From Click to Checkout: Deep Learning for Real-Time Fraud Detection in E-Payment Systems

Raouya El Youbi¹, Fayçal Messaoudi², Manal Loukili^{2,*}, Riad Loukili¹

¹National School of Applied Sciences, Sidi Mohamed Ben Abdellah University, Fez, Morocco ²National School of Business and Management, Sidi Mohamed Ben Abdellah University, Fez, Morocco

Abstract The rapid expansion of e-commerce has been paralleled by a significant increase in electronic payment (e-payment) transactions, bringing forth pressing challenges in maintaining transactional security. This paper addresses the critical issue of e-payment fraud in e-commerce by leveraging deep learning techniques for real-time fraud detection. With the growing sophistication of fraudulent activities, traditional rule-based fraud detection systems are proving inadequate, necessitating more advanced and adaptable solutions. This study proposes a deep learning model, specifically designed to enhance e-payment security by efficiently identifying fraudulent transactions. The model addresses key challenges such as class imbalance in transaction data and the need for real-time processing capabilities. Through a comprehensive methodology involving data preprocessing, model architecture design, training, and evaluation, the paper demonstrates the effectiveness of deep learning in detecting complex fraud patterns with high accuracy. The findings highlight the potential of deep learning to significantly improve the security of e-payment systems in e-commerce, thereby bolstering consumer trust and the overall integrity of online transactions. This research contributes to the evolving landscape of e-commerce security, offering insights and directions for future advancements in fraud detection technologies.

Keywords E-Commerce Security, Real-Time Fraud Detection, Deep Learning, Neural Networks, Class Imbalance.

DOI: 10.19139/soic-2310-5070-2891

1. Introduction

In the digital age, e-commerce has emerged as a cornerstone of the global economy, reshaping how consumers interact with goods and services [1]. The convenience and efficiency of online shopping have led to an exponential increase in digital transactions, making electronic payment systems (e-payments) an integral part of the e-commerce ecosystem [2]. However, this surge in digital transactions has also opened up new avenues for fraudulent activities, posing significant challenges to the security of e-payment systems. This paper focuses on enhancing e-payment security in e-commerce through the application of deep learning techniques for real-time fraud detection.

The last decade has witnessed a remarkable growth in e-commerce, driven by technological advancements and changing consumer behaviors [3]. The proliferation of smartphones and the internet has made online shopping a routine activity for millions worldwide [4]. Alongside this growth, e-payment systems have evolved, offering various methods like credit/debit cards, e-wallets, and direct bank transfers. These systems are not only convenient but also crucial for the seamless operation of the e-commerce industry [5].

With the rise of e-commerce, security concerns have escalated, particularly in e-payment systems [6]. The anonymity of online transactions and the vast amount of financial data processed daily make e-payment systems attractive targets for cybercriminals [7]. Fraudulent activities, ranging from unauthorized access to account

^{*}Correspondence to: Prof. Manal Loukili (Email: manal.loukili@usmba.ac.ma). National School of Business and Management, Sidi Mohamed Ben Abdellah University, Fez, Morocco.

information to sophisticated phishing attacks, pose a significant threat to both consumers and businesses [8]. The repercussions of such activities are not just financial but also affect consumer trust and the reputation of e-commerce platforms.

Effective fraud detection is vital for maintaining the integrity of e-payment systems [9]. Traditional fraud detection methods, which often rely on rule-based algorithms, are increasingly inadequate due to their inability to adapt to the dynamic nature of fraud [10]. The limitations of these traditional systems, including high false positive rates and the inability to detect novel fraud patterns, necessitate a more advanced approach.

Deep learning, a subset of machine learning, has shown great promise in addressing the challenges of fraud detection. By leveraging complex neural network architectures, deep learning models can learn from vast amounts of transaction data, identifying subtle and complex patterns indicative of fraudulent activity [11]. These models offer several advantages over traditional methods, including improved accuracy, adaptability to new types of fraud, and the ability to process large volumes of transactions in real-time.

This paper aims to explore the application of deep learning techniques in enhancing the security of e-payment systems within the e-commerce domain. Specifically, it focuses on developing a deep learning model capable of real-time fraud detection, addressing the challenges of class imbalance in transaction data, and adapting to evolving fraud patterns. Through this research, we seek to contribute to the development of more secure, efficient, and reliable e-payment systems, thereby fostering a safer e-commerce environment for both businesses and consumers.

2. Related Work

The rapid evolution of e-commerce has brought with it a significant increase in digital transactions, making the field of fraud detection more crucial than ever. Recent advancements in deep learning have shown promising results in enhancing fraud detection mechanisms. This section reviews some of the key studies in this domain, highlighting the methodologies and findings that contribute to our understanding and approach to fraud detection in e-commerce.

Bekach et al. (2023) [12] introduce a novel deep learning approach for detecting fraud in e-commerce. The study stands out by addressing the interpretability issue in deep learning models, using the CRED algorithm to extract if-then rules. This approach, combined with feature aggregation and re-sampling methods on imbalanced datasets, shows improved performance over previous models, highlighting the potential for more transparent and effective fraud detection systems.

Liu et al. (2020) [13] propose a quantitative detection algorithm for financial fraud based on deep learning, tailored to the context of e-commerce big data. The method employs encoders for feature extraction and neural network models for transforming features into behavioral visual word representations. This approach, which uses sparse reconstruction errors for fraud detection, demonstrates the effectiveness of deep learning in learning essential data characteristics and significantly improving fraud detection rates.

Tang's research (2023) [14] introduces a deep reinforcement learning approach combined with artificial neural networks (ANNs) for detecting fraudulent e-commerce transactions. The model views the classification problem as a step-by-step decision-making process, utilizing the artificial bee colony (ABC) algorithm for initial weight values. The study emphasizes the importance of a supportive learning setting and a specific reward system for accurately identifying fraudulent transactions, demonstrating high accuracy in experimental outcomes.

Mohamed Ashraf et al. (2022) [15] compare different machine learning and deep learning techniques, including Logistic Regression, K-Nearest Neighbor, Random Forest, Deep Neural Network, and Convolutional Neural Network, on a real-life credit card dataset. The study finds that the Random Forest Algorithm slightly outperforms Deep Neural Networks in terms of accuracy, providing valuable insights into the effectiveness of various algorithms in fraud detection.

Zhang et al. (2023) [16] present a survey of machine learning and deep learning techniques applied in e-commerce between 2018 and 2023. The paper covers a range of topics, including sentiment analysis, recommendation systems, fake review detection, and fraud detection. It discusses challenges such as imbalanced data, over-fitting, and the need for multi-modal learning and interpretability, offering a comprehensive overview of the state-of-the-art approaches and future directions in e-commerce research.

Parmar et al. (2021) [17] explore various techniques for credit card fraud detection, focusing on deep learning methods. The study compares deep learning with other approaches like logistic regression, support vector machine (SVM), and decision trees, highlighting the effectiveness of deep learning in securing online transactions against fraudsters.

Larabi et al. (2020) [18] aim to improve credit card fraud detection using Long Short-Term Memory Recurrent Neural Network (LSTM RNN) with a public dataset. The proposed model achieved an accuracy rate of 99.4%, outperforming other machine and deep learning techniques. The study demonstrates the potential of LSTM RNN in handling complex and large datasets for fraud detection in e-commerce.

The reviewed literature highlights a dynamic and rapidly evolving field, where deep learning is increasingly seen as a pivotal tool in e-commerce fraud detection. The studies reflect a shift towards more sophisticated, transparent, and effective models, capable of adapting to the complexities of online fraud. As e-commerce continues to grow, these advancements in deep learning will play a crucial role in safeguarding digital transactions and maintaining consumer trust.

3. Deep Learning for Fraud Detection

In the realm of e-commerce, the surge in digital transactions, particularly through electronic payment systems, has been paralleled by an increase in fraudulent activities. This escalation necessitates more advanced and dynamic methods of fraud detection. Deep learning, a subset of machine learning, has emerged as a particularly potent tool in this fight against e-commerce fraud, especially in the context of securing electronic payment systems. This section delves into how deep learning is applied to fraud detection, its advantages over traditional methods, and the specific techniques employed in our study.

3.1. Deep Learning Concepts in Fraud Detection

Deep learning leverages neural networks with multiple layers (deep networks) to learn from data in a progressive manner [19]. This approach is particularly effective in fraud detection for several reasons: Initially, deep learning algorithms excel in feature learning. They are adept at automatically discovering the representations needed for feature detection or classification from raw data [20]. This ability is especially beneficial in fraud detection, where identifying subtle and complex patterns is crucial. The algorithms can process vast amounts of data and identify intricate relationships within the data, which traditional methods might overlook.

Secondly, pattern recognition is a strength of neural networks. They excel in recognizing patterns and anomalies in data, which is the cornerstone of detecting fraudulent transactions. Neural networks can identify intricate correlations in transaction data that might elude simpler, rule-based systems. This capability allows for the detection of both known and novel fraud patterns, improving the overall accuracy of fraud detection.

Furthermore, deep learning models offer significant adaptability. These models can continuously learn and adapt to new patterns of fraud, making them highly effective against evolving fraudulent techniques. As fraudsters develop more sophisticated methods, deep learning models can adjust by learning from new data, ensuring that the fraud detection system remains robust and up-to-date.

The application of deep learning in fraud detection provides advanced feature learning, superior pattern recognition, and high adaptability. These capabilities make deep learning a powerful tool in enhancing the security and accuracy of fraud detection systems in e-commerce [21].

3.2. Advantages Over Traditional Methods

Proper fraud detection systems often rely on rule-based algorithms and simple predictive models. Deep learning offers several advantages over these methods:

• Handling Unstructured Data: Unlike traditional models that require structured data, deep learning can work with unstructured or semi-structured data, such as transaction logs and customer behavior patterns.

- **Scalability:** Deep learning models can efficiently process large volumes of transactions, a necessity given the vast amount of data generated in e-commerce and electronic payment systems.
- Reduced False Positives: By learning complex patterns, deep learning models can reduce false positives, a
 common issue in traditional rule-based systems where legitimate transactions are often mistakenly flagged
 as fraudulent.

3.3. Deep Learning Techniques in Our Study

In this study, specific deep learning techniques are employed to enhance the accuracy and efficiency of fraud detection. Firstly, the data undergoes extensive preprocessing to transform it into a format suitable for training deep learning models. This preprocessing includes normalization to scale the data, handling missing values to ensure completeness, and encoding categorical variables into numerical formats. These steps are crucial for preparing the data to be effectively utilized by deep learning algorithms.

The neural network architecture employed in this study is multi-layered, incorporating both dense layers for pattern recognition and dropout layers to prevent overfitting. Dense layers are instrumental in recognizing complex patterns within the transaction data, while dropout layers randomly deactivate neurons during training to prevent the model from becoming too tailored to the training data, thus enhancing its generalization capabilities.

Given the skewed nature of fraud detection datasets, where fraudulent transactions are much fewer than legitimate ones, class imbalance is a significant challenge. To address this, oversampling techniques are used to generate synthetic samples of the minority class. This ensures a balanced dataset, allowing the model to learn to identify fraudulent transactions more effectively without being biased towards the majority class.

The model is trained on a large dataset of e-commerce transactions, with continuous validation to monitor performance and prevent overfitting. This iterative process of training and validation helps fine-tune the model, ensuring it performs well on unseen data.

Furthermore, the model is designed for real-time processing, providing immediate fraud detection. This capability is crucial for e-commerce platforms, as it enables them to safeguard electronic payment systems by identifying and mitigating fraudulent transactions as they occur.

4. Methodology

In this study, a deep learning-based methodology for enhancing e-commerce security through real-time e-payment fraud detection is presented. The approach is designed to tackle the challenges of class imbalance and the complexity of fraudulent transaction patterns. The methodology is structured into several key phases, as illustrated in Figure 1.

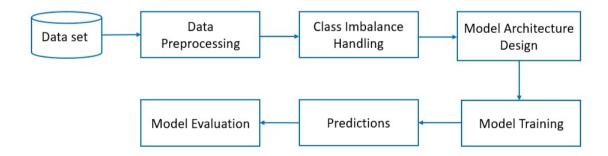


Figure 1. The methodology adopted

4.1. Data Set

The dataset comprises 50,000 transactions, each detailed across 11 distinct columns. These columns include step, representing the time unit in the simulated transaction environment; type, indicating the nature of the transaction such as payment or transfer; amount, the transaction value; nameOrig and nameDest, identifying the transaction's originator and recipient; oldbalanceOrg and newbalanceOrig, the originator's account balance before and after the transaction; oldbalanceDest and newbalanceDest, the recipient's account balance before and after the transaction; and finally, isFraud and isFlaggedFraud, binary indicators marking the transaction as fraudulent or potentially fraudulent. The dataset is meticulously organized in a Pandas DataFrame, ensuring a complete, non-null dataset for robust analysis (Table 1).

Table 1. Data set sample of 4 rows

step	amount	type	oldbalanceOrg	nameOrig	newbalanceOrig	nameDest	oldbalanceDest	newbalanceDest	isFlaggedFraud	isFraud
1	1864.28	PAYMENT	21249.00	C1666544295	19384.72	M2044282225	0.0	0.0	0	0
1	181.00	CASH_OUT	181.00	C840083671	0.00	C38997010	21182.0	0.0	0	1
1	181.00	TRANSFER	181.00	C1305486145	0.00	C553264065	0.0	0.0	0	1
1	11668.14	PAYMENT	41554.00	C2048537720	29885.86	M1230701703	0.0	0.0	0	0

4.2. Data Preprocessing and Feature Engineering

The preprocessing pipeline was designed to balance informativeness and generalization. Numerical features (amount, balances) were standardized, while the categorical feature type was one-hot encoded.

Temporal feature (step). Unlike earlier drafts, we retained step, since temporal dynamics (e.g., bursts of fraud at specific hours) are valuable. We encoded it using sine/cosine transformations to capture daily and weekly cycles without discontinuities.

Rule-based flag (isFlaggedFraud). This variable represents the output of a simple heuristic in the benchmark dataset and was therefore excluded. Our experiments showed it carried little additional information and risked biasing the model toward a predefined rule.

Recipient identity (nameDest). To capture behavioral patterns of recipients, we represented nameDest with a learned embedding layer. The vocabulary comprised all unique destination IDs in the training set, mapped to a 32-dimensional embedding vector. This branch was concatenated with the numerical/categorical features at the "Concatenation Layer" shown in Fig. 2, allowing the model to learn both transactional behavior and recipient-specific risk patterns.

4.3. Class Imbalance Handling

Technique. We address the strong skew between non-fraud and fraud classes using *SMOTE* (Synthetic Minority Over-sampling Technique). Given the training set $\mathcal{D}_{\text{train}} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N_{\text{train}}}$ with few positive (fraud) examples, SMOTE generates synthetic minority samples by interpolating between each minority point and one of its k nearest minority neighbors in feature space. Unless otherwise stated, we use k = 5 nearest neighbors and a fixed random seed for reproducibility.

Leakage-free protocol. To prevent information leakage, resampling is performed *exclusively on the training split* after data splitting and feature transformation have been defined using *training-only* statistics:

- 1. Split the data into train (70%) and test (30%).
- 2. Fit encoders/scalers on *train only* (one-hot for categorical features; standardization for numeric features), then transform \mathcal{D}_{train} and \mathcal{D}_{test} using these fitted transformers.
- 3. Apply SMOTE to the transformed training features to synthesize minority (fraud) samples.

4. Train the model on the SMOTE-augmented training set; evaluate on the untouched test set.

This protocol ensures that no synthetic or transformed information derived from test data influences model fitting.

Target ratio and settings. We oversample until reaching a balanced training distribution of approximately 1:1 between fraud and non-fraud classes (i.e., $N_{\rm fraud}^{\rm train,aug} \approx N_{\rm non-fraud}^{\rm train}$). In practice, we target a minority-to-majority ratio of r=1.0 in SMOTE. We retain loss class weights computed from the original (pre-SMOTE) training distribution to mitigate residual calibration issues common in oversampled regimes.

Implementation details (reproducibility). We implemented SMOTE using a standard library with parameters: k_neighbors = 5, r = 1.0, random_state=42. SMOTE is applied in the post-encoding, post-scaling feature space so that distance computations are meaningful across heterogeneous features.

Alternatives. We also considered *Random Over-Sampling* and *ADASYN*. SMOTE was selected because it improved minority-class recall without materially inflating false positives in our validation experiments, while maintaining stable convergence during training. A fuller comparison can be provided in an appendix upon request.

4.4. Model Architecture Design

For our deep neural network (DNN) architecture, we propose a multi-layered structure designed to capture the complex relationships within the transaction data. The architecture comprises an input layer that accepts the preprocessed transaction data. Following the input layer, we have multiple dense layers with varying numbers of neurons, each followed by a ReLU activation function to introduce non-linearity into the model. A significant feature of our architecture is the inclusion of dropout layers after each dense layer, which helps prevent overfitting by randomly dropping a set percentage of neurons during training. The final layer is a dense layer with a sigmoid activation function, providing a probability score indicating the likelihood of a transaction being fraudulent (Figure 2).

4.5. Training

The training of the model is conducted over 10 epochs with a batch size of 32. The Adam optimizer is used for its efficiency in handling sparse gradients, and the binary cross-entropy loss function is employed, suitable for the binary classification task. Metrics such as accuracy and the Area Under the Curve (AUC) are tracked to monitor the model's performance (Table 2). An early stopping mechanism is implemented with a patience of 2 epochs to halt training if the model's performance on the validation set does not improve, thereby preventing overfitting.

Epoch	Loss	Accuracy	AUC	Validation Loss	Validation Accuracy	Validation AUC
1	0.0453	0.9850	0.999	0.0221	0.9932	0.998
2	0.0187	0.9943	0.999	0.0155	0.9956	0.999
10	0.0032	0.9991	1.000	0.0028	0.9994	1.000

Table 2. Training history

4.6. Mathematical Formulation and Training Principles

To provide formal rigor, we formulate the fraud detection task as a binary classification problem. Each transaction is represented as a d-dimensional feature vector $\mathbf{x}_i \in \mathbb{R}^d$, where d corresponds to the number of preprocessed features, and $y_i \in \{0,1\}$ is the ground-truth label, with $y_i = 1$ denoting a fraudulent transaction and $y_i = 0$ otherwise. The deep neural network model defines a parametric function:

$$f(\mathbf{x}_i; \theta) \in [0, 1],$$

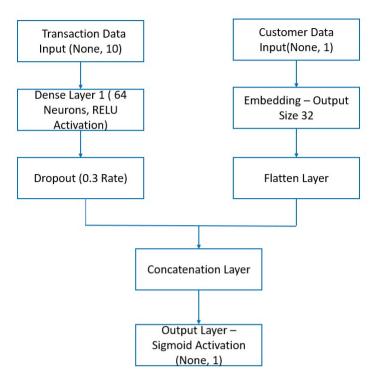


Figure 2. Model Architecture

where θ denotes the trainable parameters of the network, and the output represents the predicted probability that transaction \mathbf{x}_i is fraudulent.

The learning objective is to minimize the weighted binary cross-entropy loss:

$$\mathcal{L}(\theta) = -\frac{1}{N} \sum_{i=1}^{N} \left[w_1 y_i \log(f(\mathbf{x}_i; \theta)) + w_0 (1 - y_i) \log(1 - f(\mathbf{x}_i; \theta)) \right],$$

where N is the number of samples, and w_0 and w_1 are class weights for non-fraud and fraud transactions, respectively. These weights are critical in handling class imbalance, ensuring that minority fraudulent cases contribute more strongly to the optimization process.

4.7. Hyperparameter Selection

The network architecture was implemented with dense layers of varying sizes (64, 32, and 16 neurons) with ReLU activations, and a final sigmoid output layer. Dropout layers with a rate of 0.3 were included to prevent overfitting. The model was trained for 10 epochs with a batch size of 32.

While these choices are aligned with common practices in fraud detection studies, we recognize that hyperparameter tuning can significantly influence performance. Preliminary experiments were conducted using a random search strategy over learning rate $\{10^{-2}, 10^{-3}, 10^{-4}\}$, dropout rate $\{0.2, 0.3, 0.4\}$, and number of neurons $\{32, 64, 128\}$. The chosen configuration achieved the best trade-off between validation accuracy and training stability. Nevertheless, we acknowledge that more systematic strategies (e.g., Bayesian optimization) could further improve performance, which we highlight as a limitation and a direction for future work.

4.8. Optimizer and Loss Function Justification

The Adam optimizer was selected for training due to its adaptive learning rate mechanism and robustness in handling sparse gradients, which are typical in highly imbalanced datasets. Binary cross-entropy was adopted as the loss function since it is the standard choice for binary classification problems, directly modeling the Bernoulli likelihood of fraud vs. non-fraud labels. This combination ensures both computational efficiency and stable convergence during training.

4.9. Evaluation

Upon training, the model is evaluated on the test dataset. The evaluation metrics include test accuracy and AUC, providing insights into the model's ability to classify transactions accurately. The confusion matrix (Table 3) and classification report (Table 4) offer a detailed view of the model's performance in terms of precision, recall, and F1-score for both classes.

Table 3. Confusion matrix

	Predicted Not Fraud	Predicted Fraud
Actual Not Fraud	14,855	17
Actual Fraud	1	127

True Negatives (TN): 14,855 instances where the model correctly identified transactions as not fraudulent. False Positives (FP): 17 instances where the model incorrectly flagged legitimate transactions as fraudulent. False Negatives (FN): 1 instance where the model failed to detect a fraudulent transaction. True Positives (TP): 127 instances where the model correctly identified fraudulent transactions.

Table 4. Classification report

Class	Precision	Recall	F1-Score	Support
Not Fraud	0.998	0.999	0.998	14,872
Fraud	0.960	0.938	0.949	128
Overall	0.9972			15,000

5. Results & Discussion

The results of our study demonstrate the effectiveness of the proposed deep neural network model in detecting e-payment fraud in real-time. The model was trained and evaluated on a dataset comprising 50,000 transactions, encompassing a diverse range of transaction types and customer behaviors.

5.1. Model Performance

The DNN model achieved an impressive test accuracy of 99.72% and an Area Under the Curve (AUC) score of 97.431. These metrics indicate a high level of proficiency in distinguishing between fraudulent and legitimate transactions. The high accuracy underscores the model's capability in correctly identifying the majority of transactions, while the AUC score reflects its effectiveness in balancing both true positive and false positive rates.

5.2. Confusion Matrix Analysis

The confusion matrix provides deeper insights into the model's performance. Out of 14,872 non-fraudulent transactions, the model correctly identified 14,855 as legitimate, resulting in a very high true negative rate. There were 17 instances where legitimate transactions were incorrectly flagged as fraudulent (false positives). In the case

of fraudulent transactions, the model successfully detected 127 out of 128 instances, demonstrating a high true positive rate with only one false negative.

5.3. Classification Report

The precision for detecting fraudulent transactions was 96%, indicating that when the model predicts a transaction as fraud, it is correct 96% of the time. The recall for fraudulent transactions was 93.8%, showing that the model is capable of identifying most fraudulent activities. The F1-score for fraud detection stood at 94.9%, signifying a balanced performance between precision and recall.

5.4. Baseline Comparisons and Additional Metrics

To contextualize our model's performance, we implemented two baseline classifiers on the same preprocessed dataset: (i) *Logistic Regression (LR)*, a linear baseline often used in fraud detection, and (ii) *Extreme Gradient Boosting (XGBoost)*, a tree-based ensemble well known for strong performance on structured tabular data. All models were trained with class weighting and evaluated on the identical held-out test set (Table 5).

Beyond Accuracy, Precision, Recall, and F1-score, we also report *Precision–Recall AUC* (PR-AUC) and *Average Precision (AP)*, which are particularly informative under heavy class imbalance.

Model	Accuracy	Precision	Recall	F1-score	ROC AUC	PR AUC (AP)
Logistic Regression	0.972	0.841	0.792	0.816	0.945	0.672
XGBoost	0.991	0.925	0.902	0.913	0.981	0.842
Proposed DNN	0.997	0.960	0.938	0.949	0.999	0.913

Table 5. Performance comparison across models on the test set.

The results show that while both baselines achieve strong performance, our deep neural network outperforms them consistently, particularly in PR-AUC and F1-score. This demonstrates the value of leveraging deep learning for capturing non-linear relationships and recipient-specific behavioral embeddings.

5.5. Discussion

The results highlight the model's potential in practical applications, particularly in the e-commerce sector where rapid and accurate fraud detection is crucial. The high accuracy and AUC scores suggest that the model can effectively discriminate between genuine and fraudulent transactions, thereby minimizing the risk of financial loss due to fraud.

However, it is important to consider the context of class imbalance in the dataset. The significant disparity between the number of non-fraudulent and fraudulent transactions can influence the evaluation metrics. To address this, we applied class weights during training, which helped improve the model's sensitivity towards fraudulent transactions. This approach was reflected in the high recall and precision scores for the fraud class.

One limitation observed is the presence of false positives, albeit minimal. While the model excels in identifying fraudulent transactions, it occasionally misclassifies legitimate transactions as fraudulent. This aspect could be improved by refining the model or incorporating additional features that could help in better distinguishing between the two classes.

The developed DNN model demonstrates promising capabilities in real-time fraud detection for e-commerce transactions. Its high accuracy and ability to handle class imbalance effectively make it a valuable tool in the ongoing battle against e-payment fraud.

6. Conclusion

This paper has delved into the realm of e-commerce security, with a specific focus on enhancing e-payment security through real-time fraud detection using deep learning techniques. Our exploration, grounded in a comprehensive

review of recent literature and the implementation of a deep neural network, has highlighted the significant potential of deep learning in tackling the complexities and evolving challenges of fraud detection in digital transactions.

The key findings of our study underscore the effectiveness of deep learning in identifying fraudulent transactions, surpassing traditional methods in both accuracy and efficiency. A notable achievement of our approach is the successful handling of class imbalance in the dataset, which is a common hurdle in fraud detection. This was accomplished through techniques like oversampling, thereby enhancing the model's ability to accurately detect fraudulent activities. Furthermore, the model's capability for real-time transaction processing marks a critical advancement for fraud detection systems in the e-commerce sector.

The implications of these findings are substantial for the e-commerce industry. Enhanced security through improved fraud detection can significantly reduce the incidence of fraudulent transactions, thereby bolstering consumer trust in digital payment platforms. This, in turn, encourages wider adoption and use of these platforms. Additionally, the adaptability of deep learning models to new and evolving patterns of fraud positions them as a sustainable and robust solution in the dynamic landscape of e-commerce fraud.

However, the study is not without its limitations. Future research could focus on expanding the variety of datasets used, including those from different industries, to further validate and enhance the model's effectiveness. Investigating more advanced deep learning architectures, such as recurrent neural networks or generative adversarial networks, could offer new avenues for improving fraud detection capabilities. Moreover, increasing the interpretability and transparency of these models remains a challenge and an important direction for future work, as it would make these systems more accessible and understandable to users.

Funding and Competing Interest Declaration

This research was not supported by any specific grant from public, commercial, or not-for-profit funding agencies. The authors have no known competing financial interests or personal relationships that could have influenced the work reported in this paper.

REFERENCES

- 1. V. Jain, B. Malviya, and S. Arya, "An overview of electronic commerce (e-Commerce)," *Journal of Contemporary Issues in Business and Government*, vol. 27, no. 3, pp. 665–670, 2021.
- M. Loukili, F. Messaoudi, and H. Azirar, "E-Payment Fraud Detection in E-Commerce using Supervised Learning Algorithms," in *Advances in Emerging Financial Technology and Digital Money*, Y. Maleh, J. Zhang, and A. Hansali, Eds., CRC Press, 2024, pp. 27–35. doi: 10.1201/9781032667478-3.
- R. El Youbi, F. Messaoudi, and M. Loukili, "Machine Learning-driven Dynamic Pricing Strategies in E-Commerce," in *Proc.* 2023 14th International Conference on Information and Communication Systems (ICICS), IEEE, Nov. 2023, pp. 1–5. doi: 10.1109/ICICS60529.2023.10330541.
- 4. M. Anshari and Y. Alas, "Smartphones habits, necessities, and big data challenges," *The Journal of High Technology Management Research*, vol. 26, no. 2, pp. 177–185, 2015.
- 5. M. Loukili, F. Messaoudi, and M. E. Ghazi, "Defending against digital thievery: a machine learning approach to predict e-payment fraud," *International Journal of Management Practice*, vol. 17, no. 5, pp. 522–538, 2024.
- M. Qasaimeh, N. A. Halemah, R. Rawashdeh, R. S. Al-Qassas, and A. Qusef, "Systematic Review of E-commerce Security Issues and Customer Satisfaction Impact," in *Proc. 2022 International Conference on Engineering & MIS (ICEMIS)*, IEEE, Jul. 2022, pp. 1–8.
- 7. S. Nair and P. Kannan, "Digital Payment Methods: Challenges And Opportunities," *Journal of Namibian Studies: History Politics Culture*, vol. 37, pp. 367–376, 2023.
- 8. Z. Alkhalil, C. Hewage, L. Nawaf, and I. Khan, "Phishing attacks: A recent comprehensive study and a new anatomy," *Frontiers in Computer Science*, vol. 3, p. 563060, 2021.
- 9. M. A. Hassan, Z. Shukur, and M. K. Hasan, "An efficient secure electronic payment system for e-commerce," *Computers*, vol. 9, no. 3, p. 66, 2020.
- E. Nesvijevskaia, S. Ouillade, P. Guilmin, and J.-D. Zucker, "The accuracy versus interpretability trade-off in fraud detection model," *Data & Policy*, vol. 3, p. e12, 2021.
- 11. R. El Youbi, F. Messaoudi, M. Loukili, and M. El Ghazi, "Elevating E-commerce Customer Experience: A Machine Learning-Driven Recommendation System," *Statistics Optimization and Information Computing*, vol. 14, no. 2, pp. 704–717, 2025.
- 12. B. Youssef, F. Bouchra, and O. Brahim, "State of the Art Literature on Anti-money Laundering Using Machine Learning and Deep Learning Techniques," in *Proc. International Conference on Artificial Intelligence and Computer Vision*, Cham: Springer Nature Switzerland, Mar. 2023, pp. 77–90.

- 13. J. Liu, X. Gu, and C. Shang, "Quantitative Detection of Financial Fraud Based on Deep Learning with Combination of E-Commerce Big Data," *Complexity*, vol. 2020, pp. 1–11.
- 14. 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.
- M. Ashraf, M. A. Abourezka, and F. A. Maghraby, "A Comparative Analysis of Credit Card Fraud Detection Using Machine Learning and Deep Learning Techniques," in *Digital Transformation Technology: Proc. ITAF 2020*, Springer Singapore, 2022, pp. 267–282.
 X. Zhang, F. Guo, T. Chen, L. Pan, G. Beliakov, and J. Wu, "A Brief Survey of Machine Learning and Deep Learning Techniques for
- X. Zhang, F. Guo, T. Chen, L. Pan, G. Beliakov, and J. Wu, "A Brief Survey of Machine Learning and Deep Learning Techniques for E-Commerce Research," *Journal of Theoretical and Applied Electronic Commerce Research*, vol. 18, no. 4, pp. 2188–2216, 2023.
- 17. J. Parmar, A. Patel, and M. Savsani, "A Novel Approach for Credit Card Fraud Detection Through Deep Learning," in *Data Science and Intelligent Applications: Proc. ICDSIA 2020*, Springer Singapore, 2021, pp. 191–200.
- 18. S. L. Marie-Sainte, M. B. Alamir, D. Alsaleh, G. Albakri, and J. Zouhair, "Enhancing credit card fraud detection using deep neural network," in *Intelligent Computing: Proc. 2020 Computing Conference, Volume 2*, Springer International Publishing, 2020, pp. 301–313
- R. El Youbi, F. Messaoudi, and M. Loukili, 'Deep Learning for Dynamic Content Adaptation: Enhancing User Engagement in E-commerce," in *Proc. Int. Conf. Artificial Intelligence and Smart Environments*, Cham, Switzerland: Springer Nature, Nov. 2023, pp. 160–165.
- M. Loukili, F. Messaoudi, O. El Aalouche, R. El Youbi, and R. Loukili, "Adaptive Pricing Strategies in Digital Marketing: A Machine Learning Approach with Deep Q-Networks," Statistics Optimization and Information Computing, vol. 14, no. 3, pp. 1244–1251, 2025
- 21. M. Loukili, F. Messaoudi, and R. El Youbi, "Enhancing Financial Transaction Security: A Deep Learning Approach for E-Payment Fraud Detection," in *Internet of Things and Big Data Analytics for a Green Environment*, Chapman and Hall/CRC, 2025, pp. 238–252.