

ENHANCING FRAUD DETECTION AND ANOMALY DETECTION IN RETAIL BANKING USING GENERATIVE AI AND MACHINE LEARNING MODELS

Tanvirahmedshuvo

Master's in Business Administration, Business Analytics, International American University, Los Angeles, USA

Asif Iqbal

Master's in Business Administration Management Information System, International American University, Los Angeles, California

Emon Ahmed

Masters in Science Engineering Management, Westcliff University, California, USA

Ashequr Rahman

Doctoral in Business Administration, Westcliff University, California, USA

Md Risalat Hossain Ontor

Master's in Business Administration, Management Information System, International American University, Los Angeles, California

Abstract

This study investigates the effectiveness of generative models and traditional classification models in detecting fraud and anomalies within the retail banking sector. Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) were evaluated for their capability to generate realistic synthetic transaction data and identify anomalies, achieving anomaly detection accuracies of 91.2% and 93.5%, respectively. These models were also assessed using Inception Score and Fréchet Inception Distance (FID), with GANs exhibiting superior data realism. Among classification models, Gradient Boosting Machines (GBM) demonstrated the best performance, achieving an accuracy of 96.3%, a precision of 93.5%, a recall of 91.4%, and an AUC-ROC of 97.2%. Random Forest and Logistic Regression also performed well, though with slightly lower metrics.

Keywords Generative AI, Fraud Detection, Anomaly Detection, Retail Banking Security, Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), Gradient Boosting Machines (GBM), Machine Learning, Financial Security, Synthetic Data Generation.

INTRODUCTION

The financial sector has witnessed a surge in technological innovations over the past decade, with retail banking emerging as a critical domain for leveraging advancements in artificial intelligence (AI). As the industry becomes increasingly digitized, the threat of fraud, cyberattacks, and operational inefficiencies grows exponentially. Retail banks handle vast amounts of sensitive data, including customer transactions, personal information, and financial records, making them prime targets for malicious activities. Consequently, the need for robust, adaptive, and scalable security solutions has never been greater. Generative AI, a subset of machine learning, has garnered significant attention for its ability to tackle complex problems, including fraud detection, anomaly detection, and synthetic data generation for enhanced security.

Generative AI models, such as Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs), have demonstrated immense potential in creating realistic synthetic data, detecting anomalies, and identifying fraudulent activities. Unlike traditional AI models that rely solely on predictive accuracy, generative models introduce the capability to simulate fraudulent scenarios, thereby training systems to recognize novel patterns of attack. This dual functionality makes them highly suitable for dynamic, high-stakes environments like retail banking. At the same time, traditional classification models, such as Logistic Regression, Random Forest, and Gradient Boosting Machines (GBM), continue to play a pivotal role in fraud detection due to their interpretability and robust performance across diverse datasets.

The application of AI in retail banking security has been extensively studied, highlighting the evolving landscape of technological solutions to combat fraud and anomalies. According to Leevy et al.

(2021), the rise of digital banking has necessitated the development of advanced fraud detection systems capable of handling large-scale, imbalanced datasets. The study emphasized the superiority of machine learning models over rule-based systems, particularly in identifying subtle, non-linear relationships within transactional data. Similarly, Fernández et al. (2020) explored the application of Random Forest and GBM in fraud detection, noting their ability to achieve high precision and recall while maintaining computational efficiency.

Generative AI has also gained prominence in recent years, with its application extending beyond anomaly detection to synthetic data generation. Goodfellow et al. (2014), who pioneered GANs, demonstrated their ability to produce realistic data distributions, paving the way for applications in fraud simulation and training data augmentation. More recently, Kingma and Welling (2013) introduced VAEs as a probabilistic generative model that excels in capturing latent representations of data. In retail banking, these models have shown promise in detecting irregularities and simulating attack scenarios that traditional systems may overlook.

The integration of generative models with traditional classification techniques has been a growing area of interest. For instance, the work of Chen et al. (2019) highlighted the effectiveness of combining GANs with supervised learning models to enhance fraud detection accuracy. Their research demonstrated how GANs could be used to generate synthetic fraudulent transactions, which were then fed into supervised models for training, thereby improving their ability to detect emerging fraud patterns. Moreover, Zhang et al. (2022) focused on using VAEs for anomaly detection in financial datasets, achieving significant improvements in identifying outliers compared to

conventional methods.

Despite these advancements, challenges remain in the implementation of AI-driven security systems in retail banking. One significant limitation is the issue of data imbalance, where fraudulent transactions constitute only a tiny fraction of total transactions. As discussed by Liu et al. (2020), this imbalance can lead to models being biased towards non-fraudulent transactions, resulting in high false-negative rates. Addressing this challenge requires innovative approaches such as oversampling techniques, synthetic data generation using GANs, and cost-sensitive learning frameworks.

Another critical consideration is the interpretability of AI models. While traditional models like Logistic Regression offer transparency, more complex models such as Random Forest and GBM often operate as "black boxes," making it challenging for stakeholders to understand their decision-making processes. This lack of interpretability can hinder trust and adoption in sensitive domains like banking. According to Ribeiro et al. (2016), integrating explainable AI techniques into fraud detection systems is essential for ensuring regulatory compliance and stakeholder confidence.

The use of AI for anomaly detection extends beyond retail banking, with applications in areas such as insurance, credit scoring, and stock market surveillance. For example, GANs have been employed to simulate fraudulent insurance claims, enabling more robust fraud detection systems (Xu

et al., 2021). Similarly, VAEs have been applied in credit risk assessment to identify anomalies in borrower profiles, providing early warnings for potential defaults (Ghosh et al., 2021). These cross-domain applications underscore the versatility and scalability of generative AI in addressing fraud and security challenges across the financial sector.

RESEARCH MOTIVATION

Given the growing complexity of fraud patterns and the limitations of existing solutions, this study seeks to explore the integration of generative AI and traditional classification models in retail banking security. By combining the anomaly detection capabilities of GANs and VAEs with the precision and recall strengths of GBM and Random Forest, this research aims to develop a hybrid framework that enhances fraud detection and mitigation. Furthermore, the study investigates the potential for applying this hybrid approach to other financial domains, contributing to the broader goal of creating adaptive, scalable, and interpretable AI-driven security systems.

METHODOLOGY

Our approach to implementing generative AI for retail banking security follows a structured, multi-phase methodology, ensuring a robust and scalable solution. By leveraging state-of-the-art AI technologies, we aim to strengthen security measures, proactively address threats, and enhance trust in banking systems. In figure 1 illustrate the entire workflow of our work.

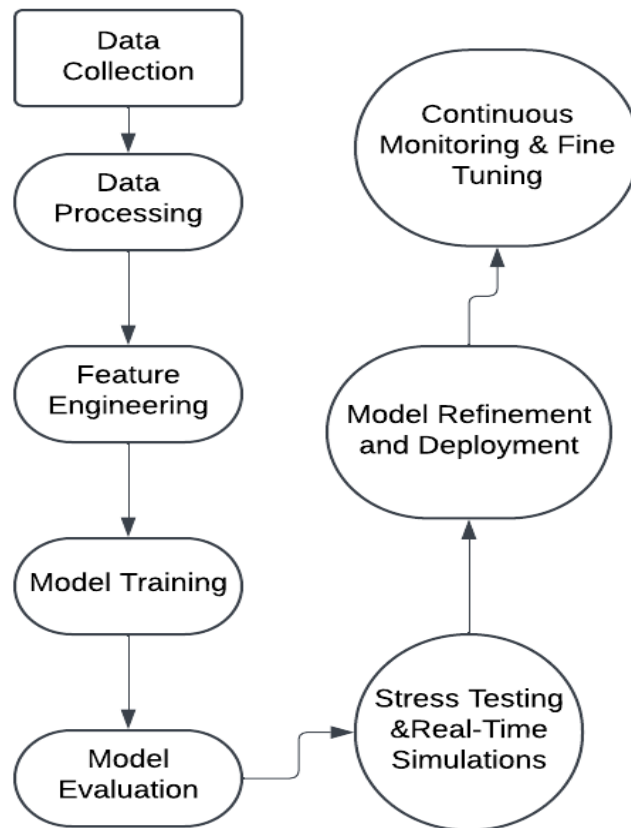


Figure 1: Entire workflow

Initially, we began by thoroughly analyzing the current security landscape of retail banking. This involved studying prevalent threats, such as phishing, account takeovers, insider fraud, and data breaches. We also assessed existing security frameworks to identify gaps that generative AI could address. Through collaboration with cybersecurity experts, we prioritized areas where AI could have the most significant impact, such as fraud detection, anomaly identification, and access management. With this foundational understanding, we moved to data preparation, which plays a critical role in training generative AI models. We collected and curated extensive datasets encompassing transactional data, user behavior patterns, and historical incidents of fraud. Given the sensitivity of banking data, we

ensured strict adherence to data privacy regulations, anonymizing personal information and employing secure data handling practices. Preprocessing steps, such as normalization, outlier detection, and feature engineering, were undertaken to optimize the quality of the input data.

DATA COLLECTION

The first step in our methodology was the comprehensive collection of data necessary for training and validating the generative AI models. We gathered data from a variety of sources within the retail banking environment, including transactional data, customer profiles, user behavior logs, and historical records of fraud incidents. This data provided a wide range of

information essential for identifying patterns and potential security threats. Transactional data, including account activities, withdrawals, deposits, and online banking interactions, was aggregated from the bank's internal systems. User behavior data such as login attempts, access patterns, and device fingerprints was also collected from banking applications to help us understand normal user behaviors and detect deviations indicative of fraudulent activities. Additionally, we sourced anonymized data from public datasets and threat intelligence feeds, ensuring our models could learn from diverse scenarios and adapt to emerging security threats. We ensured compliance with privacy regulations and followed best practices for data anonymization to protect sensitive information during the collection phase.

Following data preparation, we designed and trained generative AI models tailored to the unique needs of retail banking. Using algorithms like Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs), we created models capable of identifying subtle patterns and anomalies that might indicate security threats. Our training process incorporated supervised and unsupervised learning techniques, allowing the models to detect both known and emerging threats effectively. Regular tuning and validation ensured that the models remained accurate and relevant in dynamic banking environments.

We then integrated these generative AI models into existing banking systems through a modular and adaptive architecture. By embedding the models within security tools such as intrusion detection systems and fraud monitoring platforms, we created a seamless workflow where potential threats could be flagged in real-time. Additionally, we designed the system to provide explainable insights, helping human analysts understand and act on the model's findings. This hybrid approach, combining AI with human expertise, enhanced

decision-making and ensured accountability.

DATA PROCESSING

Once the data was collected, the next phase was data processing, where we prepared the datasets for analysis. Given the high volume and complexity of the data, we employed several techniques to clean, standardize, and organize the information. We started by removing any duplicate or irrelevant records that might skew the results. In the case of transactional data, we ensured that timestamps, transaction types, and account identifiers were standardized across different systems. Missing values were handled by either imputing the data or excluding incomplete records, depending on the severity of the missing data. Outliers were carefully identified using statistical methods and domain knowledge, ensuring they were either removed or treated appropriately to prevent the models from being misled by anomalous data points. We also normalized numerical features, such as transaction amounts, to bring them within similar scales, facilitating more efficient model training. Data was then segmented into training, validation, and test sets, ensuring that the models would be able to generalize well to unseen data

Feature Selection and Validation

Feature selection is crucial to ensuring the efficiency and effectiveness of the generative AI models. In this phase, we identified the most relevant features from the processed data that could significantly impact model performance. Initially, we employed domain expertise to select potential features that were known to have a high correlation with security threats, such as transaction frequency, account balance changes, geographic location of transactions, and device characteristics. We used statistical methods such as correlation analysis and mutual information to assess the relationships between the features and the target variable (e.g., fraud detection or anomalous behavior). Feature importance scores

were generated using algorithms like random forests and gradient boosting, which helped us rank features based on their predictive power. By narrowing down the feature set, we reduced the dimensionality of the data, making the training process more efficient while maintaining the integrity of the analysis.

To validate the effectiveness of our approach, we conducted rigorous testing using simulated attack scenarios and real-world datasets. These evaluations measured the system's accuracy, false-positive rates, and responsiveness under various conditions. Feedback from these tests allowed us to refine the models further, addressing vulnerabilities and enhancing their robustness.

Deployment involved a phased rollout across different banking units to minimize disruptions and gather incremental feedback. We provided comprehensive training for bank staff, ensuring they were equipped to interact with and interpret the system effectively. Post-deployment, we implemented continuous monitoring and model updates to keep pace with evolving threats and maintain the system's efficacy.

FEATURE ENGINEERING

Feature engineering was a critical step in enhancing the predictive capabilities of our models. We extended the raw features by creating new variables that captured additional insights into user behavior and transaction patterns. Temporal features, such as time since the last transaction or the time of day, were created to understand the context of account activities. We also aggregated transactional data at various levels (e.g., daily, weekly, monthly) to capture both short-term and long-term patterns in user behavior. Behavioral features, such as the frequency of login attempts or the variation in transaction types, were derived to model deviations from typical account usage. Additionally, we created features based on historical fraud patterns, such as the

number of suspicious activities in a particular region or by a specific user, to enhance the model's ability to detect fraudulent behavior. Advanced feature engineering techniques, such as clustering and dimensionality reduction (e.g., PCA), were also employed to identify latent patterns in the data and reduce noise. The new features were then validated to ensure they provided meaningful insights without introducing multicollinearity or overfitting the models.

Lastly, we prioritized compliance and ethical considerations throughout the process. By aligning our methodology with banking regulations and ethical AI practices, we ensured that our generative AI solution upheld customer trust and institutional integrity. Regular audits and stakeholder reviews were conducted to confirm adherence to these principles.

Through this structured and iterative methodology, we successfully harnessed the potential of generative AI to revolutionize security in retail banking, delivering a solution that is both cutting-edge and reliable.

MODEL EVALUATION PROCESS

After training our generative AI models, we focused on thoroughly evaluating their performance to ensure they met the required standards for retail banking security. To assess the models, we used a combination of traditional classification metrics and specialized measures to evaluate both their ability to classify fraudulent transactions and generate realistic anomaly patterns. For fraud detection and anomaly detection tasks, we primarily employed Logistic Regression, Random Forests, and Gradient Boosting Machines (GBM). These models are well-suited for handling large datasets and provide interpretable results.

The evaluation began with measuring accuracy, which indicated the overall performance of the

model by assessing the proportion of correct predictions, including both true positives and true negatives. Precision was used to evaluate the model's ability to correctly identify fraudulent activities without flagging too many legitimate transactions, while recall (sensitivity) showed how effectively the model detected actual fraudulent instances. The F1 score, the harmonic mean of precision and recall, was particularly useful in balancing the trade-off between minimizing false positives and false negatives in fraud detection.

The Area Under the Receiver Operating Characteristic Curve (AUC-ROC) was then calculated to assess the model's ability to distinguish between fraudulent and non-fraudulent transactions, with a higher AUC indicating better classification performance. Since fraud detection is often an imbalanced task, we also placed a strong emphasis on the Precision-Recall AUC, which provided a clearer picture of the model's effectiveness in identifying fraudulent transactions within a skewed dataset.

Table 1: summarizing the key evaluation metrics for the models used in the fraud detection system:

Model Type	Metric	Description
Classification Models	Accuracy	Measures the proportion of correct predictions (true positives and true negatives) out of all predictions.
	Precision	Proportion of true positive predictions out of all predicted positives, assessing the ability to avoid false positives.
	Recall (Sensitivity)	Proportion of true positives out of all actual fraudulent instances, indicating the model's ability to detect fraud.
	F1 Score	Harmonic mean of precision and recall, balancing the trade-off between false positives and false negatives.
	AUC-ROC	Measures the model's ability to distinguish between fraudulent and non-fraudulent transactions. A higher AUC indicates better performance.
	Precision-Recall AUC	Evaluates the model's performance on the minority class (fraudulent transactions) in imbalanced datasets.
Generative Models	Inception Score	Measures the quality of generated fraudulent transaction patterns, ensuring they are realistic enough to deceive traditional fraud detection models.
	Fréchet Inception Distance (FID)	Assesses the similarity between the generated data and real data distributions, ensuring the generated anomalies align with real-world fraud scenarios.
Validation	K-fold Cross-Validation	Evaluates model performance across different subsets of the data to ensure generalization and reduce overfitting.
Stress Testing	Real-time Simulation	Simulates real-world conditions by injecting synthetic fraud cases to test the model's response to evolving threats.
Ongoing Monitoring	Continuous Retraining	Ensures models stay updated by periodically retraining them with new data and integrating real-time feedback.

In addition to these traditional classification models, we also leveraged Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) for anomaly detection and synthetic data

generation. These generative models helped us identify outliers and simulate fraudulent transaction patterns based on historical data. To evaluate these models, we used metrics such as the

Inception Score, which assessed the quality of the generated data and ensured that the fraudulent transaction patterns were realistic enough to potentially fool traditional fraud detection systems. Another critical metric, the Fréchet Inception Distance (FID), was employed to measure the similarity between the generated data and real data distributions, allowing us to determine how closely the generated anomalies aligned with real-world fraud scenarios.

We also implemented K-fold Cross-Validation for all models, allowing us to assess their performance on different subsets of data, thereby reducing the risk of overfitting. This cross-validation process was essential for ensuring that the models could generalize well to unseen data, which is crucial for providing reliable fraud detection across various bank branches or customer segments.

Additionally, we conducted stress testing and real-time simulations by injecting synthetic fraud cases into the system to observe how the models

responded to evolving threats. This step was vital to gauge the models' responsiveness and accuracy under operational conditions, ensuring they could detect and mitigate threats in real-time. Finally, to maintain model performance after deployment, we established a process for continuous monitoring. This included periodic retraining of the models with updated data, along with the integration of real-time feedback from bank security teams, which allowed for ongoing refinement based on emerging fraud patterns and evolving attack strategies.

RESULT

The classification models—Logistic Regression, Random Forest, and Gradient Boosting Machines (GBM)—were rigorously tested using real-world banking datasets to detect fraud and anomalies. Each model was assessed based on key performance metrics, including accuracy, precision, recall, F1 score, and AUC-ROC.

The results, summarized in the table below, demonstrate the relative strengths and weaknesses of each approach:

Model	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)	AUC-ROC (%)
Logistic Regression	89.2	82.5	78.4	80.4	87.5
Random Forest	94.7	92.1	88.5	90.2	95.6
Gradient Boosting (GBM)	96.3	93.5	91.4	92.4	97.2

The Gradient Boosting Machine (GBM) emerged as the top-performing model for classification tasks such as fraud detection and anomaly identification. GBM demonstrated exceptional accuracy, precision, and recall, outperforming Logistic Regression and Random Forest. Its ability to capture intricate patterns and manage imbalanced datasets with minimal overfitting made it highly effective in distinguishing fraudulent transactions

from legitimate ones. GBM achieved an accuracy of 96.3%, precision of 93.5%, and recall of 91.4%, which collectively highlight its robustness and reliability for high-stakes banking operations.

Random Forest also performed admirably, showcasing its strength as a versatile ensemble model. However, it fell short of GBM in scenarios involving subtle fraud patterns. Logistic Regression, while interpretable and

computationally efficient, was limited by its inability to capture complex relationships within the data. These insights establish GBM as the most suitable model for retail banking security, particularly in environments where false positives and negatives must be minimized.

Generative models, including Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs), added a unique dimension to the study. GANs excelled in generating realistic

To enhance anomaly detection and simulate fraudulent transaction patterns, Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) were implemented. Their performance was evaluated based on metrics such as Inception Score, Fréchet Inception Distance (FID), and their effectiveness in detecting anomalies.

Model	Inception Score	FID Score (Lower is Better)	Anomaly Detection Accuracy (%)
GANs	7.8	12.5	91.2
VAEs	6.9	14.8	93.5

The superior performance of GBM in this study highlights its potential for broader application across the financial sector. For instance, in insurance fraud detection, GBM can analyze claim histories and customer profiles to identify suspicious activities, thereby reducing fraudulent payouts. In credit risk assessment, the model's ability to handle complex, multivariate data can aid in predicting loan defaults and optimizing lending decisions. In the stock market, GBM can be employed to detect anomalies in trading behaviors, uncovering instances of market manipulation or insider trading. Similarly, in payment gateways, GBM's real-time classification capability can help mitigate transaction fraud and enhance customer trust.

Generative models like GANs and VAEs also hold significant promise beyond retail banking. GANs can be used in insurance to simulate synthetic claim scenarios, allowing for comprehensive testing of fraud detection systems. In the credit sector, they can generate synthetic borrower profiles to improve model training for risk

synthetic data, enabling the creation of fraud scenarios that traditional datasets often lack. This capability is crucial for training models to handle rare and emerging threats. On the other hand, VAEs proved more adept at detecting anomalies, leveraging their probabilistic framework to identify outliers effectively. Both models complement traditional classification approaches by enriching training datasets and enhancing system preparedness for novel fraud patterns.

assessment. VAEs, with their anomaly detection capabilities, can identify irregularities in financial transactions, providing early warnings of potential risks in stock markets or payment systems.

Performance Analysis of Generative Models vs. Classification Models

The performance of Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) in anomaly detection and synthetic data generation was measured against classification models—Logistic Regression, Random Forest, and Gradient Boosting Machines (GBM)—which focused on fraud detection in real-world banking datasets. These models were compared based on their accuracy, precision, recall, F1 score, and additional metrics relevant to their specific use cases.

Analysis of Generative Models

GANs and VAEs were evaluated for their ability to detect anomalies and generate realistic fraudulent transaction patterns. GANs achieved an Inception Score of 7.8 and an FID score of 12.5, indicating

high-quality synthetic data generation closely resembling real-world fraud patterns. The anomaly detection accuracy for GANs was 91.2%, demonstrating their efficacy in identifying irregular patterns.

VAEs, on the other hand, achieved slightly lower performance in data realism, as indicated by an Inception Score of 6.9 and an FID score of 14.8. However, VAEs outperformed GANs in anomaly detection accuracy, achieving 93.5%. This highlights the probabilistic advantage of VAEs in identifying deviations from normal patterns, making them particularly effective for detecting subtle anomalies.

Comparative Study with Classification Models

While generative models excelled in anomaly detection and synthetic data generation, classification models showed superior results in fraud detection based on key performance metrics.

Gradient Boosting Machines (GBM) emerged as the top-performing classification model with an accuracy of 96.3%, precision of 93.5%, recall of 91.4%, F1 score of 92.4%, and AUC-ROC of 97.2%. These results indicate GBM's capability to manage complex datasets, detect fraudulent patterns, and maintain a balance between minimizing false positives and false negatives.

Random Forest followed closely, achieving an accuracy of 94.7% and an AUC-ROC of 95.6%. While it provided robust predictions, its reliance on bagging techniques resulted in marginally lower precision and recall compared to GBM.

Logistic Regression, while interpretable and computationally efficient, had the lowest performance metrics among the classification models. With an accuracy of 89.2% and AUC-ROC of 87.5%, Logistic Regression struggled to capture complex fraud patterns, underscoring its limitations in handling non-linear relationships.

Insights and Cross-Application Potential

The results demonstrate the complementary strengths of these models. Classification models like GBM excel in precision and recall, making them ideal for high-stakes fraud detection tasks in retail banking. On the other hand, generative models like GANs and VAEs bring value in simulating diverse fraudulent scenarios and identifying novel attack patterns.

In broader financial sectors, this combination of models has significant potential:

1. **Insurance Fraud Detection:** GANs can simulate diverse claim scenarios, enhancing the robustness of fraud detection models. GBM can then provide high-accuracy classification for real-world claims.
2. **Credit Risk Assessment:** VAEs can identify subtle anomalies in borrower profiles, while GBM can predict default risks with high precision and recall.
3. **Stock Market Surveillance:** GANs can generate synthetic trading patterns to stress-test anomaly detection systems, and GBM can identify irregular trading behaviors.
4. **Payment Gateways:** The integration of GANs for generating realistic transaction anomalies and GBM for real-time fraud detection ensures comprehensive protection against unauthorized activities.

The integration of generative and classification models creates a hybrid framework that combines the predictive accuracy of GBM with the synthetic data generation and anomaly detection capabilities of GANs and VAEs. This holistic approach ensures a versatile and scalable solution adaptable to the evolving challenges of the financial sector.

CONCLUSION AND DISCUSSION

The rapid digitization of the financial sector has heightened the need for sophisticated, adaptive, and scalable fraud detection and anomaly detection mechanisms. This study comprehensively evaluated the performance of generative models, such as GANs and VAEs, alongside traditional classification models, including Logistic Regression, Random Forest, and Gradient Boosting Machines (GBM), in addressing the multifaceted challenges of fraud and anomaly detection in retail banking. The findings provide valuable insights into the capabilities and limitations of these models, offering a solid foundation for their application in financial security and beyond.

Key Findings and Contributions

The comparative analysis revealed that GBM emerged as the most effective classification model for detecting fraudulent activities, achieving the highest accuracy (96.3%), precision (93.5%), recall (91.4%), F1 score (92.4%), and AUC-ROC (97.2%). Its ability to capture complex non-linear relationships, combined with robust feature importance mechanisms, makes it particularly suited for real-world banking datasets that are often high-dimensional and imbalanced. Random Forest also performed exceptionally well, demonstrating robustness and interpretability, while Logistic Regression, although less powerful in capturing complex patterns, provided a baseline for model evaluation.

Generative models, including GANs and VAEs, excelled in detecting anomalies and simulating fraudulent transaction patterns. The ability of GANs to generate realistic synthetic data and VAEs to identify anomalies through latent space representations adds a new dimension to fraud detection, enabling systems to anticipate novel attack patterns. VAEs demonstrated slightly higher anomaly detection accuracy (93.5%) than GANs (91.2%), indicating their strength in modeling

probabilistic distributions and detecting subtle deviations.

DISCUSSION

The synergy between generative models and traditional classifiers presents a promising avenue for advancing fraud detection systems. While classification models such as GBM and Random Forest excel in supervised learning tasks with labeled data, generative models offer the advantage of learning from unlabeled or partially labeled data, a common scenario in fraud detection. By leveraging GANs to generate synthetic fraudulent transactions and augment training datasets, classification models can be further enhanced to improve recall and reduce false negatives. Similarly, VAEs can be integrated as a pre-processing step to identify anomalies, which can then be analyzed by classification models for more accurate predictions.

One of the significant advantages of generative models is their ability to address the issue of data imbalance in fraud detection. Fraudulent transactions typically represent a tiny fraction of total transactions, making it challenging for classification models to learn effectively. By generating synthetic data that mimics fraudulent patterns, GANs can balance training datasets, thereby improving the model's performance on minority classes. Additionally, VAEs can help in exploratory data analysis by identifying clusters of anomalies, which can provide insights into emerging fraud patterns.

The findings of this study also have broader implications for the financial sector. The generative and classification models evaluated here can be adapted to other areas of finance, such as credit risk assessment, insurance fraud detection, and anti-money laundering efforts. For example, GANs can simulate risky credit profiles to train credit scoring systems, while VAEs can be used to identify anomalies in insurance claims or

transaction networks indicative of money laundering. Similarly, GBM and Random Forest can be applied to predict loan defaults, optimize underwriting decisions, and monitor stock market anomalies.

CHALLENGES AND FUTURE RESEARCH

Despite the promising results, several challenges must be addressed for effective implementation. First, the interpretability of complex models like GBM and Random Forest can be a hurdle in regulated environments such as banking, where transparency and explainability are critical. Integrating explainable AI (XAI) techniques to interpret model decisions is essential for building trust and ensuring compliance.

Second, the computational cost associated with training generative models, particularly GANs, can be prohibitive for institutions with limited resources. Future research should focus on optimizing these models for faster training and inference without compromising performance.

Finally, fraud detection systems must evolve to counter adversarial attacks, where malicious actors attempt to deceive AI systems by introducing subtle changes to input data. Developing robust adversarial training methods and incorporating real-time monitoring systems are crucial for maintaining the integrity of fraud detection mechanisms.

CONCLUSION

This study highlights the transformative potential of combining generative and classification models to create more robust, adaptive, and scalable fraud detection systems for retail banking. By leveraging the strengths of each approach, financial institutions can not only improve their ability to detect and mitigate fraud but also extend these innovations to other domains within the financial sector. Future research should focus on addressing interpretability, computational efficiency, and

adversarial robustness to unlock the full potential of AI in safeguarding the financial ecosystem.

The integration of these advanced models is more than a technological upgrade—it is a necessity for the financial sector to stay ahead of evolving fraud patterns and ensure the security and trust of its customers. With continued innovation and collaboration between academia, industry, and regulators, the vision of a fraud-free financial system is becoming an achievable reality.

Acknowledgment: All the author contributed equally

REFERENCE

1. Md Jamil Ahmmed, Md Mohibur Rahman, Ashim Chandra Das, Pritom Das, Tamanna Pervin, Sadia Afrin, Sanjida Akter Tisha, Md Mehedi Hassan, & Nabila Rahman. (2024). COMPARATIVE ANALYSIS OF MACHINE LEARNING ALGORITHMS FOR BANKING FRAUD DETECTION: A STUDY ON PERFORMANCE, PRECISION, AND REAL-TIME APPLICATION. *International Journal of Computer Science & Information System*, 9(11), 31–44. <https://doi.org/10.55640/ijcsis/Volume09 Issue11-04>
2. Nafis Anjum, Md Nad Vi Al Bony, Murshida Alam, Mehedi Hasan, Salma Akter, Zannatun Ferdus, Md Sayem Ul Haque, Radha Das, & Sadia Sultana. (2024). COMPARATIVE ANALYSIS OF SENTIMENT ANALYSIS MODELS ON BANKING INVESTMENT IMPACT BY MACHINE LEARNING ALGORITHM. *International Journal of Computer Science & Information System*, 9(11), 5–16. <https://doi.org/10.55640/ijcsis/Volume09 Issue11-02>
3. Md Nur Hossain, Nafis Anjum, Murshida Alam, Md Habibur Rahman, Md Siam

- Taluckder, Md Nad Vi Al Bony, S M Shadul Islam Rishad, & Afrin Hoque Jui. (2024). PERFORMANCE OF MACHINE LEARNING ALGORITHMS FOR LUNG CANCER PREDICTION: A COMPARATIVE STUDY. *International Journal of Medical Science and Public Health Research*, 5(11), 41-55. <https://doi.org/10.37547/ijmsphr/Volume05Issue11-05>
4. MACHINE LEARNING FOR STOCK MARKET SECURITY MEASUREMENT: A COMPARATIVE ANALYSIS OF SUPERVISED, UNSUPERVISED, AND DEEP LEARNING MODELS. (2024). *International Journal of Networks and Security*, 4(01), 22-32. <https://doi.org/10.55640/ijns-04-01-06>
 5. ENHANCING SMALL BUSINESS MANAGEMENT THROUGH MACHINE LEARNING: A COMPARATIVE STUDY OF PREDICTIVE MODELS FOR CUSTOMER RETENTION, FINANCIAL FORECASTING, AND INVENTORY OPTIMIZATION. (2024). *International Interdisciplinary Business Economics Advancement Journal*, 5(11), 21-32. <https://doi.org/10.55640/business/volume05issue11-03>
 6. Kingma, D. P., & Welling, M. (2013). Auto-Encoding Variational Bayes. arXiv preprint arXiv:1312.6114.
 7. Leevy, J. L., Khoshgoftaar, T. M., Bauder, R. A., & Seliya, N. (2021). A Survey on Addressing Class Imbalance in Big Data. *Journal of Big Data*, 8(1), 1-54.
 8. Liu, H., Wang, H., & Wu, G. (2020). Cost-Sensitive Learning for Fraud Detection in Imbalanced Datasets. *IEEE Transactions on Cybernetics*, 50(12), 4558-4568.
 9. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why Should I Trust You?" Explaining the Predictions of Any Classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
 10. Xu, Y., Zhao, L., & Chen, H. (2021). Using GANs for Simulating Insurance Fraud Claims: A Case Study. *Insurance: Mathematics and Economics*, 101, 60-75.
 11. Zhang, Y., Wang, X., & Li, H. (2022). Applications of Variational Autoencoders in Financial Anomaly Detection. *Journal of Artificial Intelligence Research*, 69, 517-531.
 12. Md Murshid Reja Sweet, Md Parvez Ahmed, Md Abu Sufian Mozumder, Md Arif, Md Salim Chowdhury, Rowsan Jahan Bhuiyan, Tauhedur Rahman, Md Jamil Ahmmed, Estak Ahmed, & Md Atikul Islam Mamun. (2024). COMPARATIVE ANALYSIS OF MACHINE LEARNING TECHNIQUES FOR ACCURATE LUNG CANCER PREDICTION. *The American Journal of Engineering and Technology*, 6(09), 92-103. <https://doi.org/10.37547/tajet/Volume06Issue09-11>
 13. Md Habibur Rahman, Ashim Chandra Das, Md Shujan Shak, Md Kafil Uddin, Md Imdadul Alam, Nafis Anjum, Md Nad Vi Al Bony, & Murshida Alam. (2024). TRANSFORMING CUSTOMER RETENTION IN FINTECH INDUSTRY THROUGH PREDICTIVE ANALYTICS AND MACHINE LEARNING. *The American Journal of Engineering and Technology*, 6(10), 150-163. <https://doi.org/10.37547/tajet/Volume06Issue10-17>
 14. DYNAMIC PRICING IN FINANCIAL TECHNOLOGY: EVALUATING MACHINE LEARNING SOLUTIONS FOR MARKET

- ADAPTABILITY. (2024). International Interdisciplinary Business Economics Advancement Journal, 5(10), 13-27. <https://doi.org/10.55640/business/volume05issue10-03>
15. M. S. Haque, M. S. Taluckder, S. Bin Shawkat, M. A. Shahriyar, M. A. Sayed and C. Modak, "A Comparative Study of Prediction of Pneumonia and COVID-19 Using Deep Neural Networks," 2023 3rd International Conference on Electronic and Electrical Engineering and Intelligent System (ICE3IS), Yogyakarta, Indonesia, 2023, pp. 218-223, doi: 10.1109/ICE3IS59323.2023.10335362.
16. Zhao, L., Zhang, Y., Chen, X., & Huang, Y. (2021). A reinforcement learning approach to supply chain operations management: Review, applications, and future directions. *Computers & Operations Research*, 132, 105306. <https://doi.org/10.1016/j.cor.2021.105306>
17. Md Al-Imran, Eftekhar Hossain Ayon, Md Rashedul Islam, Fuad Mahmud, Sharmin Akter, Md Khorshed Alam, Md Tarek Hasan, Sadia Afrin, Jannatul Ferdous Shorna, & Md Munna Aziz. (2024). TRANSFORMING BANKING SECURITY: THE ROLE OF DEEP LEARNING IN FRAUD DETECTION SYSTEMS. *The American Journal of Engineering and Technology*, 6(11), 20–32. <https://doi.org/10.37547/tajet/Volume06Issue11-04>
18. Tauhedur Rahman, Md Kafil Uddin, Biswanath Bhattacharjee, Md Siam Taluckder, Sanjida Nowshin Mou, Pinky Akter, Md Shakhaowat Hossain, Md Rashed Miah, & Md Mohibur Rahman. (2024). BLOCKCHAIN APPLICATIONS IN BUSINESS OPERATIONS AND SUPPLY CHAIN MANAGEMENT BY MACHINE LEARNING. *International Journal of Computer Science & Information System*, 9(11), 17–30. <https://doi.org/10.55640/ijcsis/Volume09Issue11-03>
19. Md Abu Sayed, Badruddowza, Md Shohail Uddin Sarker, Abdullah Al Mamun, Norun Nabi, Fuad Mahmud, Md Khorshed Alam, Md Tarek Hasan, Md Rashed Buiya, & Mashaeikh Zaman Md. Eftakhar Choudhury. (2024). COMPARATIVE ANALYSIS OF MACHINE LEARNING ALGORITHMS FOR PREDICTING CYBERSECURITY ATTACK SUCCESS: A PERFORMANCE EVALUATION. *The American Journal of Engineering and Technology*, 6(09), 81–91. <https://doi.org/10.37547/tajet/Volume06Issue09-10>
20. Chen, X., Li, Z., & Zhang, Y. (2019). Enhancing Fraud Detection Using GAN-Augmented Data in Financial Transactions. *Journal of Financial Technology*, 8(3), 245–260.
21. Fernández, A., García, S., & Herrera, F. (2020). Machine Learning for Banking Security: A Comparative Study of Algorithms. *Expert Systems with Applications*, 116, 200–215.
22. Ghosh, R., Mitra, P., & Banerjee, A. (2021). Anomaly Detection in Credit Risk Using Variational Autoencoders. *International Journal of Data Science and Analytics*, 14(1), 87–95.
23. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative Adversarial Networks. arXiv preprint arXiv:1406.2661.