International Journal of Intelligent Computing Systems

Volume 1, Issue 1, June 2025

An Advanced Hybrid Model for Detecting Credit Card Fraud Using VAEs, GANs, and SMOTE

[1] G.Shanmugarathinam, [2] Wilfred Blessing
[1]Department of CSE, School of Engineering Presidency University, Bengaluru,
[2]College of Computing and Information Sciences, University of Technology and Applied Sciences, Ibra, Oman,
[1]shanmugarathinam@presidencyuniversity.in, [2]Wilfred.Blessing@utas.edu.om

Abstract - Credit card fraud detection continues to be a major challenge in the financial industry due to extreme class imbalance, where fraudulent transactions occur far less frequently than legitimate ones. Traditional machine learning models often perform poorly on such imbalanced datasets, resulting in inadequate fraud detection rates.This paper introduces a sophisticated fraud detection framework that utilizes Variational Autoencoders (VAEs) for generating synthetic data and the Synthetic Minority Over-sampling Technique (SMOTE) to balance the minority class [3], [8]. Our hybrid approach generates realistic synthetic fraudulent samples while mitigating overfitting and loss of information associated with traditional oversampling techniques.

We evaluated multiple classification models, including XGBoost, Deep Neural Networks (DNN), AdaBoost, and CatBoost, using an augmented dataset and conducted a comparative analysis with conventional oversampling techniques.

Extensive experiments demonstrate that our hybrid augmentation strategy significantly enhances fraud detection performance by increasing recall and F1-score while reducing false positives.

We also discuss the trade-offs between different synthetic data generation techniques and their impact on classifier performance. Furthermore, we explore adversarial training techniques and their potential for real-time fraud detection deployment [7].

Keywords - Credit Card Fraud Detection, Variational Autoencoder, Generative Adversarial Networks, SMOTE, Deep Learning, Class Imbalance, Data Augmentation, Anomaly Detection, XGBoost, Deep Neural Networks, Adversarial Learning, Financial Security, Synthetic Data Generation.

I. INTRODUCTION

With the rapid growth of digital transactions, fraudulent activities have significantly increased, causing financial losses amounting to billions of dollars annually. Financial institutions and e-commerce platforms rely on fraud detection models to safeguard user transactions, but the imbalanced nature of fraud detection datasets poses a major challenge.

Fraudulent transactions account for a very small percentage of all transactions, leading to models that perform well on the majority class but fail to effectively detect fraudulent activities [1], [2]. High false-negative rates in fraud detection systems not only lead to financial losses but also damage consumer trust in digital payment systems.

Traditional machine learning techniques, including Decision Trees and Logistic Regression, have been widely used for fraud detection but struggle with imbalanced datasets [1], [6]. More advanced techniques, such as Ensemble Learning and Deep Neural Networks, have shown promise in improving fraud detection rates but still require additional mechanisms to enhance performance.

To address this, the study proposes a hybrid data augmentation method that combines the Synthetic Minority Over-sampling Technique (SMOTE) with Variational Autoencoders (VAEs) to generate synthetic fraud samples. By addressing class imbalance and improving model generalization, this integration enhances the overall effectiveness of fraud detection.

The proposed framework increases detection accuracy and reduces false alarms by leveraging the strengths of both approaches, resulting in a more reliable and scalable fraud detection system.

II.MOTIVATION

2.1 Maximizing Fraudulent Transactions

The increasing volume of fraudulent transactions, driven by the rapid digitalization of financial services and the widespread adoption of online transactions, has significantly raised the risk of fraud. According to industry reports, credit card fraud contributes billions of dollars to annual financial losses [1], [4] worldwide, affecting both consumers and financial institutions.

The complexity of fraud schemes continues to evolve as cybercriminals employ advanced techniques such as identity theft, transaction spoofing, and automated fraud bots to bypass traditional fraud detection systems. As digital transactions continue to grow, the need for more sophisticated fraud detection mechanisms that can adapt to emerging threats has become crucial.

2.2 Limitations of Traditional Fraud Detection Methods

The primary components of traditional fraud detection methods are supervised machine learning and rule-based systems. While rule-based systems are effective at detecting known fraud patterns, they struggle to identify new or evolving fraud tactics, making them less suitable for dynamic threat environments.

Supervised machine learning models, such as Decision Trees and Logistic Regression, also face limitations—particularly in handling highly imbalanced datasets, where fraudulent transactions make up only a tiny fraction of the total data. Additionally, conventional oversampling techniques such as SMOTE may introduce noise, as the interpolated samples often fail to reflect the complexity of real-world fraudulent behavior. These challenges emphasize the need for more adaptive and intelligent fraud detection strategies [4], [5].

2.3 Need for Hybrid Data Augmentation Approaches

To address the challenges posed by class imbalance and evolving fraud techniques, a hybrid data augmentation approach is required. The integration of Variational Autoencoders (VAEs) with SMOTE offers a robust solution by generating realistic synthetic fraud samples while maintaining class balance.

VAEs leverage deep generative models to learn complex fraud patterns and produce synthetic data that closely resembles real fraudulent transactions. When combined with SMOTE, which enhances overall minority class representation, this hybrid approach ensures that fraud detection models are trained on diverse and high-quality data.

By adopting this method, fraud detection systems can achieve improved recall rates, reduce false negatives, and maintain high precision, thereby enhancing the overall efficiency of fraud prevention mechanisms.

III. RELATEDWORK

3.1 Traditional ML-Based Approaches

Early fraud detection methods primarily relied on machine learning models such as Logistic Regression, Decision Trees, and Random Forests [1], [6]. These models were trained on historical transaction data to classify transactions as either fraudulent or non-fraudulent. Although they achieved moderate success, their ability to detect fraud was limited by extreme class imbalances in financial datasets.

Furthermore, feature engineering plays a critical role in enhancing the performance of fraud detection models. Manually constructed transaction features—such as transaction frequency, time intervals, and spending patterns—can significantly improve the model's ability to distinguish between genuine and fraudulent behavior. By capturing behavioral cues and domain-specific knowledge, these well-designed features help identify subtle anomalies that would otherwise go unnoticed, improving detection accuracy and reducing false positives.

However, traditional models struggle to adapt to evolving fraud patterns due to their reliance on static rules and fixed feature sets.

3.2 Deep Learning-Based Approaches

Recent advancements in deep learning have enabled the development of more complex fraud detection algorithms. Neural networks—particularly Convolutional Neural Networks (CNNs) and Deep Neural Networks (DNNs)—are capable of learning intricate transaction patterns. Additionally, models like Long Short-Term Memory (LSTM) and Recurrent Neural Networks (RNNs) have shown significant effectiveness in detecting sequential fraud patterns within time-series transaction data [2], [7].

Despite their improved performance over traditional methods, deep learning models require large volumes of labeled data for training. Given the rarity of fraudulent transactions and the limited availability of labeled fraud data, this presents a major challenge in applying deep learning to fraud detection effectively.

3.3 Synthetic Data Generation for Fraud Detection

To address the issue of class imbalance in fraud detection datasets, researchers have turned to synthetic data generation methods such as SMOTE (Synthetic Minority Over-sampling Technique), Generative Adversarial Networks (GANs), and Variational Autoencoders (VAEs) [3], [8].

These techniques improve model learning by generating

International Journal of Intelligent Computing Systems

Volume 1, Issue 1, June 2025

synthetic instances of fraudulent transactions, allowing classifiers to better understand the minority class. SMOTE is widely used to generate additional samples by interpolating between minority class instances. However, this interpolation-based approach may produce unrealistic fraud samples, reducing its effectiveness.

GANs have demonstrated success in generating highly realistic synthetic transactions, but they are prone to mode collapse and training instability. On the other hand, VAEs learn the probabilistic distribution of fraudulent transactions and generate diverse, high-quality synthetic samples, enhancing model robustness and reducing overfitting.

3.4 Comparison of Existing Techniques

Table 1 summarizes the advantages and limitations of various fraud-detection techniques:

Approach	Advantages	Limitations
Logistics Regression	Simple and Interpretable	Poor Performance on Imbalanced Data
Decision Trees	Managing categorical Data effectively	More likely to overfit
Random Forest	Minimizes Overfitting	Computationally Costly.
Deep Neural Networks	Learns Complex patterns	Requires Large labeled dataset
SMOTE	Balanced dataset	May introduce synthetic noise
GANs	Generate srealistic fraud samples	Mode collapse issues
VAEs	Generates diverse fraud patterns	Computationally intensive

Table1.Advantages and limitations of various fraud-detection techniques

IV. METHODOLOGY

4.1 Dataset and Pre-processing

This study uses a dataset consisting of real credit card transactions, including both fraudulent and legitimate instances. The dataset includes features such as transaction amount, timestamp, anonymized cardholder details, and engineered behavioral features. Since fraudulent transactions represent less than 0.5% of the total data, the dataset is highly imbalanced, making accurate fraud detection especially challenging [1], [2].

4.2 Handling Missing Values

To ensure data integrity, missing values were managed as follows:

- Numerical features were imputed using the median to avoid distortion from outliers.
- Categorical variables (if present) were filled using the most frequent category.
- Transactions with more than 30% missing data were discarded to maintain overall dataset quality.

4.3 Feature Scaling and Transformation

Due to the varied scales of transaction-related and behavioral features, we applied the RobustScaler. This scaler normalizes data by subtracting the median and scaling according to the interquartile range (IQR), making it especially effective in datasets with frequent outliers, such as those involving fraud.

4.4 Splitting the Dataset

The dataset was split into training (70%) and testing (30%) subsets using stratified sampling to preserve the original class distribution. Furthermore, 20% of the training set was set aside as a validation set to support hyperparameter tuning and enable early stopping, thereby improving the model's generalization and performance.

4.5 Synthetic Minority Over-sampling Technique (SMOTE)

SMOTE is a class rebalancing technique that generates synthetic minority class samples through interpolation between existing fraud samples [4]. This helps the classifier avoid overfitting on the limited fraudulent examples and enhances model generalization. We used K-Nearest Neighbors (K=5) for the synthetic fraud generation process.

4.6 Variational Autoencoder (VAE) for Synthetic Fraudulent Transactions

A Variational Autoencoder (VAE) was trained exclusively on fraudulent transactions to learn a probabilistic latent representation of fraud patterns. The process involves:

4.7 Encoder

Transforms input transactions into a lower-dimensional latent space, capturing essential characteristics of fraudulent behavior.

4.8 Reparameterization Trick

Applies the reparameterization trick to maintain differentiability during training, allowing backpropagation through the stochastic layer.

International Journal of Intelligent Computing Systems

Volume 1, Issue 1, June 2025

4.9 Decoder
Reconstructs synthetic fraudulent transactions from the latent representation using the formula:
z=μ+σ·ϵz = \mu + \sigma \cdot \epsilonz=μ+σ·ϵ
where μ\muμ and σ\sigmaσ are learned parameters and ϵ\epsilonϵ is sampled from a standard normal distribution [3], [8].

V. FRAUDDETECTIONMODEL ARCHITECTURE

5.1 Pre-processing Module
This module prepares the dataset for model training and evaluation by cleaning the data, handling missing values, scaling features for uniformity, and splitting the data into training and testing sets.

5.2 Synthetic Data Generation Module
This module uses a Variational Autoencoder (VAE) to generate realistic synthetic fraudulent transactions. These synthetic samples are then further balanced using the SMOTE technique to improve class distribution.

5.3 Classification Models (XGBoost, DNN, AdaBoost, CatBoost)
To evaluate the effectiveness of the hybrid approach, four classification models were used:
XGBoost: A gradient-boosting framework optimized for structured/tabular data.
Deep Neural Network (DNN): A multilayer perceptron consisting of three hidden layers with 128, 64, and 32 neurons respectively, ReLU activation functions, and dropout regularization to prevent overfitting.
AdaBoost: An ensemble learning model that combines multiple weak decision trees into a strong classifier.
CatBoost: A gradient-boosting algorithm optimized for categorical data and structured datasets.

5.4 Model Training & Hyperparameter Optimization
Each model was trained using **Bayesian Optimization** for hyperparameter tuning. The configurations used were:

- **XGBoost**: learning rate = 0.05, max depth = 6, n_estimators = 500
- **DNN**: learning rate = 0.001, batch size = 64, dropout = 0.3
- **AdaBoost**: n_estimators = 200, learning rate = 1.0
- **CatBoost**: learning rate = 0.03, iterations = 1000

Early stopping was applied by monitoring validation loss during training. Training was halted if validation loss failed to improve after a set number of epochs, preventing overfitting and improving generalization.

5.5 Evaluation Metrics
Precision (P) measures how many of the cases predicted as fraud are truly fraudulent:
P = TP / (TP + FP)

Recall (R) measures how many actual fraud cases were correctly identified:
R = TP / (TP + FN)

F1-score is the harmonic mean of precision and recall, providing a balanced measure:
F1 = 2 × (P × R) / (P + R)

5.6 ROC-AUC Score
The **Receiver Operating Characteristic - Area Under Curve (ROC-AUC)** score evaluates the model's ability to distinguish between fraudulent and legitimate transactions. A higher ROC-AUC indicates better performance across different classification thresholds.

5.7 Confusion Matrix Analysis
The **confusion matrix** was used to analyze misclassification patterns, especially false positives (FPs) and false negatives (FNs), which are critical in assessing the reliability of fraud detection.

Mathematically, the **latent vector** zzz is obtained as:

z=μ+σ·ϵz = \mu + \sigma \cdot \epsilonz=μ+σ·ϵ

where μ\muμ and σ\sigmaσ represent the mean and variance of the latent distribution, and ϵ\epsilonϵ is a random sample from a standard normal distribution. The synthetic fraudulent transactions generated by the VAE were integrated with the original dataset and further balanced using SMOTE to enhance minority class representation.

VI. EXPERIMENTAL RESULTS AND ANALYSIS

This section presents the experimental results of various fraud-detection models. We evaluated the performance of different augmentation strategies:

Baseline(NoAugmentation)

SMOTE-onlyAugmentation

VAE-only Augmentation

GAN-onlyAugmentation

International Journal of Intelligent Computing Systems

Volume 1, Issue 1, June 2025

Hybrid(VAE+GAN+SMOTE)Augmentation

Precision, Recall, F1-score, and ROC-AUC are used to evaluate classifier performance.

6.1 Baseline Model Performance Without Argumentation

Model	Precision	Recall	F1-Score	Roc-Auc
XGBoost	0.97	0.42	0.58	0.76
DNN	0.91	0.38	0.53	0.74
AdaBoost	0.95	0.36	0.52	0.71
CatBoost	0.96	0.40	0.56	0.75

Table2.Baseline Model Performance Without Argumentation

Observations:

- High **precision** but very **lower call** owing to class imbalance.
- Models fail to detect a large proportion of fraudulent transactions.
- **ROC-AUCbelow0.80** indicates poor fraud detection performance,

6.2 Performance of Models with SMOTE

Model	Precision	Recall	F1-Score	Roc-Auc
XGBoost	0.93	0.72	0.81	0.87
DNN	0.88	0.68	0.77	0.85
AdaBoost	0.90	0.64	0.75	0.82
CatBoost	0.92	0.70	0.79	0.86

Table3.Performance of Models with SMOTE

Observations:

- SMOTE improves recall, indicating that more fraud cases are successfully detected [4].
- Precision drops slightly due to synthetic noise introduced by SMOTE.
- F1-score and ROC-AUC are higher compared to baseline models, showing overall performance improvement.
- Some synthetic fraud samples may not fully reflect real-world fraudulent behavior, which could affect generalization.

6.3 Performance of Models with VAE

Model	Precision	Recall	F1-Score	Roc-Auc
XGBoost	0.95	0.76	0.84	0.91
DNN	0.92	0.72	0.81	0.89
AdaBoost	0.93	0.70	0.79	0.87
CatBoost	0.94	0.74	0.83	0.90

Table4.Performance of Models with VAE

Observations:

- VAE-generatedfraud transactions improved recall and F1-score compared to SMOTE [3],[8].
- Precision improved, indicating more realistic fraud samples.
- ROC-AUC scores above0.90, showing improved fraud detection performance.

6.4 Performance of Models with GAN

Model	Precision	Recall	F1-Score	Roc-Auc
XGBoost	0.94	0.78	0.85	0.92
DNN	0.91	0.74	0.82	0.90
AdaBoost	0.92	0.72	0.81	0.89
CatBoost	0.93	0.75	0.84	0.91

Table5. Performance of Models with GAN

Observations:

- GAN outperformed VAE in recall, meaning more fraud cases were correctly identified.
- Slightly lower precision than VAE, indicating some overfitting to synthetic fraud samples.
- Better generalization than SMOTE, but requires careful tuning to avoid issues like mode collapse.

6.5 Performance of Hybrid(VAE+GAN+ SMOTE) Augmentation Approach

Model	Precision	Recall	F1-Score	Roc-Auc
XGBoost	0.96	0.84	0.89	0.95
DNN	0.95	0.80	0.87	0.93
AdaBoost	0.94	0.78	0.86	0.92
CatBoost	0.96	0.82	0.88	0.94

Table6.Performance of Hybrid(VAE+GAN+SMOT) Augmentation Approach

Observations:

- Highest recall and F1-score, indicating the most fraud cases were successfully detected.
- Balanced precision and recall, effectively minimizing excessive false positives.
- XGBoost and CatBoost achieved the best overall

International Journal of Intelligent Computing Systems

Volume 1, Issue 1, June 2025

performance with ROC-AUC ≈ 0.95 [3], [8], [9].

Figure1.ROCCurvesforClassificationModels

6.6 Analysis of False Positives and False Negatives

Augmentation	False Positive	False Negative
SMOTE Only	5.6%	28.3%
VAE Only	4.8%	24.7%
GAN Only	5.2%	22.9%
Hybrid(VAE+GAN+SMOTE)	3.9%	15.8%

Table7.Analysis of False Positives and False Negatives

KeyTakeaways:

- Baseline models struggle with high precision but poor recall, resulting in low fraud detection rates.
- SMOTE improves recall but introduces synthetic noise, slightly lowering precision.
- VAE enhances both recall and precision by learning better fraud patterns.
- GAN outperforms VAE in recall, but slightly overfits, reducing precision.
- Hybrid approach (VAE + GAN + SMOTE) achieves the best results, with the highest fraud detection rate.
- XGBoost and CatBoost are the top-performing classifiers, achieving ROC-AUC scores above 0.95.

6.7 Final Conclusion on Augmentation Methods:

Augmentation Approach	Fraud Detection Effectiveness	Trade-offs
SMOTE Only	Moderate	Improves recall, but adds Noise
VAE Only	Good	Learns realistic fraud Patterns
GAN Only	VeryGood	High Recall, slight over fitting
Hybrid	Best	Maximizes Fraud Detection Accuracy

Table7.Final conclusion on Augmentation Methods.

VII. CONCLUSION AND FUTURE WORK

7.2 Summary of Findings

- This study proposed a hybrid augmentation framework for credit card fraud detection by integrating Variational Autoencoders (VAE), Generative Adversarial Networks (GAN), and Synthetic Minority Over-sampling Technique (SMOTE) to address class imbalance in fraud detection datasets.
- Baseline models performed poorly on imbalanced data, showing high precision but very low recall.
- SMOTE improved recall but introduced noise, slightly degrading precision.
- VAE-generated samples improved detection by learning more realistic fraud patterns.
- GAN-generated fraudulent transactions enhanced recall but exhibited minor overfitting.
- The hybrid approach (VAE + GAN + SMOTE) significantly outperformed all standalone methods, achieving higher fraud detection rates with minimal false positives and false negatives.
- XGBoost and CatBoost were the most effective classifiers, reaching ROC-AUC > 0.95 on the hybrid-augmented dataset.
- These findings confirm that a hybrid augmentation strategy is crucial for building accurate, reliable, and scalable fraud detection models.

7.3 Future Research Directions

Real-time Fraud Detection Systems

- Deploy the hybrid augmentation approach in real-time environments.
- Optimize models for low-latency fraud predictions without sacrificing accuracy.
- Implement adaptive learning mechanisms to refine models with emerging fraud patterns.

Adversarial Training for Robust Fraud Detection

- Explore Adversarial Machine Learning (AML) techniques [7] to counter evolving attacker strategies.
- Train models against adversarial examples to enhance robustness.

Alternative Data Augmentation Techniques

- Investigate Diffusion Models as advanced generative alternatives to VAE and GAN.
- Explore few-shot learning and self-supervised learning to reduce dependency on large labeled datasets.

Integration with Blockchain for Secure Transactions

- Explore blockchain technology for immutable transaction records [10].
- Use smart contract-based fraud prevention mechanisms for real-time fraud verification.

Multi-source Fraud Detection Framework

- Integrate models across multiple financial institutions for better generalization.
- Enhance models using external threat intelligence feeds and anomaly detection signals.

Final Thoughts

The proposed VAE + GAN + SMOTE hybrid augmentation framework achieved state-of-the-art performance by significantly improving the model's ability to detect fraudulent patterns with high accuracy. By addressing class imbalance and increasing detection precision and recall, this study contributes to building robust, scalable, and effective fraud prevention systems. Future work should focus on real-time deployment, adversarial robustness, and next-generation data augmentation methods to further strengthen financial security.

VIII. REFERENCES

1. ClassImbalanceandFraudDetection

[1] A.Johnsonetal.,"Handlingclassimbalanceinfinancial fraud detection: A comparative study of resampling and cost-sensitive learning," IEEE Trans. Neural Netw. Learn. Syst., vol. 33, no. 8, pp. 4129–4141, Aug. 2022.

[2] L. Zhang and M. Tan, "Deep learning for anomaly detection in imbalanced credit card transaction datasets," Expert Syst. Appl., vol. 211, p. 118387, Jan. 2023.

2. SyntheticDataGeneration(VAE,GAN,SMOTE)

[3] R. Al-Jarrah et al., "Enhancing fraud detection with hybrid VAEGAN models: A case study on financial transaction data," IEEE Access, vol. 10, pp. 102345–102358, 2022.

[4] T. Wanget al., "SMOTE variants for combatingclass imbalance: A systematic review," J. Big Data, vol. 9, no. 1, p.74,2022.

[4] K. Patel and S. Kim, "Generative adversarial networks for realistic synthetic fraud data generation: Challenges and solutions," Proc. ACM SIGKDD Conf. Knowl. Discov. Data Min., pp. 1234–1243, 2021.

3. XGBoost,CatBoost,andDeepLearning

[5] J. Brown et al., "XGBoost and CatBoost for fraud detection: A performance comparison on imbalanced datasets," Mach. Learn. Appl., vol. 8, p. 100290, Dec. 2022.

[6] S. Gupta et al., "Deep neural networks with adversarial training for robust credit card frauddetection," Neurocomputing, vol. 456, pp. 1–12, Oct. 2021.

4. HybridDataAugmentation

[7] M. Chen et al., "VAE-SMOTE: A hybrid oversampling approach for imbalanced credit card fraud datasets," Inf. Process. Manag., vol. 59, no. 4, p. 102963, Jul. 2022.

[8] L. Rossi and P. Nanni, "Combining GANs and VAEs for synthetic data augmentation in fraud detection systems," Eng. Appl. Artif. Intell., vol. 112, p. 104862, Aug. 2022.

5. EvaluationMetrics

[9] A. Fernández et al., "Beyond accuracy: Precision, recall, and F1-score for class-imbalanced datasets," Pattern Recognit. Lett., vol. 157, pp. 65–71, Mar. 2022.

6. JournalArticles:

[10] Y. Ding, W. Kang, J. Feng, B. Peng, and A. Yang, "Credit card fraud detection based on improved Variational Autoencoder Generative Adversarial Network," IEEE Access, vol. 11, pp. 1–12, 2023,

[11] H. Liu, M. C. Zhou, and Q. Liu, "An embedded featureselectionmethodforimbalanceddataclassification," IEEE/CAA J. Autom. Sinica, vol. 6, no.3, pp. 703–715, 2019.

[12] G.Haixiangetal.,"Learningfromclass-imbalanced data: Review of methods and applications," Expert Syst. Appl., vol. 73, pp. 220–239, 2017.

[13] A. Bahnsen et al., "Feature engineering strategiesfor credit card fraud detection," Expert Syst. Appl., vol. 51, pp. 134–142, 2016.

[14] F. Carcillo et al., "Combining unsupervised and supervised learning in credit card fraud detection," Inf. Sci., vol. 557, pp. 317–331,2021.