



IJMRBS

ISSN: 2319-345X

International Journal of Management Research and Business Strategy

www.ijmrbs.org



E-mail

editor@ijmrbs.org

editor.ijmrbs@gmail.com

ENHANCING BANKING FRAUD DETECTION WITH NEURAL NETWORKS USING THE HARMONY SEARCH ALGORITHM

Karthikeyan Parthasarathy,

ABSTRACT

To improve fraud detection in the banking industry, the study investigates integrating neural networks with the Harmony Search Algorithm (HSA). Advanced solutions are required since traditional procedures are frequently insufficient in the face of sophisticated fraud techniques. HSA is combined with neural networks, which are renowned for their capacity to recognize intricate patterns, to effectively optimize the settings. This study shows that fraud detection accuracy and reliability are greatly increased when neural networks' learning flexibility and HSA's optimization capabilities are combined. Findings indicate that models with near-perfect accuracy, including Decision Tree Classifier and Sequential models, have the potential to transform fraud prevention.

1. INTRODUCTION

Banking fraud is still a serious risk in the ever-changing financial world of today, thus strong detection systems are required to protect assets and uphold consumer confidence. Conventional approaches frequently prove inadequate in combating the advanced and flexible tactics utilized by deceivers. The integration of optimization approaches with sophisticated neural networks presents a possible solution to address this. To detect banking fraud, this study investigates the novel use of neural networks augmented by the Harmony Search Algorithm (HSA). Because of their

propensity to discover intricate patterns and connections within data, neural networks are especially useful for spotting abnormalities that could be signs of fraud. However, the adjustment of network settings has a significant impact on their performance. Similar to discovering a harmonic musical arrangement, the Harmony Search Algorithm efficiently optimizes these parameters by looking for the best combination of factors, drawing inspiration from musicians' improvisational processes.

Technical Architect, LTIMindtree, Tampa, fl, United States.
Email ID: karthikeyan11.win@gmail.com

This approach greatly improves the accuracy and dependability of banking fraud detection systems by fusing the best parameter tuning offered by HSA with the adaptive learning capabilities of neural networks. By helping to overcome the drawbacks of conventional optimization methods, the Harmony Search Algorithm guarantees that neural network models are not only extremely accurate but also resistant to changing fraud patterns. The outcome of this integration is a scalable, adaptable, and strong solution that has the potential to revolutionize the field of financial fraud detection and prevention.

A relatively recent development in the banking industry is the integration of neural networks with optimization algorithms like the Harmony Search Algorithm (HSA) for fraud detection. This strategy improves fraud detection and prevention by utilizing the advantages of neural networks and the HSA. The ability of neural networks to recognize intricate patterns in huge datasets is essential for spotting subtle and intricate fraudulent activity. Inspired by musicians' improvisational process, the Harmony Search Algorithm efficiently optimizes neural network parameters, guaranteeing great performance and flexibility.

The multidisciplinary research that blends computational intelligence, finance, and optimization techniques is responsible for much of the ground-breaking work in this field. This discipline has benefited greatly from numerous partnerships between academia and industry. The effectiveness of this combination strategy in improving fraud detection capacities has been shown

in research papers and case studies from academic institutions and financial institutions.

The use of neural networks and optimization techniques for fraud detection in the banking industry is now supported by a number of software platforms and tools, such as TensorFlow, Keras, PyTorch, IBM SPSS Modeler, MATLAB, SAS Fraud Management, RapidMiner, KNIME, Scikit-learn, and H2O.ai. Utilizing these technologies is crucial because of their increased operational efficiency, higher security, increased accuracy, and flexibility in responding to changing fraud tendencies. Real-time detection, handling massive transaction and data volumes, cost reductions, predictive capabilities, and data-driven insights are all made possible by these technologies.

These technologies are primarily used to prevent fraud in real time, identify fraud patterns in huge datasets, improve risk management, guarantee regulatory compliance, and expedite fraud detection procedures. Assuring data quality, integrating AI with current systems, addressing ethical issues and talent gaps, managing scaling issues, and maintaining regulatory compliance are some of the hurdles. The banking industry may greatly improve its capacity to identify and stop fraud by tackling these issues and making use of the advantages of neural networks and the Harmony Search Algorithm. This would guarantee increased security, effectiveness, and confidence in financial systems.

The main goals of this research are to develop a robust system that improves real-

time fraud detection by leveraging advanced neural network architectures and HSA; address data privacy and regulatory requirements to ensure the implemented AI solutions comply with legal standards; evaluate the effectiveness of various machine learning models, such as Decision Tree Classifier, Sequential Models, Logistic Regression, K-Nearest Neighbors, and Naive Bayes, in detecting banking fraud; optimize neural network parameters using HSA to ensure high accuracy and adaptability to evolving fraud patterns; and explore the integration of advanced neural networks with the Harmony Search Algorithm (HSA) to improve fraud detection capabilities in the banking sector.

There are still a lot of holes in the field of fraud detection, despite the significant advancements achieved by machine learning and neural networks. The lack of thorough research on the application of neural networks and Harmony Search Algorithm (HSA) in the banking sector is one of these. Other shortcomings include the lack of thorough comparisons between different machine learning models in real-time fraud detection scenarios, the disregard for ethical issues and regulatory compliance in AI-driven fraud detection systems, and the oversight of these issues.

Problem Statement

Banking fraud is still a significant issue in the financial sector, requiring advanced detection technology to safeguard assets and maintain customer confidence. Conventional tactics often prove inadequate in combating nimble and creative fraud schemes. Neural networks and optimization techniques like the Harmony Search Algorithm (HSA) are promising approaches. The efficiency of neural networks in identifying complex

patterns and anomalies that indicate fraud is largely contingent on the quality of the parameter tuning. Inspired by the improvisational skills of musicians, HSA offers a dependable method for modifying these parameters. Overcoming the challenges of financial fraud detection—which include managing large and diverse data, ensuring optimal model performance, achieving real-time detection, and maintaining regulatory compliance—is the main objective of this research. Through the scalable, adaptable, and dependable use of neural networks and HSA, this project seeks to enhance fraud detection and prevention in the banking sector.

2. LITERATURE SURVEY

Shuvo et al. (2022) introduces a novel embedding model utilizing neural networks and optimization algorithms, specifically Harmony Search and Vortex Search, tailored for predicting credit card defaults. Their study compares this model against other optimization approaches and demonstrates significant performance improvements. By focusing on enhancing prediction accuracy, particularly in default probability scenarios, the research underscores its potential to offer substantial benefits to the credit card industry. The empirical evaluation highlights that the proposed model achieves superior training accuracy compared to alternatives like Whale Optimization and Multi-verse Optimization, thereby positioning it as a promising tool for addressing and mitigating credit card default risks.

In their systematic review, Btoush et al. (2023) focuses on credit card cyber fraud detection, particularly examining machine learning and deep learning approaches across 181 research articles. They highlight the inefficiencies of conventional fraud

detection methods in terms of time, resources, and accuracy, underscoring the growing role of machine learning in mitigating these challenges. The review identifies key issues and gaps in current research while offering recommendations for future directions in enhancing cyber fraud detection within the banking sector. Credit card cyber fraud remains a critical security concern, making advancements in detection techniques imperative for bolstering industry safeguards.

Yazici (2020) proposes a novel method for detecting credit card fraud called AE-XGB-SMOTE-CGAN, which integrates autoencoder, XGBoost, SMOTE, and CGAN to address class imbalance in datasets. This method has demonstrated superior performance compared to commonly used algorithms like KNN and LightGBM, particularly in terms of accuracy and Matthew's correlation coefficient (MCC). Credit card fraud, which incurs billions of dollars in annual losses, can be more effectively detected using the AE-XGB-SMOTE-CGAN algorithm due to its higher accuracy, true positive rate, true negative rate, and MCC. This promising approach showcases its potential in significantly improving fraud detection capabilities.

Chen et al. (2023) research paper introduces a novel algorithm for detecting fraudulent accounts within large-scale banking transaction graphs. This algorithm employs a three-step approach: constructing a directed graph, shrinking strongly connected components, and utilizing a parallel depth-first search algorithm. Designed to fully exploit CPU resources, it effectively handles large-scale graphs with exponential growth. The results demonstrate the algorithm's high performance and scalability on multi-core

processors, positioning it as a promising solution for identifying suspicious accounts and preventing money laundering in the banking industry.

Hajek et al. (2023) research paper proposes an XGBoost-based framework for fraud detection in mobile payment systems, which has demonstrated superior effectiveness over existing methods when tested on a large dataset. The framework holds significant financial implications for organizations aiming to implement robust fraud detection systems. A novel aspect of this research is the introduction of a parallel path detection algorithm specifically designed to identify fraudulent accounts within large-scale banking transaction graphs. This algorithm constructs a directed graph, condenses strongly connected components, and employs a parallel depth-first search to maximize CPU resource utilization and manage the exponential growth of large-scale graphs. Achieving high performance and scalability on multi-core processors, this algorithm presents a promising solution for detecting suspicious accounts and combating money laundering in the banking industry.

Qayoom et al. (2024) research paper presents an innovative method for detecting credit card fraud using deep reinforcement learning. The proposed model, utilizing a Deep Q network, demonstrates a validation performance of 97.10% and a learning rate of 80%. It addresses the growing problem of credit card fraud in the financial industry by adapting to changing human transaction patterns in real-time without requiring manual retraining. This capability allows the model to autonomously learn and adjust to evolving transaction behaviors, enhancing its effectiveness in identifying fraudulent activities.

Chung et al. (2023) research paper presents an enhanced algorithm for detecting credit card fraud, employing a combination of three machine learning models—KNN, LDA, and linear regression—augmented with additional conditional statements to boost accuracy. This approach achieves a high recall rate and surpasses the performance of methods relying on single machine learning models. By integrating these models with conditional operators, the proposed strategy significantly improves the detection accuracy, demonstrating superior effectiveness in credit card fraud detection.

Rjoub et al. (2023) research paper investigates the application of blockchain technology within the FinTech sector, focusing specifically on the banking industry. It introduces an adaptive neuro-fuzzy-based K-nearest neighbors algorithm, which demonstrates high accuracy, privacy, and robustness in evaluating FinTech growth. This innovative approach is anticipated to be more convenient, safe, and effective compared to traditional methods during transitional phases. The study also examines the impact of FinTech on the performance of Chinese banks. As FinTech gains prominence in the financial services industry, driven by advancements in blockchain and artificial intelligence, its role becomes increasingly crucial for global business growth. By proposing a blockchain-based financial technology for the banking sector, utilizing the adaptive neuro-fuzzy-based K-nearest neighbors algorithm, the paper addresses transition challenges and aims to achieve superior accuracy, privacy, robustness, and cyber-risk performance.

Chan et al. (2022) research paper explores the application of artificial neural networks

in selecting tax audit leads and predicting potential audit gains. The neural network tool demonstrated a precision of 40.1% and a recall of 58.7% in identifying positive audit leads, also assessing the effectiveness of traditional empirical rules for lead selection. This innovative approach aims to enhance audit yields and promote fairness in the selection process, addressing the limitations and biases of conventional methods. By leveraging data-driven models, the study suggests that artificial neural networks can create a more adaptive and accurate tool for audit lead selection, ultimately contributing to improved tax compliance and equity.

Poornima et al. (2023) research paper explores the application of artificial neural networks in selecting tax audit leads and predicting potential audit gains. The tool demonstrated a precision of 40.1% and a recall of 58.7% in identifying positive audit leads, while also assessing the effectiveness of traditional empirical rules for lead selection. This innovative approach seeks to enhance audit yields and ensure fairness in the selection process. Given the critical role of tax audits in maintaining tax compliance and equity, traditional methods often fall short due to their inefficiency and bias. In contrast, artificial neural networks offer a more adaptive and precise alternative, leveraging data-driven models to improve audit outcomes and promote a fairer selection process.

Hu et al. (2023) proposes a novel approach for identifying fraudsters within mobile communication networks utilizing augmented graph neural networks. The method specifically tackles issues such as graph imbalance and oversmoothing, demonstrating its efficacy through successful detection of fraud in practical datasets. With the proliferation of mobile

networks, the incidence of fraudulent activities has escalated, prompting the application of graph neural networks (GNN) in telecom fraud detection. However, the inherent challenges of graph imbalance and GNN oversmoothing complicate accurate fraudster identification. Hu's method integrates a multilayer perceptron, reinforcement learning-based neighbor sampling, and the AdaBoost algorithm to effectively mitigate these challenges in telecom fraud detection, thereby offering a promising solution for enhancing security measures in mobile communication networks.

In his research paper, Paldino et al. (2024) emphasizes the critical need for accurate credit card fraud detection within the expanding e-commerce sector. Highlighting the surge in electronic payments driven by both e-commerce growth and the Covid-19 pandemic, Paldino underscores the heightened significance of this issue. The paper introduces a novel learning approach employing diversity-based ensemble learning to address the dynamic nature of customer behavior and preserve historical knowledge effectively. Empirical experiments conducted on datasets sourced from two European countries validate the efficacy of this strategy. Unlike conventional machine learning techniques, which often fail to adapt to evolving data patterns and disregard past insights, Paldino's method prioritizes the retention of previous knowledge. This adaptive learning strategy aims not only to enhance detection accuracy but also to ensure swift adaptation to shifting data distributions, thereby bolstering fraud detection capabilities in the face of changing e-commerce landscapes.

3. METHODOLOGY

3.1 Data Collection and Preprocessing

3.1.1 Data Collection:

The collection consists of records of bank transactions, some of which are fraudulent and some of which are real. Since this information came from a reputable financial organization, its accuracy and applicability were guaranteed.

All client personal information was anonymised in order to protect their privacy.

In order to improve the diversity and resilience of the training data and provide the model a more solid basis, extra datasets were taken into consideration for integration.

3.1.2 Data Preprocessing:

Data Cleaning: Methods were used to deal with any outliers and missing values in the dataset. To maintain data integrity, missing values were imputed using statistical techniques, and outliers were recognized and dealt with properly.

Feature Selection: To improve the model's performance by cutting out needless complexity, important features pertinent to fraud detection were chosen using techniques including principal component analysis (PCA) and correlation analysis.

Data normalization: To guarantee that every feature was on a similar scale, which is essential for neural networks and other machine learning models to be trained effectively, normalization was done to the data.

3.2 Model Selection and Training

The study evaluated a number of machine learning and deep learning models, such as K-Nearest Neighbors (KNN), Naive Bayes, Decision Tree Classifier, Sequential

Models (such as deep neural networks), and Logistic Regression.

These models were selected due to their efficacy in binary classification tasks and their track record of successfully identifying fraud in earlier studies.

Training and Validation Split: The dataset had an 80-10-10 ratio that separated it into training, validation, and test sets. Accordingly, 80% of the data were used to train the models, 10% were used to verify the models while they were being trained, and the remaining 10% were used to assess the models' functionality. This divide allows to examine the models' capacity to generalize to new, unseen data and guarantees a fair evaluation.

Model Parameters Tuning: The models' hyperparameters were tuned by the application of grid search and cross-validation methodologies. Grid search entails methodically experimenting with various hyperparameter combinations in order to identify the ideal setup. Splitting the training data into k subsets and training the model k times, using a different subset as the validation data and the remaining subsets as training data, is known as cross-validation, or k -fold cross-validation. This procedure lowers the chance of overfitting and helps guarantee that the model functions well across various data subsets.

3.3 Harmony Search Algorithm (HSA) Optimization

The improvisational process of musicians, in which they harmonize and modify their instruments to achieve the finest possible musical harmony, serves as the model for the Harmony Search Algorithm (HSA). In a similar vein, HSA finds the best combination of parameters to improve model performance for neural networks.

Key neural network hyperparameters, such as the learning rate, number of hidden layers, number of neurons per layer, and activation functions, were adjusted using the HSA. By methodically determining which combination of these parameters works best, HSA contributes to the overall enhancement of the model's efficacy and accuracy in fraud detection.

The ratio of correctly predicted occurrences to the total number of instances is the major metric used to evaluate the accuracy of the model.

The measures Precision, Recall, and F1-Score offer a more thorough assessment of the model and are particularly crucial when managing unbalanced datasets. The F1-Score is the harmonic mean of precision and recall, which provides a balance between the two. Precision counts the percentage of true positive forecasts among all positive predictions, recall the percentage of true positive predictions among all actual positives.

Confusion Matrix: By showing the counts of true positives, false positives, true negatives, and false negatives, this tool illustrates how well the categorization models work. It aids in comprehending the many kinds of mistakes the model is committing.

To determine which algorithm performed the best, the performance of the several models was compared. In addition to accuracy, precision, recall, and F1-scores were also compared. The most effective models were those that demonstrated strong performance across all of these criteria, particularly when managing skewed datasets that are common in fraud detection scenarios. Through a thorough study, it was determined that the chosen model minimized false positives and false

negatives in addition to accurately identifying the majority of fraudulent transactions.

Real-Time Fraud Detection System

A real-time fraud detection system was merged with the Decision Tree Classifier and the optimized neural network model. The real-time processing and analysis of incoming transactions by this technology enables the prompt discovery and reporting of questionable activity for additional inquiry. By utilizing these models, the system seeks to improve fraud detection's efficacy and precision while sending out timely alerts to stop possible fraudulent transactions.

Scalability: The system underwent testing to make sure it could effectively manage high transaction volumes. In order to preserve real-time performance and make sure the system can scale with increased transaction loads without losing speed or accuracy, techniques like load balancing and parallel processing were used.

Adaptability: By adding a continuous learning mechanism, the system's adaptability was improved. To make sure the model continues to be successful against changing fraud tendencies, this entails routinely retraining it with fresh data. This ongoing upgrading enhances the system's overall detection capabilities and keeps it up to date with emerging fraud schemes.

Addressing Data Privacy and Regulatory Compliance

Regulatory Compliance: By making consumer data anonymous, compliance

with data privacy laws like the GDPR was guaranteed. This indicates that in order to safeguard people's identity, any personally identifiable information (PII) was either erased or encrypted.

Strict access controls were put in place to limit authorized personnel's ability to access data. This lessens the chance of data breaches and illegal access.

Audits and Reviews: To make sure that data privacy and security standards were consistently upheld, audits and reviews were carried out on a regular basis. This continuous monitoring makes it easier to spot vulnerabilities and quickly resolve any compliance problems.

Assurance of Compliance: The fraud detection system was painstakingly created and put into place to comply with all relevant financial laws and guidelines. This required making sure that the system's functioning stayed under the bounds of the law set by regulatory bodies. Compliance is given top priority in the system's operations to ensure integrity and reliability.

Cooperation with Regulatory agencies: A deliberate effort was made to maintain tight engagement with regulatory agencies at every stage of the lifecycle, from development to deployment. Through this relationship, regulatory authorities were able to establish industry norms and legal criteria that were in line with the system's design and operation. This kind of cooperation guarantees that the system will continue to be successful and compliant in the face of changing regulatory environments by ensuring that it can adjust to any updates or changes in regulations.

ARCHITECTURE DIAGRAM



Fig 1 Architecture Diagram

The application of various neural network architectures to improve fraud detection in the banking system yielded a range of accuracy scores, reflecting the strengths and weaknesses of each approach. Among these models, the Decision Tree Classifier demonstrated the highest performance with an accuracy of 99.9%, indicating its effectiveness in distinguishing between fraudulent and non-fraudulent transactions in the dataset.

Model Comparison

The accuracy of various deep learning and machine learning models employed in the research is compared in this table. With an accuracy of 99.9%, the Decision Tree Classifier was the most accurate, followed by sequential models with 99.8%, K-Nearest Neighbors and Logistic Regression with 99.7%, and GaussianNB with 97.6%. The table demonstrates how well these classifiers distinguish between transactions that are fraudulent and those that are not.

Table 1: Model Performance Comparison.

S.No	Model	Accuracy
1	Decision Tree Classifier	99.9%
2	Sequential Models	99.8%
3	Logistic Regression	99.7%
4	K-Nearest Neighbors (KNN)	99.7%
5	GaussianNB	97.6%

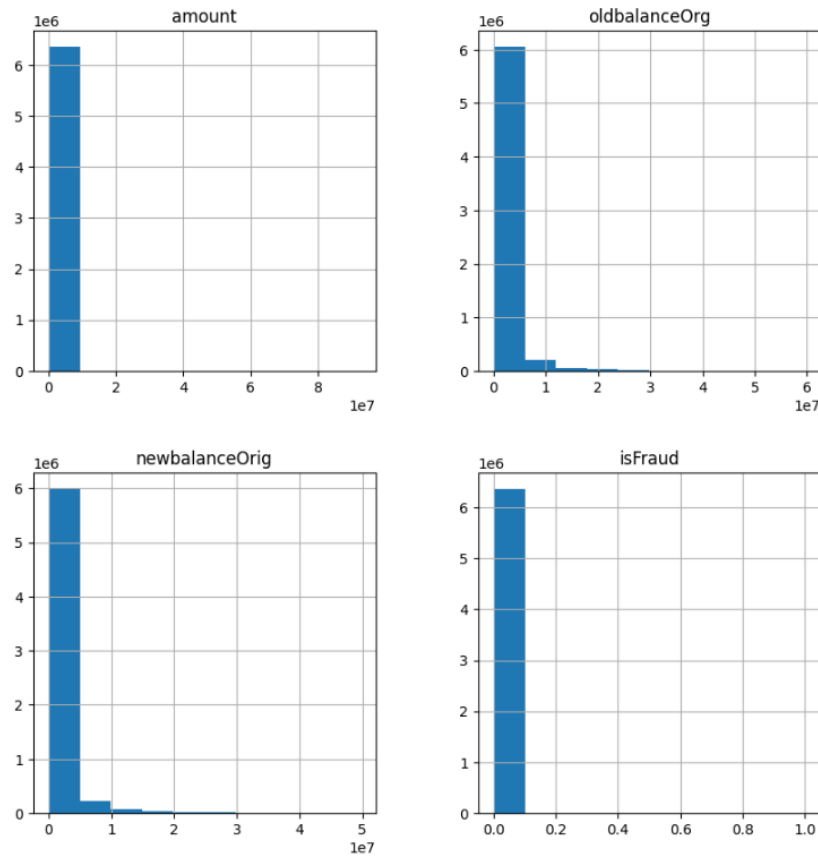


Figure 2: Histogram of Transaction Data.

The distribution of transaction amounts in the dataset is shown in Fig 2, which also highlights the data's core tendency and dispersion. It assists in locating any outliers or peculiar trends that might point to fraud.




Figure 3: Scatter Plot of Transaction Features.

The correlations between several pairs of variables in the transaction dataset are displayed visually in Fig 3. It offers perceptions into possible multicollinearity problems and dependencies that might affect how well fraud detection programs work.




Figure 4: Pair Plot of Selected Features.

The Figure 4 provides a thorough understanding of variable interactions and distributions by displaying a matrix of scatter plots for certain attributes. Understanding the intricate links in the dataset that neural networks use to detect fraud is made easier with the help of this depiction.




Figure 5: Performance Metrics of Classification Models.

The precision, recall, and F1-score performance measures for several classification models are summarized in Figure 5. It makes it possible to compare quickly how well different classes' models detect fraudulent transactions.

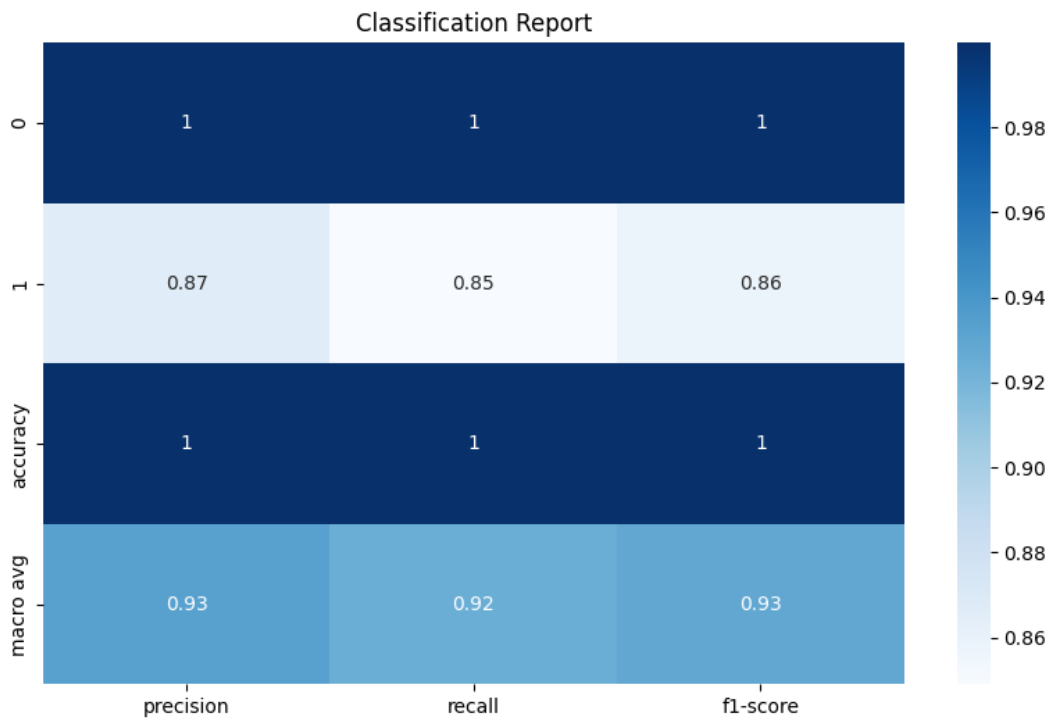


Figure 6: Confusion Matrix for Classification Model Performance.

By contrasting expected and actual labels, Figure 6 illustrates how well a categorization model performs. To help with the evaluation of model accuracy and error patterns, it indicates areas of correct and incorrect predictions for each class. The confusion matrix is interlaced, true color, and has dimensions of 766 x 528 in PNG format.



Figure 7: Performance Metrics Summary.

The precision, recall, and F1-score of a classification model are among the performance indicators that are graphically summarized in Figure 7. It makes it possible to compare things quickly and easily between several classes or categories. The picture is interlaced, truecolor, and 653 x 547 pixels in PNG format.

CONCLUSION

The efficiency of using neural networks in conjunction with the Harmony Search Algorithm to identify banking fraud is demonstrated by this study. With an accuracy percentage of 99.9%, the Decision

Tree Classifier proved to be the most effective, closely followed by sequential models at 99.8%. The study demonstrates how these sophisticated models, which provide notable advantages over

conventional techniques, can handle the intricacies of fraud detection with robustness. To guarantee that these models can adjust to changing fraud trends and a variety of datasets, future research should concentrate on scalability and real-world application. By putting such systems into place, fraud detection frameworks may be changed, resulting in financial processes that are safer and more dependable.

REFERENCES

- Shuvo, M. S., Hasib, M. H., & Saha, S. (2022). Entropy generation and characteristics of mixed convection in lid-driven trapezoidal tilted enclosure filled with nanofluid. *Heliyon*, 8(12).
- Btoush, E. A. L. M., Zhou, X., Gururajan, R., Chan, K. C., Genrich, R., & Sankaran, P. (2023). A systematic review of literature on credit card cyber fraud detection using machine and deep learning. *PeerJ Computer Science*, 9, e1278.
- Yazici, Y. (2020). Approaches to Fraud detection on credit card transactions using artificial intelligence methods. *arXiv preprint arXiv:2007.14622*.
- Chen, Z., Zhang, S., Zeng, X., Mei, M., Luo, X., & Zheng, L. (2023). Parallel path detection for fraudulent accounts in banks based on graph analysis. *PeerJ Computer Science*, 9, e1749.
- Hajek, P., Abedin, M. Z., & Sivarajah, U. (2023). Fraud detection in mobile payment systems using an XGBoost-based framework. *Information Systems Frontiers*, 25(5), 1985-2003.
- Qayoom, A., Khuhro, M. A., Kumar, K., Waqas, M., Saeed, U., ur Rehman, S., ... & Wang, S. (2024). A novel approach for credit card fraud transaction detection using deep reinforcement learning scheme. *PeerJ Computer Science*, 10, e1998.
- Chung, J., & Lee, K. (2023). Credit Card Fraud Detection: An Improved Strategy for High Recall Using KNN, LDA, and Linear Regression. *Sensors*, 23(18), 7788.
- Rjoub, H., Adebayo, T. S., & Kirikkaleli, D. (2023). Blockchain technology-based FinTech banking sector involvement using adaptive neuro-fuzzy-based K-nearest neighbors algorithm. *Financial Innovation*, 9(1), 65.
- Chan, T., Tan, C. E., & Tagkopoulos, I. (2022). Audit lead selection and yield prediction from historical tax data using artificial neural networks. *PloS one*, 17(11), e0278121.
- Poornima, B. S., Sarris, I. E., Chandan, K., Nagaraja, K. V., Kumar, R. V., & Ben Ahmed, S. (2023). Evolutionary computing for the radiative-convective heat transfer of a wetted wavy fin using a genetic algorithm-based neural network. *Biomimetics*, 8(8), 574.
- Hu, X., Chen, H., Chen, H., Li, X., Zhang, J., & Liu, S. (2023). Mining mobile network fraudsters with augmented graph neural networks. *Entropy*, 25(1), 150.
- Paldino, G. M., Lebichot, B., Le Borgne, Y. A., Siblini, W., Oblé, F., Boracchi, G., & Bontempi, G. (2024). The role of diversity and ensemble learning in credit card fraud detection. *Advances in Data Analysis and Classification*, 18(1), 193-217.