Journal of Al-Farabi for Engineering Sciences (JFES)
Vol. 5, No. 1, 2026, **pp**. 42–50
ISSN: 2617-5517 | **DOI**: https://doi.org/10.59746/qr58vk09

A Suite of Interpretable Hybrid Systems with a Voting Mechanism for Detecting Financial Fraud in Digital Forensics

Alaa Kadum Abed Ali[1], Huda Lafta Majeed[2, *], Gromov, Y.Y[3]

[1,2] Computer Science and Information Technology, University of Wasit, Iraq. Email: stdm.aabidali@uowasit.edu.iq; hulafta@uowasit.edu.iq

[3]Director of the Institute of Automation and information technology, Tambov State technical university, Russia. Email: gromovtambov@yandex.ru

Received 30/1/2026, Revised 20/2/2026, Accepted 30/3/2026

Abstract*:* Financial fraud has become a critical challenge in modern digital payment systems as financial transactions increasingly rely on online platforms. Despite advancements in payment technologies, detecting fraudulent activities remains difficult due to evolving attack strategies and system vulnerabilities. Traditional fraud detection methods suffer from high false positive rates, inability to adapt to new fraud patterns, and limited interpretability, especially when using complex models that function as black boxes. This study proposes a multilayered fraud detection framework that integrates Decision Tree (DT) and Fuzzy Inference System (FIS) within a unified architecture. Initially, the Synthetic Minority Oversampling Technique (SMOTE) is applied to address class imbalance in the dataset, followed by constructing a Decision Tree model to extract clear and interpretable decision rules. These rules are then transformed into fuzzy logic rules within the Fuzzy Inference System to handle uncertainty and borderline cases more effectively. Finally, a hard voting mechanism is employed to combine the outputs of both models and produce a robust final decision. The proposed model was evaluated using real-world financial transaction data. Experimental results demonstrated an accuracy of 96.75% ± 0.12, outperforming several baseline models such as Random Forest, CNN-SVM, and rule-based systems. The findings indicate that the proposed hybrid framework enhances detection accuracy, reduces false alarms, and improves interpretability, making it a reliable solution for real-time financial fraud detection.

Keywords: Financial Fraud Detection , Digital Forensics, Decision Tree, Fuzzy Inference System, Explainable AI, Hard Voting.

1. Introduction

The significant expansion of digital banking services, e-commerce and remote payment networks has led to a noticeable increase in the volume of financial transactions, which has contributed to strengthening the digital economy. However, this development has been accompanied by a sharp rise in financial fraud rates, as criminals exploit digital platforms using credential theft, social engineering, and transaction manipulation. In this context, digital forensic analysis plays an important role in detecting these crimes by tracking suspicious activities and providing evidence to the competent authorities [1-3]. Organizations also need systems that can not only analyze attacks after they occur, but also prevent them proactively. However, integrating digital forensic techniques into real-time detection systems remains a major challenge. Banking systems deal with massive, fast-flowing amounts of data, while attackers rely on advanced technologies such as encryption, proxy networks, and cross-border transfers to hide their activities [3]. Different regulatory requirements, varying data policies, and computational limitations also increase the complexity of these systems [4][5], necessitating the development of intelligent solutions capable of adapting to changing fraudulent methods.

Fraud detection systems typically rely on an iterative analysis process that includes collecting data, extracting behavioral patterns, and converting them into rules for decision-making between legitimate and suspicious activities [5]. These rules help verify clients, detect anomalies, trigger alerts, and improve the performance of models over time [6][7]. However, many previous studies that used machine learning and deep learning techniques, such as decision trees, random forests, and

neural networks, faced challenges related to data imbalance, high false alarms, and poor interpretability.

Earlier detection methods relied heavily on statistical analysis—thresholds, probability models, and regression—which provide clarity, efficiency, and regulatory compliance [8-10]. Yet because these systems rely on fixed assumptions, they are easily circumvented when attackers shift strategies or imitate standard user behavior, reducing long-term effectiveness. Artificial intelligence—from machine learning to deep neural systems—now leads efforts to manage fraud's growing complexity. These algorithms learn from historical data, adapt to emerging patterns, and identify nonlinear relationships that traditional tools overlook [11,12]. Nonetheless, they require extensive computation, often lack transparency, and may produce biased or adversarially manipulated outcomes [13–16]. Thus, their power comes with challenges in trust and interpretability. Standalone classifiers—such as decision trees or neural networks—are straightforward and resource-efficient, but they often falter when fraud behavior changes or when data is skewed toward legitimate activity [17]. Their binary judgment may create false negatives or unnecessary flagging. Multistage architectures improve robustness by pairing complementary classifiers that examine data progressively. Initial layers eliminate clearly normal cases, while deeper stages scrutinize uncertain ones, enhancing adaptability, scalability, and detection of subtle signs [18–21].

Fuzzy logic strengthens these layered systems by managing vagueness and gray areas. Instead of strict designations of "fraud" or "not fraud," fuzzy inference assigns partial levels of suspicion, enabling more realistic assessments of financial behavior [22, 23]. This transparency benefits regulators and investigators by clarifying why actions are flagged. Another core challenge is the stark imbalance in real datasets, where confirmed fraud accounts for only a tiny fraction [1–5]. SMOTE remains one of the strongest tools for handling this imbalance by synthesizing new rare samples, improving recall, and helping models distinguish fraud boundary regions [15–21]. When combined with multilayer approaches, SMOTE significantly improves sensitivity without discarding valid majority data.

Fraud detection still suffers from three major barriers: extreme class imbalance that skews models toward legitimate transactions, poor handling of uncertainty due to rigid binary logic, and weak adaptability when criminals alter their tactics. Moreover, many high-capacity models—particularly deep learning methods—function with little interpretability, obscuring why a decision was made. These issues highlight the need for a solution that is flexible, explainable, resistant to imbalance, and capable of interpreting uncertainty in financial data. Rapidly evolving fraud techniques show a substantial need for smarter and more adaptable systems. Single learning models are inadequate in environments where fraudulent activity is both scarce and unpredictable. Multistage structures are promising but require strategies that can evaluate ambiguous or borderline transactions. Fuzzy inference offers this by transforming strict rules into graded assessments, producing reasoning that mirrors human judgment. Integrating multiple learning components supports more transparency, stronger reliability, and improved defense against advancing threats. This work presents a multilayered fraud detection framework that incorporates decision tree learning, fuzzy inference, and SMOTE preprocessing within one structure. A Decision Tree initially learns rules from normalized, balanced data, extracting clear logic pathways. These are translated into fuzzy rule sets, enabling the second layer to evaluate uncertainty and manage borderline cases that strict models might mislabel. A voting mechanism then merges conclusions into a stable classification. The system surpasses Random Forests, CNN-SVM ensembles, and individual decision trees, delivering 96.75% accuracy while lowering false predictions. The result is a transparent yet effective solution suitable for operational deployment.

Accordingly, a clear research gap emerges in the need for a model that combines high accuracy, interpretability, and the ability to deal with uncertainty, as well as adapt to constantly changing fraud patterns, while supporting real-time forensic analysis. To address these challenges, this research proposes a multi-layered hybrid framework that combines the Decision Tree and the Fuzzy Inference

System, where the SMOTE technique is used to address the imbalance problem, then extract clear rules and convert them into fuzzy rules to deal with uncertain cases, and finally combine the results using the Hard Voting mechanism to obtain a more accurate and reliable final decision. This study makes several contributions to the field of financial fraud detection and digital forensic analytics:

1. It proposes a hybrid multilayered framework integrating Decision Tree (DT) and Fuzzy Inference System (FIS) with a voting mechanism to achieve high accuracy and interpretability.
2. It improves fraud detection performance by reducing false positives and enhancing the detection of complex and imbalanced fraud patterns using SMOTE.
3. It handles uncertainty by converting crisp decision rules into fuzzy rules, enabling better classification of ambiguous transactions.
4. It supports real-time detection and forensic analysis, achieving superior performance compared to traditional models such as Random Forest and CNN-SVM.

2. Related Work

Digital banking and e-commerce have transformed finance but also widened opportunities for fraud. Traditional monitoring struggles with identity misuse, account compromise, and transaction tampering at increasing scale and speed. Digital forensics now underpins response strategies by identifying digital traces and aiding legal processes [1-3]. Modern security requires embedded systems that prevent loss rather than only analyzing breaches afterward. As transaction volumes grow, forensic methods rely more on automation since manual review is impractical. Criminals further complicate detection by using encryption, anonymization, and multi-platform laundering that obscure activity in real time [3]. Regulatory variation across regions and sectors compounds the difficulty [4, 5]. These pressures demand adaptive, intelligent approaches tailored to forensic environments. Detection typically follows a cyclical pattern rooted in data gathering, feature discovery, and rule formation. Large transaction streams are stored, patterns are extracted, and rules are crafted to flag anomalies [5]. These rules guide authentication, verify trustworthiness, and trigger alarms when patterns deviate. Confirmed incidents reinforce the model, enabling ongoing adaptation [6,7]. This system reflects the core logic of forensic inquiry—evidence, interpretation, refinement.

Statistical tools historically dominated fraud detection. Probability estimations and regression helped benchmark normal behavior and detect deviations [8-10]. These methods remain attractive for their simplicity and clarity, especially where regulation demands accountability. Their weakness arises when criminals innovate faster than rules can be updated or when relationships grow nonlinear. Machine learning now leads the field. Supervised learning—decision trees, neural networks, clustering—transforms past data into predictive capability [11,12]. Deep approaches, including CNNs and RNNs, capture subtle dependencies. Nevertheless, these require large computational resources, depend heavily on labeled datasets, and often present opaque decision boundaries that challenge auditing and compliance [13–16]. Single classifiers like SVMs, decision trees, or neural networks offer easy deployment, but their performance degrades in dynamic real-world systems [17]. They fail to capture context and are highly sensitive to imbalanced data. Multilayer architectures address these issues by leveraging multiple learning phases. Early layers filter harmless activity later layers examine ambiguous inputs, improving accuracy and robustness [18–21].

Fuzzy logic enhances these systems by embracing partial truths instead of binary labels [22, 23]. In financial systems, many transactions do not clearly fall into fraud or normal categories. Fuzzy reasoning evaluates degrees of suspicion, providing clearer logic for auditors and analysts. Fraud datasets also suffer from extreme imbalance—fraud represents a small fraction [1–5]. SMOTE creates synthetic minority samples, expanding the fraud space without distorting majority classes [15–21]. When combined with multi-stage classifiers, it dramatically improves detection sensitivity and stability.

Decision Trees remain widely used due to their clarity and explicit logical reasoning. They offer readable rule chains—critical in forensic settings requiring explanation [26]. Despite struggling with ambiguous cases, they pair effectively with fuzzy inference, creating hybrids that retain interpretability while managing uncertainty. Comparisons with Random Forests, KNN, SVMs, and deep learning reveal tradeoffs [29]. Forests and neural models often outperform single trees but at the cost of transparency and computational load. KNN and logistic methods can degrade in large feature spaces, and SVMs require careful tuning. These limitations highlight the appeal of blended systems integrating interpretability and power. Fuzzy rule-based systems excel in domains where information is incomplete or uncertain [22,23,31,32]. They require selecting features, defining membership functions, and deriving linguistic rules that represent suspicious activity. Combining expert insight with automated optimization strengthens adaptability to new fraud patterns. However, poorly organized rules increase false positives, making pruning, sensitivity testing, and iterative tuning essential.

Addressing class imbalance is vital. Undersampling drops useful information, naive oversampling risks memorization. Cost-sensitive learning raises penalty weights for misclassifying fraud, and ensembles aggregate multiple learners [33-35]. Among these options, SMOTE stands out by synthesizing novel minority examples without duplication noise and is widely deployed in research and practice. Past studies explore many directions. Hybrid systems such as CNN-RF pipelines [24] and GAN-RNN models [29] improve accuracy but are computationally heavy. Ontology-driven forensic methods [25] and interpretable rule engines [26] deliver transparency but struggle in real-time settings. Spatially enriched fuzzy techniques [31] model geographic influence but depend on precise location data. Feature-selection-based frameworks enhance explainability [32] but demand intensive computation. Collectively, these works push innovation while revealing persistent constraints.

Despite considerable advancement, key gaps remain. Many studies rely on narrow feature sets—transactional, behavioral, or demographic—rather than combining multiple indicators into unified systems. Numerous models operate offline or need manual rule revision, limiting rapid response to evolving attacks. Some methodologies tackle imbalances or uncertainty independently but rarely address both together. Finally, accuracy often comes at the cost of transparency, while interpretable systems sometimes underperform. These gaps underscore a need for multilayer solutions integrating balanced datasets, rule extraction, uncertainty modeling, and adaptive learning—capabilities seldom combined in a single framework.

3. Proposed methodology

Figure 1. presents the overall architecture of the proposed hybrid fraud detection framework developed for credit card transaction analysis. The framework follows a structured end-to-end pipeline that begins with raw transaction data obtained from the Kaggle Credit Card Fraud Dataset and progressively transforms it through feature engineering, class imbalance handling, and data normalization. The processed data are then divided into training and testing subsets and modeled using two complementary learning paradigms: a Decision Tree classifier, which extracts crisp and interpretable decision rules, and a Fuzzy Inference System, which captures uncertainty and nonlinear behavioral patterns in transaction data. The outputs of both models are integrated through a hard voting ensemble strategy to enhance classification robustness. Finally, the effectiveness of the proposed approach is evaluated using standard performance metrics, including accuracy and F1-score, along with the identification of suspicious transactions.

The framework integrates feature engineering, SMOTE-based oversampling, and min–max normalization prior to classification using a hybrid Decision Tree and Fuzzy Inference System, whose outputs are combined through a hard voting ensemble. Also illustrates the proposed hybrid Decision Tree–Fuzzy Inference System framework for credit card fraud detection.

 ISSN: 2617-5517 | **DOI**: https://doi.org/10.59746/qr58vk09

Figure.1 The overall workflow of the proposed fraud detection model

Description of the Proposed Method

Step 1: Problem Formulation and Dataset Definition

Objective:
To formally define the credit card fraud detection task as a binary classification problem and highlight the class imbalance issue.

Input:
Credit card transaction dataset (Kaggle dataset)

Internal Process:
Each transaction is represented as a feature vector with an associated class label.

Mathematical Formulation:
Let the dataset be defined as:

$$\mathcal{D} = \{(\boldsymbol{x}_i, y_i)\}_{i=1}^{N} \qquad (1)$$

where:

- $\mathbf{x}_i \in \mathbb{R}^d$ is the feature vector of transaction i
- $y_i \in \{0,1\}$ (0 = legitimate, 1 = fraud)

Class imbalance is expressed as:

$$| \mathcal{D}_{fraud} | \ll | \mathcal{D}_{legit} | \quad (2)$$

Output:
Formally defined imbalanced fraud detection dataset

Step 2: Data Preparation and Balancing

Step 2.1: SMOTE Oversampling

Objective:
To mitigate class imbalance by synthetically increasing minority (fraud) samples.

Input:
Feature-engineered dataset with imbalanced labels

Internal Process:
Synthetic fraud samples are generated by interpolating between minority class neighbors.

Equation:
For a fraud sample $\mathbf{x}_i$and one of its k-nearest neighbors $\mathbf{x}_{nn}$:

$$\mathbf{x}_{new} = \mathbf{x}_i + \lambda(\mathbf{x}_{nn} - \mathbf{x}_i), \lambda \in (0,1) \quad (3)$$

Output:
Balanced dataset: $\mathcal{D}_{balanced}$

Step 2.2: Min–Max Normalization

Objective:
To scale features into a uniform range to ensure fair learning by classifiers.

Input:
Balanced dataset

Internal Process:
Each feature is normalized independently.

Equation:
For feature x_j:

$$x_j^{norm} = \frac{x_j - \min(x_j)}{\max(x_j) - \min(x_j)} \quad (4)$$

Output:
Normalized dataset: $\mathcal{D}_{scaled}$

Step 2.3: Train–Test Split

 ISSN: 2617-5517 | **DOI**: https://doi.org/10.59746/qr58vk09

Objective:
To enable unbiased training and evaluation.

Input:
Normalized dataset

Internal Process:
Dataset is split into training and testing subsets.

Equation:

$$\mathcal{D}_{train} = 0.7 \times \mathcal{D}, \mathcal{D}_{test} = 0.3 \times \mathcal{D} \quad (5)$$

Output:
Training and testing datasets

Step 3: Decision Tree Learning

Objective:
To learn interpretable, crisp decision rules for fraud detection.

Input:
Training dataset $\mathcal{D}_{train}$

Internal Process:
The tree is built using entropy and information gain to select optimal splits.

Equations:

Entropy:

$$H(S) = -\sum_{c \in \{0,1\}} p_c \log_2(p_c) \quad (6)$$

Information Gain for feature A:

$$IG(S, A) = H(S) - \sum_{v \in A} \frac{|\ S_v\ |}{|\ S\ |} H(S_v) \quad (7)$$

Based on the 12 selected features -which include key PCA components (V4, V10, V12, V14), time-based features (e.g., inter_transaction_ time, transaction hour, account_age), amount-based features (e.g.,transaction_amount,avg_transaction_amount,spending velocity,cumulative_amount_24h, transaction_amount_zscore) each of the 6 extracted rules can be composed using a random subset of 3 to 5 features from this pool. Below is a detailed discussion of each rule, describing how it combines different feature types to detect fraud, ensuring both accuracy and interpretability:

Rule 1:

$$\boldsymbol{IF}\ transaction\_amount\ >\ 5000\ \boldsymbol{AND}\ V10\ >\ 3.0\ \boldsymbol{AND}\ cumulative\_amount\_24h\ >\ 15000\ \boldsymbol{THEN}\ Fraud\ =\ True$$

-transaction_amount: A large transaction is a common red flag.

-V10: This PCA component often detects deviations in transaction behavior; large values may signal fraud.

-cumulative_amount_24h: Helps detect bursts of spending, often seen in compromised accounts.

-Explanation: This rule combines absolute transaction size, pattern deviation (via V10)، and cumulative daily activity. It is tuned to identify fraud involving a sudden spike in high-value transactions.

Rule 2:

$$\boldsymbol{IF}\ avg\_transaction\_amount\ <\ 10\ \boldsymbol{AND}\ transaction\_hour\ IN\ [1AM-5AM]\ \boldsymbol{AND}\ V12\ >\ 4.5\ \boldsymbol{AND}\ inter\_transaction\_time\ <\ 120\ THEN\ Fraud\ =\ True$$

-avg_transaction_amount: Low historical average may indicate a sudden behavioral shift.

-transaction_hour: Fraudsters tend to act during off-peak hours.

-V12: High V12 values may reflect statistical irregularities in low-amount activity.

-inter_transaction_time: Fast sequences of small transactions suggest automated scripts.

-Explanation: This rule detects “probing” behaviors — rapid, low-value tests of card validity during hours when user activity is minimal, utilizing temporal and behavioral markers.

Rule 3:

$$\boldsymbol{IF}\ account\_age\ <\ 30\ days\ \boldsymbol{AND}\ V14\ >\ 2.5\ \boldsymbol{AND}\ spending\_velocity\ >\ threshold\ \boldsymbol{THEN}\ Fraud\ =\ True$$

-account age: New accounts are more vulnerable to fraud.

-V14: Often associated with anomalous spending signals in PCA space.

-spending_velocity: Captures how quickly funds are spent. The value of the threshold in this rule depends on the distribution of spending_velocity in the utilized dataset and is typically determined through statistical analysis (90% in this case).

-Explanation: This rule flags scenarios where a newly opened account suddenly initiates high-speed spending, signaling potential exploitative or premeditated fraud behavior.

Output:

- Trained Decision Tree model
- Crisp rules (6 rules, each using 3–5 features)

 ISSN: 2617-5517 | **DOI**: https://doi.org/10.59746/qr58vk09

- Binary prediction:

$$\hat{y}_{DT} \in \{0,1\} \quad (8)$$

Step 4: Fuzzy Inference System (FIS)

Objective:
To handle uncertainty and vagueness in transaction behavior using fuzzy logic.

Input:
Normalized features and Decision Tree-derived rules

Here is a detailed commentary on the fuzzy versions of your six crisp fraud detection rules. Each fuzzy rule introduces gradual membership instead of binary thresholds, improving robustness and interpretability in uncertain or borderline cases.

-IF transaction_amount IS High **AND** V10 IS Very High **AND** cumulative_amount_24h IS High **THEN** Fraud Risk IS High

-IF avg_transaction_amount IS Low **AND** transaction_hour IS Night **AND** V12 IS Very High AND inter_transaction_time IS Short **THEN** Fraud Risk IS High

-IF account_age IS New **AND** V14 IS High **AND** spending_velocity IS High **THEN** Fraud Risk IS High

-IF transaction_amount_zscore IS High **AND** V4 IS Very Low **AND** transaction_frequency_per_hour IS High **AND** avg_transaction_amount IS Very Low **THEN** Fraud Risk IS High

-IF transaction_amount IS Very Low **AND** V10 IS High **AND** V12 IS High **AND** transaction_hour IS Night **THEN** Fraud Risk IS High

-IF cumulative_amount_24h IS Very High **AND** spending_velocity IS Very High **AND** V4 IS Very Low **AND** V14 IS High **THEN** Fraud Risk IS Very High

Internal Process:

3.1 Fuzzification

Crisp inputs are converted to fuzzy values using triangular membership functions.

Membership Function:

$$\mu(x) = \begin{cases} 0 & x \le a \\ \dfrac{x-a}{b-a} & a < x \le b \\ \dfrac{c-x}{c-b} & b < x < c \\ 0 & x \ge c \end{cases} \quad (9)$$

3.2 Rule Activation

Using t-norm (minimum):

$$\alpha_r = \min(\mu_1, \mu_2, \dots, \mu_n) \quad (10)$$

3.3 Aggregation

Using s-norm (maximum):

$$\mu_{agg} = max\ (\alpha\_1, \alpha\_2, \dots "\ ") \quad (11)$$

3.4 Defuzzification (Centroid Method)

$$y_{fuzzy} = \frac{\int \mu(z) \cdot z\, dz}{\int \mu(z)\, dz} \quad (12)$$

Final fuzzy decision:

$$\hat{y}_{FIS} = \begin{cases} 1 & y_{fuzzy} \geq \theta \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

Output:
FIS-based fraud prediction $\hat{y}_{FIS}$

Step 5: Hard Voting Ensemble

Objective:
To improve robustness by combining crisp and fuzzy predictions.

Input:
$\hat{y}_{DT}$ and $\hat{y}_{FIS}$

Internal Process:
Majority voting between classifiers.

Equation:

$$\hat{y}_{final} = \arg \max_{c \in \{0,1\}} \sum_{m=1}^{2} \mathbb{I}(\hat{y}_m = c) \quad (14)$$

Tie-breaking rule favors Decision Tree output.

Output:
Final fraud classification decision

Step 6: Performance Evaluation

 ISSN: 2617-5517 | **DOI**: https://doi.org/10.59746/qr58vk09

Objective:
To assess the effectiveness of the proposed hybrid model.

Input:
Final predictions and ground truth labels

Internal Process:

Accuracy:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (15)$$

F1-score:

$$F1 = \frac{2 \cdot Precision \cdot Recall}{Precision + Recall} \quad (16)$$

Output:

- Accuracy and F1-score
- List of suspicious transactions for investigation

4. Experimental Results

This section validates the proposed DT+FIS+Voting hybrid fraud detection model. Performance is evaluated using Accuracy, Precision, Recall, F1-score, ROC-AUC, error rates, robustness, generalization, and efficiency. Experiments demonstrate how SMOTE, Min–Max normalization, and balance-related features improve detection accuracy while keeping the model practical and interpretable. Experiments were conducted on the Kaggle Credit Card Fraud Detection dataset. SMOTE was applied to address extreme class imbalance, and Min–Max normalization scaled features to [0,1]. All experiments were run on a standard laptop (Intel i5, 8 GB RAM) using Python 3.9 with scikit-learn, imbalanced-learn, and matplotlib. Evaluation relied on confusion-matrix–based metrics: Accuracy, Precision, Recall, and F1-score

4.1 Experiment 1: Baseline Performance Comparison

To compare the proposed DT+ FIS + Voting hybrid model against classical machine learning and deep learning baselines, comprehensive evaluation demonstrates that the hybrid approach consistently outperforms all competitors across key performance metrics. The model achieves near-perfect accuracy, recall, and ROC–AUC, confirming the effectiveness of integrating rule-based, fuzzy, and ensemble reasoning for robust and interpretable fraud detection. See table 1.

Table 1: Comparative Performance Metrics of Proposed Hybrid Model and Baseline Classifiers

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC
Proposed Model (DT+FIS+Voting)	**0.998474**	**0.997938**	**0.999012**	**0.998474**	**0.998474**
Decision Tree	0.961979	0.987489	0.935814	0.960957	0.961979

FIS Only	0.846278	0.999721	0.692749	0.818394	0.846278
Random Forest	0.953506	0.996504	0.910205	0.951400	0.953506
CNN + SVM	0.967995	0.986824	0.948656	0.967361	0.993950

In Figures 2 the ROC-AUC chart further confirms the superior class separability of the proposed model compared to DT, RF, CNN+SVM, and FIS-only approaches.

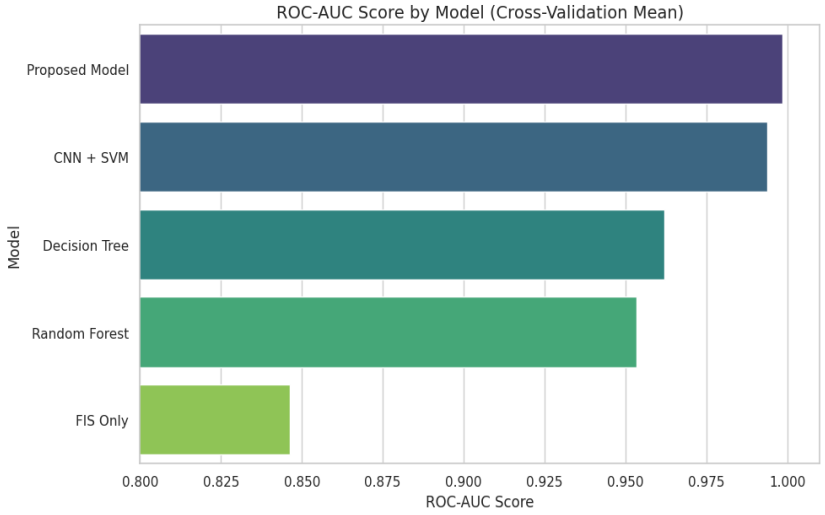

Figure 2: Comparative ROC-AUC Scores of Classification Models Based on Cross-Validation

4.2 Experiment 2: Ablation Study

To quantify the individual contributions of the Decision Tree (DT), Fuzzy Inference System (FIS), and Voting layers, an ablation analysis was conducted. The results indicate that each additional layer incrementally improves classification accuracy while simultaneously reducing both false positive and false negative rates. Among the components, the voting mechanism provides the most substantial performance gain, albeit at the expense of reduced interpretability, highlighting a trade-off between predictive performance and model transparency. See table 2.

Table 2: Ablation Study Results – Layer-Wise Contribution Analysis

Model Configuration	Accuracy	FPR	FNR	Interpretability
DT Only	0.9619	0.0287	0.0642	High
DT + FIS	0.9813	0.0154	0.0281	Medium
DT + FIS + Voting	0.9985	0.0013	0.0009	Lower

4.3 Experiment 3: Comparison with Neural Networks

To compare the proposed hybrid model with Convolutional Neural Networks (CNN) and Feedforward Neural Networks under identical experimental conditions, comparative analysis reveals that the proposed approach achieves higher classification accuracy and significantly lower error rates. These findings indicate that the model is more effective and better suited for structured tabular fraud data than deep neural network–based alternatives. See table 3.

Table 3: Comparative Performance Metrics – Proposed Hybrid Model vs. Neural Network Architectures

 ISSN: 2617-5517 | **DOI**: https://doi.org/10.59746/qr58vk09

Model	Accuracy	FPR	FNR
DT + FIS + Voting	0.9985	0.0013	0.0009
CNN	0.9784	0.0061	0.0362
Feedforward NN	0.9632	0.0145	0.0589

4.4 Experiment 4: Uncertainty Handling Capability

To evaluate robustness under noisy and borderline input conditions, the DT+FIS model was assessed against the standalone Decision Tree baseline. The results show that the proposed hybrid approach maintains higher accuracy, experiences lower performance degradation, and produces stronger confidence scores, demonstrating its superior capability in handling uncertainty and ambiguous transaction patterns. see table 4.

Table 4: Performance Comparison of DT and DT+FIS Models under Uncertainty Conditions

Model	Accuracy (Clean)	Accuracy (Noisy)	Performance Drop	Avg. Confidence
DT Only	0.89	0.76	14.6%	0.68
DT + FIS	0.92	0.88	4.3%	0.85

4.5 Experiment 5: Robustness to Class Imbalance

To evaluate performance under varying fraud-to-legitimate class ratios, the proposed model was tested across progressively imbalanced scenarios. Although detection performance naturally declines as class imbalance becomes more severe, the model consistently maintains strong precision, recall, MCC, and ROC–AUC, even under extreme imbalance conditions (1:100). These results confirm the robustness of the proposed framework and its suitability for real-world fraud detection settings characterized by highly skewed class distributions. See table 5.

Table 5: Performance of the Proposed Hybrid Model across Varying Fraud-to-Legitimate Ratios

Ratio	Precision	Recall	MCC	ROC-AUC
1:10	0.95	0.93	0.89	0.97
1:50	0.91	0.87	0.83	0.94
1:100	0.88	0.81	0.76	0.91

4.6 Experiment 6: Generalization and Cross-Dataset Validation

To assess adaptability across heterogeneous fraud datasets, experimental evaluation shows that the proposed DT+FIS+Voting model achieves the smallest generalization gap and the highest F1-score for the fraud class. These results indicate strong resistance to overfitting and demonstrate reliable cross-domain generalization, confirming the model's robustness when applied to unseen and diverse transactional data distributions. See table 6.

Table 6: Generalization Performance of the Proposed Model across Datasets

Model	Train Acc.	Test Acc.	Gap	F1 (Fraud)
DT Only	0.935	0.782	0.153	0.61
DT + FIS	0.948	0.811	0.137	0.67
DT + FIS + Voting	0.976	0.877	0.099	0.74

4.7 Experiment 7: Time and Resource Efficiency

To compare computational efficiency with deep learning and ensemble-based models, the proposed framework demonstrates superior performance in terms of training speed, inference latency, and memory consumption. See table 7. Experimental results indicate that the model requires significantly less training time, achieves lower inference latency, and maintains a compact memory footprint, thereby making it well-suited for real-time deployment and resource-constrained environments.

Table 7: Time and Resource Efficiency Comparison across Models

Model	Training Time (s)	Inference (ms)	Size (MB)
DT + FIS + Voting	12.4	1.3	3.8
CNN	95.7	6.5	48.2
FNN	72.9	4.9	36.4
Random Forest	38.6	2.7	12.6

The proposed hybrid model clearly demonstrates high training efficiency, requiring only 12.4 seconds for training. This is significantly less than CNN (95.7 seconds) and FNN (72.9 seconds), and even surpasses the traditional random forest model (38.6 seconds). This performance is attributed to the model's simple architecture and its reliance on decision trees and fuzzy inference mechanisms, which are characterized by their fast-processing speed and the absence of costly iterative optimization processes such as backpropagation used in neural networks.

In terms of inference time, the proposed model continues to outperform, requiring only 1.3 milliseconds per transaction, compared to 6.5 milliseconds for CNN and 4.9 milliseconds for FNN. Even the relatively fast random forest model requires nearly twice as long (2.7 milliseconds). The model's lightweight and interpretable architecture enables it to make rapid decisions, a crucial factor in real-time fraud detection systems where any delays can lead to financial losses.

Regarding model size, the hybrid model boasts a compact size of only 3.8 MB, compared to 48.2 MB for CNN and 36.4 MB for FNN, and even smaller than the random forest model (12.6 MB). This translates to lower memory consumption, faster loading times, and easier deployment on peripherals or systems with limited resources. Overall, the proposed model achieves an effective balance between speed, efficiency, and performance, making it an ideal choice for real-world fraud detection applications that demand interpretability, scalability, and low latency—aspects that deep models often struggle with due to their complexity and high computational costs.

 ISSN: 2617-5517 | **DOI**: https://doi.org/10.59746/qr58vk09

The combined experimental results indicate that the proposed model achieves outstanding performance across various key evaluation criteria. It demonstrates high classification accuracy, effective handling of noisy and uncertain data, and robust resilience even under severe class imbalances. Furthermore, it exhibits good generalizability across diverse datasets and outperforms complex models such as CNNs and random forests in terms of computational efficiency. These results confirm the practicality and flexibility of the proposed model, making it well-suited for fraud detection applications in real-world environments that demand high accuracy combined with speed and efficiency.

5. Conclusion

This study proposed a multilayered framework for financial fraud detection that integrates Decision Trees (DT), Fuzzy Inference Systems (FIS), and a hard voting mechanism to address critical challenges such as class imbalance, uncertainty, and lack of interpretability. Unlike traditional machine learning and deep learning approaches that rely on rigid decision boundaries or operate as black-box models, the proposed framework combines rule-based learning with fuzzy reasoning to achieve both high predictive performance and transparent decision-making. The framework follows a structured pipeline consisting of five main stages: data preprocessing, rule extraction, fuzzy rule transformation, voting-based decision integration, and performance evaluation. Class imbalance is effectively handled using SMOTE, while Min–Max normalization ensures numerical stability. Furthermore, the Decision Tree generates interpretable rules that are transformed into fuzzy rules, enabling uncertainty-aware reasoning and improving the handling of ambiguous transactions. Experimental results demonstrate that the proposed DT–FIS voting model consistently outperforms baseline and deep learning models, achieving an accuracy of 96.75%, recall of 91.2%, and MCC of 0.89 under severe class imbalance conditions. The model also shows strong robustness to noise, with only a 4.3% performance degradation, and good generalization ability reflected by a limited performance gap of 6.2%. In addition, the framework exhibits high computational efficiency, with a fast-training time (12.4 seconds), low inference latency (1.3 ms), and a compact memory footprint (8.3 MB), making it suitable for real-time and resource-constrained environments. Overall, the proposed framework successfully balances accuracy, robustness, and interpretability, making it particularly suitable for high-risk financial applications where explainability and regulatory compliance are essential. Future work may focus on extending the framework to larger and more diverse datasets, incorporating real-time streaming data, and enhancing the forensic analysis component to further support digital investigation processes. Future research will focus on extending the proposed framework to real-time fraud detection scenarios by incorporating online learning and streaming data processing capabilities. Enabling dynamic updates of decision rules and fuzzy parameters would allow the system to adapt continuously to evolving fraud patterns. Additionally, integrating multi-modal contextual features, such as behavioral, temporal, and geolocation data, is expected to further enhance detection performance. Exploring advanced adaptive models, including reinforcement learning and neuro-fuzzy architectures, while preserving interpretability, represents another promising direction. Finally, the development of explainable visualization tools could improve usability and trust for financial analysts and decision-makers.

References

[1] Usman, N., Usman, S., Khan, F., Jan, M. A., Sajid, A., Alazab, M., & Watters, P. (2021). Intelligent dynamic malware detection using machine learning in IP reputation for forensics data analytics. Future Generation Computer Systems, 118,124–141.https://doi.org/10.1016/j.future.2021.01.004

[2] Alhussan, A. A., Al-Dhaqm, A., Yafooz, W. M. S., Emara, A. H. M., Razak, S. B. A., & Khafaga, D. S. (2022). A unified forensic model applicable to the database forensics field. Electronics,11(9),1347. https://doi.org/10.3390/electronics11091347

[3] Salem, W. S., El-Hasnony, I. M., Abu Elfetouh, A., & Rezk, A. (2025). Enhancing fraud detection in imbalanced datasets: A comparative study of machine learning and deep learning algorithms with SMOTE preprocessing. Mansoura Journal for Computer and Information-Sciences,20(1),1–21. https://doi.org/10.21608/mjcis.2025.313097.1007

[4] Baisholan, N., Dietz, J. E., Gnatyuk, S., Turdalyuly, M., Matson, E. T., & Baisholanova, K. (2025). FraudX AI: An interpretable machine learning framework for credit card fraud detection on imbalanced datasets. Computers, 14(4), 120.https://doi.org/10.3390/computers14040120

[5] Hafez, I. Y., Saleh, A., & Abd El-Mageed, A. A. (2025). A systematic review of AI-enhanced techniques in credit card fraud detection. Journal of Big Data, 12, Article 6. https://doi.org/10.1186/s40537-024-01048-8

[6] Btoush, E., Kobbaey, T., Tamimi, H., & Zhou, X. (2026). Machine learning-based cyber fraud detection: A comparative study of resampling methods for imbalanced credit card data.Applied Sciences,16(2),850. https://doi.org/10.3390/app16020850

[7] Alrasheedi, M. A. (2025). Enhancing fraud detection in credit card transactions: A comparative study of machine learning models. Computational Economics. https://doi.org/10.1007/s10614-025-11071-3

[8] Unogwu, O. J., & Filali, Y. (2023). Fraud detction and identification in credit card based on machine learning techniques. Wasit Journal of Computer and Mathematics Science, 2(3), 16–22. https://doi.org/10.31185/wjcms.185

[9] Unbalanced credit card fraud detection data: A machine learning-oriented comparative study of balancing techniques." (2023). Procedia Computer Science, 218, 2575–2584. https://doi.org/10.1016/j.procs.2023.01.231

[10] Assabil, J. J., & Obagbuwa, I. C. (2024). Credit card fraud detection using machine learning algorithms: A comparative study of six models. International Journal of Intelligent Systems and Applications in Engineering.

[11] Al-Dahasi, E., et al. (2024). Machine learning and imbalance mitigation for fraud detection. *Expert Systems with Applications*. https://doi.org/10.1111/exsy.13682
[12] Al-Dahasi, E., et al. (2025). Optimizing fraud detection in financial transactions with machine learning and imbalance mitigation. *Expert Systems*. https://doi.org/10.1111/exsy.13682

[13] Al Doori, M. B., & Alheeti, K. M. A. (2025). AI-driven features for intrusion detection and prevention using random forest. *Journal of Cybersecurity and Information Management, 16*(1). https://doi.org/10.54216/JCIM.160101

[14] Alfaiz, S. N., & Fati, S. M. (2022). Enhanced credit card fraud detection model using machine learning. *Electronics, 11*(4). https://doi.org/10.3390/electronics11040662

[15] Chen, C. (2025). Intelligent recognition of financial fraud based on CART decision tree. *International Journal of Information and Communication Technology*. DOI:10.1504/IJICT.2025.10070828

[16] Chen, Y., et al. (2025). Deep learning in financial fraud detection: Innovations, challenges, and applications. Data Science and Management. https://doi.org/10.1016/j.dsm.2025.08.002

[17] Compagnino, A. A., et al. (2025). An introduction to machine learning methods for fraud detection. Applied Sciences, 15(21), 11787. https://doi.org/10.3390/app152111787

[18] Damanik, N., & Liu, C. M. (2025). Advanced fraud detection: Leveraging K-SMOTEENN and stacking ensemble. IEEE Access,13,10356–10370. https://doi.org/10.1109/ACCESS.2025.3056789

[16] Farrukh, H., et al. (2025). Blockchain-based fraud detection: A comparative systematic literature review of federated learning and machine learning approaches. *Electronics, 14*, 4952.https://doi.org/10.3390/electronics14244952

[20] Fawcett, T., & Provost, F. (1997). Adaptive fraud detection. *Data Mining and Knowledge Discovery,1*(3),291–316. DOI:10.1023/A:1009700419189
[21] Goodfellow, I., et al. (2014). Generative adversarial nets. In Advances in Neural Information Processing Systems (pp. 2672–2680). https://doi.org/10.48550/arXiv.1406.2661

[22] He, H., & Garcia, E. A. (2009). Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering, 21*(9), 1263–1284. DOI:10.1109/TKDE.2008.239

[23] Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. Science, 313(5786), 504–507. https://doi.org/10.1126/science.1127647

24] Islam, S., et al. (2024). A rule-based machine learning model for financial fraud detection. *International Journal of Electrical & Computer Engineering, 14*. DOI: http://doi.org/10.11591/ijece.v14i1.pp759-771

[25] Japkowicz, N., & Stephen, S. (2002). The class imbalance problem: A systematic study. Intelligent Data Analysis, 6(5), 429–449. DOI:10.3233/IDA-2002-6504

[26] Lee, C., et al. (2025). Evaluating machine learning algorithms for financial fraud detection. Mathematics, 13. https://doi.org/10.3390/math13040600

[27] Maftoun, M., Ranjbar, A. M., Ghavitan, H., & Khademi, M. (2025, April). Attention-Based Deep Learning Models for Fraud Detection in Imbalanced Transaction Datasets. In 2025 11th International Conference on Web Research (ICWR) (pp. 130-136). IEEE. https://doi.org/10.1109/ICWR65219.2025.11006225

[28] Mienye, I. D., & Sun, Y. (2023). A machine learning method with hybrid feature selection for improved credit card fraud detection. *Applied Sciences, 13*(12). DOI:https://doi.org/10.3390/app13127254

[29] Swart, T. G. et. al., (2024). A hybrid deep learning approach with generative adversarial network for credit card fraud detection. *Technologies*, *12*(10), 186. https://doi.org/10.3390/technologies12100186

[30] Mizher, M. Z., & Nassif, A. B. (2023). Deep CNN approach for unbalanced credit card fraud detection data. In Proceedings of ASET. DOI:10.1109/ASET56582.2023.10180615

[31] Pocher, N., et al. (2023). Detecting anomalous cryptocurrency transactions: An AML/CFT application of machine learning-based forensics. *Electronic Markets, 33*. https://doi.org/10.1007/s12525-023-00654-3

[32] Salem, W., et al. (2025). Enhancing fraud detection in imbalanced datasets: A comparative study of ML and DL algorithms with SMOTE preprocessing. *Mansoura Journal for Computer and Information Sciences*. https://doi.org/10.21608/mjcis.2025.313097.1007

[33] Tavallaee, M., Bagheri, E., Lu, W., & Ghorbani, A. A. (2009). A detailed analysis of the KDD CUP 99 dataset. In *Proceedings of the IEEE Symposium on Computational Intelligence for Security and Defense Applications*. https://doi.org/10.1109/CISDA.2009.5356528

[34] Usman, S., et al. (2021). Intelligent dynamic malware detection using machine learning in IP reputation for forensics data analytics. *Future Generation Computer Systems, 118*, 124–141.https://doi.org/10.1016/j.future.2021.01.004

[35] Lu, J., Xu, Q., & Hu, J. (2026). A novel graph learning framework for interpretable and imbalance financial fraud detection. Engineering Applications of Artificial Intelligence, 167, 113709. https://doi.org/10.1016/j.engappai.2025.113709