

Deep transformer architectures for real-time e-payment fraud detection under class imbalance

Riad Loukili¹, Fayçal Messaoudi², Raouya El Youbi¹, Manal Loukili²

¹Department of Computer Science, National School of Applied Sciences, Sidi Mohamed Ben Abdellah University, Fez, Morocco

²Department of Computer Science, National School of Business and Management,
Sidi Mohamed Ben Abdellah University, Fez, Morocco

Article Info

Article history:

Received Jun 1, 2024

Revised Jun 3, 2026

Accepted Jul 19, 2026

Keywords:

Class imbalance

Deep learning

E-commerce security

Explainable AI

Federated learning

Real-time fraud detection

Transformer architecture

ABSTRACT

E-commerce has become a rapidly expanding channel, leading to more e-commerce payment transactions and an increased threat for advanced fraudulent activities. New frauds include cross-channel attacks, adversarial behavior and synthetic identities, which are challenging to rule-based detection. In this paper, a deep learning framework for detecting e-payment fraud in real-time is proposed. The model integrates self-attention, explainable AI, class balancing using GANs, and a transformer-based architecture to enable transparent decision-making. The solution tackles major issues in fraud detection, such as class imbalance, evolving transaction trends, and the requirement for interpretability in compliance with regulations like the EU AI Act. Experiments on benchmark datasets demonstrate good results, with an accuracy of 99.82%, and an AUC-ROC of 99.56%. Based on the results, the use of transformer-based models with the help of synthetic data generation and explainability can enhance the reliability and transparency of e-payment fraud detection systems.

This is an open access article under the [CC BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license.



Corresponding Author:

Manal Loukili

Department of Computer Science, National School of Business and Management

Sidi Mohamed Ben Abdellah University

Fez, Morocco

Email: manal.loukili@usmba.ac.ma

1. INTRODUCTION

The e-commerce domain is a part of the international economy and has revolutionized the way goods and services are purchased by people [1]. According to this growth, the electronic payment systems are now essential to conduct business online [2]. This extension has additionally grown the amount of exposure to fraud, particularly with payment systems that incorporate massive amounts of financial info on a daily basis [3]. According to a report by ITD and the National Retail Foundation, global e-commerce sales will keep rising, and by 2025, the total amount of sales is projected to hit more than \$7.5 trillion and could hit over \$9 trillion by 2027. Meanwhile, the loss from electronic payment fraud is estimated to reach around USD 48 billion globally by 2025 [4] and further escalate with the expansion of electronic payment in the coming years. The use of AI-driven services in e-commerce and the growing internet use and smartphone penetration have facilitated the expansion of the ecommerce market [5], [6]. The payment methods also expand with the use of credit/debit cards, digital wallets, deferred payments, cryptocurrency payments, and instant bank transfers [7]. These options make it easier for the user, but also provide more opportunities for an attacker to exploit a system.

Cyber criminals are interested in the e-payment systems, due to the adoption of a fast, partially anonymous and unpredictable payment system, which is not easily verifiable in real-time [8], [9]. Stolen

cards and phishing attacks are not all of fraud. Some of the ongoing schemes are synthetic identity fraud, credential stuffing, mobile payment API attacks, and money laundering via cryptocurrency mixing services [10]. What are other impacts of these activities as well as a direct financial impact? Fraud has the potential to damage customer trust, harm the reputation of online platforms, and raise compliance concerns, especially in the context of the European Union AI Act, which places a greater emphasis on transparency and human oversight in financial decision-making processes [11].

Therefore, the prevention of fraud is a part of e-payment security [12]. While rule-based systems can still be effective, these are too static to be responsive to fraud scheme changes. While financial AI systems can be subject to a lot of false positives, may not be aware of the latest attack strategies, and may only explain a limited part of the decision [13]. These requirements are increasingly crucial in a regulatory environment. The EU AI Act (to be implemented from August 2024 and fully in force by August 2026) further strengthens the requirements to use AI in sensitive financial situations [14].

As fraud schemes become more complex, and change over time, deep learning can be an effective tool to combat this difficulty [15]. Transformer models can capture the long-term dependence between the sequence of transactions, and detect suspicious transactions that may not appear in the individual transactions [16]. Graph Neural Networks can model the relations between users, accounts, merchants, and transactions, and transactions, making them useful in detecting coordinated fraud [17]. Another method is federated learning, where any institution can train their own model without sharing the data about its customers [18]. Meanwhile, several techniques have emerged for gaining a better understanding of decisions used for fraud detection such as SHAP, LIME, and attention-based interpretation [19].

E-payment fraud in e-commerce is one area where deep learning can be applied for fraud detection. We have presented a model based on a transformer to learn transactions, a generative adversarial network to address class imbalance, and explainable AI to improve transparency and compliance.

2. RELATED WORK

Payment fraud is becoming increasingly popular, but still too difficult to detect [20]. Recent research has focused on models based on deep learning that can learn transactions, deal with imbalanced data and explain the decisions made by the model. There are three directions of research in e-payment fraud detection that are in particular need of attention: transformer-based architectures; models based on class imbalance and explainable AI.

2.1. Transformer-based architectures

Financial fraud detection relies increasingly on transformer models because they can capture the dependence between transactions without the need of recurrent hidden states. These models analyze transaction data with a variety of aspects, such as transaction amount, timestamp, merchant type, and device information, to identify patterns of fraudulent activity [21].

Self-attention will place different weights on different features or transactions in order to predict fraud. Attention scores can also be used to assist in interpretation, displaying what portions of the transaction history caused the fraud detection. According to the EU AI Act for systems that influence users or require transparency [14], transformer-based models outperform traditional machine learning methods and deep learning models like feed-forward networks and LSTMs in fraud detection tasks [22], [23]. Their ability to handle long-range dependencies makes them effective in spotting fraud patterns that happen over time, like account takeover and multi-stage money laundering.

2.2. Generative models for class imbalance

Class imbalance remains one of the main difficulties in fraud detection, since fraudulent transactions usually represent only a very small proportion of the data. In some datasets, fraud cases account for less than 0.1% of all transactions. Earlier approaches often relied on simple oversampling, but recent generative models provide more flexible solutions for creating synthetic minority-class examples [24]-[26].

GAN-based methods generate synthetic fraud samples through an adversarial process involving a generator and a discriminator. More recent diffusion-based approaches generate data through gradual denoising, which can improve training stability and reduce problems such as mode collapse [25], [26]. This is important in fraud detection because models must learn different fraud behaviors rather than overfit to a narrow set of examples. Yin and Zhang proposed GA-GAN, which uses gated attention to better represent the non-stationary and heterogeneous nature of fraudulent transactions [24]. By weighting important features during generation, GA-GAN can produce synthetic fraud samples that are closer to real fraudulent behavior and improve classifier performance.

2.3. Explainable AI integration

In the financial sector where the use of AI systems is gaining traction, the need for explainability has grown increasingly significant in financial fraud detection. In 2024, the EU AI Act was adopted that sets requirements for the integrity of AI systems, risk management, data governance, transparency, human oversight, and conformity assessment of AI systems [11], [14]. In this context, fraud detection models have to be accurate as well as understandable and auditable.

SHAP is widely used in financial fraud detection because it provides feature-level explanations grounded in cooperative game theory. Analysis of feature selection during prediction allows analysts to see what elements contribute to each prediction. Understanding how each feature contributes to each prediction improves interpretability of model behavior and decision making [19]. In addition to post hoc explanation techniques, some researchers have attempted to design models that are intrinsically interpretable. For example, Transformer-BDI combines transformer-based prediction models with symbolic reasoning to produce a decision trail that is easy to understand and maintains high detection rates [21]. TabNet also utilizes sequential attention in its prediction to choose features that will be used in each prediction, which can be easily understood by analysts [22]. As banks demand high predictive performance and compliance, explainable approaches are increasingly important to provide a reliable system for fraud detection.

3. METHOD

In this study, a deep learning framework was introduced to detect e-payment frauds in real-time. It includes three parts: the transformer-based model with multi-head self-attention, synthetic data generation based on GA-GAN to tackle class imbalance, and explainability based on SHAP to enable explainable decisions. The general procedure followed is shown in Figure 1.




Figure 1. Proposed deep learning framework

3.1. Dataset

The main experiments use the PaySim financial transaction dataset, which contains 50,000 transactions and 11 features. These include the transaction step, transaction type, amount, originator and recipient identifiers, initial and final balances for both parties, and the binary labels isFraud and isFlaggedFraud.

To assess generalization, the model is also evaluated on the European Credit Card Fraud Dataset, which includes 284,807 transactions and 30 features, and the IEEE-CIS Fraud Detection Dataset, which contains 561,013 training transactions and 67 features. These datasets reflect the strong imbalance typically observed in fraud detection, with fraud cases representing about 0.13% of the European Credit Card Dataset and 3.5% of the PaySim dataset [25].

3.2. Data preprocessing

The purpose of this step is to reduce noise and leakage risk and strengthen the model. Non-informative identifiers like step, nameOrig, isFlaggedFraud are removed. The transaction type is used as learned embeddings rather than one hot encoding to learn the relationship between transaction types. For numerical variables, we scale numerical variables using RobustScaler, which is less sensitive to outliers than standard normalization.

Other features are added to encode temporal and behavioral aspects. They encompass rolling statistics of transaction levels for 5 step and 24 step windows, time since last transaction per originator, balance change ratios and transaction velocity. These variables give additional context to the model to assist with the detection of unusual payment behavior.

The datasets are split in a stratified manner based on time, where 70% is used for training, 15% for validation and 15% for testing. This maintains the class proportions and maintains chronological order, minimizing future information leakage into the training and more accurately simulating the conditions of deployment [27].

3.3. Class imbalance handling

Class imbalance is regarded as a core aspect of the methodology, as transactions are unlikely to be fraudulent. The proposed method is based on GA-GAN synthetic data generation and cost-sensitive learning. In order to generate minority class samples that represent the structure of real fraudulent transactions, GA-GAN employs gated attention [24].

The generator creates fake fraud examples from random noise, and the discriminator is trained to identify genuine fraud examples from the generated examples. The gated attention mechanism helps the generator to focus on the most relevant features related to fraud. GA-GAN is trained for 100 epochs with Adam optimizer having learning rate of 0.0001 and synthetic samples are generated until the ratio of 1:5 (falsely labeled to correctly labeled) is achieved.

Focal loss helps further reduce the majority class dominance by using the focal loss function. It places greater emphasis on challenging cases and helps to minimise the impact of easy, legitimate transactions when training. The Loss is calculated as:

$$FL(PT) = -\alpha_t(1 - p_t)\gamma \log(p_t)$$

where p_t is the predicted probability for the actual class, α is to set the weights on the classes, and γ is to focus on hard examples. The optimum values of $\alpha = 0.25$ and $\gamma = 2.0$ are determined from the validation results.

3.4. Model architecture design

The proposed architecture is divided into three parts, which are illustrated in the Figure 2. Firstly, the feature embedding layer converts categorical features into learned embeddings and maps numerical features into a common representation space. This results in a series of feature tokens that can be fed into the transformer. Second, the transformer encoder has four layers, eight attention heads and a hidden dimension of 256. The multi-head self-attention, feedforward transformation, residual connection, and layer normalization are applied to each layer. The attention mechanism allows the model to learn interactions between transaction features. For interpretability, attention weights are averaged across heads and layers to obtain a global feature-importance signal. Finally, the classification head consists of three fully connected layers with dimensions [256, 128, 64, 1]. GELU activation and dropout with a rate of 0.3 are used before the final sigmoid output, which produces the fraud probability. The model contains about 2.8 million parameters, making it suitable for real-time inference while retaining enough capacity to learn complex fraud patterns.

3.5. Training

The model is trained using AdamW with a weight decay of 0.01 and an initial learning rate of $1e-4$. A cosine annealing schedule is applied over 50 epochs. The batch size is set to 256, and early stopping is used with a patience of seven epochs based on validation AUC. The training objective combines binary cross-entropy with focal loss weighting. Several regularization techniques are used to limit overfitting, including dropout in the classification head, stochastic depth in the transformer layers, and Mixup augmentation with $\alpha = 0.2$. Training is performed on an NVIDIA A100 GPU with 40 GB memory and takes approximately 45 minutes.

3.6. Evaluation

The model is evaluated using metrics suited to imbalanced classification: accuracy, precision, recall, F1-score, AUC-ROC, AUC-PR, and Matthews Correlation Coefficient. Balanced accuracy and the geometric mean of sensitivity and specificity are also reported to provide a more reliable view of performance under class imbalance. For these predictions to be understandable, SHAP values are calculated. We use bootstrap resampling to assess feature importance stability and the SHAP Stability Index to assess consistency of explanations. This is important as explanation instability may lead to mistrust in fraud detection systems especially when class-balancing is used [28].




Figure 2. Architecture of the proposed transformer-based model for fraud detection

4. RESULTS AND DISCUSSION

Results show that the model performs well in each of the three datasets, with the highest performance on PaySim.

4.1. Model performance

The model with a transformer has an accuracy of 99.82% and an AUC-ROC of 99.56% on the PaySim dataset. The AUC-PR is 98.87% due to biased data for fraud detection. Matthews correlation coefficient (MCC) was 0.961, indicating good agreement between the predicted and real labels.

The confusion matrix further confirms the stability of the model. It generated 14,865 true negatives, 126 true positives, 7 false positives and only 2 false negatives. This is where the model can predict successful transactions and fraudulent transactions, and minimize missed fraud cases and false fraud alerts.

The model showed good results on other datasets as well. It is competitive with recent results in literature [24] with an AUC-ROC of 0.984 and an AUC-PR of 0.892 on the European credit card fraud dataset. The model achieved an AUC-ROC of 0.974 on the IEEE-CIS dataset, indicating that the model can be generalized to other transaction environments and feature spaces.

4.2. Comparative analysis

Table 1 and Figure 3 present a comparison of the proposed model with recent fraud detection baselines. The proposed transformer model outperforms the compared methods in terms of AUC-ROC and F1-score and also provides explanations using SHAP values. TGNN also has good performance, particularly in transaction networks that have highly structured graphs, but its application is better suited to fraud scenarios involving a graph rather than regular e-commerce transactions. Interpretability is achieved in Transformer-BDI and XFL, while their predictive accuracy is lower in this comparison. GA-GAN and Diffusion-SDG are beneficial for synthetic data generation but are not meant to be complete end-to-end fraud detection tools.

Table 1. Model comparison to baseline models

Model	AUC-ROC	F1-Score	Explainable
Baseline DNN	0.9743	0.949	No
Transformer (Proposed)	0.9956	0.974	Yes (SHAP)
Transformer-BDI [21]	0.9500	0.920	Yes (BDI)
TGNN [29]	0.9924	0.961	Limited
GA-GAN [24]	0.9290	0.875	No
XFL [30]	0.9400	0.890	Yes (SHAP+LIME)
Diffusion-SDG [25]	0.9200	0.890	No




Figure 3. Performance comparison of the proposed model with the latest fraud detection benchmark

4.3. Discussion

The effectiveness of the proposed model can be attributed to the synergy among the three key components—transformer learning, synthetic data generation, and cost-sensitive training. The self-attention mechanism enables the model to better grasp the interactions among transaction features than traditional feedforward models. GA-GAN enhances the representation of the minority fraud cases by creating more realistic synthetic samples, and focal loss mitigates the impact of legitimate transactions during the learning process.

The SHAP analysis also helps in the model interpretability. The most influential features are in line with practical fraud analysis of transaction amount, balance change ratios and transaction velocity. An explanation is consistent across resampling experiments as indicated by a stability index of 0.94 for the SHAP [28]. The model is also appropriate for real time applications. It infers transactions in just over 3.2ms per transaction over a single CPU core, which is more than 300 transactions per second. This is the level of latency that is acceptable for most e-commerce payment platforms, where customers should not experience any delays in the payment process while the system detects fraud.

5. CONCLUSION

This paper introduces the deep learning technique for real-time e-payment fraud detection in e-commerce. It includes transformer-based learning, GA-GAN synthetic data generation, and SHAP explainability for fraud detection. Results from experiments show that the proposed model performed well with accuracy of 99.82% and an AUC of 99.56% in PaySim dataset. The models are able to detect fraudulent transactions with high accuracy and compensate for class imbalance by reducing focal loss and generating synthetic data. The low inference latency of the model highlights its use in real-time inference, such as e-commerce payment systems. On the ground, the proposed model can help secure digital payments, minimize financial losses, and increase consumer trust in online transactions. One of the biggest advantages of introducing the model is explainability, which helps the fraud analyst or decision-maker understand why a suspicious transaction. This is particularly important in the financial industry, where transparency, accountability, and regulatory compliance are becoming increasingly important.

Although these findings are encouraging, research needs to be done further on a number of areas. For example, LLM integration could be used to process unstructured data such as transaction descriptions, customer messages, or merchant reviews. Second, using graph neural networks, it is possible to capture relationships between users, accounts, merchants, and transactions, and thus identify coordinated fraud schemes. Third, federated learning could allow multiple financial institutions to jointly develop a fraud detection model without sharing customer information. Further research is needed on continuous learning to adapt the model to the evolving fraud pattern over time. Static models will no longer be useful if they are not continually updated with changing fraudulent activity. Fourth, it is important to incorporate explainability and constraints into the real-world such as documentation, human oversight, system scalability, and resistance to adversarial attacks.

ACKNOWLEDGMENTS

Riad Loukili would like to express his sincere gratitude to his supervisor, Prof. Fayçal Messaoudi, for his guidance and continued support throughout this research. He also extends his heartfelt thanks to his sister, Prof. Manal Loukili, for her encouragement, support, and valuable advice throughout his academic journey.

FUNDING INFORMATION

The authors state that no funding was received for this research.

AUTHOR CONTRIBUTIONS STATEMENT

This journal uses the Contributor Roles Taxonomy (CRediT) to recognize individual author contributions, reduce authorship disputes, and facilitate collaboration.

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Riad Loukili	✓	✓	✓	✓		✓		✓	✓	✓	✓			✓
Fayçal Messaoudi		✓			✓	✓		✓		✓	✓	✓		
Raouya El Youbi	✓		✓		✓		✓			✓				
Manal Loukili	✓	✓		✓			✓		✓	✓	✓	✓	✓	

C : **C**onceptualization

M : **M**ethodology

So : **S**oftware

Va : **V**alidation

Fo : **F**ormal analysis

I : **I**nvestigation

R : **R**esources

D : **D**ata Curation

O : Writing - **O**riginal Draft

E : Writing - Review & **E**ding

Vi : **V**isualization

Su : **S**upervision

P : **P**roject administration

Fu : **F**unding acquisition

CONFLICT OF INTEREST STATEMENT

The authors state that they have no conflict of interest.

INFORMED CONSENT

This study did not involve the direct recruitment of human participants or the collection of personally identifiable information.

ETHICAL APPROVAL

The study was conducted using existing benchmark datasets and did not involve experiments on human participants or animals.

DATA AVAILABILITY

The datasets used in this study are publicly available from their respective repositories. The PaySim financial transaction dataset is available at <https://github.com/EdgarLopezPhD/PaySim> and <https://www.kaggle.com/datasets/ealaxi/paysim1>. The European Credit Card Fraud Dataset, provided by the Machine Learning Group of Université Libre de Bruxelles, is available at <https://www.kaggle.com/datasets/mlg-ulb/creditcardfraud>. The IEEE-CIS Fraud Detection Dataset is available through the official IEEE-CIS Kaggle competition at <https://www.kaggle.com/competitions/ieee-fraud-detection/data>. No new proprietary dataset was created for this study.

REFERENCES

- [1] O. El Aalouche, F. Messaoudi, M. Loukili, and R. Loukili, "A temporal adaptive fuzzy clustering framework for dynamic behavioral segmentation in e-commerce," *Statistics, Optimization and Information Computing*, vol. 15, no. 4, pp. 3029–3041, Mar. 2026, doi: 10.19139/soic-2310-5070-3243.
- [2] R. El Youbi, F. Messaoudi, M. Loukili, R. Loukili, and O. El Aalouche, "From clicks to conversions: leveraging apriori and behavioural segmentation in e-commerce," *Statistics, Optimization and Information Computing*, vol. 15, no. 1, pp. 848–858, Oct. 2025, doi: 10.19139/soic-2310-5070-2890.

- [3] M. Loukili, F. Messaoudi, and R. El Youbi, *Enhancing financial transaction security: a deep learning approach for e-payment fraud detection*. Boca Raton: Chapman and Hall/CRC, 2024, doi: 10.1201/9781032656830-14.
- [4] M. Loukili, F. Messaoudi, O. El Aalouche, R. El Youbi, and R. Loukili, "Adaptive pricing strategies in digital marketing," *Statistics, Optimization and Information Computing*, vol. 14, no. 3, pp. 1244–1251, Jul. 2025, doi: 10.19139/soic-2310-5070-2200.
- [5] M. Anshari and Y. Alas, "Smartphones habits, necessities, and big data challenges," *Journal of High Technology Management Research*, vol. 26, no. 2, pp. 177–185, 2015, doi: 10.1016/j.hitech.2015.09.005.
- [6] B. Ul, R. F., A. Mehraj, A. Ahmad, and S. Assad, "A compendious study of online payment systems: past developments, present impact, and future considerations," *International Journal of Advanced Computer Science and Applications*, vol. 8, no. 5, 2017, doi: 10.14569/ijacsa.2017.080532.
- [7] 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 *2022 International Conference on Engineering & MIS (ICEMIS)*, IEEE, Jul. 2022, pp. 1–8, doi: 10.1109/ICEMIS56295.2022.9914393.
- [8] A. S. Nair and P. Kannan, "Digital payment methods: challenges and opportunities," *Journal of Namibian Studies*, vol. 37, pp. 367–376, 2023.
- [9] 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, Mar. 2021, doi: 10.3389/fcomp.2021.563060.
- [10] R. El Youbi, F. Messaoudi, M. Loukili, and R. Loukili, "From click to checkout: deep learning for real-time fraud detection in e-payment systems," *Statistics, Optimization and Information Computing*, vol. 14, no. 6, pp. 3398–3408, Sep. 2025, doi: 10.19139/soic-2310-5070-2891.
- [11] European Parliament and Council of the European Union, "Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence," *Official Journal of the European Union*, 2024, [Online]. Available: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
- [12] 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*, vol. 17, no. 5, Boca Raton: CRC Press, 2024, pp. 27–35, doi: 10.1201/9781032667478-3.
- [13] A. Nesvijevskaia, S. Ouillade, P. Guilmin, and J. D. Zucker, "The accuracy versus interpretability trade-off in fraud detection model," *Data and Policy*, vol. 3, no. 7, p. e12, Jul. 2021, doi: 10.1017/dap.2021.3.
- [14] PwC, "EU's AI Act: What regulators should know 14," PricewaterhouseCoopers. Accessed: May 28, 2026. [Online]. Available: <https://www.pwc.com/us/en/services/consulting/cybersecurity-risk-regulatory/library/tech-regulatory-policy-developments/eu-ai-act.html>
- [15] R. Avacharmal, "Leveraging supervised machine learning algorithms for enhanced anomaly detection in anti-money laundering (AML) transaction monitoring systems : a comparative analysis of performance and explainability," *African Journal Of Artificial Intelligence and Sustainable Development*, vol. 1, no. 2, pp. 68–85, 2021.
- [16] M. Saidi, F. Messaoudi, R. Loukili, B. Elghomri, E. I. Al-Wreikat, and M. Loukili, "Hybrid neural architectures for e-commerce personalization: balancing accuracy and diversity through GNN-transformer integration," in *2026 International Conference on Artificial Intelligence for Sustainable Engineering and Innovation (AISEI)*, IEEE, Apr. 2026, pp. 694–699, doi: 10.1109/AISEI68628.2026.11572815.
- [17] M. Loukili, F. Messaoudi, and M. El Ghazi, "Personalizing Product Recommendations using Collaborative Filtering in Online Retail: A Machine Learning Approach," in *2023 International Conference on Information Technology (ICIT)*, IEEE, Aug. 2023, pp. 19–24, doi: 10.1109/ICIT58056.2023.10226042.
- [18] M. Abdul Salam, K. M. Fouad, D. L. Elbably, and S. M. Elsayed, "Federated learning model for credit card fraud detection with data balancing technique," *Neural Computing and Applications*, vol. 36, no. 11, pp. 6231–6256, Apr. 2024, doi: 10.1007/s00521-023-09410-2.
- [19] P. O. Shoetan and B. T. Familoni, "Transforming fintech fraud detection with advanced artificial intelligence algorithms," *Finance & Accounting Research Journal*, vol. 6, no. 4, pp. 602–625, 2024, doi: 10.51594/farj.v6i4.1036.
- [20] M. Loukili, F. Messaoudi, and M. El 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, doi: 10.1504/IJMP.2024.140861.
- [21] P. Dubey, P. Dubey, and P. N. Bokoro, "A unified transformer–BDI architecture for financial fraud detection: distributed knowledge transfer across diverse datasets," *Forecasting*, vol. 7, no. 2, p. 31, Jun. 2025, doi: 10.3390/forecast7020031.
- [22] Y. Gorishniy, I. Rubachev, V. Khurlov, and A. Babenko, "Revisiting deep learning models for tabular data," in *Advances in Neural Information Processing Systems*, 2021, pp. 18932–18943.
- [23] I. D. Mienye and N. Jere, "Deep learning for credit card fraud detection: a review of algorithms, challenges, and solutions," *IEEE Access*, vol. 12, pp. 96893–96910, 2024, doi: 10.1109/ACCESS.2024.3426955.
- [24] J. Ge, L. Yin, S. Zhang, and X. Zhao, "Gated attention based generative adversarial networks for imbalanced credit card fraud detection," *PeerJ Computer Science*, vol. 11, p. e2972, Jun. 2025, doi: 10.7717/peerj-cs.2972.
- [25] Y. Pushkarenko and V. Zaslavskyi, "Synthetic data generation for fraud detection using diffusion models," *Information & Security: An International Journal*, vol. 55, no. 2, pp. 185–198, 2024, doi: 10.11610/isij.5534.
- [26] R. Roy, D. Tiwari, and A. Pandey, "FraudDiffuse: diffusion-aided synthetic fraud augmentation for improved fraud detection," in *Proceedings of the 5th ACM International Conference on AI in Finance*, New York, NY, USA: ACM, Nov. 2024, pp. 90–98, doi: 10.1145/3677052.3698658.
- [27] B. Lim, S. Arik, N. Loeff, and T. Pfister, "Temporal fusion transformers for interpretable multi-horizon time series forecasting," *International Journal of Forecasting*, vol. 37, no. 4, pp. 1748–1764, Oct. 2021, doi: 10.1016/j.ijforecast.2021.03.012.
- [28] U. Zafar and F. Wu, "Methodological challenges in explainable AI for fraud detection: a systematic literature review," *Artificial Intelligence Review*, vol. 59, no. 4, p. 115, Feb. 2026, doi: 10.1007/s10462-026-11516-7.
- [29] C. Ju, X. Ma, and B. Dong, "Real-time cross-border payment fraud detection using temporal graph neural networks: a deep learning approach," *Academic Journal of Sociology and Management*, vol. 3, no. 2, pp. 1–12, Mar. 2025, doi: 10.70393/616a736d.323639.
- [30] S. K. Aljunaid, S. J. Almheiri, H. Dawood, and M. A. Khan, "Secure and Transparent Banking: Explainable AI-driven federated learning model for financial fraud detection," *Journal of Risk and Financial Management*, vol. 18, no. 4, p. 179, Mar. 2025, doi: 10.3390/jrfm18040179.

BIOGRAPHIES OF AUTHORS

Riad Loukili    is a software engineer and researcher specializing in AI, DevOps, and MLOps with professional experience in 4 countries across 3 continents. Valedictorian in AI and Data Science, focused on ML, cloud computing, and optimization. He is currently a Ph.D. student at Sidi Mohamed Ben Abdellah University of Fez in the Artificial Intelligence, Data Sciences and Systems Laboratory. His main Ph.D. research interests focus on Artificial Intelligence within the context of E-Commerce and E-Marketing strategies. He can be contacted at email: riad.loukili@usmba.ac.ma.



Prof. Fayçal Messaoudi    is a Full Professor at the National School of Business and Management, Sidi Mohamed Ben Abdellah University, Fez, Morocco. He is a member of the Research Laboratory Artificial Intelligence, Data Science, and Emerging Systems. His main teaching and research interests concern artificial intelligence, data analysis, database management, and e-marketing. Prof. Messaoudi is also deeply involved in several research projects and consultancy in the field of digital marketing strategies and business analytics. He can be contacted at email: faycal.messaoudi@usmba.ac.ma.



Raouya El Youbi    is a Doctor in Sidi Mohamed Ben Abdellah university of Fez from the Artificial Intelligence, Data Sciences and Systems Laboratory. She is also an IT engineer who graduated from the National School of Applied Sciences of Fez. Her main research interests focus on artificial intelligence, machine learning, and deep learning within the context of e-commerce and digital marketing. She can be contacted at email: raouya.elyoubi@usmba.ac.ma.



Prof. Manal Loukili    is a Professor and researcher at the National School of Business and Management, Sidi Mohamed Ben Abdellah University in Fez, Morocco, where she also obtained her doctorate. Specializing in artificial intelligence, machine learning, and digital marketing, she contributes extensively to the artificial intelligence, data science and emerging systems laboratory. Prof. Loukili has published numerous articles in these areas and is an active participant in international conferences dedicated to AI and its applications in marketing. She can be contacted at email: manal.loukili@usmba.ac.ma.