

Enhancing Fraud Detection through Cascading Machine Learning Models and Clustering Techniques*

Giuseppe Fenza^{1,*}, Salvatore Froncillo^{2,†} and Claudio Stanzione^{3,†}

¹Department of Management & Innovation Systems, University of Salerno, 84084 Fisciano (SA), Italy

²IMT School for Advanced Studies Lucca, 55100 Lucca (LU), Italy

³Defence Analysis & Research Institute, Center for Higher Defence Studies, 00165 Rome (RM), Italy

Abstract

Bank fraud detection is a crucial challenge in the financial sector, requiring innovative methodologies to combat the evolution of fraudulent activities. Starting from traditional rule-based systems, this study introduces a pioneering approach by merging machine learning algorithms with clustering techniques. Using a cascade approach, different models adapted to different fraud patterns are employed sequentially to classify transactions; the study explores various model ensembles to find the most effective combination. The experimental results emphasize the method's effectiveness in identifying fraudulent transactions while maintaining superior recall rates; in fact, this work emphasizes the importance of recall in this field, whereas other works focus exclusively on accuracy. Conventional classification algorithms show inefficiency with the dataset used, manifesting itself in consistently low average recall rates; in contrast, the proposed methodology yields significant improvements in accuracy and recall. Meticulous analysis of false positives and negatives validates the system's robustness, promising a solid safeguard against financial losses from undetected fraud cases.

Keywords

Fraud Detection, Clustering, Machine Learning, Bank Transactions

1. Introduction

Bank fraud poses a significant threat to financial institutions and their customers worldwide, with losses amounting to billions of dollars annually [1]. The dynamic nature of fraudulent activities, coupled with the limitations of traditional rule-based systems, underscores the urgent need for more sophisticated detection mechanisms [2]. Recognizing this imperative, the integration of machine learning techniques has emerged as a promising avenue to bolster fraud detection capabilities [3]. According to UK Finance, people in the UK lost £1.2 billion to fraud in 2021, the equivalent of £2,300 every minute¹. According to SEON, fraud scams, and bank fraud schemes resulted in 485.6 billion dollars in losses globally in 2023². According to Statista, bank transfer or payment fraud resulted in losses amounting to 1.59 billion dollars in the United States

ITASEC 2024: The Italian Conference on CyberSecurity

*Corresponding author.

†These authors contributed equally.

✉ gfenza@unisa.it (G. Fenza); salvatore.froncillo@imtlucca.it (S. Froncillo); stanzione.dottorando@casd.difesa.it (C. Stanzione)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<https://www.statista.com/statistics/958997/fraud-loss-usa-by-payment-method/>

²<https://seon.io/resources/global-banking-fraud-index/>

in 2022³. These figures show the magnitude and impact of bank fraud on the financial sector and society. In this paper, a comprehensive methodology is presented, designed to address the evolving challenges of bank fraud detection. Central to the approach is the incorporation of advanced machine learning algorithms, enabling the extraction of valuable insights from complex transaction data. Unlike conventional methodologies, the approach begins with a pre-processing step that includes clustering of fraudulent transactions. This pre-processing step is instrumental in uncovering underlying patterns within fraudulent activities, thereby facilitating more targeted and effective model training. A distinguishing feature of the methodology presented in terms of results is the adoption of a cascade approach, wherein multiple trained models are sequentially applied to classify transactions. This cascade mechanism allows the exploitation of the complementary strengths of diverse machine learning models, thereby enhancing the overall detection accuracy. By iteratively refining predictions and leveraging ensemble learning techniques, the approach achieves superior performance compared to standalone models. Importantly, while conventional approaches often prioritize accuracy as the primary evaluation metric, the critical importance of recall in fraud detection is recognized [4]. A false negative, where a fraudulent transaction goes undetected, can have severe financial repercussions for both financial institutions and their customers. Therefore, the methodology places significant emphasis on optimizing recall alongside accuracy, with the overarching goal of minimizing the occurrence of false negatives and mitigating potential fraud losses. Through rigorous experimentation and evaluation, the efficacy of the proposed method in effectively detecting bank account fraud while maintaining a high recall rate is demonstrated. The findings not only contribute to the advancement of fraud detection techniques in the banking sector but also offer tangible benefits in terms of improved protection against fraudulent activities. By embracing the latest advancements in machine learning and prioritizing performance metrics that reflect real-world implications, the methodology represents a significant step forward in the ongoing battle against bank fraud. The remaining of the manuscript is structured as follows. Section 2 discusses related works on bank fraud detection. Section 3 presents the methodology of the proposed system, and Section 4 details its experimentation on a real dataset. Section 5 concludes the manuscript.

2. Related Works

Fraud detection is the process of identifying and preventing fraudulent activities, which has long been recognized as a Cybersecurity problem [5]. The problem can affect a variety of sectors, from the health care domain [6, 7] to the car insurance domain [8] to the financial domain, which is certainly the most explored sector [9, 10, 11] and where this work also focuses its attention. However, a distinction must be made within the financial domain, credit card transaction fraud [12, 13] and bank transaction fraud [14]. This paper fixes its focus strictly on bank transaction fraud, where there is not much state-of-the-art work, and mainly aims to improve accuracy, which, as extensively explained later, is not the best metric to evaluate performance in this field. In [15], the authors propose a list of the tried-and-true methods for spotting fraud, highlighting how mainly traditional machine learning methods are used. In [14], the authors propose a

³<https://www.bbc.co.uk/news/business-65545247>

system to detect bank fraud using a community detection algorithm that identifies the patterns that can lead to fraud occurrences. In contrast, in [16] Sadgali et al. analyze the performances of different machine learning models, focusing on accuracy. In [17], the authors propose the use of the Artificial Neural Network technique and Harmony Search Algorithm to detect fraud-seeking hidden patterns between normal and fraudulent customers' information. Compared to what has just been analyzed, this work differs mainly in its innovation in the proposed methodology and its focus on metrics other than those commonly used to evaluate the system.

3. Methodology

Figure 1 presents the methodology within the proposed approach. However, it is important to note that the figure depicts two phases in an experimental setting. In real-world implementation, the evaluation phase occurs continuously after the training phase is complete.

Data Collection. The first phase is the data collection phase. All banks receive and handle numerous transactions every second; behind these are frauds, transactions created ad hoc to circumvent detection systems, which analysts can often recognize. Data collection is key to unearthing any patterns that lead to understanding why and how the system is being hacked. One question mark that needs to be asked at this stage is how many transactions should be considered an elapsed period. Is a number of transactions necessary? Is it necessary to establish a time frame? This decision is up to those in charge of designing the architecture, but it should be kept in mind that it is a crucial part of the whole process. Depending on the results that are obtained and how many false positives or false negatives are obtained, it could be decided to consider a lower number or a higher number of transactions in order to be able to unearth any new patterns or eliminate patterns that no longer occur and thus lighten the whole system involving different patterns than the patterns involved.

Preprocessing Data. The preprocessing phase is primarily aimed at cleaning the data, identifying which features are most important to consider, and, most importantly, making the data anonymized to respect the privacy of those performing the transactions. Cleaning involves identifying and handling missing values, outliers, and inconsistencies, which could skew analysis. Understanding the most important features requires thoroughly examining variables like transaction amount, timestamp, and transaction type, as they can reveal patterns and anomalies essential for fraud detection and customer behavior analysis. Anonymization is imperative to safeguard sensitive information, achieved through tokenization and encryption, ensuring privacy and fairness in data analysis while adhering to regulatory standards such as GDPR. Financial institutions can derive accurate insights while upholding data privacy and integrity by meticulously addressing these aspects during preprocessing.

Filtering Frauds. The fraud filtering stage within the transaction processing system facilitates the extraction of transactions that have been categorized as fraudulent. Subsequently, this phase eliminates transactions identified as regular, directing attention solely toward examining the malicious transactions encountered thus far. This endeavor aims to discern discernible patterns that differentiate these fraudulent banking activities from legitimate ones.

Clustering Fraud. After detecting and sieving out the fraudulent activities, the next step involves pinpointing the diverse patterns that delineate the array of detected frauds. Clustering




Figure 1: Overall Framework.

serves the precise purpose of identifying common attributes binding together transactions flagged as irregular. Hence, it becomes pivotal to determine the optimal number of clusters to be considered. Several methodologies can be employed to facilitate this determination, including but not limited to the Elbow Method, cross-validation, or the Silhouette Score. Partitioning into distinct clusters will greatly aid the system user's subsequent analysis, enabling a comprehensive examination of the categorized fraud types and the specific distinguishing features of each cluster.

Fraud Detection Models' Training. The last part of the whole-system training phase involves training several machine learning models for as many used to cluster frauds. Training multiple models, each focused on a specific fraud pattern, provides a more nuanced and robust approach to fraud detection than relying solely on a single model. Adapting models to distinct fraud types or patterns goes a long way toward improving the ability to detect complex and evolving fraudulent activity. These specialized models can capture subtle nuances and variations within each fraud category, enabling more accurate identification and mitigation of fraudulent behavior. In addition, using a diverse range of models reduces the risk of false positives and negatives, improving the overall effectiveness of detection and minimizing the impact on legitimate transactions. An additional benefit of having multiple models is management in the event of

new training, so decisions can be made to retain models that perform well and eliminate models that are now adrift. In the case of a single pattern, on the other hand, the process would be more complicated, and any well-established patterns may need to be better managed.

Evaluation Phase. The evaluation phase should be considered as a streaming process phase. Indeed, the system will assess each transaction under examination, concluding the evaluation as either fraudulent or non-fraudulent. With multiple models available, it is clear to consider how to use them, either with a cascade or ensemble approach. The choice must depend on the different trained models' requirements and performance. With a cascade approach, the transaction under examination will be evaluated by each model sequentially. If labeled as non-fraudulent, it moves on to the evaluation by the next model; if labeled as fraudulent, however, the system labels it as fraudulent. With the cascade approach, therefore, it suffices for just one model to identify the transaction as fraudulent to label it as such. A significant question arises using this approach: Where should each model be positioned? Why position a certain model as first and another as last? If the clustering division is clear enough, this problem can be easily resolved, ensuring that there is no overlap between the models. With the ensemble approach, on the other hand, the transaction is evaluated by all models, and the different outputs will give the final result. In this case, it will be possible to assign a different weight to each model if one wants to emphasize a particular model, and a threshold must be set to determine whether it is fraudulent or not based on the results received from the different models.

4. System Evaluation

This chapter aims to give an overview of the experimentation conducted aimed at enhancing the proposed system by giving details of the dataset used and what use was made of it, the tools used, the different evaluation methods employed, the baseline considered on which the same tests were conducted, and finally an analysis of the results with a comparison against the baseline. A baseline of classical machine learning models was compared, and no other methods or architectures were used as no such works of this type are yet on the dataset. For the training phase, both the models used for the baseline and the models used in the proposed system were trailed with different sizes concerning the number of transactions labeled as regular and transactions labeled as fraud. Instead, the test was always on the dataset's last two months of transactions.

4.1. Dataset

The dataset used for experimentation is Bank Account Fraud (BAF) [18]. The dataset comes in six variants, mainly proposed to solve fairness problems. The variant used in this work is the Base variant. The dataset consists of 1 million instances and 32 features. The data collected are over eight months of bank transactions and what is recommended in the paper presenting the dataset, and what was also done in this work, is to use the first six months of transactions for training and the last two months for testing. There are 11029 transactions in the dataset labeled as fraud, specifically 9601 in the first six months and 1428 in the remaining two months. In contrast, the transactions labeled as regular are 903157 in the first six months and 96843 in the remaining two months. If one wants to refer to the methodology, one can label this part as part

of data collection and preprocessing. The sensitive data were anonymized by the authors of the dataset themselves and rendered in such a way that they could not identify who was behind a transaction. After a data cleaning phase where no missing values were found, the feature “device fraud count” was excluded, which always presented the value 1 in all transactions, and the feature reporting the month of the transaction was deemed unimportant for classification purposes. Instead, concerning identifying features more influential on fraud/non-fraud labeling, the best-known eXplainable AI (XAI) methods that are LIME and SHAP [19], were applied for the models also used as baselines to understand the features of interest concerning fraud. Several features were found to be more influential than others, such as the device’s operating system and housing status. However, it was decided not to use what was learned, leaving pattern identification to clustering.

4.2. Tools Used & Baseline Models

The experimentation relied on a combination of programming languages and machine-learning libraries. Python has been used for code implementation, while for data analysis and machine learning libraries, it was mostly used pandas and sklearn. The models were tested adopting different training sizes concerning regular transactions, particularly with 5k, 10k, 20k, 50k, 100k, and 200k, while all the frauds of training are ever used in each training. The baseline models used are Random Forest, XGBoost, Support Vector Machine and Logistic Regression.

Random Forest. Random Forest is a versatile machine learning algorithm that constructs multiple decision trees during training and merges their predictions to improve accuracy and prevent overfitting. It’s effective for classification and regression tasks, offering robustness and flexibility in various domains. For the experimentation the following parameters was been adopted: 100 trees, entropy criterion, and random state = 0.

XGBoost. XGBoost is a powerful and efficient implementation of gradient boosting, optimized for speed and performance. It employs a decision-tree-based ensemble learning algorithm, enhancing predictive accuracy across various regression and classification tasks, while handling missing data and preventing overfitting through regularization techniques. For the experimentation the following parameters was been adopted: max depth of 6, min child weight = 1 and learning rate = 0.3.

Support Vector Machine. Support Vector Machine is a supervised machine learning algorithm that constructs a hyperplane or set of hyperplanes in a high-dimensional space to maximize the margin between different classes while minimizing classification errors, utilizing a kernel function for nonlinear mapping. No hyperparameters were used for the experimentation.

Logistic Regression. Logistic Regression is a statistical method used for binary classification tasks, employing the logistic function to model the probability of a categorical outcome. It optimizes parameters via maximum likelihood estimation, fitting a linear decision boundary in feature space to predict class membership probabilities. Also in this case no hyperparameters were used for the experimentation.

4.3. Evaluation Phase System

The implementation details of the bank fraud detection model are outlined in this section.

4.3.1. Number of Clusters

Concerning transactions labeled as fraud in the training set, the first aspect to consider was how many clusters to use. To resolve this key issue, it was decided to use the best-known methods, the Elbow Method and the Silhouette Score. Both methods converged on the choice of 3 clusters.

4.3.2. Fraud Clustering with K-means

To begin, the K-means clustering technique is utilized to group similar fraudulent transactions into distinct clusters. K-means clustering is an unsupervised machine learning technique used for partitioning data into K clusters. It minimizes the within-cluster variance by iteratively assigning data points to the nearest cluster centroid and updating centroids based on the mean of the assigned points. This process continues until convergence, yielding clusters with minimized intra-cluster distance and distinct centroids. This clustering process allows the identification of different patterns of fraudulent behavior within the dataset.

4.3.3. Models Training

Following the clustering of fraudulent transactions, the following step involves training multiple machine learning models, each tailored to a specific cluster of fraudulent transactions. By training cluster-specific models, the goal is to capture the unique characteristics and patterns associated with different types of fraud. The different models tested as baselines were used with different sizes. Only the best-performing model was considered, which was the Random Forest. The same hyperparameters presented in the previous section were also used in this phase.

4.4. Evaluation Results

As mentioned in the methodology above, two different methods can be adopted to evaluate the results obtained from the different trained models: the cascade method or the ensemble method. In this experiment, it was decided to test both methods to decide which one to use to obtain the best performance. Implementations of the two techniques are shown below.

Cascade Approach. The cascade strategy involves the following steps:

1. **Initialization:** Initially, the entire test dataset is provided as input to the first trained model.
2. **Model Evaluation:** The first model makes predictions on the test dataset, identifying potential fraudulent transactions.
3. **Updating Test Dataset:** True positives and false positives identified by the first model are removed from the test dataset. The updated test dataset, containing only the remaining transactions, is then passed to the next model in the cascade.
4. **Sequential Model Application:** Subsequent models in the cascade repeat the process, making predictions on the updated test dataset and refining the results based on the predictions of preceding models.

As mentioned in the methodology, how the models are placed at this stage is critical. The six different possible combinations were tested during the experiment. The results were very similar, making the division between the clusters quite clear. It was then decided to draw a combination randomly; the final sequence was 2–3–1. At the conclusion of the cascade testing approach and model combination, the results from all models are aggregated to construct the confusion matrix. This matrix summarizes the classification performance, including true positives, false positives, true negatives, and false negatives. By adopting this comprehensive approach, the aim is to enhance the accuracy and effectiveness of the bank fraud detection system, leveraging cluster-specific models, cascading models, and model combination techniques. This methodology enables the construction of a comprehensive confusion matrix, providing insights into the model's performance and its ability to detect fraudulent transactions accurately.

Ensemble Approach. Within the ensemble system, the transaction undergoes dissemination across all relevant models that have been enlisted. Each model subsequently processes the transaction and generates an output representing the likelihood of its legitimacy or fraudulent nature. Following this stage, the system deliberates, assessing the array of outputs garnered from the various models. A procedural approach was established in the experimental framework wherein the outputs from three distinct models were singled out. These outputs were subjected to a statistical aggregation technique, specifically an averaging process, to derive a consolidated probability assessment. Subsequently, a threshold criterion was applied, stipulating that if the resulting average probability surpasses the 50%, the transaction is designated as fraudulent.

4.5. Experimental Results

The experimental results of the study are presented in Table 1. The results provide valuable insights into the performance of different approaches for detecting bank fraud and satisfactorily respect the proposed system. In the table concerning the proposed method, it was decided to present the results of the system with the towed models with 10k regular transactions and 9601 frauds for both methods of evaluating the proposed results, cascade and ensemble. This is exclusively because they have turned out to be the best approaches. Among the various methods evaluated, the Cascade strategy emerges as the most effective in terms of both accuracy and recall. This result is particularly significant considering the critical importance of recall in fraud detection. As pointed out in the analysis, a low recall rate implies an inability to identify fraudulent transactions, potentially leading to substantial financial losses. By prioritizing recall and accuracy, the Cascade approach demonstrates its superiority in detecting a higher percentage of fraudulent activity while maintaining a satisfactory overall accuracy level. It should be emphasized that in the field of Fraud Detection, where only 1% of transactions turn out to be fraud, other metrics are more important than accuracy, which always has its value but is affected by the high number of detections of regular transactions thus not placing the focus on false positives and false negatives. Table 1 shows how the sum of false positive and false negative rates is the lowest of all analyzed methods. Other methods, such as Random Forest, XGBoost, Support Vector Machine and Logistic Regression, show varying degrees of performance on different data sets and parameter settings. While some of these methods achieve relatively high accuracy scores, their recall and false positive rates achieve lower performance, indicating a higher probability of missing fraudulent transactions. This discrepancy underscores

the limitations of relying solely on accuracy as the primary evaluation parameter in fraud detection. By adopting a more balanced approach that considers both accuracy and recall, the proposed system demonstrates its ability to identify fraudulent transactions while minimizing false negatives effectively. The analysis highlights the impact of training set size on model performance. Across training sets of different sizes, from 5k to 200k regular transactions with the number of frauds set at 9601, the methods demonstrate steadily increasing accuracy with recall rates that, however, likewise decrease steeply just as the amounts of false positives and false negatives increase. In conclusion, the experimental results emphasize the effectiveness of the Cascade approach in detecting bank fraud. By prioritizing recall and accuracy, the method offers financial institutions a reliable and efficient solution to mitigate the risks associated with fraudulent activities.

Table 1
Performance Metrics of Different Models

Model	Acc.	Rec.	Prec.	F1	FPR	FNR	TPR	TNR
RF 5k	65.21%	89.01%	3.65%	7.02%	35.15%	10.99%	89.01%	64.85%
RF 10k	80.12%	79.41%	5.64%	10.54%	19.87%	20.59%	79.41%	80.13%
RF 20k	89.54%	67.79%	9.10%	16.04%	10.14%	32.21%	67.79%	89.86%
RF 50k	96.17%	43.98%	17.75%	25.29%	3.05%	56.02%	43.98%	96.95%
RF 100k	97.92%	23.74%	26.71%	25.14%	0.97%	76.26%	23.74%	99.03%
RF 200k	98.44%	8.33%	37.42%	13.63%	0.21%	91.67%	8.33%	99.79%
XGB 5k	70.06%	87.04%	4.14%	7.90%	30.19%	12.96%	87.04%	69.81%
XGB 10k	80.81%	78.78%	5.80%	10.80%	19.16%	21.22%	78.78%	80.84%
XGB 20k	85.15%	73.25%	6.95%	12.70%	14.67%	26.75%	73.25%	85.33%
XGB 50k	93.60%	55.53%	12.48%	20.38%	5.83%	44.47%	55.53%	94.17%
XGB 100k	96.86%	35.64%	19.32%	25.06%	2.23%	64.36%	35.64%	97.77%
XGB 200k	98.04%	20.45%	27.70%	23.53%	0.80%	79.55%	20.45%	99.20%
SVM 5k	52.93%	94.26%	2.87%	5.58%	47.69%	5.74%	94.26%	52.31%
SVM 10k	70.96%	86.27%	4.23%	8.06%	29.27%	13.73%	86.27%	70.73%
SVM 20k	85.36%	73.32%	7.05%	12.87%	14.46%	26.68%	73.32%	85.54%
SVM 50k	95.99%	40.97%	16.13%	23.15%	3.19%	59.03%	40.97%	96.81%
SVM 100k	98.04%	19.61%	27.26%	22.81%	0.78%	80.39%	19.61%	99.22%
SVM 200k	98.54%	17.28%	38.74%	15.53%	0.19%	92.11%	7.89%	99.81%
LR 5k	23.71%	93.14%	1.77%	3.48%	77.33%	6.86%	93.14%	22.67%
LR 10k	45.51%	83.61%	2.22%	4.33%	55.06%	16.39%	83.61%	44.94%
LR 20k	78.59%	54.41%	3.73%	6.97%	21.05%	45.59%	54.41%	78.95%
LR 50k	91.11%	31.72%	5.60%	9.52%	8.00%	68.28%	31.72%	92.00%
LR 100k	97.16%	14.64%	12.00%	13.19%	1.61%	85.36%	14.64%	98.39%
LR 200k	98.52%	1.68%	41.38%	3.23%	0.04%	98.32%	1.68%	99.96%
Ensemble	55.38%	91.74%	2.91%	5.64%	45.15%	8.26%	91.74%	54.85%
Cascade	82.26%	84.31%	6.63%	12.29%	17.77%	15.69%	84.31%	82.23%

5. Conclusion & Future Works

In this paper, a novel approach to bank fraud detection that integrates machine learning algorithms with clustering techniques is presented. The methodology prioritizes recall over accuracy, recognizing the severe consequences of false negatives in fraud detection. Through a cascade approach, multiple models trained on different fraud patterns are sequentially applied to classify transactions, thereby enhancing the detection accuracy. A comprehensive experiment was conducted using a dataset comprising bank transactions, demonstrating the efficacy of the approach in detecting fraudulent activities while maintaining high recall rates. The experimental results revealed that the cascade approach outperforms traditional classification algorithms, achieving a recall rate of 84.31% and an overall accuracy of 82.26%. By focusing on recall alongside accuracy, the method offers improved protection against financial losses due to undetected frauds. Moreover, the cascade strategy proved superior to the ensemble approach, providing further evidence of its effectiveness in detecting fraudulent transactions. In future works, several avenues for improvement and extension of the methodology can be explored. Firstly, refining the clustering technique to enhance the identification of distinct fraud patterns could lead to even better performance. Additionally, incorporating advanced feature engineering methods and exploring ensemble techniques tailored to the specific characteristics of fraudulent activities may further enhance detection capabilities. Furthermore, investigating the interpretability of the models and providing explanations for their decisions could enhance trust and transparency in the fraud detection process. Explainable AI techniques such as LIME and SHAP can be employed to elucidate the factors influencing model predictions, thereby aiding in decision-making and regulatory compliance. Finally, addressing the challenges of imbalanced datasets and rare event detection remains a crucial area for future research. By developing robust techniques to handle skewed class distributions and rare fraudulent events, our methodology can be further refined to deliver reliable and efficient fraud detection solutions in real-world banking scenarios.

Acknowledgments

This work was partially supported by project SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU.

References

- [1] Credit card fraud detection in the era of disruptive technologies: A systematic review, *Journal of King Saud University - Computer and Information Sciences* 35 (2023) 145–174. URL: <https://www.sciencedirect.com/science/article/pii/S1319157822004062>.
- [2] Deep learning for anomaly detection in log data: A survey, *Machine Learning with Applications* 12 (2023) 100470. URL: <https://www.sciencedirect.com/science/article/pii/S2666827023000233>.
- [3] G. Li, J. J. Jung, Deep learning for anomaly detection in multivariate time series:

Approaches, applications, and challenges, *Inf. Fusion* 91 (2022) 93–102. URL: <https://api.semanticscholar.org/CorpusID:252988274>.

- [4] G. Baader, H. Krcmar, Reducing false positives in fraud detection: Combining the red flag approach with process mining, *International Journal of Accounting Information Systems* 31 (2018) 1–16.
- [5] N. Capuano, G. Fenza, V. Loia, C. Stanzione, Explainable artificial intelligence in cybersecurity: A survey, *IEEE Access* 10 (2022) 93575–93600.
- [6] A. Y. B. R. Thaifur, M. A. Maidin, A. I. Sidin, A. Razak, How to detect healthcare fraud? “a systematic review”, *Gaceta sanitaria* 35 (2021) S441–S449.
- [7] A. Mehbodniya, I. Alam, S. Pande, R. Neware, K. P. Rane, M. Shabaz, M. V. Madhavan, Financial fraud detection in healthcare using machine learning and deep learning techniques, *Security and Communication Networks* 2021 (2021) 1–8.
- [8] B. Benedek, C. Ciumas, B. Z. Nagy, Automobile insurance fraud detection in the age of big data—a systematic and comprehensive literature review, *Journal of Financial Regulation and Compliance* 30 (2022) 503–523.
- [9] K. G. Al-Hashedi, P. Magalingam, Financial fraud detection applying data mining techniques: A comprehensive review from 2009 to 2019, *Computer Science Review* 40 (2021) 100402.
- [10] X. Zhu, X. Ao, Z. Qin, Y. Chang, Y. Liu, Q. He, J. Li, Intelligent financial fraud detection practices in post-pandemic era, *The Innovation* 2 (2021).
- [11] W. Hilal, S. A. Gadsden, J. Yawney, Financial fraud: a review of anomaly detection techniques and recent advances, *Expert systems With applications* 193 (2022) 116429.
- [12] R. Asha, S. K. KR, Credit card fraud detection using artificial neural network, *Global Transitions Proceedings* 2 (2021) 35–41.
- [13] R. Bin Sulaiman, V. Schetinin, P. Sant, Review of machine learning approach on credit card fraud detection, *Human-Centric Intelligent Systems* 2 (2022) 55–68.
- [14] D. Sarma, W. Alam, I. Saha, M. N. Alam, M. J. Alam, S. Hossain, Bank fraud detection using community detection algorithm, in: *2020 second international conference on inventive research in computing applications (ICIRCA)*, IEEE, 2020, pp. 642–646.
- [15] A. Vashistha, A. K. Tiwari, P. Singh, P. K. Yadav, S. Pandey, A robust framework for fraud detection in banking using ml and nn, *Proceedings of the National Academy of Sciences, India Section A: Physical Sciences* (2024) 1–12.
- [16] I. Sadgali, N. Sael, F. Benabbou, Performance of machine learning techniques in the detection of financial frauds, *Procedia computer science* 148 (2019) 45–54.
- [17] S. Daliri, et al., Using harmony search algorithm in neural networks to improve fraud detection in banking system, *Computational Intelligence and Neuroscience* 2020 (2020).
- [18] S. Jesus, J. Pombal, D. Alves, A. Cruz, P. Saleiro, R. P. Ribeiro, J. Gama, P. Bizarro, Turning the Tables: Biased, Imbalanced, Dynamic Tabular Datasets for ML Evaluation, *Advances in Neural Information Processing Systems* (2022).
- [19] D. Cavaliere, M. Gallo, C. Stanzione, Propaganda detection robustness through adversarial attacks driven by explainable ai, in: *World Conference on Explainable Artificial Intelligence*, Springer, 2023, pp. 405–419.