CEUR-WS.org/Vol-4029/short1.pdf

Variational Autoencoders for Detecting Anomalous and Fraudulent Transactions in Financial Systems*

Serghiy Obushnyi[1,†], Denys Virovets[1,†], Andrii Ramskyi[1,†], Maksym Zhytar[1,†] and Pavlo Skladannyi[1,2,*,†]

[1] *Borys Grinchenko Kyiv Metropolitan University, 18/2 Bulvarno-Kudriavska str., 04053 Kyiv, Ukraine*

[2] *Institute of Mathematical Machines and Systems Problems of the National Academy of Sciences of Ukraine, 42 Ac. Glushkov ave., 03680 Kyiv, Ukraine*

Abstract

The rapid expansion of digital financial transactions is leading to an increase in both the complexity and sophistication of fraudulent schemes, many of which circumvent traditional classification systems based on rules and supervised learning. In this paper, we investigate the effectiveness of variational autoencoders (VAE) as an unsupervised method for detecting both anomalous and fraudulent financial transactions. We compare three approaches: classical supervised models (logistic regression, XGBoost), the VAE trained exclusively on legitimate transactions to detect outliers due to reconstruction error, and a hybrid model that combines VAE-based anomaly detection with a classifier for fraud labeling. Furthermore, based on the VAEs, we try to generate synthetic fraudulent transactions that simulate new fraud models designed to deceive standard models. Our results demonstrate that the VAEs are effective in detecting previously unknown fraudulent behavior and offer increased flexibility and explanatory power through their latent space. This research highlights the potential of deep generative models to complement traditional financial fraud detection systems and provides a foundation for anomaly-aware hybrid architectures in the financial domain.

Keywords

Variational Autoencoder, VAE, Machine Learning, Financial Technology, Financial Transactions, Financial Forensic, Financial Fraud, Urban Development Projects, Data Science, Deep Learning in Finance, Unsupervised Learning, Fraud Detection Model

1. Introduction

In modern financial systems, transactional activity is growing exponentially, which is accompanied not only by an increase in the efficiency of customer service, but also by a rapid complication of financial fraud forms. According to international analytical agencies, the volume of losses from fraudulent actions in the banking sector is growing every year, and traditional rule-based fraud detection systems are increasingly demonstrating an inability to adapt to new, non-standard schemes. According to the Association of Certified Fraud Examiners (ACFE), the banking industry loses approximately USD 3.5 trillion annually due to ineffective control of financial crimes, and this figure continues to rise [1]. In view of this, financial institutions need methods that can detect not only known types of fraud, but also anomalous transactions, with signs of new fraud or deliberately disguised as regular operations.

Deep learning methods based on variational autoencoders (VAE) have recently attracted considerable attention from researchers, as they allow them to study the distribution of normal transactions and detect anomalies without the need for labelled data. Unlike classical classification models that require examples of fraudulent transactions for training, the VAE are able to detect new or rare patterns that were not represented in the training sample. The VAE also provide

* *DECaT'2025: Digital Economy Concepts and Technologies, April 4, 2025, Kyiv, Ukraine*

* Corresponding author.

† These authors contributed equally.

s.obushnyi@kubg.edu.ua (S. Obushnyi); seito@ukr.net (D. Virovets); a.ramskyi@kubg.edu.ua (A. Ramskyi); m.zhytar@kubg.edu.ua (M. Zhytar); p.skladannyi@kubg.edu.ua (P. Skladannyi)

0000-0001-6936-955X (S. Obushnyi); 0000-0003-4934-8377 (D. Virovets); 0000-0001-7368-697X (A. Ramskyi); 0000-0003-3614-0788 (M. Zhytar); 0000-0002-7775-6039 (P. Skladannyi)

© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

opportunities to explain the identification of transactions as fraudulent through latent space and the generation of synthetic transactions, which can be used for stress testing and building secure models.

In this paper, we conduct a comparative analysis of the fraud detection performance of three approaches: a classical XGBoost-based classification model, a variational autoencoder model, and a hybrid model that combines the VAE for anomaly detection and a classifier for fraud type refinement. To test atypical types of anomalies, we generated synthetic fraudulent transactions that simulate modern schemes that are difficult to identify with traditional models. Thus, our study aims to empirically verify the benefits of the VAE and their combinations for detecting both known and new types of financial fraud, which can be used in the design of new financial protection systems.

2. Literature review

The problem of detecting financial fraud in the context of increasing transaction complexity and dynamic development of financial services remains a central challenge in applied data analysis. Traditionally, this problem has been addressed using financial performance analysis models [2], supervised learning, in particular logistic regression [3–7], decision trees, random forests [8, 9] and gradient boosting (e.g. XGBoost) [10–13]. These approaches have demonstrated high accuracy provided that there is a complete, balanced and representative set of labeled data that reflects typical examples of fraud. However, such conditions are rarely met in real financial systems, where fraudulent transactions are rare, atypical and constantly changing form. As a result, supervised models tend to overtrain on known patterns, losing the ability to generalize in the face of the emergence of novel types of fraud.

The challenges of combating fraud are driving the growth of interest in unsupervised learning models, in particular deep generative models that can learn the structure of "normal" transactions and detect those that deviate from typical behavior. Among such models, the VAEs introduced by Kingma and Welling (2013) [14], which combine an autoencoder architecture with Bayesian optimization and a probabilistic representation in latent space, have received significant attention. In the financial literature, the VAEs are proposed to be applied to a wide range of tasks. In particular, in the studies of Bergeron et al. (2021) [15], these models were used to reconstruct implied volatility surfaces in currency options, demonstrating higher accuracy compared to classical parametric models such as Heston. Van den Oever and Borovkova (2022) [16] modified the VAE architecture by proposing the use of Student's distribution instead of Gaussian in the latent space, which allowed for better modeling of heavy-tailed characteristics of portfolio returns and improved estimation of Value at Risk (VaR). A separate niche is occupied by studies devoted to representing bank customers in latent space. Mancisidor et al. (2019) [17] applied the VAE to segment customers by behavioral patterns without the need for prior labeling, which creates the basis for credit scoring under conditions of limited information. Caprioli et al. (2024) [18] further extended the use of the VAEs by generating synthetic correlation matrices of assets while preserving the properties characteristic of empirical financial correlations, which makes this approach valuable for building stress scenarios and validating risk models.

Equally important is the ability of the VAE to help explain anomalies. Ernst (2024) [19] proposed an approach to integrate the VAE with counterfactual analysis to generate local interpretations in anomaly detection problems in healthcare and financial accounting. This approach allows to get an answer as to what minimal changes in input characteristics would lead to a change in the classification solution, which meets the requirements of transparency and algorithmic fairness in modern financial systems. Models with variational autoencoders are proposed mainly as generative models for a compressed representation of input sequential data [20]. Variants combining VAE-based resampling with classical deep learning methods has been developed and proposed to solve the problems of fraud detection in transactions [21]. To solve the problem of fraud classification, Parthasarathy, et al. (2025) introduced a hybrid platform combining variational autoencoders and transformer networks [22] and demonstrates the effectiveness of multi-model systems in fraud detection tasks.

Despite this growing interest, the direct application of the VAEs in financial analytics for detecting fraudulent transactions remains relatively underexplored. In particular, there is a lack of empirical studies assessing the VAE's capacity to respond to the emergence of new, difficult-to-detect fraudulent schemes that were not represented in the training data, and whether such a model can be integrated with classical classifiers to build hybrid detection systems. In this context, this article is intended to fill the gap by investigating the effectiveness of the variational autoencoder in detecting both known and new fraudulent transactions. The potential of a hybrid architecture of the VAE with a classifier and simulation of modern fraudulent scenarios for assessing the robustness of models in real market conditions is also explored.

3. Methodology

In our research we aim to compare three different approaches to detecting fraudulent transactions in financial systems, allowing to assess the ability of each method to adapt to complex and constantly changing conditions. We consider classical supervised learning algorithms, in particular logistic regression and XGBoost, as a baseline, as well as unknown models based on the VAEs used for anomaly detection, and finally a hybrid approach that combines the advantages of both methods.

For the experiments, we use a publicly available dataset of payment transactions, known as the Credit Card Fraud Detection dataset from ULB [23], which contains over 284 thousand records with a small number of transactions with signs of fraud. Although the data has undergone an anonymization process using PCA, which allowed us to obtain 28 components, our task is to detect anomalies in transactions. First, we standardize the "Amount" variable, since it is key for analysing the scale of transactions, and we exclude the "Time" variable due to its lack of relevant information for fraud detection. For classical models, balancing methods were applied, in particular, the use of the RandomUnderSampler algorithm, which allowed to create a representative sample for training.

A special aspect of our methodology is the generation of synthetic data that simulates modern fraud schemes capable of deceiving traditional models. We created over 300 synthetic transactions that reflect scenarios such as "low and slow" attacks, when fraudsters make numerous microtransactions, transactions with unusual patterns typical of phishing attacks or operations carried out from new devices. These transactions can be characterized by anomalous values of the "Amount" parameter, in particular non-standard rounding (e.g. 9.99, 0.01), shifted correlation between features, which forms atypical feature vectors; high distance from the centroids of clusters of real data in PCA space, simulation of the activity of new users who do not have a historical behaviour pattern, etc. Some of these model testing data were generated manually by modifying real examples, while others were generated through interpolation in the VAE latent space, focusing on areas of high reconstruction error.

4. Building, Training, and Evaluating Models

At the initial stage, we train traditional classifiers. Logistic regression, which we chose as a simple, interpretable method, and XGBoost, one of the most efficient ensemble algorithms that performs well on unbalanced datasets. These models use historical labels of fraudulent transactions to build predictions, which enables strong performance on familiar fraud patterns, but with a problem of generalization to new fraud schemes. To address this limitation, we introduce a variational autoencoder trained exclusively on normal transactions so the model learned to reflect the characteristic features of legitimate transactions in its latent space. The architecture of the VAE model consists of a two-layer encoder that compresses information to a latent space of two or four dimensions, and a corresponding decoder that restores the original data. The loss function in the model combines the root mean square reconstruction error and regularization using Kullback-Leibler divergence. Transactions with reconstruction errors exceeding a set threshold (defined as the 95th percentile on the validation set) were classified as anomalous.

To further enhance detection efficiency, we also implemented a hybrid model. This approach the VAE acts as a filter to to identify potentially suspicious transactions, after which these transactions are fed to a classifier (in our case, XGBoost), which makes the final decision on whether the transaction is fraudulent. This combination allows us to focus on analysing data where the probability of fraud is high, which is especially important in resource-constrained environments.

To compare the effectiveness of each of the three approaches, we used standard metrics such as Precision, Recall, F1-score, and ROC-AUC. Particular attention was paid to the analysis of False Positive Rate, since in financial systems false positives can lead to unwanted delays or blocking of legitimate transactions. The models were evaluated both on a real test set and on synthetic data simulating new fraudulent scenarios, which allowed us to determine their resistance to unknown threats. Thus, the chosen methodology allows us to evaluate the performance of traditional and generative models in fraud detection tasks, as well as to explore the possibilities of hybrid approaches for building more adaptive and transparent security systems.

5. Mechanisms and Benefits of Variational Autoencoders in Financial Fraud Detection

The VAE is a deep generative model that combines the structure of a classical autoencoder and a probabilistic framework of Bayesian approach to modelling latent space. Unlike a conventional autoencoder that compresses information into a fixed point in latent space, the VAE learns to approximate the probability distribution of latent variables. This key feature allows the model ability to reproduce new, yet unobserved examples from the same distribution as the original data, as well as to detect distortions that do not fit this distribution. In the context of financial fraud detection, this property is especially valuable, since the VAE allows to learn a generalized idea of normal transaction behaviour without the need for labelled examples of fraudulent transactions. After training on normal transactions, the model is able to detect anomalous transactions that have a high reconstruction error, i.e., deviate from a normal distribution.

Technically, the VAE consists of two main components: an encoder that transforms the input data into the distribution parameters (mean and variance) of the latent variables, and a decoder that reconstructs the output data from the latent space. The model optimizes a loss function that combines the reconstruction error (e.g., the mean square error) and the Kullback–Leibler (KL) divergence, which forces the latent space to approximate a given normal distribution. Unlike supervised learning methods such as logistic regression or XGBoost, which require a large number of correctly labelled transactions, the VAE operates in an unsupervised mode. This allows the model to be trained efficiently even in the absence or low quality of fraud data. Moreover, classification models are limited in detecting new or rare types of fraud, while the VAE has the ability to respond to any deviation from the learned normal distribution.

In summary, the VAE model is unique among machine learning methods in finance, combining the ability to generalize, insensitivity to the absence of labels, the ability to detect new patterns with support for flexible reconstruction of multidimensional transactions [24–26]. These properties characterize the effectiveness of the VAE for building next-generation antifraud systems, especially when combined with traditional classifiers within a hybrid architecture.

6. Results

This section presents the results of an experimental comparison of three approaches on fraud detection with transaction data: classical classification models, the VAE, and a hybrid system that combines them. The analysis covers both the detection of already known fraudulent transaction patterns and the robustness of the models to newly generated fraud scenarios.

As anticipated, the best results among the classical models were demonstrated by the XGBoost classifier, which achieved ROC-AUC of 0.982 and detected about 91% of fraudulent transactions in the test sample. However, like most supervised learning models, XGBoost was vulnerable to previously unknown types of fraud, since its accuracy is based on previously observed patterns. In contrast, the VAE (Variational Autoencoder) model, which operates in unsupervised mode, showed slightly lower accuracy (ROC-AUC of 0.955), but has the advantage of being able to detect anomalous and new fraudulent transactions without the need for labels. This makes it especially valuable in environments where attackers change behavioural patterns. Good results were achieved by a hybrid model, which combines the strengths of both approaches. It demonstrated high stability and accuracy (ROC-AUC of 0.979), effectively combining the ability of XGBoost to detect familiar patterns with the sensitivity of the VAE to novel anomalies. This indicates the promise of the hybrid approach for practical application in dynamic environments of financial systems.

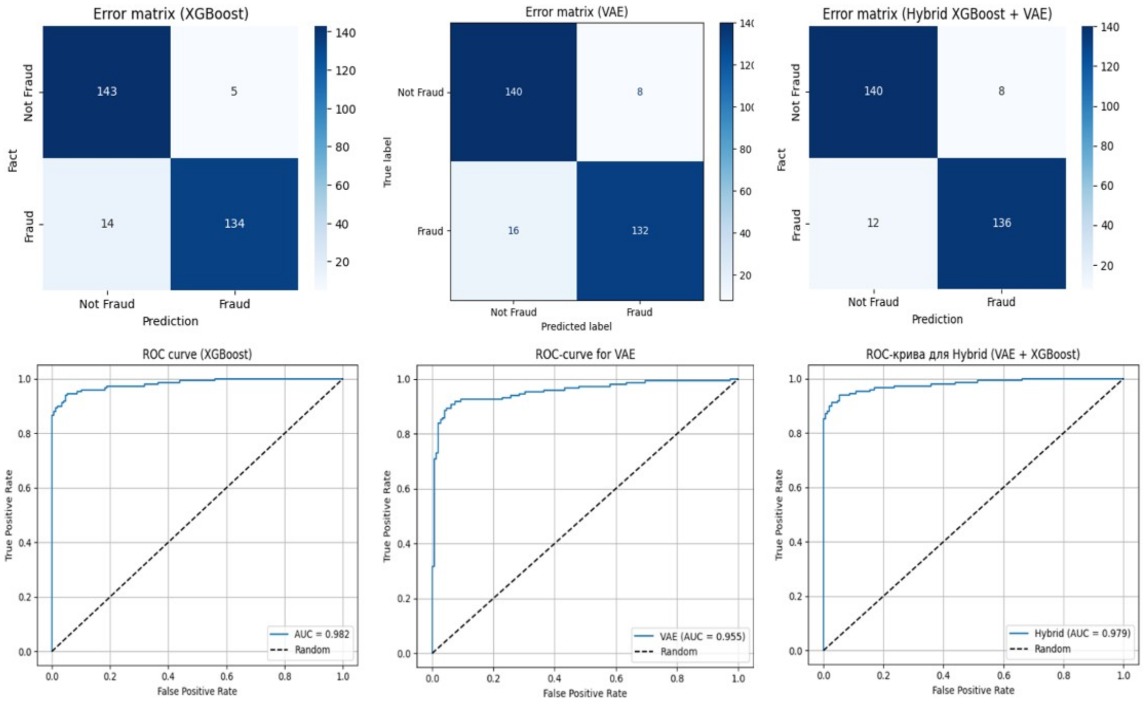

Figure 1: Error matrices and ROC curves of models

Evaluating three approaches to transactional fraud detection on a shared set of real and synthetic data revealed key advantages and limitations of each method. The most effective approach for detecting both known and novel fraud types was a compromise between XGBoost (high accuracy) and the VAE (high sensitivity). To achieve better performance, it was proposed to consider ensemble or multi-level approaches, where the VAE can act as a filter or pre-detector, and XGBoost as the final classifier.

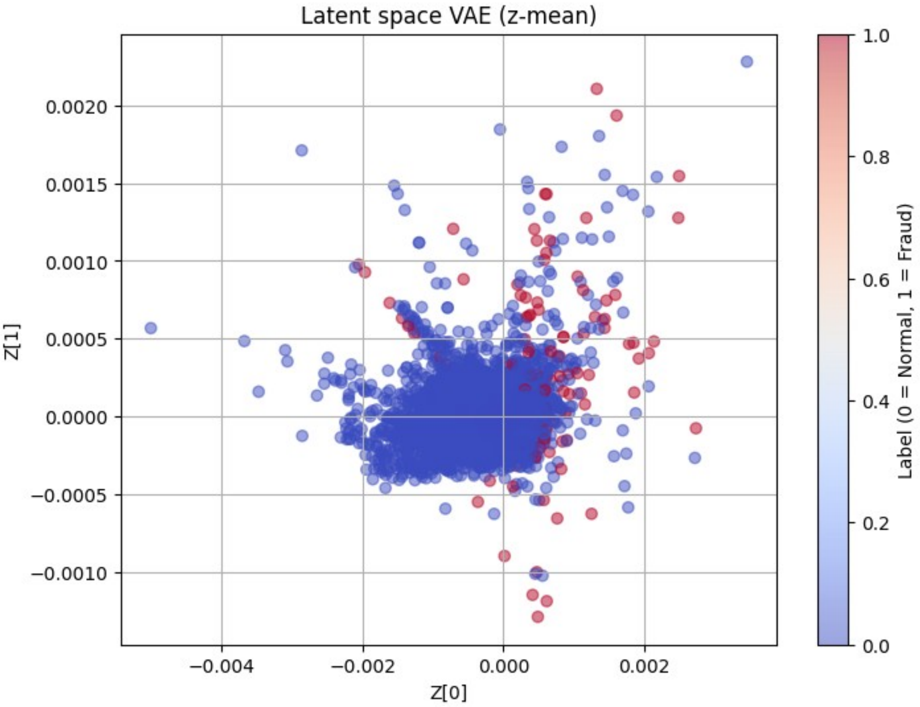

Figure 2: Latent space of the VAE with normal and fraudulent operations

The latent space of the VAE with normal and fraudulent operations shows that some fraudulent transactions are far from the centre of the cluster, but many of them overlap with normal ones. This means that some fraudulent transactions have a similar structure to normal ones, but the VAE model detects them. The VAE has effectively learned to capture behavioral anomalies as a critical feature for solving the problems of detecting emerging fraud schemes.

Despite the ability of the VAE to separate most normal transactions into a dense latent core, fraud detection results indicate that there is significant overlap between classes in the two-dimensional latent space. This limits the accuracy of the model in conditions of high similarity of some fraudulent transactions to typical samples. To improve the discriminatory power of the VAE, it is suggested to expand the latent space by increasing the dimensionality, which will allow the model to better capture subtle structural variations. Additionally, the application of clustering algorithms in the latent space (e.g., DBSCAN or Gaussian Mixture models) can help identify potential subgroups of fraudulent transactions that were not captured by the linear classification boundaries. Also promising is the approach of training the VAE on new types of fraudulent attacks using fine-tuning or few-shot learning strategies, which can increase the adaptability of the model to evolving fraud tactics.

Evaluating the performance of the models after adding synthetic transactions revealed important limitations in using generated data in fraud detection tasks. While the baseline results on the real test set demonstrated high precision and recall—especially for XGBoost and the hybrid model (VAE+XGBoost), the integration of 300 generated the VAE transactions simulating new types of fraud significantly affected the overall classification results. After adding synthetic examples, the recall of the hybrid model for the “fraud” class decreased from over 90% to 16%, and XGBoost—to 75%. Particularly revealing was the result of the VAE, which detected a significant number of false positives, incorrectly classifying normal transactions as suspicious. This indicates the limited ability of the VAE to accurately generalize the limits of the normal distribution when retraining on small or structurally different fraud samples. This contrast in results highlights the difficulty of generating credible fraud scenarios and demonstrates that models focused solely on synthetic data may lose their ability to generalize. Moreover, a hybrid model combining latent reconstruction with supervised classification was found to be vulnerable to data bias, causing false retraining on spurious anomalies.

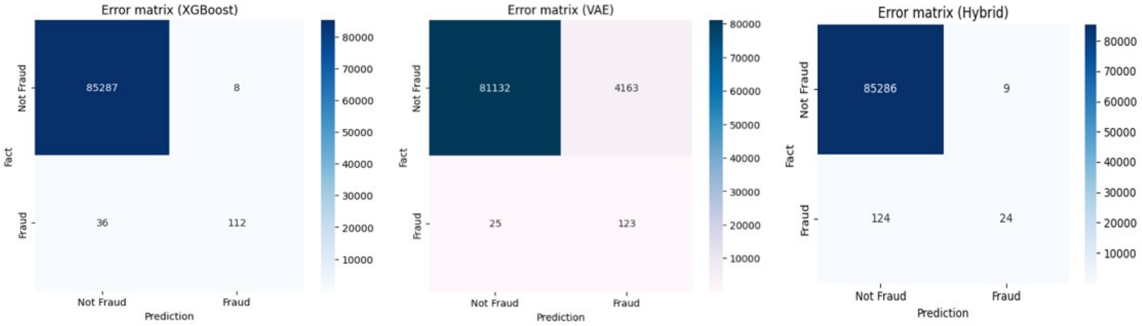

Figure 3: Error matrices of models on generated data with new fraudulent features

In summary, the experimental results demonstrate that synthetic data with signs of fraud can be useful only if their quality and structural correspondence to real patterns are carefully controlled. Future research should focus on developing more complex generative models or use active learning approaches that adapt to the dynamics of the real transaction environment. Overall, the study confirms the effectiveness of the VAE as an independent anomaly detection mechanism and its key role in hybrid architectures. In the context of dynamic changes in fraud strategies, when new attacks evolve rapidly, the ability of the VAE to detect unknown patterns becomes critically important. Its integration with traditional models allows to achieve a balance between sensitivity to new threats and stability of the forecast in in routine financial operations.

Conclusions

Detecting fraudulent transactions remains one of the key challenges for digital financial systems. In this study, we conducted a comprehensive evaluation of the effectiveness of the variational autoencoder (VAE) as a tool for detecting anomalies in financial transactions, comparing its results with the traditional XGBoost classifier, as well as implementing a hybrid model that combines both approaches. The experimental results confirmed that XGBoost demonstrates high accuracy in detecting known types of fraud, but its effectiveness decreases significantly when faced with new or changed fraud patterns. In contrast, the VAE, trained only on normal transactions, showed the ability to recognize atypical deviations, including in synthetic transactions simulating novel fraud scenarios. This allows to consider the VAE as an effective means of detecting new threats in scenarios where labelled examples are absent or limited.

The highest performance was achieved with the hybrid architecture, which combines the deep anomaly detector VAE with the accuracy of the supervised classifier XGBoost. This approach allowed not only to mitigate the level of false positives inherent in the VAE, but also to preserve its sensitivity to atypical transactions. The hybrid model effectively detected fraudulent transactions that were individually missed by each of the models, providing a synergistic effect. This emphasizes the feasibility of combining unsupervised and supervised learning to create more reliable antifraud solutions. The practical value of the proposed hybrid model lies in its ability to adapt to new forms of fraud while maintaining high performance in real time. It can be easily integrated into existing transaction monitoring systems as a pre-filtering or active learning tool.

Overall, the study confirms the effectiveness of the VAE as a modern anomaly detection tool and the feasibility of using hybrid architectures as a new approach to building adaptive antifraud systems in the financial sector. Future research should focus on improving the interpretability of the VAE solutions (for example, through counterfeit analysis), as well as expanding the use of multimodal data sources such as geolocation, time patterns, behavioural and interaction metrics to better capture the transaction context.

Declaration on Generative AI

While preparing this work, the authors used the AI programs Grammarly Pro to correct text grammar and Strike Plagiarism to search for possible plagiarism. After using this tool, the authors reviewed and edited the content as needed and took full responsibility for the publication's content.

References

[1] F. Mehdipour, et al., Banking Fraud Identification and Prevention, in: 2023 27th Int. Conf. on Circuits, Systems, Communications and Computers (CSCC), 2023, 95–100. doi:10.1109/CSCC58962.2023.00019

[2] R. Kanapickiene, Z. Grundienė, The Model of Fraud Detection in Financial Statements by Means of Financial Ratios, Procedia—Social and Behavioral Sciences, 213 (2015) 321–327. doi:10.1016/j.sbspro.2015.11.545

[3] H. Z. Alenzi, N. O. Aljehane, Fraud Detection in Credit Cards using Logistic Regression, Int. J. Adv. Comp. Sci. Appl. 11(12) (2020) 540–544. doi:10.14569/IJACSA.2020.0111268

[4] A. Mahajan, V. S. Baghel, R. Jayaraman, Credit Card Fraud Detection using Logistic Regression with Imbalanced Dataset, in: 2023 10th Int. Conf. on Computing for Sustainable Global Development (INDIACom), 2023, 339–342. doi:10.1109/INDIACom.2023.10155683

[5] H. Guan, et al., Financial Fraud Identification of the Companies based on the Logistic Regression Model, J. Competitiveness, 14(4) (2022) 155–171. doi:10.7441/joc.2022.04.09

[6] E. Brooks, D. Mercer, Logistic Regression on Banking Fraud, World J. Adv. Eng. Technol. Sci., 7(2) (2022) 334–348. doi:10.30574/wjaets.2022.7.2.0132

[7] D. Yue, X. Wu, N. Shen, C.-H. Chu, Logistic Regression for Detecting Fraudulent Financial Statement of Listed Companies in China, in: 2009 Int. Conf. Artif. Intell. Comput. Intell., 2009, 104–108. doi:10.1109/AICI.2009.421

[8] J. K. Afriyie, et al., A Supervised Machine Learning Algorithm for Detecting and Predicting Fraud in Credit Card Transactions, Data Analytics, 2 (2023) 100163. doi:10.1016/j.dajour.2023.100163

[9] Y. Salunke, S. Phalke, M. Madavi, P. Kumre, G. Bobhate, Fraud Detection: A Hybrid Approach with logistic regression, Decision Tree, and Random Forest, Cureus J. Comput. Sci., 2 (2025) eS44389-024-02350-5. doi:10.7759/s44389-024-02350-5

[10] S. Lei, K. Xu, Y. Huang, X. Sha, An XGBoost based System for Financial Fraud Detection, E3S Web Conf., 214 (2020) 02042. doi:10.1051/e3sconf/202021402042

[11] S. El Kafhali, T. Mohammed, XGBoost based Solutions for Detecting Fraudulent Credit Card Transactions, in: 2022 Int. Conf. Adv. Creat. Netw. Intell. Syst. (ICACNIS), 2022. doi:10.1109/ICACNIS57039.2022.10054965

[12] M. Wang, J. Yu, Z. Ji, Credit Fraud Risk Detection based on XGBoost-LR Hybrid Model, in: Proc. 18th Int. Conf. Electron. Bus., 2018, 336–343.

[13] M. Karbasiyan, H. Hamidi, K. Srinivasa Rao, Presenting a model to detect the fraud in banking using smart enabling tools, Int. J. Eng. Trans. C, 37(3) (2024) 529–537.

[14] D. P. Kingma, M. Welling, An Introduction to Variational Autoencoders, Foundations and Trends in Machine Learning, 12(4) (2019) 307–392. doi:10.1561/2200000056

[15] M. Bergeron, et al., Variational Autoencoders: A Hands-off Approach to Volatility, arXiv, 2021. doi:10.48550/arXiv.2102.03945

[16] S. Borovkova, M. van den Oever, Variational Autoencoders with Student-t Distribution for Large Portfolios, SSRN, 2022. doi:10.2139/ssrn.4274080

[17] R. A. Mancisidor, et al., Learning Latent Representations of Bank Customers with the Variational Autoencoder, arXiv, 2019. doi:10.48550/arXiv.1903.06580

[18] S. Caprioli, E. Cagliero, R. Crupi, Quantifying Credit Portfolio Sensitivity to Asset Correlations with Interpretable Generative Neural Networks, arXiv, 2023. doi:10.48550/arXiv.2309.08652

[19] R. Ernst, Counterfactual Generating Variational Autoencoder for Anomaly Detection, in Joint Proceedings of the xAI 2024 Late-breaking Work, Demos and Doctoral Consortium, vol. 3793, 2024, 353–360.

[20] A. Alazizi, et al., Dual Sequential Variational Autoencoders for Fraud Detection, in: Advances in Intelligent Data Analysis XVIII, LNCS, vol. 12080, 2020, 21–33. doi:10.1007/978-3-030-44584-3_2

[21] H. Tingfei, C.Guangquan, H. Kuihua, Using Variational Auto Encoding in Credit Card Fraud Detection, IEEE Access, 8 (2020) 161087–161096. doi:10.1109/ACCESS.2020.3015600

[22] K. Parthasarathy, et al., Fraud Detection with Variational Autoencoders and Transformer Networks: A Robust Deep Learning Approach for Banking Transactions, Int. J. Sci. Eng. Appl. 14(3) (2025) 51–57. doi:10.7753/IJSEA1403.1011

[23] Kaggle, Credit Card Fraud Detection Dataset. URL: https://www.kaggle.com/mlg-ulb/creditcardfraud

[24] V. Buhas, et al., Using Machine Learning Techniques to Increase the Effectiveness of Cybersecurity, in: Cybersecurity Providing in Information and Telecommunication Systems, vol. 3188, no. 2 (2021) 273–281.

[25] V. Zhebka, et al., Methodology for Predicting Failures in a Smart Home based on Machine Learning Methods, in: Cybersecurity Providing in Information and Telecommunication Systems, CPITS, vol. 3654 (2024) 322–332.

[26] M. Adamantis, V. Sokolov, P. Skladannyi, Evaluation of State-of-the-Art Machine Learning Smart Contract Vulnerability Detection Method, in: Advances in Computer Science for Engineering and Education VII, vol. 242 (2025) 53–65. doi:10.1007/978-3-031-84228-3_5