SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

A Robust Hybrid Temporal Deep Learning Framework with Attention Mechanism for Online Payment Fraud Detection

Ms. Reena More[1], Dr. Pankaj Dashore[2]

[1]Research Scholar,,School of Computer Sciences and Engineering, Sandip University Nashik (MH), India.
E-mail: reena.more3@gmail.com, Orcid: https://orcid.org/0009-0005-3852-949X
[2]Professor, School of Computer Sciences and Engineering, Sandip University Nashik (MH), India. E-mail: dashorepankaj@gmail.com

ABSTRACT

With increasing use of digital payment systems there has been a rise in frequency and complexity of fraudulent financial transactions. As a result there is a need for more precise and robust mechanisms of detecting this. In this study a hybrid deep learning approach is proposed using a temporal convolutional neural network, long short-term memory network and attention mechanism to detect fraud in online payments. The proposed model combines one-dimensional convolution for localized feature extraction, LSTM-based dependency modelling, and attention-based weighting to emphasize transaction representations that contribute most strongly to fraud classification. The dataset was preprocessed through selection of fraud-relevant transaction types, one-hot encoding of categorical variables, standardization of numerical attributes, and missing-value inspection. Class-weighted learning was employed to reduce the bias caused by the highly imbalanced distribution of legitimate and fraudulent transactions. The model was implemented using Python, TensorFlow, and Keras and evaluated using accuracy, precision, recall, F1-score, confusion matrix, and receiver operating characteristic analysis. Experimental results yielded an evaluation accuracy of 94.20% and an ROC-AUC of 0.9989. The model correctly identified 1,638 of 1,643 fraudulent transactions, achieving a fraud recall of 99.70% and a false-negative rate of 0.30%. However, fraud precision was limited to 4.85% because of a comparatively high number of false-positive predictions. These findings demonstrate strong fraud-detection sensitivity while highlighting the need for classification-threshold calibration and false-positive reduction before practical deployment.

Keywords: Online payment fraud, deep learning, TCNN, LSTM, attention mechanism, class imbalance.

1. Introduction

Rapid development of digital payment technology has driven significant change in the global financial ecosystem enabling more convenient, more accessible and faster electronic transactions. Widespread use of internet banking, mobile wallets, e-commerce and contactless payments has significantly increased the number of financial transactions taking place online every day. While through technological advances there is now greater convenience for consumers and more efficiency for institutions, this has increased the attack surface for cybercriminals. Increasingly sophisticated fraudulent activity including unauthorized transactions, identity theft, account takeover, synthetic identity fraud and payment manipulation now poses significant financial and reputational risk to consumers, merchants and financial institutions. As a result one of the key challenges in financial cybersecurity today is developing scalable, robust and accurate fraud detection systems [1], [2].

Traditionally fraud detection systems use manually written rules, statistical analysis and more recently conventional machine learning algorithms to pick out fraudulent transactions from legitimate ones. Rule-based systems are easy to implement and decision-making is interpretable. However they need regular updating to pick up new patterns of fraud. Often fraudsters will adapt their attack strategy to get around the rules meaning over time these systems become less effective. More recently there have been improvements in fraud detection using conventional machine learning e.g. decision trees, support vector machines, logistic regression and random forests. These learn patterns from historical transaction data. However often they rely on handcrafted features and assume the data distribution doesn't change over time. As a result they struggle to capture complex temporal dependencies and evolving behavior seen in real world online payment transactions [2], [8].

In recent years there have been considerable advances in artificial intelligence and in particular deep learning models. These models can automatically learn hierarchical feature representations from large transaction datasets without significant need for manual feature engineering. Through modeling nonlinear dependencies between attributes of transactions artificial neural networks have shown good performance over more traditional machine learning approaches at classifying fraud [22]. More recently there has been increasing interest in recurrent neural networks and in particular long short-term memory (LSTM) networks due to their ability to pick up long term temporal dependencies and sequential behavior in transactions. This is relevant in fraud detection as often it is through anomalous sequences of behavior malicious activity is seen rather than through individual transactions. For example hierarchical recurrent neural networks have been used to pick up structure in transactions and patterns of user behavior leading to better performance in fraud detection in

an online payment system [4].

Although using LSTM based architectures sequential dependencies can be modeled they may not efficiently capture short term local properties of transactions often seen prior to fraudulent activity. Recently temporal convolutional networks (TCNs) have shown promise as an alternative to modeling sequential data. They use causal convolutions and dilated convolution to pick up temporal features across a range of receptive fields. By processing sequences in parallel they have important computational advantages on more traditional recurrent architectures whilst preserving temporal causality. As a result by combining temporal convolution and recurrent learning both local dynamics of transactions and longer term dependencies in user behavior can be picked up. This gives a more complete picture of payment behavior.

Another important development in deep learning has been the introduction of attention mechanisms. These allow neural networks to dynamically place more weight on important time steps and informative features as part of their learning process. Instead of giving equal weight to all transactions attention allows the model to pick-up more relevant patterns of behavior more likely to indicate fraud. Using attention augmented LSTM architectures studies have seen better performance of credit card fraud models through improved discrimination of features and fewer errors of misclassification [5]. More recently in other work using attention based hybrid architectures researchers have seen similar improvements in online payment fraud detection and financial anomaly [6], [7].

Despite these advances there remain unresolved challenges in the research literature. One of the most important of these is the very skewed balance between legitimate and fraudulent transactions. Typically only a small proportion of records in the data will be fraudulent. This often leads to bias of deep learning models towards the more common legitimate transactions. As a result sensitivity of fraud detection is reduced. To address this researchers have explored e.g. cost-sensitive optimization, ensemble learning, unsupervised and reinforcement learning [8], [9], [13]. Another challenge is that fraud patterns change over time as customers' behavior evolves and more sophisticated cyberattacks emerge. This leads to concept drift over time reducing performance of models. Where models are deployed there are significant computational constraints. Financial institutions need to make decisions on potential fraud within milliseconds yet also want high accuracy of fraud detection and low false alarms [7], [10].

In recent literature there has also been increasing recognition of the value of hybrid deep learning approaches combining complementary learning mechanisms to improve generalization and robustness. Through combination of multiple predictive models ensemble learning has been shown to improve stability of classification [8], while hybrid optimization has used metaheuristic algorithms to aid parameter optimization and feature selection [18]. Recent studies have seen considerable potential of autoencoder approaches for learning latent representations of normal transaction behavior to pick up anomalous activity without the need for large amounts of labeled data [19]. In addition reinforcement learning has been used to build adaptive fraud detection systems able to continually learn from changing transaction environment [9], [20]. Often in existing work either only one deep learning architecture is used or they do not exploit within a single framework local temporal features, longer term sequential dependencies and adaptive weighting of features.

Motivated by these findings in this study a Robust Hybrid Temporal Deep Learning framework with attention mechanism for online payment fraud detection is proposed. This framework combines temporal convolutional neural network (TCNN), long short-term memory (LSTM) network and attention mechanism to capture short-term characteristics of transactions and longer term dependencies in user behavior. First using one dimensional convolution operations the TCNN learns discriminative local temporal features from sequences of transactions. These are then fed into the LSTM network to pick up longer term sequential patterns of legitimate and fraudulent payment behavior. Finally using an attention mechanism the most relevant temporal representations are highlighted prior to classification. In this way the model focuses on those aspects of transaction patterns most useful in identifying fraud. To deal with often severe class imbalance in fraud datasets the framework uses a class weighted approach during learning to boost the influence of minority fraudulent samples in the optimization.

The proposed framework is evaluated using a publicly available online payment fraud dataset under a comprehensive experimental setting. The model performance is assessed using widely accepted evaluation metrics, including accuracy, precision, recall, F1-score, receiver operating characteristic (ROC) analysis, area under the ROC curve (AUC), and confusion matrix analysis. In addition the study shows through extensive experimental analysis that combining temporal convolution, sequential learning and attention is effective for fraud detection at computational cost appropriate for use in real financial applications.

The main contributions of this study are summarized as follows:

1. A robust hybrid temporal deep learning framework integrating TCNN, LSTM, and an attention mechanism is proposed for online payment fraud detection.
2. A comprehensive preprocessing pipeline incorporating categorical encoding, feature normalization, missing-value inspection, and class-weighted learning is developed to improve model robustness under highly imbalanced transaction data.
3. The proposed architecture effectively combines local temporal feature extraction, long-range sequential

modelling, and adaptive feature weighting to enhance fraud detection performance.
4. Extensive experimental evaluation is conducted using multiple performance metrics, including accuracy, precision, recall, F1-score, ROC-AUC, and confusion matrix analysis, to validate the effectiveness of the proposed framework.
5. The study provides insights into the practical applicability of hybrid temporal deep learning architectures for scalable and financial fraud detection systems.

The rest of this paper is structured as follows. In section 2 the study considers existing deep learning and machine learning approaches for online payment fraud and highlights gaps in the research. In section 3 the study describes the proposed attention enhanced hybrid TCNN-LSTM framework. This includes theoretical foundations, mathematical modeling and algorithmic procedure. Section 4 outlines the dataset, preprocessing, experimental setup, training setup, metrics used to evaluate performance and analysis of results. In section 5 the study compares the framework against key fraud detection approaches. In section 6 the study discusses key findings, trade-offs in performance and implications. In section 7 the study concludes the study and suggests directions for future research.

2. Related Work

Increasing sophistication of online payment fraud has led to considerable research into intelligent fraud detection systems able to pick up high accuracy, low latency anomalous patterns in transactions. Initially much of this work used more traditional machine learning algorithms but more recently there has been a move towards deep learning and hybrid architectures automatically picking up complex nonlinear and temporal relationships in transaction data. Despite advances there remain limitations of many fraud detection systems due for example to class imbalance, changing strategies of fraudsters, computational efficiency and interpretability of models. In this section the recent advances in machine learning and deep learning for online payment fraud are critically examined and gaps in the literature addressed by the proposed framework are highlighted.

2.1 Machine Learning-Based Fraud Detection

Due to relatively low computational complexity and ease of implementation traditional machine learning approaches have supported automated fraud detection. Often applied to distinguish legitimate from fraudulent financial transactions are classifiers e.g. logistic regression, decision trees, support vector machines, random forests, gradient boosting and xgboost using handcrafted features of each transaction. In practice they typically require careful parameter tuning and significant feature engineering to deliver good performance.
Bahnsen et al [2] showed good feature engineering can have a greater impact on performance of fraud detection systems than simply swapping one classifier for another. They found transaction specific behavioral features significantly improved accuracy of identifying fraud with reduced false alarms. More recently Hajek et al [10] proposed an xgboost framework for detecting fraud in mobile payments. Through use of ensemble learning and transactional attributes they achieved competitive performance in classifying fraudulent from legitimate transactions. Khalid et al [8] built on this using an ensemble of machine learning classifiers to improve robustness of fraud detection against very imbalanced data across transactions.
In recent years there has also been academic focus on applying machine learning methods to online payment platforms. Lokanan [24] used a range of machine learning algorithms to predict fraudulent mobile money transactions. As part of this he extended fraud analytics to emerging digital financial services. In doing so he shows supervised learning could be applied across different electronic payment settings. Almazroi and Ayub [6] looked at a range of machine learning algorithms for detecting fraud in online payments. They highlighted the importance of data preprocessing, feature selection and training models on a balanced data set. Farouk et al [25] used machine learning on transactional financial data to detect fraud in online payments. They too show there remains scope for using supervised machine learning to pick up differences between fraudulent and legitimate payment activity. In one study Borketey [14] developed a framework showing in practice it could be feasible to apply machine learning models in financial systems. Aslam and Hussain [15] compared a range of machine learning algorithms and found generally ensemble approaches outperformed individual classifiers in terms of stability and accuracy of detection.
Taken together these papers show in many fraud detection applications traditional machine learning approaches remain relevant. However, they tend to rely on manually developed features and struggle to pick up sequential behavioral patterns across a series of financial transactions. As a result, they may struggle if over time transactional characteristics change or fraud patterns develop.

2.2 Artificial Neural Network and Deep Learning Approaches

Deep learning has proved a powerful alternative to more traditional machine learning approaches through automatic learning of hierarchical feature representations from raw transaction data. ANNs in particular have shown greater ability than more traditional machine learning approaches to capture nonlinear relationships between attributes of transactions without the need for extensive feature engineering.

R. B. and S. K. K. R [22] developed an ANN fraud detection model which outperformed a number of more traditional machine learning algorithms in terms of classification performance. More recently Al Balawi and Aljohani [17] developed a neural network credit card fraud detection system which could pick up complex patterns of fraud from transaction history. In both cases they found through use of neural networks much better accuracy of fraud detection through learning complex nonlinear decision boundaries. More recently Sanober et al [23] considered using secure deep learning for fraud detection in a wireless communication environment. They show more generally that complex fraudulent/malicious patterns could be picked up using deep neural architectures.

A limitation of feedforward neural networks is they consider each transaction in isolation and so cannot pick up temporal relationships between transactions. Often it is more the sequence of transactions which reveals fraudulent behavior rather than any individual transaction. As a result more recently, researchers have started to apply sequential deep learning architectures.

2.3 Sequential Deep Learning Models

Sequential modeling has become a key research direction in payment fraud detection; transaction history is itself a temporal sequence. RNNs, LSTMs and GRUs have been widely used to model dependencies between successive transactions.

Lin et al [4] developed a structure aware hierarchical recurrent neural network which uses hierarchical sequential learning to model both transaction relationships and user behavior. They found significant improvement in fraud detection on online payments through capture of both local dependencies between transactions and higher-level structure in user behavior.

Benchaji et al [5] developed an attention augmented LSTM to improve classification of fraud. Through combination of sequential learning and adaptive weighting of features they found LSTM outperformed more traditional feedforward neural networks. They suggest this is because through their gated memory mechanisms they better retain dependencies between transactions over longer periods.

While LSTM models do pick-up long-range temporal patterns, they tend to be more computationally complex to train than other architectures. They may also fail to pick up more local shifts in behavior over short windows of transactions. For this reason, some researchers have explored the use of convolutional temporal models to more effectively pick up local sequential patterns.

2.4 Hybrid Deep Learning Models

In recent years more attention has focused on hybrid deep learning architectures which incorporate complementary learning mechanisms to increase accuracy and robustness of fraud detection. By stacking different models of neural networks the strengths of different feature extraction approaches can be exploited while limitations of individual models are addressed.

Sharma et al [19] used deep autoencoders to pick up latent patterns of legitimate transaction behavior to inform anomaly detection. They were able to successfully reconstruct normal transactions and through analysis of reconstruction error picked up abnormal payment activity. Vimal et al [20] explored use of deep reinforcement learning for detection of payment fraud. They found through use of adaptive learning strategies fraud detection could continually improve in a changing financial environment. Tang [21] went on to explore automatic fraud detection of e-commerce transactions using a combination of deep reinforcement learning and artificial neural networks. Taken together these papers illustrate potential for adaptive learning mechanisms in dynamic environments for fraud detection. More recently Qayoom et al [9] presented a reinforcement learning framework able to adapt to changing fraud patterns through ongoing optimization of policy.

There has also been considerable scholarly interest in optimization based hybrid systems. Bhat and Nelson [18] combined a sparrows search algorithm with models of artificial intelligence to improve performance of fraud detection through better parameterization and feature selection. In another study, Abbassi et al. [13] developed an unsupervised learning fusion model for detecting fraud in online banking. They found through experimentation multiple unsupervised learning techniques together improved anomaly detection without the need for large amounts of labeled fraud data.

Taken together these papers suggest in general better generalization can be achieved through combination of different learning paradigms than using just one neural architecture. However, there are not many studies which have brought together temporal convolution, sequential memory learning and attention in a single framework for use in detecting fraud in online payments.

2.5 Recent Trends and Comprehensive Reviews

In a number of recent review studies advances in fraud detection using machine learning and deep learning have been reported. Bin Sulaiman et al [12] reviewed machine learning approaches for credit card fraud. They highlighted strengths and limitations of supervised, unsupervised and hybrid learning. Marazqah Btoush et al [11] reviewed machine learning and deep learning for cyber fraud. They found across all fraud types deep learning models outperformed more traditional algorithms where large datasets of transactions were

available.
Hove et al [16] used bibliometric analysis to show over the past decade there has been considerable growth in research into intelligent fraud detection. As part of this they found increasing use of deep learning, hybrid architectures, explainable artificial intelligence and fraud analytics. Building on this Vanini et al [7] argue more modern fraud detection systems should move beyond anomaly detection to include supported by comprehensive financial risk management. This would support adaptive decision-making and ongoing monitoring.
Taken together these studies show increasingly within the research community there is a move towards intelligent, adaptive and hybrid fraud detection systems suited to more complex financial transaction environments.

2.6 Research Gap and Motivation

In the literature the study sees increasing use of machine learning and deep learning for fraud detection. But key limitations of this approach remain unresolved.
First, many existing machine learning methods depend heavily on handcrafted feature engineering and therefore exhibit limited adaptability to rapidly evolving fraud behaviours [2], [8]. Second, although sequential models such as LSTM effectively capture long-term behavioural dependencies, they often fail to efficiently extract localized temporal transaction patterns that precede fraudulent activities [4], [5]. Third, numerous hybrid models focus on combining optimization algorithms, ensemble learning, or unsupervised learning but do not simultaneously integrate temporal convolution, sequential memory learning, and adaptive attention within a unified architecture [9], [13], [18], [19]. Fourth, severe class imbalance remains a significant challenge because fraudulent transactions constitute only a very small proportion of real-world financial datasets, causing many models to become biased toward legitimate transactions [8], [11]. Finally, several existing studies primarily optimize overall classification accuracy while providing limited discussion of deployment requirements, computational efficiency, and false-positive reduction [7], [10], [14].
To address these limitations, this work proposes a Robust Hybrid Temporal Deep Learning Framework with Attention Mechanism that combines the complementary strengths of a Temporal Convolutional Neural Network for local temporal feature extraction, an LSTM network for long-term sequential modelling, and an attention mechanism for adaptive feature weighting. Additionally, class-weighted learning is incorporated to mitigate the adverse effects of class imbalance. The proposed framework is designed to improve fraud detection accuracy while maintaining computational efficiency suitable for online payment systems.

3. Proposed Hybrid Temporal Deep Learning Framework

3.1 Overview of the Proposed Framework

Online payment fraud detection requires simultaneously capturing short-term transaction variations, long-term behavioural dependencies, and transaction-specific importance. Conventional machine learning models rely primarily on handcrafted features and often struggle to model the temporal evolution of financial transactions. Although recurrent neural networks effectively learn sequential dependencies, they are less efficient in extracting localized temporal patterns. Similarly, while convolutional architectures well capture local features, they cannot independently retain long range temporal information. In this study the authors propose a Robust Hybrid Temporal Deep Learning Framework which combines a temporal convolutional neural network, long short-term memory network and attention mechanism into one architecture for detection of fraud in online payments.
As illustrated in Figure 1, the proposed framework consists of six sequential stages: (i) transaction data collection, (ii) data preprocessing, (iii) temporal feature extraction using the TCNN, (iv) sequential dependency modelling through the LSTM network, (v) adaptive feature refinement using the attention mechanism, and (vi) binary fraud classification. This hierarchical learning strategy enables the framework to simultaneously capture local temporal variations and long-range behavioural relationships while dynamically emphasizing the most informative transaction patterns.

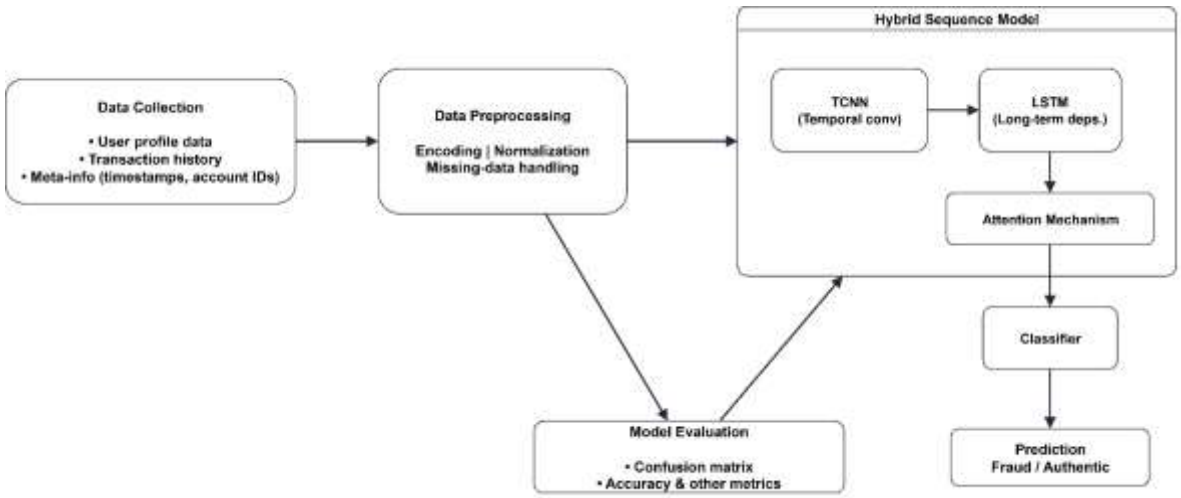

Figure 1. Overall architecture of the proposed hybrid TCNN–LSTM framework with attention mechanism for online payment fraud detection

3.2 Data Collection and Preprocessing

The proposed framework uses the publicly available dataset of online payment fraud. This includes both legitimate and fraudulent financial transactions from an online payment system. Each transaction has a range of descriptive attributes e.g. transaction type, amount, balance of sender, balance of receiver and fraud label.
Before training the model a range of preprocessing steps are applied to the raw dataset to improve data quality and make learning more efficient. Irrelevant identifiers of transactions are dropped as they add no useful information for predicting fraud. Categorical attributes of transactions are converted into numerical form using one-hot encoding. Numerical variables are standardized using the technique of StandardScaler to remove variation of scale between features. Missing values and inconsistencies in records are then checked to ensure data quality before moving on to model training.
After preprocessing, each transaction is represented as a numerical feature vector

$$x_i = [x_{i1}, x_{i2}, \dots, x_{id}]^T$$

where d denotes the total number of input features and i represents the corresponding transaction instance. The complete preprocessed dataset can therefore be expressed as

$$D = \{(x_i, y_i)\}_{i=1}^{N}$$

where N denotes the total number of transactions and

$$y_i \in \{0,1\}$$

indicates whether the transaction is legitimate (0) or fraudulent (1)

3.3 Class Imbalance Handling

One of the most significant challenges in fraud detection is the extreme imbalance between legitimate and fraudulent transactions. Since fraudulent transactions represent only a very small percentage of the dataset, conventional optimization methods tend to become biased toward the majority class.
Rather than generating synthetic samples, the proposed framework employs class-weighted learning, where higher penalties are assigned to fraudulent samples during optimization. The class weight corresponding to class c is computed as

$$w_c = \frac{N}{(K\ n_c)}$$

where
N is the total number of training samples,
K represents the number of classes,
n_c denotes the number of samples belonging to class c.

These weights are incorporated into the loss function to increase the contribution of minority fraudulent transactions during network optimization.

3.4 Temporal Sequence Construction

Fraudulent behaviour rarely appears as an isolated event. Instead, abnormal activities generally emerge through sequential changes in transaction behaviour over time. Therefore, the preprocessed transaction vectors are organized into temporal sequences before being supplied to the deep learning architecture. Each sequence is represented as

$$X_i = [x_{i-T+1}, x_{i-T+2}, \dots, x_i]$$

where T denotes the sequence length.

This sequential representation enables the network to learn both immediate transaction characteristics and long-term behavioural evolution.

3.5 Temporal Convolutional Neural Network

The first learning stage employs a Temporal Convolutional Neural Network (TCNN) to extract localized temporal features from sequential transaction data. Unlike conventional convolutional neural networks developed for image analysis, the TCNN performs one-dimensional causal convolutions that preserve the chronological ordering of financial transactions.

The convolution operation is defined as

$$z_t = ReLU\left(\sum_{k=0}^{K_c-1} W_k x_{t-k} + b_c\right)$$

where

K is the convolution kernel size,

W_k denotes the convolution weights,

b represents the bias vector,

σ(·) denotes the ReLU activation function.

The resulting temporal feature representation is expressed as

$$Z = [z_1, z_2, \dots, z_T]$$

The TCNN effectively captures short-term behavioural variations including sudden transaction amount changes, balance fluctuations, and abnormal transaction frequencies.

3.6 Long Short-Term Memory Network

While the TCNN is good at capturing localized temporal information it isn't well suited to capturing longer term dependencies in behavior. As a result, the temporal features from the TCNN are passed into an LSTM network. This captures sequential dependencies across longer histories of transactions.

The forget gate is computed as

$$f_t = \sigma\left(W_f z_t + U_f h_{t-1} + b_f\right)$$

The input gate is

$$i_t = \sigma(W_i z_t + U_i h_{t-1} + b_i)$$

The candidate memory is

$$C_t' = \tanh(W_c z_t + U_c h_{t-1} + b_c)$$

The memory state is updated according to

$$C_t = f_t \odot C_{t-1} + i_t \odot C_t'$$
$$o_t = \sigma(W_o z_t + U_o h_{t-1} + b_o)$$

Finally, the hidden state becomes

$$h_t = o_t \odot \tanh(C_t)$$

where o_t denotes the output gate.

These gated operations enable the network to preserve long-term transaction dependencies while selectively discarding irrelevant historical information.

3.7 Attention Mechanism

Not all transactions will contribute equally to predicting fraud. To reflect this the study introduces an attention mechanism which dynamically highlights the most informative temporal representations from the LSTM network. In their framework of attention Vaswani et al. [3] showed through experiment the value of adaptively weighting to highlight informative representations. Building on this the proposed framework applies different importance weights to the representations generated by the LSTM before classification.

The attention score is calculated as

$$e_t = v_a^T \tanh(W_a h_t + b_a)$$

The normalized attention weight is

$$\alpha_t = \frac{\exp(e_t)}{\sum_{j=1}^{T} \exp(e_j)}$$

The context vector is obtained by

$$c = \sum_{t=1}^{T} \alpha_t h_t$$

The resulting context vector summarizes the most discriminative behavioural information required for fraud classification.

3.8 Fraud Classification and Model Optimization

The attention-weighted feature vector is supplied to a fully connected layer followed by a sigmoid activation function to estimate the fraud probability

$$\hat{y}_i = \sigma(W_p c_i + b_p)$$

where $\hat{y}$ represents the predicted fraud probability.

The network is optimized using the weighted binary cross-entropy loss

$$L = -\left(\frac{1}{N}\right) \sum_{i=1}^{N} [w_1 y_i \ln(\hat{y}_i) + w_{0(1-y_i)\ln}(1 - \hat{y}_i)]$$

where w_i denotes the class weight assigned to each transaction.

The Adam optimizer is employed to minimize the objective function through iterative backpropagation, while EarlyStopping and adaptive learning-rate scheduling improve convergence and reduce overfitting.

3.9 Algorithm of the Proposed Framework

Finally in algorithm 1 the full procedure of the proposed hybrid TCNN-LSTM framework with attention mechanism is summarized. Here the feature extraction, temporal learning, optimization and fraud classification steps are outlined in order.

Algorithm 1
Input:
Transaction dataset D; sequence length T; number of epochs E; batch size B; learning rate η; classification threshold τ
Output:
Trained model M; fraud probabilities $\hat{p}$; predicted labels $\hat{y}$

1: Load the online payment transaction dataset D
2: Inspect D for missing values, duplicate records, and irrelevant attributes
3: Remove irrelevant transaction identifiers and duplicate records
4: Encode categorical transaction attributes using one-hot encoding
5: Standardize numerical attributes:
$x' = \frac{x-\mu}{\sigma}$
6: Separate the preprocessed data into feature matrix X and target vector y
7: Partition X and y into training, validation, and test subsets
using stratified sampling
8: Compute the class weight for each class c:
$w_c = \frac{N}{K\, n_c}$
9: Construct the model input sequences:
$X_{seq} = \{X_1, X_2, \ldots, X_m\}$
where each X_i contains T ordered transaction representations
10: Initialize model M with:
Conv1D layer: 64 filters, kernel size = 3,
causal padding, ReLU activation
Dropout layer: rate = 0.3
LSTM layer: 64 hidden units
Attention layer: Dense + Softmax

GlobalAveragePooling1D layer
Dense output layer with sigmoid activation
Initialize Adam optimizer with learning rate η
for epoch = 1 to E do
  Divide the training data into mini-batches of size B
  for each mini-batch (X_b, y_b) do
    Extract localized representations using TCNN:
    $Z_b = ReLU\big(Conv1D(X_b)\big)$
    Apply dropout to Z_b
    Pass Z_b through the LSTM layer to obtain hidden states:
    $H_b = \{h_1, h_2, \dots, h_T\}$
    Compute attention scores for the LSTM representations:
    $e_t = v_a^T \tanh(W_a h_t + b_a)$
    Normalize the attention scores:
    $\alpha_t = \frac{\exp(e_t)}{\sum_{j=1}^{T} \exp(e_j)}$
    Obtain the attention-weighted representation:
    $c = \sum_{t=1}^{T} \alpha_t h_t$
    Apply global pooling and the final classification layer
    Estimate the fraud probability:
    $\hat{p} = Sigmoid\big(W_p c + b_p\big)$
    Compute the class-weighted binary cross-entropy loss:
    $L = -\left(\frac{1}{B}\right) \sum_{i=1}^{B} \big[w_1 y_i \ln(\hat{p}_i) + w_{0(1-y_i)\ln}(1 - \hat{p}_i)\big]$
    Compute gradients of L through backpropagation
    Update all trainable parameters using Adam
  end for
  Evaluate M on the validation subset
  Record validation loss, accuracy, and AUC
end for
Select the final trained model M
for each test sample X_i do
  Compute fraud probability:
  $\hat{p}_i$ = M(X_i)
  if $\hat{p}_i$ >= τ then
    $\hat{y}_i$ = 1 ▷ Fraud
  else
    $\hat{y}_i$ = 0 ▷ Legitimate
  end if
end for
Compute Accuracy, Precision, Recall, F1-score,
ROC-AUC, FPR, FNR, and Confusion Matrix
Return M, $\hat{p}$, $\hat{y}$ and the evaluation results

4. Experimental Setup and Result Analysis

This section presents the dataset, the implementation environment, model setup, criteria used to evaluate the model and findings from experiments using the proposed hybrid TCNN-LSTM with attention mechanism. The analysis is of outputs from running the model; this includes training and validation curves, confusion matrix, class level performance, and receiver operating characteristic curve.

4.1 Dataset Description

The proposed framework was evaluated using the publicly available Online Payment Fraud Detection Dataset obtained from Kaggle. This dataset consists of simulated online financial transactions labeled as either legitimate or fraudulent. For each record there is transaction time step, category, amount transferred, account details of sender and receiver, account balances before and after the transaction and the fraud label isFraud. Where isFraud=0 the transaction is legitimate, if isFraud=1 then it is fraudulent. The principal dataset attributes are summarized in Table 1.

Table 1. Principal attributes of the online payment fraud detection dataset

Attribute	Description
step	Time step representing the number of hours from the beginning of the simulation
type	Type of online transaction, such as CASH_OUT, TRANSFER, or PAYMENT

amount	Monetary value of the transaction
nameOrig	Anonymized identifier of the originating account
oldbalanceOrg	Sender's account balance before the transaction
newbalanceOrig	Sender's account balance after the transaction
nameDest	Anonymized identifier of the destination account
oldbalanceDest	Receiver's account balance before the transaction
newbalanceDest	Receiver's account balance after the transaction
isFraud	Binary fraud label: 0 for legitimate and 1 for fraudulent

The dataset is highly imbalanced because fraudulent transactions constitute only a small proportion of the complete transaction population. Such imbalance reflects the characteristics of real payment systems but can cause a learning algorithm to become biased toward the majority legitimate class.

4.2 Data Preprocessing

Several preprocessing operations were performed before model training. Only transactions belonging to the TRANSFER and CASH_OUT categories were retained because these transaction types were most closely associated with fraudulent activities in the selected dataset. The categorical type variable was transformed using one-hot encoding through the pd.get_dummies() function. The drop_first=True option was applied to remove one redundant dummy variable.

Continuous attributes, including step, amount, oldbalanceOrg, newbalanceOrig, oldbalanceDest, and newbalanceDest, were standardized using StandardScaler. For a feature x, the standardized value was calculated as

$$x' = \sigma x - \mu,$$

where μ and σ represent the feature mean and standard deviation, respectively. Standardization placed all continuous variables on comparable scales and supported stable gradient-based optimization. Missing values were examined using the isnull().sum() operation before model training.

The resulting input was reshaped to the form

$$(sequence\ length, number\ of\ features),$$

which provided the three-dimensional structure required by the Conv1D and LSTM layers. Therefore, in the implemented configuration, the layers operated over the ordered representation of transaction features.

4.3 Experimental Environment

The proposed model was implemented in Python 3.10 using TensorFlow and Keras. Pandas was used for data manipulation, whereas Matplotlib supported the visualization of model performance. The experiments were executed under Windows 11 using an Intel Core i7 processor, an NVIDIA GeForce RTX 3050 GPU, and 64 GB of RAM. The implementation environment is presented in Table 2.

Table 2. Hardware and software environment used for model implementation

Component	Specification
Operating system	Windows 11
Programming language	Python 3.10
Deep learning libraries	TensorFlow and Keras
Data-processing library	Pandas
Visualization library	Matplotlib
Processor	Intel Core i7 at 3.10 GHz
Graphics processing unit	NVIDIA GeForce RTX 3050
System memory	64 GB RAM

4.4 Model Configuration and Training Parameters

The hybrid architecture comprised a one-dimensional temporal convolutional layer, an LSTM layer, a custom Dense–Softmax attention mechanism, global average pooling, and a sigmoid classification layer. The Conv1D component contained 64 filters with a kernel size of 3, causal padding, and ReLU activation. The LSTM layer contained 64 hidden units. A dropout rate of 0.3 was applied to reduce overfitting.

Class weights were computed and supplied during training to increase the optimization penalty assigned to the minority fraud class. The network was trained for 10 epochs using the Adam optimizer, a learning rate of 0.001, a batch size of 1024, and binary cross-entropy loss. Accuracy and AUC were monitored during training.

The complete hyperparameter configuration is presented in Table 3.

Table 3. Training parameters of the proposed hybrid TCNN–LSTM–Attention framework

Hyperparameter	Selected value
Conv1D filters	64

Conv1D kernel size	3
Conv1D padding	Causal
Conv1D activation	ReLU
LSTM units	64
Attention mechanism	Custom Dense + Softmax
Pooling layer	GlobalAveragePooling1D
Dropout rate	0.3
Output activation	Sigmoid
Optimizer	Adam
Learning rate	0.001
Loss function	Binary cross-entropy
Training metrics	Accuracy and AUC
Maximum epochs	10
Batch size	1024
Class-imbalance treatment	Computed class weights
Validation data	(X_{test}_seq, y_{test})

The implementation used (X_{test}_seq, y_{test}) as validation data after each epoch. Consequently, this subset is more appropriately described as an evaluation or validation subset, rather than a completely unseen final test set.

4.5 Training and Validation Performance

Given the marked class imbalance in the dataset accuracy alone is not sufficient as a measure of performance for fraud detection. As a result precision, recall, F1-score, false positive rate, false negative rate, confusion matrix and ROC-AUC were examined. Using the ROC curve the relationship between how for different classification thresholds the true positive rate relates to false positive rate was explored. The area under this curve was used to assess discriminative ability of the model.

The training and validation accuracy and loss curves obtained over 10 epochs are presented in Figure 2.

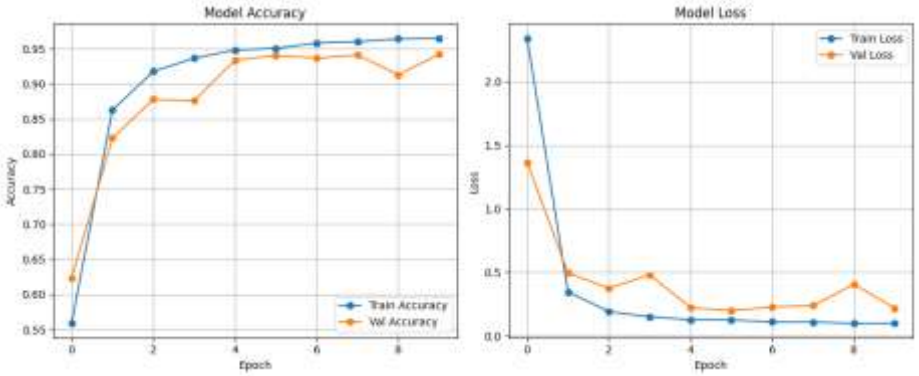

Figure 2. Training and validation accuracy and loss of the proposed hybrid TCNN–LSTM–Attention model over 10 epochs

Over the first epoch training accuracy moved quickly from 55.8% to over 86% in the second epoch. It then steadily improved to 96.4% at the final epoch. Over the first two epochs validation accuracy moved from 62.2% to over 82% and by the final epoch was 94.2%.

The fairly small difference between final accuracies on training and validation suggests in the process the model developed meaningful representations of transactions. However there is some slight instability of generalization performance across the validation accuracy through the epochs particularly around epochs 3 and 8. The dip at epoch 8 shows additional training didn't reliably improve classification on validation.

The training loss decreased sharply from approximately 2.33 to 0.34 during the first two epochs and continued declining to approximately 0.09 by the final epoch. Validation loss also decreased substantially from approximately 1.36 to approximately 0.20. However, validation loss fluctuated during training, with noticeable increases around epochs 3 and 8. These fluctuations suggest that the model began to exhibit mild overfitting after the initial convergence period.

Overall, the accuracy and loss curves demonstrate successful optimization. However, checkpoint selection based on minimum validation loss would likely provide a more stable final model than simply retaining the parameters from the final epoch.

4.6 Confusion Matrix Analysis

The confusion matrix obtained from the evaluation dataset is shown in Figure 3.

Figure 3. Confusion matrix of the proposed hybrid model for legitimate and fraudulent transaction classification

The model correctly classified 520,334 legitimate transactions and 1,638 fraudulent transactions. Only five fraud cases were incorrectly classified as legitimate, demonstrating that the framework achieved very high sensitivity to fraudulent activity. However, 32,105 legitimate transactions were incorrectly classified as fraudulent. The resulting overall evaluation accuracy was 94.20%. The fraud-class false-negative rate was only 0.30%. Conversely, the false-positive rate was 5.81%.

These findings show that the framework was configured to prioritize the identification of fraudulent transactions. This behaviour is valuable when the cost of missing a fraudulent transaction is considerably higher than the cost of reviewing a legitimate transaction. Nevertheless, the large absolute number of false alerts may impose substantial manual verification costs in a practical payment-processing environment.

4.7 Class-Wise Precision, Recall, and F1-Score

The class-wise precision, recall, and F1-score obtained from the execution are illustrated in Figure 4.

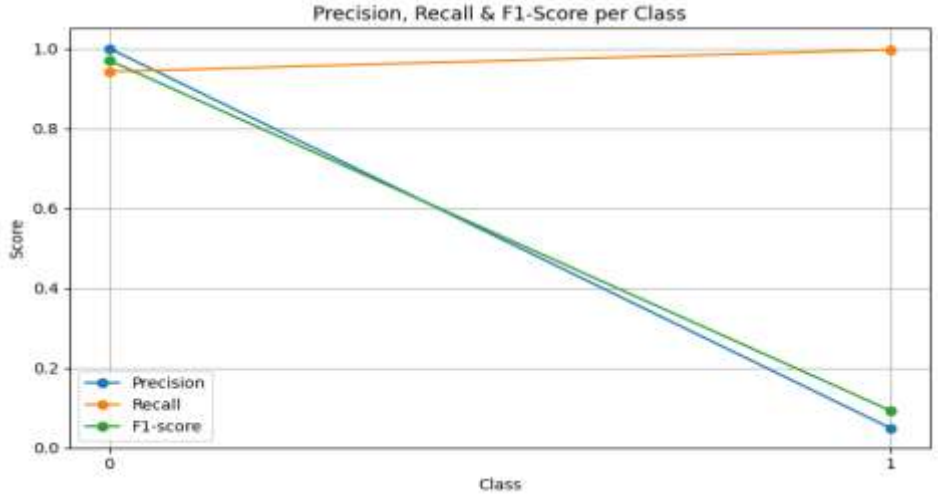

Figure 4. Class-wise precision, recall, and F1-score of the proposed fraud detection model

The detailed values derived from the confusion matrix are reported in Table 4.

Table 4. Class-wise performance of the proposed framework

Performance measure	Legitimate class	Fraud class
Precision	100.00%	4.85%
Recall	94.19%	99.70%
F1-score	97.01%	9.26%
Number of samples	552,439	1,643

For the legitimate class, the model achieved approximately 100% precision, 94.19% recall, and a 97.01% F1-score. These values indicate that nearly every transaction classified as legitimate was genuinely legitimate.

However, because 32,105 legitimate transactions were assigned to the fraud class, legitimate-class recall was lower than its precision.

For the fraud class, the model obtained 99.70% recall, showing that it detected almost every fraudulent transaction. In contrast, fraud precision was only 4.85%, meaning that a large proportion of the transactions flagged as fraudulent were actually legitimate. Consequently, the fraud-class F1-score was limited to 9.26%.

The difference between fraud recall and fraud precision reflects the use of class weighted training and choice of classification threshold. In order to avoid missing fraud cases the model penalized this heavily. As a result the model ended up with a very sensitive decision boundary. This reduced the risk of false negatives but increased the risk of false positives.

4.8 ROC and AUC Analysis

The ROC curve obtained from the predicted probabilities is presented in Figure 5.

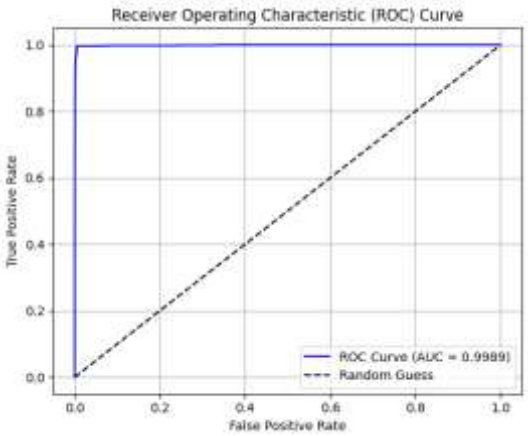

Figure 5. Receiver operating characteristic curve of the proposed hybrid TCNN–LSTM–Attention framework

The proposed model achieved an ROC-AUC value of 0.9989. The curve remained close to the upper-left corner, indicating that the predicted probabilities provided excellent discrimination between fraudulent and legitimate transactions across the complete range of classification thresholds.

The high ROC-AUC sits comfortably with the low precision of fraud. ROC-AUC reflects the ability of the model to rank the two classes across all possible decision thresholds. Fraud precision is calculated at a specific operating threshold. So it does appear the model could pick up difference between the two classes at probability ranking level but at the current threshold there are too many false positive decisions.

This suggests the decision threshold should be calibrated on the validation data. Raising the threshold might reduce number of legitimate transactions incorrectly classified as fraud but could increase number of missed fraud cases. The choice of operating threshold should depend on relative financial cost of false positive vs false negative predictions.

4.9 Overall Result Analysis

The experimental findings demonstrate that the hybrid TCNN–LSTM–Attention architecture successfully learned discriminative patterns from the online payment transaction data. The model achieved a training accuracy of approximately 96.4%, an evaluation accuracy of 94.20%, and an ROC-AUC of 0.9989.

The primary strength of the model is its ability to detect fraudulent transactions. Of the 1,643 fraud cases in the evaluation dataset, 1,638 were correctly identified, resulting in a fraud recall of 99.70%. Only five fraudulent transactions remained undetected.

However, this high sensitivity was achieved at the expense of fraud precision. The model incorrectly flagged 32,105 legitimate transactions, reducing fraud precision to 4.85%. Therefore, the model is more suitable as a high-sensitivity screening mechanism than as an autonomous final decision system under its current classification threshold. Flagged transactions would require an additional verification stage, secondary classifier, or manual review process.

The ROC-AUC result suggests that the underlying probability estimates contain strong discriminatory information. Consequently, the principal limitation may not be the learned representation itself, but rather the selected classification threshold and the optimization emphasis created by class weighting. Threshold calibration, cost-sensitive decision analysis, probability calibration, and precision–recall evaluation should therefore be considered before deployment.

Furthermore, although the framework is intended for fraud detection, execution time, inference latency, transaction throughput, model size, and memory requirements were not recorded in the available experiment. The predictive results demonstrate the classification capability of the model, but applicability should be

quantitatively validated through computational-performance measurements in future experiments.
Overall, the results indicate that the proposed framework provides excellent fraud coverage and class-ranking ability, but further refinement is required to reduce false alarms and improve fraud-class precision without substantially lowering fraud recall.

5. Comparative Analysis

In this section the proposed hybrid TCNN-LSTM-attention framework is considered against a range of machine learning, sequential deep learning, attention, ensemble, anomaly detection and adaptive fraud detection approaches reported in the literature. The class imbalance handling, feature extraction strategy, temporal modeling capability, learning mechanism and limitations of each approach are compared. It would be misleading to directly compare numerical results from published studies as often they are generated on different datasets, using different preprocessing, different class distributions, different train/test splits, different classification thresholds and different evaluation metrics. Instead here the focus is on differences between approaches and the performance of the framework is reported on the dataset used in this study.

5.1 Comparison with Traditional Machine Learning Approaches

Traditionally fraud detection will use classifiers e.g. logistic regression, decision trees, support vector machines, random forests and gradient boosting algorithms. Where input features are carefully engineered these can deliver good performance. Bahnsen et al [2] showed through transaction specific feature engineering the study could considerably improve performance of fraud detection. More recently XGboost and other ensemble methods have been used to pick up non-linear dependencies between attributes of transactions [8], [10].
A key strength of more traditional machine learning approaches is they tend to have low computational complexity. Tree methods for example could also produce feature importance which could be used to inform interpretation. However they rely on good quality manually engineered features. They tend to treat each transaction as an independent observation so may miss dependencies in an ordered representation of transactions.
In contrast, using the proposed framework features are automatically learned using causal conv1d filters, memory operations in the LSTM and weighting using attention. Localized patterns are captured by the TCNN part, dependencies and importance of different elements in the representation are picked up by the LSTM and attention modules. As a result the architecture relies less on handcrafted indicators of fraud. This comes at the cost of more complex and less interpretable model compared with e.g. logistic regression or decision tree based systems.

5.2 Comparison with Standalone Neural Networks

Feedforward artificial neural networks are commonly used for fraud classification as they are able to pick up non-linear relationships between transaction variables [17], [22]. Typically, they have more representation capacity than linear classifiers but do not include explicit sequential memory of inputs. The proposed architecture builds on the standard neural network approach by adding components of TCNN, LSTM and attention. Rather than feeding the preprocessed features directly into dense layers the model first uses the temporal convolution to pick up local interactions between the features. This representation is then fed into the LSTM layer. Finally, attention is used to weight informative positions. By adding this hierarchy, the model obtains more representational capacity than a simple feedforward neural network. However, this comes at the cost of more trainable parameters and longer training time. Where improved sensitivity of fraud justifies this cost over a feedforward network the model should be preferred.

5.3 Comparison with LSTM-Based Fraud Detection

LSTM networks are well suited to modeling sequential dependencies through their gated memory mechanism which retains useful information and dampens less useful signals from the past. Lin et al [4] used a hierarchical architecture of recurrent networks to capture behavioral and structural links in online payment data. Benchaji et al [5] added an attention mechanism on top of LSTM to help pick up more informative patterns in transactions. In isolation LSTM models will pick up long range dependencies but may struggle to pick up short range local patterns. Also, they perform their recurrent calculations one after another increasing complexity of training.
This limitation is addressed in the proposed model by placing the TCNN before the LSTM layer. Through the causal operation of Conv1D local interactions between features are captured before they are fed into the LSTM. Rather than the original transaction variables the LSTM therefore acts on a more refined representation. In this way within a single architecture the framework is able to incorporate both local pattern extraction and modeling of dependencies. Compared with a more standard TCNN-LSTM architecture the model includes an attention mechanism which places more weight on those representations in the LSTM most important for fraud classification.

5.4 Comparison with Attention-Based Architectures

Attention enhanced fraud detection systems improve over non-attention approaches through preventing the model treating all of the hidden representations as equally useful. Benchaji et al. [5] showed benefit of combining attention and recurrent learning for fraud detection. With attention the model is able to focus on suspicious patterns and reduce influence of less relevant data.

The proposed model follows this principle using a bespoke attention mechanism between dense and softmax layers. The attention weights adjust the output from the LSTM layer before global aggregation and binary classification. By forcing the classifier to focus on representations of minority class behavior this contributes to high recall of fraud in the model.

However, attention weights should not automatically be assumed to represent detailed causal explanation of why a transaction is fraudulent. They reflect weighting of the input to the model but do not in themselves show why a transaction is fraudulent. To justify fraud at transaction level further methods of explainability would be needed.

5.5 Comparison with Ensemble and Optimization-Based Methods

Ensemble approaches use a number of classifiers to improve robustness and reduce instability of individual models. Khalid et al. [8] showed value in using a combination of machine learning classifiers for credit card fraud. Optimization approaches have subsequently been used to find appropriate model and feature subsets. For example, Bhat and Nelson [18] used an optimization enabled artificial intelligence approach for fraud detection.

In comparison to ensemble systems, the proposed framework includes complementary learning operations in one end-to-end neural architecture. It doesn't require predictions from a range of classifiers each trained separately. As a result, the framework has a simpler inference pipeline compared with large heterogeneous ensembles.

Unlike the proposed framework, ensemble models could be more robust across different datasets as errors in individual models could partly cancel out when aggregated. The architecture uses a single trained network so could be more at risk of concept drift or distributional change unless it is periodically retrained.

Nor in the study is an external metaheuristic optimization algorithm used. Rather than using differential evolution, sparrow search or similar population-based approach, hyperparameters were fixed prior to training. As a result, no claim should be made that the choice of architecture is the globally optimal model.

5.6 Comparison with Unsupervised and Reinforcement Learning Methods

Unsupervised approaches to fraud detection are useful where there is limited labeled examples of fraud. Using autoencoders a representation of normal transaction behavior can be built. Anomalous observations are then flagged using reconstruction error [19]. Abbassi et al [13] explored using fusion of unsupervised learning methods for fraud detection in online banking.

One benefit of unsupervised learning is it could pick up previously unseen anomalies without the need for large labeled dataset of fraud. However, an anomalous transaction doesn't necessarily mean it is fraudulent. And false alarms could be raised under a reconstruction based model where customer behavior changes.

In the proposed framework class labels are used and the separation between legitimate and fraudulent transactions is explicitly maximized. This allows the model to build decision boundaries specific to fraud. A downside is good quality labeled examples of fraud are needed for this to work.

As an alternative reinforcement learning has been explored as a way of adapting fraud detection policy to changing environment of transactions [9], [20]. Such a system could be more adaptive than a fixed supervised classifier. Unlike the proposed model once deployed it doesn't continuously update. To pick up changes in fraud strategy or transaction patterns the model would need to be retrained/fine-tuned.

5.7 Architectural Comparison

The methodological characteristics of representative fraud detection approaches and the proposed framework are summarized in Table 5.

Table 5. Methodological comparison of representative fraud detection approaches and the proposed hybrid framework

Study or approach	Principal method	Automatic feature learning	Local pattern extraction	Dependency modelling	Attention	Imbalance handling	Principal limitation
Bahnsen et al. [2]	Machine learning with engineered	Limited	No	No	No	Application-dependent	Strong reliance on handcrafted features

	features						
Lin et al. [4]	Hierarchical recurrent neural network	Yes	Limited	Yes	Architecture-dependent	Not the primary contribution	Recurrent computational complexity
Benchaji et al. [5]	Attention-enhanced LSTM	Yes	Limited	Yes	Yes	Dataset-dependent	Limited explicit local feature extraction
Khalid et al. [8]	Ensemble machine learning	Partial	No	No	No	Ensemble or resampling strategy	Requires multiple classifiers and feature engineering
Qayoom et al. [9]	Reinforcement learning	Yes	Architecture-dependent	Adaptive policy learning	Architecture-dependent	Reward-dependent	Complex training and reward design
Abbassi et al. [13]	Unsupervised learning fusion	Yes	Model-dependent	Limited	No	Does not require balanced labels	Anomalies may not represent actual fraud
Bhat and Nelson [18]	Optimization-supported AI model	Model-dependent	Model-dependent	Model-dependent	No	Technique-dependent	Increased optimization complexity
Sharma et al. [19]	Autoencoder-based anomaly detection	Yes	Limited	Limited	No	Suitable for scarce fraud labels	Reconstruction error may produce false alarms
Proposed framework	TCNN–LSTM with attention	Yes	Yes	Yes	Yes	Computed class weights	Low fraud precision at the selected threshold

Table 5 shows that the proposed framework is distinguished by its integration of local feature extraction, dependency modelling, attention-based weighting, and class-weighted learning within one architecture. However, its current operational limitation is the high number of legitimate transactions classified as fraudulent.

5.8 Comparative Findings

The comparative analysis indicates that the proposed framework offers four important methodological capabilities:

1. automatic learning of transaction representations without extensive handcrafted feature engineering;
2. extraction of localized feature patterns using causal one-dimensional convolution;
3. dependency modelling using LSTM memory operations;
4. adaptive weighting of informative representations through attention.

Compared with conventional machine learning, the framework provides stronger representation learning but requires greater computational resources and offers less direct interpretability. Compared with standalone LSTM models, it adds explicit localized feature extraction. Compared with autoencoder-based anomaly detection, it directly learns from labelled fraudulent transactions. Compared with reinforcement learning, it has a simpler training objective but lacks online policy adaptation.

The principal experimental strength of the proposed model is its fraud recall of 99.70% and false-negative rate of only 0.30%. Its main weakness is the fraud precision of 4.85%, resulting from 32,105 false-positive predictions. Accordingly, the framework demonstrates strong potential as a fraud-screening model, but threshold calibration, probability calibration, ablation analysis, and false-positive reduction are necessary before it can be considered suitable for autonomous decision-making.

6. Discussion

The experimental findings demonstrate that the proposed TCNN–LSTM–Attention framework is highly sensitive to fraudulent transactions. The model achieved an evaluation accuracy of 94.20%, fraud recall of 99.70%, a false-negative rate of only 0.30%, and an ROC-AUC of 0.9989. Of the 1,643 fraudulent transactions, 1,638 were correctly detected, indicating that the proposed architecture effectively learned discriminative patterns associated with fraudulent behaviour. The integration of causal convolution, LSTM-based dependency learning, attention-based weighting, and class-weighted optimization contributed to the strong fraud detection capability of the model.

However, the model generated 32,105 false positive predictions. This gives 4.85% precision for fraud and 9.26% F1-score for the fraud class. This suggests the model favors recall of fraud over precision. By reducing risk of undetected fraud, the model would therefore pick up more legitimate transactions triggering false alerts. This could increase cost of manual checking of alerts and inconvenience legitimate customers. The high ROC-AUC and low precision of fraud suggest the model performs well at ranking fraudulent transactions. However further calibration of class weight and threshold would be needed to suit its current use.

As a result, the framework could be more useful as part of a high sensitivity fraud screening process than as an automatic blocking mechanism of transactions in its current form. With further threshold optimization, probability calibration, adjustment of class weight and secondary verification the framework could reduce false positives without reducing recall of fraud. To confirm applicability and generalization of the framework there would be a need for independent evaluation on a test set, ablation of components and measurement of latency of inference.

7. Conclusion and Future Work

In this study a hybrid TCNN-LSTM framework with attention was proposed for fraud detection in online payments. This framework uses localized feature extraction using causal one dimensional convolution, dependency modeling using LSTM, feature weighting using attention and class weighted learning to account for the highly skewed distribution between legitimate and fraudulent transactions. In the experiments an accuracy of 94.20% of the evaluation and ROC-AUC of 0.9989 were achieved. Of the 1643 fraudulent transactions 1638 were correctly flagged by the model giving 99.70% recall of fraud and just 0.30% false negative rate. The results show the proposed architecture would significantly reduce undetected fraud.

Despite its high sensitivity for fraud detection the framework wrongly labeled 32,105 otherwise legitimate transactions as fraudulent. As a result, the framework was only able to achieve 4.85% precision on fraud and 9.26% F1-score on the fraud class. In its current form more focus is given to fraud recall than minimizing false alarms. As a result, the proposed framework would be more appropriate for use in high sensitivity screening of transactions followed up with verification rather than for autonomous blocking of transactions. The high ROC-AUC shows there is good class ranking ability in the model so with appropriate adjustment of threshold and probability precision could be improved without materially reducing fraud recall.

In future the study should consider tuning the decision threshold and class weights to better trade-off fraud recall against precision. To evaluate the model the study should use separate training, validation and test sets; repeated stratified cross-validation; and ablation of different configurations of TCNN, LSTM, TCNN-LSTM and with attention. As part of future work the research could consider balanced mini-batch learning, focal loss, probability calibration, explainable AI and adaptive learning to pick up changing patterns of fraud and concept drift. To assess suitability of the framework in an online payment system the research should investigate inference latency, throughput of transactions, memory usage and size of model on both CPU and GPU.

References

[1] M. Fiore, F. De Santis, F. Perla, P. Zanetti, and F. Palmieri, “Using generative adversarial networks for improving classification effectiveness in credit card fraud detection,” *Information Sciences*, vol. 479, pp. 448–455, 2019, doi: 10.1016/j.ins.2017.12.030.

[2] C. Bahnsen, D. Aouada, and B. Ottersten, “Feature engineering strategies for credit card fraud detection,” *Expert Systems with Applications*, vol. 51, pp. 134–142, 2016, doi: 10.1016/j.eswa.2015.12.030.

[3] A. Vaswani et al., “Attention is all you need,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 5998–6008.

[4] W. Lin, L. Sun, Q. Zhong, C. Liu, J. Feng, X. Ao, and H. Yang, “Online credit payment fraud detection via structure-aware hierarchical recurrent neural network,” in *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, 2021, pp. 3670–3676, doi: 10.24963/ijcai.2021/505.

[5] I. Benchaji, S. Douzi, B. El Ouahidi, and J. Jaafari, “Enhanced credit card fraud detection based on attention mechanism and LSTM deep model,” *Journal of Big Data*, vol. 8, art. no. 151, 2021, doi: 10.1186/s40537-021-00541-8.

[6] A. A. Almazroi and N. Ayub, “Online payment fraud detection model using machine learning techniques,” *IEEE Access*, vol. 11, pp. 137188–137203, 2023, doi: 10.1109/ACCESS.2023.3339226.

[7] P. Vanini, S. Rossi, E. Zvizdic, and T. Domenig, “Online payment fraud: From anomaly detection to risk

management," *Financial Innovation*, vol. 9, art. no. 66, 2023, doi: 10.1186/s40854-023-00470-w.
[8] A. R. Khalid, N. Owoh, O. Uthmani, M. Ashawa, J. Osamor, and J. Adejoh, "Enhancing credit card fraud detection: An ensemble machine learning approach," *Big Data and Cognitive Computing*, vol. 8, no. 1, art. no. 6, 2024, doi: 10.3390/bdcc8010006.
[9] A. Qayoom, M. A. Khuhro, K. Kumar, M. Waqas, U. Saeed, S. ur Rehman, Y. Wu, and S. Wang, "A novel approach for credit card fraud transaction detection using deep reinforcement learning scheme," *PeerJ Computer Science*, vol. 10, art. no. e1998, 2024, doi: 10.7717/peerj-cs.1998.
[10] P. Hajek, M. Z. Abedin, and U. Sivarajah, "Fraud detection in mobile payment systems using an XGBoost-based framework," *Information Systems Frontiers*, vol. 25, pp. 1985–2003, 2023, doi: 10.1007/s10796-022-10346-6.
[11] E. A. L. Marazqah Btoush, X. Zhou, R. Gururajan, K. C. Chan, R. Genrich, and P. Sankaran, "A systematic review of literature on credit card cyber fraud detection using machine and deep learning," *PeerJ Computer Science*, vol. 9, art. no. e1278, 2023, doi: 10.7717/peerj-cs.1278.
[12] R. Bin Sulaiman, V. Schetinin, and P. Sant, "Review of machine learning approach on credit card fraud detection," *Human-Centric Intelligent Systems*, vol. 2, pp. 55–68, 2022, doi: 10.1007/s44230-022-00004-0.
[13] H. Abbassi, S. El Mendili, and Y. Gahi, "Real-time online banking fraud detection model by unsupervised learning fusion," *HighTech and Innovation Journal*, vol. 5, no. 1, pp. 185–204, 2024.
[14] B. Borketey, "Real-time fraud detection using machine learning," *Journal of Data Analysis and Information Processing*, vol. 12, pp. 189–209, 2024, doi: 10.4236/jdaip.2024.122011.
[15] A. Aslam and A. Hussain, "A performance analysis of machine learning techniques for credit card fraud detection," *Journal of Artificial Intelligence*, vol. 6, no. 1, pp. 1–21, 2024, doi: 10.32604/jai.2024.047226.
[16] D. Hove, O. Olugbara, and A. Singh, "Bibliometric analysis of recent trends in machine learning for online credit card fraud detection," *Journal of Scientometric Research*, vol. 13, no. 1, pp. 43–57, 2024, doi: 10.5530/jscires.13.1.4.
[17] S. Al Balawi and N. Aljohani, "Credit-card fraud detection system using neural networks," *The International Arab Journal of Information Technology*, vol. 20, no. 2, pp. 194–200, 2023.
[18] C. R. Bhat and M. Nelson, "Artificial intelligence-based credit card fraud detection for online transactions optimized with Sparrow Search Algorithm," *International Journal of Performability Engineering*, vol. 19, no. 9, pp. 624–632, 2023, doi: 10.23940/ijpe.23.09.p7.624632.
[19] A. Sharma, G. R. B. R., B. Ramamurthy, and H. B. R., "Credit card fraud detection using deep learning based on auto-encoder," *ITM Web of Conferences*, vol. 50, art. no. 01001, 2022.
[20] S. Vimal, K. Kayathwal, H. Wadhwa, and G. Dhama, "Application of deep reinforcement learning to payment fraud," *arXiv preprint arXiv:2112.04236*, 2021.
[21] Y. Tang, "Automatic fraud detection in e-commerce transactions using deep reinforcement learning and artificial neural networks," *International Journal of Advanced Computer Science and Applications*, vol. 14, no. 7, 2023.
[22] A. R. B. and S. K. K. R., "Credit card fraud detection using artificial neural network," *Global Transitions Proceedings*, vol. 2, pp. 35–41, 2021.
[23] S. Sanober, I. Alam, S. Pande, F. Arslan, K. P. Rane, B. K. Singh, A. Khamparia, and M. Shabaz, "An enhanced secure deep learning algorithm for fraud detection in wireless communication," *Wireless Communications and Mobile Computing*, vol. 2021, art. no. 6079582, 2021, doi: 10.1155/2021/6079582.
[24] M. Lokanan, "Predicting mobile money transaction fraud using machine learning algorithms," *Qeios*, 2022, doi: 10.32388/ELVM4L.
[25] M. Farouk, N. S. Ragab, D. Salama, and O. E., "Fraud_Detection_ML: Machine learning based online payment fraud detection," *Journal of Computing and Communication*, vol. 3, no. 1, pp. 116–131, 2024.