JOURNAL OF MANAGEMENT AND BUSINESS SOLUTIONS

Money Laundering Detection based on Data Mining and Classification using Deep Learning

1. Mehdi. Shakeri Behbahani : Department of Management, Na.C., Islamic Azad University, Najafabad, Iran.
2. Mehdi. Sadeghzadeh*: Department of Computer Engineering, SR.C., Islamic Azad university, Tehran, Iran.
3. Naser. Khani: Department of Management, Na.C., Islamic Azad University, Najafabad, Iran.
3. Akbar. Nabiollahi: Department of Computer Engineering, Na.C., Islamic Azad University, Najafabad, Iran.

*corresponding author's email: mehdi.sadeghzadeh@iau.ac.ir

ABSTRACT

In recent years, and with the rapid advancement of technology, an emerging phenomenon called electronic banking has emerged that affects the lives of all people. Along with the many benefits that electronic banking has brought to people, the field of fraud and money laundering has also entered this field, which helps fraud in banking transactions to be carried out very accurately and systematically in the field of electronic banking. What is important is a precise, fast and calculated confrontation with profiteers who move large amounts of money in this way and enter it into their personal accounts with the utmost intelligence. In this article, while introducing data mining and its techniques, we analyze and discuss the measures taken in the field of fraud detection in various sectors, especially banking. After that, we introduce 3 types of the most widely used data mining algorithms and implement them on bank transactions in the CCFD benchmark dataset. Finally, after examining the performance of the algorithms, the best algorithm in terms of its performance in fraud detection is introduced.

Keywords: Data mining, fraud detection, bank transactions, deep learning, deep neural networks

Introduction

Money laundering has become one of the most complex and challenging financial crimes in the contemporary global economy. The rapid expansion of digital financial systems, electronic banking infrastructures, online payment platforms, and cross-border financial transactions has significantly increased the volume and velocity of financial activities, creating new opportunities for economic growth while simultaneously expanding the avenues available for illicit financial operations. Criminal organizations increasingly exploit technological advancements to disguise the origins of illegally obtained funds and integrate them into legitimate financial systems, making detection and prevention considerably more difficult than in traditional banking environments (1, 2). Consequently, financial institutions, regulatory authorities, and law enforcement agencies face mounting pressure to develop sophisticated mechanisms capable of identifying suspicious transactions and preventing the circulation of illicit capital within the global financial ecosystem.

Article history:
Received 25 January 2024
Revised 21 May 2024
Accepted 28 May 2024
Initial Publish 05 June 2024
Published online 01 January 2024

How to cite this article:
Shakeri Behbahani, M., Sadeghzadeh, M., Khani, N., & Nabiollahi, A. (2027). Money Laundering Detection based on Data Mining and Classification using Deep Learning. *Journal of Management and Business Solutions, 5*(1), 1-16. https://doi.org/10.61838/jmbs.344

© 2027 the authors. This is an open access article under the terms of the Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0) License.

The economic and social consequences of money laundering extend far beyond individual financial losses. Money laundering facilitates organized crime, terrorism financing, corruption, tax evasion, drug trafficking, and numerous other illegal activities that undermine economic stability and public trust in financial institutions. Traditional anti-money laundering (AML) frameworks have historically relied on rule-based monitoring systems and manual investigations conducted by compliance officers. While these approaches have provided a foundational mechanism for identifying suspicious transactions, they often struggle to cope with the enormous volumes of data generated by modern financial systems and the increasingly sophisticated strategies employed by money launderers (3, 4). As a result, there has been a growing interest in leveraging advanced analytical techniques and artificial intelligence technologies to improve the effectiveness and efficiency of AML operations.

The emergence of data mining technologies has fundamentally transformed the manner in which financial institutions analyze transactional information. Data mining enables organizations to extract meaningful patterns, relationships, and anomalies from large datasets, thereby facilitating the identification of potentially fraudulent or suspicious activities. Early studies in fraud detection demonstrated that data mining techniques could successfully uncover hidden structures within transaction data that might otherwise remain undetected through conventional auditing procedures. These approaches became particularly valuable in financial environments characterized by large-scale transactional databases and complex customer behavior patterns (5). The ability to automatically analyze millions of transactions and identify unusual behavioral patterns represented a major advancement in financial crime detection.

However, the application of data mining to money laundering detection presents several methodological challenges. One of the most significant difficulties is the highly imbalanced nature of financial transaction datasets. In real-world financial systems, suspicious or illicit transactions typically constitute only a tiny fraction of the overall transaction volume. As a result, machine learning algorithms trained on such datasets may become biased toward the majority class and fail to identify rare but critically important money laundering events. This phenomenon, commonly referred to as the class imbalance problem, has been extensively documented in the machine learning literature and remains one of the central obstacles to effective fraud detection systems (6, 7). The disproportionate distribution between legitimate and suspicious transactions often leads to models with high overall accuracy but poor performance in detecting illicit activities.

Researchers have proposed various techniques to address the class imbalance problem. Among these approaches, sampling strategies have attracted considerable attention due to their effectiveness in improving classifier sensitivity toward minority classes. Under-sampling methods reduce the number of majority-class observations, while over-sampling techniques generate additional minority-class instances to achieve a more balanced dataset. Comparative investigations have demonstrated that under-sampling methods can often outperform over-sampling approaches in highly imbalanced environments because they reduce computational complexity while maintaining sufficient discriminatory information for classification tasks (8). Furthermore, probability calibration techniques specifically designed for unbalanced classification problems have shown promising results in enhancing predictive performance and reducing classification bias (9).

The growing availability of computational resources and advanced machine learning frameworks has accelerated the adoption of artificial intelligence technologies within the financial sector. Traditional machine learning algorithms such as decision trees, support vector machines, logistic regression, random forests, and k-nearest neighbors have been widely applied to fraud detection and AML tasks. These methods have demonstrated varying degrees of

success depending on the characteristics of the underlying datasets and the specific nature of suspicious financial activities. Comprehensive reviews of money laundering detection systems indicate that machine learning models generally outperform conventional rule-based approaches because they can adapt to evolving criminal behaviors and discover complex patterns that may not be explicitly encoded within predefined rules (3, 4).

More recently, deep learning has emerged as a particularly powerful branch of artificial intelligence with significant implications for financial crime detection. Deep learning models utilize multiple layers of artificial neural networks to automatically learn hierarchical representations of data. Unlike traditional machine learning algorithms that often require extensive feature engineering, deep learning architectures can extract increasingly abstract and informative features directly from raw data. This capability enables them to model complex nonlinear relationships and identify subtle patterns that may be indicative of money laundering activities (10, 11). The success of deep learning in domains such as image recognition, natural language processing, speech recognition, and anomaly detection has encouraged researchers to explore its applicability within financial crime analytics.

Artificial neural networks constitute the foundation of deep learning systems. These computational models are inspired by the structure and functioning of biological neural networks and consist of interconnected processing units organized into layers. Modern neural network architectures benefit from advances in optimization techniques, activation functions, and regularization strategies that significantly enhance their predictive capabilities. Theoretical and practical developments in neural network design have established robust frameworks for constructing models capable of handling complex classification problems across diverse application domains (12). In particular, the introduction of Rectified Linear Unit (ReLU) activation functions has substantially improved training efficiency and convergence behavior in deep neural networks, enabling the development of deeper and more expressive architectures (13).

One of the critical challenges associated with deep learning models is overfitting, especially when training datasets are limited or highly imbalanced. To address this issue, researchers have developed several regularization techniques that improve model generalization. Among the most influential of these methods is dropout, which randomly deactivates subsets of neurons during training to prevent excessive dependence on specific network pathways. Empirical evidence has consistently demonstrated that dropout significantly enhances predictive performance and reduces overfitting in deep neural networks across a wide range of applications (14). Such advancements have made deep learning increasingly attractive for AML applications where model robustness and adaptability are essential.

The effective training of deep neural networks also depends on sophisticated optimization algorithms capable of efficiently navigating high-dimensional parameter spaces. Stochastic Gradient Descent (SGD) and its variants have become the standard optimization methods for large-scale machine learning systems because of their computational efficiency and scalability. These optimization strategies enable neural networks to learn complex patterns from massive datasets while maintaining practical training times (15). Concurrently, the development of open-source machine learning frameworks and software libraries has facilitated the implementation and deployment of advanced deep learning models in both research and industrial settings. Platforms such as Scikit-learn provide accessible tools for machine learning experimentation and evaluation, while TensorFlow supports large-scale deep learning applications and distributed computing environments (16, 17).

In the context of fraud and money laundering detection, deep learning approaches have demonstrated increasingly promising outcomes. Research examining fraud detection in banking and telecommunications

environments has shown that deep neural networks can achieve superior performance relative to conventional machine learning methods. For example, generative adversarial network-based models have been successfully employed to detect fraudulent telecommunications transactions, illustrating the capacity of deep learning systems to identify sophisticated fraudulent behaviors in highly dynamic environments (18). Similar findings have encouraged the application of deep learning methodologies to AML systems, where the complexity and evolving nature of money laundering schemes require highly adaptive analytical solutions.

Recent studies specifically focusing on anti-money laundering have reinforced the value of machine learning and deep learning techniques. Alotibi et al. demonstrated the effectiveness of deep neural networks and other machine learning classifiers in detecting money laundering activities within cryptocurrency environments, highlighting substantial improvements in classification performance and reductions in false-positive rates (19). Similarly, Samaddar proposed a deep learning-based AML framework and reported superior performance compared with several traditional machine learning algorithms, emphasizing the potential of deep architectures for identifying suspicious financial transactions (20). These findings suggest that deep learning models may provide a more effective mechanism for monitoring increasingly complex financial systems.

Beyond purely technical considerations, contemporary scholarship has also begun to explore the integration of criminological theories into machine learning-based AML frameworks. Rather than relying exclusively on statistical patterns, these approaches seek to incorporate theoretical insights regarding criminal behavior, decision-making processes, and money laundering mechanisms. Such integrations may enhance model interpretability and predictive effectiveness by aligning algorithmic detection processes with established criminological knowledge. Empirical evidence indicates that theory-informed machine learning frameworks can outperform conventional predictive models across multiple evaluation metrics, thereby offering a promising direction for future AML research and practice (1).

Despite these advances, significant challenges remain in developing robust and scalable money laundering detection systems. Financial institutions continue to grapple with issues related to data quality, class imbalance, model interpretability, false-positive rates, regulatory compliance, and the evolving tactics employed by criminal organizations. Consequently, there remains a critical need for empirical studies that evaluate the effectiveness of deep learning models in detecting suspicious financial transactions and compare their performance against established classification approaches. Such investigations can contribute to both the academic literature and practical AML implementation by identifying methodologies capable of enhancing detection accuracy while minimizing operational costs and compliance burdens.

Therefore, the aim of this study is to evaluate the effectiveness of deep learning-based classification models for money laundering detection and compare their performance with conventional machine learning approaches in identifying suspicious financial transactions.

Methods and Materials

In this study, the library research method is used to design an algorithm that is capable of automatically identifying illegal transactions in a financial payment service. In this regard, the detection of fraud in transactions is formulated as a binary classification problem in which a vector of features and a class label correspond to each of the transactions. Typically, the datasets used in such studies are highly unbalanced because illegal transactions only account for a small fraction of the total transactions. However, the main focus is on the class with a smaller

population (illegal transactions). This further highlights the need for proper management of the dataset. In this paper, the resampling method will be used to overcome this problem.

Since the detection of illegal transactions is the main objective of this study, limited sampling will be used to create a balance between the two different classes. In this method, the class with the largest population is sampled in a limited manner so that the final ratio of both classes is equal to each other. Finally, the new dataset will be used as a benchmark dataset for training and testing various machine learning algorithms. In the following, we will introduce the deep learning method and describe the algorithm.

Deep learning, as a machine learning method, seeks to solve problems that other learning methods such as support vector machines typically have difficulty solving due to the nature of their shallow architecture. In this method, a set of features is extracted from the training data using a multilayer architecture of nonlinear processors that are statistically robust.

In general, one of the main problems in machine learning is the selection of an appropriate feature space. In such a way that the input data has the desired characteristics to solve a specific problem. For example, in supervised byte classification, it is often necessary to separate two classes from each other with the help of a hyperplane. In the case where such a feature is not directly obtainable in the input space, it is assumed that this data can be mapped to an intermediate feature space in which the classes are linearly separable. This intermediate space can be either directly specified with manually selected features.

Or indirectly defined with a kernel function or automatically learned. In the first two cases, the design of the feature space is the responsibility of the user. This can be expensive in terms of computational time and knowledge required, especially when the input space is of high dimension. In the third case, automatic feature learning with deep architectures consisting of multiple layers of nonlinear processors can be considered as an acceptable option.

In other words, in deep learning, different layers of processors try to represent the data at a higher, more abstract level by mapping it. By combining a sufficient number of such mappings, very complex functions can be learned. In particular, in the classification problem, higher representation layers highlight aspects of the input data that are important for distinguishing classes and removing outliers. Thus, the algorithm is considered as a deep artificial neural network using the deep learning method, which consists of six hidden layers. Figure 1 shows the deep architecture of this algorithm. The linear rectifier function (ReLU) is used for the activity function of neurons in the hidden layers, and the neurons in the output layer also use the sigmoid function as their activity function, and the decision threshold in the output layer is considered to be 0.5. In addition, the initial weights of the neurons follow a uniform random distribution.

In order to prevent overfitting of the proposed deep network, a deletion strategy has been used to train this network. In this strategy, each time the weights are updated, a part of the neurons of each hidden layer is randomly selected and deleted. It should be noted that this deletion is temporary and after the end of the training stage, all neurons will be present in the network and only their weights are modified in such a way that the final network provides an output approximately equivalent to the trained network.

For more information about this method, refer to the article by Srivastava et al. (2014). Finally, the stochastic gradient descent (SGD) method is used to train the proposed deep network.

Findings and Results

In this section, we will examine and evaluate the performance of the proposed deep network. For this purpose, we will use the CCFD benchmark dataset, which is available at https://datahub.io/machine-learning/creditcard. This dataset consists of 284,800 transactions registered by European customers over two days in September 2013. Among them, there are 492 transactions that are labeled as illegal, accounting for 0.172% of the total transactions in the benchmark dataset, indicating a severe imbalance in this data. To maintain confidentiality, the main features and background information about them are not provided. Only the numerical features of this dataset are available, and that too after anonymization. These numerical features, obtained by applying principal component analysis (PCA) to the main features of the benchmark dataset, are labeled with labels 1 to V28. The only features that are not changed by PCA are the time and number of transactions. The time feature indicates the number of seconds that have passed between each transaction and the first transaction in the benchmark dataset. The class feature expresses the response variable or label of each transaction and takes the value of one if the transaction is illegal and zero if it is legal. It should be noted that the scikit-learn and tensorflow libraries in Python were used to perform the numerical experiments presented in this section.

The training time of the proposed deep network on a computer with a quad-core processor with 3.3 GHz processing power and 4 GB of RAM is about one minute. In addition, readers can contact the authors of the article to access the final trained network.

To prepare the data, first all 30 mentioned features, including features V1 to V28, time and transaction rate are separately normalized in such a way that the average of each feature of each of its elements is reduced and divided by the standard deviation of that feature. Then, using limited sampling, a balanced dataset consisting of 984 transactions is obtained. For this purpose, first 492 legal transactions are randomly selected from the legal transaction set and then combined with 492 illegal transactions in the benchmark dataset to obtain a balanced dataset with a ratio of 50% of the two types of transactions. Next, this balanced dataset is randomly divided into training and testing datasets in a ratio of 7:3. Thus, the number of 688 transactions will be used for learning and 296 transactions will be used for testing the algorithm.

We enter the pre-processed data into the Deep Neural Network algorithm in the RapidMiner environment, as shown in Figure 1.

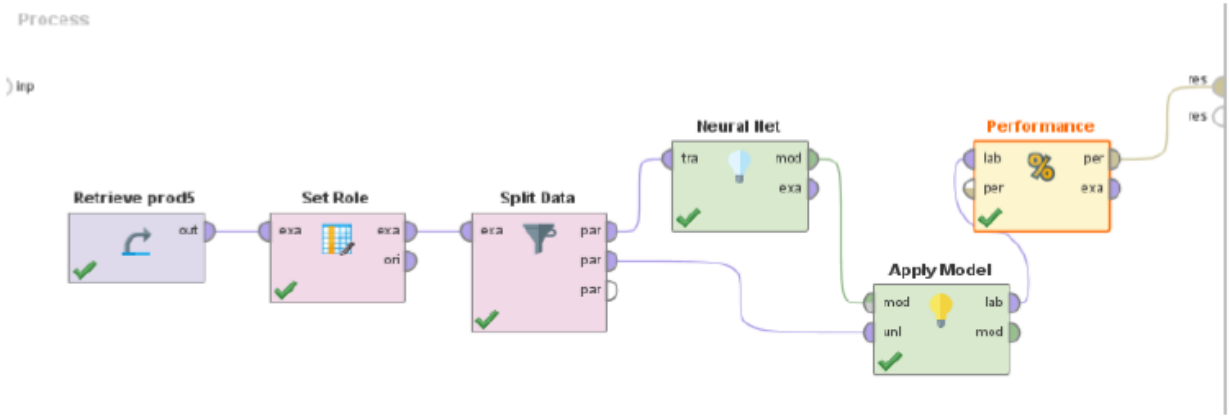

Figure 1: Implementation of Deep Neural Networks Algorithm on Data

The above layout shows the use of Deep Neural Networks algorithm on our data. In this algorithm, our hidden layer size is 2. The result of running the above algorithm is given below:

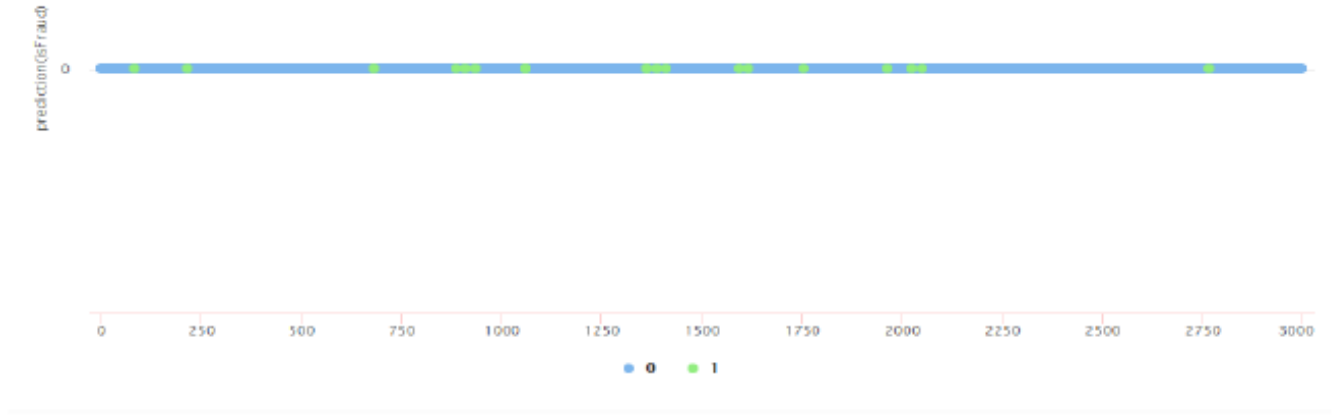

Figure 2 Output of applying the deep neural network algorithm to the data

At the end of the work, we also call the performance table to measure the performance of the algorithm:

Table 1: Result of applying the deep neural network algorithm to the data

Class precision	True1	True0	
99.33%	20	2979	Pred0
0.00%	0	0	Pred1
	0.00%	100%	Class recall

According to Table 1, the high performance of the decision tree algorithm correctly identified 2979 transactions that were not fraudulent. There was also no transaction that was mistakenly considered as a transaction that was not fraudulent. 20 transactions were mistakenly identified as fraudulent. Finally, there were no transactions that were fraudulent and the algorithm correctly identified fraud.

The accuracy of the decision tree algorithm on our data was calculated to be 99.33%.

Given the limitations we found in this algorithm, we enter our data into the cycle. The process of implementing the algorithm in the Rapidminer environment and applying the nearest neighbor algorithm is as follows:

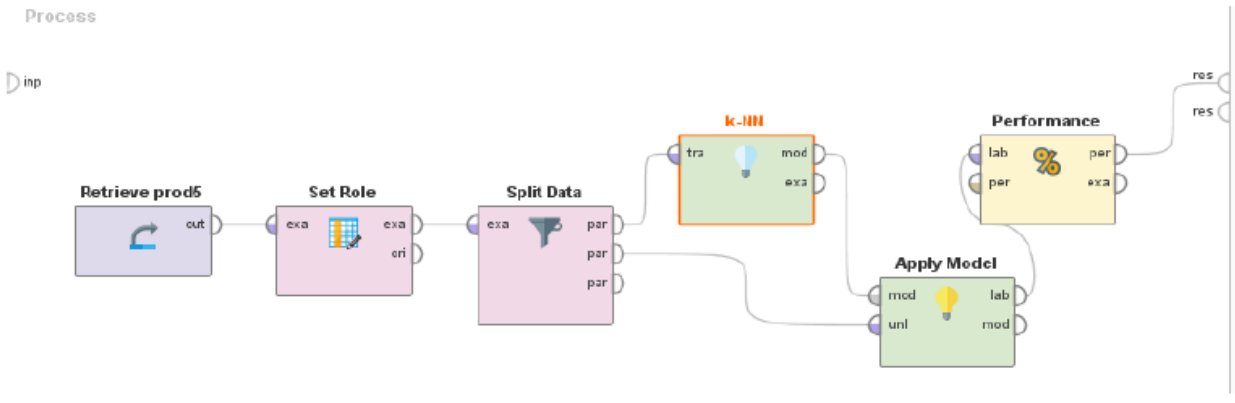

Figure3: Implementation of the k-nearest neighbor algorithm on data in Rapidminer

In choosing the best value of K, considering the different values we tried for it, the number 6 represented the best result for our work. Because the prediction rate in the cheating and non-cheating cases was better than the other cases.

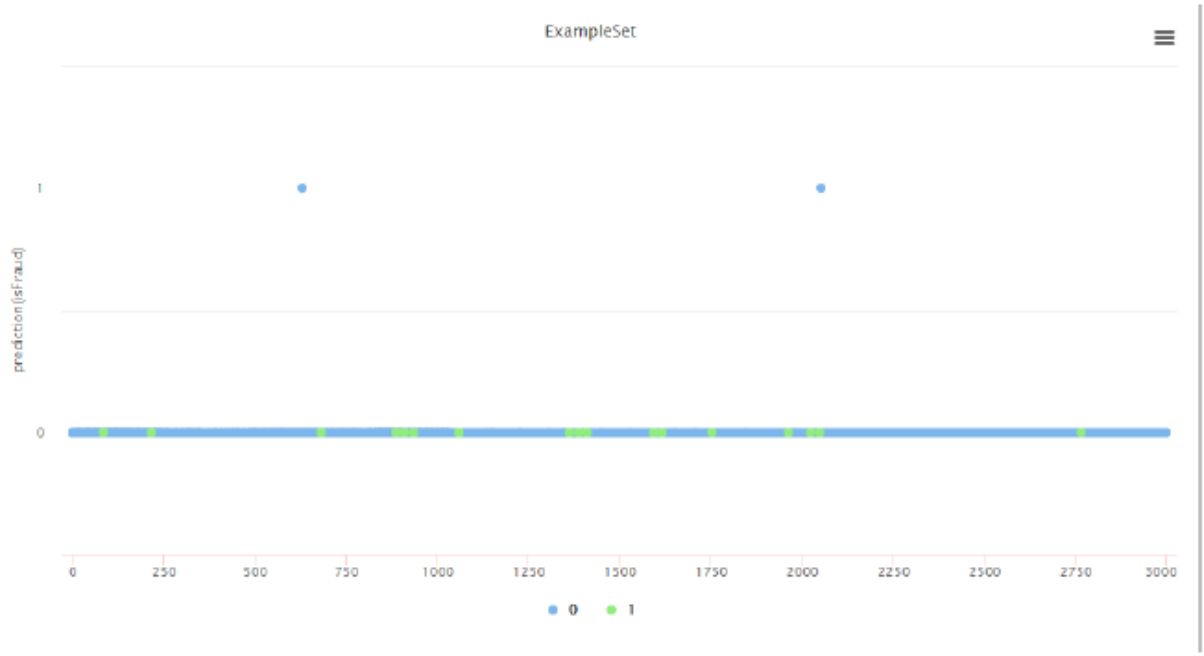

Figure 4: Output of implementing the nearest neighbor algorithm on the data

In the image above, out of the 3000 transactions that we entered into the algorithm, the horizontal axis shows the number of transactions, or rather the unique number of each transaction. The vertical axis shows the prediction of the algorithm regarding fraud or non-fraud of the transaction. The blue color indicates that there was no fraud and the green color indicates the presence of fraud. For example, in row 1 of the diagram, where we have two blue dots, the prediction was that there should be fraud in these two transactions, but in reality, we did not have fraud. But in the row of blue dots, it is said that in our algorithm's prediction there was no fraud and in reality, we did not have such a case. But as you can see, there are some green dots at different distances in the same row, which indicate that our algorithm's prediction is based on the presence of fraud, but in reality, we did not have any fraud in this number of transactions and the algorithm made a mistake. For a general consensus on the performance of the nearest neighbor algorithm, we also add the performance of the tool, the results of which are listed below.

Table 2: Results of implementing the nearest neighbor algorithm on the data

Class precision	True1	True0	
99.33%	20	2977	Pred0
0.00%	0	2	Pred1
	0.00%	99.93%	Class recall

According to the performance table above, the decision tree algorithm correctly identified 2977 transactions that were not fraudulent, 2 transactions were incorrectly considered as transactions that were not fraudulent, 20 transactions were incorrectly identified as transactions that were fraudulent, and finally, there were no transactions that were fraudulent and the algorithm correctly detected fraud. The accuracy of the decision tree algorithm on our data was calculated to be 99.27%.

We have reviewed some information about the Support Vector Machine algorithm with an example, we will put our data into the cycle. The process of implementing the algorithm in the Rapid Miner environment and applying the mentioned algorithm is shown below.

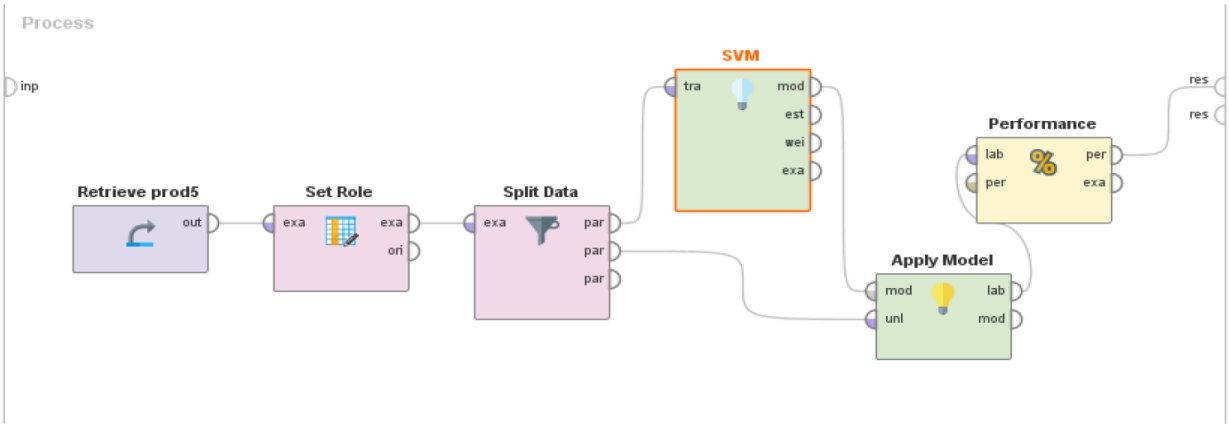

Figure 5: Implementation of the support vector machine algorithm on data in Rapidminer

After running the algorithm, the following results were obtained for each transaction, which we will interpret below:

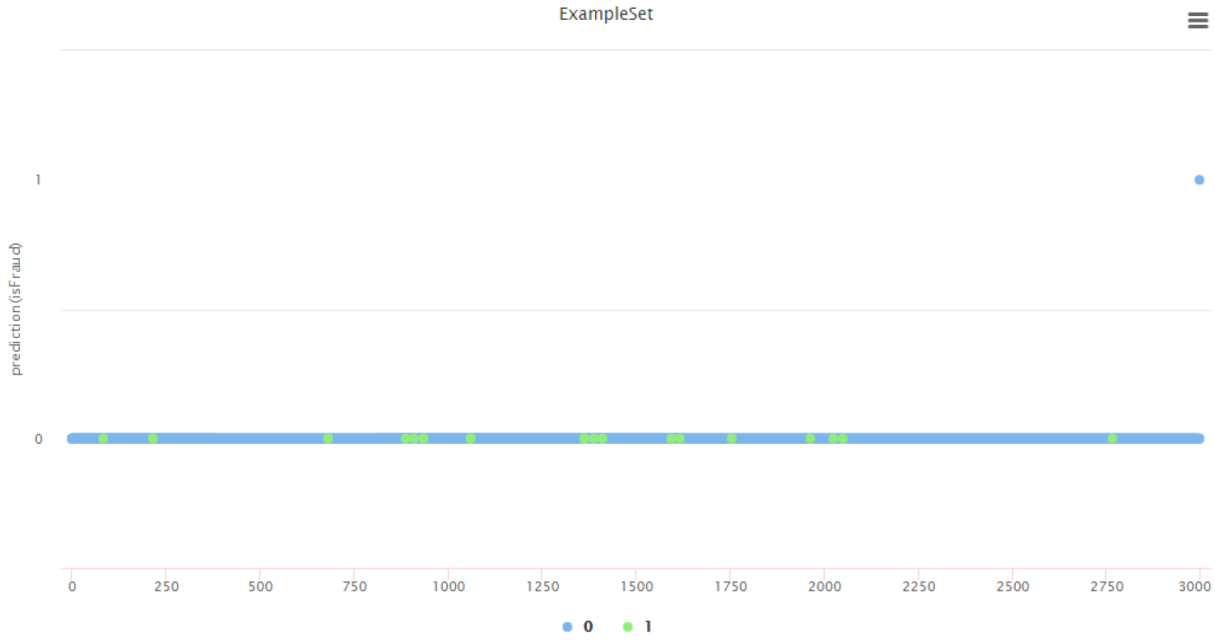

Figure 6: Output of implementing the support vector machine algorithm on the data

Note the image above: The support vector machine algorithm has been applied to our 3000 transactions. In the vertical mode, the prediction of the aforementioned algorithm is checked in fraud detection. See row 1. In this row, our algorithm's prediction is based on the confirmation of fraud detection in transactions. We see the number of a blue circle, transaction (2997). These blue circles in row 1 indicate that the algorithm predicted that fraud had occurred in these transactions, but in fact there was no fraud in this one case. But we note that there is no green circle in this row. This means that there is no information in any transaction that the algorithm's prediction dictates the existence of fraud and in fact there was fraud.

See row 0. In this row, the algorithm's prediction is that there is no fraud. The blue circles in this row also say that in fact there is no fraud. The algorithm's work is correct, and the green circles say that there is fraud, which is contrary to the truth.

At the end of this algorithm, we also implement the performance table to be informed in detail about the statistics and figures of the accuracy of the algorithm's work.

Table 3: Result of implementing the support vector machine algorithm on the data

Class precision	True1	True0	
99.33%	20	2978	Pred0
0.00%	0	1	Pred1
	0.00%	98.97%	Class recall

According to the performance table above, the decision tree algorithm correctly identified 2978 transactions that were not fraudulent, incorrectly included 1 transaction as a transaction that was not fraudulent, 20 transactions were incorrectly identified as fraudulent, and finally, there were no transactions that were fraudulent and the algorithm correctly identified fraud.

The accuracy of the decision tree algorithm on our data was calculated to be 99.30%.

Table 1 shows the performance of these 3 algorithms on the test data using the sensitivity or recall, precision, F-criterion, and accuracy indices. Sensitivity indicates the ability to detect an illegal transaction, provided that it is actually illegal. In other words, sensitivity is the ratio of correctly detected illegal transactions (the number of positives is correct) to the total number of illegal transactions (the sum of true positives and false negatives). Accuracy is the ratio of correctly detected illegal transactions (the number of true positives) to the total number of transactions that are detected illegal (the sum of true positives and false positives). The F-criterion is the harmonic mean of sensitivity and precision, and accuracy is the ratio of predictions that are correctly detected. As can be seen in this table, the performance of the deep neural network algorithm is better than the nearest neighbor and support vector machine algorithms in all the mentioned criteria.

Table 4. Comparison of the performance of different algorithms

Algorithm	Accuracy	Precision	F-Score	Sensitivity
Deep Neural Network Algorithm	0.959	0.979	0.958	0.939
Nearest Neighbor Algorithm	0.932	0.950	0.931	0.912
Support Vector Machine Algorithm	0.939	0.945	0.938	0.932

Finally, Figure 7 attempts to evaluate the performance of algorithms on the entire available unbalanced dataset using the Precision-Recall Curve (PRC). This curve shows the ratio between precision and recall at different thresholds and is useful in cases where the population ratio of classes is very unbalanced. According to this figure, the deep neural network algorithm, with its area under the graph, has a better performance compared to the other two algorithms.

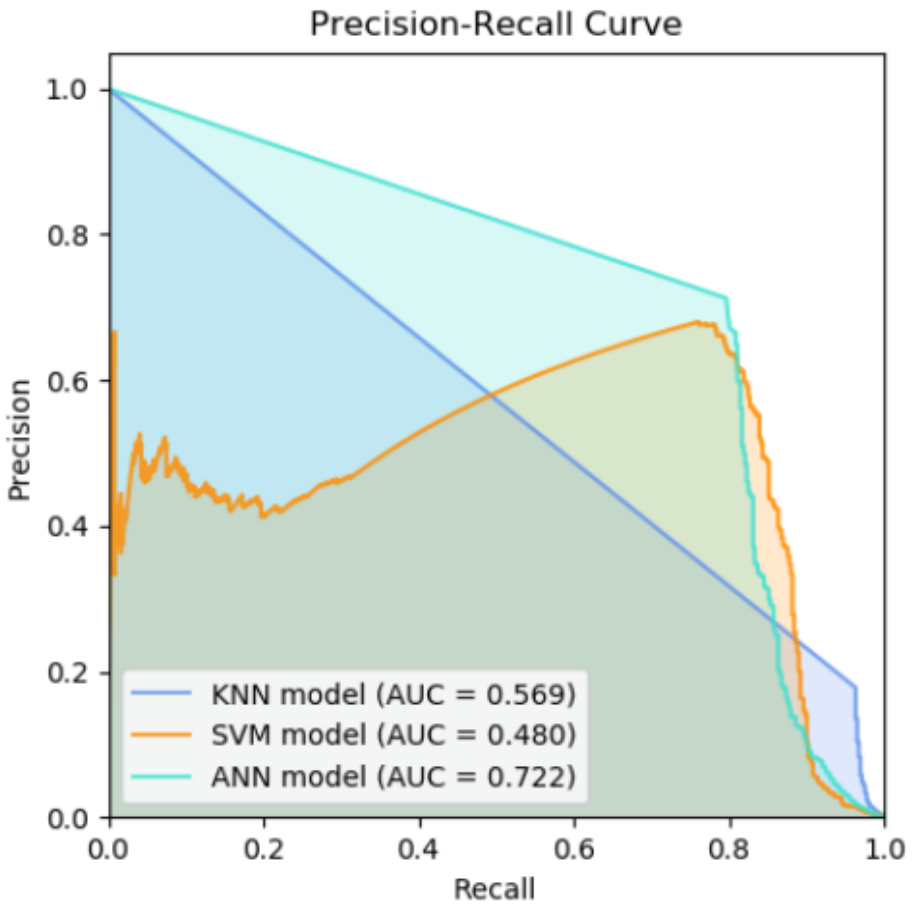

Figure 7: Precision-Recall Curve for Different Algorithms on the Entire Dataset

Discussion and Conclusion

The findings of the present study demonstrated that the proposed deep learning model achieved superior performance in detecting money laundering transactions compared with the K-nearest neighbor (KNN) and support vector machine (SVM) algorithms. Specifically, the deep neural network obtained the highest levels of accuracy, precision, F-score, and sensitivity among the evaluated classification techniques. Furthermore, the precision–recall analysis revealed that the deep learning model generated the largest area under the curve, indicating a stronger ability to discriminate between legitimate and suspicious transactions in an imbalanced financial dataset. These findings suggest that deep learning architectures are capable of extracting complex nonlinear patterns from transaction data and effectively identifying hidden indicators of illicit financial activities.

The superior performance of the deep neural network can be attributed to the hierarchical representation-learning capabilities inherent in deep learning systems. Unlike traditional machine learning algorithms that rely heavily on manually engineered features, deep neural networks automatically learn multiple levels of abstraction from raw data, enabling them to capture intricate behavioral patterns associated with money laundering activities. This interpretation aligns with the theoretical foundations of deep learning proposed by Arnold et al., who argued that multilayer neural architectures can progressively transform raw inputs into increasingly meaningful representations suitable for classification tasks (10). Similarly, LeCun et al. emphasized that deep learning systems excel in identifying highly complex relationships that are difficult to model using shallow machine learning approaches (11).

In the context of anti-money laundering detection, where suspicious transactions often exhibit subtle and dynamic characteristics, such representation-learning capabilities provide a substantial advantage.

Another important finding of the study is the ability of the deep learning model to maintain strong predictive performance despite the severe class imbalance inherent in financial transaction datasets. Money laundering detection systems typically face a situation in which suspicious transactions constitute only a very small proportion of all financial activities. This imbalance can significantly hinder the effectiveness of conventional classification algorithms because they tend to favor the majority class and overlook rare but critical suspicious events. The challenge of imbalanced datasets has long been recognized as one of the most significant obstacles in fraud detection research (6, 7). The results of the present study indicate that the combination of data balancing strategies and deep learning techniques can effectively mitigate this challenge and improve the identification of minority-class transactions.

The findings further support the effectiveness of sampling-based approaches in handling imbalanced financial datasets. Prior to model training, under-sampling techniques were employed to create a more balanced dataset. This strategy likely contributed to the enhanced sensitivity of the proposed model by ensuring that suspicious transactions received adequate representation during the learning process. These results are consistent with the work of Drummond and Holte, who demonstrated that under-sampling methods often outperform alternative balancing techniques in highly skewed classification environments (8). Likewise, Dal Pozzolo et al. showed that carefully calibrated sampling procedures can substantially improve classifier performance in fraud detection contexts characterized by severe class imbalance (9). Therefore, the success of the proposed model appears to result not only from the deep learning architecture itself but also from the effective preprocessing procedures applied to the data.

The performance comparison between the deep neural network and the KNN algorithm provides additional insight into the strengths of deep learning for anti-money laundering applications. Although KNN is recognized as a simple and effective non-parametric classifier in many pattern recognition tasks, its performance tends to deteriorate when confronted with high-dimensional datasets and complex decision boundaries. Financial transaction data often involve intricate relationships among multiple variables, making local similarity-based classification insufficient for capturing sophisticated laundering patterns. The lower performance of KNN observed in the present study therefore appears consistent with theoretical expectations and previous findings indicating that traditional distance-based algorithms struggle to accommodate the complexity of modern financial crime detection problems (4).

Similarly, the support vector machine algorithm demonstrated weaker performance than the deep neural network. Although SVM has historically been regarded as a powerful classification technique and has been successfully applied in numerous fraud detection studies, its effectiveness may be limited when dealing with large-scale financial datasets characterized by complex nonlinear interactions and evolving transaction patterns. Deep learning models possess greater flexibility for modeling such interactions through multiple hidden layers and adaptive feature extraction mechanisms. Consequently, the superior results achieved by the deep neural network support the growing consensus that deep learning techniques are increasingly outperforming traditional machine learning approaches in complex financial analytics tasks (11, 20).

The findings of this study are strongly aligned with recent empirical research on anti-money laundering systems. Alotibi et al. reported that deep learning methods achieved remarkable classification performance in cryptocurrency-

related money laundering detection and frequently outperformed conventional machine learning techniques (19). Their conclusions emphasized the ability of deep neural architectures to identify hidden patterns associated with illicit financial behavior, a finding that closely corresponds with the results obtained in the present investigation. Likewise, Samaddar demonstrated that deep learning-based AML models yielded superior performance compared with several widely used machine learning classifiers, reinforcing the notion that multilayer neural architectures offer considerable advantages in suspicious transaction detection (20). The convergence of these findings across different financial environments suggests that deep learning possesses robust generalizability for AML applications.

The present results also support the broader literature on technology-driven anti-money laundering systems. Demetis argued that advanced analytical technologies can significantly improve the effectiveness of AML operations by enabling the automated detection of suspicious financial behaviors that might remain unnoticed within traditional compliance frameworks (3). Similarly, the report by Grint et al. emphasized that emerging technologies are becoming essential components of modern AML compliance systems due to the increasing scale and complexity of financial crime (2). The superior performance of the proposed deep learning model provides empirical evidence supporting these arguments and illustrates how artificial intelligence can strengthen institutional capabilities in combating financial crime.

An additional implication of the findings concerns the reduction of false negatives. In anti-money laundering contexts, false negatives are particularly problematic because they represent illicit transactions that successfully evade detection. Such failures may facilitate continued criminal activity and expose financial institutions to regulatory penalties, reputational damage, and substantial financial losses. The deep learning model demonstrated stronger sensitivity and recall performance than the competing algorithms, indicating a greater capacity to identify suspicious transactions. This outcome is especially important because AML systems are generally evaluated not merely on overall accuracy but on their ability to detect rare criminal activities with minimal oversight. The enhanced recall achieved by the deep neural network therefore represents a meaningful practical contribution to financial crime prevention efforts.

The findings also resonate with recent developments in criminology-informed machine learning research. Ramadhan proposed that integrating criminological perspectives into machine learning frameworks can significantly improve money laundering detection performance across multiple evaluation metrics (1). Although the present study focused primarily on data-driven classification rather than theory-based behavioral modeling, the strong performance of the deep learning approach suggests that future systems may benefit even further from combining sophisticated computational architectures with criminological insights. Such integration could improve both predictive accuracy and model interpretability, thereby addressing one of the most frequently cited challenges associated with artificial intelligence applications in financial regulation.

The implementation of modern machine learning infrastructures likely contributed to the success of the proposed approach. Advances in machine learning software ecosystems, including Scikit-learn and TensorFlow, have substantially improved researchers' ability to develop, optimize, and deploy sophisticated classification systems (16, 17). Likewise, improvements in neural network design, optimization procedures, and activation functions have enhanced the efficiency and stability of deep learning models (12, 13). The use of dropout regularization and stochastic gradient descent optimization further strengthened the generalization capability of the proposed model, reducing the likelihood of overfitting and improving predictive performance on unseen data (14, 15).

The results additionally align with evidence from other fraud detection domains beyond anti-money laundering. Zheng et al. demonstrated that deep learning-based systems could successfully detect telecommunications fraud and substantially reduce financial losses in real-world banking environments (18). The consistency between their findings and the present study suggests that deep learning possesses broad applicability across diverse forms of financial fraud. Whether applied to telecommunications fraud, payment fraud, credit card fraud, or money laundering, deep neural networks appear capable of identifying hidden behavioral structures that distinguish legitimate activities from illicit ones.

Overall, the findings provide compelling evidence that deep learning represents a highly effective approach for anti-money laundering detection. The superior performance observed across multiple evaluation metrics indicates that deep neural networks are better equipped than conventional machine learning algorithms to address the complexities of modern financial crime detection. As financial systems continue to generate larger volumes of transactional data and criminal methodologies become increasingly sophisticated, the adoption of advanced deep learning technologies may become an indispensable component of future AML infrastructures.

Despite the promising findings, several limitations should be acknowledged. First, the study relied on a single benchmark dataset, which may limit the generalizability of the results to other banking environments, financial institutions, or geographical contexts. Second, the dataset was highly anonymized, restricting the ability to interpret specific transaction characteristics and understand the behavioral mechanisms underlying model decisions. Third, the evaluation focused on a limited number of classification algorithms and did not include more recent approaches such as gradient boosting, ensemble learning, graph neural networks, or transformer-based architectures. Finally, the study was conducted in an experimental setting rather than a real-time operational environment, meaning that practical deployment challenges such as scalability, latency, concept drift, and regulatory compliance were not fully examined.

Future studies should evaluate deep learning-based money laundering detection systems using multiple real-world datasets obtained from diverse financial institutions and jurisdictions. Researchers may also investigate the integration of behavioral, criminological, and network-analysis features into deep learning architectures to enhance predictive performance and interpretability. Comparative analyses involving state-of-the-art artificial intelligence techniques, including graph-based learning, attention mechanisms, explainable AI models, and hybrid ensemble frameworks, would provide valuable insights into the next generation of AML systems. Longitudinal studies examining model adaptation to evolving laundering strategies and changing transaction patterns would further contribute to understanding the sustainability and robustness of AI-driven AML solutions.

Financial institutions should consider incorporating deep learning technologies into their existing anti-money laundering frameworks to improve suspicious transaction detection capabilities. Effective implementation should be accompanied by robust data governance policies, continuous model monitoring, and periodic retraining procedures to ensure sustained performance in dynamic financial environments. Regulatory agencies and compliance departments may benefit from investing in advanced analytical infrastructures and workforce training programs that facilitate the integration of artificial intelligence into AML operations. Organizations should also establish multidisciplinary teams consisting of data scientists, compliance experts, risk managers, and financial crime investigators to maximize the effectiveness and accountability of AI-supported money laundering detection systems.

Acknowledgments

We would like to express our appreciation and gratitude to all those who helped us carrying out this study.

Authors' Contributions

All authors equally contributed to this study.

Declaration of Interest

The authors of this article declared no conflict of interest.

Ethical Considerations

All ethical principles were adheried in conducting and writing this article.

Transparency of Data

In accordance with the principles of transparency and open research, we declare that all data and materials used in this study are available upon request.

Funding

This research was carried out independently with personal funding and without the financial support of any governmental or private institution or organization.

References

1. Ramadhan S. Harnessing Machine Learning for Money Laundering Detection: A Criminological Theory-Centric Approach. Journal of Money Laundering Control. 2025;28(1):184-201. doi: 10.1108/JMLC-04-2024-0083.
2. Grint R, O'Driscoll C, Patton S. New Technologies and Anti-Money Laundering Compliance Report. Financial Conduct Authority, 2017.
3. Demetis DS. Fighting Money Laundering with Technology: A Case Study of Bank X in the UK. Decision Support Systems. 2018;105:96-107.
4. Nawaz S, Ghous H. Detection of Money Laundering Using Data Mining Models: A Review. International Research Journal of Modernization in Engineering Technology and Science. 2021;3(1).
5. Kou Y, Lu CT, Sirwongwattana S, Huang YP, editors. Survey of Fraud Detection Techniques. Proceedings of the IEEE International Conference on Networking, Sensing and Control; 2004.
6. Japkowicz N, Stephen S. The Class Imbalance Problem: A Systematic Study. Intelligent Data Analysis. 2002;6(5):429-49.
7. He H, Garcia EA. Learning from Imbalanced Data. IEEE Transactions on Knowledge and Data Engineering. 2009;21(9):1263-84.
8. Drummond C, Holte RC, editors. C4.5, Class Imbalance, and Cost Sensitivity: Why Under-Sampling Beats Over-Sampling. Proceedings of the Workshop on Learning from Imbalanced Datasets II; 2003.
9. Dal Pozzolo A, Caelen O, Johnson RA, Bontempi G, editors. Calibrating Probability with Undersampling for Unbalanced Classification. Proceedings of the IEEE Symposium on Computational Intelligence and Data Mining; 2015.
10. Arnold L, Rebecchi S, Chevallier S, Paugam-Moisy H, editors. An Introduction to Deep Learning. Proceedings of the European Symposium on Artificial Neural Networks; 2011.

11. LeCun Y, Bengio Y, Hinton G. Deep Learning. Nature. 2015;521:436-44.

12. Hagan MT, Demuth HB, Beale MH, De Jesus O. Neural Network Design. 2nd ed. USA: Martin Hagan; 2014.

13. Nair V, Hinton GE, editors. Rectified Linear Units Improve Restricted Boltzmann Machines. Proceedings of the 27th International Conference on International Conference on Machine Learning; 2010.

14. Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. Journal of Machine Learning Research. 2014;15:1929-58.

15. Bottou L. Stochastic Gradient Descent Tricks. In: Montavon G, Orr GB, Muller KR, editors. Neural Networks: Tricks of the Trade. Berlin: Springer; 2012. p. 421-36.

16. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-Learn: Machine Learning in Python. Journal of Machine Learning Research. 2011;12:2825-30.

17. Abadi M, Barham P, Chen J, Chen Z, Davis A, Dean J, et al., editors. TensorFlow: A System for Large-Scale Machine Learning. Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation; 2016.

18. Zheng YJ, Zhou XH, Sheng WG, Xue Y, Chen SY. Generative Adversarial Network Based Telecom Fraud Detection at the Receiving Bank. Neural Networks. 2018;102:78-86.

19. Alotibi J, Almutanni B, Alsubait T, Alhakami H, Baz A. Money Laundering Detection Using Machine Learning and Deep Learning. International Journal of Advanced Computer Science and Applications. 2022.

20. Samaddar S. Deep Learning Model for Anti-Money Laundering Detection Techniques. 2024.