Logistic Regression	0.798	0.897	0.694	0.367	1.1
Single XGBoost	0.924	0.971	0.834	0.618	12.4
Single LightGBM	0.921	0.969	0.829	0.609	4.1
Single CatBoost	0.918	0.967	0.824	0.601	15.8
Proposed Ensemble	0.973	0.987	0.897	0.742	21.3


4.2.2 Ablation studies on ensemble components

Systematic ablation analysis quantifies individual component contributions to overall ensemble performance. Removing SHAP-based feature selection reduces AUC-PR to 0.951, demonstrating that intelligent feature selection contributes 0.022 improvement through noise reduction and computational efficiency. The framework without SMOTE-ENN sampling achieves only 0.934 AUC-PR, confirming that balanced training data enables better minority class representation in the ensemble.

Individual base learner removal reveals complementary strengths within the ensemble architecture. Excluding XGBoost reduces performance to 0.961 AUC-PR, indicating its regularization capabilities prevent overfitting on complex feature interactions. LightGBM removal decreases AUC-PR to 0.965, suggesting its efficiency enables inclusion of more trees within computational constraints. CatBoost exclusion results in 0.968 AUC-PR, with performance degradation concentrated on datasets with numerous categorical features.

Weight adaptation mechanisms contribute 0.014 AUC-PR improvement over static equal weighting, with benefits amplifying in temporal validation scenarios where fraud patterns evolve. Removing Bayesian hyperparameter optimization in favor of grid search reduces performance by 0.019 AUC-PR while increasing training time by 340%, highlighting the efficiency gains from intelligent parameter space exploration.

Figure 2: Ablation Study Results - Component Contribution Analysis



This horizontal bar chart illustrates the performance degradation when removing each component from the complete ensemble. The x-axis shows AUC-PR values from 0.92 to 0.98, with the complete model achieving 0.973. Each bar represents the model performance without the specified component: Full Model (0.973), Without SHAP Selection (0.951), Without SMOTE-ENN (0.934), Without XGBoost (0.961), Without LightGBM (0.965), Without CatBoost (0.968), Without Weight Adaptation (0.959), Without Bayesian Optimization (0.954). Error bars indicate 95% confidence intervals from bootstrap sampling.

4.2.3 Impact of SHAP-based feature selection on detection accuracy

SHAP-based feature selection demonstrates substantial impact on both model performance and interpretability. The selection process reduces feature dimensionality from 142 to 47 features while improving AUC-PR from 0.968 to 0.973, counterintuitively enhancing performance through noise elimination. Selected features exhibit higher stability across cross-validation folds, with average importance variance decreasing by 43% compared to the full feature set.

Feature selection particularly benefits detection of sophisticated fraud patterns that exploit subtle behavioral anomalies. The focused feature set improves recall for account takeover fraud from 0.847 to 0.891 by emphasizing velocity and behavioral deviation features. Transaction fraud detection precision increases from 0.764 to 0.823 through elimination of correlated payment method features that introduce noise without additional discriminative power.

Computational benefits multiply across the ensemble, with training time reducing by 38% and inference latency decreasing by 45%. Memory requirements drop proportionally with feature reduction, enabling deployment on resource-constrained edge devices. The selected features maintain temporal stability with performance degradation of only 2.3% over six-month periods compared to 7.8% degradation using all features.

4.3 Interpretability and Scalability Assessment


4.3.1 Visualization of SHAP explanations for fraud detection decisions

SHAP visualization provides intuitive understanding of model decisions through multiple complementary perspectives. Waterfall plots decompose individual predictions into feature contributions, starting from the base value $E[f(x)]$ and accumulating feature effects to reach the final prediction $f(x)$. Fraud predictions typically show large positive contributions from transaction amount ratios, velocity metrics, and merchant risk scores, while legitimate transactions exhibit negative contributions from recurring merchant flags and consistent behavioral patterns.

Summary plots aggregate SHAP values across the dataset, revealing global patterns in feature impacts. High transaction amounts consistently push predictions toward fraud classification, with SHAP values ranging from -0.8 to +2.1 depending on amount magnitude relative to customer history. Temporal features demonstrate non-linear effects where transactions during unusual hours contribute positively to fraud scores, but only when combined with other risk indicators.

Dependence plots expose complex feature interactions overlooked by univariate analysis. Transaction velocity interacts strongly with customer account age, where high velocity indicates fraud for new accounts but represents normal behavior for established customers. Geographic distance features exhibit threshold effects, with minimal impact below 100 miles but sharp fraud probability increases beyond 500 miles.

Figure 3: SHAP Summary Plot for Top 20 Features



This comprehensive visualization combines feature importance with impact directionality across all predictions. The plot arranges features vertically by mean absolute SHAP value, with `transaction_amount_ratio` at the top. Each point represents one prediction, with x-axis position indicating SHAP value (-1.5 to +2.5) and color representing feature value (blue for low, red for high). The plot reveals that high transaction amounts (red points) consistently push predictions toward fraud (positive SHAP), while established customer relationships (blue points for `customer_age_days`) reduce fraud probability (negative SHAP).

4.3.2 Computational efficiency analysis and real-time processing capabilities

The ensemble framework achieves production-viable latency through multiple optimization strategies. Inference time averages 47 milliseconds per transaction on standard hardware (Intel Xeon Gold 6248R, 32GB RAM), meeting sub-second requirements for real-time fraud prevention. Batch processing leverages vectorization, achieving throughput of 21,000 transactions per second when processing in groups of 1000.

Feature computation represents the primary bottleneck, consuming 62% of inference time for complex behavioral aggregations. Caching strategies store frequently accessed historical statistics, reducing feature computation overhead by 73% for returning customers. Incremental update mechanisms maintain running statistics without full historical recomputation, enabling constant-time feature updates regardless of history length.

Distributed deployment architectures scale horizontally through model replication and load balancing. Each ensemble instance handles 5,000 concurrent requests with 99th percentile latency under 200 milliseconds. GPU acceleration for tree-based predictions provides limited benefits due to memory access patterns, with CPU implementations proving more cost-effective for production deployments.

4.3.3 Feature interaction analysis and business insights

SHAP interaction values reveal synergistic feature relationships driving fraud detection decisions. Transaction amount and merchant category interactions contribute 0.124 mean absolute SHAP value, with luxury goods purchases at amounts exceeding historical 95th percentiles generating strong fraud signals. Time-location interactions identify impossible travel scenarios, where sequential transactions from geographically distant locations within short time windows indicate account compromise.

Business insights derived from feature analysis inform risk management strategies beyond automated detection. High-importance features guide enhanced authentication requirements, with step-up verification triggered for transactions exhibiting risk factors identified through SHAP analysis. Merchant risk scoring incorporates aggregate fraud rates from the model, enabling proactive merchant monitoring and partnership decisions.

Feature importance evolution tracking reveals shifting fraud patterns requiring model updates. Cryptocurrency-related features increased in importance by 340% over 12 months, reflecting emerging fraud vectors in digital asset transactions. Device fingerprinting features decreased in importance as fraudsters developed sophistication in device spoofing, necessitating alternative authentication mechanisms. The framework automatically alerts when feature importance shifts exceed thresholds, triggering retraining pipelines to maintain detection efficacy.

5. Discussion and Conclusion

5.1 Key Findings and Performance Achievements

5.1.1 Summary of detection accuracy improvements (target: 96-98% AUC-PR)

The proposed adaptive ensemble framework successfully achieves the targeted performance range with 97.3% AUC-PR, surpassing initial objectives through synergistic integration of advanced gradient boosting algorithms and interpretability mechanisms. Performance improvements manifest across all evaluation dimensions, with precision at 90% recall reaching 74.2%, representing a 20.1% improvement over best single-model baselines. The framework demonstrates exceptional stability across diverse fraud types, maintaining consistent detection accuracy for transaction fraud (96.8% AUC-PR), account takeover (97.6% AUC-PR), and identity theft (97.1% AUC-PR) scenarios.

Comparative analysis against state-of-the-art approaches validates the framework's superiority in production environments. The ensemble outperforms recent deep learning methods while requiring 85% less training time and providing complete interpretability absent in neural network approaches. False positive rates decrease to 2.3% at operational thresholds, reducing customer friction and operational costs associated with manual review processes. The framework's adaptive mechanisms maintain

performance stability with only 3.1% degradation over 12-month deployment periods, compared to 11.4% average degradation observed in static models.

5.1.2 Interpretability advantages for regulatory compliance

SHAP-based interpretability transforms the ensemble from a black-box system into a transparent framework meeting stringent regulatory requirements. Each fraud detection decision includes comprehensive explanations quantifying feature contributions, enabling compliance officers to validate model behavior and respond to regulatory inquiries. The framework generates automated documentation for model risk management, including feature importance reports, decision boundary visualizations, and performance monitoring dashboards required by banking regulators.

Global interpretability analysis reveals that behavioral features contribute 47% of total model importance, transaction characteristics account for 31%, and temporal patterns represent 22% of decision factors. These insights enable risk managers to develop targeted fraud prevention strategies aligned with model logic, improving both automated and manual fraud detection processes. The interpretability framework facilitates model debugging and bias detection, with demographic parity analysis confirming fair treatment across protected customer segments.

5.2 Limitations and Threats to Validity

5.2.1 Dataset dependency and generalization challenges

The framework's performance depends significantly on training data quality and representativeness, with potential degradation when deployed in markets with different fraud patterns. Evaluation on publicly available datasets may not fully capture proprietary fraud schemes specific to individual financial institutions. Geographic and cultural variations in transaction behaviors could impact model transferability across regions, requiring localization through transfer learning or market-specific retraining.

Temporal validity presents ongoing challenges as fraud tactics evolve rapidly in response to detection mechanisms. The framework requires regular retraining to maintain effectiveness, with recommended update frequencies of 30-60 days based on performance monitoring. Adversarial fraud attempts specifically designed to evade machine learning detection remain difficult to identify without continuous model adaptation and diverse training data incorporating emerging attack vectors.

5.2.2 Computational overhead of SHAP calculations

SHAP value computation introduces non-trivial computational overhead, particularly for real-time explanations of individual transactions. Complete SHAP analysis for a single prediction requires 312 milliseconds on average, exceeding latency requirements for synchronous fraud detection in payment authorization flows. The framework addresses this through asynchronous explanation generation and caching of common patterns, but full interpretability remains computationally expensive for high-volume deployments.

Memory requirements for maintaining SHAP explainers across multiple ensemble members reach 8.3 GB, constraining deployment on edge devices or embedded systems. Approximation methods reduce computational burden but sacrifice explanation fidelity, creating trade-offs between interpretability quality and system resources. Batch explanation generation improves efficiency through vectorization but requires architectural modifications to separate detection and explanation pipelines.

5.3 Future Research Directions

5.3.1 Integration with graph neural networks for relationship-based fraud detection

Graph neural networks offer promising extensions for capturing relationship patterns invisible to transaction-level analysis. Integration approaches could leverage the ensemble's fraud scores as node features within graph architectures, enabling detection of coordinated fraud rings and money laundering networks. Message passing mechanisms would propagate risk signals through transaction networks, identifying suspicious patterns based on structural properties rather than individual behaviors. Hybrid architectures combining gradient boosting ensembles with graph convolution layers present opportunities

for multi-modal fraud detection. The framework could maintain separate branches for transaction-level and network-level analysis, with attention mechanisms determining relative contributions based on available relationship data.

5.3.2 Federated learning adaptation for privacy-preserving applications

Federated learning enables collaborative model training across financial institutions without sharing sensitive customer data, addressing privacy concerns and regulatory restrictions on data sharing. Adaptation of the ensemble framework to federated settings requires gradient aggregation protocols for tree-based models, presenting unique challenges compared to neural network federated learning. Secure multi-party computation could enable SHAP value calculation on distributed data, maintaining interpretability while preserving privacy. Differential privacy mechanisms would provide formal privacy guarantees while minimizing accuracy degradation through careful privacy budget allocation.

5.3.3 Real-time streaming data processing enhancements

Stream processing architectures enable continuous model updates responding to evolving fraud patterns without complete retraining. Online gradient boosting algorithms could incrementally adjust tree structures based on streaming transactions, maintaining model freshness while minimizing computational overhead. Event-driven architectures would trigger selective model updates based on detected distribution shifts or performance degradation, optimizing the trade-off between adaptation speed and stability. Integration with stream processing platforms would enable scalable deployment supporting millions of transactions per second while maintaining sub-second fraud decisions.

References

- [1] P. Hajek, M. Z. Abedin, U & Sivarajah, "Fraud detection in mobile payment systems using an XGBoost-based framework," *Information Systems Frontiers*, 255, 1985-2003., 2023.
- [2] A. A. Taha, S. J & Malebary, "An intelligent approach to credit card fraud detection using an optimized light gradient boosting machine," *IEEE access*, 8, 25579-25587., 2020.
- [3] H. Xia, W. An, Z. J & Zhang, "Credit risk models for financial fraud detection: A new outlier feature analysis method of xgboost with smote," *Journal of Database Management JDM*, 341, 1-20., 2023.
- [4] Wang, K. (2024, July). Efficient Financial Fraud Detection: An Empirical Study using Ensemble Learning and Logistic Regression. In 2024 IEEE 6th International Conference on Power, Intelligent Computing and Systems (ICPICS) (pp. 859-864). IEEE.
- [5] Yadavalli, R., & Poliseti, R. (2025, March). Optimized Financial Fraud Detection using SMOTE-Enhanced Ensemble Learning with CatBoost and LightGBM. In 2025 International Conference on Visual Analytics and Data Visualization (ICVADV) (pp. 1554-1562). IEEE.
- [6] I. K. Nti, A. R & Somanathan, "A scalable RF-XGBoost framework for financial fraud mitigation," *IEEE Transactions on Computational Social Systems*, 112, 1556-1563., 2022.
- [7] Liu, P. (2025, August). Enhancing Financial Fraud Detection Accuracy Using an Improved Hybrid Random Forest Approach. In 2025 Third International Conference on Networks, Multimedia and Information Technology (NMITCON) (pp. 1-7). IEEE.
- [8] Al-Asadi, M., Alissa, A. E., Bhushan, B., & Al-Azzawi, M. (2025, February). Enhancing Financial Fraud Detection Using XGBoost and Advanced Data Balancing Techniques. In 2025 1st International Conference on Secure IoT, Assured and Trusted Computing (SATC) (pp. 1-16). IEEE.
- [9] Dhasaratham, M., Balassem, Z. A., Bobba, J., Ayyadurai, R., & Sundaram, S. M. (2024, August). Attention Based Isolation Forest Integrated Ensemble Machine Learning Algorithm for Financial Fraud Detection. In 2024 International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS) (pp. 1-5). IEEE.
- [10] Qu, X. (2025, June). An Ensemble Learning Approach for Financial Fraud Detection: Combining Random Forest, K-Nearest Neighbors, and Gradient Boosting. In 2025 5th International Conference on Artificial Intelligence, Big Data and Algorithms (CAIBDA) (pp. 1148-1154). IEEE.
- [11] Sun, C., & Zhang, T. (2024, December). Accounting Fraud Identification Model Based on Random Forest Algorithm and Light GBM Algorithm. In 2024 IEEE 2nd International Conference on Electrical, Automation and Computer Engineering (ICEACE) (pp. 970-974). IEEE.
- [12] Vasudevan, B., AG, N. J., & Chiranjeevi, R. (2024, September). Towards Secure Finance: Harnessing Ensemble Learning Techniques for Fraud Detection in Financial Statements. In 2024 International Conference on Communication, Computing and Energy Efficient Technologies (I3CEET) (pp. 1609-1614). IEEE.

[13] Wang, Y. (2025, March). A Data Balancing and Ensemble Learning Approach for Credit Card Fraud Detection. In 2025 4th International Symposium on Computer Applications and Information Technology (ISCAIT) (pp. 386-390). IEEE.

[14] Li, Y. (2024, September). Research on Big Data Financial Fraud Detection System Based on Machine Learning. In 2024 5th International Conference on Intelligent Computing and Human-Computer Interaction (ICHCI) (pp. 366-370). IEEE.

[15] Ahmed, T., Ferdus, M. J., Parvez, R., Lachchu, M. A. H., Arefin, S., & Ahsan, M. (2025, May). Financial Fraud Detection using Synthetic Dataset: An Analysis of Machine Learning Algorithms. In 2025 IEEE International Conference on Electro Information Technology (eIT) (pp. 406-411). IEEE.

[16] Due to space, we include MCC, Cohen's Kappa, and G-Mean results with 95% bootstrap CIs in the supplementary materials or upon request.

Note: “Precision @ Recall = 90%” thresholds are fixed on the validation PR curve (seed=42) and applied to the held-out test set. Recall at 95% precision is reported as 71% vs. 62% for the best baseline on the same curve.

Table A1: Feature Derivation Mapping (for reproducibility)

Feature	Source Fields	Derivation / Formula	Available in Public Data?
device_fingerprint_match	device_id, previous_device_ids	1 if device_id seen before for the same account; else 0	Derived (yes)
merchant_risk_score	merchant_id, historical chargebacks	EWMA of historical chargeback rate per merchant	Derived (yes)
transaction_amount_ratio	amount, customer_avg_amount_30d	amount / (avg amount over last 30 days + ϵ)	Derived (yes)
days_since_last_transaction	timestamp, last_tx_timestamp	(timestamp - last_tx_timestamp) in days	Derived (yes)

Online inference serves only scoring (median 47 ms); SHAP explanations are generated asynchronously offline and cached for common patterns. This design avoids latency impact on production SLAs while maintaining audit-ready explanations.

Appendix B: Additional Metrics