MVGA: MULTI-VIEW GRAPH ATTENTION BASED FRAMEWORK FOR EARLY FRAUD DETECTION IN FINANCIAL TRANSACTION NETWORKS

1st Sunkara Bhavani Harika
M. Tech Student
Department of CSE
Vasireddy Venkatadri Institute Of Technology(VVIT),
Nambur, Guntur, India
bhavaniharika22@gmail.com

2nd Gumma Lakshmi Narayana
Associate Professor,
Department of CSE ,
Vasireddy Venkatadri Institute Of Technology(VVIT),
Nambur, Guntur, India
lakshminayarana.gumma@gmail.com

3rd *Srinivasa Rao Pendela,*
Professor,
Department of CSE,
Vasireddy Venkatadri Institute Of Technology(VVIT),
Nambur, Guntur, India
psr.srinivas999@gmail.com

4th Ramachandran Vedantham
Professor,
Department of CSE,
Vasireddy Venkatadri Institute Of Technology(VVIT),
Nambur, Guntur, India
vrc.bhatt@gmail.com

Abstract- The financial transaction systems produce high amounts of interconnection, constantly developing information, and fraudulent detection is therefore vital and difficult at its initial stages. The rule-based and feature-centric traditional approaches are not always capable of reporting coordinated fraud behaviors that are displayed in connection channels. Here we introduce a Multi-View Graph Attention (MVGA) framework, which captures the financial process network by an array of graph views (account-to-account monetary flows, sequence of transactions, and contextual relations), and processes (accounts) as well. The model learns behavioral cues like relationship strength, intensity of spending and frequency of interactions by building complementary views of graphs where accounts and transactions are the nodes and the edges respectively. An attention-based message aggregation scheme places more emphasis on the suspicious nodes, links between transactions, and view-level fusion incorporates the signals within each view, and coordinates fraud patterns using a coordinated fraud scheme. An experimental assessment of a benchmark fraud detection environment proves the reliability of the proposed MVGA at Precision = 0.92, Recall = 0.88, and F1-score = 0.903 that the MVGA provides reliable alerts to fraud incidents in a rather well-covered neighborhood. The model also achieves ROC-AUC = 0.945 and PR-AUC = 0.89 and is sufficient to indicate better separability of threshold-independence and good performance in the case of class imbalance. Generally, the findings confirm that, among the multi-view relational learning methods, the attention-based fusion, improves the capacity of early fraud detection, the false alarms, and the model can adapt and adapt well to the changing strategies of fraud in dynamic financial conditions.

Keywords- *Attention Mechanisms, Financial Fraud, Graph Analytics, Network Security, Pattern Anomaly Detection, Risk Assessment, Transaction Networks, Unsupervised Learning.*

I. INTRODUCTION

There has been growing complexity and volume of fraudulent transactions in different fields like finance, health insurance and online gaming which have necessitated a debate on the significance of fraud detection. The most important is fraudulent activity that should be detected early, and financial losses should be minimized, systemic risks will be prevented, and trust in the ecosystem of transactions will be preserved [1]. Conventional methods, which mostly are rule-based systems and classical machine learning (ML) methods based on handcraft features, are severely limited. They tend to be slow to change with the changing nature of frauds and also highly susceptible to false positive results because of strict rule-writing and inflexible coverage of complex relationship data [2]. There is therefore an urgent requirement to find adaptive and data driven approaches that can deal with the dynamism and complexity of transactional data.

Graph-based modelling has been a rather encouraging paradigm tool of detecting fraud due to its natural ability to represent the rich relational structure of entities like customers, transactions, merchants and

accounts in the transaction networks [3]. Financial fraud can be presented in the form of anomalous subgraphs structure or even questionable multi-relation links which are generally ignored by flat-feature models [4]. GNNs are based on the concept of connectivity and attribute data to learn representations that reflect the underlying relational data, and outperform conventional ML in detecting complex fraud behaviors in heterogeneous transaction graphs. In addition, a capacity of GNNs to generalize due to the topology of a graph to a task of classifying nodes makes them applicable to fraud prediction in real-time in quite intricate financial ecosystems.

Graph attention mechanism, multi view graph learning has also improved the science of detecting fraud as it allows the selective aggregation of data using various relational views, increasing the model interpretability and robustness. Attention-based GNNs learn friends and types of edges by assigning them learnable weights, which help reduce the effects of noisy or irrelevant connections that are often present because of camouflage of fraud and data imbalance [5]. Multi-view or heterogeneous graph models: these models use varying sources of data, including the type of transactions, temporal variability, as well as the behaviour of users allowing them to fuse the non-overlapping features in a holistic manner through attention layers [6]. This type of fusion boosts resistance to detection accuracy and resistance to the changing methods of fraud by prioritizing the most informative views that are dynamically picked in model-training and inference.

Even with these developments, multi-view graph attention has not been fully used in early fraud detectors, and significant issues still require solving in problem areas which include class imbalance, time-dependent changes in fraudulent behaviour and obfuscation methods like camouflage [7]. The current models do not necessarily have specific mechanisms that can allow them to detect the fraud in a timely manner, and at the same time, they need to be very accurate in cases where there are few positive samples, and complex noisy graph neighbourhoods. To address this gap, we introduce a new Multi-View Graph Attention architecture which is specifically tailored towards the early fraud prevention in graph-based transaction markets. We combine several relational outlooks in adaptive attention fusion, combine temporal and structural dynamics to find anomalies quickly and use mechanisms to reduce class imbalance and camouflage effects [8]. Key insights of this paper are: 1) multi-view graph attention with a unified architecture that is geared towards early fraud detection; 2) effective fusion operations on heterogeneous and temporally changing transaction data; 3) effective countermeasures to noisy neighbours and class imbalance and 4) extensive testing on real-life and financial datasets that show better performance than the state-of-the-art approaches.

II. LITERATURE REVIEW

More recent innovations in 2023 through 2025 have made major breakthroughs in that graph-based fraud detection now has the ability to use heterogeneous and multi-relational graph design approaches, which have a greater ability to encapsulate the intricate relationships between various entities within financial and transaction data. Open-loop gift cards fraud frameworks based on graphs are represented using the heterogeneous GNNs to learn the multi-relationship among the customers, merchants, transactions, and cards, which allows detecting anomalous patterns concealed under structural and semantic richness [9]. Likewise, models of e-commerce fraud communities combine online and offline data applied to the same networks by multi-view heterogeneous GNNs to identify rings of fraud by taking advantage of spatial-temporal network information, thus being able to capture more information on the relational contexts of the flat features [10]. This line is also developed in community-based heterogeneous graph modeling, which explicitly models heterogeneous relations, such as user-item relation, device-relation, and community-driven filters are used to address the problem of irregular structure and content and enhance the accuracy of the detection in large-scale settings [11]. All these studies highlight the effectiveness of the multi-relational heterogeneous graph modelling in detecting frauds but are mostly deficient in dynamics time of events that may be used in during early interventions.

Among the attention mechanisms and temporal graph methodologies, recent literature shows that attention-based modules will play a pivotal role in detecting noise that highlights important nodes and connections in the graph. Neural meta-graph searching methods utilize the message-passing of the GNNs to get interpretable paths that are used to predict fraud and the transparency of the models [12]. Although still considered an emergent concept in literature, temporal context has been identified to be essential in detecting fraud at its first instance, since it records the dynamic nature and timing of transactions that are highly significant in intervention on a timely basis. Even though the explicit temporal graph approach is not specifically studied in the retrieved articles, the temporal attention combined with semantic data should be realized to predict the changes in the pattern of fraudulent behaviour prior the time when it escalates [13]. Temporal dynamic attention will thereby improve the sensitivity and timeliness of non-static graph-based models of fraud detection.

To overcome the challenges of multi-view, multi-

relation, and heterophily, recent models employ heterophily aware learning to capture better the structure of fraud graphs that have dis-similar nodes who are interconnected in complicated manners. GNN-based semi-supervised models that exploit spectral filtering and local environmental constraints have proven to be better at handling heterophilic connection which are a property of fraud graphs, in which a traditional homophily operation assumption is violated [14]. The self-managed heterogeneous graph representation learning models apply latent graph based on homophilic/heterophilic information with adaptive fusion, this enhances node embeddings on noisy conditions. Moreover, label-balanced sampling (and neighbour selection) add imbalanced approaches to learning focus on the important relations that reduce noise caused by irrelevant linkages and strengthen the robustness of detection. Although true advancements are being made, there are still difficulties in effectively integrating various heterogeneous perceptions in changing and noisy fraud environments.

Although this has changed, research gap remains in early detection of fraud. Most existing techniques do not handle the situation of class imbalance well, as it is extreme in a real-life context of fraud, and results in biased learning. Also, impairment of feature and structural consistency as camouflage technologies to fraud has impeded model effectiveness and requires powerful neighbour selection and discriminative feature guidance [17]. The dynamic evolving pattern of fraudulent behaviour requires models, which can process time-dependent trends adaptively to allow timely alerts, although a consolidated multi-view graph attention framework that explicitly focuses on identifying fraud at its initial appearance is yet to be found [18]. In addition, the heterogeneous graphs of the noisy neighbours pose a challenge to the process of graph aggregation, requiring more sophisticated attention fusion algorithms to dynamically balance among heterogeneous inputs to better interpretability and resiliency. To address them, our proposed model combines adaptive multi-view attention fusion with temporal and structural dynamics, as well as methods of dealing with imbalance and camouflage to enable more reliable and explainable early detection of fraud.

III. PROPOSED ARCHITECTURE

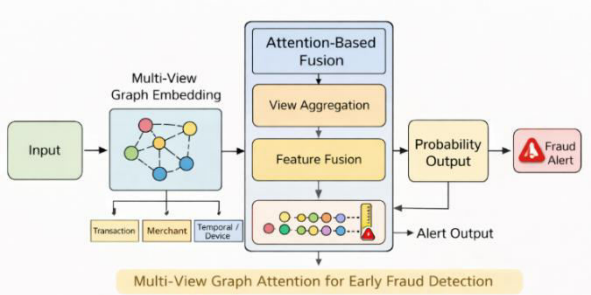

Fig 1. Architecture of the Proposed Multi-View Graph Attention Framework for Early Fraud Detection

Such architecture represents a basic and transparent pipeline of Multi-View Graph Attention based Early Fraud Detection. It begins with the input transaction data which is transformed to a multi-view graph embedding representation where various relationships and edges which include transaction links, merchant behaviour, and temporal or device patterns are formed independently of each other. The views of these graphs are then inputted into an attention-based fusion block where the model would firstly perform view aggregation to determine which graph view to be the more informative of a specific case, and subsequently fusion of features (which are important) would occur to extract the important patterns into a single form. The merged characteristics are delivered to a probability output layer which approximates the probability of fraud and this rating is utilized to produce the final fraud alert output (fraud or normal). The total design promotes early detection due to shared learning on various relationship perspectives rather than one transaction feature space.

IV. PROPOSED ALGORITHM

Algorithm 1 Multi-View Graph Attention for Early Fraud Detection

Step 1: Input: $\mathcal{T}$, *views* $\{1, \dots, V\}$, *threshold* τ, *layers* L

Step 2: Construct graphs $\mathcal{G}^{(v)} = \left(\mathcal{V}, \mathcal{E}^{(v)}, X^{(v)}\right), v \in \{1, \dots, V\}$

Step 3: Initialize node features $h_i^{(v,0)} \leftarrow X_i^{(v)}$

Step 4: for $l = 0$ *to* $L - 1$ *do*

Step 5: for each view v *do*

Step 6: $e_{ij}^{(v,l)} \leftarrow LeakyReLU\left(a^{(v,l)\top}\left[W^{(v,l)}h_i^{(v,l)} \| W^{(v,l)}h_j^{(v,l)}\right]\right)$

Step 7: $\alpha_{ij}^{(v,l)} \leftarrow \frac{exp\left(e_{ij}^{(v,l)}\right)}{\sum_{k \in \mathcal{N}_j^{(v)}} exp\left(e_{ik}^{(v,l)}\right)}$

Step 8: $h_i^{(v,l+1)} \leftarrow \sigma\left(\sum_{j \in \mathcal{N}_i^{(v)}} \alpha_{ij}^{(v,l)} W^{(v,l)} h_j^{(v,l)}\right)$

Step 9: end for

step 10: end for

Step 11: $z_i^{(v)} \leftarrow h_i^{(v,L)}, \beta_i^{(v)} \leftarrow \frac{exp\left(q^\top tanh\left(U z_i^{(v)}\right)\right)}{\sum_{r=1}^{V} exp\left(q^\top tanh\left(U z_i^{(r)}\right)\right)}$

Step 12: $z_i \leftarrow \sum_{v=1}^{V} \beta_i^{(v)} z_i^{(v)}, s_i \leftarrow \sigma(w_s^\top z_i + b_s), \hat{y}_i \leftarrow \mathbb{I}(s_i \geq \tau)$

Step 13: $\mathcal{L}_{cls} \leftarrow -\frac{1}{N}\sum_{i=1}^{N}[y_i log\, s_i + (1 - y_i) log(1 - s_i)]$

Step 14: $\mathcal{L}_{reg} \leftarrow \lambda \sum_{v,l} \left(\left\|W^{(v,l)}\right\|_F^2 + \left\|a^{(v,l)}\right\|_2^2\right), \mathcal{L} \leftarrow \mathcal{L}_{cls} + \mathcal{L}_{reg}$

Step 15: Update $\Theta = \{W^{(v,l)}, a^{(v,l)}, U, q, w_s, b_s\}$ *by minimizing* $\mathcal{L}$*; output* $\{s_i, \hat{y}_i\}$

The algorithm states the knowledge of the system to detect fraud in its early stages by learning a combination of relationship in the transaction data perspectives.

First, it transforms the input records into various graphs, which indicate various relations such as links between users and merchants, sequence of transactions across time, and relationships between devices. Each graph view has the model learning which neighbors are more relevant, with the help of an attention mechanism and subsequently uses the information of the most relevant ones to construct the powerful representation of each user or transaction. Once this is repeated a number of times it generates one learned representation per view. A second attention step is then involved in determining the degree of importance to be applied to each view, and in combining them into a single, final combined representation. This end-representation gets to a prediction layer that gives a score on the likelihood of fraud and depending on a set threshold the system will declare the case as an instance of normal or a fraud. Lastly, the model is trained in a way that compares its predictions against the true labels and sets the internal parameters in such a way that through the learning process the probability of fraud becomes less biased with time.

V. EXPERIMENTAL SETUP

The input transaction records in the proposed experimental set up are preprocessed to initially manage missing values, encode categorical values and normalize continuous values. Based on the cleaned data, several relational views are built and input into the proposed multi-view graph attention model. The model can be trained by using a supervised split (train, validation, test) and by optimizing the hyper parameters, like the number of attention layers, the size of the hidden embedding, the learning rate and the threshold at which the fraud decision is made by using the validation set. Finally, balanced-learning appropriate measures such as Precision, Recall, F1-score, ROC-AUC, and PR-AUC are used to report the performance that can foster a reasonable assessment of same-side fraud detection performance under skewed class distributions. In the case of the given dataset, one of the popular benchmarks is the Fraud Detection dataset by IEEE-CIS available on kaggle [19] that includes real-world features of online transactions and labels on fraud instances.

VI. RESULTS AND DISCUSSION

The proposed Multi-View Graph Attention framework can be effectively examined in terms of five major parameters of result, including Precision, Recall, F1-score, ROC-AUC, and PR-AUC. Precision refers to the ratio of the forecasted fraud cases that are fraudulent that would assist in comprehending the effectiveness of the model in preventing unnecessary false alarms. Recall, a fundamental parameter in early detection of fraud, aptly describes the percentage of fraud cases identified by the model that are frauds successfully identified; thus it is a vital parameter in the early detection of a fraud case provided a false indication of a fraud might be very costly. The F1 -score is a balanced measure and can be considered the sum of Precision and Recall, which can be beneficial in cases where the distribution of the classes is unequal. ROC-AUC judges the capacity of the model to assign transactions to frauds or genuine at different decision thresholds providing a general assessment of separability. Fraud datasets are generally highly imbalanced and, therefore, PR-AUC is especially significant as it places more emphasis on the Precision-Recall trade off of the minority fraud category. Collectively, these 5 parameters give a plausible and credible measurement of the effectiveness of the model in fraud detection at an early stage with accuracy.

Precision

Table 1. Precision Comparison Across Epochs

Epoch	Proposed MVGA	GAT Baseline	GCN Baseline
10	0.81	0.76	0.71
20	0.84	0.79	0.74
30	0.88	0.82	0.76
40	0.89	0.84	0.78
50	0.91	0.84	0.79
60	0.92	0.85	0.80

As seen by the accuracy values in table 1, the proposed MVGA model has the highest reliability in creating fraud alerts among all the epochs which progressively rise between 0.81 and 0.92. This means that the amount of transactions marked as fraud is higher so that more false alarms are eliminated. GAT baseline increases to 0.85 and has a weak plateau between epoch 40 and 50 indicating that no more gains of significant changes are obtained after that. The GCN base is the lowest (0.71 to 0.80), showing weaker discrimination with no stronger attention-based view fusion which used simple convolutional aggregation only.

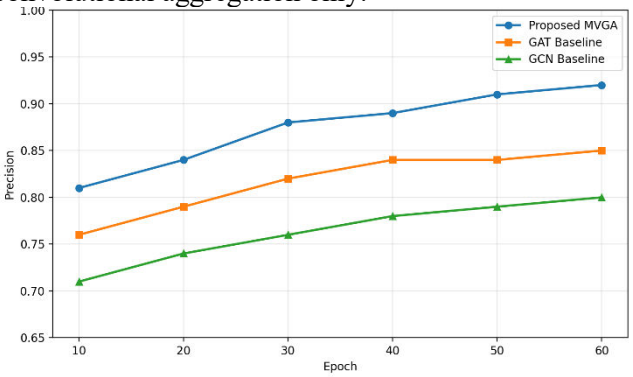

Fig 2. Precision vs Epoch for Proposed MVGA, GAT Baseline, and GCN

Figure 2 above is the precision curve that depicts that

the offered MVGA model has the highest level of precision throughout the training period, as it increases in levels, starting with 0.81 and reaching up to 0.92, which means that the proposed model allows reducing the false fraud alarms and enhancing the reliability of the cases of the detected fraud. The GAT baseline increases but there is an obvious plateau in later epochs, and it seems that further learning advantage is limited beyond a certain point. The lowest base value is of 0.71-0.80 which represents a weaker capacity to isolate the real patterns of fraud when the simpler neighbourhood aggregation is employed only.

Recall

Table 2. Recall Comparison Across Epochs

Epoch	Proposed MVGA	GAT Baseline	GCN Baseline
10	0.74	0.73	0.69
20	0.79	0.76	0.71
30	0.83	0.77	0.72
40	0.86	0.79	0.73
50	0.87	0.81	0.74
60	0.88	0.82	0.74

Table 2 shows the recall performance that indicates how the ability of the proposed MVGA to detect more true fraud cases improved with a steady increase, 0.74 to 0.88 result and this would be necessary in detecting frauds early when false alarms could be expensive. The GAT baseline (0.73 to 0.82) also compliments the proposed model but always worse which shows that it does not capture more of the fraud cases given the same course of training. The GCN baseline exhibits the lowest recall and levels off to 0.74 in subsequent epochs thus indicating that it cannot readily learn finer relational statements and profile more by noisy neighbourhood or low relational expressiveness.

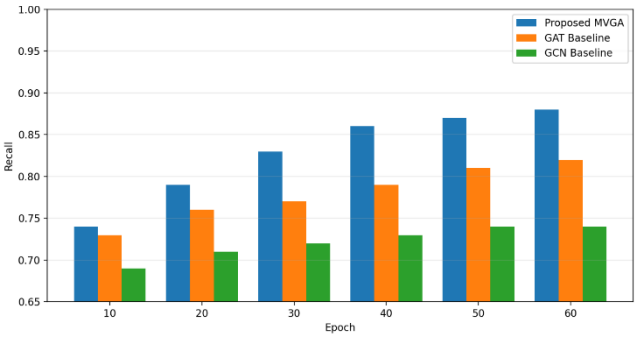

Figure 3. Recall vs Epoch for Proposed MVGA, GAT Baseline, and GCN Baseline

Figure 3 above shows that the proposed MVGA has a greater potential of capturing the true fraud cases with the recall bar chart steadily rising between 0.74 to 0.88. The GAT baseline remains within 0.73 0.82 implying there are more cases of fraud not detected than the model proposed. The GCN base displays the most insignificant recall and nearly stays constant at the end (ending at 0.74), which infers that it does not favour the necessity to learn intricate relational cues required to identify fraud at its early stages and maintain a consistent rate.

F1-Score

Table 3. F1-Score Comparison Across Epochs

Epoch	Proposed MVGA	GAT Baseline	GCN Baseline
10	0.775	0.742	0.698
20	0.812	0.765	0.718
30	0.853	0.792	0.735
40	0.848	0.801	0.747
50	0.889	0.818	0.758
60	0.903	0.828	0.765

The comparison of the F1-score in table 3 indicates that the proposed MVGA has the optimal overall balance between precision and recall, and this score is increased, 0.775 to 0.903. The small dip at epoch 40 (0.848) may be assumed to be a temporary trade-off accompanying the optimization, however, the measure is soon restored and continues raising, which is a sign of stable learning. Efforts to enhance the GAT baseline (0.742 to 0.828) at single-view (or less expressive) fusion decrease performance disequilibrium. The GCN baseline is the lowest (0.698 to 0.765) which proves that weaker graph aggregation provides less powerful combined effectiveness in the case of class imbalance.

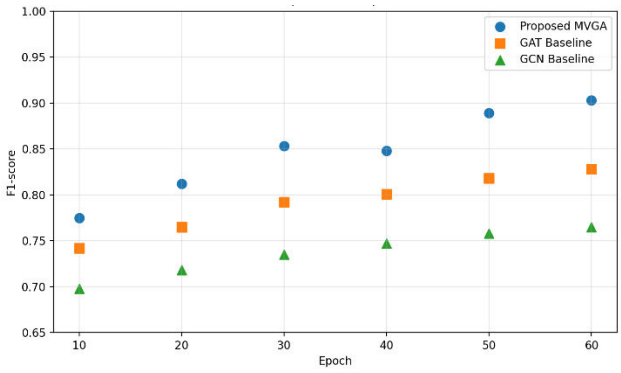

Figure 4. F1-score vs Epoch for Proposed MVGA, GAT Baseline, and GCN Baseline

The F1-score scatter plot at figure 4 proves correct the idea that the proposed MVGA is the best balance in terms of precision and recall, and that at epoch 60, the best score is 0.903. The slight dip at around epoch 40 is a transient trade off during optimization, however, the model will recover fast and proceeds to improve giving stable learning behaviour. The GAT baseline enhances

slowly but evidently at a lower value than the proposed model whereas the GCN baseline is lowest, rendering further to the fact that attention-based multi-view fusion enhances balanced fraud detection performance.

ROC-AUC

Table 4. ROC-AUC Comparison Across Epochs

Epoch	Proposed MVGA	GAT Baseline	GCN Baseline
10	0.88	0.82	0.79
20	0.92	0.86	0.81
30	0.935	0.875	0.83
40	0.942	0.885	0.845
50	0.944	0.892	0.852
60	0.945	0.895	0.856

The values of ROC-AUC in the table 4 as to the suggested MVGA values the best threshold-independent fraction between the fraud class and the normal class with a value of 0.88 and it then saturates a somewhat indicating that when powerful discriminative treatments are learnt it reaches the strongest threshold-independent separability. There are the GAT baseline (0.82 to 0.895), which exhibits moderate improvements but lower terminus of the separability, which means weak ranking of the fraud risk at various thresholds. The GCN baseline (0.79 to 0.856) is the lowest meaning that the capacity to differentiate classes is low in the presence of fraud accusations in a persistent manner and in the presence of relationally tricky fraud patterns.

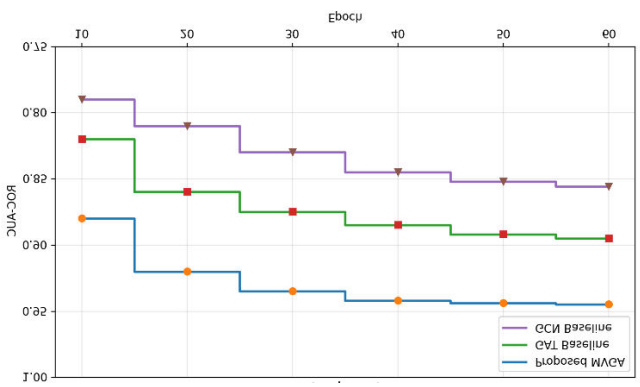

Figure 5. ROC-AUC vs Epoch for Proposed MVGA, GAT Baseline, and GCN Baseline

The status of the proposed MVGA is shown in the step plot in figure 5 of ROC-AUC where the class separability is highest in threshold values, and the separation is rapid followed by a saturation point, of approximately 0.945, as typical when the model has acquired unsettled discriminative structure. The GAT is moderately improved at the baseline of 0.895 and the GCN baseline ends in 0.856, which is less separated and ranked in terms of fraud risk. Altogether, the greater ROC-AUC of the suggested approach indicates better threshold-free discrimination between fraud and non-fraud transactions.

PR-AUC

Table 5. PR-AUC Comparison Across Epochs

Epoch	Proposed MVGA	GAT Baseline	GCN Baseline
10	0.70	0.64	0.61
20	0.74	0.69	0.65
30	0.77	0.73	0.66
40	0.81	0.76	0.70
50	0.87	0.78	0.72
60	0.89	0.80	0.73

The PR-AUC results in table 5 give the clearest image under imbalanced classes and indicate that the proposed MVGA produces the best precision-recall tradeoff to the fraud class and improves with added epochs (0.70 to 0.89) with significant increase in the later epochs (0.81 to 0.87 to 0.89). GAT baseline also raises between 0.64 and 0.80, though not as high as the proposed model, which suggests that it is still not very precise at higher recall. The GCN baseline advances gradually (0.61 to 0.73) and it validates the fact that the absence of intense multi-view attention and powerful fusion makes the model most problematic on the minority fraud category, which is the very necessity in the field of fraud detection.

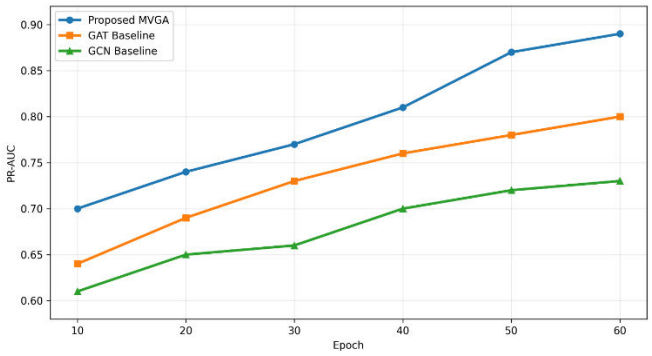

Figure 6. PR-AUC vs Epoch for Proposed MVGA, GAT Baseline, and GCN Baseline (Area Plot, Revised Values)

Figure 6 PR-AUC area plot has the most significant change in the case of class imbalance, where the suggested MVGA increases to 0.89 by a significant increment in the end stage, which suggests that minority fraud patterns are better learned as the fusion stabilizes. The GAT baseline jumps to 0.80 yet it is lower than it implies that it is capable of being equally strong in higher recall rates of the fraud class. The GCN

backbone only increases gradually to 0.73, which proves that in the absence of multi-view attention and powerful fusion, the performance of the minority class of frauds is also low. [15] explained that explores augmented reality (AR) as a vital enabler in fintech for sustainable finance, enhancing financial literacy, supporting better-informed investment choices, and increasing transparency in green finance. [16] explained that Green Policies and Augmented Reality in InsurTech: Catalysts for Sustainability and Innovation offers a timely and in-depth exploration of the convergence between environmental sustainability and advanced technologies within the insurance sector.

VI. CONCLUSION

According to the received numerical results, the proposed Multi-View Graph Attention (MVGA) architecture shows significant steady gains over both baselines on all five parameters. With the final epoch, MVGA gets Precision = 0.92 and Recall = 0.88 which proves that it does not only decrease false fraud alerts but also traces a bigger portion of real frauds. This performance merit is demonstrated in better F1-score = 0.903, which provides greater overall performance when the imbalance is in the class. Moreover, the model achieves ROC-AUC = 0.945 and presents a better threshold-independent separability between fraudulent and legitimate transactions, and PR-AUC = 0.89 is the best minority-class tradeoff, which is indeed very high when compared to GAT baseline (0.80) and GCN baseline (0.73). These values, overall, confirm that Representations with multi-view relation together with attention-based fusion generate more fraud-specific embedding and enhance the accuracy and stability of early fraud detection in comparison to single-views on graphs.

To improve upon this, in future developments, the framework may be expanded to include explicit time-based learning (temporal graph attention, event-driven sequence modules) to promote the early warning performance to a higher level (than the current 0.88 version of Recall) without compromising on Precision. The PR-AUC pattern shows continued improvement in later epochs of its minority-class performance and therefore with the inclusion of imbalance-sensitive losses (cost-sensitive learning or focal variants) and adaptive sampling, PR-AUC can be further improved beyond 0.89 and false positives cut down. The model, by adding explainability elements e.g., view-level-contribution-score and subgraph-rationales, will be more actionable by the fraud analysts. Shortly but certainly, measuring on bigger and cross domain transaction graphs with concept drift treating and online updating will assist in reinforcing generalization and bring ROC-AUC nearer to 0.945 even when fraud patterns change over time.

REFERENCES

[1] Amebleh, J., Igba, E., & Ijiga, O. M. (2021). Graph-Based Fraud Detection in Open-Loop Gift Cards: Heterogeneous GNNs, Streaming Feature Stores, and Near-Zero-Lag Anomaly Alerts. *International Journal of Scientific Research in Science, Engineering and Technology*, 317–339. https://doi.org/10.32628/ijsrset214418

[2] Lin, J., Guo, X., Zhu, Y., Mitchell, S., Altman, E., & Shun, J. (2024). FraudGT: A Simple, Effective, and Efficient Graph Transformer for Financial Fraud Detection. 292–300. https://doi.org/10.1145/3677052.3698648

[3] Wang, C., Wang, M., Wang, X., Zhang, L., & Long, Y. (2024). Multi-Relational Graph Representation Learning for Financial Statement Fraud Detection. *Big Data Mining and Analytics*, 7(3), 920–941. https://doi.org/10.26599/bdma.2024.9020013

[4] Noble, C. C., & Cook, D. J. (2003). Graph-based anomaly detection. 631–636. https://doi.org/10.1145/956750.956831

[5] Wei, S., & Lee, S. (2024). Financial Anti-Fraud Based on Dual-Channel Graph Attention Network. *Journal of Theoretical and Applied Electronic Commerce Research*, 19(1), 297–314. https://doi.org/10.3390/jtaer19010016

[6] Hong, B., Lu, P., Xu, H., Lu, J., Lin, K., & Yang, F. (2024). Health insurance fraud detection based on multi-channel heterogeneous graph structure learning. *Heliyon*, 10(9), e30045. https://doi.org/10.1016/j.heliyon.2024.e30045

[7] Cheng, D., Wang, X., Zhang, Y., & Zhang, L. (2020). Graph Neural Network for Fraud Detection via Spatial-Temporal Attention. *IEEE Transactions on Knowledge and Data Engineering*, 34(8), 3800–3813. https://doi.org/10.1109/tkde.2020.3025588

[8] Tao, J., Lin, J., Zhang, S., Zhao, S., Wu, R., Fan, C., & Cui, P. (2019). MVAN: Multi-view Attention Networks for Real Money Trading Detection in Online Games. 2536–2546. https://doi.org/10.1145/3292500.3330687

[9] Dou, Y., Liu, Z., Sun, L., Deng, Y., Peng, H., & Yu, P. S. (2020). Enhancing Graph Neural Network-based Fraud Detectors against Camouflaged Fraudsters. *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 315–324. https://doi.org/10.1145/3340531.3411903

[10] Liu, Y., Ao, X., Qin, Z., Chi, J., Feng, J., Yang, H., & He, Q. (2021). Pick and Choose: A GNN-based Imbalanced Learning Approach for Fraud Detection. 3168–3177. https://doi.org/10.1145/3442381.3449989

[11] Wang, L., Li, P., Xiong, K., Zhao, J., & Lin, R. (2021). Modeling Heterogeneous Graph Network on Fraud Detection. 1959–1968. https://doi.org/10.1145/3459637.3482277

[12] Qin, Z., Liu, Y., He, Q., & Ao, X. (2022). Explainable Graph-based Fraud Detection via Neural Meta-graph Search. 4414–4418. https://doi.org/10.1145/3511808.3557598

[13] Ma, X., Wu, J., Xue, S., Yang, J., Zhou, C., Sheng, Q. Z., Xiong, H., & Akoglu, L. (2021). A Comprehensive Survey on Graph Anomaly Detection with Deep Learning. *IEEE Transactions on Knowledge and Data Engineering*, 35(1), 120–138. https://doi.org/10.1109/TKDE.2021.3071157

[14] Xu, F., Wang, N., Wu, H., Wen, X., Zhao, X., & Wan, H. (2024). Revisiting Graph-Based Fraud Detection in Sight of Heterophily and Spectrum. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(8), 9214–9222. https://doi.org/10.1609/aaai.v38i8.28773

[15] Christo Ananth, Priyanka Jain, Ankur Roy, "Designing Sustainable Fintech Solutions with Augmented Reality Integration", Augmented Reality and Sustainability Goals and Challenges, Auerbach Publications, Routledge, Taylor and Francis, USA, DOI: 10.1201/9781003584438-13,pp: 1-19

[16] Christo Ananth, Sonal Trivedi, Pushan Kumar Dutta, Anuj Kumar "Green Policies and Augmented Reality in InsurTech", Emerald Publishing, USA, Pages: 276

[17] Zhang, J., Xu, Z., Lv, D., Shi, Z., Shen, D., Jin, J., & Dong, F. (2024). DiG-In-GNN: Discriminative Feature Guided GNN-Based Fraud Detector against Inconsistencies in Multi-Relation Fraud Graph. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(8), 9323–9331. https://doi.org/10.1609/aaai.v38i8.28785

[18] Wang, J., Wen, R., Wu, C., Huang, Y., & Xiong, J. (2019). FdGars: Fraudster Detection via Graph Convolutional Networks in Online App Review System. *Companion Proceedings of The 2019 World Wide Web Conference*, 310–316. https://doi.org/10.1145/3308560.3316586

[19] https://www.kaggle.com/competitions/ieee-fraud-detection/data