FEDERATED GRAPH FRAUD DETECTION WITH SPECTRAL ADAPTATION AND CONFIDENCE-AWARE AGGREGATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Graph-based Fraud Detection (GFD), which aims to identify fraud patterns within interaction networks, has attracted significant interest in trustworthy machine learning. However, real-world privacy regulations distribute graph data across decentralized clients, where fraudsters exhibit severe spectral heterogeneity both within and across these silos. Consequently, existing centralized fraud detectors relying on static filter banks fail to capture intricate patterns. More critically, fraudsters deliberately isolate themselves to evade detection, rendering any graph filter that operates via message-passing schemes completely ineffective. To address these challenges, we propose a federated learning framework featuring two major techniques: (1) **FedSACA**, which employs dynamic spectral channel attention to handle intra-client heterogeneity, and performs per-frequency confidence-aware aggregation on filter coefficients to tackle inter-client heterogeneity. (2) **FedSACA+**, which extends FedSACA with a feature-induced structure augmentation subject to the frequency confidence evaluation, thereby preventing the introduction of potential noise. Extensive experiments on multiple datasets demonstrate the superior performance of our proposed techniques.

1 INTRODUCTION

With the rapid progression of information technology, fraudulent behavior has expanded across financial systems (Lin et al., 2021), review networks (Rayana & Akoglu, 2015; Dou et al., 2020) and social media platforms (Deng et al., 2022), making the detection of such activities critically important. Graphs, which demonstrate the real-world complex interactions, serve as a prevalent tool for fraud detection. To fully utilize this relational information, Graph Neural Networks (GNNs) leverage the graph structure to learn rich representations, making them perfect for resolving the difficulties of Graph Fraud Detection (GFD).

However, most existing GFD methods focus on centralized training (Xu et al., 2024; Zhuo et al., 2024; Pan et al., 2025), implicitly assuming access to a complete, global view of the interaction network. In real-world high-stakes scenarios, such as interbank transaction monitoring or cross-platform social analysis (Dou et al., 2020; Liu et al., 2019; 2018), graph data is inevitably fragmented across distributed clients due to strict privacy regulations (Yang et al., 2019; Kairouz & McMahan, 2021) and competitive barriers. This data isolation renders centralized GFD methods impractical. Federated Learning (FL) (Li et al., 2020a; McMahan et al., 2017b), combined with GNNs, specifically Federated Graph Learning (FGL) (Liu et al., 2026; Zhang et al., 2021b), has emerged as a compelling paradigm, enabling collaborative model training without exchanging raw graph data or sensitive topology (Zhang et al., 2021b; Yao et al., 2023). However, applying FGL to fraud detection introduces two significant topology-related challenges posed to previous GFD methods.

From the perspective of fraud pattern distributions in the federated setting, we characterize these deceptive topologies as underlying frequency signatures and identify the first challenge: **Fraud Spectral Heterogeneity**. We decouple this phenomenon into two dimensions: (1) **Intra-Client Heterogeneity**: Within a single client, fraud patterns are usually not confined to a single band but exhibit signals spanning **diverse**, **non-contiguous frequencies**. This arises from various connectivity patterns among fraudsters and benign nodes within the local graph, as reflected in the fraud

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

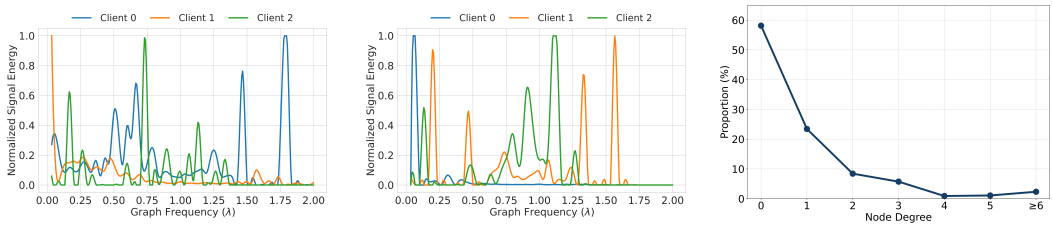

(a) YelpChi **(b) Amazon** **(c) Fraudulent Node Degree**

Figure 1: Phenomenon Demonstration. (a-b) Normalized signal energy distributions of fraudsters across graph frequencies λ on YelpChi and Amazon. Fraudsters exhibit **diverse**, **non-contiguous** spectral profiles within a single client and the distributions are **conflicting** across clients, highlighting severe **fraud spectral heterogeneity**. The signal energy calculation is detailed in Appendix. (c) The fraudulent node degree statistics on FDCompCN. The proportion exhibits a cliff-like drop, with nearly 60% of fraudsters being completely isolated (degree = 0). This validates the severe **structural isolation of fraudulent nodes** in real-world networks.

spectral distributions (Figure 1a 1b). Unfortunately, due to the spectral filter rigidity, existing detectors (Tang et al., 2022; Xu et al., 2024) are inadequate to capture such complicated patterns. Even the recent method (Guo et al., 2025) still relies on unstable pseudo-labels to switch spectral filters. (2) **Inter-Client Heterogeneity**: As demonstrated in Figure 1a 1b, the fraud spectral distributions across clients are often **conflicting** rather than consistent. During aggregation, such spectral conflicts wash out distinct local fraud patterns, yielding a compromised, flattened global filter. Consequently, upon local deployment, this rigid frequency response fails to adapt to various client distributions, leading to poor generalization. Driven by these two dimensions, we raise question **I**): *How can we effectively capture diverse fraud patterns under such* ***distinct intra- and inter-client heterogeneity*** *without accessing raw data or sensitive topology?*

From the perspective of fraudster concealment, the second challenge is the structural isolation of fraudulent nodes (Figure 1c). Specifically, to evade detection, fraudsters abandon disposable accounts following transient malicious activities. This deliberate isolation severely disrupts the GNN message-passing scheme, rendering the filters entirely ineffective on disconnected nodes as it severs signal propagation pathways. Furthermore, naively mitigating this isolation by relying solely on forced new connections inevitably propagates deceptive noise from fraudsters, thereby undermining detection effectiveness. The above analysis leads to question **II**): *How to address* ***fraudulent node isolation*** *under privacy constraints while minimizing the introduction of noise?*

To address the first challenge, we propose **Federated** Graph Fraud Detection with **Spectral Adaptation and Confidence-Aware Aggregation** (FedSACA), which rethinks filter adaptation and fraud spectral-domain aggregation in the federated setting. (1) To resolve Intra-Client Heterogeneity, we conceptualize filtering as a flexible spectral channel attention process. By equipping each frequency band with a self-learnable confidence score, our local models adaptively emphasize or suppress individual frequencies to fit their own fraud spectral distributions. (2) To tackle Inter-Client Heterogeneity, given that achieving consensus on the overall spectral distribution is practically unattainable, we focus instead on acquiring reliable frequency-specific knowledge. To explicitly mitigate spectral conflicts, we first align local updates by identifying universally recognized fraud frequency bands. Moreover, inspired by expert-driven parameter estimation (Zhou et al., 2022), we aggregate filter coefficients by dynamically scaling the contribution of each client per frequency band based on its confidence score. This ensures the global filter remains highly reliable across all bands, thus adapting to diverse local spectral distributions. Through this strategy, FedSACA enhances global model generalization while empowering adaptive local filters to capture distinct fraud patterns without compromising privacy.

To tackle the second issue, we extend FedSACA to **FedSACA+** by introducing a feature-induced structure augmentation. Motivated by the K Nearest Neighbor (kNN) method (Cover & Hart, 1967), each client introduces edges based on local feature distances before training. While local feature proximity can reliably reconnect isolated nodes, unregulated edge additions risk propagating deceptive noise from fraudsters. Our key insight is that augmented connections aid detection only if their frequency responses reflect the spectral distributions of originally isolated fraud patterns. Specifically, we first incorporate the augmented edges as an independent relation initialized with a strong negative confidence across all frequency bands to suppress the potential noise propagation. The

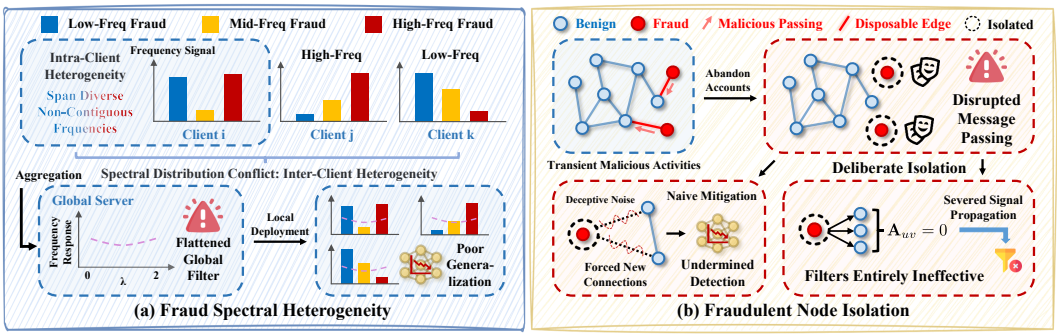

Figure 2: Problem illustration. (a) **Fraud Spectral Heterogeneity**: Fraud patterns span diverse, non-contiguous frequencies within clients and exhibit severe spectral distribution conflicts across clients, leading to a flattened global filter and poor local generalization. (b) **Fraudulent Node Isolation**: Deliberate account abandonment disrupts GNNs message passing, rendering filters entirely ineffective. Furthermore, naively forcing new connections propagates deceptive noise that undermines detection.

confidence is then adaptively updated by the local spectral attention toward the specific distribution that benefits detection. In synergy, FedSACA+ maintains strong fraud detection performance under both fraud spectral heterogeneity and fraudulent node isolation. Our principal contributions are summarized as follows.

- We are the first to reveal fraud pattern heterogeneity from an innovative spectral perspective, and to recognize fraudulent node isolation as a fundamental detection obstacle.
- We propose FedSACA, which overcomes fraud spectral heterogeneity by ensuring that filters remain reliable across all frequencies. Furthermore, FedSACA+ restores filter effectiveness on isolated fraudsters via noise-free structure augmentation.
- We conduct experiments under both severe heterogeneity and extreme structural sparsity, validating the superiority of our proposed framework.

2 Related Work

Graph-based Fraud Detection (GFD) primarily relies on spatial (Dou et al., 2020; Peng et al., 2021; Wang et al., 2023) and spectral (Wu et al., 2023; Xu et al., 2024; Zhuo et al., 2024) Graph Neural Networks (GNNs) to capture complex heterophilic fraud patterns. However, in privacy-preserving scenarios, these centralized methods fail due to data isolation. While Federated Graph Learning (FGL) (Xie et al., 2021; Tan et al., 2023; Zhou et al., 2025) enables collaborative training and explores topology recovery (Zhang et al., 2021a; Liu et al., 2025), applying standard FGL aggregation to adaptive spectral filters (Defferrard et al., 2016; He et al., 2021b) blindly averages conflicting frequency responses. This destroys local adaptivity and exacerbates inter-client spectral conflicts. Moreover, existing structure augmentation techniques (Chen et al., 2020; Zhang et al., 2024) designed to alleviate graph sparsity often naively connect isolated, camouflaged fraudsters to benign nodes, propagating deceptive noise. To address these critical gaps, we propose FedSACA to bridge the underexplored domain of federated spectral fraud detection via a confidence-aware spectral attention framework that dynamically reconciles diverse fraud frequency profiles. Furthermore, its extension, FedSACA+, rigorously controls feature-induced structure augmentation through frequency confidence evaluation, safely repairing fraudulent node isolation without introducing noise. A comprehensive review of the related literature is provided in Appendix A.

3 Preliminaries

3.1 Graph Fraud Detection

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ be a global heterogeneous graph, where $\mathcal{V}$ is the set of nodes and $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times d}$ is the node feature matrix with dimension d. Real-world fraud patterns are typically multi-relational, thus the total edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ comprises multiple relation-specific edge sets $\mathcal{E}^{(r)}$ based on

different relation types $r \in \mathcal{R}$, where R is the relation set. For each relation r, the topological structure is represented by the adjacency matrix $\mathbf{A}^{(r)} \in \mathbb{R}^{|\mathcal{V}|\times|\mathcal{V}|}$, where $\mathbf{A}^{(r)}_{uv} = 1$ denotes the existing edge $e^{(r)}_{uv} \in \mathcal{E}^{(r)}$ and $\mathbf{A}^{(r)}_{uv} = 0$ otherwise. The normalized form of the adjacency matrix is defined as $\hat{\mathbf{A}}^{(r)} = (\mathbf{D}^{(r)})^{-1/2}\mathbf{A}^{(r)}(\mathbf{D}^{(r)})^{-1/2}$, where the diagonal degree matrix $\mathbf{D}^{(r)}$ is given by $D^{(r)}_{uu} = \sum_v \mathbf{A}^{(r)}_{uv}$. The objective of graph fraud detection is to learn a mapping function $f\colon (\cdot,\cdot) \to \mathbf{Y}$ that outputs node-level binary labels:

$$f\colon (\mathbf{X}, \{\hat{\mathbf{A}}^{(r)}\}_{r\in\mathcal{R}}) \to \mathbf{Y}, \tag{1}$$

where $\mathbf{Y} \in \{0,1\}^{|\mathcal{V}|}$, with label 1 indicating fraudulent nodes and label 0 for benign nodes (Jing et al., 2021).

3.2 Federated Graph Learning

In the federated setting, the global graph $\mathcal{G}$ is distributed across N clients. Each client $i \in \{1,2,\ldots,N\}$ holds a local subgraph $\mathcal{G}_i = (\mathcal{V}_i, \mathcal{E}_i, \mathbf{X}_i)$ with $n_i = |\mathcal{V}_i|$ nodes and adjacency matrices $\{\hat{\mathbf{A}}^{(r)}_i\}_{r\in\mathcal{R}_i}$. The goal of Federated Graph Learning (FGL) is to collaboratively optimize the global model parameters $\mathbf{\Theta}$ without sharing raw local graph data. The collaborative objective minimizes:

$$\min_{\mathbf{\Theta}} \mathcal{F}(\mathbf{\Theta}) = \min_{\mathbf{\Theta}} \sum_{i=1}^{N} w_i \mathcal{L}_i(\mathbf{\Theta}), \tag{2}$$

where $\mathcal{F}(\mathbf{\Theta})$ represents the global optimization objective, measuring the weighted average loss across all clients. $\mathcal{L}_i(\cdot)$ denotes the local loss function, and $w_i \geq 0$ is the weight assigned to client i, satisfying $\sum_i w_i = 1$. During communication round t, each client i receives the global model $\mathbf{\Theta}^t$ and performs local optimization $\mathbf{\Theta}^{t+1}_i \leftarrow \mathbf{\Theta}^t - \eta\nabla_{\mathbf{\Theta}}\mathcal{L}_i(\mathbf{\Theta}^t)$, where η is the learning rate and $\nabla_{\mathbf{\Theta}}\mathcal{L}_i$ denotes the local gradient. Subsequently, the server aggregates the updated parameters $\mathbf{\Theta}^{t+1}_i$ by weighted averaging to construct the new global model:

$$\mathbf{\Theta}^{t+1} = \sum_{i=1}^{N} w_i \mathbf{\Theta}^{t+1}_i. \tag{3}$$

3.3 Spectral Graph Filter

Spectral Graph Neural Networks (GNNs) capture graph signals $\mathbf{X}$ by applying spectral filters g_θ. For each relation r, the Laplacian matrix is $\mathbf{L}^{(r)} = \mathbf{D}^{(r)} - \mathbf{A}^{(r)}$, while the normalized Laplacian matrix is defined as $\tilde{\mathbf{L}}^{(r)} = \mathbf{I} - \hat{\mathbf{A}}^{(r)}$, where $\mathbf{I}$ is the identity matrix. Employing eigendecomposition, we obtain $\tilde{\mathbf{L}}^{(r)} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$, in which $\mathbf{U}$ is an orthogonal matrix whose columns are the eigenvectors, and $\mathbf{\Lambda} = \mathrm{diag}(\lambda_0,\ldots,\lambda_{|\mathcal{V}|-1})$ is a diagonal matrix of eigenvalues satisfying $0 = \lambda_0 \leq \lambda_1 \leq \cdots \leq \lambda_{|\mathcal{V}|-1} \leq 2$. The spectral convolution is formulated as:

$$g_\theta(\tilde{\mathbf{L}}^{(r)})\mathbf{X} = \mathbf{U}g_\theta(\mathbf{\Lambda})\mathbf{U}^\top\mathbf{X}. \tag{4}$$

To avoid the prohibitive $\mathcal{O}(|\mathcal{V}|^3)$ eigendecomposition cost, modern GNNs approximate the filter using a truncated Chebyshev polynomial of order $K-1$. While standard implementations apply this polynomial to a strictly scaled Laplacian mapped by $\lambda_{|\mathcal{V}|-1}$, we directly utilize the normalized adjacency matrix $\hat{\mathbf{A}}^{(r)}$. Since the eigenvalues λ of $\tilde{\mathbf{L}}^{(r)}$ are bounded in $[0,2]$, the eigenvalues of $\hat{\mathbf{A}}^{(r)} = \mathbf{I} - \tilde{\mathbf{L}}^{(r)}$ are $1-\lambda \in [-1,1]$, which exactly satisfies the required domain of the Chebyshev polynomial. The efficient spectral filtering formulation is consequently computed as:

$$\mathbf{H}^{(r)} = \sum_{k=0}^{K-1} \theta^{(r)}_k T_k(\hat{\mathbf{A}}^{(r)})\mathbf{X}, \tag{5}$$

where $\mathbf{H}^{(r)} \in \mathbb{R}^{|\mathcal{V}|\times d}$ represents the filtered output graph signals. The vector $\theta^{(r)} = [\theta^{(r)}_0,\ldots,\theta^{(r)}_{K-1}] \in \mathbb{R}^K$ denotes the learnable filter coefficients. $T_k(\cdot)$ is the Chebyshev polynomial basis, which follows the recurrence $T_0(\hat{\mathbf{A}}^{(r)}) = \mathbf{I}$, $T_1(\hat{\mathbf{A}}^{(r)}) = \hat{\mathbf{A}}^{(r)}$, and $T_k(\hat{\mathbf{A}}^{(r)}) = 2\hat{\mathbf{A}}^{(r)}T_{k-1}(\hat{\mathbf{A}}^{(r)}) - T_{k-2}(\hat{\mathbf{A}}^{(r)})$. Although this polynomial formulation acts as a universal approximator, enforcing a shared $\theta^{(r)}$ across all clients in the non-IID federated setting inevitably causes severe spectral parameter conflicts, which directly motivates our adaptive confidence-aware design.

4 METHODOLOGY

4.1 FEDSACA

Motivation. Fraudsters often manipulate their local connectivity to blend in with benign nodes, making purely spatial aggregation unreliable. Spectral filters, by capturing fraud patterns in the frequency domain, offer a principled way to separate fraud from normal behavior. However, severe fraud spectral heterogeneity due to diverse fraud frequencies in Figure 1 significantly disrupts this process in the federated setting. To address this challenge existing in both intra- and inter-client dimensions, we propose FedSACA, which introduces spectral attention for local adaptation and further ensures global filter reliability by per-frequency confidence-aware aggregation.

Spectral Adaptation. To overcome intra-client heterogeneity, local models must dynamically adapt to diverse and non-contiguous frequency profiles that evade static filter banks. We redesign the standard Chebyshev polynomial filter by equipping it with an adaptive confidence mechanism. Specifically, we introduce an unconstrained, self-learnable mask $\mathbf{m}_i \in \mathbb{R}^K$ alongside the base filter coefficient $\theta_i \in \mathbb{R}^K$ for each client i. Rather than applying $\mathbf{m}_i$ directly, we map it to a normalized **frequency confidence score** $\mathbf{s}_i \in (0,1)^K$ via a sigmoid activation:

$$\mathbf{s}_i = \sigma(\mathbf{m}_i) = \frac{1}{1+\exp(-\mathbf{m}_i)}. \tag{6}$$

Here, each element $s_{i,k} \in \mathbf{s}_i$ quantifies how much client i relies on the k-th order frequency band to identify fraud patterns. Let $\mathbf{V}_{i,k} = T_k(\hat{\mathbf{A}}_i)\mathbf{X}_i$ denotes the corresponding k-th order frequency basis representation. By treating each frequency band as an independent spectral channel, the subsequent filtering naturally operates as a **spectral channel attention** process:

$$\mathbf{H}_i = \text{SpectralAttention}(\mathbf{m}_i, \theta_i, \mathbf{V}_i) = \sum_{k=0}^{K-1} s_{i,k} \cdot (\theta_{i,k}\mathbf{V}_{i,k}) = \sum_{k=0}^{K-1} (s_{i,k} \cdot \theta_{i,k})\, T_k(\hat{\mathbf{A}}_i)\mathbf{X}_i, \tag{7}$$

where the output $\mathbf{H}_i \in \mathbb{R}^{|\mathcal{V}_i| \times d}$ denotes the locally filtered node representations. In this paradigm, the confidence score $s_{i,k}$ acts as a self-learnable spectral channel attention weight that captures the client-specific spectral preference, independently gating each frequency band. Concurrently, each element $\theta_{i,k} \in \theta_i$ serves as the base filter coefficient that scales the frequency basis $\mathbf{V}_{i,k}$ to adapt to local fraud spectral distributions. Unlike traditional polynomial filters that solely rely on the base coefficient $\theta_{i,k}$, our frequency response is dynamically governed by the effective combined weight $s_{i,k} \cdot \theta_{i,k}$. $s_{i,k} \in (0,1)$ preserves the ability to recover any suppressed frequency band if subsequent federated updates support it, as illustrated in Figure 3 in details.

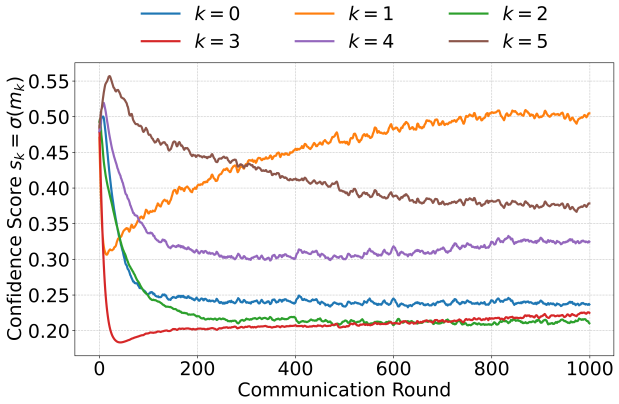

Figure 3: Evolution curves of global confidence scores s_k for the *rur* relation on YelpChi. The curves demonstrate a highly dynamic **spectral attention** process. While some informative bands (e.g., $k = 1$) are adaptively recovered after suppression, others (e.g., $k = 5$) are gradually pruned as the global model converges, effectively preventing from overfitting to specific frequency band.

To train the adaptive spectral filter, the local parameter set $\mathbf{\Theta}_i = \{\theta_i, \mathbf{m}_i\}$ is optimized by minimizing a local objective consisting of three distinct components. First, we employ the standard weighted cross-entropy loss $\mathcal{L}_{\text{CE}}^{(i)}$ over the local labeled training node set $\mathcal{V}_i^{\text{train}}$ to drive fraud detection:

$$\mathcal{L}_{\text{CE}}^{(i)} = -\frac{1}{|\mathcal{V}_i^{\text{train}}|} \sum_{v \in \mathcal{V}_i^{\text{train}}} \Big[w_{\text{pos}}\, y_v \log p_v + (1-y_v) \log(1-p_v) \Big], \tag{8}$$

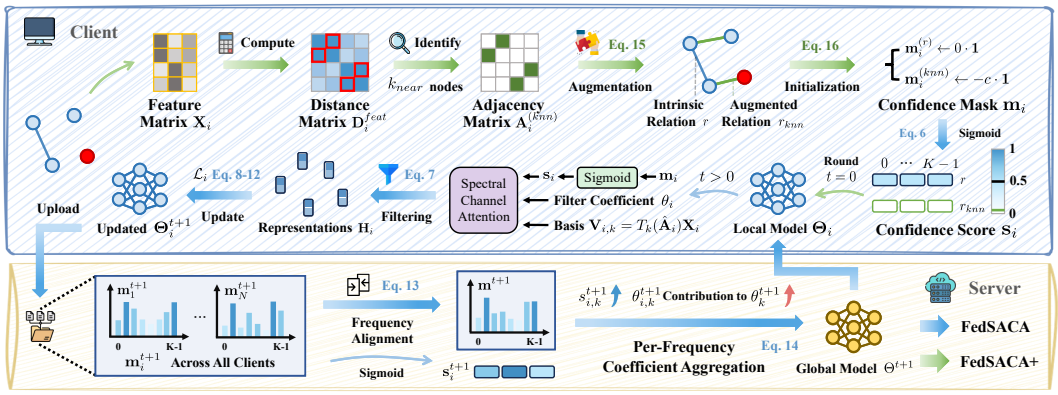

Figure 4: Architecture Illustration. We use blue and green arrows to represent the training pipelines of FedSACA and its extension FedSACA+, respectively. In FedSACA, local models utilize self-learnable **frequency confidence masks** $\mathbf{m}_i$ to adapt to diverse fraud spectral distributions. Then, the global server aggregates filter coefficients through **per-frequency confidence-aware aggregation**, which ensures the global filter reliability across all bands. In FedSACA+, to alleviate fraudulent node isolation, we construct an **augmented relation** by connecting edges based on local feature distances. Furthermore, we initialize the confidence mask of this extended relation with a **strong negative constant bias** to suppress the potential noise propagation.

where $y_v \in \{0, 1\}$ denotes the actual label of node v, and p_v is the predicted fraud probability derived from the filtered representations $\mathbf{H}_i$. The weight $w_{\text{pos}} = n_{\text{neg}}/n_{\text{pos}}$ is computed from the counts of benign nodes n_{neg} and fraudulent nodes n_{pos} within $\mathcal{V}_i^{\text{train}}$. This loss amplifies the contribution of scarce fraud samples, preventing the model from being dominated by the benign majority. Second, we impose an ℓ_1-norm regularization $\mathcal{L}_{\text{reg}}^{(i)}$ on the confidence scores. Since this sparsity constraint applies to the entire heterogeneous graph, we compute the sum across all relation types $r \in \mathcal{R}$:

$$\mathcal{L}_{\text{reg}}^{(i)} = \sum_{r\in\mathcal{R}} \|\mathbf{s}_i^{(r)}\|_1 = \sum_{r\in\mathcal{R}} \sum_{k=0}^{K-1} s_{i,k}^{(r)}, \tag{9}$$

where $\mathbf{s}_i^{(r)}$ denotes the confidence score vector specific to relation r. This regularization actively penalizes uninformative frequencies by pushing their scores toward zero, thereby forcing the local model to capture fraud patterns through only a few dominant spectral components and reducing the introduction of noise. Third, we leverage a proximal alignment term $\mathcal{L}_{\text{prox}}^{(i)}$:

$$\mathcal{L}_{\text{prox}}^{(i)} = \frac{\mu}{2}\|\mathbf{\Theta}_i - \mathbf{\Theta}^t\|_2^2, \tag{10}$$

where $\mathbf{\Theta}^t = \{\theta^t, \mathbf{m}^t\}$ denotes the broadcast global parameter set in communication round t, and $\mathbf{\Theta}_i$ encompasses the local parameters across $\mathcal{R}$. The hyperparameter μ controls the alignment strength. This term restricts the local model from overfitting to its noise, ensuring stable integration into the global consensus. Consequently, the overall local objective $\mathcal{L}_i$ is constructed as:

$$\mathcal{L}_i = \mathcal{L}_{\text{CE}}^{(i)} + \lambda\mathcal{L}_{\text{reg}}^{(i)} + \mathcal{L}_{\text{prox}}^{(i)}, \tag{11}$$

where λ governs the relative penalty of the frequency sparsity. During the local training phase, the local parameter set $\mathbf{\Theta}_i$ is jointly updated via standard gradient descent:

$$\mathbf{\Theta}_i^{t+1} \leftarrow \mathbf{\Theta}_i - \eta\nabla_{\mathbf{\Theta}_i}\mathcal{L}_i. \tag{12}$$

After completing the local training epochs, we denote the final optimized parameter set as $\mathbf{\Theta}_i^{t+1}$.

Confidence-Aware Aggregation. Following the local optimization phase in communication round t, clients upload their updated parameter sets $\mathbf{\Theta}_i^{t+1} = \{\theta_i^{t+1}, \mathbf{m}_i^{t+1}\}$ to the global server. Conventional FGL algorithms rigidly average these parameters weighted by local dataset sizes. However, such aggregation mixes conflicting frequency responses and thus fails to address inter-client heterogeneity. To address this, we introduce confidence-aware per-frequency aggregation, comprising frequency alignment and coefficient estimation as two key processes. First, we aggregate the mask $\mathbf{m}_i^{t+1}$ across all clients to mitigate spectral conflicts:

$$\mathbf{m}^{t+1} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{m}_i^{t+1}, \tag{13}$$

where $\mathbf{m}^{t+1}$ constitutes the updated global mask for round $t+1$. This arithmetic mean identifies universally recognized fraud frequency bands, aligning local updates before coefficient aggregation. Second, we aggregate the filter coefficients by weighting the contribution of each client per frequency band according to its locally learned confidence. Using the confidence score $\mathbf{s}_i^{t+1} = \sigma(\mathbf{m}_i^{t+1})$, the updated global filter coefficient θ_i^{t+1} for round $t+1$ is computed as:

$$\theta^{t+1} = \frac{\sum_{i=1}^{N} \mathbf{s}_i^{t+1} \odot \theta_i^{t+1}}{\sum_{i=1}^{N} \mathbf{s}_i^{t+1}}, \tag{14}$$

where $\odot$ denotes element-wise multiplication. This formulation ensures that a client coefficient $\theta_{i,k}^{t+1}$ with high confidence in a specific frequency band (i.e., $s_{i,k}^{t+1} \to 1$) dominates the corresponding global coefficient θ_k^{t+1}, while clients with low confidence contribute negligibly to it. Consequently, the global filter selectively inherits informative spectral knowledge from each client, thus remaining reliable across all frequencies. Together, these two steps allow FedSACA to establish a robust spectral consensus under severe inter-client heterogeneity while preserving local spectral adaptation.

4.2 FEDSACA+

Motivation. While FedSACA effectively resolves fraud spectral heterogeneity, a critical topology bottleneck remains due to the adversarial fraud behavior. Specifically, fraudsters deliberately sever their own connections after brief malicious activity, creating isolated fraudulent nodes that lack the neighborhood connectivity essential for message-passing schemes. Therefore, irrespective of how precisely filters adapt to spectral distributions, they remain powerless without frequency signals from these disconnected targets. Our solution, FedSACA+, addresses fraudulent node isolation through feature-induced structure augmentation. Moreover, this process is rigorously controlled by frequency confidence evaluation to suppress the potential noise propagation.

To eliminate the isolation of fraudulent nodes, we first construct augmented connections based on the local node feature $\mathbf{X}_i$. For each node $u \in \mathcal{V}_i$, we extract its corresponding feature $\mathbf{x}_{i,u} \in \mathbb{R}^d$ from $\mathbf{X}_i$ and compute Euclidean distance $\mathbf{D}_i^{feat}$ to identify its k_{near} number of nearest neighbor nodes $v \in \mathcal{V}_i$ in the feature space $\mathbb{R}^d$:

$$\mathbf{D}_i^{feat} = \|\mathbf{x}_{i,u} - \mathbf{x}_{i,v}\|_2. \tag{15}$$

To ensure strict compatibility with the spectral convolution, we define an undirected, symmetrized binary adjacency matrix $\mathbf{A}_i^{(knn)} \in \{0,1\}^{|\mathcal{V}_i| \times |\mathcal{V}_i|}$ without self-loops for the augmented relation r_{knn}:

$$\mathbf{A}_{i,uv}^{(knn)} = \begin{cases} 1, \text{ if } \big(\mathbf{x}_{i,v} \in \text{kNN}(\mathbf{x}_{i,u}) \vee \mathbf{x}_{i,u} \in \text{kNN}(\mathbf{x}_{i,v})\big) \text{ and } u \neq v \\ 0, \text{ otherwise} \end{cases}, \tag{16}$$

where $\text{kNN}(\mathbf{x}_{i,u})$ denotes the set of the k_{near} nearest feature vectors to node u in client i, and $\vee$ represents the logical OR operator. This feature-induced topology expands the original intrinsic relation set $\mathcal{R}$ to an augmented set $\mathcal{R}^* = \mathcal{R} \cup \{r_{knn}\}$, yielding a newly normalized adjacency matrix $\hat{\mathbf{A}}_i^{(knn)}$. Crucially, this new relation r_{knn} undergoes the exact same FedSACA pipeline alongside all other intrinsic relations. The sole algorithmic distinction lies in its initialization. While each entry of the mask for all intrinsic relations $r \in \mathcal{R}$ is initialized neutrally (i.e., $\mathbf{m}_i^{(r)} \leftarrow 0 \cdot \mathbf{1}$, $\sigma(0) = 0.5$), we initialize the mask for this augmented relation, $\mathbf{m}_i^{(knn)} \in \mathbb{R}^K$, with a strong negative constant bias to suppress the potential noise propagation:

$$\mathbf{m}_i^{(knn)} \leftarrow -c \cdot \mathbf{1}, \tag{17}$$

where $c > 0$ is a large scaling factor (e.g., $c = 3$) and $\mathbf{1}$ is the all-ones vector. Consequently, each entry of the initial confidence score vector $\mathbf{s}_i^{(knn)}$ is squashed to nearly zero (e.g., $\sigma(-3) \approx 0.047$). By fully leveraging the existing mathematical properties of FedSACA, this simple initialization naturally functions as a strict structural gate. However, since the gradient channels of $\mathbf{s}_i^{(knn)} \in (0,1)^K$ remain open, $\mathbf{m}_i^{(knn)}$ can still be adaptively reactivated if the local objective $\mathcal{L}_i$ reveals indispensable utility in specific frequency bands of $\hat{\mathbf{A}}_i^{(knn)}$. By building upon FedSACA in this noise-free way, FedSACA+ transforms vulnerable structure augmentation into a self-correcting hypothesis test, robustly addressing fraudulent node isolation without incurring significant extra model complexity.

5 EXPERIMENTS

Table 1: Overall performance comparison on YelpChi, Amazon, FDCompCN, and Elliptic datasets. The best and second-best results are highlighted in **bold** and underlined, respectively. OOT: out of time.

Methods	YelpChi		Amazon		FDCompCN		Elliptic	
	AUC	F1-Macro	AUC	F1-Macro	AUC	F1-Macro	AUC	F1-Macro
GCN	$54.21_{\pm 0.85}$	$51.34_{\pm 1.12}$	$78.45_{\pm 0.65}$	$60.12_{\pm 1.05}$	$52.14_{\pm 1.25}$	$45.32_{\pm 1.85}$	$74.56_{\pm 0.75}$	$64.23_{\pm 1.15}$
GraphSAGE	$70.36_{\pm 0.64}$	$61.25_{\pm 1.05}$	$84.52_{\pm 0.55}$	$73.45_{\pm 0.95}$	$55.67_{\pm 1.15}$	$48.21_{\pm 1.65}$	$77.89_{\pm 0.65}$	$68.75_{\pm 1.05}$
GAT	$55.45_{\pm 0.52}$	$52.88_{\pm 0.96}$	$79.11_{\pm 0.45}$	$61.34_{\pm 0.85}$	$52.89_{\pm 1.05}$	$45.98_{\pm 1.55}$	$75.12_{\pm 0.55}$	$64.88_{\pm 0.95}$
CARE-GNN	$80.07_{\pm 0.19}$	$66.59_{\pm 0.67}$	$94.98_{\pm 0.18}$	$88.34_{\pm 0.58}$	$62.36_{\pm 1.47}$	$55.42_{\pm 1.69}$	$86.33_{\pm 0.24}$	$76.93_{\pm 0.40}$
GAGA	$78.50_{\pm 0.35}$	$66.20_{\pm 0.82}$	$95.10_{\pm 0.35}$	$89.40_{\pm 0.70}$	$62.80_{\pm 0.85}$	$54.50_{\pm 1.35}$	$90.20_{\pm 0.45}$	$82.50_{\pm 0.80}$
PMP	OOT		OOT		OOT		OOT	
GPA	$84.72_{\pm 0.21}$	$65.33_{\pm 0.73}$	$94.98_{\pm 0.35}$	$90.28_{\pm 0.46}$	$63.39_{\pm 1.75}$	$54.07_{\pm 1.88}$	$86.79_{\pm 0.33}$	$77.21_{\pm 0.37}$
BWGNN	$85.33_{\pm 0.17}$	$73.37_{\pm 0.65}$	$96.76_{\pm 0.22}$	$91.18_{\pm 0.54}$	$63.09_{\pm 1.38}$	$59.36_{\pm 1.71}$	$97.62_{\pm 0.19}$	$92.15_{\pm 0.30}$
SplitGNN	$79.37_{\pm 0.77}$	$66.00_{\pm 0.79}$	$82.80_{\pm 0.45}$	$61.44_{\pm 0.59}$	$58.01_{\pm 1.96}$	$52.15_{\pm 2.13}$	$88.39_{\pm 0.43}$	$78.91_{\pm 0.82}$
GHRN	$85.99_{\pm 0.52}$	$65.80_{\pm 0.78}$	$97.41_{\pm 0.43}$	$91.06_{\pm 0.55}$	$58.42_{\pm 1.89}$	$51.85_{\pm 2.07}$	$97.29_{\pm 0.34}$	$90.68_{\pm 0.57}$
SEC-GFD	$83.72_{\pm 0.23}$	$71.40_{\pm 0.68}$	$96.91_{\pm 0.28}$	$91.12_{\pm 0.63}$	$63.88_{\pm 1.36}$	$59.08_{\pm 1.76}$	$96.82_{\pm 0.20}$	$92.11_{\pm 0.32}$
AHFAN	$87.51_{\pm 0.67}$	$72.11_{\pm 0.87}$	$97.69_{\pm 0.32}$	$90.59_{\pm 0.50}$	$61.88_{\pm 1.84}$	$59.08_{\pm 1.23}$	$96.13_{\pm 0.56}$	$91.51_{\pm 0.44}$
FedPAM	$74.52_{\pm 0.61}$	$61.20_{\pm 0.85}$	$88.15_{\pm 0.45}$	$80.20_{\pm 0.72}$	$55.50_{\pm 1.15}$	$48.15_{\pm 1.45}$	$84.95_{\pm 0.55}$	$76.55_{\pm 0.88}$
FedGH	$75.14_{\pm 0.52}$	$61.85_{\pm 0.76}$	$88.62_{\pm 0.38}$	$80.85_{\pm 0.65}$	$56.08_{\pm 1.05}$	$48.65_{\pm 1.32}$	$85.60_{\pm 0.48}$	$77.15_{\pm 0.75}$
Fed-Kalter	$77.85_{\pm 0.44}$	$64.12_{\pm 0.63}$	$91.25_{\pm 0.31}$	$83.56_{\pm 0.52}$	$58.45_{\pm 0.88}$	$51.24_{\pm 1.10}$	$88.42_{\pm 0.35}$	$79.68_{\pm 0.62}$
FedSACA	$89.70_{\pm 0.14}$	$74.23_{\pm 0.62}$	$98.69_{\pm 0.09}$	$91.21_{\pm 0.41}$	$73.88_{\pm 1.35}$	$65.36_{\pm 1.61}$	$98.03_{\pm 0.14}$	$92.07_{\pm 0.31}$
FedSACA+	$\mathbf{91.84}_{\pm \mathbf{0.12}}$	$\mathbf{76.38}_{\pm \mathbf{0.56}}$	$\mathbf{99.30}_{\pm \mathbf{0.02}}$	$\mathbf{92.19}_{\pm \mathbf{0.21}}$	$\mathbf{76.16}_{\pm \mathbf{1.03}}$	$\mathbf{67.97}_{\pm \mathbf{1.40}}$	$\mathbf{99.18}_{\pm \mathbf{0.08}}$	$\mathbf{93.60}_{\pm \mathbf{0.15}}$

5.1 EXPERIMENTAL SETUP

Datasets. We evaluate our approach through four datasets related to GFD: YelpChi (Rayana & Akoglu, 2015), Amazon (McAuley & Leskovec, 2013), FDCompCN (Wu et al., 2023), Elliptic (Weber et al., 2019). YelpChi and Amazon are review fraud datasets characterized by severe fraud spectral heterogeneity. FDCompCN is a financial statement fraud graph, and Elliptic is a Bitcoin transaction network for illicit detection. The graph structures of these two datasets are sparse.

Baselines. To ensure a comprehensive evaluation, we adapt all originally centralized methods to the federated setting by deploying them under FedAvg (McMahan et al., 2017a). We compare our model with state-of-the-art approaches, which are categorized into four groups. The first group contains conventional GNNs, including GCN (Kipf & Welling, 2017), GraphSAGE (Hamilton et al., 2017), and GAT (Veličković et al., 2017). The second group encompasses spatial-domain fraud detection methods, such as CARE-GNN (Dou et al., 2020), GAGA (Wang et al., 2023), PMP (Zhuo et al., 2024) and GPA (Qiu et al., 2025). The third group focuses on spectral-domain baselines, including BWGNN (Tang et al., 2022), SplitGNN (Wu et al., 2023), GHRN (Gao et al., 2023), SEC-GFD (Xu et al., 2024) and AHFAN Wang et al. (2025). The last group comprises FGL methods tackling misalignment and heterogenity, such as FedPAM (Fu et al., 2026), FedGH (Zhuang et al., 2026), and Fed-Kalter (Tang & Chen).

Implementation Details. We adopt the Dirichlet distribution (Hsu et al., 2019) to partition the graph across $N = 5$ clients, which is a typical strategy to construct non-IID data in FGL (He et al., 2021a). For each dataset, we employ a 40%/20%/40% split for training, validation, and testing. The local GNN models are trained using the Adam optimizer (Kinga et al., 2015). The global training process spans 1000 communication rounds, with each client performing 3 local epochs per round. We truncate the Chebyshev polynomial expansion at a maximum order of $K - 1 = 5$. To prevent the base weights from dominating the confidence-aware optimization, we assign a larger learning rate 0.02 to the mask parameters and 0.01 to the general parameters. The ℓ_1-norm sparsity penalty λ and the proximal alignment term μ are set to 0.005 and 0.01 respectively. Furthermore, we set the number of nearest neighbor nodes to $k_{near} = 5$ and the constant bias to $c = 3$. The results represent the average of 5 runs with different random seeds.

5.2 EXPERIMENTAL RESULTS

Performance Comparison. We present the overall federated graph fraud detection results across four datasets, with the average test AUC and F1-Macro scores detailed in Table 1. The results demonstrate that FedSACA+ outperforms all baseline methods. Conventional GNNs like GCN and spatial-based GFD methods like CARE-GNN exhibit significantly degraded performance, as their spatial aggregation schemes are highiy vulnerable to topological fragmentation and fail to handle

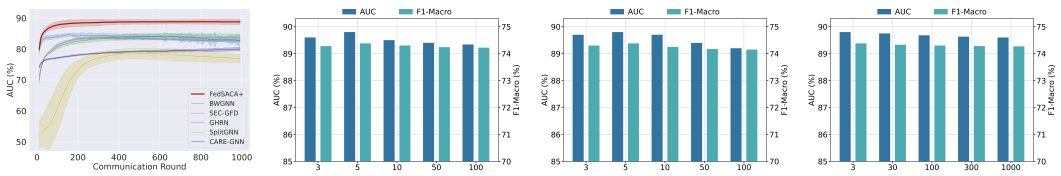

(a) Convergence Curves **(b)** Order $K-1$ **(c)** Number k_{near} **(d)** Factor c

Figure 5: Convergence and Hyperparameter Analysis of FedSACA+. (a) Test AUC convergence curves of FedSACA+ and five other methods over communication rounds on YelpChi. (b)-(d) Evaluation of AUC and F1-Macro with respect to the Chebyshev polynomial order $K-1$, the number of nearest neighbor nodes k_{near}, and the scaling factor c, respectively.

fraud spectral heterogeneity. While centralized spectral-based methods like BWGNN and SEC-GFD attempt to leverage frequency signals, their static filter designs fail to reconcile inter-client spectral conflicts in the federated setting, leading to suboptimal global models. Furthermore, existing FGL methods like FedSage+ suffer from the propagation of deceptive noise when handling fraudulent node isolation. Notably, among the four datasets listed, FedSACA+ achieved the most significant improvements in AUC and F1-Macro scores over FedSACA on the FDCompCN dataset, where the issue of the structural isolation of fraudulent nodes is the severest. This result underscores the effectiveness of FedSACA+ in addressing fraudulent node isolation. Furthermore, given that FedSACA+ consistently outperforms FedSACA across the other three datasets as well, it further demonstrates the robustness of the structure augmentation under the control of frequency confidence.

Convergence Analysis. Figure 5a illustrates the test AUC convergence curves over communication rounds on YelpChi. The results demonstrate that existing approaches exhibit significant performance fluctuations or overfitting under fraud spectral heterogeneity and fraudulent node isolation. In contrast, our proposed FedSACA+ maintains both stable convergence and superior performance, showcasing its robustness to federated fraud detection.

5.3 Hyper-parameter Study

We conduct a comprehensive investigation of three key hyperparameters on YelpChi: the Chebyshev polynomial order $K-1$, the number of nearest neighbor nodes k_{near}, and the scaling factor c (Figure 5b to 5d). For the order $K-1$, performance initially increases as higher-order filters capture more complex fraud signals, peaking at $K-1=5$. For the structure augmentation, FedSACA+ maintains superior performance across a wide range of k_{near} and c values. This performance advantage across various parameter configurations demonstrates the robustness of our proposed techniques.

6 Conclusion

This paper presents a comprehensive study addressing two critical yet overlooked challenges in federated graph fraud detection: fraud spectral heterogeneity and fraudulent node isolation. Existing centralized GFD methods and standard FGL frameworks fail to address these issues, often resulting in spectral parameter conflicts or the propagation of deceptive noise. In response, we propose **FedSACA**, a novel framework that proposes spectral channel attention and confidence-aware aggregation. This design allows for dynamic local adaptation and reconcilable aggregation, thus overcoming fraud spectral heterogeneity. To further safely reconnect isolated fraudulent nodes, we extend it to **FedSACA+**, introducing a feature-induced structure augmentation rigorously subject to the frequency confidence evaluation. Extensive experiments across multiple datasets demonstrate the superior robustness and effectiveness of our proposed techniques.

AI use statement

In this work, we have not used generative AI tools for any tasks with required disclosure, such as generating synthetic datasets, developing theoretical models, formulating or proving mathematical claims, proposing hypotheses, designing methodology, implementing methods, or interpreting results. Additionally, we used generative AI tools for tasks with recommended disclosure, specifically

to summarize existing literature and polish the manuscript to improve overall readability. We have reviewed all AI-assisted work. Specifically, the authors manually verified all AI-polished text to ensure scientific accuracy and alignment with our original intended meaning. We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

REFERENCES

Fahao Chen, Peng Li, Toshiaki Miyazaki, and Celimuge Wu. Fedgraph: Federated graph learning with intelligent sampling. *IEEE PDS*, 33(8):1775–1786, 2021.

Yu Chen, Lingfei Wu, and Mohammed Zaki. Iterative deep graph learning for graph neural networks: Better and robust node embeddings. *Advances in neural information processing systems*, 33:19314–19326, 2020.

Dawei Cheng, Yao Zou, Sheng Xiang, and Changjun Jiang. Graph neural networks for financial fraud detection: a review. *Frontiers of Computer Science*, 19(9):199609, 2025.

Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE transactions on information theory*, 13(1):21–27, 1967.

Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. *Advances in neural information processing systems*, 29, 2016.

Leyan Deng, Chenwang Wu, Defu Lian, Yongji Wu, and Enhong Chen. Markov-driven graph convolutional networks for social spammer detection. *IEEE Transactions on Knowledge and Data Engineering*, 35(12):12310–12322, 2022.

Yingtong Dou, Zhiwei Liu, Li Sun, Yutong Deng, Hao Peng, and Philip S Yu. Enhancing graph neural network-based fraud detectors against camouflaged fraudsters. In *Proceedings of the 29th ACM international conference on information & knowledge management*, pp. 315–324, 2020.

Lele Fu, Tianchi Liao, Sheng Huang, Bowen Deng, Shirui Pan, and Chuan Chen. Unsupervised federated graph learning. *Advances in Neural Information Processing Systems*, 38:97431–97457, 2026.

Xingbo Fu, Binchi Zhang, Yushun Dong, Chen Chen, and Jundong Li. Federated graph machine learning: A survey of concepts, techniques, and applications. *arXiv preprint arXiv:2207.11812*, 2022.

Yuan Gao, Xiang Wang, Xiangnan He, Zhenguang Liu, Huamin Feng, and Yongdong Zhang. Addressing heterophily in graph anomaly detection: A perspective of graph spectrum. In *Proceedings of the ACM web conference 2023*, pp. 1528–1538, 2023.

Ronghui Guo, Xiaowang Zhang, Zhizhi Yu, Minghui Zou, Sai Zhang, and Zhiyong Feng. Naastgnn: Neighborhood adaptive aggregation and spectral tuning for graph anomaly detection. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25, International Joint Conferences on Artificial Intelligence Organization*, pp. 2847–2855, 2025.

Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.

Chaoyang He, Keshav Balasubramanian, Emir Ceyani, Carl Yang, Han Xie, Lichao Sun, Lifang He, Liangwei Yang, Philip S Yu, Yu Rong, et al. Fedgraphnn: A federated learning system and benchmark for graph neural networks. In *ICLR*, 2021a.

Mingguo He, Zhewei Wei, Hongteng Xu, et al. Bernnet: Learning arbitrary graph spectral filters via bernstein approximation. *Advances in Neural Information Processing Systems*, 34:14239–14251, 2021b.

Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*, 2019.

Mengda Huang, Yang Liu, Xiang Ao, Kuan Li, Jianfeng Chi, Jinghua Feng, Hao Yang, and Qing He. Auc-oriented graph neural network for fraud detection. In *Proceedings of the ACM web conference 2022*, pp. 1311–1321, 2022.

Rongrong Jing, Hu Tian, Gang Zhou, Xingwei Zhang, Xiaolong Zheng, and Daniel Dajun Zeng. A gnn-based few-shot learning model on the credit card fraud detection. In *2021 IEEE 1st International Conference on Digital Twins and Parallel Intelligence (DTPI)*, pp. 320–323. IEEE, 2021.

Peter Kairouz and H Brendan McMahan. Advances and open problems in federated learning. *Foundations and trends in machine learning*, 14(1-2):1–210, 2021.

D Kinga, Jimmy Ba Adam, et al. A method for stochastic optimization. In *ICLR*, 2015.

Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.

Li Li, Yuxi Fan, Mike Tse, and Kuo-Yi Lin. A review of applications in federated learning. *Computers & Industrial Engineering*, 149:106854, 2020a.

Pengbo Li, Hang Yu, Xiangfeng Luo, and Jia Wu. Lgm-gnn: A local and global aware memory-based graph neural network for fraud detection. *IEEE Transactions on Big Data*, 9(4):1116–1127, 2023.

Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2:429–450, 2020b.

Wangli Lin, Li Sun, Qiwei Zhong, Can Liu, Jinghua Feng, Xiang Ao, and Hao Yang. Online credit payment fraud detection via structure-aware hierarchical recurrent neural network. In *IJCAI*, pp. 3670–3676, 2021.

Jiaqi Liu, Zihan Tan, Guancheng Wan, Wenke Huang, He Li, and Mang Ye. Fedsdr: Federated graph learning with structural noise detection and reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3379–3389, 2026.

Jingxin Liu, Renda Han, Wenxuan Tu, Haotian Wang, Junlong Wu, and Jieren Cheng. Federated node-level clustering network with cross-subgraph link mending. In *Forty-second International Conference on Machine Learning*, 2025.

Ziqi Liu, Chaochao Chen, Xinxing Yang, Jun Zhou, Xiaolong Li, and Le Song. Heterogeneous graph neural networks for malicious account detection. In *Proceedings of the 27th ACM international conference on information and knowledge management*, pp. 2077–2085, 2018.

Ziqi Liu, Chaochao Chen, Longfei Li, Jun Zhou, Xiaolong Li, Le Song, and Yuan Qi. Geniepath: Graph neural networks with adaptive receptive paths. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pp. 4424–4431, 2019.

Julian John McAuley and Jure Leskovec. From amateurs to connoisseurs: modeling the evolution of user expertise through online reviews. In *Proceedings of the 22nd international conference on World Wide Web*, pp. 897–908, 2013.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, pp. 1273–1282, 2017a.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pp. 1273–1282. PMLR, 2017b.

Junjun Pan, Yixin Liu, Xin Zheng, Yizhen Zheng, Alan Wee-Chung Liew, Fuyi Li, and Shirui Pan. A label-free heterophily-guided approach for unsupervised graph fraud detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 12443–12451, 2025.

Hao Peng, Ruitong Zhang, Yingtong Dou, Renyu Yang, Jingyi Zhang, and Philip S Yu. Reinforced neighborhood selection guided multi-relational graph neural networks. *ACM Transactions on Information Systems (TOIS)*, 40(4):1–46, 2021.

Tian Qiu, Wenda Li, Zunlei Feng, Jie Lei, Tao Wang, Yi Gao, Mingli Song, and Yang Gao. Association-focused path aggregation for graph fraud detection. *Advances in Neural Information Processing Systems*, 38:6118–6141, 2025.

Shebuti Rayana and Leman Akoglu. Collective opinion spam detection: Bridging review networks and metadata. In *Proceedings of the 21th acm sigkdd international conference on knowledge discovery and data mining*, pp. 985–994, 2015.

Yue Tan, Yixin Liu, Guodong Long, Jing Jiang, Qinghua Lu, and Chengqi Zhang. Federated learning on non-iid graphs via structural knowledge sharing. In *AAAI*, pp. 9953–9961, 2023.

Bisheng Tang and Xiaojun Chen. Federated graph learning via structure-aware fusion using a kalman framework with learnable dynamics. In *Forty-third International Conference on Machine Learning*.

Jianheng Tang, Jiajin Li, Ziqi Gao, and Jia Li. Rethinking graph neural networks for anomaly detection. In *ICML*, 2022.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.

Guancheng Wan, Wenke Huang, and Mang Ye. Federated graph learning under domain shift with generalizable prototypes. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pp. 15429–15437, 2024.

Xiang Wang, Hao Dou, Dibo Dong, and Zhenyu Meng. Graph anomaly detection based on hybrid node representation learning. *Neural Networks*, 185:107169, 2025.

Xiyuan Wang and Muhan Zhang. How powerful are spectral graph neural networks. In *International conference on machine learning*, pp. 23341–23362. PMLR, 2022.

Yuchen Wang, Jinghui Zhang, Zhengjie Huang, Weibin Li, Shikun Feng, Ziheng Ma, Yu Sun, Dianhai Yu, Fang Dong, Jiahui Jin, et al. Label information enhanced fraud detection against low homophily in graphs. In *Proceedings of the ACM Web Conference 2023*, pp. 406–416, 2023.

Mark Weber, Giacomo Domeniconi, Jie Chen, Daniel Karl I Weidele, Claudio Bellei, Tom Robinson, and Charles E Leiserson. Anti-money laundering in bitcoin: Experimenting with graph convolutional networks for financial forensics. *arXiv preprint arXiv:1908.02591*, 2019.

Bin Wu, Xinyu Yao, Boyan Zhang, Kuo-Ming Chao, and Yinsheng Li. Splitgnn: Spectral graph neural network for fraud detection against heterophily. In *Proceedings of the 32nd ACM international conference on information and knowledge management*, pp. 2737–2746, 2023.

Han Xie, Jing Ma, Li Xiong, and Carl Yang. Federated graph classification over non-iid graphs. In *NeruIPS*, 2021.

Fan Xu, Nan Wang, Hao Wu, Xuezhi Wen, Xibin Zhao, and Hai Wan. Revisiting graph-based fraud detection in sight of heterophily and spectrum. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pp. 9214–9222, 2024.

Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.

Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19, 2019.

Yuhang Yao, Weizhao Jin, Srivatsan Ravi, and Carlee Joe-Wong. Fedgcn: Convergence-communication tradeoffs in federated training of graph convolutional networks. *Advances in neural information processing systems*, 36:79748–79760, 2023.

Ke Zhang, Carl Yang, Xiaoxiao Li, Lichao Sun, and Siu Ming Yiu. Subgraph federated learning with missing neighbor generation. In *NeurIPS*, volume 34, pp. 6671–6682, 2021a.

Ke Zhang, Carl Yang, Xiaoxiao Li, Lichao Sun, and Siu Ming Yiu. Subgraph federated learning with missing neighbor generation. *Advances in Neural Information Processing Systems*, 34:6671–6682, 2021b.

Qirong Zhang, Jin Li, Qingqing Ye, Yuxi Lin, Xinlong Chen, and Yang-Geng Fu. Dwssa: Alleviating over-smoothness for deep graph neural networks. *Neural Networks*, 174:106228, 2024.

Xiaokang Zhou, Wei Liang, I Kevin, Kai Wang, Katsutoshi Yada, Laurence T Yang, Jianhua Ma, and Qun Jin. Decentralized federated graph learning with lightweight zero trust architecture for next-generation networking security. *IEEE Journal on Selected Areas in Communications*, 2025.

Yanqi Zhou, Tao Lei, Hanxiao Liu, Nan Du, Yanping Huang, Vincent Zhao, Andrew M Dai, Quoc V Le, James Laudon, et al. Mixture-of-experts with expert choice routing. *Advances in Neural Information Processing Systems*, 35:7103–7114, 2022.

Shuman Zhuang, Zhihao Wu, Wei Huang, Luojun Lin, Jia-Li Yin, Lele Fu, and Hong-Ning Dai. Prior refinement is better: Diffusion-driven graph harmonization for federated graph learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 29241–29250, 2026.

Wei Zhuo, Zemin Liu, Bryan Hooi, Bingsheng He, Guang Tan, Rizal Fathony, and Jia Chen. Partitioning message passing for graph fraud detection. *arXiv preprint arXiv:2412.00020*, 2024.

A Related Work

A.1 Graph-based Fraud Detection

Recent research in graph-based fraud detection has largely focused on Graph Neural Networks (GNNs) (Cheng et al., 2025; Li et al., 2023; Huang et al., 2022), which are typically divided into spatial and spectral methods to address the heterophilic nature of fraud networks. Spatial methods focus on relation-aware aggregation and neighbor sampling. For instance, CARE-GNN (Dou et al., 2020) employs reinforcement learning to select informative neighbors under different relations. Extending this single-layer selection, RioGNN (Peng et al., 2021) leverages a recursive framework to dynamically refine neighbor sampling, while GAGA (Wang et al., 2023) introduces a group aggregation mechanism with learnable encodings to address low homophily. Alternatively, spectral methods are dedicated to filter design to capture distinct fraud signals in the frequency domain. SplitGNN (Wu et al., 2023) splits graphs via an edge classifier and employs band-pass filters against heterophily. SEC-GFD (Xu et al., 2024) utilizes hybrid filters to integrate multi-frequency information, and PMP (Zhuo et al., 2024) further bridges spatial and spectral domains by partitioning message passing to learn node-adaptive spectral filters. However, these centralized methods fail in the federated setting. Cross-client edge isolation hinders their spatial sampling, while their spectral filters cannot reconcile diverging fraud frequency profiles across clients. This motivates us to design a confidence-aware spectral framework, bridging the literature gap where spectral analysis remains underexplored in federated fraud detection.

A.2 Federated Graph Learning

Federated Graph Learning (FGL) aims to collaboratively train GNNs across distributed clients without sharing raw data (Zhou et al., 2025; Wan et al., 2024; Chen et al., 2021; Fu et al., 2022). Early FGL frameworks typically integrate standard federated optimization algorithms (McMahan et al., 2017b; Li et al., 2020b) with representative GNNs (Kipf & Welling, 2017; Veličković et al., 2017; Xu et al., 2019; Hamilton et al., 2017) to aggregate model parameters. To handle the inherent heterogeneity across non-IID clients, subsequent research has explored clustered and decoupled learning techniques. For instance, GCFL (Xie et al., 2021) dynamically groups clients based on their gradient sequences to maximize collaboration among graphs with similar distributions. In contrast, FedStar (Tan et al., 2023) explicitly decouples feature and structure encoders, sharing only the structural

knowledge across clients to extract domain-invariant patterns. In addition, to address the structural deficiency caused by data isolation, recent FGL studies focus on topology recovery and refinement. FedSage+ (Zhang et al., 2021a) trains local generative models to reconstruct missing cross-subgraph neighbors. FedNCN (Liu et al., 2025) further restores such links on the server side using clustering prior knowledge, combining graph similarity estimation with N-Cut-based pruning.

A.3 Adaptive Spectral Filters and Structure Augmentation

To capture complex graph signals, polynomial-based spectral GNNs, such as ChebNet (Defferrard et al., 2016), BernNet (He et al., 2021b), and JacobiConv (Wang & Zhang, 2022), serve as adaptive spectral filters capable of approximating arbitrary frequency responses. This adaptivity perfectly suits divergent fraud frequency profiles across clients in the federated setting. However, leveraging traditional aggregation strategies (McMahan et al., 2017b; Li et al., 2020b) can average these conflicting spectral objectives, thereby destroying their local adaptivity and leading to poor generalization. Meanwhile, to overcome graph sparsity, structure augmentation is widely adopted. For instance, IDGL (Chen et al., 2020) iteratively refines graph topologies based on feature and embedding similarities, while DWSSA (Zhang et al., 2024) adds connections via a *K* Nearest Neighbor (kNN) algorithm. However, such methods can unintentionally link camouflaged fraudulent nodes to benign ones, which is particularly concerning in the federated setting. To take full advantage of adaptive spectral filters and safely repair structural isolation, we propose FedSACA, along with its extension FedSACA+. These two techniques perfectly preserve client-specific spectral adaptivity while rigorously controlling the feature-induced structure augmentation, enabling robust federated fraud detection.

B Calculation of Fraud Signal Energy

To explicitly quantify and visualize the fraud spectral heterogeneity presented in Figure 1a, we analyze the distribution of fraud signals in the spectral domain using Graph Signal Processing (GSP) techniques. The calculation fundamentally relies on the Graph Fourier Transform (GFT) and consists of the following key steps:

Graph Laplacian and Eigendecomposition. Following the definition in the Preliminary section, for a specific relation r, the normalized Laplacian matrix is defined as $\tilde{\mathbf{L}}^{(r)} = \mathbf{I} - \hat{\mathbf{A}}^{(r)}$. To obtain the graph frequencies, we perform an exact eigendecomposition on $\tilde{\mathbf{L}}^{(r)}$:

$$\tilde{\mathbf{L}}^{(r)} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^{\top},$$

where $\mathbf{\Lambda} = \text{diag}(\lambda_0, \ldots, \lambda_{|\mathcal{V}|-1})$ is the diagonal matrix of eigenvalues representing graph frequencies bounded in $[0, 2]$, and $\mathbf{U} = [\mathbf{u}_0, \ldots, \mathbf{u}_{|\mathcal{V}|-1}]$ is the orthogonal matrix of eigenvectors serving as the Graph Fourier basis.

Graph Fourier Transform (GFT). Let $\mathbf{y} \in \{0, 1\}^{|\mathcal{V}|}$ denote the ground-truth fraud label set, where the v-th element $y_v = 1$ indicates a fraudulent node and $y_v = 0$ indicates a benign node. To observe how fraud patterns are distributed across different frequencies, we project this spatial label signal onto the spectral domain via GFT:

$$\hat{\mathbf{y}} = \mathbf{U}^{\top}\mathbf{y}.$$

Here, $\hat{\mathbf{y}} \in \mathbb{R}^{|\mathcal{V}|}$ is defined as the spectral representation of the fraud labels. Specifically, each GFT coefficient for a specific frequency λ_k is computed as $\hat{y}(\lambda_k) = \mathbf{u}_k^{\top}\mathbf{y}$, which captures the amplitude of the fraud signal at that frequency.

Spectral Energy Calculation. The spectral energy of the fraud signal at a given frequency λ_k is formulated as the squared magnitude of its corresponding GFT coefficient:

$$E(\lambda_k) = |\hat{y}(\lambda_k)|^2.$$

This metric reflects the precise contribution of the frequency λ_k to the overall fraud behavior.

Binning and Normalization. To construct the continuous-like spectral energy distributions (Figure 1a), we uniformly partition the frequency spectrum $[0, 2]$ into M equal-width bins. Let $\mathcal{B}_m$

denote the frequency interval of the m-th bin, where $m \in \{1, 2, \ldots, M\}$. The energy of the m-th bin is the sum of energies of all frequencies falling into this interval:

$$E_{bin}(m) = \sum_{\lambda_k \in \mathcal{B}_m} E(\lambda_k).$$

Finally, to facilitate a fair comparison of distribution shapes across clients with varying graph sizes and fraud node counts, the binned energy is min-max normalized to $[0, 1]$:

$$\bar{E}_{bin}(m) = \frac{E_{bin}(m) - \min_{j\in\{1,\ldots,M\}} E_{bin}(j)}{\max_{j\in\{1,\ldots,M\}} E_{bin}(j) - \min_{j\in\{1,\ldots,M\}} E_{bin}(j)}.$$

C Additional Experiments

Methods	YelpChi			Amazon		
	AP	G-Mean	Recall	AP	G-Mean	Recall
GCN	$45.12_{\pm 1.25}$	$48.33_{\pm 1.42}$	$49.15_{\pm 1.38}$	$58.12_{\pm 1.15}$	$65.34_{\pm 1.22}$	$66.10_{\pm 1.18}$
GraphSAGE	$55.45_{\pm 1.05}$	$59.21_{\pm 1.12}$	$58.74_{\pm 1.08}$	$70.15_{\pm 0.95}$	$75.88_{\pm 1.05}$	$76.24_{\pm 1.02}$
GAT	$48.34_{\pm 0.95}$	$51.05_{\pm 1.15}$	$50.88_{\pm 1.10}$	$59.45_{\pm 0.85}$	$68.12_{\pm 0.96}$	$68.55_{\pm 0.92}$
CARE-GNN	$59.34_{\pm 0.85}$	$63.45_{\pm 0.92}$	$62.98_{\pm 0.88}$	$85.34_{\pm 0.65}$	$88.12_{\pm 0.75}$	$88.65_{\pm 0.72}$
GAGA	$58.12_{\pm 0.92}$	$62.88_{\pm 1.05}$	$62.15_{\pm 0.98}$	$86.15_{\pm 0.55}$	$89.24_{\pm 0.68}$	$89.55_{\pm 0.65}$
PMP		OOT			OOT	
GPA	$58.95_{\pm 0.88}$	$61.22_{\pm 0.95}$	$61.50_{\pm 0.90}$	$86.05_{\pm 0.60}$	$89.55_{\pm 0.72}$	$89.88_{\pm 0.68}$
BWGNN	$63.45_{\pm 0.75}$	$68.75_{\pm 0.82}$	$69.12_{\pm 0.78}$	$88.75_{\pm 0.52}$	$92.15_{\pm 0.65}$	$92.45_{\pm 0.60}$
SplitGNN	$59.88_{\pm 0.95}$	$62.14_{\pm 1.02}$	$62.55_{\pm 0.96}$	$78.44_{\pm 0.75}$	$80.12_{\pm 0.85}$	$80.55_{\pm 0.82}$
GHRN	$64.12_{\pm 0.68}$	$70.33_{\pm 0.75}$	$70.85_{\pm 0.72}$	$89.12_{\pm 0.48}$	$93.45_{\pm 0.55}$	$93.88_{\pm 0.52}$
SEC-GFD	$62.55_{\pm 0.82}$	$68.12_{\pm 0.88}$	$68.44_{\pm 0.85}$	$88.90_{\pm 0.55}$	$92.88_{\pm 0.62}$	$93.15_{\pm 0.58}$
AHFAN	$65.34_{\pm 0.72}$	$71.45_{\pm 0.81}$	$71.98_{\pm 0.77}$	$89.45_{\pm 0.45}$	$94.10_{\pm 0.52}$	$94.50_{\pm 0.48}$
FedPAM	$55.67_{\pm 0.95}$	$58.34_{\pm 1.05}$	$58.90_{\pm 1.02}$	$75.34_{\pm 0.85}$	$82.15_{\pm 0.95}$	$82.60_{\pm 0.92}$
FedGH	$56.12_{\pm 0.92}$	$59.10_{\pm 0.98}$	$59.45_{\pm 0.95}$	$76.12_{\pm 0.82}$	$83.10_{\pm 0.92}$	$83.55_{\pm 0.88}$
Fed-Kalter	$58.33_{\pm 0.85}$	$61.25_{\pm 0.92}$	$61.70_{\pm 0.88}$	$79.55_{\pm 0.75}$	$86.45_{\pm 0.85}$	$86.90_{\pm 0.82}$
FedSACA	$66.83_{\pm 0.55}$	$83.02_{\pm 0.61}$	$83.15_{\pm 0.58}$	$92.74_{\pm 0.42}$	$95.82_{\pm 0.51}$	$95.85_{\pm 0.48}$
FedSACA+	$\mathbf{69.24}_{\pm \mathbf{0.12}}$	$\mathbf{85.28}_{\pm \mathbf{0.15}}$	$\mathbf{85.35}_{\pm \mathbf{0.14}}$	$\mathbf{95.12}_{\pm \mathbf{0.11}}$	$\mathbf{98.15}_{\pm \mathbf{0.13}}$	$\mathbf{98.21}_{\pm \mathbf{0.10}}$

Table 2: Additional performance comparison on YelpChi and Amazon datasets in terms of Average Precision (AP), Geometric Mean (G-Mean), and Recall.

Methods	FDCompCN			Elliptic		
	AP	G-Mean	Recall	AP	G-Mean	Recall
GCN	$15.34_{\pm 1.45}$	$45.12_{\pm 1.55}$	$46.88_{\pm 1.52}$	$60.12_{\pm 1.25}$	$65.34_{\pm 1.35}$	$66.88_{\pm 1.32}$
GraphSAGE	$18.45_{\pm 1.32}$	$48.34_{\pm 1.42}$	$49.15_{\pm 1.38}$	$65.45_{\pm 1.15}$	$69.12_{\pm 1.22}$	$70.15_{\pm 1.18}$
GAT	$16.12_{\pm 1.25}$	$46.55_{\pm 1.35}$	$47.10_{\pm 1.30}$	$61.34_{\pm 1.05}$	$66.15_{\pm 1.15}$	$67.50_{\pm 1.10}$
CARE-GNN	$22.45_{\pm 1.15}$	$54.12_{\pm 1.25}$	$55.60_{\pm 1.20}$	$72.15_{\pm 0.85}$	$76.88_{\pm 0.95}$	$77.40_{\pm 0.92}$
GAGA	$21.88_{\pm 1.05}$	$53.45_{\pm 1.15}$	$54.80_{\pm 1.10}$	$78.45_{\pm 0.75}$	$82.12_{\pm 0.85}$	$83.05_{\pm 0.82}$
PMP		OOT			OOT	
GPA	$22.15_{\pm 0.95}$	$54.05_{\pm 1.05}$	$55.20_{\pm 1.02}$	$75.34_{\pm 0.82}$	$79.45_{\pm 0.90}$	$80.12_{\pm 0.88}$
BWGNN	$25.66_{\pm 0.85}$	$58.15_{\pm 0.95}$	$59.40_{\pm 0.92}$	$88.15_{\pm 0.65}$	$91.34_{\pm 0.72}$	$91.88_{\pm 0.68}$
SplitGNN	$20.12_{\pm 1.05}$	$51.34_{\pm 1.15}$	$52.88_{\pm 1.12}$	$74.12_{\pm 0.95}$	$78.55_{\pm 1.05}$	$79.10_{\pm 1.02}$
GHRN	$26.88_{\pm 0.75}$	$59.45_{\pm 0.85}$	$60.15_{\pm 0.82}$	$87.55_{\pm 0.55}$	$90.12_{\pm 0.65}$	$90.65_{\pm 0.60}$
SEC-GFD	$25.10_{\pm 0.88}$	$58.70_{\pm 0.98}$	$59.88_{\pm 0.95}$	$88.05_{\pm 0.62}$	$91.15_{\pm 0.70}$	$91.50_{\pm 0.66}$
AHFAN	$28.15_{\pm 0.72}$	$62.34_{\pm 0.81}$	$63.50_{\pm 0.78}$	$89.12_{\pm 0.52}$	$92.45_{\pm 0.62}$	$92.88_{\pm 0.58}$
FedPAM	$18.55_{\pm 1.05}$	$48.90_{\pm 1.15}$	$49.50_{\pm 1.10}$	$72.34_{\pm 0.95}$	$76.15_{\pm 1.05}$	$77.10_{\pm 1.02}$
FedGH	$19.10_{\pm 0.95}$	$49.15_{\pm 1.05}$	$50.12_{\pm 1.02}$	$73.15_{\pm 0.88}$	$77.10_{\pm 0.98}$	$78.05_{\pm 0.95}$
Fed-Kalter	$20.45_{\pm 0.85}$	$51.88_{\pm 0.95}$	$52.45_{\pm 0.90}$	$75.45_{\pm 0.75}$	$80.12_{\pm 0.85}$	$81.05_{\pm 0.82}$
FedSACA	$34.05_{\pm 0.65}$	$70.12_{\pm 0.72}$	$71.43_{\pm 0.68}$	$94.02_{\pm 0.35}$	$95.15_{\pm 0.41}$	$95.21_{\pm 0.38}$
FedSACA+	$\mathbf{36.56}_{\pm \mathbf{0.14}}$	$\mathbf{72.82}_{\pm \mathbf{0.18}}$	$\mathbf{73.85}_{\pm \mathbf{0.15}}$	$\mathbf{96.45}_{\pm \mathbf{0.08}}$	$\mathbf{97.34}_{\pm \mathbf{0.10}}$	$\mathbf{97.42}_{\pm \mathbf{0.12}}$

Table 3: Additional performance comparison on FDCompCN and Elliptic datasets in terms of Average Precision (AP), Geometric Mean (G-Mean), and Recall.

D Complexity Analysis

We analyze the space and time complexity of FedSACA+ and compare it with GPA, SEC-GFD, and BWGNN. Let N be the number of clients. Client i holds n_i nodes, d-dimensional features, and a relation set $\mathcal{R}_i$. Let $|E_i| = \sum_{r\in\mathcal{R}_i} |E_i^{(r)}|$ be the number of intrinsic edges in client i. FedSACA uses a Chebyshev order K. FedSACA+ adds one kNN relation r_{knn} with k_{near} nearest neighbors, yielding $\mathcal{R}^* = \mathcal{R}\cup\{r_{\text{knn}}\}$. We use h to denote the hidden dimension of the local classifier. For centralized baselines, let $n = \sum_i n_i$ and M denote the total number of nodes and edges, respectively. For GPA, I is the number of sampled paths, K_p the number of selected paths, J the path length, C_p the path embedding dimension, and L the number of attention layers. For SEC-GFD and BWGNN, C denotes their spectral order; in SEC-GFD, L denotes the number of SGC layers and k denotes the kNN size used by the local environmental constraint.

FedSACA+ performs Chebyshev filtering on each relation. The sparse multiplication $T_k(\hat{\mathbf{A}}_i^{(r)})\mathbf{X}_i$ costs $O(|E_i^{(r)}|d)$, and K orders give $O(Kd\sum_{r\in\mathcal{R}^*}|E_i^{(r)}|) = O(Kd(|E_i| + n_i k_{\text{near}}))$. The kNN graph construction is a one-time preprocessing cost of $O(n_i^2 d)$ in the worst case, but it only adds $O(n_i k_{\text{near}})$ edges during training. The server aggregates only K-dimensional coefficients and masks, so its time and memory are $O(N|\mathcal{R}^*|K)$, independent of the graph size. This is much cheaper than centralized path or spectral methods, which must process the global graph.

GPA samples I paths of length J per node and selects K_p paths. Its dominant cost is path interaction: self-attention over K_p paths costs $O(K_p^2 C_p + K_p C_p^2)$ per layer, leading to $O(nL(K_p^2 C_p + K_p C_p^2))$ training time. It also stores path and attention tensors, giving $O(nK_pJ + K_p^2)$ intermediate storage. Thus GPA is quadratic in the number of selected paths and linear in the total number of nodes, which can be expensive on large fraud graphs.

SEC-GFD uses $2C$ hybrid frequency bands. Its filtering cost is $O(CMd)$, and the local environmental constraint with L-layer SGC costs $O(LMh)$. However, its kNN view requires pairwise feature similarity, which costs $O(n^2 d)$ in the worst case and $O(nk)$ storage. Therefore, SEC-GFD has a quadratic node dependence, which is prohibitive for large-scale graphs.

BWGNN applies $C + 1$ Beta wavelet filters. Its sparse filtering cost is $O(CMh)$, and it stores $O(Cnh)$ frequency-band representations. It avoids the $O(n^2)$ kNN term of SEC-GFD, but it is still centralized and cannot handle federated graph fragmentation or isolated fraudulent nodes.

In summary, FedSACA+ has client time linear in the local edge count plus a small kNN overhead, server time linear in the number of clients and relations, and server memory independent of graph size. In contrast, GPA is quadratic in K_p, SEC-GFD is quadratic in n, and BWGNN, though linear in M, is centralized and lacks privacy preservation.

Methods	Client Mem	Server Mem	Client Storage	Server Storage
FedSACA+	$\mathbf{O(n_i d + \lvert E_i\rvert + n_i k_{\text{near}} + \lvert\mathcal{R}^*\rvert K + dh)}$	$\mathbf{O(N\lvert\mathcal{R}^*\rvert K)}$	$O(n_i d + \lvert E_i\rvert + n_i k_{\text{near}})$	$O(\lvert\mathcal{R}^*\rvert K)$
GPA	$O(nd + M + nK_pJ + LC_p^2 + JdC_p)$	–	$O(nd + M + nK_pJ + K_p^2)$	–
SEC-GFD	$O(nd + M + Cnh + nk + Cdh)$	–	$O(nd + M + Cnh + nk)$	–
BWGNN	$O(nd + M + Cnh + Cdh)$	–	$O(nd + M + Cnh)$	–

Table 4: Space complexity analysis. N: clients; n_i: nodes in client i; d: feature dimension; K: Chebyshev order in FedSACA; $\mathcal{R}^*$: relation set including the kNN relation; $|E_i|$: intrinsic edges in client i; k_{near}: kNN size in FedSACA+; h: hidden dimension; $n = \sum_i n_i$: total nodes in centralized baselines; M: total edges in centralized baselines; C: spectral order in baselines; K_p: selected paths in GPA; J: path length; C_p: path embedding dimension; L: attention/SGC layers. Best in bold.

Methods	Client Time	Server Time	Training Time	Inference Time
FedSACA+	$\mathbf{O(Kd(\lvert E_i\rvert + n_i k_{\text{near}}) + n_i dh + \lvert\mathcal{R}^*\rvert K)}$	$\mathbf{O(N\lvert\mathcal{R}^*\rvert K)}$	$\mathbf{O(Kd(\lvert E_i\rvert + n_i k_{\text{near}}) + n_i dh + \lvert\mathcal{R}^*\rvert K)}$	$\mathbf{O(Kd(\lvert E_i\rvert + n_i k_{\text{near}}) + n_i dh)}$
GPA	$O(n[IJ + K_pJ(d + C_p) + L(K_p^2 C_p + K_p C_p^2) + C_p h])$	–	$O(n[IJ + K_pJ(d + C_p) + L(K_p^2 C_p + K_p C_p^2) + C_p h])$	$O(n[IJ + K_pJ(d + C_p) + L(K_p^2 C_p + K_p C_p^2) + C_p h])$
SEC-GFD	$O(CMd + LMh + n^2 d + ndh)$	–	$O(CMd + LMh + n^2 d + ndh)$	$O(CMd + LMh + ndh)$
BWGNN	$O(CMh + ndh)$	–	$O(CMh + ndh)$	$O(CMh + ndh)$

Table 5: Time complexity analysis. FedSACA+ reports per-client local epoch and per-round server aggregation. Centralized baselines are reported per training/inference epoch. I: sampled paths in GPA; other symbols follow Table 4. Best in bold.