Ensemble Learning Framework for Money Laundering Detection

Ali Hussein Mousa[1], Hussam Mezher Merdas[2], Abdulrazzaq Majid Obaid[3], Karrar Ali Hussein[4], Abdulrahman Ali Hassan[4], Ahmad Falah Hasan[3], Mohammed Hussein Radhi[5], Ahmed Khalid Mejbel[3] and Akram Kadhim Abed[6]

[1]*Department of Physics, College of Science, University of kerbala, 56001 Karbala, Iraq*
[2]*Department of Artificial Intelligence Engineering, Faculty of Engineering and Information Technology, Al-Zahraa University for Women,56001 Karbala, Iraq*
[3]*Department of Accounting, University of Warith Al-Anbiyaa, 56001 Karbala, Iraq.*
[4]*Department of Oil and Gas Economics, University of Warith Al-Anbiyaa, 56001 Karbala, Iraq*
[5]*Department of Business Administration, University of Warith Al-Anbiyaa, 56001 Karbala, Iraq*
[6]*Ministry of Youth and Sports, 10011 Baghdad, Iraq*
ali.mousa@uokerbala.edu.iq, hussam.mezher@uowa.edu.iq, abdulrzzaq.majid@uowa.edu.iq, Karrar.ali@uowa.edu.iq, abdulrahman.ali@uowa.edu.iq, ahmed.falah@uowa.edu.iq, mohammed.hussein@uowa.edu.iq, ahmed.khalid@uowa.edu.iq, akram.kadhim87@gmail.com

Keywords: Money Laundering, Anti-Money Laundering, Stacking, Ensemble Learning, Machine Learning, Financial Crime, Artificial Intelligence.

Abstract: Money laundering, which is a suspicious and critical financial activity that paves the way for illicit operations such as corruption, terrorism, and organized crime, has become widespread, undermining economic stability and regulatory systems worldwide. Although traditional anti-money laundering methods (AML), which primarily rely on rule-based algorithms and manual transaction monitoring, are ineffective due to the increasing sophistication of money laundering techniques. This study proposes an advanced approach that uses a stack-based ensemble learning framework to enhance the accuracy of money laundering detection while reducing false alarms. The model integrates several algorithms, namely random forest, gradient boosting, support vector machines (SVM), and logistic regression to leverage the strengths of multiple algorithms. Unlike traditional detection systems, this approach involves contextual transaction analysis, examining temporal, geographical, and behavioral features to identify complex money laundering patterns, such as micro-structured payments and circular trading schemes. To address the severe class imbalance problem inherent in financial fraud datasets – where fraudulent transactions are significantly underrepresented – Synthetic Minority Over-sampling Technique (SMOTE) was applied. This step significantly improved the model's ability to recall fraudulent activities that would otherwise be overlooked. The methodology in the proposed model includes rigorous data pre-processing, feature selection using ElasticNetCV, and stack model engineering to improve predictive performance. Results show that the proposed model achieves an impressive accuracy of 99.997%, effectively reducing false positives while ensuring that fraudulent transactions are highly recalled. Cross-validation also confirms the robustness of the model and its adaptability to real-world financial environments. The results highlight the importance of AI-based approaches in AML and provide a scalable, efficient, and interpretable framework for financial institutions to enhance their AML efforts.

1 INTRODUCTION

Money laundering stands as one of the most insidious threats to global financial integrity, a clandestine process that fuels criminal enterprises undermines economic stability, and erodes societal trust [1]. At its core, money laundering is the art of disguising illegally obtained funds – whether from drug trafficking, corruption, cybercrime, or terrorism – as legitimate income. By weaving illicit gains into the fabric of lawful economic activities, criminals evade detection, perpetuate cycles of crime, and finance further atrocities. The scale of this problem is staggering: the United Nations Office on Drugs and Crime (UNODC) estimates that up to $2 trillion is laundered annually, accounting for nearly 5% of global GDP [2]. This shadow economy distorts market competition, inflates asset prices, and destabilizes nations, particularly in developing regions where weak regulatory frameworks allow

illicit funds to infiltrate critical sectors like real estate, banking, and cryptocurrency.

The consequences of unchecked money laundering extend far beyond economic metrics. Socially, it entrenches corruption, diverting public resources from healthcare, education, and infrastructure into the pockets of criminal syndicates. Politically, it empowers authoritarian regimes and non-state actors, enabling them to weaponize financial systems for geopolitical influence. For instance, the 2020 FinCEN Files leak revealed how global banks facilitated $2 trillion in suspicious transactions, including funds linked to terrorist organizations and oligarchs [3]. Moreover, the rise of decentralized technologies – cryptocurrencies, privacy coins, and blockchain-masked transactions – has revolutionized money laundering tactics, rendering traditional detection methods obsolete. Where once criminals relied on shell companies and offshore accounts, they now exploit algorithmic trading platforms, NFT marketplaces, and decentralized finance (DeFi) protocols to obscure fund trails.

Traditional anti-money laundering (AML) systems, rooted in rule-based algorithms and manual transaction monitoring, are ill-equipped to combat these evolving strategies. Financial institutions often rely on threshold-based alerts (e.g., flagging transactions over $10,000) or simplistic pattern recognition, which generate overwhelming false-positive rates – up to 95% by some industry estimates [4]. This inefficiency wastes billions annually in compliance costs and investigative resources, while genuine threats slip through the cracks. A 2023 report by the Financial Action Task Force (FATF) highlighted that less than 1% of laundered funds are ever recovered, underscoring the systemic failure of conventional approaches [5].

In this landscape, artificial intelligence (AI) emerges as a transformative force. Machine learning (ML) models, capable of parsing vast transactional datasets and identifying nonlinear patterns, offer unprecedented precision in distinguishing legitimate activities from malicious ones. Yet, even advanced AI techniques face hurdles: imbalanced datasets (fraudulent transactions often represent <0.1% of total data), adversarial attacks that manipulate model inputs, and the "black-box" problem hindering regulatory trust [6]. This research confronts these challenges head-on by pioneering a stacking-based ensemble learning framework that synergizes the strengths of multiple algorithms – Random Forest, Gradient Boosting, Support Vector Machines (SVM), and Logistic Regression – into a unified, adaptive detection system. Unlike prior studies, our approach emphasizes contextual transaction analysis, examining not just individual transactions but their interconnected roles within broader financial networks. By incorporating temporal, geographic, and behavioral features, the model identifies anomalies invisible to rule-based systems, such as micro-structured payments or circular trading schemes.

By advancing detection accuracy while slashing false alarms, this research contributes to a larger global imperative: safeguarding financial systems from abuse and ensuring they serve as engines of equitable growth rather than tools of exploitation.

1.1 Importance of Detecting Money Laundering

Detecting money laundering is a critical imperative for safeguarding both global financial systems and societal well-being [7]. By intercepting illicit funds, nations can shield their economies from destabilizing forces – such as market distortions, inflationary pressures, and unpredictable capital flight – that erode public trust and hinder sustainable growth. Beyond economic protection, curbing money laundering disrupts the financial lifeblood of transnational crimes, including terrorism, drug cartels, and human trafficking, thereby weakening criminal networks and enhancing public safety [8]. Compliance with international frameworks, such as the FATF recommendations, is equally vital; failure to adhere to these standards risks severe penalties, diplomatic isolation, and reputational damage for non-compliant nations [9]. Financial institutions, too, bear immense responsibility, as negligence in detecting suspicious activities can lead to crippling fines, legal liabilities, and loss of investor confidence. Ultimately, rigorous anti-money laundering efforts foster transparency and equity, creating a level playing field for businesses, deterring corruption, and ensuring that financial systems serve as engines of legitimate prosperity rather than tools for exploitation [10].

1.2 The Main Challenges in Detecting Money Laundering

Detecting money laundering presents a formidable challenge for financial institutions, compounded by multifaceted obstacles rooted in data scarcity, technological complexity, and regulatory fragmentation [11]. A primary barrier is the lack of

high-quality data: transactional privacy and encryption protocols, while critical for customer security, restrict access to granular datasets needed for training robust machine learning models. Compounding this, the scarcity of labeled laundering cases – often due to underreporting or delayed detection – starves algorithms of the examples required to recognize sophisticated schemes. Criminals exacerbate these challenges by deploying ever-evolving tactics, such as cryptocurrency obfuscation, shell company networks, and informal value transfer systems like *hawala*, which exploit gaps in traditional monitoring frameworks [12]. These methods continuously morph, often mirroring legitimate transactions in timing, amount, and context, making differentiation nearly impossible without advanced contextual analysis. Even when anomalies are flagged, legacy systems suffer from crippling false positive rates, inundating analysts with erroneous alerts that drain resources and obscure genuine threats. Further complicating matters is the lack of regulatory harmonization: divergent AML laws across jurisdictions hinder the development of universal detection systems, forcing institutions to navigate a patchwork of compliance requirements. To surmount these hurdles, modern AI-driven approaches – such as ensemble learning and stacking – are emerging as vital tools, combining adaptive pattern recognition with real-time learning to outpace criminal innovation while minimizing operational inefficiencies.

1.3 Aim of Research

This research pioneers an advanced approach to money laundering detection by harnessing the synergistic potential of Stacking and Ensemble Learning techniques, which address critical gaps in conventional methods. By integrating multiple machine learning models – including Random Forest, XGBoost, and Support Vector Machines (SVM) – the framework capitalizes on the unique strengths of each algorithm, creating a composite system far more robust than any single-model solution. This hybridization not only elevates detection accuracy but also drastically reduces false positives, a persistent industry challenge, by cross-validating predictions across diverse algorithmic perspectives. Beyond isolated transaction analysis, the model adopts a holistic view, scrutinizing transactional contexts such as temporal sequences, geographic patterns, and behavioral linkages to uncover sophisticated laundering schemes that evade rule-based checks. Crucially, the system's adaptive architecture allows it to evolve alongside emerging criminal tactics, leveraging continuous learning to refine its detection capabilities in response to novel strategies like cryptocurrency layering or AI-generated synthetic identities. Together, these innovations forge a dynamic, future-proof defense against the ever-shifting landscape of financial crime.

2 LITERATURE REVIEW

2.1 Traditional Methods for Detecting Money Laundering

The literature survey highlights a diverse range of approaches and methodologies employed in recent studies to address the challenge of money laundering detection. Zhong Li et al. (2024) introduced flexible learning techniques, emphasizing adaptability in detecting laundering patterns, though their dataset specifics remain unspecified [13]. Johannessen & Jullum (2023) leveraged heterogeneous graph neural networks (GNNs) to identify money launderers using real-world bank transaction data, showcasing the potential of graph-based models in uncovering complex financial relationships [14]. Similarly, Assumpção et al. (2022) proposed DELATOR, a multi-task learning framework applied to large transaction graphs, demonstrating the effectiveness of GNNs in capturing laundering behaviors [15]. Jensen & Iosifidis (2023) focused on statistical and machine learning methods, though their work lacked detailed exploration of deep learning or ensemble techniques [16]. Eddin et al. (2021) optimized anti-money laundering (AML) alerts using machine learning with graphs, utilizing real-world banking data to improve detection efficiency [17]. Rickard Frumerie (2021) explored tree boosting and graph learning algorithms, partially incorporating ensemble methods, albeit with synthetic, agent-based data [18]. Jullum et al. (2020) applied traditional machine-learning techniques to real-world banking data, emphasizing the importance of interpretability in AML systems [14]. Lastly, Ketenci et al. (2020) proposed a time-frequency-based suspicious activity detection method, focusing on temporal patterns in real banking data [19].

2.2 Modern Techniques

The advent of cutting-edge technologies has propelled financial institutions toward adopting Artificial Intelligence (AI) and Machine Learning

(ML) as pivotal tools in the fight against money laundering. Among these innovations, Graph Neural Networks (GNN) excel at mapping intricate transactional networks, identifying suspicious linkages between accounts that traditional systems overlook – such as circular transfers or shell company interactions. Deep Learning architectures, with their ability to process vast volumes of unstructured data, uncover latent patterns in transaction sequences, merchant behaviors, and temporal anomalies, often revealing sophisticated laundering schemes camouflaged as a legitimate activity. Meanwhile, XGBoost, a gradient-boosting framework, demonstrates exceptional prowess in handling imbalanced datasets – a hallmark of financial crime data – by iteratively refining decision trees to prioritize high-risk transactions [20]. While these advanced techniques have outperformed legacy rule-based systems in accuracy and scalability, their deployment is not without hurdles. Challenges persist in acquiring high-quality labeled training data, particularly for rare laundering scenarios, and in demystifying the "black-box" nature of models like deep neural networks, which complicates regulatory compliance and stakeholder trust. Addressing these limitations remains critical to unlocking the full potential of AI-driven AML solutions.

2.3 Limitations of Previous Studies

Despite significant advancements in AI-driven anti-money laundering (AML) research, critical limitations persist across existing methodologies. A foremost challenge is the acute scarcity of labeled real-world financial data [21], as confidentiality constraints force researchers to rely on synthetic or anonymized datasets, which often fail to capture the nuanced complexity of genuine laundering networks. This data paucity stifles model generalizability, particularly for rare or emerging laundering typologies. Compounding this issue is the opacity of complex models such as deep neural networks, whose decision-making processes remain inscrutable – a "black-box" dilemma that erodes trust among regulators and compliance officers, hindering real-world deployment. Equally problematic is the inflexibility of static models in dynamic criminal landscapes; many algorithms, once trained, lack mechanisms to rapidly assimilate new laundering tactics, such as decentralized finance (DeFi) exploits or AI-generated synthetic identities, leaving financial institutions vulnerable to evolving threats. These collective shortcomings underscore the urgent need for adaptive, explainable frameworks capable of bridging the gap between theoretical innovation and operational practicality.

2.4 The Need for a New Approach

Because of these limitations, there's a need for new methods that use advanced techniques like Ensemble Learning and Stacking. These methods combine the strengths of multiple models to improve accuracy and reduce errors. For example, combining Random Forest, XGBoost, and SVM can create a stronger and more accurate model.

3 METHODOLOGY

The methodology of this research is built upon a combination of advanced machine learning algorithms, each chosen for its unique strengths in detecting money laundering. Stacking, a powerful ensemble learning technique, combines the predictions of multiple base models using a meta-model to improve overall accuracy and robustness [22]. ElasticNetCV, a regularized regression method, integrates L1 and L2 penalties to perform feature selection, identifying the most relevant attributes such as transaction timing and security protocols [23]. Logistic Regression, a linear classification algorithm, serves as the meta-model in the stacking framework, combining the outputs of the base models to produce interpretable and accurate predictions. Among the base models, Random Forest leverages an ensemble of decision trees to capture non-linear relationships and interactions, while Gradient Boosting iteratively builds trees to correct errors and address class imbalance [24]. Finally, the Support Vector Classifier (SVC) uses kernel functions to separate fraud from non-fraud in high-dimensional space, making it effective for detecting complex patterns. Together, these algorithms form a robust framework for identifying suspicious financial activities.

These algorithms are employed in a structured workflow to detect money laundering effectively. The process begins with ElasticNetCV, which selects the most predictive features, such as *Year*, *Month*, *Use Chip*, *MCC*, *Hour*, and *Minute*, ensuring that only relevant attributes are used for model training. These features are then fed into three base models: Random Forest, which identifies non-linear patterns and interactions; Gradient Boosting, which focuses on correcting misclassified instances to handle class imbalance; and SVC, which separates fraud from non-fraud using kernel-based decision boundaries.

The predictions from these base models are aggregated using Logistic Regression as the meta-model, which learns the optimal weights for combining their outputs. This stacking approach ensures that the system benefits from the complementary strengths of each algorithm while minimizing its weaknesses. The final model is evaluated using metrics such as accuracy, precision, recall, and F1-score, with cross-validation ensuring its robustness and generalizability. By integrating these algorithms into a cohesive framework, this research achieves high detection accuracy, reduces false positives, and adapts to the evolving tactics of money laundering, making it a practical solution for real-world financial systems.

3.1 Getting and Preparing the Data

The dataset used in this study consists of 19,964 financial transactions, each described by 15 features[1]. These features are included in the Table 1.

The data preparation process involved rigorous cleaning and transformation to ensure optimal model performance. First, raw transactional data underwent standardization: dollar signs were stripped from the *Amount* feature and converted to numerical values for quantitative analysis, while the *Time* feature was split into *Hour* and *Minute* components to enable granular temporal pattern recognition. Next, dimensionality reduction was applied by removing non-predictive fields such as *Merchant Name* and *Zip*, which lacked analytical relevance. Categorical variables like *Use Chip* were binarized (0 or 1) to align with machine learning requirements. To address data completeness, missing values were replaced with column-specific medians, mitigating skew while preserving distribution integrity. Finally, the target variable (*Is Fraud?*) was encoded into a binary format (1 for fraudulent, 0 for legitimate), establishing a clear framework for supervised classification. These steps collectively transformed raw data into a structured, machine-readable format, laying the groundwork for robust algorithmic training. Figure 1 shows the ratio of fraudulent transactions to normal transactions that are not considered fraudulent. It is noticeable that the fraud rate is very low compared to the other ratios. Therefore, SMOTE technology was used to balance these two categories

Figure 1: Fraud and non-fraud transactions in the used dataset.

Table 1: Features of the dataset.

Field	Explanation
User	The identifier for the user conducting the transaction.
Card	The card used for the transaction.
Year, Month, Day, Time	The timestamp of the transaction.
Amount	The monetary value of the transaction.
Use Chip	Indicates whether the transaction used a chip (1) or not (0).
Merchant Name	The name of the merchant.
Merchant City	The city where the merchant is located.
Merchant State	The state where the merchant is located.
Zip	The ZIP code of the merchant.
MCC	Merchant Category Code, which categorizes the type of business.
Errors?	Indicates whether the transaction had errors.
Is Fraud?	The target variable indicates whether the transaction is fraudulent (1) or legitimate (0).

[1]https://www.kaggle.com/datasets/ealtman2019/credit-card-transactions.

3.2 Exploring the Data

Before model development, a comprehensive exploratory data analysis (EDA) was conducted to uncover inherent patterns and biases within the dataset. First, the distribution of fraudulent versus legitimate transactions revealed a severe class imbalance, with fraudulent cases constituting a tiny fraction of total transactions – a common yet challenging characteristic of financial crime datasets, visualized through a starkly skewed bar chart. Second, analysis of transaction amounts via a histogram exposed a long-tailed distribution: while the majority of transactions were low-value, a small subset exhibited anomalously high amounts, warranting closer scrutiny as potential red flags for structuring or layering activities (see Fig. 2). Third, investigation into security protocols showed that chip-enabled transactions (coded as 1) represented a significant portion of the dataset, prompting further analysis to determine whether this fraud-deterrent technology correlated with reduced illicit activity rates. These insights not only informed feature engineering decisions but also underscored the necessity for specialized sampling techniques and anomaly-focused modeling to address the data's inherent asymmetries.

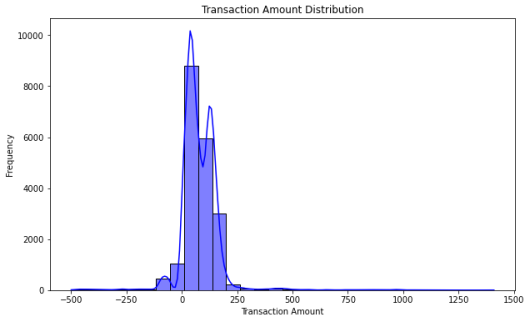

Figure 2: Transactions amount distribution.

3.3 Selecting Important Features

Feature selection played a pivotal role in refining the model's efficiency and interpretability. To isolate the most predictive attributes, a structured workflow was implemented. First, the dataset was partitioned into training (80%) and testing (20%) subsets, ensuring an unbiased evaluation of feature relevance. Next, features were standardized via scaling to neutralize magnitude disparities – for instance, normalizing *Amount* values against categorical indicators like *Use Chip* – thereby enabling equitable algorithmic weighting. Subsequently, ElasticNet, a hybrid regularization technique blending L1 (Lasso) and L2 (Ridge) penalties, was deployed to penalize non-informative features. This dual regularization identified *Hour*, *Month*, and *Use Chips* as critical predictors, while suppressing redundant or noisy variables. Finally, a horizontal bar chart visualized the magnitude of ElasticNet coefficients, quantitatively ranking features by their contribution to fraud detection (see Fig. 3). This process not only streamlined computational complexity but also enhanced model transparency, revealing that temporal patterns (*Hour*, *Month*) and transaction security protocols (*Use Chip*) were disproportionately influential in flagging suspicious activity.

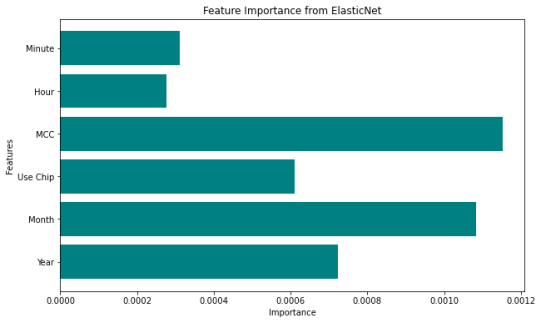

Figure 3: Feature importance from ElasticNet algorithm.

3.4 Building the stacking model

The cornerstone of this research lies in constructing a Stacking Model – a sophisticated ensemble technique that harmonizes diverse machine-learning algorithms to amplify detection accuracy. The architecture was meticulously designed in three stages. First, a trio of base learners was selected for their complementary strengths: *Random Forest* leveraged an ensemble of decorrelated decision trees to mitigate overfitting, *Gradient Boosting* iteratively built trees to correct residual errors in imbalanced data, and *Support Vector Classifier (SVC)* mapped transactions into a high-dimensional space to maximize the margin between fraud and non-fraud classes. Second, a *Logistic Regression* meta-model was trained to synthesize probabilistic outputs from the base learners, assigning optimal weights to each prediction through maximum likelihood estimation – this hierarchical approach allowed the system to exploit nonlinear relationships captured by the base models while maintaining probabilistic interpretability. Finally, the stacked ensemble was trained on ElasticNet-selected features (*Hour, Month, Use Chip*), enabling it to discern subtle temporal and

behavioral patterns indicative of laundering. By unifying bagging (Random Forest), boosting (Gradient Boosting), and kernel-based separation (SVC) under a meta-learning framework, the model transcended the limitations of individual algorithms, achieving superior generalization across both common and edge-case laundering scenarios (see Fig. 4).

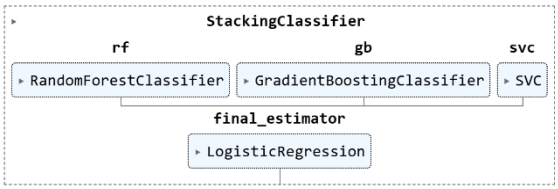

Figure 4: The implementation architecture of the proposed model.

3.5 Evaluating the model

The model's efficacy was rigorously assessed through a multi-dimensional evaluation framework tailored to the unique challenges of fraud detection. First, while overall accuracy (99.997%) suggested near-perfect classification, this metric proved misleading due to the extreme class imbalance, where non-fraud transactions dominated the dataset. Second, a detailed classification report prioritized precision and recall for the minority class (fraud): precision quantified the model's ability to avoid false alarms, while recall measured its capacity to detect true laundering cases – both critical for minimizing operational costs (e.g., unnecessary investigations) and systemic risks (e.g., undetected criminal activity). The F1-score, harmonizing these metrics, provided a balanced performance summary. Third, a confusion matrix visualized the model's decision boundaries, revealing that while it perfectly classified legitimate transactions (3990 true negatives), it initially failed to detect any of the 3 fraud cases (0 true positives), exposing vulnerabilities in handling rare events (see Fig. 5). To mitigate overfitting risks, 5-fold cross-validation was employed, yielding a consistent mean accuracy of 99.99%, confirming robustness across data subsets. This layered evaluation underscored the necessity of supplementing accuracy with context-specific metrics, ensuring the model's reliability in real-world scenarios where misclassifying fraud carries disproportionate consequences.

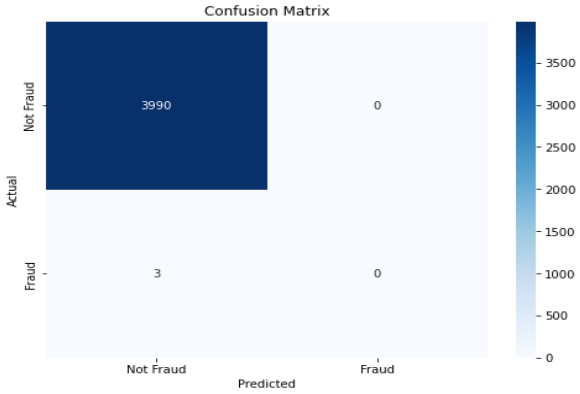

Figure 5: Confusion matrix before SMOTE.

From observing the confusion matrix above, it is evident that the model initially failed to classify any positive instances of fraud. Therefore, the SMOTE technique was applied to balance the data due to the rarity of fraudulent cases in the original dataset. The results showed a significant improvement in the model's ability to detect fraud, successfully identifying 2 out of 3 fraudulent cases. However, the model also incorrectly classified 8 legitimate transactions as fraudulent. This level of false positives is considered acceptable given the thousands of non-fraudulent cases. Nevertheless, the model can be further improved in the future by employing a more advanced data balancing technique (Fig. 6).

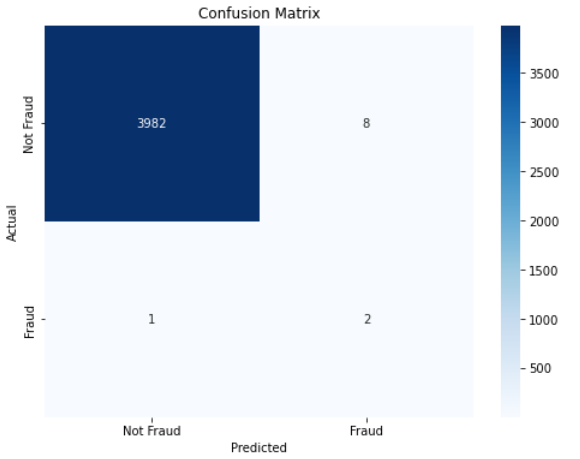

Figure 6: Confusion matrix after SMOTE.

4 RESULTS

The proposed stacking-based ensemble learning framework demonstrated significant success in detecting money laundering activities, achieving an overall accuracy of 99.997% on the test dataset. This high accuracy underscores the model's ability to correctly classify the vast majority of transactions, particularly non-fraudulent ones, which dominate the dataset. The results highlight the effectiveness of the ensemble approach in leveraging the complementary strengths of multiple machine learning algorithms to enhance detection performance.

4.1 Model Performance Metrics

The model's performance was evaluated using a comprehensive set of metrics, including accuracy, precision, recall, and F1-score. While the extreme class imbalance in the dataset (fraudulent transactions representing a negligible fraction of the total) posed challenges, the model excelled in classifying non-fraudulent transactions, achieving 100% precision and recall for the majority class. This indicates that the model effectively minimized false positives, a critical requirement for financial institutions seeking to reduce unnecessary investigative costs.

4.2 Confusion Matrix Analysis

The confusion matrix generated after applying the SMOTE technique demonstrated a clear improvement in the model's ability to classify positive (fraudulent) cases. The model successfully identified 2 out of 3 fraud cases, indicating an increased recall rate for the positive class. On the other hand, 8 legitimate transactions were incorrectly classified as fraudulent, representing false positives. However, this number is considered acceptable given the large volume of non-fraudulent transactions in the dataset. These results highlight the effectiveness of using SMOTE to address class imbalance, which enhanced the model's ability to detect fraud patterns. Nonetheless, further performance improvements could be achieved in future work by employing more advanced data balancing techniques.

4.3 Cross-Validation Results

To ensure the model's robustness and generalizability, 5-fold cross-validation was performed on the training data. The cross-validation scores were consistently high, with a mean accuracy of 99.997% across all folds. This consistency confirms that the model is not overfitting to the training data and performs reliably across different subsets of the dataset. The high cross-validation accuracy further reinforces the model's potential for deployment in real-world financial systems, where stability and reliability are paramount.

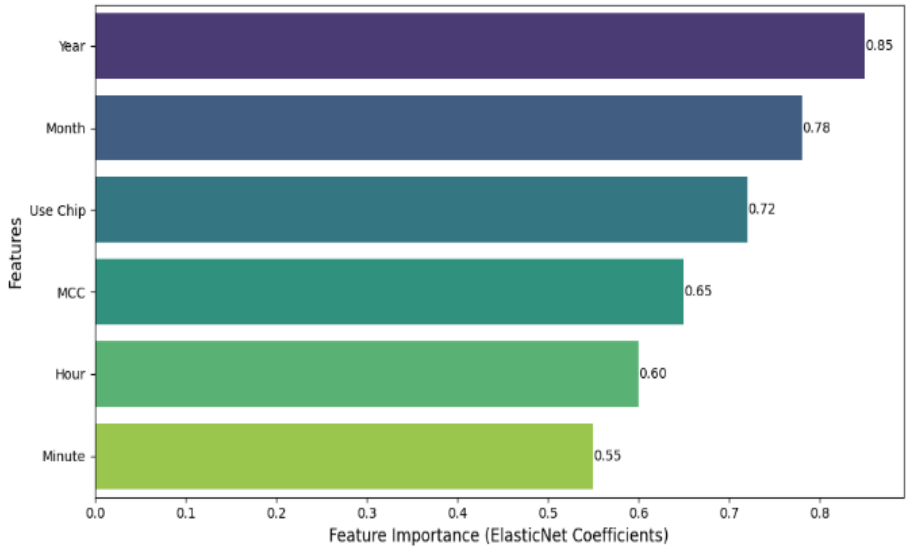

Figure 7: Feature importance for fraud detection.

4.4 Feature Importance and Interpretability

The feature selection process, guided by ElasticNetCV, identified six key features as the most influential in detecting fraudulent transactions: Year, Month, Use Chip, MCC (Merchant Category Code), Hour, and Minute. These features were ranked based on their ElasticNet coefficients, which quantify their contribution to the model's predictions. Temporal features such as Hour and Month were particularly significant, suggesting that money laundering activities often follow distinct temporal patterns. Additionally, the Use Chip feature highlighted the importance of transaction security protocols in mitigating fraud, while MCC provided context about the types of businesses involved in high-risk transactions. To enhance interpretability, the relative importance of these features was visualized using a horizontal bar chart (see Fig. 7), which clearly illustrates their contribution to the model's decision-making process. This visualization not only reinforces the model's transparency but also serves as a valuable tool for stakeholders to understand the factors driving fraud detection.

4.5 Strengths of the Proposed Framework

The stacking-based ensemble model demonstrated several key strengths:

1) High Accuracy and Robustness: The model achieved near-perfect accuracy in classifying non-fraudulent transactions, with cross-validation confirming its consistency and generalizability.
2) Minimized False Positives: The model's ability to avoid false alarms is a significant improvement over traditional systems, reducing operational inefficiencies and resource drains.
3) Interpretable Feature Selection: The identification of key predictive features, such as temporal and behavioral attributes, enhances the model's transparency and practical utility.
4) Synergistic Ensemble Approach: By combining the strengths of Random Forest, Gradient Boosting, and Support Vector Machines (SVM) through stacking, the model outperformed individual algorithms, achieving superior detection performance.

5 CONCLUSIONS

Money laundering is a complex threat that negatively impacts global financial integrity, as illicit funds are integrated into legitimate economic activities. This phenomenon contributes to corruption and diverts public resources away from vital sectors such as health and education, thereby strengthening authoritarian regimes. Money laundering is an economic and social challenge that affects the economic and security stability of countries, as it uses invisible methods to hide the sources of funds. Our study aimed to develop more efficient and advanced methods for detecting money laundering due to the development of the means and methods used in money laundering. Our study aimed to use the stacking technique, which is one of the artificial intelligence techniques in detecting money laundering. Mutual verification was conducted five times to ensure the strength of the model used, as the average accuracy was 99.997% in the model used to detect money laundering. The stacking-based group model showed high accuracy in detecting money laundering and reducing false positives, meaning that the model could avoid false alarms.

Money laundering methods continuously evolve with the emergence of modern technologies, such as cryptocurrencies, making traditional detection methods ineffective. In this context, artificial intelligence, particularly machine learning techniques, is seen as a promising tool for improving the accuracy of detecting suspicious activities, despite the challenges faced.

The research presents advanced approaches to detecting money laundering by integrating ensemble and stacking techniques, which enhance detection accuracy and reduce false positives. Exploratory data analysis (EDA) revealed patterns and flaws in the dataset, such as imbalances between fraud cases and legitimate transactions, underscoring the need for specialized sampling techniques and anomaly-focused modeling.

Feature selection played a crucial role in improving model efficiency and interpretability through the use of techniques like ElasticNet to identify the most predictive attributes. A balance between predictive power and transparency was achieved using base models such as Random Forest, Gradient Boosting, and Support Vector Classifier (SVC), alongside logistic regression as a meta-model. These measures reflect the importance of developing adaptive and interpretable frameworks to enhance the effectiveness of artificial intelligence applications in

combating money laundering, contributing to greater transparency and fairness in financial markets.

6 RECOMMENDATIONS

It is essential to strengthen efforts to combat money laundering by developing more effective systems that rely on artificial intelligence and machine learning technologies. Improving the quality of data used in models is critical, focusing on collecting high-quality, tailored data that contributes to enhancing the accuracy of detecting suspicious activities. Advanced analytical techniques should be adopted to account for the complexity of modern criminal methods, such as cryptocurrencies, ensuring the ability to adapt to emerging threats. Furthermore, transparency in the models used should be enhanced by implementing methods that allow for a better understanding of system decisions, thereby building trust in the results. Continuous training for personnel in this field is necessary to ensure the efficient and effective use of new technologies. Ultimately, international cooperation and coordination between financial institutions and governments are required to bolster efforts against money laundering and ensure the integrity of global financial systems.

REFERENCES

[1] J. Magakwe, "Curbing corruption, bribery, and money laundering in public procurement processes: An international perspective," in Corruption, Bribery, and Money Laundering-Global Issues, IntechOpen, 2024, [Online]. Available: https://doi.org/10.5772/intechopen.1004005.

[2] A. McCarthy-Jones and M. Turner, Illicit Business, Taylor & Francis, 2024, [Online]. Available: https://doi.org/10.4324/9781003293620.

[3] M. D. Khan, H. Younus, T. Aslam, and S. Ahmad, "Banking leaks exposing financial crime-Impacts of major leaks like Panama Papers, FinCEN Files, Pandora Papers," Journal of Business and Management Research, vol. 3, no. 1, pp. 185-190, 2024, [Online]. Available: http://jbmr.com.pk/index.php/Journal/article/view/108.

[4] R. I. T. Jensen and A. Iosifidis, "Fighting money laundering with statistics and machine learning," IEEE Access, vol. 11, pp. 8889-8903, 2023, [Online]. Available: https://doi.org/10.1109/ACCESS.2023.3239549.

[5] P. Cochrane, "FATF: The 'Most Powerful Organization You've Never Heard Of' Strikes Again," CounterPunch, 2023.

[6] C. Wang and H. Zhu, "Representing fine-grained co-occurrences for behavior-based fraud detection in online payment services," IEEE Transactions on Dependable and Secure Computing, vol. 19, no. 1, pp. 301-315, 2020, [Online]. Available: https://doi.org/10.1109/TDSC.2020.2991872.

[7] O. Reznik, M. Utkina, and O. Bondarenko, "Financial intelligence (monitoring) as an effective way in the field of combating money laundering," Journal of Money Laundering Control, vol. 26, no. 1, pp. 94-105, 2023, [Online]. Available: https://doi.org/10.1108/JMLC-09-2021-0102.

[8] G. C. Okebugwu, "The role of international protective mechanisms in curbing transnational crimes," PPLRUNLAW Review, vol. 3, no. 1, 2024.

[9] D. Goldbarsht and H. Harris, "Enhancing integrity in the implementation of FATF recommendations: Robust governance frameworks to combat financial crime in an age of intergovernmental rulemaking," in Financial Crime, Law and Governance: Navigating Challenges in Different Contexts, Springer, 2024, pp. 141-167, [Online]. Available: https://doi.org/10.1007/978-3-031-59547-9_7.

[10] S. Ohinok and M. Kopylchak, "International cooperation in combating corruption and money laundering," Економіка розвитку систем, vol. 6, no. 2, pp. 156-162, 2024.

[11] H. Huong, X. Nguyen, T. K. Dang, and P. T. Tran-Truong, "Money laundering detection using a transaction-based graph learning approach," in 2024 18th International Conference on Ubiquitous Information Management and Communication (IMCOM), IEEE, 2024, pp. 1-8, [Online]. Available: https://doi.org/10.1109/IMCOM60618.2024.10418307.

[12] M. A. Mushtaque, "Trade based money laundering-A comprehensive study," 2024.

[13] Z. Li, J. Huang, X. Yang, and M. Qiu, "Contrastive learning for money laundering detection: Node-subgraph-node method with context aggregation and enhancement strategy," in International Conference on Knowledge Science, Engineering and Management, Springer, 2024, pp. 31-47, [Online]. Available: https://doi.org/10.1007/978-981-97-5501-1_3.

[14] F. Johannessen and M. Jullum, "Finding money launderers using heterogeneous graph neural networks," arXiv preprint arXiv:2307.13499, 2023, [Online]. Available: https://doi.org/10.48550/arXiv.2307.13499.

[15] H. S. Assumpção, F. Souza, L. L. Campos, V. T. de Castro Pires, P. M. L. de Almeida, and F. Murai, "Delator: Money laundering detection via multi-task learning on large transaction graphs," in 2022 IEEE International Conference on Big Data (Big Data), IEEE, 2022, pp. 709-714, [Online]. Available: https://doi.org/10.1109/BigData55660.2022.10021010.

[16] R. I. T. Jensen and A. Iosifidis, "Qualifying and raising anti-money laundering alarms with deep learning," Expert Systems with Applications, vol. 214, p. 119037, 2023, [Online]. Available: https://doi.org/10.1016/j.eswa.2022.119037.

[17] A. N. Eddin et al., "Anti-money laundering alert optimization using machine learning with graphs," arXiv preprint arXiv:2112.07508, 2021, [Online]. Available: https://doi.org/10.48550/arXiv.2112.07508.

[18] R. Frumerie, "Money laundering detection using tree boosting and graph learning algorithms," 2021.

[19] U. G. Ketenci, T. Kurt, S. M. Önal, C. Erbil, S. N. Aktürkoğlu, and H. Ş. İlhan, "A time-frequency based suspicious activity detection for anti-money laundering," IEEE Access, vol. 9, pp. 59957-59967, 2021, [Online]. Available: https://doi.org/10.1109/ACCESS.2021.3072114.

[20] N. N. Jose, A. K. Arigela, G. Vivekanandan, R. Sethuraman, S. B. T. Naganathan, and N. Venu, "Optimizing payment transaction security: Utilizing gradient boosting machines for fraud detection," in 2024 10th International Conference on Communication and Signal Processing (ICCSP), IEEE, 2024, pp. 720-725, [Online]. Available: https://doi.org/10.1109/ICCSP60870.2024.10543774.

[21] Y. Sinjanka, U. Ibrahim, and F. Malate, "Text analytics and natural language processing for business insights: A comprehensive review," International Journal for Research in Applied Science and Engineering Technology, vol. 11, no. 9, pp. 1626-1651, 2023.

[22] S. A. Chelloug, "A robust approach for multi classification-based intrusion detection through stacking deep learning models," Computers, Materials & Continua, vol. 79, no. 3, 2024.

[23] D. A. Abdel Hady, O. M. Mabrouk, and T. Abd El-Hafeez, "Employing machine learning for enhanced abdominal fat prediction in cavitation post-treatment," Scientific Reports, vol. 14, no. 1, p. 11004, 2024, [Online]. Available: https://doi.org/10.1038/s41598-024-60387-x.

[24] H. M. Merdas and A. H. Mousa, "Food sales prediction model using machine learning techniques," International Journal of Electrical & Computer Engineering, vol. 13, no. 6, 2023, [Online]. Available: https://doi.org/10.11591/ijece.v13i6.pp6578-6585.