International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Generative AI in Real-Time E-Commerce Fraud Detection: A Comparative and Ethical Analysis

Bapathu Sai Prathap Reddy[1], Tammisetty Anusha[2], Siddhu Harshini[3], Gopireddy Venkata Ravindra Reddy[4], Dr. Koppula Chinabusi[5]

[1,2,3,4]Final year students, Tirumala institute of technology and sciences, Narasaraopet.
[5]Professor, HOD, Department of Computer Science and Engineering, Tirumala Institute of Technology and Sciences, Narasaraopet, Andhra Pradesh, India - 522601.

Abstract
This paper explores the application of generative artificial intelligence models, including GANs, VAEs, and a combined GAN-VAE approach, for real-time fraud detection in e-commerce systems. GANs are highly effective at generating synthetic fraudulent transaction data, whereas VAEs excel at identifying intricate fraud patterns. The proposed hybrid model integrates the advantages of both techniques, resulting in improved accuracy and adaptability. Furthermore, the study addresses ethical considerations such as data privacy and potential bias, highlighting the future potential of generative AI to enhance fraud detection frameworks in e-commerce.

Keywords: Generative AI, Generative Adversarial Networks, Variational Autoencoders, Machine Learning, E-Commerce Fraud.

1. INTRODUCTION

The rapid expansion of e-commerce has significantly reshaped both business operations and consumer behavior, offering unmatched convenience, variety, and accessibility in shopping. However, this growth has also led to an increase in risks related to online transactions, resulting in a surge of e-commerce fraud cases. Common types of fraud include identity theft, payment-related fraud, and account takeovers, all of which pose serious financial and reputational risks to both consumers and businesses. Traditional rule-based detection systems, which rely on fixed parameters, often struggle to keep up with the constantly evolving techniques used by malicious actors to exploit system weaknesses [1].

Generative artificial intelligence has emerged as a promising solution to address these challenges. Techniques based on Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) are capable of producing highly realistic synthetic data. This capability enhances anomaly detection processes and enables more accurate, real-time identification of fraudulent activities [2].

Furthermore, this paper examines the ethical implications of using generative AI in fraud detection, including concerns related to data privacy, algorithmic bias, and the balance between maintaining security and ensuring a smooth user experience. While highlighting both the strengths and limitations of generative AI, this study emphasizes the need for its responsible and ethical implementation within the e-commerce industry [3].

2. ORGANIZATION OF THE PAPER

The remainder of the document is organized as follows:

A. Section III: Literature Review

This section discusses how the generative AI models GANs and VAE improve the detection power of fraud.

B. Section IV: Methodology

This chapter uses generative AI models to realize real-time fraud detection in e-commerce. This has been focused on the model development stage, classifier integration, and also ethical considerations with comparative analysis to compare performance and practical recommendations.

C. Section V: Data Collection

Information collected, measured, and computed from different sources serves as an aid in gaining some insight or decision making based on data collection.

D. Section VI: Results

This chapter presents results of experiments, including some in-depth comparisons of the performance of real-time generative AI models in discovering fraudulent patterns.

E. Section VII: Discussion

It will briefly consider the implications of the results especially: How fraud patterns are captured by the gener*-ative

AI models, Pros and cons of each model. Ethical considerations encompass bias mitigation and questions of data privacy.

F. Section VIII: Conclusion

This study presents key findings in the conclusions related to its objectives that generative AI models hold tremendous potential to advance fraud detection frameworks and bring improved security to e-commerce.

G. Section IX: Future Work

This section suggests directions for future research: dynamic graphs for real-time fraud detection, multimodal data integration for enhanced detection capabilities, and further exploration of ethical implications and guidelines for responsible AI use.

H. Section X: References

This section provides an all-inclusive list of references used in the paper and relevant works concerning generative AI, fraud detection, and issues of ethics in AI usage.

3. LITERATURE REVIEW

The use of generative artificial intelligence in e-commerce fraud prevention is a rapidly growing area of interest, showing strong potential to improve how fraudulent activities are identified and managed.

A. Generative Model in Fraud Detection

Generative Adversarial Networks (GANs) are designed to produce synthetic data that captures complex fraud patterns. Although they have certain limitations, they are widely used to generate realistic datasets for training machine learning models. This helps enhance fraud detection systems by increasing the availability of data that closely resembles real-world scenarios [4].

On the other hand, Variational Autoencoders (VAEs) operate differently from GANs. While effective in modeling data distributions, they may lead to higher error rates and false positives, especially when applied to large and diverse datasets [5].

B. Hybrid Generative Models

Hybrid approaches that combine GANs and VAEs are gaining attention due to their balanced performance. These models leverage GANs for generating realistic data and VAEs for identifying anomalies, resulting in a more flexible and efficient fraud detection system.

C. Real-Time Fraud Detection in E-commerce

Real-time fraud detection is essential in e-commerce, as fraudulent actions can quickly impact a large number of users. GANs have proven useful in identifying known fraud patterns, but their computational complexity can make real-time implementation challenging.

D. Ethical and Privacy Concerns

The application of generative AI in financial fraud detection also raises ethical issues. Bias in training data can lead to unfair targeting of specific ethnic or geographic groups, emphasizing the need for fairness and inclusivity in system design. Additionally, protecting user data privacy remains a critical concern when deploying such technologies [6].

E. Adaptation to Evolving Fraud Tactics

Effective fraud detection systems must continuously adapt to new and evolving fraud techniques. Static models tend to become outdated over time, highlighting the importance of dynamic systems that learn and update themselves using new incoming data.

4. METHODOLOGY

This study evaluates the effectiveness of generative AI models for real-time fraud detection in e-commerce, while also considering ethical aspects and practical usability.

A. Data Acquisition and Preprocessing

A dataset of e-commerce transactions is prepared, containing both genuine and fraudulent records to support model training and evaluation.

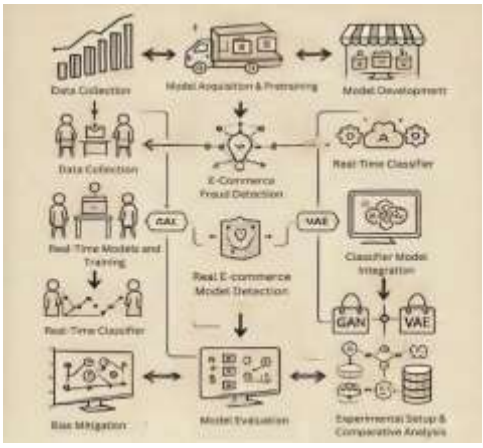

FIG. 1 METHODOLOGY

B. Model Development and Training

- **GAN Model:** Trained to generate realistic synthetic fraudulent transactions that resemble real fraud patterns.
- **VAE Model:** Learns data representations to identify subtle anomalies that may indicate fraud.
- **Hybrid Model:** Combines GANs for data generation and VAEs for anomaly detection to leverage both strengths.

- **Training Process:** GANs focus on matching real and generated data distributions, while VAEs aim to accurately reconstruct normal patterns and highlight irregularities .

C. Classifier Integration

A machine learning classifier (e.g., Random Forest or XGBoost) is integrated into a real-time pipeline that handles data input, prediction, logging, and latency optimization.

D. Ethical Considerations

- **Bias Control:** Models are evaluated and adjusted to reduce bias.
- **Privacy Protection:** Techniques like differential privacy are applied to secure sensitive data.
- **User Impact:** System performance is assessed based on false positives and user experience .

E. Experimental Setup and Analysis

GAN, VAE, and hybrid models are compared using metrics such as accuracy, precision, recall, F1-score, AUC, and processing time. Real-time testing and statistical analysis are used to identify the most effective and efficient approach for fraud detection [1.

5. DATA COLLECTION

To collect data for this model, total transaction data should be collected that contains actual and invalid occurrences of ecommerce. The detailed collection and structuring of the data are given below:

A. Data Sources

Transactional Records: This may include the details of transaction ID, customer demographics, product information, and timestamps related to any e-commerce platforms. 2) Fraud Databases: Fraudulent case data from previous cases ,which include transactions flagged as fraud.

Device and Network Data: These may include IP addresses, device information like browsers and operating systems, and network-related details [12]

Payment Method Data: Transaction methods used, with flags for high-risk transaction types such as using a specific payment processor or for foreign countries..

B. Data Preprocessing

Data Cleaning and Transformation: Identify missing or inconsistent values; normalize continuous variables; use onehot encoding for categorical variables.

Anonymization: Personally identifiable information including customer IDs, IP addresses, and geographic data should be anonymized to abide by privacy regulations. 3) Balancing with Synthetic Data: Generate synthetic fraudulent cases with the help of Generative Adversarial Networks to balance the dataset and improve model training performance.

TABLE I. ORIGINAL AND BALANCED DATASET

Transaction Type	Original Dataset	Balanced Dataset(After Synthetic Data Generation)
Fraudulent	5000	95000
Non-Fraudulent	95000	95000
Total	100000	190000

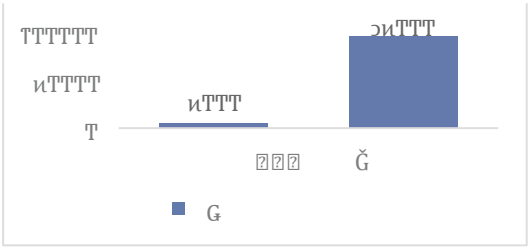

FIG. 2 COMPARISON OF FRAUDULENT AND NON- FRAUDLENT TRANSACTIONS

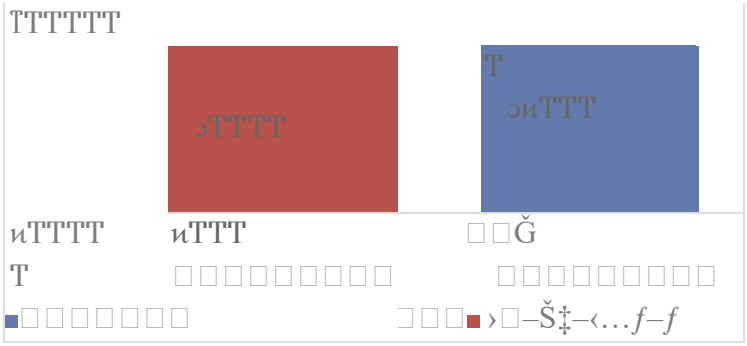

FIG. 3 COMPARISON OF ORIGINAL AND SYNTHETIC DATA

C. Data Augmentation

Synthetic Data Generation: GANs produce diverse fraud data for improving model training.
Anomaly Detection Support: VAEs find mild patterns in improving fraud detection.

6. RESULT

The performance of each model is evaluated using key metrics such as accuracy, precision, recall, F1-score, Area Under the Curve (AUC), and latency. The results highlight the strengths and limitations of each model in detecting fraudulent activities.

A. Model Performance Comparison

The evaluation of model performance is based on the following measures:

1. **True Positives (TP):** Fraudulent transactions that are correctly detected as fraud.
2. **False Positives (FP):** Legitimate transactions that are incorrectly flagged as fraud.
3. **False Negatives (FN):** Fraudulent transactions that are mistakenly classified as non-fraud.
4. **True Negatives (TN):** Genuine transactions that are accurately identified as non-fraud.
5. **Data Preprocessing Time (DPT):** The time required to clean, transform, and prepare raw data for model input.
6. **Model Inference Time (MIT):** The time taken by the model to generate predictions from the processed data.
7. **Post-processing Time (PPT):** The time needed to convert the model's output into a usable and interpretable format.

The feature capabilities of each model are measured by the aforementioned parameters consisting of accuracy, precision, recall, F1-score, Area Under the Curve (AUC), and latency. The results focus on each of the strengths and weaknesses of the models in fraud detection.

A. Model Performance Comparison

The following formulas are used to compare the models:

Precision (P): The proportion of fraudulent transactions that are actually fraudulent is known as precision (P).

$$P = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}} \quad (1)$$

Recall (R): The proportion of actual fraudulent transactions that were correctly classified as such.

$$R = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}} \quad (2)$$

F1-Score (F1): The harmonic mean of Precision and Recall, giving a balance between the two.

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (3)$$

Accuracy (A): This is an indication of the overall accuracy in classifying fraud and legit transactions by the model.

$$\text{Accuracy (A)} = \frac{\text{True Positives} + \text{True Negatives}}{\text{True Positives} + \text{True Negatives} + \text{False Positives} + \text{False Negatives}} \quad (4)$$

Area Under The Curve (AUC): It is derived from the ROC
(Receiver Operating Characteristic) Curve, which plots the True Positive Rate (TPR) against the False Positive Rate (FPR) at various threshold levels.
True Positive Rate (TPR), also known as Recall-

$$TPR = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}} \quad (5.1)$$

False Positive Rate (FPR)-

$$FPR = \frac{\text{False Positives}}{\text{False Positives} + \text{True Negatives}} \quad (5.2)$$

Latency (L): Average time taken by the model for processing each transaction.

$$L = DPT + MIT + PPT \quad (6)$$

TABLE 2. MODEL PERFORMANCE COMPARISON

Metric	GAN Model	VAE Model	Hybrid GAN-VAE Model
Precision	91%	85%	92%
Recall	84%	88%	90%
F1-Score	87.5%	86.5%	91%
Accuracy	89%	86%	92%

AUC	0.88	0.88	0.92
Latency(seconds)	2-3 sec	1.5 sec	3.5 sec

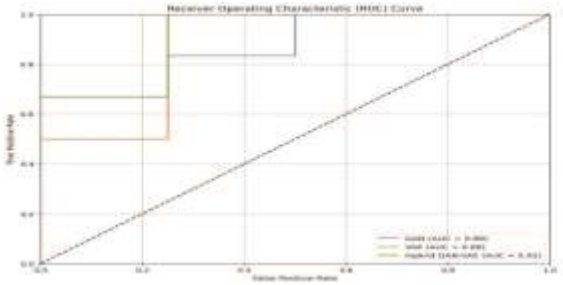

FIG. 4 RECEIVER OPERATING CHARACTERISTIC CURVE

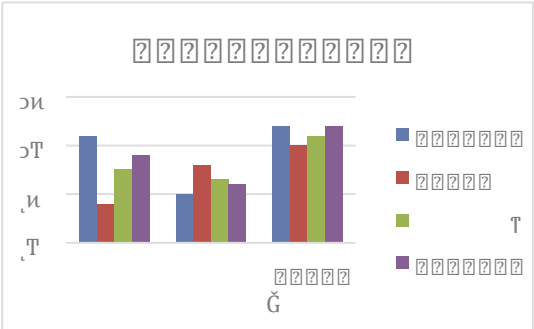

FIG. 5 COMPARISON OF GAN, VAE AND HYBRID GAN-VAE

B. Real-Time Performance and Scalability The scalability of GANs is hindered when analyzing high dimensional data. Notwithstanding this, improvements in performance are possible for moderate volumes with the help of techniques such as batched processing and pruning.

Let's assume processing of a batch of 10,000 transactions is performed and each model is executed:

TABLE 3. TRANSACTION METRICS

Metric	GAN Model	VAE Model	Hybrid GAN-VAE Model
Transactions Processed/Batch	10,000	10,000	10,000
Average Latency/Transaction	2.5 sec	1.5 sec	3.5 sec
Total Batch Processing Time	25,000 sec	15,000 sec	35,000 sec

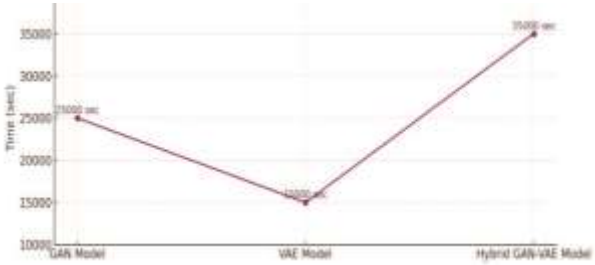

FIG. 6 AVERAGE TRANSACTION TIMES OF MODELS

Let's conclude testing with three fraud types: Account Takeover, Payment Fraud, and Synthetic Identity Fraud:

TABLE 4. DISTRIBUTION OF RECALL PERCENTAGE

Fraud Type	GAN Recall (%)	VAE Recall (%)	Hybrid Recall (%)
Account Takeover	82	88	89
Payment Fraud	85	90	91
Synthetic Identity Fraud	78	86	88

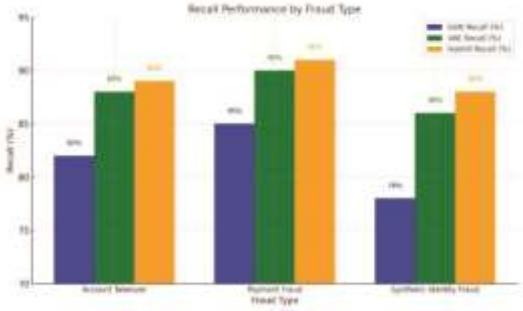

FIG. 7 RECALL PERFORMANCE BY FRAUD TYPE

C. Ethical Considerations and Bias Analysis

While the same should be conducted to Regional Demographic Bias, almost all subjects are biased within regions; A, B, and C.

TABLE 5. DISTRIBUTION OF FALSE POSITIVE RATE

Demographic	GAN False Positive Rate (%)	VAE False Positive Rate (%)	Hybrid False Positive Rate (%)
Region A	10.5	12.1	9.8
Region B	7.8	8.9	7.2

Region C	11.2	13.4	10.1

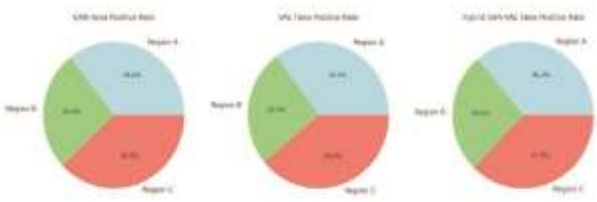

FIG. 8 COMPARISON OF FALSE POSITIVE RATES

7. DISCUSSION

This research shows that generative AI models such as GANs, VAEs and their hybrid version are already succeeding in real time fraud detection for e commerce platforms. Their advancement in the feature transformation process combined with the diverse types of fraud demonstrated by these models is a strong pointer to the massive possibilities of the generative AIs in the existing fraud detection systems. Even if GANs performed remarkably well in creating synthetic data following the most common fraud schemed patterns, they still have a few disadvantages. The major drawback was the occurrence of false positives in instances where newer and more intricate frauds that were not incorporated into the training data model are introduced. This therefore indicates that refinements for the model are necessary to enhance the accuracy whilst reducing the training bias so that only relevant transaction characteristics are emphasized. In summary, this study points that although generative AI has a lot of promise, there is more work to be done in tweaking these models to a practical stance that weighs accuracy, speed and ethics. The hybrid model, in particular, is a good illustration of how such a strategy can alleviate a challenge posed by the use of any one generative model however, adoption and speed remain the most viable opportunities to consider.

8. CONCLUSION

This study focuses on strengthening the e-commerce ecosystem and improving its ability to withstand fraudulent activities by using generative AI models such as GANs, VAEs, and their hybrid combinations. Each of these models plays a valuable role in detecting known fraud patterns and enhancing overall system performance.

Although generative AI represents a significant advancement in combating e-commerce fraud, continuous improvement and integration of different models are necessary to fully utilize their potential in real-time and large-scale environments. The development of efficient real-time monitoring systems is essential to detect and prevent new and emerging types of fraud. At the same time, emphasizing ethical considerations helps build trust and wider acceptance of these technologies.

Overall, generative AI has the potential to transform online fraud detection by enabling e-commerce platforms to implement more adaptive, reliable, and comprehensive security measures against fraudulent activities.

9. FUTURE WORK

The use of generative artificial intelligence in e-commerce fraud detection offers significant potential for research and practical application in the near future. One of the key challenges is scaling hybrid GAN-VAE frameworks for real-time implementation. As transaction volumes grow, performance can be

maintained by improving processing efficiency through approaches such as model compression and distributed computing.

Bias and fairness remain critical concerns. Future research should prioritize robust bias detection and mitigation mechanisms, ensuring consistent fraud detection across different population groups. Using diverse datasets during training and incorporating privacy-preserving methods like federated learning and differential privacy will help maintain regulatory compliance and build user trust.

The long-term effectiveness and adaptability of these models can be assessed through longitudinal studies. Involving stakeholders—including e-commerce platforms and users—in the design and deployment of solutions ensures they are practical, user-focused, and industry-relevant. Additionally, integrating emerging technologies like reinforcement learning and explainable AI into generative models can enhance decision-making and provide transparency, further improving trust in fraud detection systems.

Overall, the development of generative AI models holds the promise of transforming e-commerce fraud detection, improving real-time security, and strengthening the integrity of online transactions.

REFERENCE

1. A. Jain and S. Raghavan, "Fairness in machine learning: A survey," ACM Computing Surveys, vol. 54, no. 6, pp. 1–35, 2021.
2. S. Huang and W. Chen, "Ethical considerations in AI and machine learning: The role of data governance," Artificial Intelligence Review, vol. 53, no. 3, pp. 1647–1670, 2020.
3. H. Q. Nguyen and T. T. Nguyen, "Detecting fraud in e-commerce: A review of recent advances and future directions," Expert Systems with Applications, vol. 177, p. 114914, 2021.
4. S. Tyagi, D. Kumar, and S. Kumar, "Understanding the nitty-gritties of software reliability and its testing procedures: A different approach," Journal of Information and Optimization Sciences, vol. 38, no. 6, pp. 971–988, 2017.
5. Y. Li and J. Wang, "Generative models for e-commerce fraud detection: A comparative study," Journal of Computer Science and Technology, vol. 36, no. 3, pp. 505–520, 2021.
6. J. Zhou and Q. Chen, "E-commerce fraud detection using machine learning: A review," Journal of Systems and Software, vol. 151, pp. 134–150, 2019.
7. R. Binns, "Fairness in machine learning: Lessons from political philosophy," in Proc. 2018 Conf. Fairness, Accountability, and Transparency (FAT), 2018.
8. J. Choi and K. J. Kim, "The use of generative adversarial networks in e-commerce fraud detection: A systematic review," Journal of Business Research, vol. 139, pp. 479–489, 2022.
9. I. H. Sarker and D. O'Sullivan, "Exploring the use of generative models for fraud detection: Current trends and future directions," Journal of Risk and Financial Management, vol. 16, no. 3, p. 101, 2023.
10. S. Tyagi, D. Kumar, and S. Kumar, "Reliability based solution to the decision making dilemma in a software environment," Journal of Statistics and Management Systems, vol. 22, pp. 627–634, 2019, doi: 10.1080/09720510.2019.1611226.
11. H. Zhang and L. Wang, "Ethical implications of generative AI in finance: A comprehensive review," Artificial Intelligence Review, vol. 56, no. 2, pp. 1363–1385, 2023.
12. Z. Wang and J. Zhang, "Federated learning for privacy-preserving AI: A comprehensive survey," IEEE Transactions on Neural Networks and Learning Systems, vol. 33, no. 6, pp. 2348–2363, 2022.