International Journal of Science and Research Archive

eISSN: 2582-8185

Cross Ref DOI: 10.30574/ijsra

Journal homepage: https://ijsra.net/

(REVIEW ARTICLE)

Check for updates

Cybersecurity threats in banking: Unsupervised fraud detection analysis

Karthik Meduri *

University of the Cumberlands, KY, USA.

International Journal of Science and Research Archive, 2024, 11(02), 915–925

Publication history: Received on 19 February 2024; revised on 28 March 2024; accepted on 30 March 2024

Article DOI: https://doi.org/10.30574/ijsra.2024.11.2.0505

Abstract

Customers all over the world now enjoy remarkable levels of accessibility and convenience thanks to the digital transformation of the banking industry. However, technology has also brought up new difficulties, including cybersecurity. The incapacity of conventional rule-based fraud detection strategies to keep up with the rapid evolution of cyber threats has generated interest in flexible and efficient approaches like unsupervised learning. The potential of unsupervised learning to improve fraud detection in the banking sector is examined in this article. The article addresses the disadvantages of traditional methods, the benefits of unsupervised learning, and how cybersecurity measures may be affected. A thorough framework for putting unsupervised fraud detection strategies into practice, including data preprocessing, feature engineering, isolation forest implementation, thresholding, and assessment, is provided in the methodology section. To further improve anomaly detection frameworks, future efforts propose integrating advanced machine learning techniques, dynamic thresholding, enhanced feature engineering, and continuous model monitoring. In summary, this essay offers useful insights on using modern machine learning algorithms to reduce cybersecurity threats and ensure the security of digital transactions within the banking industry.

Keywords: Cybersecurity; Banking; Fraud Detection; Unsupervised Learning; Machine Learning; Digital Transactions

1. Introduction

The digital transformation of the banking business has enabled customers worldwide to experience a new level of accessibility and convenience. However, this change is followed by a new set of difficulties, particularly in the cybersecurity field [1]. Financial institutions must be alert to the threats presented by attackers, who are always developing new strategies to take advantage of flaws in banking systems. Cyber risks, such as malware or phishing attacks, always develop in various ways [2]. These risks put financial organizations and their customers at high risk of identity theft, fraud, and data breaches [2]. Despite efforts to mitigate these risks, traditional rule-based fraud detection techniques have demonstrated shortcomings in keeping up with the dynamic nature of cyber threats. The identification and prevention of fraudulent activity is a major challenge in the field of financial security [3]. Efficient identification of fraudulent activities safeguards the financial interests of organizations and consumers' faith in the banking system. Maintaining transaction integrity, protecting private information, and guaranteeing adherence to legal requirements are essential.

Conventional techniques for detecting fraud mostly depend on established rules and signatures to identify suspicious activity. Though relatively effective, these techniques often find it difficult to adjust to new and advanced fraudulent practices [4]. As a result, the necessity for more flexible and efficient fraud detection methods is becoming more obvious. Unsupervised learning is a field within machine learning that presents an alternative to traditional rule-based techniques. Unsupervised learning works on unlabeled data, enabling computers to recognize patterns and anomalies without explicit instruction, unlike supervised learning, which requires labeled data for training [5]. Unsupervised learning's fundamental adjustability makes it ideal for real-time threat response and detecting fraudulent activity that

* Corresponding author: Karthik Meduri

has never been seen before [5]. The potential benefits of unsupervised learning to improve fraud detection in the banking industry are explored in this article. We discuss the disadvantages of traditional rule-based techniques, the advantages of unsupervised learning, and how these can affect financial organizations' cybersecurity defenses. Using machine learning and data analytics, we aim to illuminate a viable strategy for tackling the obstacles presented by cyber threats inside the banking industry.

2. Cyber Threats in Banking

Cybersecurity risks are a continuing concern to the banking industry because criminals are always coming up with new ways to take advantage of weaknesses in digital systems [2]. Sensible cyberattacks targeting banking systems have increased recently, putting financial organizations and their clients at serious risk. Many types of cyber threats target banking systems today, such as ransomware, phishing, malware attacks, and distributed denial-of-service (DDoS) attacks [6]. These attacks aim to steal money from victims, tamper with processes, and compromise confidential financial information [7]. A few notable incidents are ransomware attacks that destroy financial infrastructures, data breaches that lead to the loss of consumer information, and fraudulent transactions made possible by compromised credentials. Finding weaknesses in banking systems is essential to successfully reducing cyber threats. Common vulnerabilities include outdated software systems, weak authentication procedures, insufficient network security measures, and human errors. Strong security measures, including multi-factor authentication, encryption, intrusion detection, and frequent security audits, must be implemented to address these weaknesses.

Cybersecurity threats to the banking sector are evolving, and banks must take decisive action to safeguard their operations and client information [6]. Cyber threat classifications are essential as they provide insight into the potential nature and severity of threats that the financial sector might face, including categorizing threats by source, impact, and attack methods [8]. A thorough understanding of these classifications allows financial institutions to tailor countermeasures effectively. Indeed, robust countermeasures such as deploying firewalls and antivirus software, conducting employee training on cybersecurity best practices, and implementing comprehensive incident response plans are imperative. Moreover, collaborations with law enforcement and cybersecurity experts play a critical role in combating these threats. As the article [9] suggests, "the survival of every business largely depends on its customer base. Hence, this study propels the financial institutions as a reminder to strengthen their security and privacy concerns strategically". Therefore, by addressing cybersecurity risks proactively and in cooperation with broader security communities, banking institutions can bolster their defenses, enhance customer trust, and maintain the integrity of the financial systems amidst an ever-changing threat landscape. An active and adaptable strategy for addressing these issues is provided by integrating AI technologies, such as machine learning, automated incident response systems, natural language processing, and predictive analytics [10].

Figure 1 A hacker illustration representing cybersecurity risks and the requirement for strong anomaly detection protocols in financial systems [20]

3. Traditional Fraud Detection Methods

Predictive models are trained on labeled datasets using traditional fraud detection techniques, mainly based on supervised learning methodologies. Decision trees and logistic regression are two of these techniques that are widely used; each has unique advantages and disadvantages.

3.1. The Logistic Regression Model

As a statistical method for binary classification, logistic regression is ideally suited for fraud detection since the result is usually binary (fraudulent or non-fraudulent). Through the logistic function, the likelihood of the binary outcome is modeled in logistic regression as a function of independent factors (features) [11].

3.2. The benefits of using logistic regression

Interpretability: The log-odds ratio of the event's likelihood is represented by the logistic regression coefficients, which make the data interpretable and shed light on the significance of each parameter in fraud prediction [11]. Simple and Effective: Logistic regression is a popular option for binary classification jobs since it is easy to apply and computationally efficient [11]. Handles Linear Relationships: The logit of the outcome and the predictor variables are assumed linear by logistic regression, making it appropriate for capturing linear relationships in the data [11].

3.3. Logistic regression's limitations

Restricted Complexity: The only linear relationships that can be modeled by logistic regression are those between characteristics and the log odds of the result. Complex nonlinear correlations or interactions between features could be difficult to capture. Outlier Sensitivity: Data outliers can distort coefficient estimates and impair model performance in logistic regression [11]. Assumption of Independence: Time-series data and correlated observations are two real-world examples where the assumption of independence made by logistic regression may not hold.

3.4. Decision Trees

A non-parametric supervised learning technique called decision trees is applied to regression and classification problems. Decision trees create a tree-like structure where each internal node represents a decision based on a feature, and each leaf node represents a class label (fraudulent or non-fraudulent) [12]. Decision trees divide the feature space into segments recursively based on the values of input features.

Nonlinear Relationships: Decision trees are a good tool for modeling complex patterns in data because they can capture nonlinear relationships and interactions between features. Interpretability: Decision trees offer stakeholders an easy-to-understand visual representation of the decision-making process, which makes them interpretable. Decision trees can prioritize features according to how well they predict a result, which can shed light on the variables that motivate dishonest behavior [12].

3.5. Decision Tree Restrictions

Overfitting: Decision trees are more likely to overfit when the dataset is noisy, or the tree depth is not appropriately limited. Poor generalization performance on unseen data might result from overfitting [4]. Instability: The learned model may exhibit high variance and instability due to decision trees' sensitivity to even minute changes in the training set. Lack of Smoothness: Because decision borders in decision trees are inherently discontinuous, they can occasionally produce subpar performance, particularly in complicated decision trees.

In conclusion, decision trees and logistic regression are well-liked techniques for supervised fraud detection, each with advantages and disadvantages. Decision trees are superior at identifying nonlinear correlations and offering insights into the significance of features, but logistic regression is more efficient and interpretable. To increase the resilience and efficacy of fraud detection systems, it might be necessary to investigate alternate strategies or ensemble methodologies to overcome these approaches' shortcomings.

3.6. Federated Learning

Federated learning is a decentralized machine learning technique in which several parties train a model without directly exchanging data [13]. Instead, model updates are computed locally at each participating institution or on the device used by each user and then combined to create a global model. This method preserves data confidentiality and privacy while allowing a model to be trained across dispersed data sources.

Advantages

- Privacy Preservation: By allowing models to be trained on decentralized data sources instead of centralized data, federated learning protects user confidentiality and privacy [13].
- Data Diversity: Federated learning can capture a more comprehensive and broad variety of patterns and behaviors by utilizing data from numerous sources, which may enhance model performance and generalization [13].
- Scalability: Because model updates are computed locally and aggregated asynchronously, federated learning reduces data migration and centralized computation requirements, making it scalable to large and geographically distant datasets [13].

Difficulties

- Communication Overhead: Communication between the central server and participating devices or institutions is necessary to update models through federated learning. This communication can result in delays and communication overhead, especially in networks with low capacity or high latency [13].
- Consistency of Data: To ensure that the global model comes together, federated learning must consider the heterogeneity of data distributions across various devices or institutions. This may call for using techniques like data standardization or adaptive learning rates [13].
- Security dangers: Model poisoning and inference attacks, in which malicious parties try to influence the training procedure or deduce private information from model updates, are among the new security dangers introduced by federated learning [13].

3.7. Unsupervised Fraud Detection Techniques

Using innovative machine learning algorithms, unsupervised fraud detection approaches provide creative ways to fight financial crime without needing labeled datasets. This section covers anomaly detection techniques and gives an overview of unsupervised learning for fraud detection in cybersecurity and financial fraud detection. In fraud detection, unsupervised learning is a paradigm change that enables algorithms to find patterns and deviations in data without explicit supervision. Unsupervised learning techniques like clustering and anomaly detection reveal fraudulent behavior based on deviations from expected behavior [5], unlike supervised learning, which uses labeled instances of fraud for training. Unsupervised fraud detection relies heavily on anomaly detection techniques, which find anomalies or abnormalities in transactional data that point to fraudulent behavior [5]. These techniques cover many technologies, such as network-based detection systems, machine learning algorithms, and statistical approaches. Anomaly detection techniques help banks minimize financial losses and reputational harm via early identification and prevention of fraudulent transactions by highlighting odd patterns or behaviors. These systems are more accurate and efficient than conventional rule-based methods, enabling proactive fraud strategy detection and reaction [14]. Maintaining the effectiveness of fraud detection strategies requires that fraud detection tools be able to adjust to the constantly evolving and dynamic methods used by fraudsters [15]. This helps to protect financial assets and maintain customer trust. In conclusion, unsupervised fraud detection methods offer a revolutionary strategy for preventing financial fraud and boosting banking cybersecurity. Banking institutions can effectively detect and mitigate fraudulent activities, ensuring the integrity and security of digital transactions by leveraging the power of modern machine learning algorithms.

4. Methodology

4.1. Data Cleaning

Handling Missing Values: In real-world datasets, missing values frequently occur. They can result from several reasons, including insufficient data collection, technical issues during data entry, or the simple missing of important information. Missing values may affect machine learning model performance. Thus, managing them is important [16]. Imputation, which replaces missing values with estimated values based on further observations. Deletion, which eliminates rows or columns containing missing values, and considering missing values as a distinct category are methods for dealing with missing values.

Handling Outliers: Significantly different data points from the rest of the dataset are known as outliers. Outliers in financial data may point to unusual transactions or probably fraudulent activity. To keep outliers from distorting statistical analyses or impairing model performance, it is crucial to recognize and manage them. Trimming (removing extreme values), winsorization (changing extreme values with less extreme ones), and data transformation (making the data more normally distributed) are methods for handling outliers.

Noise reduction: Errors or random fluctuations in the data that can hide important patterns are called noise. When it comes to banking data, errors in measurement or recording could lead to noise. The signal-to-noise ratio can be increased by lowering or eliminating noise, making it simpler for machine learning algorithms to find important patterns. Among the methods for reducing noise are those for smoothing (such as moving averages or filters), dimensionality reduction (such as principal component analysis), or utilizing algorithms that are resistant to noise.

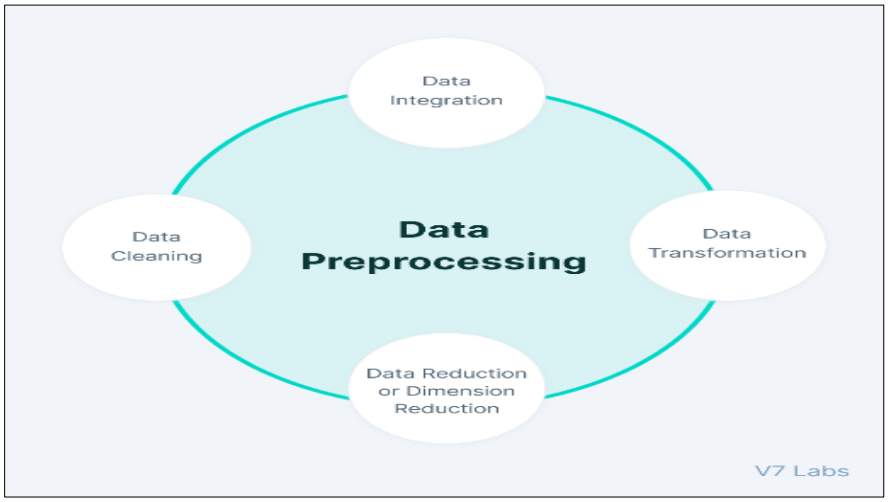

Figure 2 Data Preprocessing Procedure [21]

4.1.1. Standardization

Standardizing Numerical Features: Machine learning methods, especially those that depend on distance metrics (e.g., K-nearest neighbors, support vector networks), might face difficulties when dealing with numerical features that differ in scale and unit inside a dataset. To ensure that numerical features have equal scales, standardization rescales them with a mean of 0 and a standard deviation of 1. Algorithms become less sensitive to the magnitude of input features and converge more quickly due to this process. Analogously, normalization reduces numerical features to a range of 0 to 1.

4.1.2. Encoding Categorical Variables

Conversion to Numerical Representations: Gender and product type are categorical variables that must be transformed into numerical representations because many machine learning methods request numerical input data. A common method for completing this task is one-hot encoding, in which a binary vector representing the presence or absence of each category replaces each categorical variable. With this, it is made sure that categorical variables are represented in a way that is understandable to algorithms while also avoiding the introduction of accidental ordinal correlations between categories.

4.2. Feature Engineering

4.2.1. Extracting Relevant Features

Transaction Amount: One of the most important factors in detecting fraud is frequently the amount of money involved in a transaction. Unusual transaction amounts, whether big or small, might indicate fraud.

Frequency: A user's or account's transaction frequency may be a good indicator of typical activity. Variations in transaction frequency that occur suddenly could be signs of fraud.

Location: The physical location of transactions might offer important background information. For instance, transactions in a nation or area other than typical could cause suspicion.

Time: Transaction timing, such as the day of the week, month, and time of day, can highlight trends in user behavior. Fraudsters can commit fraud by taking advantage of specific timing patterns or anomalies.

Considering individual features and combinations that might provide more insights while extracting features is important. Finding relevant features that capture significant trends in the data requires domain expertise and exploratory data analysis.

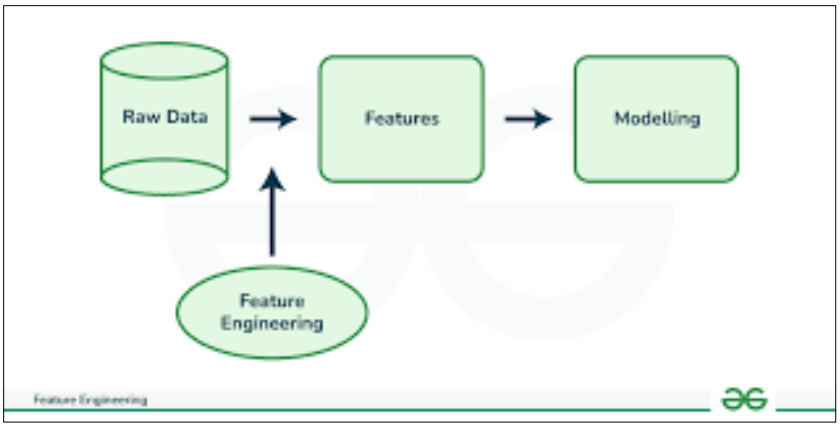

Figure 3 Feature Engineering steps [22]

4.2.2. Dimensionality Reduction

Principal Component Analysis, or PCA, is a widely used method for minimizing the dimensionality of high-dimensional data while maintaining an important portion of variance [17]. This is accomplished by converting the initial characteristics into a new collection of main components, which are orthogonal. The ranking of these components is based on how much variance they explain, which permits dimensionality reduction without sacrificing the majority of the crucial information present in the data [17].

The nonlinear dimensionality reduction method known as t-distributed Stochastic Neighbor Embedding, or t-SNE, is very helpful for visualizing high-dimensional data in lower-dimensional space [18]. By mapping high-dimensional data points to a lower-dimensional space and modeling pairwise similarities, t-SNE seeks to maintain local structure compared to PCA's focus on maintaining global structure [18]. Rather than being a stage in the preparation process for machine learning algorithms, t-SNE is frequently utilized for exploratory data analysis and visualization.

By concentrating on the most informative features or lowering the chance of overfitting, dimensionality reduction techniques like PCA and t-SNE can help streamline complicated datasets, cutting down on computational expenses and possibly even enhancing the performance of machine learning models. When using dimensionality reduction approaches, it's crucial to carefully weigh the trade-offs and consider the impact on model interpretability.

4.3. Isolation Forest Implementation

4.3.1. Hyperparameter selection

An effective technique for finding anomalies is isolation forest, especially when working with high-dimensional information [16]. For best results, though, hyperparameters must be adjusted. The following are the main hyperparameters in Isolation Forest: Number of Trees: The forest's tree count is determined by this property. The model's capacity to identify abnormalities may be enhanced by adding more trees, but doing so will add to its computational complexity. The dataset and the intended trade-off between processing power and performance determine the ideal number of trees [16]. Contamination Level: The percentage of outliers (anomalies) in the dataset is determined by this parameter. Usually, cross-validation or domain expertise is used to fine-tune it. In anomaly detection, regulating the balance between precision and recall requires setting a suitable degree of contamination [16].

4.3.2. Training the model

The preprocessed banking data is used to train the Isolation Forest model after the hyperparameters have been chosen. The model constructs an isolated tree forest and discovers the underlying patterns in the data during training [16]. By choosing a feature and a split value at random, each tree in the forest divides the feature space until every data point is isolated in a separate leaf node or reaches the maximum tree depth. Isolation Forest effectively detects anomalies as data points that require fewer partitions to separate, suggesting that they differ from the bulk of the data by utilizing the characteristics of isolation trees [16].

4.3.3. Anomaly score calculations

Every data point in the dataset is given an anomaly score following training, which measures how far it deviates from the average. Usually, the average path length (APL) in the isolation trees is used to compute anomaly scores [16]. Higher anomaly ratings are given to data points more likely to be abnormal and need shorter average path lengths to isolate. Anomaly scores range from higher to lower, indicating the likelihood of a given data point being an outlier or anomaly or closer to most of the data and less likely to be abnormal [16].

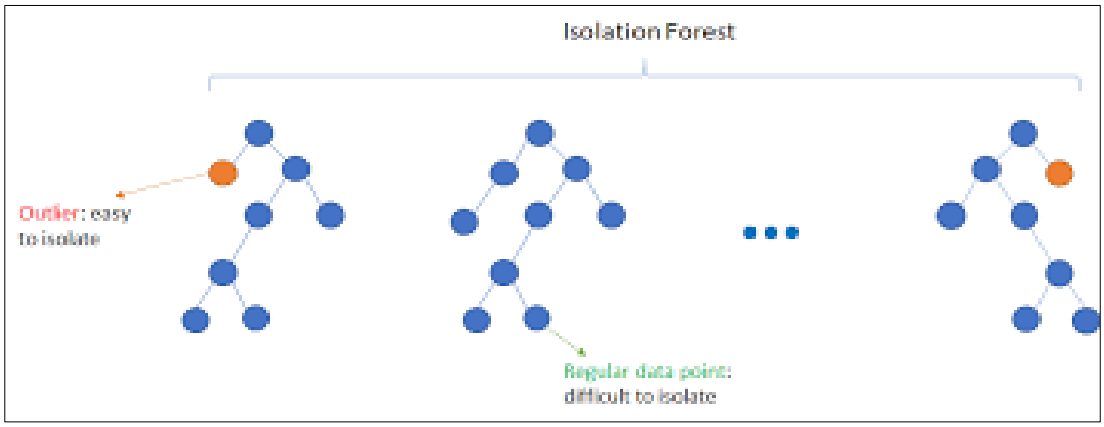

Figure 4 Depiction of Isolation Forest method [23]

4.4. Thresholding and Anomaly identification

4.4.1. Threshold selection

The anomaly score threshold is a cutoff point differentiating between abnormalities and regular data points. Data points below the threshold are considered normal, whereas those over the threshold are categorized as anomalies. Achieving a suitable threshold requires striking a compromise between precision and recall, two important variables. Precision is the percentage of anomalies that are accurately identified out of all the data points that are classed as anomalies. A high precision means that the model can correctly identify anomalies without mistakenly classifying normal data points as anomalies, indicating good false positive prevention. Recall, synonymous with sensitivity, measures the percentage of real anomalies the model properly detects. A high recall means that most true anomalies are captured by the model with a minimal number of errors. The threshold selection influences the accuracy-recall trade-off; raising the threshold results in higher precision but lower recall, and vice versa.

4.4.2. Impact of threshold selection on performance

According to [19], anomaly detection systems must effectively manage the imbalance between classes in datasets. The recommended approach in the article implements cost-sensitive methods to address this imbalance problem, which typically focuses on correctly identifying the rarer positive cases, such as fraudulent activities, without introducing significant bias [19]. In fraud detection scenarios, this implies carefully adjusting learning parameters and sometimes introducing sophisticated methods like imbalanced graph learners to ensure the model remains effective and unbiased [19]. With this in mind, in the context of threshold settings for anomaly detection, it would be appropriate to consider both the ratio of positive to negative samples and use evaluation metrics such as Macro AUC, Macro recall, and G-mean to gauge performance. Hence, while selecting thresholds, one must aim for a balanced strategy that addresses both class imbalance and the inherent challenges of fraud detection, using methodologies like the one discussed in the paper that go beyond simple threshold adjustments [19].

4.4.3. Threshold selection strategies

Domain Knowledge: The choice of a suitable threshold can be influenced by past knowledge of the data and the problem domain. Determining a threshold that is in line with the intended ratio of precision to recall can be facilitated by thoroughly understanding the typical distribution of anomaly scores and the business environment.

Receiver Operating Characteristic (ROC) curve: Plotting the genuine positive rate (recall) versus the false positive rate at different threshold levels is known as the Receiver Operating Characteristic (ROC) curve [16]. The area under the

ROC curve or AUC-ROC provides a measure of the model's overall performance across various thresholds. Selecting a threshold that maximizes the trade-off between precision and recall can be aided by ROC curve analysis.

Figure 5 ROC Curve [24]

Cost-Benefit Analysis: Determining a threshold that reduces the overall cost to the company can be aided by evaluating the expenses related to false positives and negatives [19]. This method balances the trade-off by considering the real-world effects of various decision outcomes.

4.5. Evaluation

4.5.1. Precision

The percentage of accurately recognized anomalies, or true positives, among all data points classified as anomalies is known as precision [12]. Regarding fraud detection, precision indicates the system's ability to accurately identify possibly fraudulent transactions without mistakenly classifying normal transactions as anomalies. Precision is calculated as the ratio of true positives (TP) to the sum of true and false positives [12]. The system is better at preventing needless warnings for valid transactions when there are fewer false positives, shown by a higher precision.

$$\text{Precision} = \frac{TP}{TP+FP}$$

4.5.2. Recall

The percentage of true anomalies (also known as true positives) the model successfully detects is measured by recall [12]. Recall in the context of fraud detection describes the system's capacity to identify most real fraudulent transactions while ignoring some. Recall is the ratio of true positives to the sum of true positives and false negatives [12]. A higher recall describes fewer missed fraudulent transactions without missing too many.

$$\text{Recall} = \frac{TP}{TP+FN}$$

4.5.3. F1-score

The F1-score is a metric that balances precision and recall by taking the harmonic mean of the two [16]. When there is an unequal class distribution between normal and abnormal examples, it is especially helpful. The F1-score is at its worst at 0 and highest at 1 (perfect recall and precision). The F1-score is calculated as:

$$\text{F1-score} = \frac{2 * Precision * Recall}{Precision+Recall}$$

4.5.4. ROC-AUC (Receiver Operating Characteristic - Area Under Curve)

The trade-off between true positive rate (recall) and false positive rate (FPR) at various threshold levels is measured by ROC-AUC [16]. Plotting the true positive rate (TPR) versus the false positive rate (FPR) at different threshold values is known as the ROC curve. The model's overall performance can be expressed as a single scalar value using the area under the ROC curve (AUC-ROC). A higher AUC-ROC value indicates better discriminating between positive and negative classes; a value of 1 denotes perfect discrimination, while 0.5 denotes random guessing. ROC-AUC is a helpful tool when

comparing various models or algorithms and assessing the effectiveness of binary classifiers, such as anomaly detection systems.

4.6. Sample size considerations

4.6.1. Data availability

The availability of banking transaction data directly impacts the effectiveness of anomaly detection models. More data improves generalization and performance by enabling the model to catch a larger range of patterns and anomalies. Training accurate models that differentiate between normal and abnormal behaviors requires sufficient data. More data will enable the model to identify a wider range of patterns and variations in transaction behavior, improving its capacity to identify anomalies. However, it is essential to guarantee the accuracy and applicability of the data. Data should be free from biases and inaccuracies that could negatively impact model performance and should be indicative of the underlying distribution of transactions [15].

4.6.2. Anomaly Prevalence

The sample size needed to identify abnormalities in a dataset effectively depends on how common they are. Larger sample sizes could be necessary for rare anomalies, like complex fraud schemes or extremely unusual transaction patterns, to guarantee proper representation in the dataset. A model may find it difficult to understand the patterns connected to uncommon abnormalities if they are not represented, which could increase the false negative rates [15].

4.6.3. Desired Statistical Power

The ability of a model to identify real effects or anomalies in the data when they happen is known as statistical power. Greater statistical power can be achieved through larger sample sizes, which lowers the possibility of false negative results and increases the possibility of identifying real anomalies. Building dependable anomaly detection models that precisely and confidently identify anomalies requires sufficient statistical power [15].

4.6.4. Computational Constraints

Computational challenges can occur with large datasets, especially with memory and processing capacity availability. Large amounts of transaction data processing and analysis may call for scalable computer architecture and effective algorithms. Large datasets may need methods like subsampling, which involves choosing a subset of the data or distributed computing, which involves processing data across several machines, to remain within computational bounds. It is crucial to balance model performance and computational efficiency so that the model can manage the available data and still fulfill operational needs [15].

4.6.5. Future Directions

It is crucial to continuously improve anomaly detection frameworks to keep ahead of new risks as the world of fraud and banking transactions changes. Going ahead, many directions can be investigated to improve the current framework:

Integration of Advanced Machine Learning Techniques: Although Isolation Forests are a powerful tool for detecting anomalies, there may be more depth to detecting complex patterns and anomalies in banking data if other advanced machine learning techniques, like deep learning models (like autoencoders), are integrated. These methods could improve the framework's performance and have demonstrated encouraging results in various anomaly detection jobs.

Dynamic Thresholding Techniques: The anomaly detection system's responsiveness and adaptability may be enhanced by creating dynamic thresholding techniques that adjust to shifting trends in banking transactions. For example, time-series analysis and reinforcement learning techniques might allow the framework to automatically modify anomaly score limits in response to changing transaction behavior and new fraud trends.

Improved Feature Engineering through Sophisticated Data Representation Methods: More informative representations of banking transactions might be produced using advanced feature engineering strategies and data representation approaches like graph-based representations or embedding techniques. These methodologies can potentially capture complex interrelationships and dependencies between transactional data points, resulting in enhanced anomaly detection performance and more discriminative feature representations.

Continuous Model Monitoring and Feedback Loop: A strong framework for continuous model monitoring and feedback loop mechanisms must be established to ensure that the anomaly detection system remains effective over time. The

platform may be made to be flexible and responsive to changing fraud threats in the banking industry by putting in place automatic model retraining pipelines, feedback methods for adding new labeled data, and real-time monitoring of model performance indicators.

5. Conclusion

In conclusion, clients worldwide now enjoy unmatched levels of accessibility and convenience thanks to the digital transformation of the banking industry. However, this development has also brought up new difficulties, especially cybersecurity. Financial institutions must be on guard against the constantly changing risks posed by cybercriminals, who are always coming up with new ways to break into financial systems. More flexible and efficient fraud detection methods are now more important than ever because traditional rule-based procedures have shown themselves unable to keep up with the constantly evolving nature of cyber threats. A possible alternative is unsupervised learning, a branch of machine learning that allows computers to identify patterns and abnormalities in data without explicit guidance. The disadvantages of traditional rule-based methods and the benefits of unsupervised learning for fraud detection in the banking sector were discussed in this article. We discussed the changing nature of cybersecurity risks in the banking industry, the weaknesses of conventional fraud detection techniques, and the possible advantages of using unsupervised learning algorithms. Financial organizations can improve cybersecurity defenses and prevent fraud by implementing unsupervised fraud detection approaches like isolation forests. Banks can keep ahead of developing risks and maintain the integrity and security of their digital transactions by maintaining effective model monitoring and feedback procedures, integrating advanced machine learning approaches, and continuously upgrading anomaly detection frameworks. Using unsupervised learning has much potential to safeguard financial assets, increase fraud detection skills in the banking industry, and maintain consumer confidence in the digital age of banking. Financial institutions must adopt modern technologies and maintain an active cybersecurity strategy to stay one step ahead of cybercriminals as long as cyber threats persist.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

[1] M. Bhasin, "Frauds in the Banking Sector: Experience of a Developing Country Asian Journal of Social Sciences and Management Studies Frauds in the Banking Sector: Experience of a Developing Country," vol. 3, no. 1, pp. 1–9, 2016.

[2] A. Q. Stanikzai and M. A. Shah, "Evaluation of Cyber Security Threats in Banking Systems," 2021 IEEE Symposium Series on Computational Intelligence (SSCI), Dec. 2021, doi: https://doi.org/10.1109/ssci50451.2021.9659862.

[3] H. Taherdoost, "A Review on Risk Management in Information Systems: Risk Policy, Control and Fraud Detection," Electronics, vol. 10, no. 24, p. 3065, Dec. 2021, doi: https://doi.org/10.3390/electronics10243065.

[4] X. Niu, L. Wang, and X. Yang, "A Comparison Study of Credit Card Fraud Detection: Supervised versus Unsupervised.," arXiv (Cornell University), Apr. 2019, [Online]. Available: http://export.arxiv.org/pdf/1904.10604

[5] [5] R. J. Bolton and D. J. Hand, "Unsupervised Profiling Methods for Fraud Detection," Jan. 2002, [Online]. Available: http://procon.bg/article/unsupervised-profiling-methods-fraud-detection

[6] A. A. Darem, A. A. Alhashmi, T. M. Alkhaldi, A. M. Alashjaee, S. M. Alanazi, and S. A. Ebad, "Cyber Threats Classifications and countermeasures in banking and financial sector," IEEE Access, vol. 11, pp. 125138–125158, Jan. 2023, doi: 10.1109/access.2023.3327016.

[7] D. Ghelani, T. K. Hua, and S. K. R. Koduru, "Cyber security threats, vulnerabilities, and security solutions models in banking," Authorea (Authorea), Sep. 2022, doi: 10.22541/au.166385206.63311335/v1.

[8] A. Q. Stanikzai and M. A. Shah, "Evaluation of cyber security threats in banking systems," 2021 IEEE Symposium Series on Computational Intelligence (SSCI), Dec. 2021, doi: 10.1109/ssci50451.2021.9659862.

[9] A. B. Jibril, M. A. Kwarteng, M. Chovancová, and R. Denanyoh, "Customers' perception of cybersecurity threats toward e-banking adoption and retention: A conceptual study," 15th International Conference on Cyber Warfare and Security - ICCWS 2020, Jan. 2020, doi: 10.34190/iccws.20.020.

[10] H. Gonaygunta, G. Sandeep Nadella, K. Meduri, D. Priyanka Pramod Pawar, and Kumar, "The Detection and Prevention of Cloud Computing Attacks Using Artificial Intelligence Technologies," International Journal of Multidisciplinary Research and Publications (IJMRAP), vol. 6, no. 8, pp. 191–193, 2024.

[11] H. Gonaygunta, "MACHINE LEARNING ALGORITHMS FOR DETECTION OF CYBER THREATS USING LOGISTIC REGRESSION," International Journal of Smart Sensors and Ad Hoc Networks, pp. 36–42, Jan. 2023, doi: 10.47893/ijssan.2023.1229.

[12] S. Khatri, A. Arora, and A. P. Agrawal, "Supervised Machine Learning Algorithms for Credit Card Fraud Detection: A Comparison," 10th International Conference on Cloud Computing, Data Science & Engineering (Confluence), Jan. 2020, doi: 10.1109/confluence47617.2020.9057851.

[13] D. Kumar, P. Pawar, H. Gonaygunta, and S. Singh, "Impact of Federated Learning on Industrial IoT - A Review," IJARCCE, vol. 13, no. 1, Dec. 2023, doi: https://doi.org/10.17148/ijarcce.2024.13105.

[14] V. Agaskar, M. Babariya, S. Chandran, and N. Giri, "Unsupervised learning for credit card fraud detection," International Research Journal of Engineering and Technology (IRJET), vol. 4, no. 3, Mar. 2017.

[15] G. Cabanès, Y. Bennani, and N. Grozavu, "Unsupervised Learning for Analyzing the Dynamic Behavior of Online Banking Fraud," 2013 IEEE 13th International Conference on Data Mining Workshops, Dec. 2013, doi: 10.1109/icdmw.2013.109.

[16] F. T. Liu, K. M. Ting, and Z. Zhou, "Isolation Forest," Dec. 2008, doi 10.1109/icdm.2008.17.

[17] S. Wold, K. H. Esbensen, and P. Geladi, "Principal component analysis," Chemometrics and Intelligent Laboratory Systems, vol. 2, no. 1–3, pp. 37–52, Aug. 1987, doi: 10.1016/0169-7439(87)80084-9.

[18] G. C. Linderman, M. Rachh, J. G. Hoskins, S. Steinberger, and Y. Kluger, "Efficient Algorithms for t-distributed Stochastic Neighborhood Embedding," Europe PMC (PubMed Central), Dec. 2017, [Online]. Available: http://europepmc.org/articles/pmc6402590

[19] X. Hu et al., "Cost-Sensitive GNN-Based Imbalanced learning for mobile social network fraud detection," IEEE Transactions on Computational Social Systems, pp. 1–16, Jan. 2024, doi: 10.1109/tcss.2023.3302651.

[20] N. Morpus, "What are the Methods and Motives for Hacking?," VMware Security Blog, Jul. 26, 2022. https://blogs.vmware.com/security/2022/07/what-are-the-methods-and-motives-for-hacking.html

[21] P. Baheti, "A Simple Guide to Data Preprocessing in Machine Learning," www.v7labs.com, Aug. 31, 2021. https://www.v7labs.com/blog/data-preprocessing-guide

[22] "What is Feature Engineering?" GeeksforGeeks, Mar. 20, 2023. https://www.geeksforgeeks.org/what-is-feature-engineering/

[23] N. E. Kın, "What are Isolation Forests?," Medium, Jun. 07, 2022. https://engineering.teknasyon.com/what-are-isolation-forests-151d8e98ef5f

[24] AI Team, "How to explain the ROC AUC score and ROC curve?," www.evidentlyai.com. https://www.evidentlyai.com/classification-metrics/explain-roc-curve

[25] "14,926 Cyber Security Polygon Royalty-Free Photos and Stock Images," Shutterstock. https://www.shutterstock.com/search/cyber-security-polygon (accessed Feb. 28, 2024)