Explainable Graph Neural Networ Framework for Digital Transaction Frau Detection

Nalam chaitanya Gopinath[1], Mrs. Swathi Koganti[2], Ampapurapu Devendra Teja[3]
[1,3]*Department of Computer Science and Engineering Vignan's Foundation for Science, Technology and Research Guntur, India*
[2]*Assistant Professor, Department of Computer Science and Engineering Vignan's Foundation for Science, Technology and Research*
Guntur, India

***Abstract*—Detecting credit card fraud is a major financial security challenge. This is a result of shifting fraud patterns and extremely unbalanced transaction data. In order to enhance detection performance, this study presents a hybrid fraud detection framework that combines gradient boosting classification with graph-based representation learning. First, transaction features, including the Time attribute considered as a numerical variable, are preprocessed using standard scaling and Borderline-SMOTE to tackle class imbalance. A static k-nearest neighbor similarity graph is created from the scaled feature space. A Temporal Graph Attention Network (TGAT) architecture is used to learn distinct node embeddings from graph relationships, even without explicit temporal encoding. These embeddings are combined with the original transaction features and fed into an XGBoost classifier for final fraud prediction. Optimal decision thresholds are determined through precision-recall analysis.**

Additionally, SHAP-based explainable artificial intelligence is used to clarify global feature importance and local prediction behavior, which improves model transparency. Experimental evaluation with stratified cross-validation shows that the TGAT-XGBoost hybrid model performs strongly across Accuracy, Precision, Recall, F1-score, ROC-AUC, and PR-AUC metrics. This indicates its effectiveness for reliable and understandable credit card fraud detection.

***Index Terms*—XGBoost, Class Imbalance Handling, Borderline-SMOTE, Explainable Artificial Intelligence (XAI), Graph Representation Learning, Temporal Graph Attention Network (TGAT), Static Similarity Graph, Credit Card Fraud Detection, SHAP Interpretation, and Precision–Recall Optimization.**

I. INTRODUCTION

The rapid development of digital payment systems and online financial transactions has significantly raised the risk of credit card fraud, which has led to significant financial losses for both consumers and financial institutions. Because of the highly unbalanced nature of transaction datasets, changing fraud patterns, and the requirement for precise real-time decision-making, detecting fraudulent transactions is intrinsically difficult. Conventional machine learning techniques, such as logistic regression, decision trees, and support vector machines, may perform worse when there is a significant class imbalance because they frequently fail to capture intricate relationships between transactions.

Recent advances in graph-based deep learning have demonstrated strong potential for fraud detection by modeling relationships among entities rather than relying solely on independent feature representations. By learning contextual representations from nearby transactions, Graph Neural Networks (GNNs), especially attention-based architectures, improve the ability to distinguish between fraudulent and legitimate activity. Among these, attention mechanisms intended for temporal interaction modelling are introduced by the Temporal Graph Attention Network (TGAT). Alternative similarity-based graph formulations are necessary because explicit temporal graph structures are either unavailable or challenging to construct in many real-world financial datasets.

This study suggests a hybrid fraud detection framework that blends gradient boosting classification with graph representation learning in order to overcome these difficulties. To address class imbalance, transaction data are first preprocessed using Borderline-SMOTE and feature scaling, which treat the Time attribute as a numerical feature instead of an explicit timestamp. The scaled feature space is then used to create a static k-nearest neighbor similarity graph, and informative node embeddings are learned from graph relationships using a TGAT architecture. After being fused with the original transaction features, these learned embeddings are fed into an XGBoost classifier, which uses optimized decision thresholds derived from precision–recall analysis to perform the final fraud prediction.

Model interpretability is crucial for regulatory compliance and financial decision support, in addition to predictive performance. In order to improve transparency and trust in the suggested system, this work integrates SHAP-based explainable artificial intelligence to provide explanations at the local transaction level as well as at the global feature importance level. Stratified cross-validation experimental evaluation shows strong performance on several metrics, such as F1-score, ROC-AUC, PR-AUC, Accuracy, Precision, and Recall.

II. LITERATURE SURVEY

Traditional machine learning, deep learning, graph-based representation learning, and ensemble classification methods have all been widely applied in the investigation of credit card fraud detection. In order to overcome issues like extreme class imbalance, intricate transaction relationships, and restricted model interpretability in financial decision systems, recent research has focused on enhancing detection performance. To achieve precise and comprehensible fraud detection in this context, the suggested framework combines static graph-based TGAT embeddings with XGBoost classification and SHAP-based explainable artificial intelligence.

Nagarajan et al. proposed a hybrid CNN–SVM and XGBoost-based framework for online transaction fraud detection, achieving improved classification accuracy, F1-score, and AUC in contrast to independent models. In 2025, Sathis Kumar and colleagues presented a hybrid fraud detection framework that incorporates Graph Neural Networks with XGBoost to model relational dependencies among financial transactions. The graph-based component captures interaction patterns between entities, while gradient boosting refines the final prediction to improve fraud identification capability and reduce false alarms. [2]. Pattnayak et al. (2025) explored fraud detection in business transactions using machine learning on large-scale financial datasets, emphasizing ensemble methods such as Random Forest for capturing complex behavioral fraud patterns beyond rule-based systems. [3]. Suryanarayana et al. (2025) offered a method for identifying credit card fraud using the CatBoost algorithm to improve identification of fraudulent transactions in digital financial systems. The approach applies data preprocessing, feature selection, and gradient-boosting–based classification to enhance fraud recognition and reduce detection time. [4]. A hybrid framework combining XGBoost and Deep Neural Networks was presented by Usheelamma et al. in 2025 to analyze digital transaction patterns and support fraud identification in financial systems. The approach integrates feature extraction from XGBoost with deep learning–based pattern modeling to enable anomaly detection, transaction categorization, and real-time analytical insights. [5]. Afriyie et al. (2023) investigated methods for supervised machine learning, such as Random Forest, Decision Tree, and Logistic Regression, for identifying fraudulent credit card transactions using simulated financial data. The study emphasized preprocessing, imbalance handling, and comparative evaluation to identify the most suitable predictive technique for fraud identification. [6]. Chang et al. (2024) used sophisticated machine learning techniques to investigate credit card payment fraud detection while addressing severe class imbalance through resampling techniques like SMOTE and random un-

Der sampling. Logistic regression, K-nearest neighbors, support vector machines, and decision trees were among the classifiers used in the study to analyze transactional behavior and support reliable fraud identification in digital payment environments. [7]. Mienye et al. (2025) proposed a hybrid fraud detection framework integrating Graph Attention Network relational learning and Variational Autoencoder-based anomaly detection and XGBoost

stacking to address severe class imbalance and evolving fraud behavior. [8]. Sha et al. (2025) proposed a technique to model multi-relational transaction networks involving users, merchants, and financial entities using heterogeneous graph neural networks to detect credit card fraud. The framework integrates graph attention mechanisms, temporal decay modeling, and imbalance-handling strategies to capture complex relational and time-dependent fraud patterns in large financial datasets. [9]. Nti et al. (2024) proposed a scalable financial fraud mitigation framework that integrates Random Forest–based feature engineering with XGBoost classification to analyze large-scale transaction data and detect anomalous financial activities. The method incorporates adaptive over-sampling, feature selection, and parameter optimization to improve fraud identification in highly imbalanced datasets. [10]. Sadaf and Manivannan (2024) presented a fraud detection framework for UPI transactions using the Gradient Boost-ing algorithm with feature engineering and imbalance-aware learning to identify suspicious financial activities in digital payment systems. [11]. Majumder and Mishra (2024) created a credit card fraud detection system that addresses significant class imbalance in financial transaction datasets by combining XGBoost and SMOTE. The study emphasizes how machine learning-based automated detection can lower financial losses and enhance the detection of fraudulent activity in large-scale digital payment environments. [12]. Alrasheedi (2025) carried out a comparison of deep learning, ensemble learning, and machine learning models for detecting credit card fraud using both balanced and imbalanced transactional datasets. [13]. A framework for detecting mobile payment fraud was proposed by Hajek et al. in 2023 integrating random under-sampling, unsupervised outlier detection, and XGBoost to handle severe class imbalance in large-scale transaction data. The semi-supervised ensemble method transforms anomaly scores and combines them with supervised learning to enhance fraud identification and financial risk analysis in mobile pay-ment environments. [14]. Sivakumar et al. (2025) suggested a fraud detection system based on Graph Neural Networks (GNNs) that models' relationships among users, accounts, and transactions to identify abnormal financial behavior in digital transaction environments. [15]. Amulya et al. (2025) suggested using Decision Tree, Random Forest, Gradient Boosting, and AdaBoost models in a machine learning-based UPI fraud detection system to identify fraudulent digital payment trans-actions. [16]. Narkhede and Jagtap (2025) examined the use of Python machine learning techniques for online payment fraud detection, emphasizing issues in digital payment ecosystems like identity theft, account takeover, and card-not-present fraud. [17]. Challapalli et al. (2025) presented a financial fraud detection framework integrating Big Data analytics with hybrid deep learning architectures, specifically CNN and LSTM, to model spatial and temporal transaction patterns for improved fraud identification. [18]. Vashistha et al. (2024) suggested a framework for detecting banking fraud using neural network and machine learning models with SMOTE-based imbalance handling and comparative evaluation across multiple classifiers, including Random Forest, Decision Tree. [19]. Dalal et al. (2022) presented a framework for financial payment services to predict fraud using a hyperparameter-tuned, optimized XGBoost model combined with supervised learning and cross-validation strategies. [20]

III. METHODOLOGY FOR IDENTIFYING CREDIT CARD FRAUD

A hybrid temporal graph–ensemble learning is introduced by the suggested framework that integrates Temporal Graph Attention Networks (TGAT) with XGBoost classification, along with SMOTE-based imbalance handling and SHAP-based explainability. The following sequential steps make up the system's workflow, which is depicted in Fig. 2.

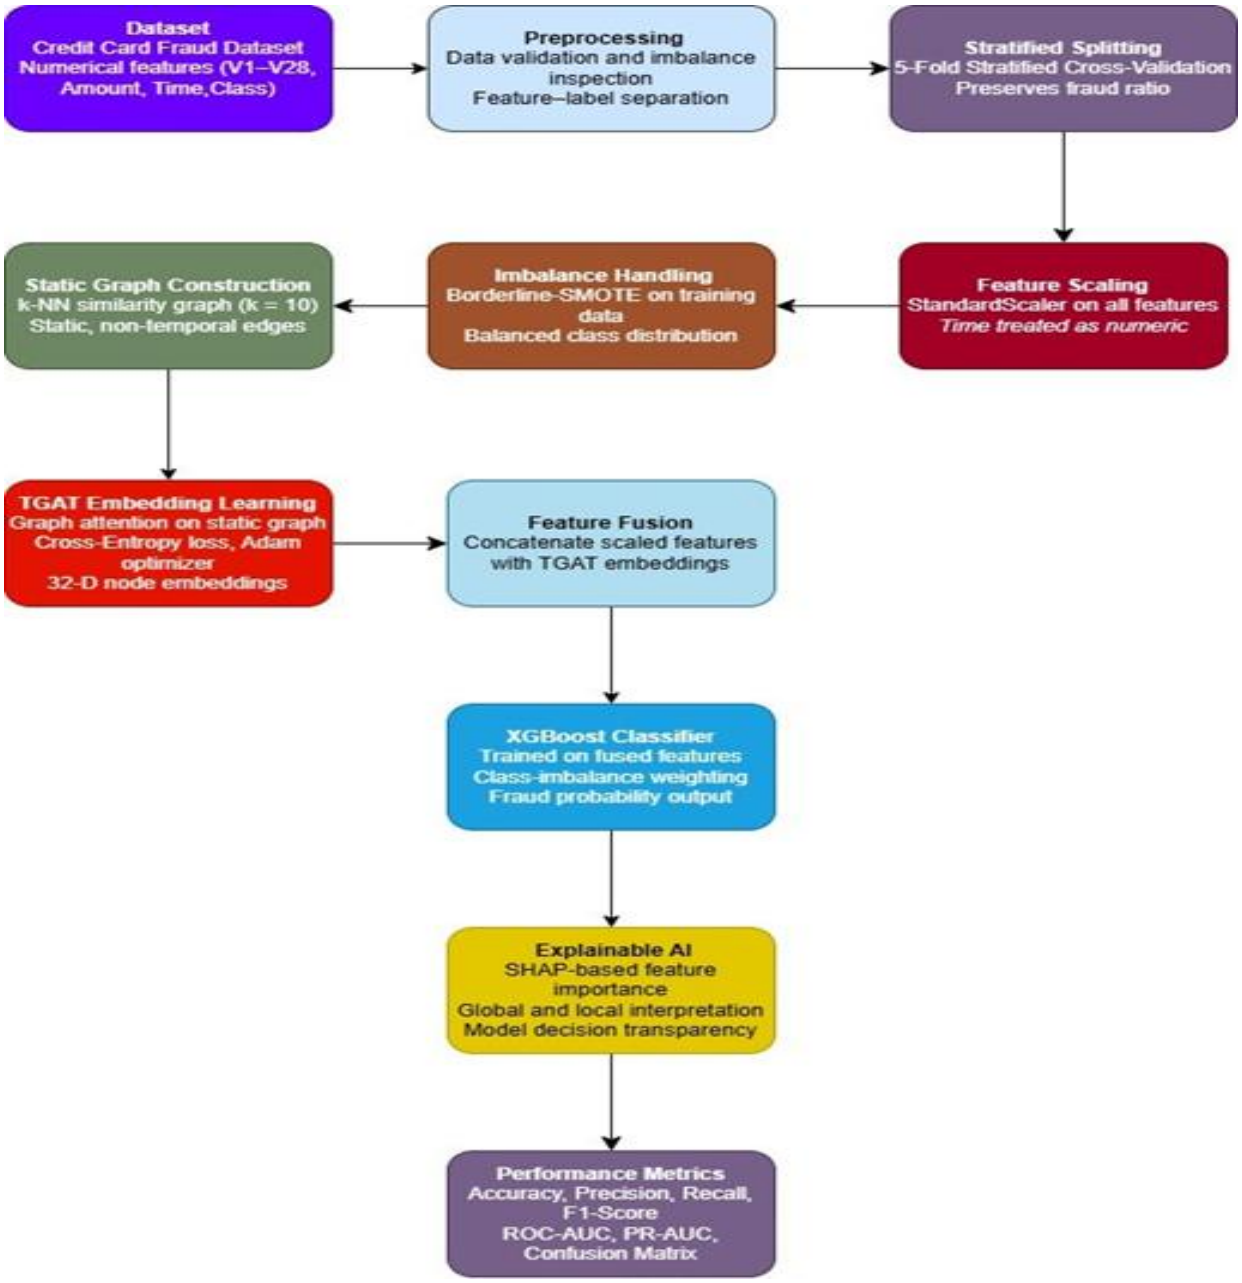

Fig. 1. Proposed Methodology for Fraud Detection

IV. METHODOLOGY

This section describes the complete workflow of the pro-posed hybrid fraud detection framework that integrates graph representation learning, ensemble classification, and explain-able artificial intelligence. The objective is to accurately identify fraudulent credit card transactions from highly imbalanced financial data while maintaining interpretability and robust-ness.

A. Dataset Representation

A publicly accessible credit card fraud detection dataset, comprising 284,807 transactions over a two-day period, is used by the suggested model. Only 492 transactions are classified as fraudulent among these, which corresponds to a highly unbalanced fraud ratio of roughly 0.172%. Anonymized numerical variables V1–V28, transaction Amount, and Time, as well as a binary class label designating fraudulent or legitimate activity, are used to describe each transaction. While

maintaining the discriminative statistical patterns required for fraud identification, this anonymization protects privacy.
The dataset is formally represented as:

D = {(xi, yi)} N (1)
where the fraud label is represented by yi ∈ 0, 1} and the feature vector of the ith transaction is denoted by xi ∈ Rd. This supervised formulation enables the learning of a predictive mapping f : xi → yi.

B. Data Preprocessing
Scale variances, redundant features, and extreme class im-balances in raw financial transaction data can all have a detrimental effect on model learning. As a result, before building and classifying graphs, a structured preprocessing pipeline is used.
The dataset is first separated into a label vector y and a feature matrix X, where:

X = {xi}N, y = {yi}N, yi ∈ {0, 1}. (2)
Unbalance-aware resampling is necessary because fraud-ulent samples are incredibly uncommon. All preprocessing procedures are carried out **inside each training fold only** to ensure objective evaluation and stop information leaks.

C. Stratified Cross-Validation Splitting
The suggested framework uses Stratified k-Fold Cross-Validation with k = 5 to obtain accurate performance estima-tion under severe class imbalance. In every training and testing partition, this approach maintains the initial fraud ratio.
Let the dataset indices be partitioned as:

$$D = \bigcup_{j=1}^{k} D_j, \quad D_j \cap D_l = \emptyset \ (j \neq l). \tag{3}$$

For the j-th fold, the training and testing sets are defined as:

$$D_{train}^{(j)} = D \setminus D_j, \quad D_{test}^{(j)} = D_j \tag{4}$$

Stratification guarantees:

$$\frac{Fraud_{train}}{|D_{train}|} \approx \frac{Fraud_{test}}{|D_{test}|} \approx \frac{Fraud_{total}}{N}. \tag{5}$$

This ensures fair comparison and prevents biased performance estimates.

D. Financial transaction attributes often exhibit different numerical ranges, which may bias gradient-based learning and distance-based similarity measures. To address this issue, all features are standardized using z-score normalization

$$x'_i = \frac{x_i - \mu}{\sigma}, \tag{6}$$

where μ and σ represent the estimated mean and standard deviation from the training subset only. This prevents information leakage and ensures fair model evaluation.
In this work, the Time attribute is treated strictly as a numeri-cal feature rather than a temporal timestamp. Consequently, the downstream graph structure captures similarity relationships instead of chronological

E. E. Class Imbalance Handling using Borderline-SMOTE
Credit card fraud datasets are extremely imbalanced, where fraudulent samples constitute only a small fraction of total transactions. Training directly on such data leads to biased classifiers favoring the majority class. The Borderline-SMOTE technique is used on the training data to artificially create minority samples close to the decision boundary in order to lessen this restriction.
An instance of synthetic fraud is produced as:

$$x_{new} = x_i + \lambda(x_{nn} - x_i), \quad \lambda \sim U(0, 1), \tag{7}$$

where xnn represents one of its closest minority neighbors and xi is a minority-class sample.

F. Static k-Nearest Neighbor Graph Construction
To capture inter-transaction relationships beyond indepen-dent tabular learning, a similarity-based graph is constructed using the k-nearest neighbor algorithm. Every transaction is shown as a node, with edges joining neighbors with similar features.

$$G = (V, E), \tag{8}$$

$$E_{ij} = \begin{cases} 1, & \text{if } x_j \in kNN(x_i), \\ 0, & \text{otherwise.} \end{cases} \tag{9}$$

Because temporal ordering is not used, this structure forms a static similarity graph. Such relational

modeling enables detection of coordinated or behaviorally similar fraud patterns that are not visible in independent feature space.

G. TGAT-Based Graph Embedding Learning

Graph relationships are encoded using a Temporal Graph Attention Network (TGAT) operating on the constructed static graph. By giving nearby nodes adaptive importance weights,the model can focus on the most instructive relational interactions thanks to the attention mechanism.

The following is how attention coefficients are calculated:

$$\alpha_{ij} = \frac{\exp\ \text{LeakyReLU}\ \ a^{\top}[Wh_i \mid Wh_j]}{\sum_{k\in N\ (j)} \exp\left(\text{LeakyReLU}\left(a^{\top}[Wh_i \mid Wh_k]\right)\right)}. \quad (10)$$

Node embeddings are updated via:

$$h_i' = \sigma\left(\sum_{j\in N\,(j)} \alpha_{ij} W h_j\right). \quad (11)$$

Model parameters are optimized using binary cross-entropy loss

$$L = -\sum_{i=1}^{N} [y_i \log \hat{y}_i + (1 - y_i) \log(1 - \hat{y}_i)]. \quad (12)$$

This process produces compact low-dimensional embeddings capturing relational fraud behavior.

H. Hybrid Feature Fusion

To combine statistical and relational knowledge, TGAT embeddings zi are concatenated with scaled original features:

$$\tilde{x}_i = [x_i' \mid z_i]. \quad (13)$$

This hybrid representation enriches discriminative information and improves downstream classification capability.

I. XGBoost-Based Fraud Classification

Fraud prediction is performed using XGBoost, an ensemble of gradient-boosted decision trees capable of handling nonlinear feature interactions and imbalanced data distributions.

$$\hat{y}_i = \sum_{k=1}^{K} f_k(\tilde{x}_i), \quad f_k \in F. \quad (14)$$

The training objective is:

$$L_{XGB} = \sum_i \ell(y_i, \hat{y}_i) + \sum_k \Omega(f_k), \quad (15)$$

$$\Omega(f) = \gamma T + \frac{1}{2}\lambda |w|^2. \quad (16)$$

Regularization controls model complexity and prevents overfitting

J. Explainable Artificial Intelligence using SHAP

To ensure interpretability and regulatory transparency, Fea-ture contributions are measured using SHAP (Shapley Additive Explanations). The definition of the Shapley value for feature j is:

$$\phi_j = \sum_{S\subseteq F\setminus\{j\}} \frac{|S|!(|F| - |S| - 1)!}{|F|!} f_{S\cup\{j\}}(x) - f_S(x). \quad (17)$$

This provides both global feature importance and local transaction-level explanation.

V. EXPERIMENTAL SETUP

The proposed fraud detection model was developed in Python with the aid of PyTorch Geometric, Scikit-learn, XGBoost, and SHAP libraries. The model was assessed using 5-fold stratified cross-validation, and a 80:20 ratio of the data was reserved for training and testing, respectively, while maintaining the extreme class imbalance of the credit card dataset. During preprocessing, numerical variables were standardized using z-score normalization, and Borderline-SMOTE Oversampling was only applied to the training folds. training folds to produce the class of minority fraud without causing data leakage. A static k-nearest neighbor similarity graph with k = 10 was constructed to learn 32-dimensional TGAT embeddings, which were concatenated with scaled features and classified using an XGBoost model with class-imbalance weighting. Finally, SHAP-based explainable artificial intelligence was employed for model interpretation, and performance was evaluated

using Accuracy, Precision, Recall, F1-score, ROC-AUC, and PR-AUC.

VI. RESULTS AND DISCUSSION

A. Analysis of Confusion Matrix for Proposed Ensemble Model

The TGAT + XGBoost model's confusion matrix further demonstrates its effectiveness in detecting fraudulent and actual transactions. It is evident that the model has successfully classify 56,848 actual transactions as non-fraud with only 15 actual transactions being incorrectly classified as fraud, which demonstrates the model's exceptionally low false positive rate and high reliability for actual transaction filtering.

In the case of the fraud class, the model has been able to correctly identify 79 fraudulent transactions with only 19 fraud cases being missed, which is a clear indication of the model's high fraud detection capability even with a highly imbalanced class. The confusion matrix clearly verifies the stability, accuracy, and real-world applicability of the proposed hybrid approach for financial fraud detection in real-world scenarios, where both the false negative and unnecessary transaction filtering rates need to be kept to a minimum.

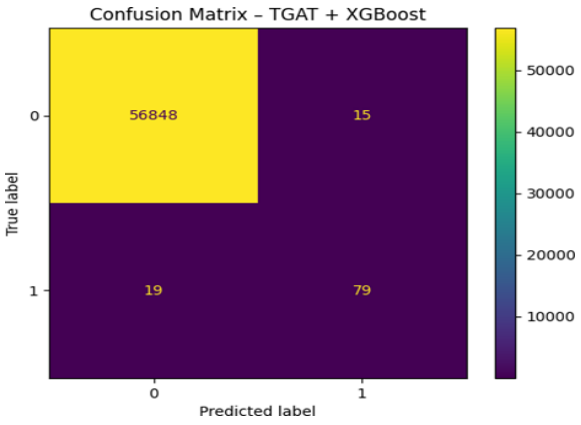

Fig. 2. Analysis of Confusion Matrix for Proposed Ensemble Model

B. ROC CurveAnalysis for Proposed Hybrid Ensemble Model

The ROC curve, or Receiver Operating Characteristic further supports the discriminative power of the examined connection between the false positive and true positive rates rate at various decision thresholds using the suggested TGAT + XGBoost hybrid model. As evident, the ROC curve remains well above the diagonal reference line, thereby ensuring effective classification and a proper distinction between fraudulent and genuine transactions.

The AUC value of approximately 0.9859 obtained in this study clearly indicates excellent class separability with very less misclassification. As an AUC value close to 1.0 indicates near-perfect predictive power, this further supports that the proposed TGAT + XGBoost hybrid model ensures high sensitivity, low false alarm rates, and overall excellent fraud detection performance.

C. Precision-recall curve Analysis for proposed Hybrid Ensemble Model

A more useful method for evaluating the effectiveness of the suggested TGAT + XGBoost hybrid model in the presence of a high degree of class imbalance, because it illustrates how recall and precision are traded off for various decision thresholds. As seen, the PR curve preserves a high level of precision for a wide range of recall values, which clearly indicates the high capacity of the model to correctly classify fraudulent transactions while avoiding false positives. The high PR-AUC score of about 0.7655 demonstrates the model's outstanding performance in correctly identifying the minority class, which is crucial for fraud detection applications, where positive instances are extremely rare.

D. SHAP-Based Feature Importance and Model Interpretability

The SHAP summary plot demonstrates the global feature importance of the proposed TGAT + XGBoost fraud detection

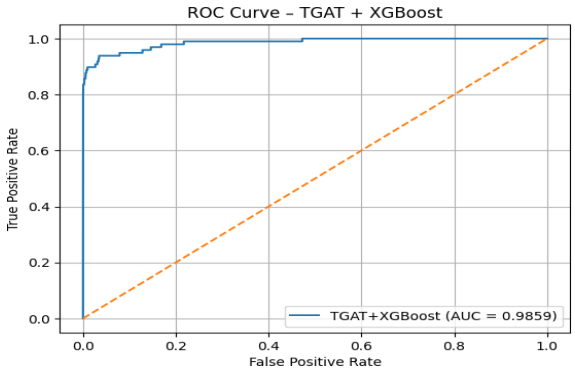

Fig. 3. ROC Curve Comparison of Hybrid Models

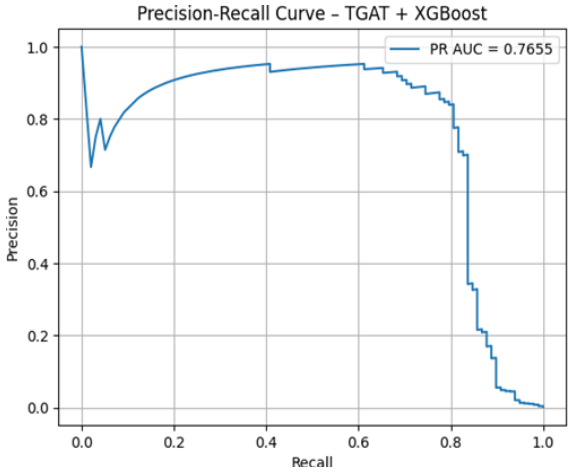

Fig. 4. Precision-recall Comparison of Hybrid Models

system by displaying how each input feature contributes to the prediction outcomes for every transaction. The features are ordered by mean absolute SHAP values, with Feature 31, Feature 4, and Feature 12 being the most important, while the horizontal distribution and blue-to-red color gradient indicate the magnitude, direction, and value-dependent importance for the fraud classification.

The plot is obtained from the model's final fused feature space, where the original 30 tabular features (V1 through V28, Time, and Amount) are concatenated with 32-dimensional TGAT embeddings, resulting in 62 total input features for the XGBoost model. Hence, the higher-index features (such as Feature 31, Feature 52, and Feature 56) represent graph-derived embedding dimensions, whose importance indicates that the relational feature representations learned by the TGAT model have a significant impact on improving fraud discrimination and, therefore, validate the efficacy and interpretability of the proposed hybrid.

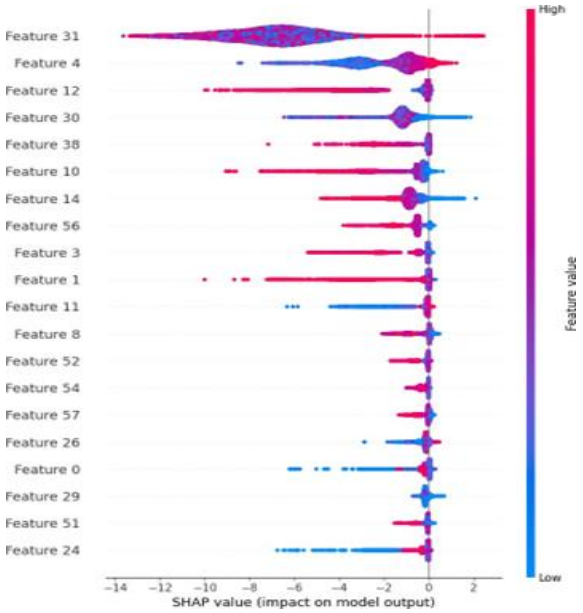

Fig. 5. SHAP summary plot showing global feature importance of the TGAT+ XGBoost fraud detection model.

E. Cross-Validation Analysis for Proposed Ensemble Model

The suggested TGAT + XGBoost fraud detection sys-tem's effectiveness was evaluated using 5-fold stratified cross-validation to accurately measure its performance under ex-treme class imbalance. The system resulted in a mean accuracy of 0.9994, which is an indication of high overall consistency in classification. Moreover, the precision (0.8404) and recall (0.8061) values obtained by the system are an demonstration of the system's high degree of transaction classification accu-racy as fraudulent while maintaining a high level of control over false positives. The system's F1-score of 0.8229 shows that the system's sensitivity and precision are well-balanced. Additionally, the system's high capacity to discern between fraudulent and legitimate transactions is demonstrated by its ROC-AUC score of 0.9859.

```
====== FINAL CROSS-VALIDATION RESULTS ======
Mean Accuracy : 0.9994
Mean Precision: 0.8404
Mean Recall   : 0.8061
Mean F1-Score : 0.8229
Mean ROC-AUC  : 0.9859
```

Fig. 6. Cross-Validation Analysis for Proposed Ensemble Model

VII. CONCLUSION AND UPCOMING WORK

A hybrid credit card fraud detection system based on relational embeddings and TGAT was presented in this study. SHAP-based explainable AI and XGBoost classification. With a static similarity graph and combined tabular-graph features, the system demonstrated high accuracy, good precision-recall curves, and ROC-AUC values very close to 0.99 for 5-fold stratified cross-validation, and successfully identified the minority fraud instances. The SHAP value analysis also validated the feature contributions and enhanced the model interpretability, and thus demonstrated how the suggested framework for financial fraud detection could be used in practical situations. The future work can be directed towards more accurate temporal graph modeling, online learning for real-time fraud detection, better optimization strategies for imbalance prob-
lems, and more advanced explainable AI techniques.

REFERENCES

[1] P. Nagarajan and B. Kamali, "Fraud detection in online transaction using hybrid CNN and XGBoost," 2024 International Conference on Emerging Research in Computational Science (ICERCS), pp. 1–6, 2024, doi: 10.1109/ICERCS63125.2024.10895440.

[2] N. R. Sathis Kumar et al., "Fraud detection in financial transaction using GNN with XGBoost," Proc. 4th OPJU Int. Technology Conf. (OTCON), IEEE, 2025, doi: 10.1109/OTCON65728.2025.11070687.

[3] J. Pattnayak, A. Choudhary, H. K. Manjula, P. Salgotra, S. T. Gupta, and T. Sirisha, "Fraud detection in business transactions: A machine learning-based approach," Proc. 2025 World Skills Conf. on Uni-versal Data Analytics and Sciences (WorldSUAS), IEEE, 2025, doi: 10.1109/WorldSUAS66815.2025.11199082.

[4] V. Suryanarayana, K. Maddileti, D. Satyanarayana, R. L. Jyothi, K. Sreekanth, P. Mande, R. N. Miriyala, and O. Sudhakar, "An efficient implementation of credit card fraud detection using CatBoost algorithm," Indonesian Journal of Electrical Engineering and Computer Science, vol. 38, no. 3, pp. 1914–1923, 2025, doi: 10.11591/ijeecs.v38.i3.pp1914-1923.

[5] S. K. H. Susheelamma et al., "A hybrid model for predicting and analyzing digital transaction patterns using deep neural networks and XGBoost," Proc. 2nd Int. Conf. on New Frontiers in Communication, Automation, Management and Security (ICCAMS), IEEE, 2025, doi: 10.1109/ICCAMS65118.2025.11234561.

[6] J. K. Afriyie, K. O. Adu-Boahen, E. B. Assante, and J. M. Osei, "A supervised machine learning algorithm for detecting and predicting fraud in credit card transactions," Decision Analytics Journal, vol. 6, p. 100163, 2023, doi: 10.1016/j.dajour.2023.100163.

[7] V. Chang, B. Ali, L. Golightly, M. A. Ganatra, and M. Mohamed, "Investigating credit card payment fraud with detection methods using advanced machine learning," Information, vol. 15, no. 8, p. 478, 2024, doi: 10.3390/info15080478.

[8] I. D. Mienye, E. Esenogho, and C. Modisane, "Detecting imbalanced credit card fraud via hybrid graph attention and variational autoencoder ensembles," Applied Mathematics, vol. 5, p. 131, 2025, doi: 10.3390/ap-pliedmath5040131.

[9] "A scalable RF–XGBoost framework for financial fraud mitigation," IEEE Transactions on Computational Social Systems, vol. 11, no. 2,pp. 1556–1562, 2024, doi: 10.1109/TCSS.2022.3209827.

[10] R. Sadaf and R. Manivannan, "Enhanced detection of fraud in Unified Payments Interface (UPI) transactions using gradient boosting method," International Journal of Interpreting Enigma Engineers, vol. 2, no. 2, 2025.

[11] T. Majumder and D. Mishra, "Financial fraud detection for credit card using XGBoost SMOTE," Nanotechnology Perceptions, vol. 20, no. S15, pp. 32–50, 2024.

[12] Masad A. Alrasheedi, "Enhancing fraud detection in credit card transac-tions: A comparative study of machine learning models," Computational Economics, Springer, accepted July 2025.

[13] P. Ha´jek, M. Z. Abedin, and U. Sivarajah, "Fraud detection in mobile payment systems using an XGBoost-based framework," Information Systems Frontiers, vol. 25, pp.

1985–2003, 2023.

[14] V. Sivakumar, N. Saw, A. K. Verma, V. Harikrishnan, and P. Kumar, "Fraud detection in financial transaction using GNN," Proc. Int. Conf. on Visual Analytics and Data Visualization (ICVADV), IEEE, 2025, doi: 10.1109/ICVADV63329.2025.10960835.

[15] K. Amulya, M. Bhumika, T. S. Murari, S. Komaravalli, and V. V. Gopal, "UPI fraud detection using machine learning," International Journal of Scientific Development and Research (IJSDR), vol. 10, no. 4, 2025.

[16] D. N. Narkhede and S. Jagtap, "Online payment fraud detection using machine learning in Python," International Journal of Progressive Research in Engineering Management and Science (IJPREMS), vol. 5, no. 4, pp. 74–76, 2025

[17] C. Sujana, A. Sasi Himabindu, D. Srinivasa Rao, S. Rasamsetty, A. M. S. Arunkumar, P. Muthukumar, and S. K. Shanmugam, "Big data and artificial intelligence revolutionizing financial fraud detection systems," Journal of Theoretical and Applied Information Technology, vol. 103, no. 18, 2025.

[18] A. Vashistha, A. K. Tiwari, P. Singh, P. K. Yadav, and S. Pandey, "A robust framework for fraud detection in banking using ML and NN," Proc. Natl. Acad. Sci., India, Sect. A Phys. Sci., vol. 94, no. 2, pp. 201–212, 2024, doi: 10.1007/s40010-024-00871-1.

[19] Surjeet Dalal, Bijeta Seth, Magdalena Radulescu, Carmen Secara, and Claudia Tolea, "Predicting Fraud in Financial Payment Services through Optimized Hyper-Parameter-Tuned XGBoost Model," Mathematics, vol. 10, 2022.

[20] J. Liu, H. Wang, and Z. Du, "Temporal Graph Neural Networks for Evolving Financial Fraud Detection," IEEE Transactions on Knowledge and Data Engineering, vol. 36, no. 4, pp. 3892–3905, 2024, doi: 10.1109/TKDE.2024.3279123.