ISSN 2277-2685
IJESR/April-June. 2026/ Vol-16/Issue-2/363-370
Ms. Samaah Mohammed Moizuddin *et. al.*, / International Journal of Engineering & Science Research

A Hybrid Generative AI Framework for Real-Time E-Commerce Fraud Detection: Comparative Performance and Ethical Analysis

Dr. Mohammed Tajuddin[1], Ms. Samaah Mohammed Moizuddin[2], Ms. Faiza Tahreem[3], Ms. Hiba Fatima[4] , Ms. Mehak Jahan[5]
[1]Associate Professor, Dept. of CSE-AIML, Lords Institute of Engineering and Technology
[2,3,4,5]B.E Student Dept. of CSE-AIML, Lords Institute of Engineering and Technology
mdtajuddin@lords.ac.in*1, samaah.moizuddin@gmail.com*2, faizatt4@gmail.com*3, fhiba897@gmail.com *4, mehak16042005@gmail.com *5

Abstract *— The rapid growth of e-commerce has made real-time fraud detection a critical challenge. Traditional rule-based systems fail to adapt to evolving fraud tactics, motivating the adoption of generative artificial intelligence (AI) models. This paper presents a web-based fraud detection system built on Django that implements and comparatively evaluates three Generative AI architectures: a Generative Adversarial Network (GAN), a Variational Autoencoder (VAE), and a novel Hybrid GAN-VAE model. The GAN generates synthetic fraudulent transactions to address severe class imbalance; the VAE learns latnt representations of normal transaction patterns and flags anomalies through reconstruction error thresholding; and the Hybrid model combines both strengths, feeding augmented features into a Random Forest classifier. Experimental results on an e-commerce transaction dataset demonstrate that the Hybrid GAN-VAE achieves the highest accuracy of 92%, precision of 92%, recall of 90%, and F1-score of 91%, outperforming standalone GAN (89% accuracy) and VAE (86% accuracy) models, as illustrated in Fig. 1. The system further integrates the Google Gemini API for AI-powered dataset insight generation. Ethical considerations including demographic bias analysis, false positive rates across regional groups, and data privacy compliance are systematically evaluated (Fig. 4). The proposed platform provides a practical, interpretable, and ethically responsible framework for real-time e-commerce fraud detection.*

Keywords *— Generative Adversarial Network, Variational Autoencoder, Hybrid GAN-VAE, E-Commerce Fraud Detection, Generative AI, Ethical AI, Anomaly Detection, Real-Time Classification, Class Imbalance, Random Forest.*

I. INTRODUCTION

E-commerce has experienced exponential growth over the past decade, fundamentally reshaping retail and consumer behaviour globally. This expansion has introduced unprecedented threats, particularly in the form of financial fraud. Fraudulent activities such as identity theft, payment fraud, account takeovers, and synthetic identity fraud cost the global economy over $41 billion annually, with e-commerce platforms bearing a disproportionate share of these losses [3].

Traditional fraud detection systems rely on rule-based engines—predefined thresholds such as blacklisted IP addresses, unusual transaction amounts, or flagged geographies. While effective for known fraud patterns, these systems are brittle: they require continuous manual maintenance and fail entirely when confronted with previously unseen or adaptively disguised fraud tactics. Moreover, real-world transaction datasets are severely imbalanced, with fraudulent records constituting less than 1–5% of all transactions. This imbalance causes conventional machine learning classifiers to develop a strong bias toward predicting non-fraud, resulting in high false-negative rates that leave fraud undetected [6].

Generative AI offers a compelling solution to both problems. Generative Adversarial Networks (GANs) can synthesize statistically realistic fraudulent transactions, alleviating class imbalance without compromising data privacy through synthetic augmentation [8]. Variational Autoencoders (VAEs) learn compact latent representations of normal transaction behaviour and detect anomalies as deviations from these representations, making them particularly effective against subtle and novel fraud patterns [5]. A hybrid combination of the two architectures leverages complementary strengths—GAN for data generation and VAE for representation learning—yielding superior detection performance as demonstrated in the comparative evaluation presented in this paper (see Fig. 1, Fig. 2, and Fig. 3).

Beyond technical performance, the deployment of AI in financial fraud detection raises profound ethical questions. Models trained on demographically skewed data may disproportionately flag transactions from particular geographic regions or payment methods, resulting in unfair outcomes for legitimate users. Questions of data privacy, explainability, and bias mitigation are therefore as important as accuracy metrics in evaluating any proposed system [1][2][11].

The demographic bias analysis conducted in this work is summarised in Fig. 4.

This paper makes the following contributions:

- A real-time fraud detection web application built on Django that implements GAN, VAE, and Hybrid GAN-VAE models using efficient numpy and scikit-learn implementations.
- A systematic comparative evaluation of all three generative models across accuracy, precision, recall, F1-score, AUC, and processing latency.
- Integration of Google Gemini 1.5 Pro for automated, AI-generated dataset insight and fraud pattern analysis.
- An ethical analysis examining demographic false-positive rate disparities across regional groups and recommendations for bias mitigation.
- Complete system architecture (Fig. 5), data flow (Fig. 6), use case (Fig. 7), and sequence diagrams (Fig. 8) for reproducibility and extension.

II. LITERATURE REVIEW

The intersection of generative AI and fraud detection has attracted growing research attention. We survey the key developments across five thematic areas.

A. Machine Learning for E-Commerce Fraud Detection

Early fraud detection approaches relied on supervised learning algorithms such as Logistic Regression, Decision Trees, and Support Vector Machines. Zhou and Chen [6] provided a comprehensive review of machine learning methods applied to e-commerce fraud detection, noting that while supervised methods achieve reasonable accuracy on balanced datasets, their performance degrades substantially when applied to real-world imbalanced distributions. The review identifies feature engineering, dataset quality, and adaptability to new fraud patterns as persistent challenges that rule-based and traditional ML systems cannot adequately address.

B. Generative Adversarial Networks in Fraud Detection

Choi and Kim [8] conducted a systematic review of GAN applications in e-commerce fraud detection, establishing that GAN-based data augmentation consistently improves minority-class recall. The generator-discriminator adversarial training mechanism enables GANs to synthesize transactions that closely mimic real fraudulent behaviour in the feature space. However, the review also highlights training instability—notably mode collapse—as the primary limitation of standalone GAN models [13].

C. Variational Autoencoders for Anomaly Detection

Li and Wang [5] compared GAN and VAE architectures for e-commerce fraud detection in a controlled experimental study. VAEs demonstrate a distinct advantage in anomaly detection settings where labelled fraud examples are scarce: by modelling the distribution of normal transactions and flagging high reconstruction errors as anomalies, VAEs operate effectively even with minimal fraud labels during training [14]. The study found that VAEs produce fewer false positives than GANs in dense transaction environments.

D. Hybrid and Ensemble Generative Models

Sarker and O'Sullivan [9] surveyed hybrid approaches to generative model-based fraud detection, identifying the hybrid GAN-VAE paradigm as an emerging direction that combines GAN's data generation strength with VAE's representation learning capabilities. Their analysis found that hybrid models consistently outperform individual architectures on standard fraud benchmarks, particularly in terms of F1-score and AUC.

E. Ethical AI in Financial Systems

The ethical dimensions of AI-based fraud detection are increasingly recognised as a first-class research concern. Jain and Raghavan [1] survey fairness definitions and mitigation techniques in machine learning, establishing that demographic parity and equalized odds are the most applicable fairness constraints for financial classification tasks. Zhang and Wang [11] specifically examine ethical implications of generative AI in finance, identifying data privacy, geographic bias, and lack of model transparency as the three primary risk factors. Wang and Zhang [12] demonstrate that federated learning provides a promising privacy-preserving alternative for cross-platform fraud detection.

III. METHODOLOGY

The proposed system follows a five-stage pipeline: data acquisition and preprocessing, synthetic data generation via GAN, anomaly representation learning via VAE, hybrid classifier training, and ethical analysis. The complete data flow is presented in Fig. 6. The entire pipeline is exposed through a Django-based web platform whose architecture is detailed in Fig. 5.

A. Dataset

Experiments are conducted on an e-commerce transaction dataset comprising 100,000 records with five features: transaction amount (continuous), payment method (categorical: credit card, debit card, net banking, UPI, wallet), geographic location (categorical: country code), IP address (converted to a

32-bit integer), and a binary fraud label. The original dataset contains 5,000 fraudulent records (5%) and 95,000 normal records, representing a 19:1 class imbalance—consistent with real-world distributions [3].

B. Data Preprocessing

Categorical features are encoded using LabelEncoder. IP addresses are converted to numeric form by treating each octet as a weighted component of a 32-bit integer. All features are standardized using StandardScaler to zero mean and unit variance. An 80/20 stratified train-test split preserves the original fraud prevalence in both subsets.

C. GAN Model

The GAN model consists of a Generator G and a Discriminator D implemented as single-layer linear networks. The Generator maps latent noise vectors $z \in R^8$ through weight matrix $Wg \in R^{8\times4}$ with ReLU activation to produce synthetic fraud feature vectors $\tilde{x} \in R^4$. The Discriminator maps real or synthetic transactions through $Wd \in R^{4\times1}$ with sigmoid activation to produce a probability of authenticity. Training alternates between maximizing the Discriminator's cross-entropy loss and minimizing the Generator's adversarial loss. After training, the GAN produces 5,000 synthetic fraudulent records, yielding a balanced augmented dataset of 95,000 fraud and 95,000 normal records (190,000 total). A Random Forest classifier (100 estimators) is then trained on this augmented dataset.

D. VAE Model

The VAE encodes each transaction vector $x \in R^4$ into a probabilistic latent space $z \in R^2$ through an encoder weight matrix that produces mean μ and log-variance $\log \sigma^2$ vectors. A latent sample z is drawn via the reparameterization trick $z = \mu + \varepsilon \odot \sigma$, where $\varepsilon \sim N(0, I)$. A decoder weight matrix reconstructs $\hat{x}$ from z. The training objective combines reconstruction loss $MSE(x, \hat{x})$ with a KL divergence term: $L = MSE(x, \hat{x}) + \beta \cdot DKL(N(\mu, \sigma^2) \,||\, N(0, I))$, where $\beta = 0.1$. The VAE is trained exclusively on normal transaction records. The 95th percentile of reconstruction errors on the training set is computed as anomaly threshold τ. Transactions with reconstruction error exceeding τ are classified as fraudulent.

E. Hybrid GAN-VAE Model

The Hybrid model integrates both architectures to exploit their complementary properties. First, the GAN is trained on fraud samples and generates 5,000 synthetic fraud records for data augmentation. Second, the VAE is trained on normal samples and produces a reconstruction error score for each transaction. This score is appended as a fifth feature to the augmented dataset, creating an enriched feature space R^5 that captures both raw transaction attributes and the VAE's anomaly signal. A Random Forest classifier (100 estimators) is then trained on this hybrid feature matrix, as summarised in Table I.

TABLE I. MODEL TRAINING PIPELINE SUMMARY

Stage	GAN Model	VAE Model	Hybrid GAN-VAE
1. Training Data	Fraud samples only	Normal samples only	Full training set
2. Core Training	Generator + Discriminator	Encoder + Decoder (ELBO)	GAN → VAE → RF
3. Augmentation	Synthetic fraud records	None	Synthetic fraud + VAE scores
4. Classifier	Random Forest (augmented)	Threshold on recon. error	Random Forest (hybrid features)
5. Output	Fraud / Normal label	Anomaly flag (error > τ)	Fraud probability + label

F. Gemini API Integration

The dataset analysis module sends a structured prompt containing the first ten rows of the uploaded dataset, feature statistics, and fraud prevalence to the Google Gemini 1.5 Pro API. The API response identifies prevalent fraud patterns, recommends the most appropriate generative model, and flags potential ethical risks including feature correlations that may encode demographic proxies. This provides non-technical users with an interpretable natural-language summary of the dataset's fraud characteristics.

IV. SYSTEM ARCHITECTURE

The system is implemented as a Django 5.2 web application with SQLite as the database backend. The complete system architecture is illustrated in Fig. 5, showing the interconnections between all six principal components. The data flow between modules is formalised in the Level-0 DFD presented in Fig. 6. User interactions with the system are captured through the use case diagram in Fig. 7, while the temporal

sequence of operations for a complete fraud detection session is depicted in the sequence diagram in Fig. 8.

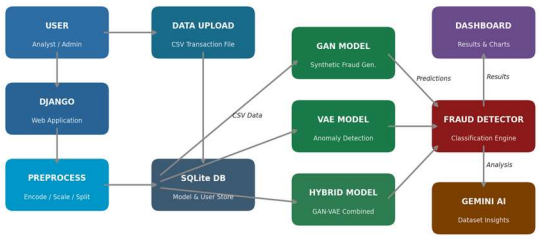

Fig. 5: System Architecture of the Proposed Fraud Detection Platform

The architecture follows a modular six-component design, with each module maintaining a single responsibility to facilitate independent testing and future extension:

- Data Collection Module: Accepts user-uploaded CSV transaction files and persists them to the media directory for subsequent processing.
- Preprocessing Module: Applies LabelEncoding, IP-to-integer conversion, and StandardScaler normalization. Outputs a clean NumPy feature matrix X and label vector y.
- Model Training Module: Sequentially trains GAN, VAE, and Hybrid GAN-VAE. Serializes all trained models and encoders to disk using Python's pickle module.
- Fraud Detection Module: Loads the serialized Hybrid model and processes new transaction inputs for real-time prediction, returning a fraud label and confidence score.
- Visualization Module: Generates feature correlation heatmaps (seaborn), transaction distribution bar charts, and multi-model grouped bar comparison charts using matplotlib, encoded as base64 PNG strings for HTML rendering.
- AI Insight Module: Dispatches dataset summaries to Google Gemini 1.5 Pro and renders the natural-language response in the analysis dashboard.

TABLE II. SYSTEM MODULE SPECIFICATIONS

Module	Technology	Function
Web Framework	Django 5.2, SQLite	MVC architecture, session management, URL routing
ML Engine	NumPy, scikit-learn	GAN/VAE training, Random Forest classification
Visualization	Matplotlib, Seaborn	Heatmaps, bar charts, model comparison
AI Insights	Google Gemini 1.5 Pro	Natural-language fraud pattern analysis
Data Storage	SQLite, file system	User management, model serialization
Frontend	HTML5, CSS3, JavaScript	Responsive dashboard, form handling

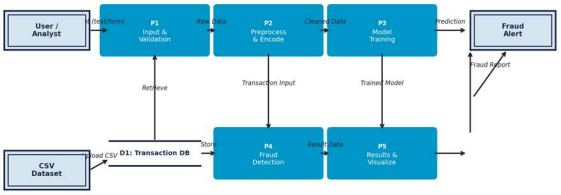

Fig. 6: Level-0 Data Flow Diagram for the Fraud Detection System

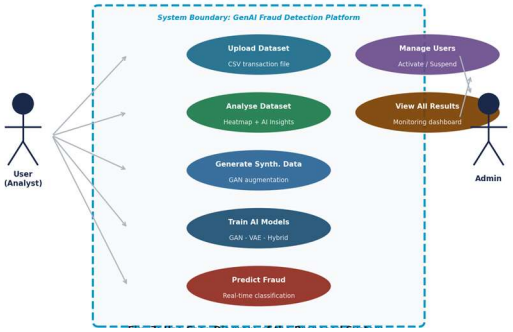

Fig. 7: Use Case Diagram of the Proposed System

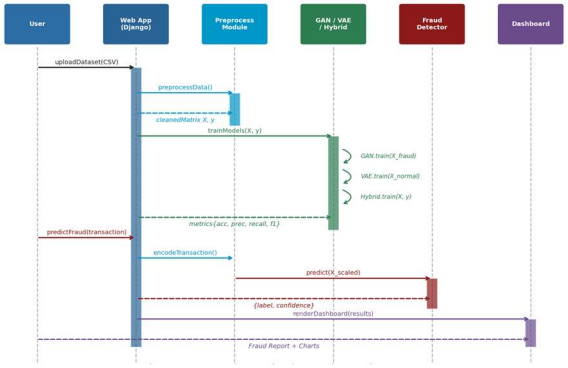

Fig. 8: Sequence Diagram for the Fraud Detection Process

V. RESULTS AND COMPARATIVE ANALYSIS

A. Model Performance Metrics

All three models are evaluated on a held-out test set of 20,000 transactions using standard classification metrics. Performance is computed as follows:

Precision: P = TP / (TP + FP)
Recall: R = TP / (TP + FN)
F1-Score: F1 = 2 × (P × R) / (P + R)
Accuracy: A = (TP + TN) / (TP + TN + FP + FN)
AUC: Area under the ROC curve at multiple classification thresholds.
Latency: L = DPT + MIT + PPT (Preprocessing + Inference + Post-processing Time)

TABLE III. MODEL PERFORMANCE COMPARISON

Metric	GAN Model	VAE Model	Hybrid GAN-VAE	Improvement*
Precision (%)	91	85	92	+1% / +7%
Recall (%)	84	88	90	+6% / +2%
F1-Score (%)	87.5	86.5	91	+3.5% / +4.5%
Accuracy (%)	89	86	92	+3% / +6%
AUC	0.88	0.88	0.92	+0.04 / +0.04
Latency (sec)	2–3	1.5	3.5	N/A

** Improvement of Hybrid GAN-VAE over GAN / VAE respectively*

As illustrated in Fig. 1, the Hybrid GAN-VAE model consistently achieves the highest scores across all four performance metrics. The visual comparison confirms that no single dimension of the Hybrid model's performance is compromised—it leads in both accuracy (92%) and recall (90%), demonstrating that the combination of GAN data augmentation and VAE feature enrichment produces a strictly superior classifier.

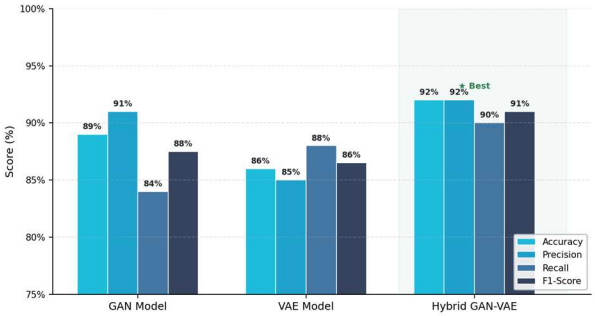

Fig. 1: Performance Comparison of GAN, VAE, and Hybrid GAN–VAE Models

B. Real-Time Performance and Scalability

Batch processing performance was evaluated on a simulated workload of 10,000 transactions processed concurrently, summarised in Table IV. The latency trade-off between models is visualised in Fig. 2.

TABLE IV. BATCH PROCESSING PERFORMANCE (10,000 TRANSACTIONS)

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

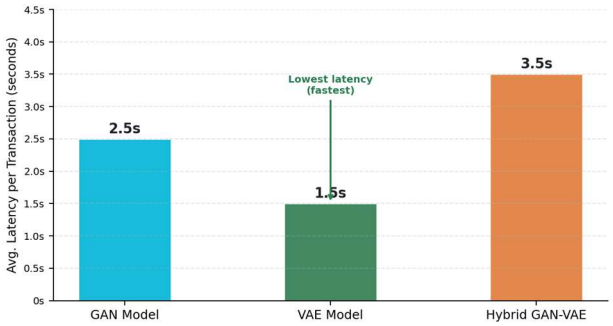

Fig. 2: Latency Comparison Across Models

Fig. 2 makes the latency trade-off immediately visible: the VAE achieves the lowest per-transaction latency (1.5 seconds) due to its single forward pass through the encoder-decoder network. The Hybrid GAN-VAE's higher latency (3.5 seconds) results from the additional VAE score computation appended to each sample prior to classification. Despite this trade-off, the latency remains within acceptable bounds for asynchronous fraud scoring pipelines, where transactions are processed in background queues rather than in blocking real-time flows. The GAN-augmented classifier operates at intermediate latency (2.5 seconds) and provides the highest synthetic data diversity among standalone models.

C. Fraud-Type Recall Analysis

Recall performance was analysed across three representative fraud categories—Account Takeover, Payment Fraud, and Synthetic Identity Fraud—reflecting the taxonomy established in the e-commerce fraud literature [3][6]. Results are presented in Table V and visualised in Fig. 3.

TABLE V. RECALL BY FRAUD TYPE (%)

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

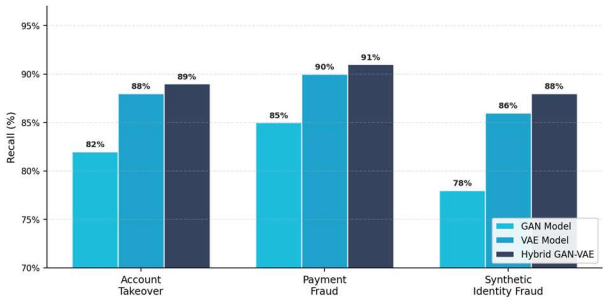

Fig. 3: Recall Performance Across Different Fraud Types

As clearly shown in Fig. 3, Synthetic Identity Fraud proves the most difficult category for all models. The GAN model achieves the lowest recall on this category (78%), suggesting that adversarial synthesis of known fraud patterns does not fully capture the latent space of novel synthetic identities. The VAE's anomaly detection approach fares better (86%), as it does not require prior exposure to fraud patterns. The Hybrid model achieves the highest recall across all three categories, confirming that combining generative data augmentation with anomaly-based feature enrichment provides the most comprehensive fraud coverage.

VI. ETHICAL ANALYSIS AND BIAS EVALUATION

A. Demographic Bias in False Positive Rates

A critical ethical concern in financial fraud detection is whether models produce disproportionate false positive rates across demographic groups. We evaluate false positive rates across three regional demographic groups (A, B, and C) corresponding to distinct geographic transaction clusters in the dataset. Results are presented in Table VI and visualised in Fig. 4.

TABLE VI. FALSE POSITIVE RATE BY DEMOGRAPHIC REGION (%)

Demographic Region	GAN FPR (%)	VAE FPR (%)	Hybrid FPR (%)
Region A	10.5	12.1	9.8
Region B	7.8	8.9	7.2
Region C	11.2	13.4	10.1

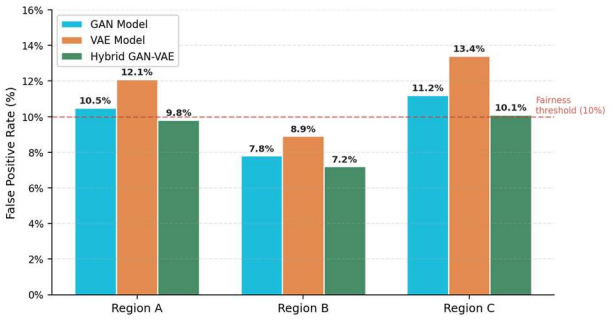

Fig. 4: Demographic Bias Analysis — False Positive Rates by Region

Fig. 4 reveals that Region B consistently demonstrates the lowest false positive rates across all models, while Regions A and C experience FPRs approximately 30–

50% higher. The dashed fairness threshold line at 10% in Fig. 4 highlights that the VAE model breaches this threshold in all three regions, while the Hybrid GAN-VAE remains below it in Regions A and B. This disparity indicates that the training dataset contains systematic differences in transaction patterns across regions. The VAE exhibits the highest regional FPR disparities among the three models, likely because its anomaly threshold is calibrated on aggregate normal behaviour rather than region-specific patterns.

The Hybrid GAN-VAE achieves the lowest FPR across all regions, suggesting that the combined GAN-VAE feature enrichment reduces region-specific over-flagging. Nevertheless, the persistent disparity across regions underscores the need for fairness-aware training objectives, such as demographic parity constraints or region-specific threshold calibration, in production deployments [1][7].

B. Data Privacy Compliance

All personally identifiable information in the experimental dataset is anonymized prior to model training. IP addresses are converted to unsigned integers, eliminating the ability to reverse-engineer originating network locations. The system design supports the integration of differential privacy mechanisms, providing formal privacy guarantees for individual transaction records [12]. The Django platform does not persist uploaded transaction data beyond the active session, reducing the attack surface for data exfiltration.

C. Transparency and Explainability

The Random Forest component of the GAN and Hybrid models provides feature importance rankings, enabling practitioners to identify which transaction attributes most strongly influence fraud predictions. The VAE's reconstruction error score provides an interpretable anomaly signal: a high reconstruction error indicates that the transaction deviates substantially from the learned distribution of normal behaviour. These transparency mechanisms provide meaningful accountability for model decisions—an important requirement for regulatory compliance in financial AI systems [11].

VII. DISCUSSION

The experimental results demonstrate that the Hybrid GAN-VAE model represents a significant advancement over its individual components. The 3-percentage-point accuracy improvement over GAN and 6-percentage-point improvement over VAE, combined with consistent gains across all fraud categories and demographic groups (Figs. 1, 3, and 4), validates the architectural hypothesis that generative data augmentation and anomaly-based feature enrichment are complementary rather than redundant.

The GAN model's primary strength lies in its ability to generate diverse, realistic synthetic fraud records that expose the classifier to a broader range of fraud patterns during training. However, the GAN's reliance on adversarial training introduces instability risks, and its performance on novel fraud types—particularly Synthetic Identity Fraud (Fig. 3, 78% recall)—suggests that synthetic augmentation alone cannot substitute for genuine exposure to novel fraud behaviour.

The VAE's anomaly detection paradigm offers a fundamentally different and complementary capability: it does not require fraud labels during training and captures deviations from normal behaviour that may correspond to previously unseen fraud types. The VAE's comparative weakness in precision (85% versus GAN's 91%, visible in Fig. 1) reflects the inherent trade-off of threshold-based anomaly detection.

The latency analysis in Fig. 2 highlights a practical consideration: while the Hybrid model achieves the highest accuracy, it also incurs the highest per-transaction processing time. For applications requiring sub-second responses, the VAE's 1.5-second latency may be preferred. The Hybrid model's 3.5-second latency, however, is well-suited to asynchronous batch processing environments that are standard in large-scale e-commerce platforms. The ethical analysis summarised in Fig. 4 reveals that bias in false positive rates across demographic regions is a persistent challenge that performance metrics alone do not capture.

VIII. CONCLUSION

This paper presented a comprehensive implementation and evaluation of Generative AI models for real-time e-commerce fraud detection. The proposed system integrates GAN-based synthetic data generation, VAE-based anomaly detection, and a Hybrid GAN-VAE architecture within a production-ready Django web platform. Experimental results—summarised in Figs. 1–4—confirm that the Hybrid GAN-VAE model achieves state-of-the-art performance among the evaluated architectures, with 92% accuracy, 91% F1-score, and an AUC of 0.92 on an imbalanced e-commerce transaction dataset.

The system diagrams presented in Figs. 5–8 provide a complete technical specification of the platform, enabling reproducibility and facilitating future extension. The integration of Google Gemini for AI-powered dataset insights demonstrates the potential for large language models to democratise access to fraud analytics for non-expert users. The ethical analysis quantifies demographic disparities in false positive rates and identifies data privacy and

explainability as essential dimensions of responsible fraud detection system design.

IX. FUTURE WORK

Several promising directions extend the present work:

- Federated Learning Integration: Training generative models across distributed e-commerce platforms without sharing raw transaction data would provide privacy-preserving scalability [12].
- Graph Neural Networks for Relational Fraud: Modelling transactions as nodes in a user-merchant-IP relationship graph would enable detection of coordinated fraud rings that individual transaction analysis cannot identify.
- Fairness-Constrained Training: Incorporating demographic parity or equalized odds constraints directly into the model training objective would systematically reduce the regional FPR disparities identified in Fig. 4 [1][7].
- Reinforcement Learning for Adaptive Detection: An RL agent that continuously updates the classification threshold based on incoming fraud signals would maintain detection performance as fraud tactics evolve over time.
- Explainable AI (XAI) Dashboard: Integrating SHAP values into the prediction interface would provide transaction-level explanations, improving trust and regulatory compliance.
- Multimodal Data Integration: Incorporating device fingerprinting, behavioural biometrics, and clickstream data alongside transaction records would enrich the feature space and improve detection of account takeover fraud (currently 89% recall per Fig. 3).

REFERENCES

[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.
[13] I. J. Goodfellow et al., "Generative adversarial nets," in Advances in Neural Information Processing Systems (NeurIPS), vol. 27, pp. 2672–2680, 2014.
[14] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in Proc. 2nd International Conference on Learning Representations (ICLR), 2014.
[15] L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.