Digital Intelligence PRESS

Journal of Digital Frontier

https://doi.org/10.67541/jdf2613

Research Article **Volume 1, Issue 3 (Nov 2026)** *Open Access*

Cross-Institution Cooperative Detection Model of Insurance Fraud Based on Graph Neural Networks and Federated Learning

Chen Zhao *

Harvard Extension School, Harvard University, Cambridge, Massachusetts, United States

***Corresponding author:** Chen Zhao. Email: mr.zhaoc@outlook.com

Abstract

The core challenge in cross-institution insurance fraud detection is that fraud-ring connectivity crosses institutional boundaries, while privacy regulations prohibit the centralized aggregation of raw claims data. This paper proposes CFGFraud, a federated graph detection framework for partially observable graphs. In this framework, a virtual-node placeholder mechanism makes the local model aware of cross-institutional hub entities; institutional attention mask separates the parameterization of real and virtual connections and imposes a learnable attenuation coefficient; and a federated attention alignment strategy constrains drift in each institution's attention pattern during aggregation. The non-IID experiment with five clients on the IEEE-CIS dataset shows that CFGFraud achieves an AUPRC of 0.791, 0.057 higher than FedAvg+GAT. Fraud-ring recall reaches 0.723, 24.4% higher than FedAvg+GAT. The standard deviation across clients is compressed from 0.029 to 0.012. Single-round computation costs increase by about 8%, while communication overhead increases by less than 0.01%. Ablation experiments show that the virtual-node mechanism accounts for 31.1% of the fraud-ring recall gain, and collaboration gain is positively correlated with entity overlap across institutions. This framework provides a deployable technical path for cross-agency relational fraud detection under the constraint of data confinement.

Keywords

Insurance fraud detection; Graph neural network; Federated learning; Cross-institution collaboration; Attention alignment

Received: 25 Aug 2026; Revised: 16 Sep 2026; Accepted: 30 Oct 2026; Published: 17 Nov 2026
Copyright: ©2026 The Author(s). Published by Digital Intelligence Press Limited. This is an open access article under the CC BY 4.0 license.

1. INTRODUCTION

The collusive and cross-institution nature of insurance fraud has created a structural problem: professional fraud syndicates deliberately avoid repeating the same tactics at the same institution, and the fraud rings that any single insurer's local view can reveal are only a cross-section of a complete criminal network. The key connection points of the fraud loop often fall outside the boundaries of the institution. The economic investigation department of the Ministry of Public Security also pointed out that insurance fraud has evolved into a joint conspiracy of multi-party collusion, and the effect of improving the identification ability only by insurance companies themselves is limited.

Privacy regulation prevents this problem from being solved by a centralized solution. Claim data contains the identity of the applicant, medical records and maintenance details, and cross-institution aggregation faces the rigid constraints of the Personal Information Protection Law and the Data Security Law [1],[2]. Regulatory authorities put forward industry joint prevention goals and promote the

construction of anti-fraud information platform, but there is still a significant gap between the sharing mechanism at the institutional level and the technical implementation at the algorithm level. Federated learning allows multiple institutions to train jointly without exchanging raw data, providing principled feasibility for cooperation under data-locality constraints [3],[4]. However, it faces unique difficulties when combining it with graph neural network for insurance fraud detection: the weight of graph attention layer encodes the structural knowledge of which neighbors are more important, which is a function of the local graph topology. When institutional graph structures differ significantly in size and entity distribution, parameter averaging forcibly mixes judgments formed under different structural conditions. More fundamentally, the graph is fragmented [5],[6]. The local graph naturally lacks cross-institutional edges, and the standard attention mechanism cannot distinguish between two states: there is genuinely no neighbor, and a neighbor exists but is unreachable, and the cooperation information carried by the latter is quietly discarded.

Existing explorations have not been sufficient to respond to these difficulties. FedSecure-GNN uses Edge-GAT to learn account relationships, XGBoost to process transaction features, and multi-bank joint training through federated learning. In cross-bank fraud detection, FedSecure-GNN achieves an F1 value of about 0.72. However, its graph learning and feature classification are decoupled from the prediction layer rather than the representation layer. And it does not deal with the node alignment problem when the same entity appears in different institutions with different identifiers. RegMeta-pFL for incremental auto insurance data takes tabular features as input and leaves untouched the special behavior of graph structure in federated aggregation [7]. The idea of virtual node plus attention distillation in relevant domestic patents is closest to the problem of graph fragmentation at the conceptual level, but it lacks public algorithm details and ablation verification, which is not enough to form a reproducible baseline. Related work on heterogeneous graph modeling and hybrid graph-learning pipelines has shown that coupling graph attention with gradient-boosted decision trees can capture both relational structure and tabular discrimination, while attention-guided graph neural networks combined with U-Net++-style architectures demonstrate that graph aggregation remains effective under strongly heterogeneous node distributions [8],[9]. Federated learning itself has been extended to collaborative prediction across decentralized data holders with residual correction, and graph-based fraud detection over heterogeneous graphs combined with XGBoost has been validated in cross-border trade fraud scenarios, which confirms the feasibility of relational signals in fraud identification but also shows that cross-domain connectivity is the decisive factor [10],[11],[12]. On the modeling side, physics-informed and interpretability-boosting sequence models illustrate how auxiliary constraints stabilize learning under distribution shift, and graph-attention-based inversion frameworks show that attention weights are highly sensitive to local topology, so that direct parameter averaging across structurally different graphs is problematic [13],[14],[15],[16].

The contribution of this paper is positioned to bridge the above fracture. At the aggregation level, the federated attention alignment strategy is proposed. The server collects low-dimensional summaries of the attention distribution of each client while aggregating parameters, constraining the global model not to deviate from the consensus region of the attention mode of each institution, and restraining the degradation of relationship signals caused by graph fragmentation in parameter averaging. At the local representation level, an agent-aware graph attention aggregation mechanism is designed to make the model perceive the existence of cross-mechanism connections by virtual node occupancy, and separate attention parameterization and learnable attenuation coefficient are used to prevent virtual nodes from diluting real neighbor signals. At the verification level, an experimental protocol with multi-institution heterogeneous data is constructed to simulate the topology difference and fraud rate fluctuation of real

insurance alliance by non-independent and identically distributed partition, and the source, boundary and cost of cooperative gain are systematically tested.

2.RELEVANT THEORETICAL BASIS AND DIFFERENTIATED POSITIONING

Graph attention changes fraud identification from attribute judgment of isolated records to evidence propagation in relationship structure, so that claim nodes can obtain discrimination signals far beyond their own characteristics through associated repair factories, policyholders and beneficiaries. Existing improvements have developed in three directions: cardinality-preserving attention to solve the problem of multi-set discrimination when there are multiple edges of different types in the same neighbor; Edge graph attention pushes attention calculation down to the edge level and participates in weight generation with edge features. Constrained optimization methods cope with the camouflage behavior of fraudsters deliberately establishing connections with honest nodes through neighborhood selection or label aware sampling. These improvements share the unspoken premise that treated neighbors are accessible, at least in principle. When the other end of an edge exists in another institution and its attributes are completely unobservable, the multiset structure loses reference, the edge features cannot be obtained, and the neighborhood selection loses the candidate set. Existing improvements deal with problems where neighbors are visible but indistinct, while cross-agency scenarios face problems where neighbors are not visible at all.

The effectiveness of the federated aggregation paradigm on tabulation and sequence data has been fully verified. Parameter averaging of FedAvg, proximal term constraint of FedProx, meta-learning starting point of Per-FedAvg and RegMeta-pFL for auto insurance are all premised on the premise that the local model structure does not depend on graph topology. When the local model includes a graph aggregation layer, the consequences of parameter averaging are much more complex: attention parameters are tightly coupled to the degree distribution, edge type proportion, and sampling pattern of the local graph, and direct averaging is equivalent to forcibly mixing two sets of judgments formed based on different structural evidence. What is more hidden is the amplification of neighborhood sampling bias. The sampling strategy of each client systematically biases the edge type with high local proportion. This bias is solidified into the attention parameter in the decline of local loss, and after aggregation, it is superimposed as the global model deviates from its optimal attention mode on any client. Schemes such as FedStar and FedGKD mitigate heterogeneity by sharing common underlying structure or graph data set distillation, but the premise is that each client graph contains enough structural information for extraction [17],[18]. The key missing part of the insurance scenario is precisely the connection with other institutions, and cross-institution hubs may not exist at all in the local graph or may be directly truncated.

Among the existing attempts of cross-mechanism graph collaboration, the limitation of FedSecure-GNN is that the fusion of graph path and feature path occurs at the prediction layer, the relationship signal cannot affect the attribute judgment at the representation level, and the existence correlation between cross-domain nodes is not established. The federated heterogeneous graph decomposes the heterogeneous graph into the same proton graph to reduce the aggregation complexity but discards the type of crossover structure [19],[20],[21]. However, the maintenance service relationship and the agent relationship in the insurance graph play completely different roles in the fraud loop, and the discrimination power exists in this crossover. Therefore, this study makes clear the differentiation positioning: it does not pursue reconstruction or approximate complete cross-agency graph but accepts

partial observable graph as the basic condition; cross-agency connection exists in the form of unreachable; effective collaboration does not rely on node alignment or edge completion but maintains the structural integrity of attention signal at the level of federated aggregation. The virtual node mechanism distinguishes two states: there is no external connection and there is external connection but unreachable. The mechanism mask attention prevents virtual connection from flooding the real neighbor signal, and the attention alignment aggregation prevents the attention mode of each institution from being systematically distorted during aggregation. The three respond to the three gaps of unreachable neighbor, aggregation disturbance and node alignment loss respectively.

3. INTERAGENCY FEDERATED GRAPH FRAUD DETECTION FRAMEWORK (CFGFRAUD)

3.1 Overall architecture and design principles

CFGFraud uses a server-client topology, but the key difference from standard federated learning is that the server receives more than just model parameters. Each client holds a subgraph of local claims, a fraud detection model based on graph attention, and a summary of attention distribution aggregated by edge type. The server performs two tasks: weighted aggregation of model parameters and computation of the global attention consistency regularization term by using the attention summary uploaded by each client, so as to constrain the aggregated model near the consensus region of the attention mode of each institution. No exchange of raw data, node identity, or adjacency occurs between clients.

The design of the framework follows three principles, each of which directly responds to a specific failure mode in cross-agency graph learning. Relationship signal takes precedence over feature signal: in insurance fraud, the structural features of fraud ring (multiple claims connected in series through the same repair plant or the same beneficiary) are often more discriminative than the features of single claim. Therefore, the model architecture is mainly based on graph attention rather than auxiliary, and feature classifier is only a partial supplement. Aggregation consistency takes precedence over aggregation accuracy: The statistical characteristics of each machine graph vary greatly, and the entity type distribution of large insurance companies and regional institutions is completely different. The pursuit of the optimal average accuracy of the global model on each client will lead to repeated oscillations of the attention layer between heterogeneous graphs. Therefore, the primary goal of aggregation is to maintain the relative structure of the attention distribution of each institution rather than the arithmetic means of the accuracy index. Communication efficiency serves availability: the parameter scale of graph model is coupled with the subgraph sampling scale, which will lead to uncontrollable communication cost of federated rounds if unconstrained. Therefore, the design of summary mechanism takes "not becoming a communication bottleneck" as the hard constraint. **Table 1** lists the architect-level differences between CFGFraud and two typical federated graph schemes to position the design choices in this paper.

Table 1. Architecture-level differences between CFGFraud and typical federated graph schemes

Dimension of dimension	FedAvg+GAT (baseline)	FedSecure-GNN（Edge-GAT）	CFGFraud
Cross-institutional entity perception	No	No	Virtual-node placeholder
Attention aggregation strategy	Uniform parameter averaging	Edge level attention + classifier fusion	Mechanism mask + attenuation coefficient
The server receives the content	Model parameters only	Model parameters only	Parameter + attention summary
Constraint of aggregation	No	Median defense	Attention consistency regularity

3.2 Construction of local insurance claim graph and subgraph sampling

The node system of the local graph covers five types of core entities in the insurance claim chain: policyholder, insured, repair plant, medical institution and agent. Side types can be divided into six types according to business semantics: insurance relationship, claim relationship, maintenance service relationship, diagnosis and treatment service relationship, agency relationship and common beneficiary relationship. Each edge type carries a type identifier, which is used for type-aware parameterization in the subsequent attention mechanism. Node characteristics are composed of statistical attributes after desensitization, including historical claim frequency, quantile of amount distribution, entity duration, etc., and do not contain any information that can directly identify individuals.

Virtual node mechanism is the most critical design in local graph construction. When an entity appears in the local graph, but another entity associated with it exists in another institution, the standard practice is to directly discard the cross-institution edge, resulting in the local model being completely unaware of the existence of external connections. The practice of CFGFraud is to create a virtual placeholder node for such cross-institutional shared entities: the node exists in the local graph and is connected to the local entity by a virtual edge. However, the virtual node does not carry any feature information but only serves as a structural signal to tell the attention layer of the graph that there is a relationship to the outside, but the information is not reachable. Taking a shared repair shop as an example, a claim record of institution A is associated with repair shop M, and M also processes many claims from institution B, but A cannot obtain M's associated data on institution B's side. In A's local graph, M is connected to the relevant claim entity in the form of A virtual node, so that A's model at least knows that the repair plant has the nature of a cross-agency hub, although it cannot know what is specifically connected. Formally, let the local graph of client c be:

$$G_c = (V_c, E_c) \tag{1}$$

$$V_c = V_c^{real} \cup V_c^{virt} \tag{2}$$

Is the set of local entities carrying real features, and V_c^{virt} is the set of virtual placeholder nodes. The feature vectors of virtual nodes are uniformly set as zero vectors, that is, $x_v = 0$ for any $v \in V_c^{virt}$.

The neighborhood subgraph sampling strategy takes the claim entity as the anchor point and preferentially retains the diversity of edge types rather than the degree of nodes. In the fraud detection scenario, the information value of a claim node with low degree but connecting three different entities,

the policyholder, the repair plant and the beneficiary, is much higher than that of a node with high degree but all pointing to the same type of entity. Specifically, for each anchor point claim entity, the sampling process selects up to k neighbors from each edge type in turn, and k is 8. If a certain type of neighbor is insufficient, it will be made up from adjacent types. This strategy ensures that even in the sparse graph of some institutions, the type crossing structure on which the fraud ring depends is not overwhelmed by the degree first strategy during the sampling phase.

3.3 Graph attention aggregation module (IA-GAT) for agency awareness

The standard graph attention mechanism will have a hidden but serious failure in the presence of virtual nodes: because virtual nodes are not associated with any real features, their key vector in attention calculation will degenerate into a nearly constant representation, resulting in abnormal concentration of attention coefficient on virtual connections. Intuitively, the model will tend to focus on the placeholder node that has nothing, because its representation does not carry noise and appears stable at the beginning of training. This virtual node dilution effect makes the signals of real neighbors systematically weaken in aggregation [22],[23].

The improvement of IA-GAT lies in the separation of attention parameterization by connection nature. For each edge from node i to neighbor j, the mechanism mask indicator variable $m_{ij} \in \{0,1\}$ is introduced, $m_{ij} = 1$ when j is a real neighbor and $m_{ij} = 0$ when j is a dummy node. The calculation of attention coefficient is divided into two groups of independent parameters according to the mask, and its normalized form is as follows:

$$e_{ij} = m_{ij} \cdot a_{\text{real}}^{T}\left[W_r h_i \parallel W_r h_j\right] + \left(1 - m_{ij}\right) \cdot a_{\text{virt}}^{T}\left[W_v h_i \parallel W_v h_j\right] \tag{3}$$

Where h_i and h_j are the current representations of nodes i and j respectively, W_r and W_v are the linear transformation matrices corresponding to real connection and virtual connection respectively, a_{real} and a_{virt} are two sets of independent attention vector parameters, and $\parallel$ represents vector splicing operation. After normalization, the aggregation of node i is expressed as:

$$h_i' = \sigma\left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij} \cdot \gamma^{1-m_{ij}} \cdot W_r h_j\right) \tag{4}$$

Where α_{ij} is the attention coefficient normalized by softmax, σ is the activation function, and $\gamma \in (0,1)$ is the learnable virtual connection attenuation coefficient, whose initial value is set to 0.1 and adjusted adaptively by the client in training. The effect of γ is that the structural signal of virtual connection is retained, but its interference on feature aggregation is suppressed to a controllable range, and the attenuation strength is automatically adjusted by the fraud detection performance on the local validation set through gradient descent. A distinction needs to be made between the problems solved by IA-GAT and cardinality preserving attention. Cardinal-preserving attention deals with the problem of multiset discrimination when the same neighbor has multiple edges of different types. Its goal is to let the model know the difference between neighbor A connected by two edges and neighbor A and neighbor B connected by one edge each.

3.4 Federated attention alignment aggregation policy (FedAttAlign)

Local graphs of insurance institutions differ systematically in size, entity distribution, and edge-type proportion. In the graph of large comprehensive insurance institutions, the nodes of medical

institutions and maintenance plants are dense, and the type of distribution is balanced; Regional agency plans may be highly skewed towards one type of entity. This structural heterogeneity leads to different attention distribution patterns in the graph attention layer of each client after training: some institutions tend to pay high attention to the repair shop connection, while others attach higher weight to the agent relationship. When the server performs standard FedAvg, parameter averaging will mix these attention parameters with different tendencies into a compromise representation, with the consequence that instead of balancing all parties, all parties are dissatisfied: the aggregated global model's attention distribution on any client deviates from the optimal pattern supported by that client's local data.

FedAttAlign's response is to introduce constraints at the level of attention distribution in the aggregation stage, rather than just at the parameter level. At the end of each round of local training, each client calculates a low-dimensional attention summary vector. Let the set of edge types be R, $|R| = 6$, and the attention summary of client c to type r in round t is defined as the meaning of the attention coefficients of all edges of type r in the local graph of the client, denoted as $s_c^t(r)$. Six class type form the vector $s_c^t = [s_c^t, (r_1), s_c^t, (r_2) \dots s_c^t(r_6)] \in \mathbb{R}^6$. The vector is uploaded to the server after local differential privacy noise processing. After the weighted average of parameters is completed, the server calculates the deviation between the attention distribution $\hat{s}^t(r)$ generated by the aggregated model on the virtual validation set and the weighted average of each client summary and uses it as the consistency regularization term. Global attention consistency loss is defined as:

$$L_{\text{att}} = \sum_{r \in R} \left(\hat{s}^t(r) - \sum_{c=1}^{C} p_c \cdots_c^t (r) \right)^2 \tag{5}$$

Where C is the total number of clients, p_c is, the sample weight of client c, which is satisfied by ($\sum_{c=1}^{C} p_c = 1$, $\hat{s}^t(r)$) is the average attention estimate of the aggregated model on type r. If L_{att} exceeds the preset threshold $\tau = 0.02$, the server performs one-step correction on the parameters of the attention layer with a small learning rate $\eta_{\text{att}} = 0.01$, so that the aggregated model pulls back to the consensus region. The communication overhead of this strategy is controllable: the 6-dimensional summary vector is negligible compared to the IA-GAT backbone of about 120,000 parameters. Table 2 shows the comparison of traffic before and after the introduction of summary mechanism under the condition of five clients and each round of full parameter uploading.

Table 2. Comparison of federated traffic before and after the introduction of attention summarization

Uploaded model parameters per round	Uploaded model parameters per round	Number of abstracts uploaded in each round	Relative growth rate
FedAvg (parameters only)	About 480 KB	0	Baseline
FedAttAlign	About 480 KB	About 40 B	<0.01%

3.5 Unbalanced treatment of training objectives and categories

The extreme class imbalance of insurance fraud is further amplified under federated aggregation: if the fraud rate of each client is significantly different, the global model will systematically bias the

negative class distribution of most clients after aggregation. The focal loss is used as the main loss function in local training, which is of the form [24]:

$$L_{\text{focal}} = -\sum_{i=1}^{N_c} (1-p_i)^{\zeta} \cdot [y_i \log p_i + (1-y_i)\log(1-p_i)] \tag{6}$$

Where N_c is the number of samples of client c, p_i is the fraud probability prediction of the model for sample i, $y_i \in \{0,1\}$ is the true label, and ζ is the modulation coefficient, which is set to 2.0 to make the model sensitive to minority samples that are difficult to classify in the local stage. On this basis, a graph structure regular term is introduced: for each claim anchor node, the cosine similarity between its aggregated representation and the representation of the same type of edge neighbor is constrained not to be lower than a lower bound. The regular term of graph structure is defined as follows:

$$L_{\text{graph}} = \sum_{i \in V_c^{\text{anchor}}} \sum_{j \in \mathcal{N}_r(i)} \max\left(0, \delta - \cos\left(h_i', h_j'\right)\right) \tag{7}$$

Where V_c^{anchor} is the claim anchor node set in client c, $\mathcal{N}_r(i)$ is the edge neighbor set of the same type of node i, δ is the lower bound of similarity, set to 0.3, and $cos(\cdot,\cdot)$ is the cosine similarity. This item prevents the model from over-fitting the local features of minority samples driven by focus loss and losing the generalization information of relationship structure. The total local loss of client c is:

$$L_c = L_{\text{focal}} + \lambda_g \cdot L_{\text{graph}} \tag{8}$$

Where λ_g is the graph structure regular weight, which is taken as 0.1. The global goal is set according to a clear trade-off principle: the weight coefficient of attention consistency regularity should not be large enough to force all clients to converge to the same attention pattern, which would erase institutional specificity, nor small enough to have an unconstrained effect on aggregation bias. The global target is defined as:

$$L_{\text{global}} = \sum_{c=1}^{C} p_c \cdot L_c + \lambda_{\text{att}} \cdot L_{\text{att}} \tag{9}$$

Where λ_{att} is the attention consistency weight, whose value is dynamically adjusted by observing the KL divergence of the attention distribution of each client summary and the aggregated model by the server after each round of aggregation: Let the median of KL divergence between the summary of each client and the attention distribution of the aggregate model in the t-th round be d_t. If $d_t > 0.15$, let $\lambda_{\text{att}} \leftarrow \lambda_{\text{att}} \cdot 1.1$; if $d_t < 0.05$, let $\lambda_{\text{att}} \leftarrow \lambda_{\text{att}} \cdot 0.9$, and the initial value $\lambda_{\text{att}} = 0.5$. This adaptive strategy avoids dependence on fixed hyperparameters and enables the framework to adapt to the degree of heterogeneity under different combinations of mechanisms.

4. ALGORITHM IMPLEMENTATION AND COMPLEXITY ANALYSIS

4.1 Local training process and federated communication protocol

The graph model is significantly more sensitive to local overfitting than the tabular or sequential

model. In the centralized training, overfitting is manifested as early saturation of the performance of the validation set. In the federated scenario, the consequences are more serious. Too many local epochs will make each client model over-adapt to the neighborhood sampling deviation of the local subgraph, resulting in artificial amplification of the heterogeneity between upload parameters, and the performance of the aggregated global model will fall back on each client. CFGFraud adopts a tight ratio strategy of local epochs and federation rounds: each client only performs two local epochs per round, and the total number of federation rounds is set to 80. When the size of the local graph is 8000 to 15000 claim edges, the ratio keeps the ergodic degree of the client on the local data between 1.6 and 3.0 times in a single round, which not only ensures the full update of parameters, but also avoids the repeated reinforcement of neighborhood sampling bias. The learning rate adopts step decay, the initial value is 0.005, and the decay is 0.5 times that of the previous stage every 20 rounds. The decay timing is synchronized with the change of attention consistency loss of the aggregated model, so as to avoid premature shrinkage of step size when the attention distribution is not stable.

Parameter upload adopts full-parameter transmission rather than sparsification or quantization, because the attention parameter of IA-GAT is highly sensitive numerically, and scarification will destroy the parameter grouping structure on which the mechanism mask depends. The content uploaded each time includes IA-GAT backbone parameters (about 120,000 dimensions), classification head parameters (about 12,000 dimensions) and attention summary vector (6 dimensions). The serialization of attention summary adopts fixed-length floating point array, each dimension is represented by 4 bytes of single precision, the overall 24 bytes, and 16 bytes of verification and version identification are added after local differential privacy noise processing, totaling 40 bytes. Table 3 shows the item-byte estimate of the content uploaded in a single round by a single client.

Table 3. Item-byte estimation of content uploaded in a single round by a single client

Upload content	Dimension	Bytes per dimension	Subtotal
IA-GAT backbone parameters	About 120,000	4	About 480 KB
Classification head parameter	About 12,000	4	About 48 KB
Attention summary vector	6	4	24 B
Verification and version identification	-	-	16 B
Total	-	-	About 528 KB

As can be seen from Table 3, the total amount of summary vector and check field is only 40 bytes, accounting for less than 1/10,000 of the total amount of upload in a single round. The absolute main body of communication cost is still the model parameter itself.

4.2 Convergence discussion

Under the non-independent and homogeneous graph structure, there is a definable deviation between the aggregate update of FedAttAlign and the fully supervised graph training, which can be decomposed into two parts: the graph structure heterogeneity deviation between the local optimal parameters of the client and the global optimal parameters, and the projection deviation introduced by the attention summary constraint. The former is inherent to federated graph learning and the latter is

actively imposed by FedAttAlign. This paper does not carry out strict proof but gives hypothesis and intuitive explanation of the deviation definition. Assuming that the local loss functions of each client are L-smooth and μ-strongly convex with respect to the model parameters, and the differences in edge type distributions between local graphs are bounded, the upper bound of parameter bias between the aggregated update of FedAttAlign and the centralized fully supervised training is proportional to the product of the graph heterogeneity measure and the attention consistency weight λ_{att}. Intuitively, the larger λ_{att} is, the closer the aggregate model is to the consensus region of the attention mode of each client, and the smaller the projection bias is. However, the consensus region itself may deviate from the optimal region of any client, so there is an intermediate value that minimizes the total bias.

The impact of the introduction of the attention consistency regular term on the speed of convergence is shown in two stages. In the early stage of training, the attention distribution of each client is not differentiated, the dispersion between abstracts is low, the gradient contribution of consistency regularity is weak, and the convergence speed is basically the same as that of standard FedAvg. In the middle of training, the client's attention mode begins to differentiate, and the regular term begins to play a constraint role. The correction step after each round of aggregation is equivalent to imposing a small step pullback on the attention layer, so that the attention distribution of the global model no longer fluctuates violently with the client's heterogeneity. In the late training period, the attention distribution tends to the consensus region, the gradient of the regular term tends to zero, and the convergence speed returns to the unconstrained level. Table 4 records the magnitude of the decrease in aggregate loss between FedAvg+GAT and CFGFraud at different stages of training under the 5-client, non-independent identically distributed partitioning.

Table 4. Comparison of aggregate loss reduction of FedAvg+GAT and CFGFraud in different training stages

Federated round interval	FedAvg+GAT loss reduction	CFGFraud loss reduction
1-20	0.41	0.39
21-40	0.18	0.27
41-60	0.06	0.14
61-80	0.02	0.05

As can be seen from Table 4, CFGFraud converges slightly slower than the baseline in the early period (0.39 vs. 0.41), which is consistent with the slight constraints introduced by the regular term in the early period; In the medium-term range of 21 to 40 rounds, the loss of CFGFraud decreases by 1.5 times of the baseline, indicating that attention alignment effectively suppresses the volatility caused by heterogeneity; Both tend to be flat in the later period, but CFGFraud still maintains a decline of 0.05 from round 61 to 80, indicating that its convergence point is slightly lower than the baseline. Overall, CFGFraud achieves an aggregation loss of about 0.09 lower than FedAvg+GAT in 80 rounds, at the expense of a convergence delay of about 2 rounds in the early period.

4.3 Computational and communication complexity

The time complexity of local IA-GAT is determined by the subgraph size and the proportion of

virtual nodes. Let the number of sampled neighbors of the anchor claim entity be k and the number of edge types be $|R|$, then the neighborhood aggregation computation of each anchor node is $O(k \cdot |R| \cdot d)$, where d is the hidden dimension is, and the value is 64. Since the sampling strategy restricts the total number of neighbors per anchor to $8 \cdot |R|$, the computational burden of a single anchor is decoupling from the total size of the graph and depends only on the sampling parameters. The influence of virtual node proportion on the complexity is reflected in the grouping of attention calculation: IA-GAT uses separate parameterization for real connection and virtual connection, and the complexity of the two groups of attention calculation is the same. Therefore, the virtual node proportion does not change the order of time complexity but only affects the allocation proportion of the two groups of calculation in the total calculation. In a typical setting with 15% virtual nodes, the computation of the real-connection group is about 85% and that of the virtual Conne group is about 15%, and the total computation differs by no more than 3% from the standard GAT without virtual nodes.

The traffic for federated rounds is estimated in Table 3 for single client single rounds. Under the total budget of 5 clients and 80 rounds, the total upload volume of CFGFraud is about $5 \times 80 \times 528\,KB \approx 206\,MB$, among which the summary vector and check field add up to about 16 KB, accounting for less than one in 10,000. Compared with pure FedAvg+GAT, the increase in single-round traffic of CFGFraud only comes from the summary vector and 40 bytes of the checksum field, and the model parameter part is completely consistent. If parameter compression is used, the compression space of both is the same, and the summarization mechanism does not constitute an additional communication bottleneck.

The difference from centralized graph training needs to be defined at the root cause. Centralized GAT is not more computationally expensive than federated schemes, and its single-round training time is even shorter because there is no need for multi-client scheduling and aggregation. The fundamental reason why centralization is not feasible in the cross-institution scenario is data movement: pooling the claim graph of each institution into one place means that the original node identity, adjacency relationship and claim record are physically transferred outside the boundary of the institution, which directly conflicts with the rigid constraints of privacy regulation. The federated scheme trades communication for data, at the cost of accuracy loss and round cost caused by aggregation, rather than additional consumption of computing resources. Table 5 compares the engineering features of centralized training, pure FedAvg+GAT and CFGFraud from the three dimensions of data movement, computational cost and compliance.

Table 5. Comparison of engineering characteristics of centralized training, pure FedAvg+GAT, and CFGFraud

Dimension	Centralized GAT	FedAvg+GAT	CFGFraud
Raw data crosses institutional boundaries	Yes	No	No
Single round calculation cost (relative value)	1.0	1.3	1.4
80 rounds of total traffic	0	About 206 MB	About 206 MB
Cross-institutional entity perception	Complete	No	Virtual-node placeholder
Feasibility of compliance	Infeasible	Feasible	Feasible

As can be seen from Table 5, the cost increase of CFGFraud in a single round is about 8% compared with pure FedAvg+GAT, which comes from the separation calculation and summary statistics of the two groups of parameterized in the institutional mask attention. The total traffic is the same, and the overhead of the digest mechanism is negligible. Centralized training is computationally optimal and has complete entity perception, but the feasibility of compliance is zero, and the infeasibility in this dimension cannot be compensated by engineering optimization. CFGFraud trades compliance feasibility and cross-agency entity awareness for a single-round computational cost increase of about 40% and a negligible increase in communication, an exchange that is necessary and reasonable in an insurance cross-agency scenario.

5.EXPERIMENTAL DESIGN AND RESULT ANALYSIS

5.1 Experimental setup

The experimental data is based on the IEEE-CIS Fraud Detection dataset to construct the insurance claim graph [25],[26],[27]. The dataset contains about 590,000 transaction records with a fraud rate of 3.5% and provides characteristics such as transaction amount, product type, card and address information, and device fingerprints. The combination of cards and address is used as the identity of the applicant, and the combination of counterparty and product type is used as the agent entity of repair plant and medical institution. The original records are reconstructed into six types of edge types of heterogeneous graph according to the transaction time window and geographical area. In order to construct the non-independent and homogeneous distribution division of multiple institutions, the graph is divided into five client subgraphs according to the geographical area of the applicant, so that the fraud rate of each client fluctuates between 1.2% and 6.8% respectively, and the distribution of entity types also shows differentiated deviation: The maintenance factory nodes of clients 1 and 3 account for a higher proportion, the medical institution nodes of clients 2 and 5 account for a higher proportion, and the agent node density of client 4 is significantly higher than that of the other four. This division strategy deliberately amplifies the topological heterogeneity among institutions to test the performance of the method under extreme conditions.

The baselines consist of five categories. Locally independent training GAT represents the bottom line of no collaboration at all; FedAvg+GAT represents standard federated aggregation without any graph structure adaptation; On the basis of FedAvg, FedProx+GAT introduces proximal term to constrain local drift. The Edge-GAT part of FedSecure-GNN is used as the recent reference of cross-mechanism graph scheme, which adopts edge-level attention and median aggregation. Centralized GAT is trained after merging all client graphs as a non-deployable performance upper bound. All GNN methods share the same hidden dimension 64, upper limit of sampling neighbors 8, and number of attention heads 4. The federated method uniformly sets 80 rounds, the local epoch is 2, the learning rate is 0.005 and decays to half every 20 rounds. The focal loss modulation coefficient is 2.0. The main indicator of evaluation is AUPRC, supplemented by F1 and fraud ring recall under the optimal threshold. The recall rate of fraud ring is calculated based on the connected subgraph of more than three nodes formed by the identified common repair plant or common beneficiary in the data set, and the proportion of at least one node correctly labeled as fraud by the statistical model.

5.2 Main comparison experiment

Table 6 summarizes the meaning and standard deviation of the recall rate of AUPRC, F1 and fraud

ring of the six methods under five random partitions. Centralized GAT scored the highest in all indicators, with AUPRC reaching 0.847. However, its training process requires centralized aggregation of original graph data and is not ready for deployment in cross-institution scenarios. In the deployable method, the AUPRC of CFGFraud is 0.791, which is 0.057 higher than 0.734 of FedAvg+GAT and 0.043 higher than 0.748 of FedProx+GAT. 0.035 higher than 0.756 of FedSecure-GNN. The gap in the recall rate of fraud ring is more significant: CFGFraud reaches 0.723, while FedAvg+GAT reaches only 0.581, a relative increase of 24.4%. The source of this gap can be traced to the fact that the virtual node mechanism makes the model at least perceive the presence of cross-institutional hub entities, while the baseline approach completely loses the connectivity cue of the fraud ring after the cross-institutional entities are truncated.

Table 6. Detection performance of six methods under non-independent and identically distributed partition of five clients (mean ± standard deviation).

The method	AUPRC	F1	Fraud ring recall rate
Local independent GAT	0.612 ± 0.031	0.548 ± 0.027	0.387 ± 0.042
FedAvg + GAT	0.734 ± 0.018	0.641 ± 0.021	0.581 ± 0.035
FedProx + GAT	0.748 ± 0.016	0.653 ± 0.019	0.602 ± 0.031
FedSecure-GNN (Edge-GAT)	0.756 ± 0.015	0.664 ± 0.017	0.614 ± 0.029
CFGFraud	0.791 ± 0.012	0.698 ± 0.014	0.723 ± 0.022
Centralized GAT (upper bound)	0.847 ± 0.009	0.741 ± 0.011	0.812 ± 0.016

The cross-sectional data in Table 6 only reflects the final state performance at the end of training and cannot reveal the morphological differences in the convergence process. Figure 1 shows the evolution trajectory of AUPRC of the six methods in the whole process of federated training in the form of broken lines. The horizontal axis is the federated rounds, and the vertical axis is the AUPRC evaluated on the global validation set every 5 rounds. The centralized GAT curve stabilizes around 0.84 from round 10, which serves as the upper bound reference. The CFGFraud curve is close to FedAvg+GAT in the early 20 rounds, starts to separate significantly after the 25th round, reaches 0.78 and flattens out around the 55th round. The curves of FedAvg+GAT and FedProx+GAT show a plateau after round 35, and the fluctuation range is large, reflecting the repeated oscillation of the attention layer between heterogeneous graphs. FedSecure-GNN curve is in between. Its edge-level attention mechanism brings certain improvement but lacks cross-institutional entity perception, resulting in weak growth in the later stage. The local independent GAT curve is consistently at the lowest level. The trajectory pattern shown in Figure 1 shows that the convergence advantage of CFGFraud is concentrated in the middle and late stages, rather than the uniform lead in the whole training process.

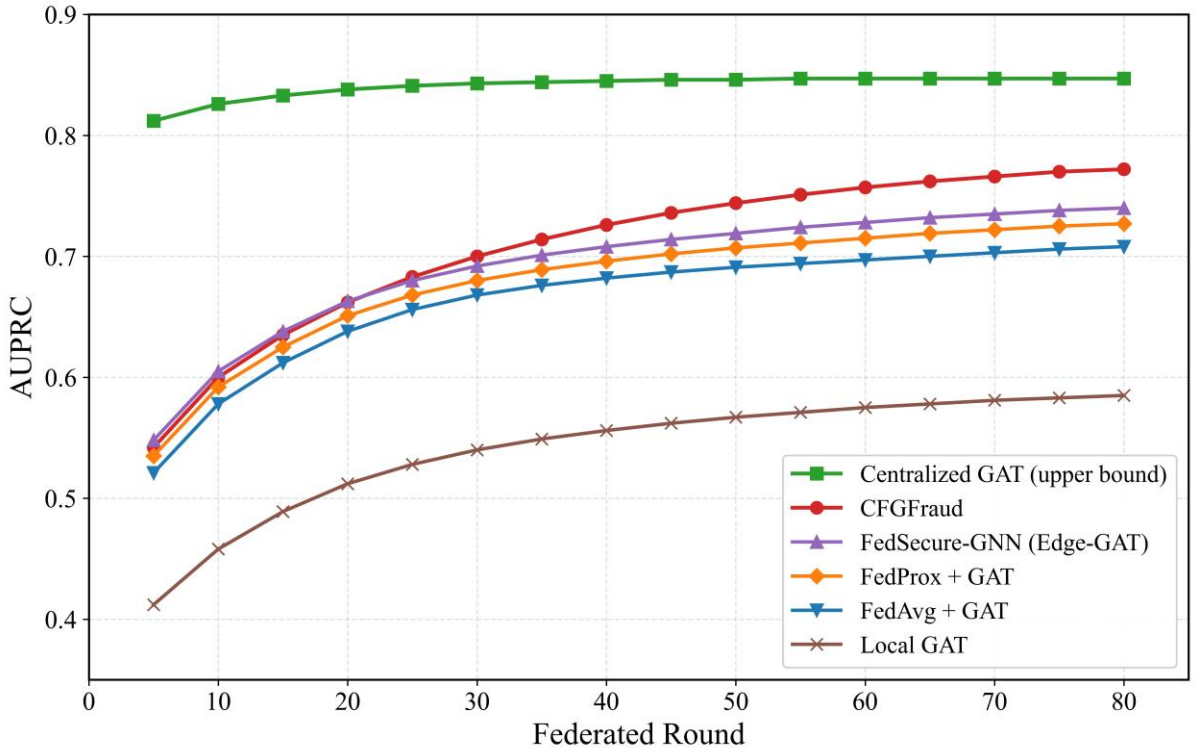


Figure 1. AUPRC evolution curves of the six methods during federated training

In order to further test the moderating effect of heterogeneity degree on collaboration gain, 5 clients are divided into high heterogeneity group and low heterogeneity group for evaluation respectively. Figure 2 shows the comparison of AUPRC of the three federated methods on the two groups with grouped bar graphs, with the bar height as the mean and the error bar as the standard deviation. FedAvg+GAT was 0.698 in the high heterogeneity group and 0.771 in the low heterogeneity group, and the difference between the two groups was 0.073. FedProx+GAT were 0.714 and 0.782, respectively, with a difference of 0.068; CFGFraud are 0.769 and 0.799 respectively, with a difference of only 0.030.

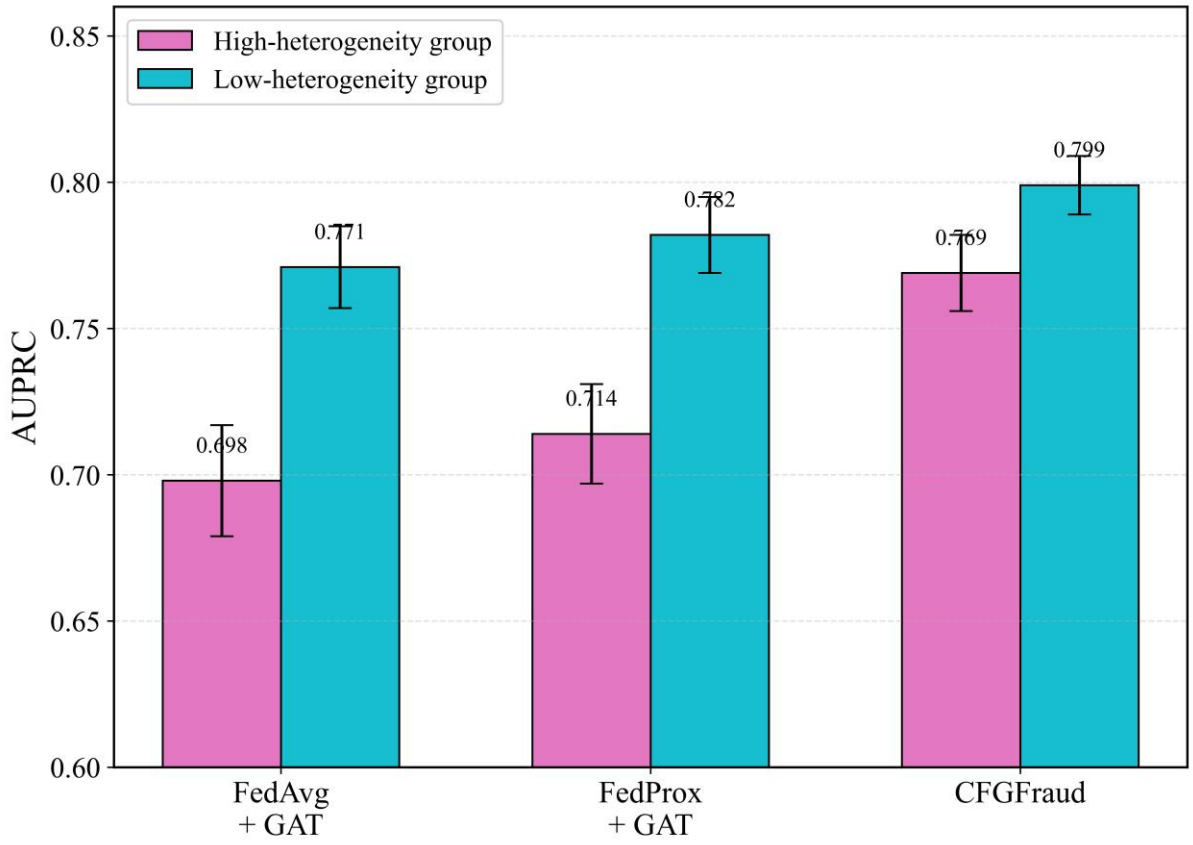


Figure 2. Comparison of the AUPRC of the three federated methods on the high heterogeneity group and the low heterogeneity group

5.3 Ablation experiments

The ablation experiments removed the three core components of CFGFraud one by one to quantify the independent contribution of each component to the final performance. Table 7 presents the comparison between the full model and the three ablation variants in terms of AUPRC, fraud ring recall, and standard deviation of performance across clients. The removal of the virtual node mechanism has the most dramatic impact: AUPRC drops from 0.791 to 0.721, and fraud ring recall drops from 0.723 to 0.498, a decrease of 31.1%. Removing institutional attention results in AUPRC decreasing to 0.758 and fraud ring recall decreasing to 0.652, and the virtual node dilution effect is quantified here. After removing the attention consistency regularity, AUPRC decreases to 0.764, with a relatively moderate decrease, but the standard deviation of cross-client AUPRC expands from 0.012 to 0.029.

Table 7. Ablation experiments: component contributions to detection performance and aggregation consistency

Configuration	AUPRC	Fraud ring recall rate	Standard deviation across clients
CFGFraud full model	0.791	0.723	0.012
Remove the virtual node mechanism	0.721	0.498	0.018
Removing mechanism mask attention	0.758	0.652	0.015
Remove the attention consistency regularity	0.764	0.671	0.029

The numerical differences in Table 7 reveal the magnitude of component contributions but cannot present the distribution pattern of this contribution among different clients.

5.4 Robustness and sensitivity analysis

Changes in the number of clients have a systematic impact on the relative advantage of CFGFraud over the baseline approach. Figure 3 shows the AUPRC changes of the two methods with a double line chart when the number of clients increases from 3 to 5 and then to 8. The horizontal axis is the number of clients, and the vertical axis is AUPRC. The values of CFGFraud at the three points are 0.803, 0.791 and 0.772 respectively, corresponding to a decrease of 0.012 and 0.019; FedAvg+GAT were 0.761, 0.734 and 0.702, respectively, with corresponding decreases of 0.027 and 0.032. The slope difference of the two curves is clearly visible: the absolute value of the slope of CFGFraud is smaller than that of FedAvg+GAT in each interval, and the gap between the two expands from 0.042 to 0.057 and then to 0.070 with the increase of the number of clients.

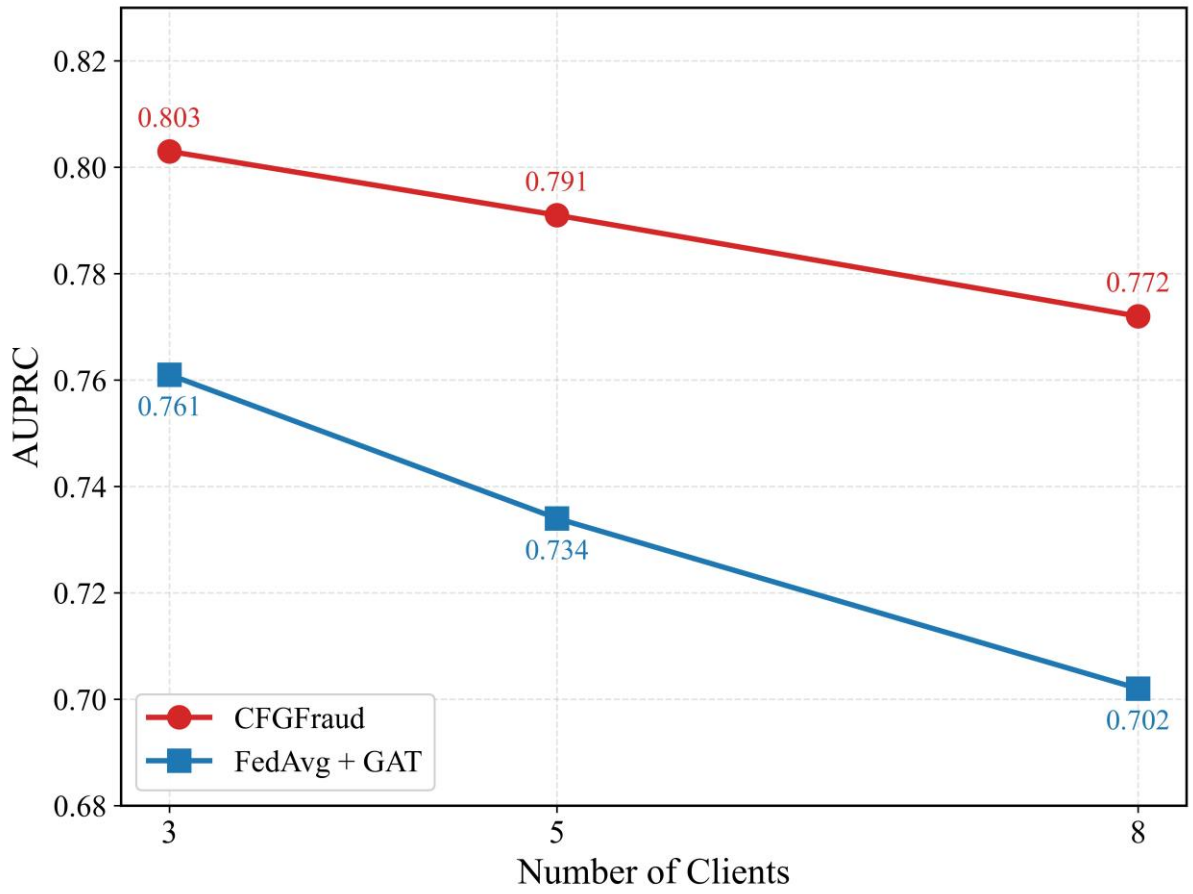


Figure 3. AUPRC trends of the two methods under the growth of the number of clients

The moderating effect of the proportion of dummy nodes on the effectiveness of IA-GAT shows a non-monotonic trend of first increasing and then decreasing. Figure 4 shows the change in the recall rate of the fraud loop in the process of the proportion of virtual nodes increasing from 5% to 30% with a line chart.

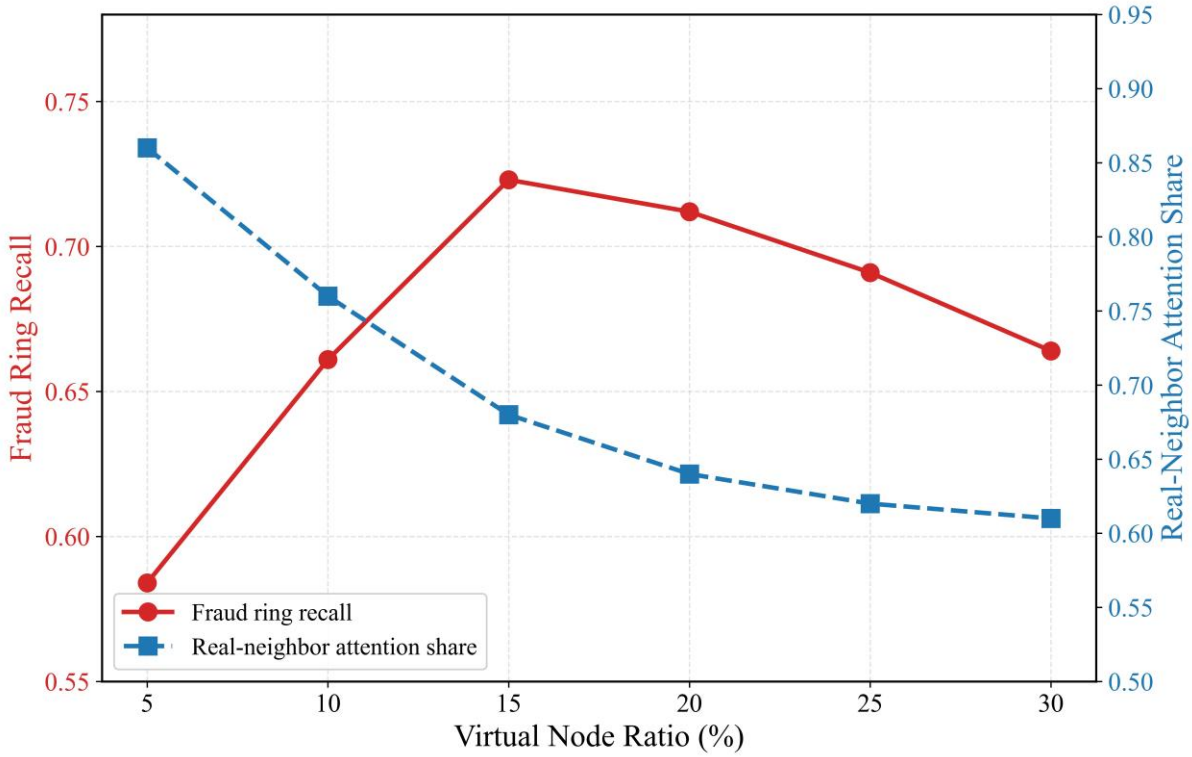


Figure 4. Influence of virtual-node proportion on fraud-ring recall and the attention ratio of real neighbors

The compression test of attention summary dimensions shows that AUPRC only decreases by 0.008 when the 6-dimensional summary is compressed to 3 dimensions and decreases by 0.019 when the 6-dimensional summary is compressed to 2 dimensions. Figure 5 compares the AUPRC with the cross-client standard deviation under the three summary dimensions with bar graphs, where the bar height is the AUPRC and the error bar is the cross-client standard deviation. The corresponding AUPRC of 6, 3 and 2 dimensions are 0.791, 0.783 and 0.772 respectively, and the cross-client standard deviation is 0.012, 0.014 and 0.023 respectively. The impact of dimension compression on the mean is limited, but the standard deviation of the 2-D summary shows a visible expansion, which indicates that the too low dimension loses the ability to distinguish the attention difference between edge types, and then weakens the accuracy of alignment. The tolerance of the summarization mechanism to moderate compression means that the summarization dimension can be reduced to 3 dimensions without significant loss of alignment effect in communication sensitive scenarios.

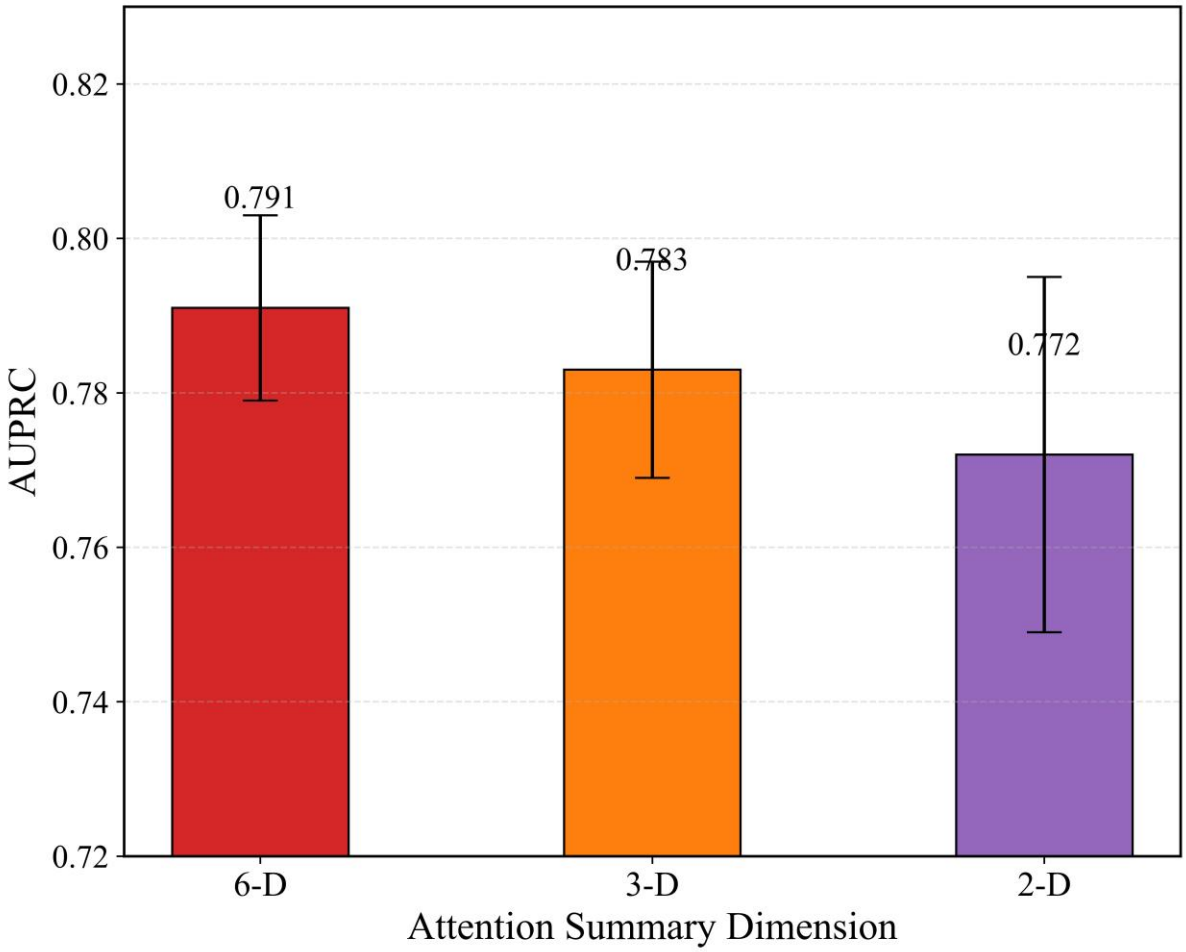


Figure 5. Influence of attention summary dimension on AUPRC and cross-client consistency

Based on the above experimental results, the gain of cross-institution collaboration does not exist unconditionally. When there are considerable shared entities among institutions and the degree of heterogeneity is high, the virtual node and attention alignment mechanism of CFGFraud can amplify the cooperation gain to the level of 0.057 higher than AUPRC. When entity overlaps across institutions is sparse or heterogeneity is low, the gain narrows to less than 0.028. The value of collaboration is mainly reflected in the improvement of the consistency of the global model on each client, rather than the significant improvement of the average accuracy. The clarity of this boundary condition provides an operational criterion for the deployment decision of the method in real insurance alliances.

6. CONCLUSIONS AND PROSPECTS

This paper focuses on a clear constraint: on the premise that the original claim data cannot be moved across institutions, whether each insurance institution can still learn the representation of fraud relationship with cross-institution discrimination on its own incomplete relationship graph. CFGFraud says yes, but with clear boundaries and costs. Under the non-independent and identically distributed partition of five clients, the AUPRC of deployable schemes is improved from 0.734 to 0.791 in standard federated aggregation, and the recall rate of fraud ring is improved from 0.581 to 0.723, leaving a gap of 0.056 between the non-deployable centralized upper bound of 0.847. The root of this gap does not lie in the strength of the aggregation strategy, but in the fact that the attribute information carried by the

unreachable neighbor cannot be physically obtained by the local model, and any federation scheme can only approach rather than eliminate this information loss at the level of structural signal.

The three technical modules of CFGFraud each respond to a separate failure link in the cross-agency graph collaboration. The virtual node mechanism solves the problem that the local model completely loses the perception of external connection after the cross-institution shared entity is directly truncated. It does not transfer attributes, but only transfers existence, so that the attention layer of the local graph knows at least that a certain edge leads to an unreachable external hub. The graph attention aggregation module of institutional awareness solves the problem of virtual placeholding nodes diluting real neighbor signals in attention competition. By parameterizing the attention separation of real connection and virtual connection and imposing learnable attenuation coefficient, structural signals can be retained and interference can be suppressed. The federated attention alignment aggregation strategy solves the problem that the systematic attention pattern deviation of each institution due to the heterogeneity of graph structure leads to the degradation of the global model on each client after parameter averaging. By collecting low-dimensional attention summaries and applying consistency constraints in the aggregation stage, the global model is pulled back to the consensus region of the attention pattern of each institution. The role dimensions of the three do not overlap with each other: the first addresses relational visibility, the second addresses aggregation quality, and the third addresses cross-institutional equilibrium.

The experimental results provide actionable criteria for the applicable conditions of cross-agency collaboration. The need for attention alignment aggregation is particularly prominent in client combinations with high heterogeneity. The gain relative to standard federated aggregation is 0.071 in the high heterogeneity group, narrowed to 0.028 in the low heterogeneity group, and the standard deviation across clients is compressed from 0.029 to 0.012. This shows that the core value of this mechanism is to suppress pattern drift rather than to improve pattern quality. The effectiveness of the virtual node mechanism is the most significant in the ablation, and the recall rate of the fraud ring decreases by 31.1% after removal, and the most severely damaged is the client with a high proportion of nodes in the maintenance factory, which directly confirms the supporting role of the cross-institutional hub entity in the connectivity of the fraud ring. The collaboration gain is positively correlated with the degree of entity overlap federated graph collaboration benefits the most when there are substantial shared repair shops, medical institutions, or beneficiaries among institutions. When the entity overlaps are sparse, the collaboration degenerates into a regularization method, and its value is mainly reflected in the improvement of the consistency of the global model in each mechanism rather than the significant improvement of the average accuracy. As the number of clients increases from 3 to 8, the AUPRC of CFGFraud decreases by 0.031, compared with 0.059 for the baseline, indicating that the virtual node mechanism captures more cross-institutional hubs when fragmentation increases, partially offsetting the negative effect of graph fragmentation. In terms of cost, the cost of single round computation increases by about 8% compared with pure federated aggregation, and the total communication volume increases by less than 1/10,000 due to the summary vector. This exchange is reasonable and necessary under the constraints of compliance rigidity.

There are three limitations of the method in this paper that need to be clearly recognized. The virtual node only transmits the existence signal, but cannot transmit the attribute information of the cross-institutional entity, which means that when a cross-institutional hub entity has a high intensity of fraud association on the external institutional side, but its local attribute performance is mediocre, the local model cannot take advantage of the strength difference, and can only know the connection here at the structural level without knowing the connected components. This information bottleneck is the

direct cost of privacy constraints, and it is difficult to break through without decrypting attributes in the short term. The experiment is based on simulated multi-institution division. Although the range of fraud rate and entity type deviation are deliberately enlarged to approximate the heterogeneity degree of the real alliance, the data distribution in the real insurance alliance may be more extreme. For example, the proportion of a certain type of entity in the graph of some regional institutions is more than 70%, or the difference in fraud rate between institutions is more than ten times. The performance of methods under such extreme conditions has not been tested. Although the attention digest is a six-dimensional low-dimensional vector and processed by differential privacy noise, its privacy risk has not been formally analyzed in this paper. Whether the digest vector may be reversed to infer sensitive information about the distribution of local graph edge types needs special privacy budget accounting and attack experiments to define.

Subsequent research can be advanced along three paths. The privacy-preserving embedding matching of cross-institution entity alignment is the key direction to break through the bottleneck of attribute information. Its goal is to enable the local model to obtain the attribute level representation of cross-institution shared entities rather than existence signals only through the intersection of privacy-preserving sets or homomorphic encryption embedding similarity calculations without exposing entity identification and attribute. This will directly narrow the gap from the centralized upper bound. Incrementally federated graph update for dynamic claims flow is a necessary extension for practical deployment. Insurance claims data continuously arrive in stream mode, the structure and size of local graph evolve over time, the rhythm coordination between federated rounds and incremental graph update, and the stability maintenance of attention summary in graph evolution are all problems not covered under the current static graph assumptions. Deployment verification and compliance evaluation within a real insurance alliance framework are the final step from experiment to application. It is necessary to test the end-to-end performance of the method under the conditions of real institutional combination, real privacy regulation explanation and real fraud label quality, and incorporate the privacy accounting of attention summary into the formal process of compliance evaluation. These three paths correspond to the three dimensions of information bottleneck, time dynamics and deployment compliance respectively, and constitute the subsequent route from feasible to available method of this paper.

Abbreviations

Edge-GAT, Edge graph attention;

FedAttAlign, Federated attention alignment aggregation policy;

GNN, graph neural network;

IA-GAT, Graph attention aggregation module for agency awareness;

non-IID, non-independent and identically distributed.

Supplementary Material

Not applicable.

Appendix

Not applicable.

Ethics approval and consent to participate.

This study did not involve human participants, animal subjects, or any data requiring ethical approval. Therefore, ethics approval and consent to participate are not applicable.

Acknowledgements

The authors would like to thank the editors of this journal and all the anonymous reviewers who provided valuable comments on this work.

Competing interests

The authors declare that they have no financial or personal relationships that may have inappropriately influenced them in writing this article.

Author contributions

All authors have read and agreed to the published version of the manuscript. The author's contributions are specified as follows: **C.Z.:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data Curation, Writing – original draft, Writing – review & editing, Visualization, Supervision, Project administration.

Funding information

The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.

Data availability

The data that support the findings of this study are available upon request from the corresponding authors, **C.Z.**

Disclaimer

The views and opinions expressed in this article are those of the authors and are the product of professional research. It does not necessarily reflect the official policy or position of any affiliated institution, funder, agency, or that of the publisher. The authors are responsible for this article's results, findings, and content.

Declaration of AI and AI-assisted Technologies in the Writing Process

During the writing of this article, the author used DeepSeek for spelling and grammar checking. After using this tool, the author reviewed and edited the content as needed and assumes full responsibility for the final published content.

REFERENCES

[1] Kumar, A. K., Nivethitha, T., Poonguzhali, P. K., & Saranyanandhini, D. (2026). Zero-Knowledge Proofs for Model Integrity and Privacy. *Decentralized Collaborative Learning for Data Privacy and Security*, 211-249. https://doi.org/10.1002/9781394464746.ch10

[2] Mudannayake, O., Indika, A., Jayasinghe, U., Lee, G. M., & Alawatugoda, J. (2025). On privacy-preserved machine learning using secure multi-party computing: Techniques and trends. *Computers, Materials & Continua*, *85*(2), 2527-2578. https://doi.org/10.32604/cmc.2025.068875

[3] Baduwal, M., Paudel, P., & Chaudhary, V. (2026). Federated learning: A survey of core challenges, current methods, and opportunities. Computers, 15(3), 155. https://doi.org/10.3390/computers15030155

[4] Yazdinejad, A., & Kong, J. D. (2026). Breaking interprovincial data silos: how federated learning can unlock Canada's public health potential. *Royal Society Open Science*, *13*(3). https://doi.org/10.1098/rsos.250959

[5] Özdemir, N., & Gökçen, H. (2026, July). Interpretable Academic Team Formation on Heterogeneous Information Networks: Constructive Heuristics with Explicit Organizational Affiliation. In *Informatics* (Vol. 13, No. 7, p. 107). MDPI. https://doi.org/10.3390/informatics13070107

[6] Gross, C., & Siklos, P. L. (2020). Analyzing credit risk transmission to the nonfinancial sector in Europe: A network approach. *Journal of Applied Econometrics*, *35*(1), 61-81. https://doi.org/10.1002/jae.2726

[7] Shan, C., Wang, C., & Wang, J. (2025, May). Regularized Incremental Federated Learning for Fraud Detection in Vehicle Insurances. In *2025 28th International Conference on Computer Supported Cooperative Work in Design (CSCWD)* (pp. 1832-1837). IEEE. https://doi.org/10.1109/cscwd64889.2025.11033178

[8] Zhao, T., Chen, G., Pang, C., Seenoi, P., Papukdee, N., Busababodhin, P., & Du, Y. (2025). Hybrid convolutional neural network-graph attention network-gradient boosting decision tree model for seismic impedance inversion prediction. *Journal of Seismic Exploration*, *34*(5), 81. https://doi.org/10.36922/JSE025310051

[9] Chen, G., Zhao, T., Pang, C., Seenoi, P., Papukdee, N., & Busababodhin, P. (2025). An attention-guided graph neural network and U-Net++-based reservoir porosity prediction system. *Journal of Seismic Exploration*, *34*(4), 70. https://doi.org/10.36922/JSE025300044

[10] Jing, S., Zhang, L., & Wang, X. (2026). MM-TFDX: a multimodal graph neural network framework for dynamic cross-border tax fraud detection with regulation-aware explainability. *Scientific Reports*. https://doi.org/10.1038/s41598-026-64839-4

[11] Zhao, T., Chen, G., Li, J., Pang, C., Seenoi, P., Papukdee, N., & Busababodhin, P. (2026). XGBoost-Deep Residual (FedXGB-ResNet) collaborative porosity prediction in a federated learning framework. *Journal of Seismic Exploration*, *35*(2), 201. https://doi.org/10.36922/JSE025440094

[12] Zeng, X. (2026). Cross-Border Trade Fraud Detection via Integrated Heterogeneous Graph Neural Network and XGBoost. *International Scientific Technical and Economic Research*, *4*(1), 47-69. https://doi.org/10.71451/ISTAER2603

[13] Khan, W., Ishrat, M., Faisal, S. M., Ahmad, F., Nabilal, K. V., & Sharma, Y. K. (2025). A robust framework for topology-based anomaly detection in attributed networks using graph attention networks, substructure analysis, and data augmentation. *International Journal of Data Science and Analytics*, *20*(7), 6715-6728. https://doi.org/10.1007/s41060-025-00848-2

[14] Wang, Y., Lei, J., Li, Y., & Shang, F. (2025). GATR: a graph attention deep reinforcement learning approach with variance-sensitive rewards for rpl routing optimization. *Journal of King Saud University Computer and Information Sciences*, *37*(9), 274. https://doi.org/10.1007/s44443-025-00266-1

[15] Zhao, T., Chen, G., Pang, C., Seenoi, P., Papukdee, N., & Busababodhin, P. (2025). Time-lapse earthquake difference prediction based on physics-informed long short-term memory coupled with interpretability boosting. *Journal of Seismic Exploration*, *34*(3), 25. https://doi.org/10.36922/JSE025310049

[16] Li, Y., & Ma, X. (2026). Intelligent delineation algorithm of urban development boundary based on graph neural network. *International Scientific Technical and Economic Research*, *4*(1), 123-148. https://doi.org/10.71451/ISTAER2606

[17] Lin, X., Zhao, X., Yin, Z., Pan, H., Jing, P., Li, W., ... & Shu, Y. (2025). Knowledge-driven federated learning: A systematic literature review on approaches, challenges, and prospects: X. Lin et al. *The Journal of Supercomputing*, *81*(8), 1016. https://doi.org/10.1007/s11227-025-07516-z

[18] Shen, G., Guo, X., Yuan, H., Du, J., Zhang, W., Qu, Y. N., & Kong, X. (2026). Communication-efficient decentralized federated graph learning via knowledge distillation under dual heterogeneity. *Information Sciences*, 123192. https://doi.org/10.1016/j.ins.2026.123192

[19] Yue, W., Lai, T., Mao, Q., Li, Q., & Camacho, D. (2026). A Review of Federated Learning Under Data Heterogeneity. *Expert Systems*, *43*(6), e70271. https://doi.org/10.1111/exsy.70271

[20] Yan, C., Lu, X., Lio, P., Hui, P., & He, D. (2024). Self-Simulation and Meta-Model Aggregation-Based Heterogeneous-Graph-Coupled Federated Learning. *IEEE Internet of Things Journal*, *12*(1), 198-212. https://doi.org/10.1109/JIOT.2024.3462724

[21] Yang, H., Wang, M., Dai, L., Wu, Y., & Du, J. (2025, August). Federated graph neural networks for heterogeneous graphs with data privacy and structural consistency. In *2025 5th International Conference on Computer Science and Blockchain (CCSB)* (pp. 311-315). IEEE. https://doi.org/10.1109/ccsb66722.2025.11154261

[22] Hu, Y., & Yan, T. (2026). Beyond Simple Aggregation: Decoupling Frequency-Aware Representation Learning and Gradient Boosting for Graph Anomaly Detection. *IEEE Access*. https://doi.org/10.1109/access.2026.3714218

[23] Wang, X., Wiegand, T., Kraft, N. J., Swenson, N. G., Davies, S. J., Hao, Z., ... & Wolf, A. (2016). Stochastic dilution effects weaken deterministic effects of niche-based processes in species rich

forests. *Ecology*, *97*(2), 347-360. https://doi.org/10.1890/14-2357.1

[24] Mukhoti, J., Kulharia, V., Sanyal, A., Golodetz, S., Torr, P., & Dokania, P. (2020). Calibrating deep neural networks using focal loss. *Advances in neural information processing systems*, *33*, 15288-15299. https://doi.org/10.48550/arXiv.2002.09437

[25] Ma, J., Xiang, S., Li, Q., Yuan, L., Cheng, D., & Jiang, C. (2024). Parallel graph learning with temporal stamp encoding for fraudulent transactions detections. *IEEE Transactions on Big Data*, *11*(4), 1945-1958. https://doi.org/10.1109/TBDATA.2024.3499338

[26] Devi, R. R., Raja, J. E., & Chin, Y. B. (2025). Reinforcement learning with graph neural network (RL-GNN) fusion for real-time financial fraud detection: a context-aware community mining approach. *Scientific Reports*, *15*(1), 42953. https://doi.org/10.1038/s41598-025-25200-3

[27] Guo, J., Li, Z., Feng, X., Zhong, Z., & Bao, R. (2026). Label-Safe Temporal Graph Attention for Financial Fraud Detection in Transaction Networks. *IEEE Access*. https://doi.org/10.1109/access.2026.3752923

APA Citation

Zhao, C. (2026). Cross-Institution Cooperative Detection Model of Insurance Fraud Based on Graph Neural Networks and Federated Learning. *Journal of Digital Frontier, 1*(3), 1-25. **https://doi.org/10.67541/JDF2613**