

FraudCenGCL: Role-Aware Graph Contrastive Learning for Low-Homophily Fraud Detection

Seonkyu Lim^{*§}, Jeongwhan Choi^{†§} and Jaehoon Lee^{‡§¶}

*Korea Financial Telecommunications & Clearings Institute**, Seoul, South Korea

Korea Advanced Institute of Science and Technology (KAIST)[†], Daejeon, South Korea

LG AI Research[‡], Seoul, South Korea

sklim@kftc.or.kr^{*}, jeongwhan.choi@kaist.ac.kr[†], jaehoon.lee@lgresearch.ai[‡],

Abstract—Detecting financial fraud is essential for ensuring the security and reliability of financial systems. While graph neural networks (GNNs) have shown promise in modeling transaction networks, they often struggle to capture subtle structural patterns associated with fraudulent behavior. To address this challenge, we propose **FraudCenGCL**, a graph contrastive learning framework that enhances fraud detection by integrating account-level features with graph-based centrality measures. By combining behavioral attributes with structural indicators such as degree, closeness, and betweenness centrality, our framework enriches node representations that reflect both transactional and topological characteristics. We evaluate **FraudCenGCL** on real-world interbank transfer data from the housing finance information network (HOFINET) in South Korea. Experimental results demonstrate that our framework consistently outperforms existing GCL baselines across six representative backbones and multiple evaluation metrics. These findings demonstrate the effectiveness of incorporating structural information to improve the performance and practicality of fraud detection systems (FDS).

Index Terms—fraud detection, graph contrastive learning, bank transfers, graph analysis

I. INTRODUCTION

Financial fraud poses an ongoing threat to the security and trustworthiness of financial systems, often exploiting the structure of transaction flows to evade detection. Many fraudulent behaviors do not manifest through isolated transactions but rather emerge from coordinated patterns across account networks. As such, effective fraud detection requires analyzing not only transactional behaviors but also the structural roles of accounts within the transaction graph. However, traditional rule-based systems and conventional machine learning models typically treat transactions as independent events, failing to capture the positional and relational context that is critical for identifying complex fraud patterns [1]–[5].

Graph neural networks (GNNs) have been widely adopted for analyzing transaction data by modeling accounts and transfers as nodes and edges, respectively [6]–[9]. In fraud detection tasks, GNN-based models learn account relationships through message passing and have shown promise in capturing transactional dependencies. However, standard GNNs often

[§]These authors contributed equally to this research.

[¶]This work was carried out while he was at Yonsei University.

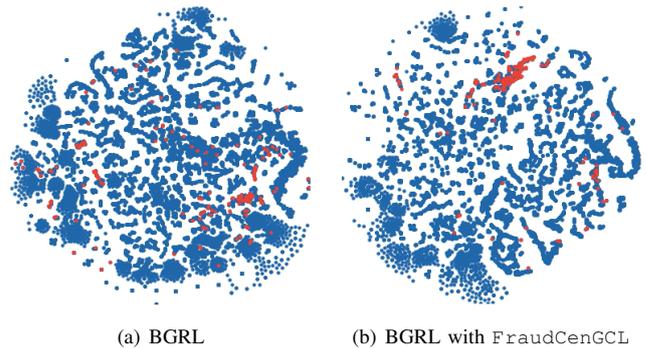


Fig. 1. Effect of centrality-guided augmentation in **FraudCenGCL**. The embeddings are learned using BGRL [12] on our interbank transfer dataset. (a) Without centrality features, the embeddings show poor class separation, with fraudulent (red) and benign (blue) nodes highly entangled. (b) With degree centrality used as an augmented view in **FraudCenGCL**, fraudulent nodes form more compact clusters and exhibit clearer separation from benign ones.

struggle with representing diverse structural roles of nodes and may fail to differentiate subtle but critical behaviors in fraud-related interactions. In particular, financial transaction networks often exhibit low homophily, where fraudulent accounts are more frequently connected to benign accounts than to each other. This camouflage behavior undermines the neighborhood similarity assumption of GNNs, making them vulnerable in fraud detection settings [10], [11]. These limitations become more pronounced in financial networks, where fraudulent patterns often span multiple hops or emerge through complex topological configurations.

Contrastive learning has recently emerged as a promising self-supervised approach for capturing complex patterns in financial transactions [13]–[15]. By maximizing agreement between similar instances while pushing dissimilar ones apart, contrastive learning enables models to learn discriminative embeddings without requiring extensive label supervision. This property makes it particularly effective in financial domains, where labeled fraud cases are scarce and subtle behavioral differences are critical. Moreover, contrastive learning is well-suited for graph-structured data, as it can enhance representation learning by leveraging both local and global structural information.

Building on the success of contrastive learning, recent works have combined it with GNNs to improve graph representation learning [16]–[20]. These graph contrastive learning (GCL) models typically learn embeddings by maximizing consistency between multiple graph representations, including augmented views (e.g., GRACE [20], MVGRL [18]) and corrupted inputs (e.g., DGI [12]). However, most existing approaches focus on either node-level attributes or global graph structures, often neglecting the structural roles of individual nodes. This limitation is critical in fraud detection, where the position and influence of an account within the transaction network often signal suspicious activity.

As shown in Fig. 1, the embeddings produced by the BGRL model [12] without centrality augmentation (a) exhibit significant overlap between fraudulent and benign nodes, indicating limited class separation. In contrast, when degree centrality is introduced as an additional contrastive view (b), the embeddings form more compact and clearly separated clusters. These results suggest that incorporating structural role information such as centrality improves the discriminative power of node representations. Incorporating these roles into contrastive learning provides a promising direction for addressing the unique challenges of fraud detection in low-homophily financial transaction networks.

In this paper, we propose **fraud** detection with **Centrality** enhanced **Graph Contrastive Learning (FraudCenGCL)**, a novel framework that improves fraud detection in financial transaction networks by incorporating account centrality features into contrastive learning. *FraudCenGCL* enriches node representations by combining behavioral features with graph-derived centrality measures such as degree, closeness, and betweenness. This design enables the model to capture both transactional and structural patterns of accounts within existing contrastive learning pipelines, without requiring additional labels or architectural changes.

We evaluate our framework using the housing finance information network (HOFINET) dataset, a large-scale real-world dataset of interbank transfers operated by the Korea Financial Telecommunications and Clearings Institute (KFTC), the clearing house of South Korea. The HOFINET dataset is proprietary and subject to strict regulatory constraints, which makes it inaccessible to the public and rarely available for research use. Unlike synthetic or benchmark datasets, it directly reflects production-level financial transaction flows across institutions, offering a uniquely realistic environment for evaluating fraud detection models.

Experimental results show that *FraudCenGCL* consistently outperforms baseline GCL models across six representative backbones, with significant gains in precision, F1-score, and AUPRC for detecting fraudulent entities.

The main contributions of this work are summarized as follows:

- We propose *FraudCenGCL*, a contrastive learning framework that integrates centrality-based structural features with behavioral features, thereby improving the rep-

resentation of accounts in financial transaction networks (Section IV-B, IV-E, IV-F).

- Our framework supports multiple centrality measures (e.g., degree, closeness, betweenness) and can be applied to diverse GCL backbones such as BGRL, GRACE, GBT, DGI, and MVGRL without modifications to their core architectures (Section IV-C, IV-D).
- Through extensive experiments on real-world interbank transfer data, we show that *FraudCenGCL* achieves consistent performance gains across all backbones and evaluation metrics. These results indicate that incorporating structural role information improves the practical effectiveness of fraud detection models in low-homophily financial networks (Section IV-B, IV-F).

II. RELATED WORK

Graph-based models have been widely applied to analyze relationships between accounts in financial transaction networks and identify fraudulent entities from legitimate ones. Representative graph neural networks (GNNs) such as GCN [6], GAT [7], and GraphSAGE [8] have shown strong performance in learning node connectivity and importance. Building on these foundations, GNN-based fraud detection models have been proposed to capture disguised behaviors and temporal patterns in financial data [21]–[25].

Contrastive learning is effective in learning similarities and differences between data points, making it suitable for detecting subtle fraud patterns in financial datasets. SimCLR uses data augmentations to bring samples of the same class closer in embedding space [13], while MoCo improves training stability through memory-based negative sampling [14]. SupCon combines contrastive and supervised learning [15], and PIRL [26] and BYOL [27] learn representations without relying on negative samples. These methods are promising for fraud detection, where labeled data is limited and behavioral distinctions are subtle.

Recent works have integrated contrastive learning into graph-based models to enhance representation quality. DGI [16] captures global-local consistency. InfoGraph [17] aligns whole-graph and subgraph representations. MVGRL [18] and GRACE [19] apply multi-view and feature perturbation strategies, respectively. GraphCL [20] focuses on structural diversity via augmentations. BGRL [12] and GBT [28] further improve contrastive learning by avoiding negatives or reducing redundancy.

However, most existing contrastive graph models emphasize node features or global structures, while overlooking the structural roles of individual nodes. Prior studies have also shown that conventional GNNs face difficulties in disassortative or low-homophily graphs, where neighborhood aggregation fails to capture long-range dependencies [11]. This challenge is particularly important in financial transaction networks, where the positional and structural influence of accounts often reflects suspicious behavior. To address this gap, we propose *FraudCenGCL*, which integrates centrality measures into

the contrastive learning process to jointly model behavioral attributes and structural roles of accounts.

III. PROPOSED METHOD

We propose `FraudCenGCL`, a contrastive learning framework that jointly processes behavioral and structural features for fraud detection in financial transaction networks. It incorporates account centrality measures as an additional structural view that complements behavior-driven features. As shown in Figure 2, the framework processes both aggregated transaction features (from tabular data) and graph-structured information (from graph data) to generate node embeddings that better distinguish fraudulent accounts from legitimate ones.

A. *FraudCenGCL Framework*

`FraudCenGCL` consists of four key components: (1) graph construction from interbank transfer data, (2) dual-source feature engineering that includes both behavioral and centrality-based features, (3) GNN encoder that performs contrastive learning over representations derived from heterogeneous features, and (4) optimization via a contrastive loss function. These components jointly enable the model to learn node representations that capture both transactional patterns and topological roles, which are crucial for downstream fraud detection.

1) *Graph Construction*: We construct a directed graph where each node represents a unique account and each edge corresponds to a fund transfer between accounts. This transformation from tabular data captures the transactional topology, including the directionality and intensity of account-to-account interactions, enabling structural analysis of financial flows.

2) *Graph Node Features*: We extract two types of features for each account node. First, behavioral features are derived from aggregated transaction-level statistics, including frequency, average amount, and variance. These features describe the individual activity patterns of accounts. Second, we compute structural features using centrality measures such as degree, closeness, and betweenness centrality. These features quantify the importance and position of each node within the transaction graph.

To ensure compatibility for joint representation learning, we apply separate projection layers to map behavioral and structural features into a shared latent space. These layers align feature dimensions and semantics, enabling joint representation learning over distinct semantic signals.

3) *GNN Encoder*: The GNN encoder acts as the core component of our framework, processing both account features and centrality features to generate embeddings. This encoder uses the graph structure to learn relationships between nodes and their neighborhoods. Through this process, each account node learns a representation that reflects not only its inherent characteristics but also its structural position within the transaction network.

Our framework implements a contrastive learning approach based on separate feature channels that generates distinct embeddings for account and centrality features. These feature

sources provide complementary semantic signals about the accounts: one reflecting behavioral patterns and the other capturing structural roles within the transaction network. The resulting embeddings are then concatenated to form a unified representation for each node, which is subsequently used for contrastive optimization and fraud classification.

4) *Loss Functions and Training*: To train the model, we apply a contrastive loss that encourages embeddings from distinct feature representations of the same account to be similar, while separating those of different accounts. This objective shapes the embedding space to reflect structural and behavioral distinctions, facilitating fraud classification.

The learned embeddings are then used in a downstream classification task. By incorporating both transactional behaviors and structural roles, the resulting node representations enable precise identification of suspicious accounts within financial networks.

B. *Graph Centrality Analysis*

Centrality measures quantify the relative importance of nodes in a graph. In the context of financial networks, they help identify influential accounts based on their position and connectivity [29]–[32]. We focus on three representative measures: degree, closeness, and betweenness centrality [33]. For this analysis, we model the transaction network as a directed graph, where each node corresponds to an account and edges represent fund transfers from withdrawal to deposit accounts.

To formally define the centrality measures, let $G = (V, E)$ be a directed transaction graph where $v \in V$ denotes an account and $(u, v) \in E$ represents a fund transfer from u (withdrawal) to v (deposit). We define the centrality measures as follows:

- **Degree centrality** counts the total number of incoming and outgoing transactions associated with a node:

$$\text{Degree_Cen}(v) = \sum_{u \in V} (\mathbf{1}_{(u,v) \in E} + \mathbf{1}_{(v,u) \in E}), \quad (1)$$

- **Closeness centrality** measures how efficiently a node can reach and be reached by other nodes:

$$\text{Closeness_Cen}(v) = \sum_{u \in V} \left(\frac{1}{d(u,v)} + \frac{1}{d(v,u)} \right), \quad (2)$$

where $d(u, v)$ is the length of the shortest path from u to v .

- **Betweenness centrality** reflects how often a node lies on the shortest paths between other pairs of nodes:

$$\text{Betweenness_Cen}(v) = \sum_{u_1, u_2 \in V} \mathbf{1}_{v \in p(u_1, u_2)}, \quad (3)$$

where $p(u_1, u_2)$ denotes the shortest path from u_1 to u_2 .

By formally defining these centrality measures and integrating them into our framework, `FraudCenGCL` is equipped to capture both the behavioral signatures and structural roles of accounts.

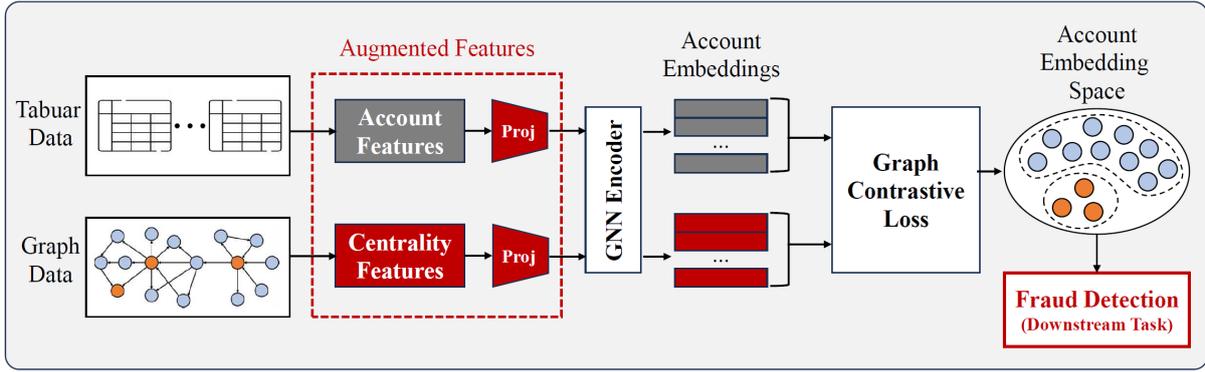


Fig. 2. Workflow of the FraudCenGCL framework, integrating centrality and account features into contrastive learning by combining multiple semantic sources in fraud detection. The framework leverages tabular and graph data to generate embeddings that capture both behavioral and structural patterns of accounts, enhancing the discrimination of fraudulent accounts.

IV. EXPERIMENTS

In this section, we evaluate the proposed FraudCenGCL framework by addressing the following research questions:

- **RQ.1** How does FraudCenGCL perform compared to existing GCL baselines in fraud detection?
- **RQ.2** What are the structural characteristics of fraudulent accounts based on centrality measures?
- **RQ.3** How do different centrality measures contribute to performance gains?
- **RQ.4** How does the choice of GNN encoder affect model performance?
- **RQ.5** How does FraudCenGCL shape the structure of the learned embedding space?

A. Experimental Setting

All experiments were conducted using PYTHON 3.8 and PYTORCH 1.12 with CUDA 11.4 support. We utilized NETWORKX 2.8.4 for graph processing, and all models were trained on an NVIDIA TESLA T4 GPU.

1) *Datasets*: We use the housing finance information network (HOFINET) dataset, which contains real-world interbank transfer transactions. Table II summarizes the dataset statistics, and Table I provides the details of its feature fields.

Each node in the transaction graph corresponds to a unique account, defined as a combination of a bank code and an account number. A directed edge is created from a withdrawal account (source) to a deposit account (target) for each transaction, resulting in a multi-edge directed graph that reflects the structure of the financial transfer network.

Fraudulent labels are assigned based on institution-reported suspicion indicators included in the dataset. An account is labeled as fraudulent if it appears as either sender or receiver in any transaction marked suspicious. All other accounts are considered benign. This strategy captures both direct and indirect participation in potentially fraudulent activity.

We construct node features by aggregating transaction-level information into account-level representations. Behavioral features include statistical descriptors such as mean, maximum,

TABLE I
THE DETAILS OF THE HOFINET DATASET.

Field	Description
Transaction Date	The date of the transaction.
Transaction Time	The time of the transaction.
Amount	The amount of money.
Media Type	Transaction medium (e.g., Mobile).
Fund Type	Type of funds (e.g., salary).
Withdrawal Bank Code	Bank identification code.
Withdrawal Account Number	Account number for withdrawal.
Deposit Bank Code	Bank identification code.
Deposit Account Number	Account number for deposit.
Suspicious Indicator	Information indicating suspicious activity.

TABLE II
STATISTICS OF HOFINET TRANSFER TRANSACTION DATASET.

Time Range	# Accounts	# Transfers	# Suspicious
Mar. 2024	30,106	145,023	253 (0.1745%)

standard deviation, transaction count, and entropy-based measures for fields such as transaction amount, fund type, and media type. Structural features are computed using graph-based centrality measures, specifically degree, closeness, and betweenness centrality, which capture each account's position and importance in the network topology.

2) *Homophily Analysis*: To better understand the structural properties of HOFINET, we measure the homophily ratio ϕ , defined as:

$$\phi = \frac{|\{(u, v) \in E : y_u = y_v\}|}{|E|}, \quad (4)$$

which quantifies the proportion of edges that connect nodes with the same class label [34]. The dataset exhibits an imbalanced distribution: benign accounts show a high homophily ratio ($\phi \approx 0.99$), while fraudulent accounts display a much lower ratio ($\phi \approx 0.41$). This indicates that fraudulent accounts are more frequently connected to benign accounts rather than to each other, demonstrating the low-homophily nature of real-world fraud networks. Such characteristics highlight why

TABLE III

PERFORMANCE EVALUATION OF THE FRAUDCENGL FRAMEWORK ACROSS VARIOUS GCL BACKBONES. ALL METRICS ARE COMPUTED ON THE FRAUD (POSITIVE) CLASS. IMPROV. (%) INDICATES THE RELATIVE IMPROVEMENT OF THE MODEL WITH FRAUDCENGL OVER THE ORIGINAL BACKBONE.

Model	Precision	Recall	F1-score	AUROC	AUPRC
BGRL	0.1177	0.4211	0.1839	0.7514	0.0852
+ FraudCenGCL	0.3214	0.6207	0.4235	0.8533	0.3096
Improv.	173.21%	47.41%	130.29%	13.57%	263.38%
DGI-IND	0.1453	0.7083	0.2411	0.8562	0.3217
+ FraudCenGCL	0.3177	0.8182	0.4576	0.9144	0.5986
Improv.	118.62%	15.51%	89.79%	6.80%	86.12%
DGI-TRS	0.0382	0.7143	0.0725	0.7427	0.1770
+ FraudCenGCL	0.1849	0.7857	0.2993	0.8738	0.4475
Improv.	384.33%	10.00%	313.08%	17.66%	152.78%
GBT	0.0552	0.7083	0.1024	0.8690	0.2352
+ FraudCenGCL	0.4314	0.7333	0.5432	0.9293	0.5084
Improv.	681.61%	3.53%	430.43%	6.94%	116.17%
GRACE	0.0178	0.7917	0.0349	0.7410	0.0696
+ FraudCenGCL	0.1539	0.9565	0.2651	0.9751	0.3201
Improv.	763.36%	20.82%	660.36%	31.60%	359.59%
MVGRL	0.0800	0.3077	0.1270	0.6655	0.1223
+ FraudCenGCL	0.1774	0.8800	0.2953	0.9458	0.3230
Improv.	121.76%	186.00%	132.56%	42.12%	164.14%

GNNs relying on neighborhood similarity face difficulties in this domain, and motivate the design of FraudCenGCL.

3) *Baselines*: To assess the effectiveness of FraudCenGCL, we evaluate it on six widely used graph contrastive learning (GCL) models by applying our framework on top of their architectures:

- 1) **BGRL** [12]: A bootstrapped method that aligns node representations between original and perturbed graphs without using negative samples.
- 2) **DGI** [16]: A mutual information-based method that contrasts local patch embeddings with a global summary vector using a discriminator. The original version was transductive (DGI-TRS), and an inductive (DGI-IND) variant with sampling was later introduced for large graphs.
- 3) **GBT** [28]: A Barlow Twins-based model that reduces redundancy in embeddings across augmented views.
- 4) **GRACE** [20]: A dual-view contrastive method that uses feature masking and edge dropping for augmentation.
- 5) **MVGRL** [18]: A multi-view learning approach that contrasts node and graph-level representations across structurally different views.

4) *Evaluation Metrics*: We evaluate model performance using five widely used metrics for binary classification, particularly suitable for imbalanced data:

- **Precision**: The proportion of correctly predicted fraud cases among all instances predicted as fraud. High precision indicates a low false positive rate.

- **Recall**: The proportion of actual fraud cases that are correctly identified by the model. High recall reflects the model’s sensitivity to positive cases.
- **F1-score**: The harmonic mean of precision and recall, providing a balanced evaluation between false positives and false negatives.

We report all three metrics only for the fraud (positive) class to reflect model performance on rare but critical cases.

- **AUROC** (Area Under the Receiver Operating Characteristic Curve): A threshold-independent metric that quantifies the model’s ability to rank fraud cases above non-fraud cases.
- **AUPRC** (Area Under the Precision-Recall Curve): A ranking-based metric that emphasizes the model’s ability to identify rare positive instances, making it especially informative in highly imbalanced settings where AUROC may be misleading.

All metrics range from 0 to 1, with higher values indicating better performance.

B. Performance Comparison (RQ.1)

Table III shows the performance of FraudCenGCL applied to six representative GCL backbones: BGRL, DGI-IND, DGI-TRS, GBT, GRACE, and MVGRL. In all cases, FraudCenGCL improves performance in fraud detection.

FraudCenGCL achieves substantial gains across all evaluation metrics. F1-score improvements range from 89.79% (DGI-IND) to 660.36% (GRACE), while AUPRC gains reach 359.59% (GRACE) and 263.38% (BGRL). Precision increases markedly for GRACE (+763.36%) and GBT (+681.61%), indicating that the framework effectively reduces false positives. Recall gains are most pronounced for MVGRL (+186.00%) and GRACE (+20.82%), showing that FraudCenGCL enhances the ability to recover true fraud cases. In terms of AUROC, all models improve, with MVGRL (+42.12%) and GRACE (+31.60%) showing the largest increases.

Lightweight backbones such as GRACE and GBT show the greatest relative improvements, with F1-score increases of four to six times compared to their baselines. By contrast, inductive models such as DGI-IND, which exhibit stronger baseline results, achieve smaller but still meaningful improvements. Overall, FraudCenGCL enhances the precision-recall balance and ranking capability across diverse GCL architectures, demonstrating robustness in fraud detection under varying model settings.

C. Account Centrality Analysis (RQ.2)

Figure 3 shows the distributions of degree, closeness, and betweenness centrality for accounts, separated by fraudulent and benign classes. The box plots reveal distinctive structural patterns between the two groups, providing insights into how fraudulent accounts are positioned differently within the transaction network.

For degree centrality, benign accounts show a long-tailed distribution with a few extreme hubs (maximum about 0.41), whereas fraudulent accounts never exceed a much smaller

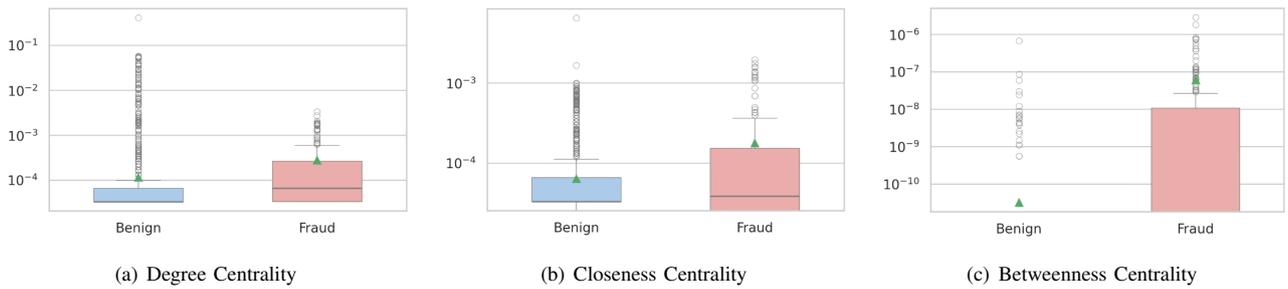


Fig. 3. Box plots of centrality measures for account graphs. Each plot compares the distribution of (a) degree, (b) closeness, and (c) betweenness centrality between benign and fraud nodes. The plots highlight distinct structural characteristics of fraudulent accounts, such as higher variance and presence of extreme outliers.

range (maximum about 0.0033). Despite lacking such hubs, fraudulent accounts exhibit higher average and median values (mean about 2.8×10^{-4} , median about 6.6×10^{-5} versus benign mean about 1.1×10^{-4} , median about 3.3×10^{-5}), indicating that mid-level connectivity is more common among them. This implies that benign accounts are dominated by many minimally connected nodes and a few hubs, while fraudulent accounts are more evenly distributed at moderate connectivity levels. Closeness centrality also highlights differences. Fraudulent accounts have higher mean and upper quartile values (mean about 1.8×10^{-4} , upper quartile about 1.5×10^{-4}) compared to benign accounts (mean about 6.4×10^{-5} , upper quartile about 6.6×10^{-5}). This may indicate that some fraudulent accounts are strategically positioned to access the network efficiently. Betweenness centrality shows the clearest distinction. Nearly all benign accounts remain at zero (mean about 3.2×10^{-11} , maximum about 6.8×10^{-7}). Fraudulent accounts, however, display a broader spread with significantly larger values (mean about 6.1×10^{-8} , maximum about 2.9×10^{-6}). This reveals that a subset of fraud nodes serve as intermediaries or bridges in transaction flows.

These findings are consistent with prior research [35], which emphasizes that nodes with relatively low degree but high betweenness often play organizational or intermediary roles in money laundering schemes. Our analysis shows that fraudulent accounts avoid extreme hub positions but instead exploit mid-level connectivity, higher reachability, and bridging roles. Based on these observations, we incorporate degree, closeness, and betweenness centrality into the `FraudCenGCL` framework as auxiliary structural features. This integration enriches node representations with structural semantics, enabling the model to better capture fraud-specific behavioral patterns and improve embedding quality and fraud detection performance.

D. Impact of Centrality Measures (RQ.3)

We analyze how individual and combined centrality measures affect the performance of `FraudCenGCL`, using BGRL as the backbone (Table IV). Incorporating all three centrality types (degree, closeness, and betweenness) achieves the best overall balance, with the highest F1-score (0.4235) and strong improvements in precision (0.3214) compared to the baseline.

TABLE IV
PERFORMANCE OF THE BGRL MODEL USING DIFFERENT COMBINATIONS OF CENTRALITY MEASURES: DEGREE CENTRALITY (DC), CLOSENESS CENTRALITY (CC), AND BETWEENNESS CENTRALITY (BC).

Centralities	Precision	Recall	F1-score	AUROC	AUPRC
Baseline	0.1177	0.4211	0.1839	0.7514	0.0852
DC+CC+BC	0.3214	0.6207	0.4235	0.8533	0.3096
DC	0.2754	0.6786	0.3918	0.9430	0.3331
CC	0.2421	0.7419	0.3651	0.9172	0.3171
BC	0.2447	0.7419	0.3680	0.9171	0.3177
DC+CC	0.2453	0.5200	0.3333	0.8611	0.2433
DC+BC	0.2623	0.5926	0.3636	0.9168	0.2008
CC+BC	0.2391	0.4400	0.3099	0.8802	0.2788

However, the AUROC (0.8533) is lower than that of individual degree centrality.

When each centrality measure is used in isolation, the performance gains vary. Degree centrality provides the best AUROC (0.9430) and AUPRC (0.3331), showing its strength in ranking quality. Closeness and betweenness centrality yield the highest recall (0.7419), but their precision is relatively low, resulting in modest F1-scores (0.3651 and 0.3680). Two-measure combinations also show unstable behavior: for example, DC+CC and DC+BC lead to lower F1-scores (0.3333 and 0.3636), reflecting a trade-off where recall improves only slightly while precision drops.

Each centrality offers a distinct structural signal. Degree centrality captures local connectivity, closeness centrality measures network reachability, and betweenness centrality highlights intermediary roles within the transaction graph. When used together, these perspectives improve the overall robustness of the model, as seen in the highest F1-score from DC+CC+BC. However, partial inclusion of only two centrality measures can degrade performance, indicating that complementary information is best exploited when all three are integrated.

In summary, the integration of all three centrality measures results in more robust and balanced representations for fraud detection. Partial inclusion, on the other hand, can lead to degraded performance, highlighting the importance of holistic structural modeling in graph-based learning.

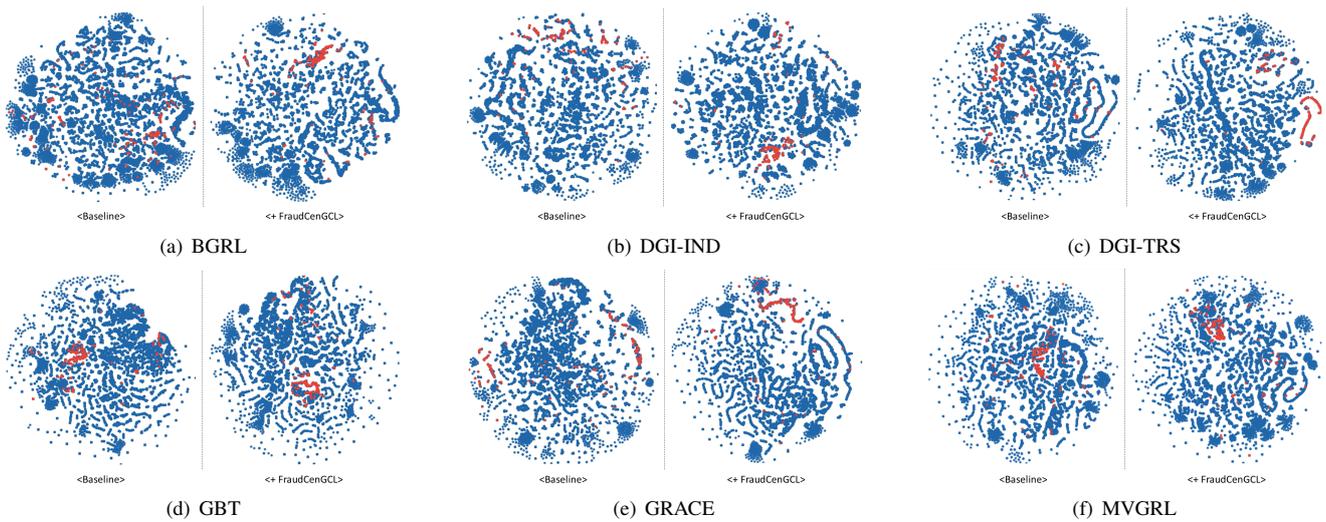


Fig. 4. t-SNE visualizations of node embeddings before and after applying `FraudCenGCL` across six backbone models. `FraudCenGCL` consistently improves the compactness and separability of fraudulent nodes (red), enabling more discriminative representations for fraud detection.

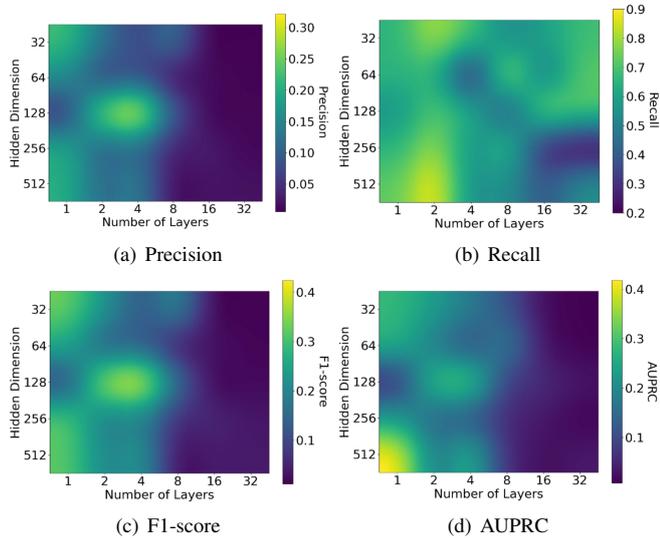


Fig. 5. Performance heatmaps of the BGRL model across different GNN layer depths and hidden dimensions for (a) Precision, (b) Recall, (c) F1-score, and (d) AUPRC. Brighter regions indicate better performance.

E. GNN Architecture Analysis (RQ.4)

We evaluate the impact of the GNN encoder architecture on the performance of `FraudCenGCL`, using BGRL as the backbone model. Figure 5 presents heatmaps of four evaluation metrics (Precision, Recall, F1-score, and AUPRC) across different combinations of GNN layers and hidden dimensions.

The results indicate that shallow architectures with moderately large hidden dimensions, particularly around 128, consistently yield higher F1-score and AUPRC. These configurations strike a balance between expressiveness and generalization, enabling the model to capture structural and behavioral patterns without incurring oversmoothing. Precision also reaches its peak in this region, while recall tends to favor larger

hidden dimensions such as 512 with two layers. This highlights that different metrics emphasize different trade-offs, where wider but shallow networks maximize precision and F1-score, whereas very large dimensions with two layers maximize recall but reduce precision.

In summary, under the BGRL backbone, GNN encoders with one to four layers and hidden dimensions in the range of 128 to 512 achieve more stable and robust performance. These observations provide practical guidance for selecting architectural configurations in graph-based fraud detection tasks, where balancing sensitivity and specificity is critical.

F. Impact on Representation Space (RQ.5)

To qualitatively assess the impact of `FraudCenGCL` on the learned embedding space, we visualize node representations before and after applying the framework using t-SNE across six backbone models (Figure 4). Fraudulent accounts are shown in red and benign accounts in blue.

Overall, `FraudCenGCL` consistently improves both the compactness of fraudulent node clusters and their separation from benign nodes. The improvement is most pronounced in lightweight backbones such as GBT and GRACE, where fraudulent nodes form distinctly tighter and more isolated clusters after applying our framework.

For BGRL and DGI-TRS, the improvements are moderate but evident, as fraudulent nodes exhibit reduced scattering and more coherent local clusters. Even in stronger baselines such as DGI-IND and MVGRL, where the initial separation of classes was already relatively good, `FraudCenGCL` yields noticeable refinements in cluster boundaries and enhances the clarity of fraudulent node groupings.

These qualitative findings are consistent with the quantitative results reported in Section IV-B. While t-SNE visualization does not directly measure classification performance, the observed structural refinement of the embedding space

provides clear evidence that centrality-guided augmentation enhances representation learning and improves the discriminability of fraudulent accounts.

V. CONCLUSION AND FUTURE WORK

In this paper, we present FraudCenGCL, a graph contrastive learning framework that integrates behavioral features with centrality-based structural roles for detecting fraud in financial transaction networks. Experiments on large-scale interbank transfer data from the housing finance information network (HOFINET) in South Korea show consistent improvements across six representative GCL backbones in terms of precision, recall, F1-score, AUROC, and AUPRC.

Our analysis highlights that fraudulent accounts differ structurally from benign ones by exhibiting mid-level connectivity, higher reachability, and bridging roles. Incorporating all three centrality measures yields the most robust gains, while shallow GNNs with moderately large hidden dimensions provide a favorable trade-off between precision and recall.

As a limitation, our current framework relies on account-level labeling derived from transaction-level suspicion flags. In future work, we plan to extend the framework to directly model suspicious transactions, integrate temporal dynamics, and explore cross-institutional graph settings. These directions will help capture evolving patterns of fraudulent activity at a finer granularity in real-world financial networks.

REFERENCES

- [1] U. Fiore, A. De Santis, F. Perla, P. Zanetti, and F. Palmieri, "Using generative adversarial networks for improving classification effectiveness in credit card fraud detection," *Information Sciences*, vol. 479, pp. 448–455, 2019.
- [2] E. Kurshan and H. Shen, "Graph computing for financial crime and fraud detection: Trends, challenges and outlook," *International Journal of Semantic Computing*, vol. 14, no. 04, pp. 565–589, 2020.
- [3] Y. Lucas and J. Jurgovsky, "Credit card fraud detection using machine learning: A survey," *arXiv preprint arXiv:2010.06479*, 2020.
- [4] A. Ali, S. Abd Razak, S. H. Othman, T. A. E. Eisa, A. Al-Dhaqm, M. Nasser, T. Elhassan, H. Elshafie, and A. Saif, "Financial fraud detection based on machine learning: a systematic literature review," *Applied Sciences*, vol. 12, no. 19, p. 9637, 2022.
- [5] D. Cheng, Y. Zou, S. Xiang, and C. Jiang, "Graph neural networks for financial fraud detection: a review," *Frontiers of Computer Science*, vol. 19, no. 9, p. 199609, 2025.
- [6] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *arXiv preprint arXiv:1609.02907*, 2016.
- [7] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio, "Graph attention networks," *arXiv preprint arXiv:1710.10903*, 2017.
- [8] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," *Advances in neural information processing systems*, vol. 30, 2017.
- [9] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?" *arXiv preprint arXiv:1810.00826*, 2018.
- [10] Y. Wang, J. Zhang, Z. Huang, W. Li, S. Feng, Z. Ma, Y. Sun, D. Yu, F. Dong, J. Jin *et al.*, "Label information enhanced fraud detection against low homophily in graphs," in *Proceedings of the ACM Web Conference 2023*, 2023, pp. 406–416.
- [11] H. Pei, B. Wei, K. C.-C. Chang, Y. Lei, and B. Yang, "Geom-gcn: Geometric graph convolutional networks," *arXiv preprint arXiv:2002.05287*, 2020.
- [12] S. Thakoor, C. Tallec, M. G. Azar, M. Azabou, E. L. Dyer, R. Munos, P. Veličković, and M. Valko, "Large-scale representation learning on graphs via bootstrapping," *arXiv preprint arXiv:2102.06514*, 2021.
- [13] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.
- [14] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 9729–9738.
- [15] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, "Supervised contrastive learning," *Advances in neural information processing systems*, vol. 33, pp. 18 661–18 673, 2020.
- [16] P. Veličković, W. Fedus, W. L. Hamilton, P. Liò, Y. Bengio, and R. D. Hjelm, "Deep graph infomax," *arXiv preprint arXiv:1809.10341*, 2018.
- [17] F.-Y. Sun, J. Hoffmann, V. Verma, and J. Tang, "Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization," *arXiv preprint arXiv:1908.01000*, 2019.
- [18] K. Hassani and A. H. Khasahmadi, "Contrastive multi-view representation learning on graphs," in *International conference on machine learning*. PMLR, 2020, pp. 4116–4126.
- [19] Y. Zhu, Y. Xu, F. Yu, Q. Liu, S. Wu, and L. Wang, "Deep graph contrastive representation learning," *arXiv preprint arXiv:2006.04131*, 2020.
- [20] Y. You, T. Chen, Y. Sui, T. Chen, Z. Wang, and Y. Shen, "Graph contrastive learning with augmentations," *Advances in neural information processing systems*, vol. 33, pp. 5812–5823, 2020.
- [21] Y. Dou, Z. Liu, L. Sun, Y. Deng, H. Peng, and P. S. Yu, "Enhancing graph neural network-based fraud detectors against camouflaged fraudsters," in *Proceedings of the 29th ACM international conference on information & knowledge management*, 2020, pp. 315–324.
- [22] Y. Liu, X. Ao, Z. Qin, J. Chi, J. Feng, H. Yang, and Q. He, "Pick and choose: a gnn-based imbalanced learning approach for fraud detection," in *Proceedings of the web conference 2021*, 2021, pp. 3168–3177.
- [23] D. Cheng, X. Wang, Y. Zhang, and L. Zhang, "Graph neural network for fraud detection via spatial-temporal attention," *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 8, pp. 3800–3813, 2020.
- [24] G. Zhang, J. Wu, J. Yang, A. Beheshti, S. Xue, C. Zhou, and Q. Z. Sheng, "Fraudre: Fraud detection dual-resistant to graph inconsistency and imbalance," in *2021 IEEE international conference on data mining (ICDM)*. IEEE, 2021, pp. 867–876.
- [25] S. Xiang, M. Zhu, D. Cheng, E. Li, R. Zhao, Y. Ouyang, L. Chen, and Y. Zheng, "Semi-supervised credit card fraud detection via attribute-driven graph representation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 12, 2023, pp. 14 557–14 565.
- [26] I. Misra and L. v. d. Maaten, "Self-supervised learning of pretext-invariant representations," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 6707–6717.
- [27] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. Guo, M. Gheshlaghi Azar *et al.*, "Bootstrap your own latent—a new approach to self-supervised learning," *Advances in neural information processing systems*, vol. 33, pp. 21 271–21 284, 2020.
- [28] P. Bielak, T. Kajdanowicz, and N. V. Chawla, "Graph barlow twins: A self-supervised representation learning framework for graphs," *Knowledge-Based Systems*, vol. 256, p. 109631, 2022.
- [29] T. U. Kuzubaş, I. Ömercikoğlu, and B. Saltoğlu, "Network centrality measures and systemic risk: An application to the turkish financial crisis," *Physica A: Statistical Mechanics and its Applications*, vol. 405, pp. 203–215, 2014.
- [30] A. Temizsoy, G. Iori, and G. Montes-Rojas, "Network centrality and funding rates in the e-mid interbank market," *Journal of Financial Stability*, vol. 33, pp. 346–365, 2017.
- [31] Y. Xu and J. Corbett, "Using network method to measure financial interconnection," National Bureau of Economic Research, Tech. Rep., 2019.
- [32] R. Martínez-Ventura, R. Mariño-Martínez, and J. Miguélez-Márquez, "Redundancy of centrality measures in financial market infrastructures," *Latin American Journal of Central Banking*, vol. 4, no. 4, p. 100098, 2023.
- [33] L. C. Freeman *et al.*, "Centrality in social networks: Conceptual clarification," *Social network: critical concepts in sociology*. Londres: Routledge, vol. 1, pp. 238–263, 2002.
- [34] J. Zhu, Y. Yan, L. Zhao, M. Heimann, L. Akoglu, and D. Koutra, "Beyond homophily in graph neural networks: Current limitations and effective designs," *Advances in neural information processing systems*, vol. 33, pp. 7793–7804, 2020.

- [35] R. Dreżewski, J. Sepielak, and W. Filipkowski, “The application of social network analysis algorithms in a system supporting money laundering detection,” *Information Sciences*, vol. 295, pp. 18–32, 2015.