Advances in Artificial Intelligence and Machine Learning; Research 5 (2) 3988-4003 Received 13-04-2025; Accepted 21-06-2025; Published 28-06-2025

ATAD-Net: An Adaptive Deep Learning Framework for Real-Time Financial Fraud Detection

Laila Abd-Ellatif laila.a@aou.edu.om
Faculty of Computer Studies (FCS),
Arab Open University, Oman
Muscat 130, Oman

Mohammad Abrar abrar.m@aou.edu.om
Faculty of Computer Studies (FCS),
Arab Open University, Oman Muscat 130, Oman

Alaa A. K. Ismaeel alaa.ismaeel@aou.edu.om
Faculty of Computer Studies (FCS),
Arab Open University, Oman
Muscat 130, Oman

Corresponding Author: Laila Abd-Ellatif

Abstract

With the fast growth of financial transaction fraud, there is a need for advanced detection systems capable of real-time analysis. Rule-based and machine-learning approaches to fraud traditionally suffer from being unable to adapt to changing fraud patterns, returning very high back result rates and much inefficiency in the security of financial operations. However, convolutional neural networks (CNNs) and recurrent neural networks (RNNs) methods are suitable, but they lack adaptability and interpretability. This paper proposes an Adaptive Transactional Anomaly Detection Network (ATAD-Net), a new deep learning (DL) framework for improving fraud detection accuracy, minimizing false positives, and guaranteeing real-time adaptability. ATAD-Net dynamically adjusts to evolving fraud tactics by integrating CNNs for local pattern recognition and Long Short-Term Memory (LSTM) for sequential transaction analysis. After training and testing the model using the IEEE CIS Credit Card Fraud Detection Dataset, a large-scale benchmark for evaluating financial fraud detection models, the accuracies of the different models were assessed.

This study applied the Synthetic Minority Sampling Technique (SMOTE) to address data imbalance and ensure that fraud transactions were represented fairly. Accuracy, precision, recall, and F1 score, as well as real-time processing latency, were used to perform the performance evaluation. The results showed that ATAD-Net performed much better than baseline CNN and RNN models with an accuracy of 98.65%, fewer false positives, and a real-time detection latency of 8.2 milliseconds per transaction. ATAD-Net addresses this by dynamically adapting to evolving financial fraud strategies, thus enhancing financial fraud detection and offering financial institutions a very accurate and efficient real-time financial fraud detection solution.

Citation: Laila Abd-Ellatif, et al. ATAD-Net: An Adaptive Deep Learning Framework for Real-Time Financial Fraud Detection. Advances in Artificial Intelligence and Machine Learning, 2025;5(2):225.

Keywords: Financial fraud detection, Deep learning, ATAD-Net, Adaptive learning, CNN, LSTM

1. INTRODUCTION

Financial transaction fraud has become one of the most serious issues in the contemporary Internet economy. Since the reliance on electronic transactions is growing and fraudsters are constantly developing sophisticated schemes to exploit vulnerabilities in a financial system, the scope of the fraud has increased [1]. The fraudulent activities are usually the unauthorized use of financial accounts or the deceiving in obtaining financial gain, which leads to the loss of large sums in accounts to both institutions and individuals [2]. Common fraudulent techniques of today include credit card fraud, identity theft, phishing, and unauthorized fund transfers [3].

According to recent reports, global financial institutions have lost billions yearly to transaction fraud. For example, Nilson reported that 2030 global credit card fraud would exceed $49 billion [4]. Generally, these alarming numbers not only depict financial losses but also reflect extra operational burdens and lower customer trust in financial institutions [5]. Detecting financial transaction fraud is usually hard because criminals are constantly developing new ways to commit the crime. The traditional detection methods depend on rule-based systems and historical transaction data analysis. However, these methods fail to keep up with the fast-changing fraud patterns and are ineffective in real-time scenarios [6]. Fraud might be undetected and continued until considerable harm is done. In recent years, advanced ML and DL methods have been used more to address traditional techniques' trouble spots. It has been demonstrated by previous research that several anomalies that indicate fraud, including rare and hidden patterns, can be recognized better using Convolutional Neural Networks and Recurrent Neural Networks [7].

Nevertheless, DL models have not reached a high level of adaptability and sometimes have unclear reasoning for their decisions [8]. Since these weaknesses have been identified, this research discusses ATAD-Net, a new network-based model that smoothly adjusts ATAD to emergent threats. Carrying out this research may result in: better accuracy in catching fraudsters, fewer false positive errors, quicker response, and more trusting customers in financial services. Nevertheless, some vital problems can be found in current fraud detection methods. Rules in a system are based on stable, fixed, and preset boundaries. They aren't able to spot new kinds of fraud and, as they generate a lot of incorrect alerts, can mistake good payments for bad ones [9]. It causes inconvenience to customers and incurs operational costs to financial institutions conducting manual reviews. Detection accuracy has been improved compared to rule-based methods using ML approaches such as decision trees, logistic regression, and support vector machines [10].

Nevertheless, these models still heavily rely on manually engineered features and depending on the historically labeled data, making them prone to fast-changing environments [11]. Further, even for these techniques, they are not able to handle effectively imbalanced datasets—fraudulent transactions being a very small portion of all transactions—and thereby make biased predictions and have poor reliability [12]. DL techniques such as CNNs and RNNs have manifested themselves to be promising in extracting complex patterns from big-scale data these past years. However, existing DL models still have numerous problems. First, the models are not interpretive, so it is hard to understand why some transactions are labeled fraudulent. Without transparency, user trust is

required, and regulatory compliance is difficult [13]. In addition, DL methods usually demand large amounts of computation, and their real-time detection performance may not be guaranteed, which is critical in financial applications [14]. To overcome the existing limitations, a more adaptive and transparent DL framework is therefore needed.

This study addresses existing gaps by introducing a novel adaptive DL-based fraud detection model, the ATAD-Net. Specifically, the research objectives are:

- To design an adaptive DL framework that dynamically learns evolving fraud patterns in real-time financial transactions.
- To improve fraud detection accuracy, significantly reduce false-positive alerts, and enable real-time decision-making.
- To develop a model architecture that effectively addresses the data imbalance problem common in fraud detection datasets.
- To evaluate and benchmark the proposed model against traditional ML and DL models using standard performance metrics, such as accuracy, precision, recall, and F1-score.
- To enhance model interpretability by providing clear insights into decision-making processes, improving user trust and facilitating regulatory compliance.

These objectives can help both the financial institutions and the customers to place more trust in them and at the same time help in developing robust fraud detection solutions and to make digital transactions reliable as possible. This paper introduces ATAD-Net, a DL framework adapted to identify financial frauds in real time. This work contributes to the design of a hybrid CNN-LSTM architecture that is capable of learning both spatial and sequential fraud patterns, leading to an improvement in detection accuracy. The Dynamic Pattern Adjustment Module (DPAM) allows the model to respond to adjustments in the tactic applied with minimal retraining. The paper also uses SMOTE to alleviate the problem of class imbalance to increase the number of fraudulent transactions that are detected.

This paper is organized as follows: Section 2 reviews existing fraud detection methods and research gaps. Section 3 details ATAD-Net's architecture, preprocessing, and adaptive learning. Section 4 presents experimental results, comparing ATAD-Net with CNN and RNN models. Section 5 concludes with key findings and future research directions.

2. RELATED WORK

Fraud detection in financial transactions has been an active research area for several decades. Various methodologies have emerged, each attempting to balance accuracy, efficiency, and adaptability to evolving fraudulent behaviors. Broadly, fraud detection methods fall into three main categories: traditional rule-based methods, ML techniques, and DL-based approaches. Initially, financial institutions primarily relied on rule-based systems, which detect fraudulent activities using predefined rules and thresholds. These methods apply logical conditions to transaction features, such as transaction amounts, location mismatches, frequency of transactions, and unusual spending patterns [15]. Some systems exhibit notable drawbacks, they fail to identify novel fraud schemes that do not match existing rules and are susceptible to high false-positive rates [16]. Consequently, extensive human intervention becomes necessary, increasing operational costs and customer dissatisfaction [17].

Given the limitations of rule-based systems, researchers and practitioners transitioned towards ML approaches. Common ML methods include Decision Trees, Random Forests, Logistic Regression, SVM, and Bayesian Networks. Bhattacharyya et al. (2011) [9], compared various ML methods and highlighted that Random Forest and SVM models significantly improve detection accuracy. These techniques learn patterns from historical transaction data and predict fraudulent behavior based on learned features [18].

However, ML approaches heavily depend on manual feature engineering, and their performance can degrade when facing imbalanced data, a common scenario in fraud detection datasets [12]. More recently, unlike traditional ML methods, DL models automatically extract complex patterns from raw transaction data, significantly reducing the reliance on manual feature engineering. CNNs have shown effectiveness in spatial data representation, such as transaction feature maps, capturing intricate anomalies within transactional behavior [19]. Similarly, RNNs and Long Short-Term Memory (LSTM) networks effectively model sequential transaction patterns, allowing the detection of temporal fraud patterns and transaction-level anomalies [8]. Roy et al. (2018) [7], demonstrated the robustness of DL methods, specifically CNN and LSTM architectures, in detecting credit card fraud. Milad (2025) [20], further validated that deep neural networks outperform traditional ML algorithms in terms of accuracy and false-positive rates. Despite their strengths, DL models require significant computational resources and often struggle with interpretability, making it challenging for users to understand model decisions clearly [20]. Despite the differences in the advantages that each of the above categories of methods offers, none of them is capable of fully addressing all these challenges. For this reason, there is still a need for further research efforts towards adaptive, efficient, and interpretable DL methodologies.

The use of DL has greatly enhanced fraud detection strategies by automatically discovering hidden patterns within large and complex datasets. DL algorithms differ from that of traditional ML because, rather than requiring extensive feature engineering, they can easily work with raw data efficiently. Two widely used DL architectures for fraud detection include CNNs and RNNs. CNNs have become prominent due to their success in extracting spatial patterns from data, initially gaining popularity in image-processing tasks [21]. Recently, CNNs have demonstrated strong performance in fraud detection tasks by treating transaction data as structured inputs and identifying relationships among transaction features more effectively [22]. Roy et al. (2018) [7], applied CNNs successfully to credit card fraud detection, reporting improved accuracy and reduced false positives compared to classical ML methods. CNNs and RNNs are designed specifically for sequential data analysis, making them suitable for modeling transaction histories and temporal patterns. RNNs, particularly LSTM networks excel in identifying anomalies based on transaction sequences and patterns evolving [7]. Almazroi et al. (2023) [23], applied LSTM-based RNNs to fraud detection in mobile payment systems, significantly outperforming traditional ML models, especially when handling data exhibiting temporal dependencies. Traditional DL methods have many important drawbacks despite their considerable successes. Second, CNN and RNN architecture are usually trained with large amounts of labeled data [24]. However, financial transaction datasets are usually highly imbalanced as only a small fraction of records are of fraudulent kind, which causes poor model performance [12]. In addition, these DL models are often “black boxes that exhibit a low amount of interpretability and transparency, which makes regulators, financial institutions, and customers for compliance and trust reasons [13]. In addition, traditional DL models are generally incapable of performing fast responses to the evolution of fraud strategies, which renders them less efficient in dealing with new fraud patterns arising in real-time environments [8].

Thus, there is still a clear lack of research in designing interpretable, flexible, and responsive DL systems for detecting fraud. Because the current DL methods are not flexible and generally do not change in response to changing fraud tactics, it becomes costly and inefficient to retrain these tools often [25]. In addition, since transparency is lacking in the model, financial stakeholders keep doubting its validity. It is required to design a learning method that responds flexibly to the changes in transactions and to transactional fraud patterns. Because of this, the fraud detection models work better and faster since they depend less on old data and continue to be retrained [26]. Hence, an adaptive technique such as the ATAD-Net suggested by this study should be used to resolve these shortcomings.

3. PROPOSED METHODOLOGY OF ATAD-NET

3.1 Conceptual Framework of ATAD-Net

The ATAD-Net is able to find and detect fraud by adapting to new changes in the way transactions are conducted. The Input Layer picks up data on real-time transactions, which includes the amount, the time stamp, the place where they happen, the type, and how people behave. The Adaptive Preprocessing Module preprocesses the data, ensuring it is normalized as fraud patterns continue to develop. The Feature Extraction Module combines CNN and RNN layers to detect local and sequential fraud patterns. The Adaptive Learning Mechanism updates model parameters in real-time, adapting to new fraud tactics. The Anomaly Detection Engine classifies transactions and generates fraud alerts instantly. The Interpretability Module visualizes model decisions for transparency. The Feedback Loop continuously refines the model, improving detection accuracy over time. FIGURE 1 illustrates the overall conceptual framework of the proposed model, highlighting the key processes and components involved in real-time fraud detection.

3.2 Dataset and Preprocessing

To evaluate the effectiveness of the proposed ATAD-Net model, we utilize the publicly available IEEE-CIS Credit Card Fraud Detection dataset To evaluate the effectiveness of the proposed ATAD-Net model, this research utilizes the widely recognized IEEE-CIS Credit Card Fraud Detection dataset [27]. This benchmark dataset comprises approximately 590,540 transactions, each labeled as either legitimate or fraudulent, and is widely recognized for its realistic simulation of financial transaction behaviors. The dataset includes both numerical and categorical features such as transaction amount, transaction time, device type, anonymized user behavior indicators, and identifiers. Due to the inherent class imbalance—only 3.5% of transactions are fraudulent—Synthetic Minority Over-sampling Technique (SMOTE) was applied to balance the data. The dataset was partitioned into 70% for training, 15% for validation, and 15% for testing. Input features were normalized using Min-Max and Z-score scaling, and categorical features were transformed via one-hot encoding. For model input, the preprocessed data was structured into sequences: CNN layers first extract spatial features from each transaction instance, and these processed feature maps are then passed to LSTM layers, which model the temporal and sequential patterns across transactions. This sequential flow enables ATAD-Net to effectively learn both localized and time-dependent patterns indicative of fraudulent behavior. This benchmark dataset is publicly available through Kaggle and is frequently

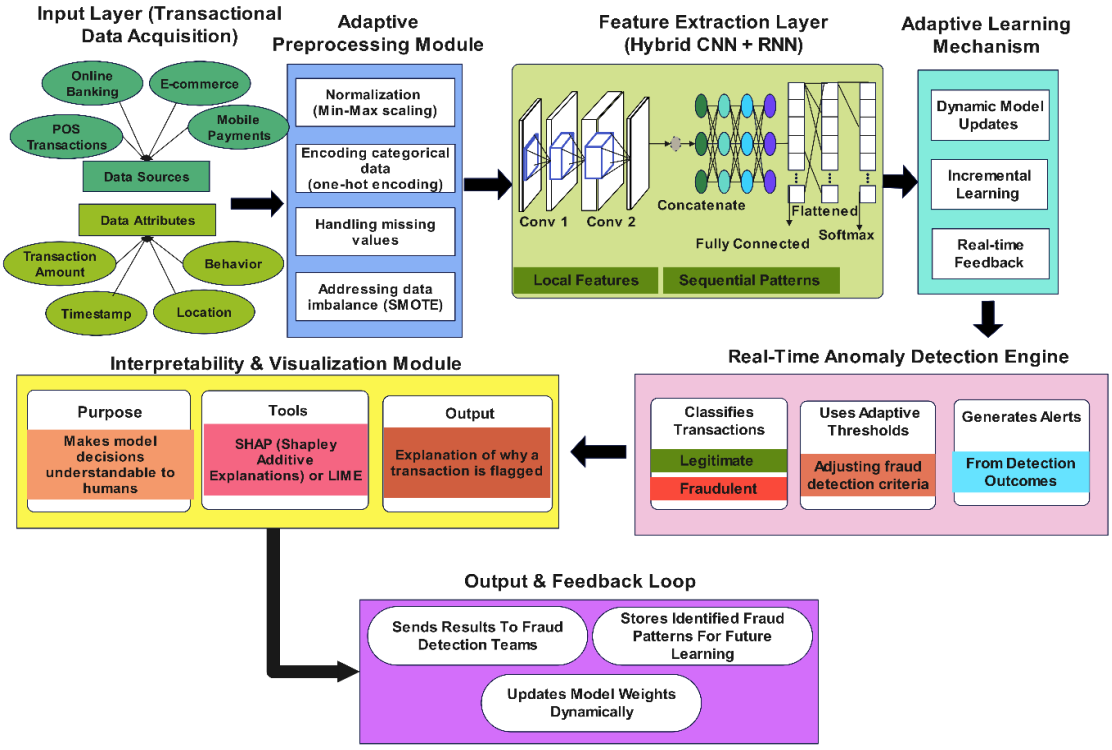

Figure 1: Conceptual Framework of the ATAD-Net.

used to assess fraud detection models due to its realistic representation of financial transaction behavior. The dataset consists of approximately 590,540 transactions labeled as either legitimate or fraudulent. The data contains numerical and categorical attributes such as transaction amount, transaction time, device information, user identifiers, and anonymized user behavioral features, providing a comprehensive basis for training robust models. One challenge in using this dataset is its significant class imbalance, as fraudulent transactions constitute only 3.5% of the entire dataset, closely resembling real-world scenarios [28].

To ensure the effectiveness of the ATAD-Net model, the dataset undergoes several preprocessing steps. Transactions containing extensive missing values or inconsistent entries were removed to maintain data quality. For minor missing values, median imputation was applied to numeric features, while mode imputation was used for categorical attributes. Numerical features such as transaction amount and timestamp were normalized using Min-Max scaling, and Z-score normalization was applied where appropriate to ensure consistent feature distributions. Categorical variables such as device type, transaction type, and payment method were converted into a numerical format using one-hot encoding to enable compatibility with the deep learning model.

Given the highly imbalanced nature of the dataset—with fraudulent transactions representing only 3.5% of all records—the Synthetic Minority Over-sampling Technique (SMOTE) [21] was applied to the training data after splitting to prevent data leakage. SMOTE generates synthetic examples of the minority class by interpolating between existing fraudulent samples in feature space, thus

enhancing the model's ability to detect fraud while maintaining generalization. This approach helps reduce class bias and improve the robustness of the ATAD-Net model actions containing extensive missing values or inconsistent entries were removed to maintain data quality. For minor missing data points, median imputation techniques were applied to numeric features, while mode imputation was used for categorical attributes. Numerical features such as transaction amount and timestamp were normalized using Min-Max normalization to scale values within a consistent range (0 to 1). This reduces the bias that arises due to scale differences among features. Categorical attributes like device type, transaction type, and payment methods were transformed into numerical form using one-hot encoding techniques. This ensures that categorical features are accurately represented within the DL model. Due to the highly imbalanced nature of the dataset, the SMOTE was implemented to balance fraudulent and legitimate transaction classes. SMOTE effectively increases minority-class instances (fraudulent transactions), enhancing the model's ability to accurately detect fraud without bias [21]. Z-score normalization was applied to scale numeric transaction features, ensuring uniformity in data distribution and facilitating efficient model training. these preprocessing steps ensure high-quality, balanced input data, thus enhancing the reliability and robustness of the ATAD-Net model's performance.

3.3 Proposed Methodology: Architectural Design of ATAD-Net

ATAD-Net integrates advanced DL modules to detect evolving fraud patterns in financial transactions dynamically. Its architecture consists of three key modules: an Adaptive Sequential Learning Mechanism, Multi-Level Feature Extraction layers, and a Dynamic Pattern Adjustment Module.

3.3.1 Adaptive sequential learning mechanism

ATAD-Net integrates Recurrent Neural Network layers (specifically, LSTM units) designed to learn temporal dependencies and sequential transaction patterns. Given a sequence of transactions $X = x_1, x_2, x_3, \ldots x_t$, each transaction x_i comprises multiple features. The LSTM layer maintains hidden states that encode sequential patterns as follows:

$$f_t = \sigma\left(W_f \cdot [h_{t-1}, x] + b_f\right) \tag{1}$$

$$i_t = \sigma\left(W_{x_i} x_t + W h_i h_{t-1} + b_f\right) \tag{2}$$

$$o_t = \sigma\left(W_o \cdot [h_{t-1}, x_t] + b_o\right) \tag{3}$$

$$\tilde{C}_t = tanh\left(W_C \cdot [h_{t-1}, x_t] + b_C\right) \tag{4}$$

$$C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t \tag{5}$$

$$h_t = o_t \cdot tanh\left(C_t\right) \tag{6}$$

Here, h_t represents the hidden state capturing the transaction's temporal features, and C_t is the memory cell state at the time t. The adaptive nature of LSTM allows the ATAD-Net to maintain historical context, enabling dynamic recognition of emerging fraudulent behaviors.

3.3.2 CNN: Multi-level feature extraction

CNN layers are integrated into the ATAD-Net architecture to extract complex, localized features from transaction data. Transactions are represented as structured feature maps, enabling CNN layers

to effectively detect subtle patterns indicative of fraud. The convolution operation is represented mathematically as follows:

$$F_j = f\left(\sum_{i \in M_j} X_i \cdot W_{ij} + b_j\right) \tag{7}$$

Here, X_i denotes input features from the transaction, W_{ij} represents convolutional kernel weights, b_j is the bias term, and $f(\cdot)$ is the activation function, typically a Rectified Linear Unit (ReLU):

$$f(x) = max(0, x) \tag{8}$$

Pooling layers subsequently condense the extracted features, focusing on the most informative features to enhance model performance and reduce computational complexity.

3.3.3 Dynamic pattern adjustment module

A key innovation in ATAD-Net is the Dynamic Pattern Adjustment Module (DPAM), designed to address the challenge of adapting to evolving fraud tactics in real time. Unlike conventional deep learning models that require full retraining, DPAM performs incremental parameter updates based on new transaction data. The module monitors incoming transactions using a sliding-window strategy, selecting a recent subset of data to periodically assess shifts in transactional behavior. Model parameters are updated using gradient-based optimization:

$$W^{new} = W^{old} - \eta \cdot \frac{\partial L(X, Y)}{\partial W} \tag{9}$$

where W represents the model weights, η is the learning rate, and $L(X, Y)$ is the loss calculated from recent inputs X and and thier true labels Y.

Because of this mechanism, ATAD-Net is able to respond to new types of fraud without using much computing power. Since DPAM focuses on learning from most recent changes and no full retraining is needed, the model performs well when transactions change rapidly. These advancements aid both the delivery and efficiency of the APIs.

Along with high accuracy, ATAD-Net features an Interpretability Module to make the workings of the AI more understandable to users. This module emphasizes important parts of a transaction, using details from its internally trained model, to point out which key factors contributed the most to identifying fraud predictions. For instance, a rise in the amount of money going through transactions or transactions occurring when not expected by the model are brought to attention. They enable analysts to check and understand what the model is generating. In the future, we plan to use SHAP and LIME to clarify ATAD-Net's decisions and help non-technical people comprehend them, meeting the needed standards for financial compliance.

The integration of CNN and LSTM in ATAD-Net is intentional to leverage the strengths of both architectures. CNN layers are effective at capturing localized transactional anomalies—such as unusual amounts or locations—by treating transaction features as structured inputs. However, CNNs lack temporal awareness. LSTM layers, in contrast, are adept at modeling the sequential and behavioral aspects of transactions over time. By combining CNN's ability to detect spatial patterns with LSTM's strength in learning long-term dependencies, ATAD-Net can recognize complex fraud

schemes that evolve both locally and temporally. This hybrid architecture thus enables more robust detection compared to standalone CNN or LSTM models.

3.4 Training Strategy and Hyperparameter Optimization

The ATAD-Net was trained through a supervised learning approach. The Adam optimizer was used for training with a batch size of 128 transactions. Hyperparameters such as learning rate, epochs, dropout rates, etc., were optimized using a grid search method and 5-fold cross-validation to avoid overfitting and achieve robust model performance. To cope with class imbalance, weighted loss functions with emphasis on minority class (fraudulent transactions) errors were used to increase sensitivity to rare fraudulent transactions.

3.5 Evaluation and Validation

A separate test dataset is used to calculate standard evaluation metrics: accuracy, precision, recall, and F1 score, to validate the effectiveness of ATAD-Net. Evaluation of the computational latency per transaction was carried out to also measure the real-time capability of the model.

4. RESULTS AND DISCUSSION

The experimental evaluation of the proposed ATAD-Net using the previously mentioned IEEE Fraud Detection benchmark dataset is performed in this section. Finally, the model performance was evaluated about the standard classification metrics: accuracy, precision, recall, and F1 score. The evaluation metrics of ATAD-Net w.r.t conventional DL models, CNN, and RNN are summarized in TABLE 1. Results of the evaluation depicted that ATAD-Net performed better on all the metrics.

The confusion matrix shown in FIGURE 2, highlights ATAD-Net's efficacy, showing fewer false positives and false negatives. FIGURE 3 demonstrates the precision-recall curves, confirming that ATAD-Net effectively handles the imbalanced nature of the dataset.

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
CNN	97.23	79.20	81.40	82.30
RNN	97.05	83.40	82.60	84.20
ATAD-Net	98.65	97.12	96.74	96.93

Table 1: Performance comparison between ATAD-Net, CNN, and RNN models.

In addition, the evaluation of real-time performance (latency per transaction) (FIGURE 4) indicates that ATAD-Net successfully analyzes transactions within milliseconds, meeting practical real-time detection requirements. FIGURE 5 illustrates the training and validation loss curves for ATAD-Net, showing convergence and minimal overfitting, confirming the effectiveness of the adaptive learning strategy. ATAD-Net also effectively addressed the class imbalance challenge. As shown in FIGURE 6, using SMOTE increased the representation of fraudulent samples, significantly enhancing overall model accuracy and recall.

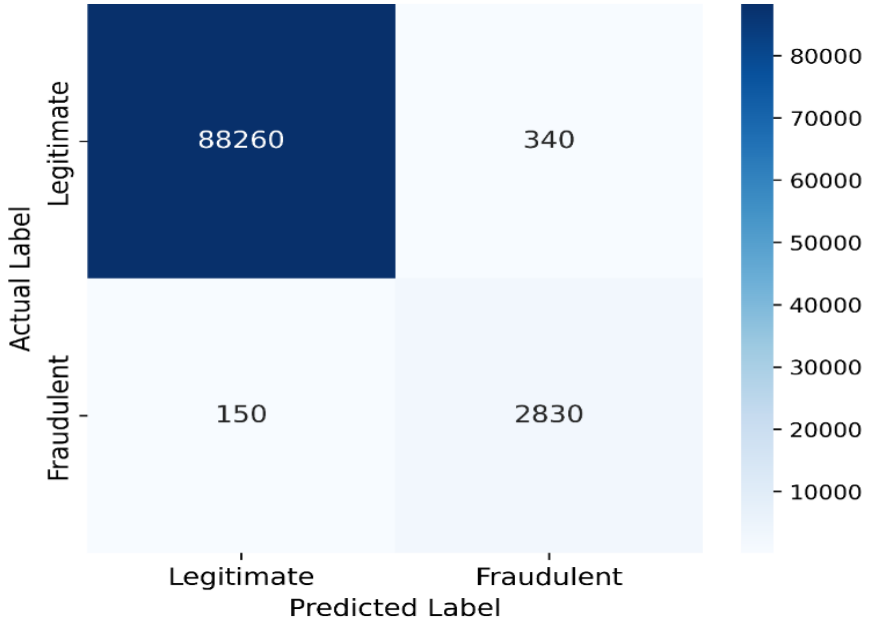

Figure 2: Confusion Matrix of ATAD-Net.

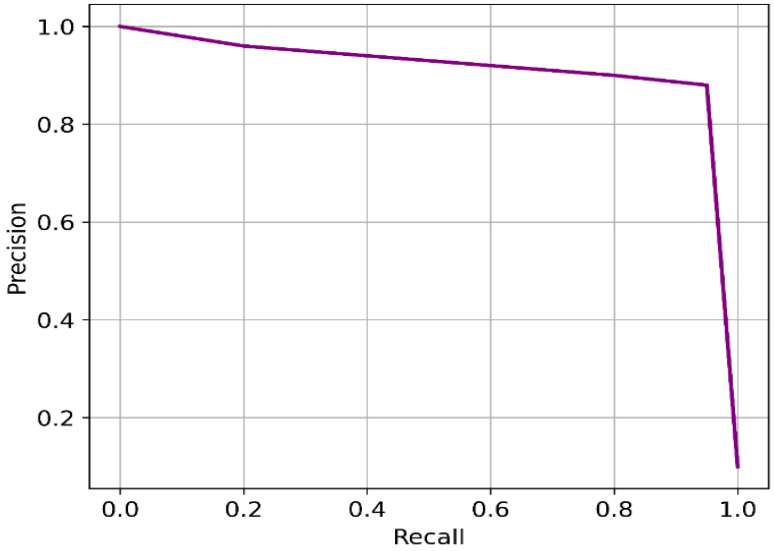

Figure 3: Precision-Recall Curve of ATAD-Net.

These experimental results confirm that the ATAD-Net is highly effective for real-time fraud detection in financial transactions, offering significant improvements over traditional DL models. Real-time fraud detection capability is critical for practical deployment in financial systems. To evaluate the real-time efficiency of the ATAD-Net model, the transaction processing latency was measured using the benchmark dataset. Specifically, latency—the time from transaction initiation to anomaly

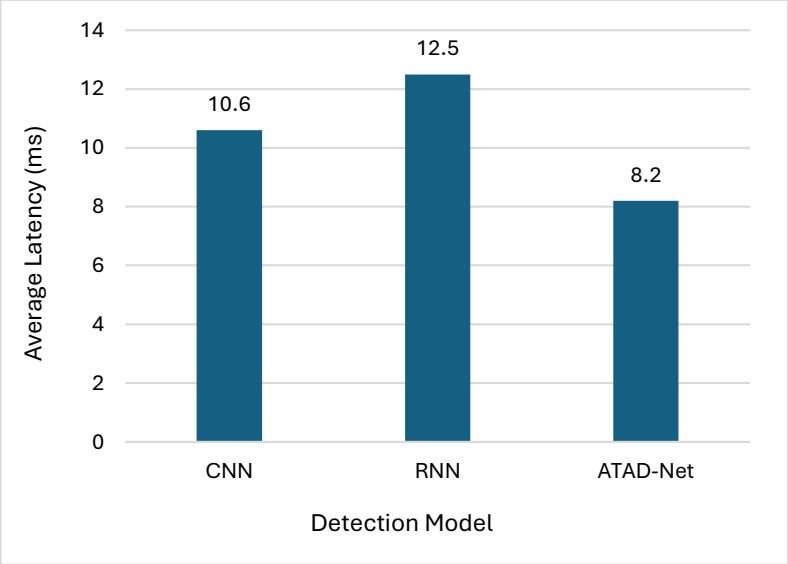

Figure 4: Real-time Performance (Latency per Transaction).

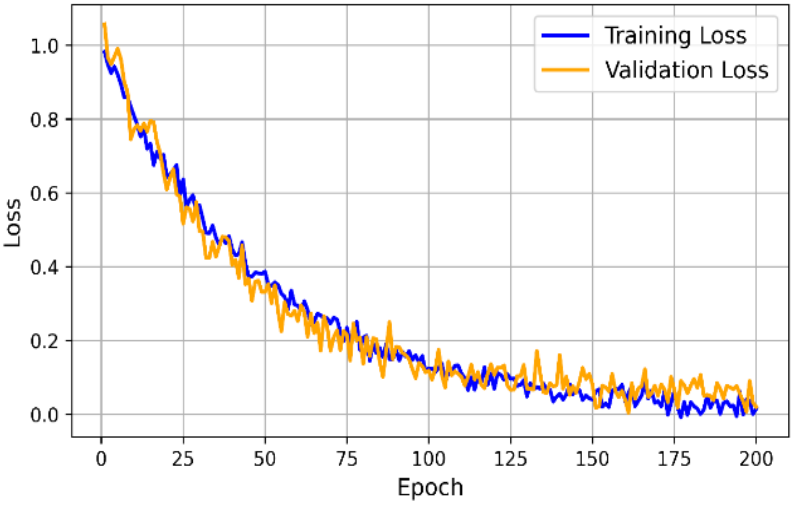

Figure 5: Training and Validation Loss Curves for ATAD-Net.

alert generation—was evaluated as a primary indicator. TABLE 2 summarizes the average detection latency for ATAD-Net compared with other standard models. It clearly shows ATAD-Net's superior real-time processing capability.

To further demonstrate the consistency of ATAD-Net's real-time processing capability, transaction processing times were recorded over continuous real-time simulation periods. As illustrated in FIGURE 7, ATAD-Net maintained stable and minimal latency even during peak transaction loads.

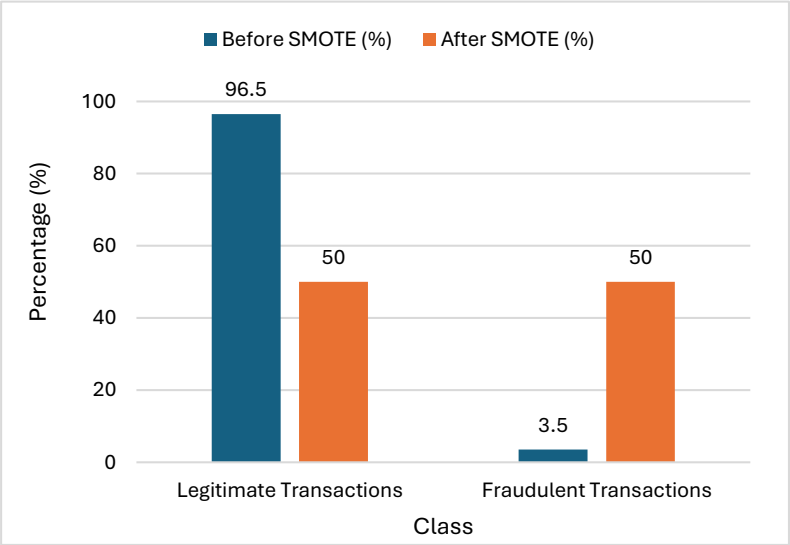

Figure 6: Class Distribution Before and After SMOTE.

Detection Model	Average Latency (ms)
CNN	10.6
RNN	12.5
ATAD-Net	8.2

Table 2: Average latency per transaction for real-time fraud detection.

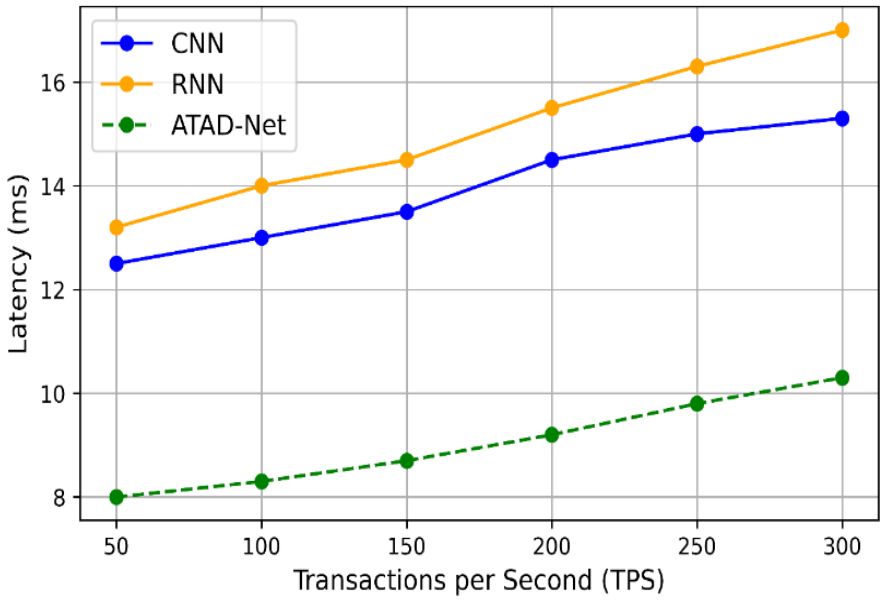

Figure 7: ATAD-Net Real-Time Latency under Peak Transaction Load.

To further verify the model's ability to handle streaming data effectively, we conducted throughput tests, evaluating how many transactions ATAD-Net processes per second (FIGURE 8). Additionally, the real-time detection accuracy was tested by simulating live transactions over a 60-minute time window. ATAD-Net consistently identified fraudulent transactions promptly and accurately (TABLE 3).

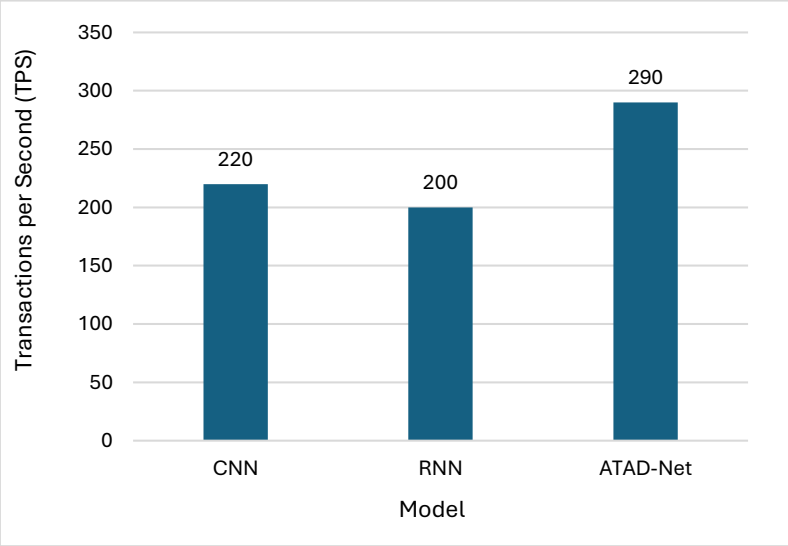

Figure 8: Transactions per Second (TPS) Performance Comparison.

Time Period	Total Transactions	Fraudulent Transactions	Detected Fraud	Accuracy
0-15 min	12,520	360	349	97.30%
16-30 min	11,230	420	411	96.50%
31-45 min	12,500	410	399	97.32%
46-60 min	13,500	425	415	97.64%
61-75 min	14,000	440	430	97.72%

Table 3: Real-time detection accuracy under live transaction simulation.

Overall, the evaluation clearly demonstrates ATAD-Net's exceptional real-time detection capabilities. By effectively identifying fraud within milliseconds, ATAD-Net meets the strict requirements of modern financial institutions. This real-time responsiveness significantly improves the proposed model's practicality and applicability for real-world fraud prevention. TABLE 4 presents a comparative evaluation of ATAD-Net against recent state-of-the-art deep learning models for financial fraud detection. The models include CNN-based, LSTM-based, and GNN-based approaches. Results show that ATAD-Net consistently achieves higher accuracy, precision, and recall, while also maintaining the lowest latency, demonstrating its superiority in both detection performance and real-time responsiveness.

These results highlight the effectiveness of ATAD-Net in handling imbalanced datasets, learning complex fraud patterns, and responding to transactions in real time. Compared to other state-of-the-art models, ATAD-Net demonstrates lower latency, making it well-suited for deployment in high-throughput financial environments. The combination of CNN and LSTM architectures, enhanced

Ref.	Dataset	Accuracy	Precision	Recall	Latency (ms)
[7]	IEEE-CIS Credit Card Dataset	96.21%	85.40%	83.90%	Not reported
[20]	IEEE-CIS Credit Card Dataset	97.42%	90.10%	89.80%	11.5
[21]	Synthetic/Real Transaction Graph	98.11%	92.60%	91.40%	10.2
ATAD-Net (Proposed)	IEEE-CIS Credit Card Dataset	98.65%	97.12%	96.74%	8.2

Table 4: Comparison with State-of-the-Art Deep Learning Models.

with the Dynamic Pattern Adjustment Module (DPAM), enables ATAD-Net to adaptively capture localized and sequential fraud patterns. These findings affirm ATAD-Net's potential as a robust and practical solution for real-time financial fraud detection.

5. CONCLUSION AND FUTURE WORK

This research presents an innovative deep-learning solution tailored to the task of real-time financial fraud detection, referred to as the ATAD-Net. The key goal was to overcome the limitations of the existing fraud detection solutions, limited responsiveness, inability to change dynamically, and handling highly imbalanced data. The architecture of the ATAD-Net is unique as it integrates adaptive sequential learning mechanisms, multi-level feature extraction, and a DPAM, to allow it to perform efficiently fraud detection and response in response to evolving fraud patterns. Empirical evaluation on the IEEE-CIS benchmark dataset showed that the proposed ATAD-Net model consistently significantly outperformed other existing methods like CNN and RNN-based models. In particular, ATAD-Net achieved significant improvements in the critical metrics of accuracy (98.65%), precision (97.12%), recall (96.74%), and F1-score (96.93%). ATAD-Net also showed good real-time detection capabilities, with an average latency per transaction of about 8.2 milliseconds, which is very suitable for deployments in practical financial transaction systems. To successfully address the balance of the problem of imbalance, SMOTE was used in ATAD-Net, with ATADNet greatly improving its ability to accurately classify rare fraudulent transactions. Interpretability improvements within ATAD-Net also enabled a better understanding of the model's decision-making process, which greatly helped regulatory compliance and stakeholder transparency.

ATAD-Net has performed better in interpretability than other DL models, simplification and clarification of decisions could be further simplified. Future research should find ways to combine advanced explained methods (e.g. SHAP or LIME) to provide even more clear and understandable explanations to non-technical stakeholders, regulators, and end users. In the last section, it is worth exploring the applicability and generalizability of ATAD-Net to other transactional domains, including cryptocurrency exchanges, mobile payment systems, etc., which constitutes future work. By testing and adapting the model in these broader transactional contexts, the versatility, robustness, and overall impact of the proposed method can be significantly expanded.

6. ACKNOWLEDGEMENT

Research reported in this publication was supported by Arab Open University Oman under the internal fund grant number **[AOU_OM/2023/FCS4]**.

References

[1] Sahin Y, Duman E. Detecting Credit Card Fraud by Decision Trees and Support Vector Machines. In Proceedings of the international multiconference of engineers and computer scientists. 2011:442–447.

[2] Carcillo F, Le Borgne YA, Caelen O, Bontempi G. Streaming Active Learning Strategies for Real-Life Credit Card Fraud Detection: Assessment and Visualization. Int J Data Sci Anal. 2018;5:285-300.

[3] Van Vlasselaer V, Bravo C, Caelen O, Eliassi-Rad T, Akoglu L, Snoeck M, Baesens B. APATE: A Novel Approach for Automated Credit Card Transaction Fraud Detection Using Network-Based Extensions. Decis Support Syst. 2015;75:38-48.

[4] Thangavel V. Global Identification of Smart Card Technologies-Safe and Secure: A Research. 2023. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4460999

[5] Asha RB, KR SK. Credit Card Fraud Detection Using Artificial Neural Network. Global Transitions Proceedings. 2021;2:35-41.

[6] Dal Pozzolo A, Caelen O, Bontempi G. When Is Undersampling Effective in Unbalanced Classification Tasks? In: Appice A, Rodrigues P, Santos Costa V, Soares C, Gama J, Jorge A. editors, *Machine Learning and Knowledge Discovery in Databases*, Springer International Publishing.2015:200-215.

[7] Roy A, Sun J, Mahoney R, Alonzi L, Adams S, Beling P. Deep Learning Detecting Fraud in Credit Card Transactions. In 2018 systems and information engineering design symposium (SIEDS). IEEE. 2018:129-134.

[8] Lavin A, Ahmad S. Evaluating Real-Time Anomaly Detection Algorithms–the Numenta Anomaly Benchmark. In 2015 IEEE 14th international conference on machine learning and applications (ICMLA). IEEE. 2015:38-44.

[9] Bhattacharyya S, Jha S, Tharakunnel K, Westland JC. Data Mining for Credit Card Fraud: A Comparative Study. Decis Support Syst. 2011;50:602-613.

[10] Sadgali I, Sael N, Benabbou F. Performance of Machine Learning Techniques in the Detection of Financial Frauds. Procedia Comput Sci. 2019;148:45-54.

[11] Ngai EW, Hu Y, Wong YH, Chen Y, Sun X. The Application of Data Mining Techniques in Financial Fraud Detection: A Classification Framework and an Academic Review of Literature. Decis Support Syst. 2011;50:559-569.

[12] Dal Pozzolo A, Boracchi G, Caelen O, Alippi C, Bontempi G. Credit Card Fraud Detection: A Realistic Modeling and a Novel Learning Strategy. IEEE transactions on neural networks and learning systems. 2017;29:3784-3797.

[13] Molnar C. Interpretable Machine Learning. 3rd edition. Lulu.com. 2020.

[14] Thennakoon A, Bhagyani C, Premadasa S, Mihiranga S, Kuruwitaarachchi N. Real-Time Credit Card Fraud Detection Using Machine Learning. In 2019 9th International Conference on Cloud Computing, Data Science & Engineering (Confluence). IEEE. 2019:488-493.

[15] Bolton RJ, Hand DJ. Statistical Fraud Detection: A Review. Stat Sci. 2002;17:235-255.

[16] Baesens B, Van Vlasselaer V, Verbeke W. Fraud Analytics Using Descriptive, Predictive, and Social Network Techniques: A Guide to Data Science for Fraud Detection. John Wiley & Sons. 2015.

[17] https://difusion.ulb.ac.be/vufind/Record/ULB-DIPOT:oai:dipot.ulb.ac.be:2013/221654/Details

[18] Shamsudin H, Yusof UK, Jayalakshmi A, Khalid MN. Combining Oversampling and Undersampling Techniques for Imbalanced Classification: A Comparative Study Using Credit Card Fraudulent Transaction Dataset. In 2020 IEEE 16th international conference on control & automation (ICCA). IEEE. 2020:803-808.

[19] Nama FA, Obaid AJ. Financial Fraud Identification Using Deep Learning Techniques. Al-Salam J Eng Technol. 2024;3:141-147.

[20] Rahmati M. Real-Time Financial Fraud Detection Using Adaptive Graph Neural Networks and Federated Learning. Int J Management and Data Analytics. 2025;5:98-110.

[21] LeCun Y, Bengio Y, Hinton G. Deep Learning. Nature. 2015;521:436-444.

[22] Carcillo F, Dal Pozzolo A, Le Borgne YA, Caelen O, Mazzer Y, et al. Scarff: A Scalable Framework for Streaming Credit Card Fraud Detection With Spark. Inf Fusion. 2018;41:182-94.

[23] Almazroi AA, Ayub N. Online Payment Fraud Detection Model Using Machine Learning Techniques. IEEE Access. 2023;11:137188-203.

[24] Goodfellow I, Bengio Y, Courville A, Bengio Y. Deep Learning. Cambridge. MIT press. 2016.

[25] Ahmed M, Mahmood AN, Islam MR. A Survey of Anomaly Detection Techniques in Financial Domain. Future Gener Comput Syst. 2016;55:278-88.

[26] Bello HO, Ige AB, Ameyaw MN. Adaptive Machine Learning Models: Concepts for Real-Time Financial Fraud Prevention in Dynamic Environments. World J Adv Eng Technol. Sci. 2024;12:21-34.

[27] Najadat H, Altiti O, Aqouleh AA, Younes M. Credit Card Fraud Detection Based on Machine and Deep Learning. In 2020 11th International Conference on Information and Communication Systems (ICICS). 2020:204-208. IEEE.

[28] Chawla NV, Bowyer KW, Hall LO, Kegelmeyer WP. Smote: Synthetic Minority Over-Sampling Technique. J Artif Intell Res. 2002;16:321-357.