



Rethinking Client-oriented Federated Graph Learning

Zekai Chen
zackchen02@163.com
Beijing Institute of Technology
Beijing, China

Xunkai Li
cs.xunkai.li@gmail.com
Beijing Institute of Technology
Beijing, China

Yinlin Zhu
ylzhuawesome@163.com
Sun Yat-sen University
Guangzhou, Guangdong, China

Rong-Hua Li
lironghuabit@126.com
Beijing Institute of Technology
Beijing, China

Guoren Wang
wanggrbit@gmail.com
Beijing Institute of Technology
Beijing, China

Abstract

As a new distributed graph learning paradigm, Federated Graph Learning (FGL) facilitates collaborative model training across local systems while preserving data privacy. We review existing FGL approaches and categorize their optimization mechanisms into: (1) Server-Client (S-C), where clients upload local model parameters for server-side aggregation and global updates; (2) Client-Client (C-C), which allows direct exchange of information between clients and customizing their local training process. We reveal that C-C shows superior potential due to its refined communication structure. However, existing C-C methods broadcast redundant node representations, incurring high communication costs and privacy risks at the node level. To this end, we propose FedC4, which combines graph Condensation with C-C Collaboration optimization. Specifically, FedC4 employs graph condensation technique to refine the knowledge of each client's graph into a few synthetic embeddings instead of transmitting node-level knowledge. Moreover, FedC4 introduces three novel modules that allow the source client to send distinct node representations tailored to the target client's graph properties. Experiments on eight public real-world datasets show that FedC4 outperforms state-of-the-art baselines in both task performance and communication cost. Our code is now available on <https://github.com/Ereshkigal1/FedC4>.

CCS Concepts

• Computing methodologies → Artificial intelligence.

Keywords

Graph Neural Network; Federated Learning; Dataset Condensation

ACM Reference Format:

Zekai Chen, Xunkai Li, Yinlin Zhu, Rong-Hua Li, and Guoren Wang. 2025. Rethinking Client-oriented Federated Graph Learning. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management (CIKM '25)*, November 10–14, 2025, Seoul, Republic of Korea. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3746252.3761228>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CIKM '25, November 10–14, 2025, Seoul, Republic of Korea.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-2040-6/2025/11
<https://doi.org/10.1145/3746252.3761228>

1 Introduction

Graph Neural Networks (GNNs) have emerged as powerful tools for capturing the underlying topological and attribute patterns in graph-structured data, achieving notable success in real-world scenarios such as biomedical [15], recommendation [21], and finance [14]. However, most existing GNN approaches rely on the assumption of centralized data storage, where the entire graph data is collected and accessed by a single institution. In contrast, in many real-world scenarios, graph data is inherently distributed across multiple institutions, with each institution only able to access its own private data due to privacy or storage constraints [10, 32]. This assumption-reality gap introduces significant challenges in enabling GNNs to effectively harness the full potential of collective intelligence in the decentralized setting.

To this end, Federated Graph Learning (FGL) has emerged as a promising solution, combining the advantages of federated learning with GNNs. In this framework, multiple local graph systems collaborate to train a global GNN, with coordination managed by a trusted central server. Existing FGL optimization strategies can be classified into two paradigms: (1) **Server-Client (S-C)**, where clients upload local parameters for server-side aggregation and global updates. Approaches in this paradigm aim to enhance the global model's generalizability through modifications to training modules or architecture [10, 23]. (2) **Client-Client (C-C)**, which facilitates the exchange of information like node embeddings between clients, enabling each client to customize its training process based on information from other clients. Research in this paradigm is dedicated to developing cross-client collaboration and local training processes for improved performance [29, 32].

Despite the effectiveness of both paradigms, we compare them in terms of information diversity to highlight their progressive relationship in personalization potential. As shown in Fig. 1(a), in the S-C paradigm, each of the K clients uploads local information (e.g., model parameters) to a central server, which aggregates it into a global model in traditional S-C (**Level 1**) or up to K distinct models in personalized S-C (**Level 2**). While this allows some personalization, the information diversity is limited to K pieces, and does not support direct local optimization. In contrast, the broadcast-based C-C paradigm not only covers all of the S-C processes but also enables clients to share their local information, enhancing local optimization with global features and topology. Despite there being K^2 types of inter-client communication relations, this results in only K distinct pieces of information (**Level 3**).

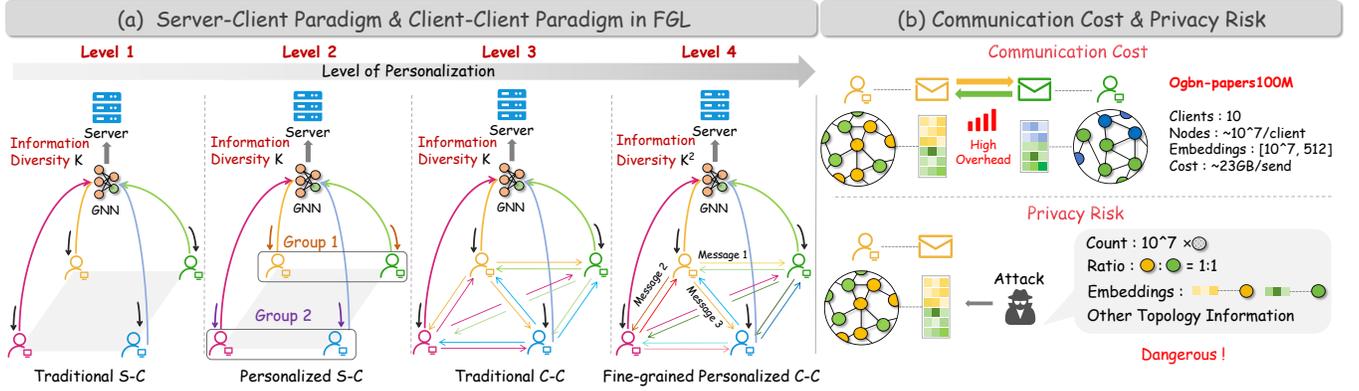


Figure 1: (a) Progressive relationship in personalization potential in S-C and C-C: *Traditional S-C*: Clients send local updates to a central server for global aggregation. *Personalized S-C*: The server clusters clients and provides group-specific aggregation. *Traditional C-C*: Clients exchange identical information with all other clients. *Fine-Grained Personalized C-C*: Clients exchange distinct and personalized information with different clients. (Different colored arrows represent unique information exchanged between clients.) (b) A key challenge is the substantial communication overhead and heightened risk of privacy leakage.

Hence, the broadcast-based C-C overshadows its personalization potential due to the lack of a refined communication design. Intuitively, a more fine-grained personalized C-C approach allows each source client to tailor information for each target client instead of broadcasting the same information, increasing the diversity of inter-client communication to K^2 distinct pieces (**Level 4**). This motivates us to rethink the design of the FGL optimization strategy under the C-C paradigm to achieve fine-grained personalization.

Moreover, to devise an effective FGL optimization strategy within the fine-grained personalized C-C paradigm, a critical challenge lies in addressing the substantial communication overhead and the heightened risk of privacy leakage, as illustrated in Fig. 1(b). First, all existing C-C technologies transmit data at the node level, meaning that as the volume of graph data grows, the communication cost becomes increasingly prohibitive. From a privacy perspective, traditional C-C methods such as FedSage+ [32] and FedGCN [29] rely on transferring node hidden representations or propagated features. However, extensive studies have shown that these approaches inherently expose significant security vulnerabilities [16]. Although FedDEP [31] enhances security with noiseless differential privacy, attackers can still infer sensitive information (e.g., node quantity) from side-channel data.

Building on these insights, we propose a framework combining graph Condensation with C-C Collaboration optimization (FedC4). Specifically, FedC4 employs a graph condensation technique to refine the knowledge of each client’s private graph into a few synthetic node embeddings instead of directly transmitting node-level knowledge, so as to achieve low-cost and high-privacy knowledge sharing among clients. Moreover, to achieve fine-grained personalized local optimization, FedC4 introduces three novel modules that allow the source client to send distinct node representations tailored to the target client’s graph properties, thereby enhancing its local optimization with global features and topology.

Our main contributions can be summarized as: (1) **New Observation**. This work introduces a novel viewpoint by categorizing existing FGL strategies based on personalization potential. (2) **New Paradigm**. We introduce a novel framework, Federated Graph Learning

with Graph Condensation and Client-Client Collaboration (**FedC4**), designed to enhance the C-C paradigm in FGL. By incorporating personalized and fine-grained inter-client knowledge sharing and local condensation, this framework effectively address challenges in C-C FGL. (3) **SOTA Performance**. We conduct comprehensive experiments on 8 datasets including transductive, inductive, heterogeneous, and large-scale graph settings. FedC4 outperforms the SOTA baselines with an average performance gain of 1.73% and an efficiency improvement of up to 1000×.

2 Preliminaries and Related Work

2.1 Preliminaries

Graph Neural Network. GNNs are a class of deep learning models that operate on graph-structured data, leveraging both topology and node features to learn representations. A graph $G = (A, X, Y)$ consists of an adjacency matrix A , a node feature matrix X , and node labels Y . GNNs iteratively aggregate and update node features through a message-passing process. For example, the ℓ -th layer of a GCN [9] propagates information as:

$$H^{(\ell)} = \text{ReLU}(\hat{A}H^{(\ell-1)}W^{(\ell)}), \quad (1)$$

where $H^{(0)} = X$, \hat{A} is the normalized adjacency matrix with self-loops, and $W^{(\ell)}$ is the trainable weight matrix.

After L layers, GNNs capture L -hop neighborhood information, and the final node embeddings can be written as:

$$H = f(A, X), \quad (2)$$

where f is the GNN model. The final embeddings $h_i^{(L)}$ are passed to a classifier F for task-specific predictions:

$$z_i = F(h_i). \quad (3)$$

Federated Graph Learning. FGL extends federated learning to graph-structured data, enabling clients to train local GNNs on private subgraphs while preserving data privacy. The optimization

objective in FGL for a node classification task can be formulated as:

$$\min_{\theta} \frac{1}{N} \sum_{i=1}^N \frac{|V^i|}{|V|} \mathcal{L}(f_{\theta}(V^i), Y^i), \quad (4)$$

where V^i represents the set of nodes in the subgraph owned by client i , Y^i are the corresponding labels, f_{θ} denotes the GNN parameterized by θ , and \mathcal{L} is the task-specific loss.

In the FGL framework, each client independently trains a local GNN and then shares model parameters. The global model is then updated by aggregating these local updates in server-side, and then shared with all clients for the next round of training.

Graph Condensation. GC reduces large-scale graph data into a smaller, representative graph $S = (A', X', Y')$, where $A' \in \mathbb{R}^{N' \times N'}$ represents the adjacency matrix, $X' \in \mathbb{R}^{N' \times d}$ denotes node features, and $Y' \in \{1, \dots, C\}^{N'}$ indicates node labels. The objective is to train a GNN f_{θ} on S with performance comparable to training on the original graph G , while significantly reducing computational costs.

GC optimizes the synthesized graph S by minimizing the task-specific loss associated with the model's predictions:

$$\min_S \mathcal{L}(\text{classifier}(f_{\theta}(S)), Y), \quad (5)$$

where \mathcal{L} is the loss function (e.g., cross-entropy). Synthesized graph S is designed to retain the structural and feature distributions of original graph G , ensuring its utility in downstream tasks.

2.2 Related Work

Graph Neural Networks. GNNs have advanced from early attempts to extend convolution to graphs [1], which often involved high parameter counts. GCN [9] addressed this using a first-order Chebyshev filter to capture local neighborhoods, while GAT [19] introduced attention mechanisms for weighted aggregation. GraphSAGE [6] further enhanced message aggregation with learnable functions. More research on GNNs can be found in surveys [22, 34].

Federated Graph Learning. FGL combines the principles of federated learning with GNNs and has emerged as a prominent research direction in recent years. The S-C paradigm, exemplified by methods such as FedLIT [24] and FedGTA [10], aggregates client parameters on a central server. In contrast, the C-C paradigm enables direct message exchanges between clients. Representative methods include FedSage+ [32], which facilitates subgraph training via embedding exchanges, FedDEP [31], which incorporates dependency-aware communication, and FedGCN [29], which exchanges compressed intermediate features during GNN training. However, these methods still face challenges in fine-grained personalization, communication costs, and privacy risk, highlighting the need for more effective solutions.

Graph Condensation. GC aims to efficiently distill a small-scale synthetic graph from a large-scale graph while ensuring that models trained on the condensed graph achieve performance comparable to those trained on the original graph. GCond [8] leverages gradient matching to optimize adjacency matrices and node features, while SGDD [26] incorporates Laplacian energy matching to enhance structural representation. MCond [4] focuses on feature condensation for multimodal graph data, and SFGC [33] introduces a global consistency mechanism to improve condensation performance.

3 Proposed Framework: FedC4

3.1 Framework Overview

As illustrated in Fig. 2, FedC4 consists of four key components: Local Graph Condensation, Customizer Module (CM), Node Selector Module (NS), and Graph Rebuilder Module (GR). While the figure describes the high-level process of the framework, the detailed design and functionalities of the GNN models, as well as the individual modules, will be elaborated in subsequent sections.

We also notice a recent work, FedGC [25], which integrates FGL concepts into the GC process by decoupling the gradient matching procedure: clients compute local gradients, and the server aggregates them to optimize a global condensed graph. Essentially, FedGC applies FGL techniques to support the task of GC.

In contrast, FedC4 leverages GC techniques to enhance C-C FGL. It targets fine-grained personalization, communication efficiency, and privacy protection under the C-C paradigm. These goals are achieved by combining GC with three dedicated modules: CM, NS, and GR. This design makes FedC4 fundamentally different from FedGC in both objective and architecture: FedGC uses FL to facilitate GC, whereas FedC4 uses GC to empower C-C FGL.

3.2 Local Graph Condensation

As mentioned in Sec. 1, the C-C paradigm introduces significant communication costs and privacy challenges. GC addresses these issues by synthesizing compact and secure graphs S that retain the structural and semantic properties of the original graph \mathcal{G} .

GC initialization begins by randomly initializing the feature matrix X' from a Gaussian distribution. Labels Y' are selected to match the label distribution of \mathcal{G} . The adjacency matrix A' is dynamically generated by passing X' through a randomly initialized trainable MLP, allowing the condensed graph to adaptively learn its structure during optimization.

To ensure that the condensed graph preserves the original graph's properties, the optimization process minimizes the gradient matching loss of the GNN parameterized by θ :

$$\mathcal{L}_{\text{mat}} = \sum_{i=1}^{|\theta|} \|\nabla_{\theta_i} \mathcal{L}^G - \nabla_{\theta_i} \mathcal{L}^S\|^2, \quad (6)$$

where \mathcal{L}^G and \mathcal{L}^S denote the cross-entropy losses of the original and condensed graphs, respectively. \mathcal{L}_{mat} is then used to optimize the MLP. To further refine A' , a sparsification step is applied by thresholding edge weights:

$$A'_{ij} = \begin{cases} A'_{ij}, & \text{if } A'_{ij} > \delta, \\ 0, & \text{otherwise,} \end{cases} \quad (7)$$

where δ is the sparsification threshold, consistent with same hyperparameter selection described in GCond [8]. The complete process of local condensation is also described in GCond [8].

3.3 Customizer (CM)

In the C-C paradigm, the CM module addresses the limitation of broadcasting identical content by facilitating collaboration through the selective sharing of local embedding statistics, such as distributions and prototype embeddings. During the first round, local statistics are shared with all clients to establish a baseline understanding.

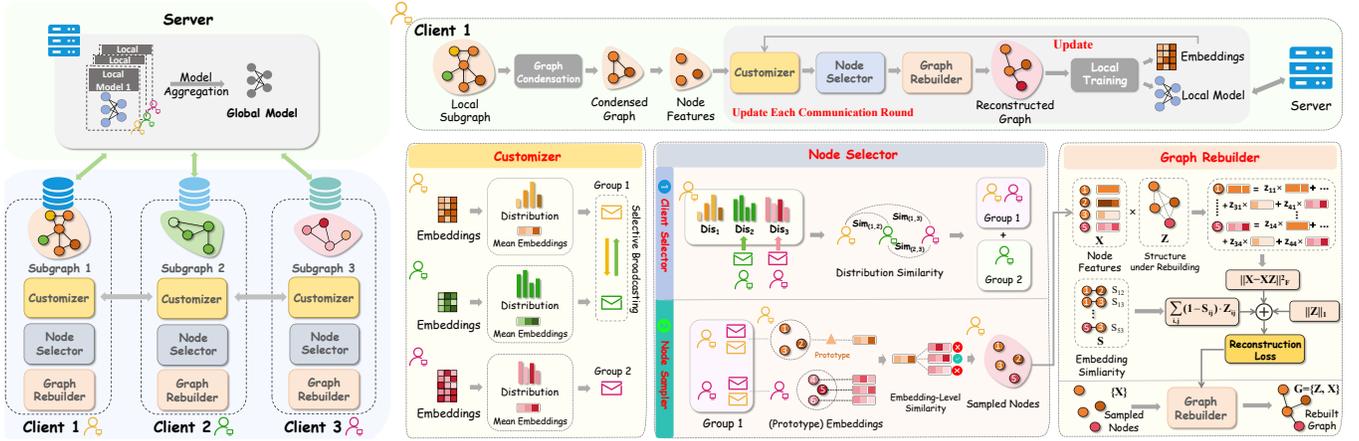


Figure 2: Overview of our proposed FedC4 framework. The left side illustrates the overall workflow of server-side broadcasting and aggregation. The right side illustrates client-side operations including local condensation and three core modules.

In subsequent rounds, these statistics are selectively broadcast to same-cluster clients (C_{same}), as determined by the clustering results from the Node Selector module in the previous round, enabling personalized and selective collaboration.

For a client c , local node embeddings $\mathbf{H}_c = \{\mathbf{h}_i \mid i \in \mathcal{V}_c\}$ and node number \mathcal{V}_c are used to compute the embedding distribution Dis_c and prototype embedding μ_c :

$$Dis_c = \{\|\mathbf{h}_i\| \mid i \in \mathcal{V}_c\}, \quad \mu_c = \frac{1}{|\mathcal{V}_c|} \sum_{i \in \mathcal{V}_c} \mathbf{h}_i, \quad (8)$$

where $\|\mathbf{h}_i\|$ is the norm of embedding \mathbf{h}_i . These statistics are normalized for consistency across clients and then broadcast according to the communication strategy:

$$\mu_{\text{global}} = \frac{1}{C} \sum_{c=1}^C \mu_c, \quad \sigma_{\text{global}} = \sqrt{\frac{1}{C} \sum_{c=1}^C \|\mu_c - \mu_{\text{global}}\|^2}. \quad (9)$$

To ensure consistent scale and distribution of embeddings across all clients, we normalize the individual client statistics using the global metrics as follows:

$$\mu'_c = \frac{\mu_c - \mu_{\text{global}}}{\sigma_{\text{global}} + \epsilon}, \quad Dis'_c = \left\{ \frac{\|\mathbf{h}_i\| - \mu_{\text{dis}}}{\sigma_{\text{dis}}} \mid \|\mathbf{h}_i\| \in Dis_c \right\}, \quad (10)$$

where ϵ is a small constant to prevent division by zero, and μ_{dis} and σ_{dis} are the mean and standard deviation of the norms of embeddings across all clients, respectively. The updated values Dis'_c and μ'_c are then selectively broadcast to clients based on clustering results C_{target} from the Node Selector, enabling fine-grained and personalized collaboration:

$$\{Dis'_c, \mu'_c \mid c' \in C_{\text{target}} \setminus \{c\}\}, \quad (11)$$

where C_{target} refers to all clients (C) in the first round and same-cluster clients (C_{same}) in subsequent rounds.

By equipping clients with normalized global statistics and leveraging selective communication based on previous clustering results,

the CM module ensures informed local decision-making, improving fine-grained collaboration in the C-C paradigm. The complete process of CM module is described in Algorithm 1 in Appendix.

3.4 Node Selector (NS)

In addition to selective broadcasting, it is essential to tailor the messages for each client. By identifying representative nodes that capture key information from local graphs, NS ensures that exchanged data is both relevant and contextually aware. This targeted strategy addresses the limitations of uniform broadcasting, enabling more personalized and efficient collaboration.

For each client $c \in C$, the module computes the Sliced Wasserstein Distance (SWD) between its distribution Dis_c and those of other clients $Dis_{c'}$, where $c' \in C \setminus \{c\}$:

$$SWD_{c,c'} = \int_{S^{d-1}} \sup_{\theta \in S^{d-1}} \left| \int_{\mathbb{R}^d} \theta \cdot x dDis_c(x) - \int_{\mathbb{R}^d} \theta \cdot y dDis_{c'}(y) \right| d\theta, \quad (12)$$

where S^{d-1} represents the unit sphere in \mathbb{R}^d . SWD effectively reduces the computational complexity in high-dimensional spaces compared to the traditional WD by projecting the distributions onto a line and measuring their one-dimensional discrepancies.

Clients are then grouped into clusters $\{C_1, C_2, \dots\}$ based on SWD similarity. Within each cluster C_c , the cosine similarity between client c 's node embeddings \mathbf{h}_i and prototype embedding $\mu_{c'}$ of other clients c' in C_c is computed:

$$S(\mathbf{h}_i, \mu_{c'}) = \frac{\mathbf{h}_i \cdot \mu_{c'}}{\|\mathbf{h}_i\| \|\mu_{c'}\|}. \quad (13)$$

Nodes exhibiting similarity scores exceeding a designated threshold τ are incorporated into the representative set \mathcal{S}_c . The selection criteria for τ and its consequent effects on the accuracy curve are comprehensively elucidated in hyper-parameter study in Sec. 5.4.

By selecting nodes aligned with shared semantic patterns, the NS module identifies and selects representative nodes for effective and fine-grained collaboration. The complete process of NS module is described in Algorithm 2 in Appendix.

3.5 Graph Rebuilder (GR)

The GR module generates appropriate structures for all candidate nodes selected by NS by leveraging the concept of self-expressive graph reconstruction [5], rather than simply linking newly introduced nodes to the original condensed graph. To adapt to the FGL settings, GR incorporates both embedding similarity loss and sparsification loss, enabling the generation of graph structures that are well-suited for distributed learning. Specifically, GR constructs adaptive and reliable graph topologies for the candidate nodes selected by NS, ensuring that the reconstructed graph effectively preserves essential structural and semantic information while maintaining sparsity for efficiency. By constructing adaptive graph topologies tailored to the specific features and relationships of the candidate nodes, GR ensures fine-grained personalized structural information.

For each client $c \in C$, GR takes the node feature matrix \mathbf{X}_c of the sampled nodes as an extra input. Embedding similarity between candidate nodes can be computed as:

$$S_{ij} = \frac{\mathbf{h}_i \cdot \mathbf{h}_j}{\|\mathbf{h}_i\| \|\mathbf{h}_j\|}, \quad (14)$$

where S_{ij} captures the similarity between nodes i and j . Using this similarity, the module updates the reconstructed adjacency matrix \mathbf{Z} by minimizing the reconstruction loss:

$$\mathcal{L}_{\text{Rec}} = \alpha \|\mathbf{X}_c - \mathbf{X}_c \mathbf{Z}\|_F^2 + \beta \|\mathbf{Z}\|_1 + \sum_{i,j} (1 - S_{ij}) \cdot \mathbf{Z}'_{ij}, \quad (15)$$

The reconstruction loss \mathcal{L}_{Rec} consists of three terms: the first term, $\|\mathbf{X}_c - \mathbf{X}_c \mathbf{Z}\|_F^2$, ensures that the reconstructed adjacency matrix \mathbf{Z} preserves the structural relationships of the original graph in the feature space. The second term, $\|\mathbf{Z}\|_1$, promotes sparsity in \mathbf{Z} , retaining only the most significant edges and avoiding overly dense graphs. The third term, $\sum_{i,j} (1 - S_{ij}) \cdot \mathbf{Z}'_{ij}$, aligns the reconstructed graph with node similarity S_{ij} , penalizing edges inconsistent with the semantic relationships captured by the embeddings. The hyperparameters α and β control the balance between these terms.

The reconstructed graph is then used to train local GNN, with the updated model uploaded for global aggregation. Although partially dependent on the smoothing effects of graph neural networks, the excellent performance on heterogeneous graph datasets in Sec. 5.2 demonstrates that the GR module has robust capabilities for handling heterogeneous graphs. The detailed process of GR module is described in Algorithm 3 in Appendix.

4 Theoretical Analysis

4.1 Privacy Protection

In this section, we theoretically demonstrate that GC inherently provides privacy protection when transmitting synthetic graph embeddings in C-C paradigms.

Before presenting the proof, we clarify the privacy scenario GC aims to protect against. The objective is to minimize the risk of inferring individual node attributes from synthetic graph embeddings. This is particularly relevant in scenarios where node data may contain sensitive information. The privacy analysis presented here focuses on ensuring that changes in individual nodes of the original graph \mathcal{G} do not significantly influence the derived synthetic graph \mathcal{S} , thus maintaining a high level of privacy.

Proof. Given the synthetic graph $\mathcal{S} = (\mathcal{V}_{\text{syn}}, \mathcal{E}_{\text{syn}}, \mathbf{X}_{\text{syn}})$ and the original graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ with $n = |\mathcal{V}|$ nodes, the training loss is expressed as:

$$\mathcal{L}(\mathcal{S}, \mathcal{G}) = \frac{1}{n} \sum_{i=1}^n \ell(f_{\theta}(\mathbf{x}_i), y_i), \quad (16)$$

where ℓ is the sample-wise loss, f_{θ} is the model parameterized by θ , and (\mathbf{x}_i, y_i) denotes the feature-label pair. Removing node v_j modifies the graph to \mathcal{G}^{-j} , the adjusted loss is:

$$\mathcal{L}(\mathcal{S}, \mathcal{G}^{-j}) = \frac{1}{n-1} \sum_{\substack{i=1 \\ i \neq j}}^n \ell(f_{\theta}(\mathbf{x}_i), y_i). \quad (17)$$

The gradient difference due to node removal expands to:

$$\begin{aligned} \Delta \nabla &= \nabla \mathcal{L}(\mathcal{S}, \mathcal{G}) - \nabla \mathcal{L}(\mathcal{S}, \mathcal{G}^{-j}) \\ &= \frac{1}{n} \nabla \ell(f_{\theta}(\mathbf{x}_j), y_j) - \frac{1}{n(n-1)} \sum_{\substack{i=1 \\ i \neq j}}^n \nabla \ell(f_{\theta}(\mathbf{x}_i), y_i). \end{aligned} \quad (18)$$

Assuming the gradient of each sample is bounded as $\|\nabla \ell(f_{\theta}(\mathbf{x}_i), y_i)\| \leq C$, gradient difference is bounded as:

$$\|\Delta \nabla\| \leq \frac{C}{n} + O\left(\frac{1}{n^2}\right). \quad (19)$$

For large n , the second term is negligible, yielding:

$$\|\Delta \nabla\| \leq O\left(\frac{1}{n}\right). \quad (20)$$

This bounded gradient difference propagates through the GNN layers, impacting the node embeddings as:

$$\Delta_j = \|\mathbf{H}_{\text{syn}} - \mathbf{H}_{\text{syn}}^{-j}\| \leq \|\mathbf{H}^{-1}\| \cdot \|\Delta \nabla\|, \quad (21)$$

where $\|\mathbf{H}^{-1}\|$ is the norm of the inverse of the Hessian matrix. Substituting the bound on $\Delta \nabla$, we obtain:

$$\Delta_j \leq O\left(\frac{m}{n}\right), \quad (22)$$

where $m = |\mathcal{V}_{\text{syn}}|$ is the size of the synthetic graph. \square

Conclusion. The bound on Δ_j demonstrates that the influence of a single node v_j in the original graph on the synthetic graph embeddings diminishes as the size of the original graph increases. This theoretical derivation validates the intrinsic privacy-preserving nature of GC, proving that the technique itself is secure. Furthermore, as highlighted in the Sec. 1, GC transitions the potential privacy exposure from the level of the entire graph's original nodes to that of the compressed synthetic nodes, effectively reducing the granularity of sensitive information leakage. It is important to note that the provided proof primarily applies to scenarios where nodes do not overlap across different subgraphs or communities. In cases where nodes belong to overlapping clusters, additional considerations and methods may be necessary to ensure privacy protection. These results establish GC as a secure solution for privacy-sensitive FGL scenarios within the defined application boundaries.

Table 1: Performance comparison across datasets with varying compression ratios and baselines.

Property	Dataset	Ratio	FL		FL+Graph Reduction			FL+GC			FGL C-C			FGL S-C	Ours
			FedAvg	FedDC	Random	Herdling	Coarsening	GCond	DosCond	SFGC	FedSage+	FedGCN	FedDEP	Best S-C	FedC4
Small	Cora	8%			81.35±0.33	81.87±0.25	82.63±0.56	83.61±1.14	82.24±0.41	83.58±0.59					86.02±0.56
		4%	84.70±0.61	84.79±0.69	80.61±1.04	80.63±0.96	81.81±0.62	82.97±0.85	81.83±0.49	83.15±0.77	83.69±0.73	84.32±0.55	85.01±0.66	84.97±0.70	84.98±0.65
		2%			79.03±1.25	79.74±1.11	80.73±0.70	81.36±1.33	81.49±0.62	81.42±0.98					84.33±0.80
	Citeseer	8%			65.97±0.27	68.60±0.30	65.57±0.19	72.03±1.02	72.58±0.65	73.11±0.25					75.95±0.66
		4%	72.11±0.73	71.94±0.52	65.89±0.99	67.52±0.87	66.83±0.40	67.27±1.35	68.02±0.81	71.00±0.29	70.77±0.88	74.37±0.49	73.83±0.86	74.89±0.93	74.84±0.74
		2%			65.73±1.56	65.73±1.02	67.91±0.45	67.46±2.24	70.41±1.22	69.23±0.66					73.91±1.17
Medium	Arxiv	0.4%			58.91±0.82	56.63±0.77	59.96±0.16	65.98±0.44	63.41±0.23	66.17±0.52					67.66±0.38
		0.2%	62.16±0.53	61.20±0.66	57.60±1.03	56.22±0.94	59.11±0.29	65.38±0.82	63.25±0.37	65.62±0.68	57.07±0.55	64.17±0.69	64.31±0.79	63.14±0.62	66.95±0.76
		0.1%			57.29±1.89	55.98±1.22	57.38±0.41	64.81±1.17	62.33±0.52	65.09±0.57					65.42±1.36
	Physics	1.4%			92.25±0.45	92.72±0.51	90.24±0.28	92.88±0.69	91.56±0.38	92.97±0.55					94.61±0.72
		0.7%	92.74±0.43	94.02±0.37	91.37±0.75	92.37±0.33	89.73±0.22	92.13±1.01	90.91±0.42	92.41±0.63	93.01±0.68	92.88±0.65	93.72±0.80	94.12±0.53	93.65±0.82
		0.35%			91.55±1.63	91.53±0.82	88.42±0.40	91.48±1.64	90.52±0.68	91.86±0.79					93.41±1.33
Inductive	Flickr	1%			45.21±0.15	44.80±0.18	46.39±0.09	47.11±0.22	45.54±0.38	47.03±0.25					48.62±0.27
		0.5%	46.09±0.23	43.73±0.21	44.75±0.31	44.46±0.47	46.14±0.13	47.06±0.28	45.01±0.78	46.86±0.42	45.27±0.38	46.44±0.52	47.36±0.44	48.11±0.30	48.33±0.22
		0.1%			44.39±0.42	44.08±0.69	45.82±0.22	46.75±0.39	44.86±1.02	46.42±0.70					47.89±0.31
	Reddit	1%			87.82±0.73	89.37±0.86	88.12±0.66	91.55±0.49	91.39±0.33	91.44±0.38					92.38±0.21
		0.5%	91.80±0.42	91.85±0.38	87.57±0.66	89.01±1.03	87.88±0.91	91.43±0.52	91.28±0.25	91.32±0.42	OOT	91.84±0.55	OOT	91.80±0.33	91.73±0.38
		0.1%			86.89±1.22	88.34±1.83	87.27±1.14	91.07±0.87	91.03±0.38	91.10±0.53					91.34±0.56
Large	Products	0.4%			82.65±0.55	82.41±0.21	82.00±0.18	83.39±0.42	82.52±0.33	83.02±0.27					85.33±0.45
		0.2%	82.31±0.47	78.41±0.38	82.38±0.71	81.99±0.39	81.72±0.32	83.08±0.85	82.33±0.47	82.79±0.56	OOM	84.20±0.53	OOM	OOM	84.53±0.72
		0.1%			81.77±1.31	82.30±0.60	81.45±0.48	82.81±1.42	81.94±0.73	82.57±0.80					84.17±0.99
Hetero.	Empire	1.4%			65.26±0.61	62.11±0.18	61.76±0.37	65.80±0.38	68.73±0.25	68.97±0.16					71.05±0.27
		0.7%	45.24±0.32	39.45±0.28	65.01±0.45	61.15±0.38	61.09±0.22	65.37±0.42	66.17±0.33	66.27±0.18	69.68±0.70	68.82±0.61	67.55±0.53	49.08±0.57	68.19±0.44
		0.35%			64.90±1.08	60.54±0.75	60.14±0.46	64.66±0.81	65.39±0.58	66.02±0.26					67.30±0.58

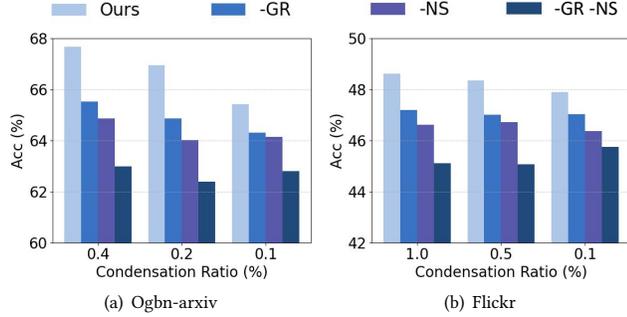


Figure 3: Ablation Study of FedC4 on NS and GR Modules.

4.2 Communication Cost

We analyze the theoretical communication cost of FedC4 and compare it with that of the traditional S-C, personalized S-C and traditional C-C paradigms, as shown in Table 2.

Table 2: Comparison of Communication Costs

Method	Communication Cost
S-C Paradigm (lv.1, lv.2)	$O(2Cp)$
C-C Paradigm (lv.3)	$O(C^2Nd)$
FedC4 (lv.4)	$O(C \log C N' d)$

Here, C is the number of clients, p is the model size, N is the number of nodes, N' is the number of condensed nodes ($N' \ll N$), and d is the embedding dimension. Traditional S-C scales linearly with C , C-C grows quadratically, while FedC4 reduces communication to $O(C \log C N' d)$ through GC and selective broadcasting in CM, achieving better scalability. The local graph condensation process is executed entirely on each individual client, incurring absolutely no additional communication costs at all.

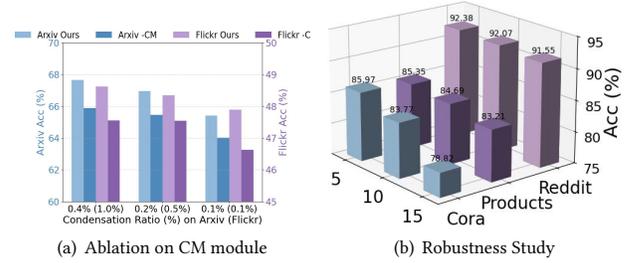


Figure 4: (a) Ablation Study of FedC4 on CM Module. (b) Accuracy Comparison with varying clients (5, 10, and 15) on Cora (8%), Products (0.4%), and Reddit (1%).

5 Experiments

In this section, we evaluate the effectiveness of our framework, FedC4, using 8 graph datasets. We introduce baseline GC and FGL approaches, followed by a detailed experimental setup. The study addresses the following questions: **Q1:** Does FedC4 outperform state-of-the-art FGL frameworks and GC methods across diverse graph settings? **Q2:** What drives the performance improvements in FedC4? **Q3:** How does FedC4 perform in terms of efficiency and scalability compared to other FGL methods? **Q4:** Does FedC4 ensure privacy within the C-C paradigm in FGL scenarios?

5.1 Experimental Setup

Datasets. We evaluate the performance of the proposed framework on 8 datasets, including transductive and inductive settings, as well as large-scale and heterogeneous graphs. For the transductive setting, we use *Cora*, *Citeseer* [27], *Ognb-arxiv* [28], and *Physics* [18] datasets. For the inductive setting, we employ *Flickr* and *Reddit* [30] datasets. Additionally, we include the large-scale *Ognb-products* [18]

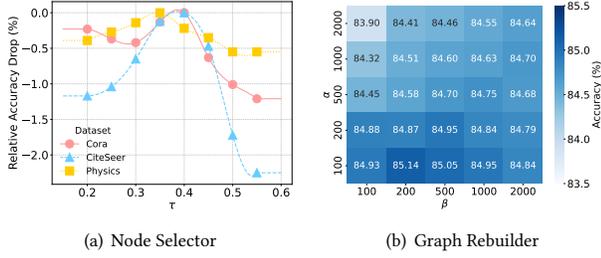


Figure 5: Hyper-parameter study for (a) τ in Node Selector Module. (b) α , β in Graph Rebuilder Module.

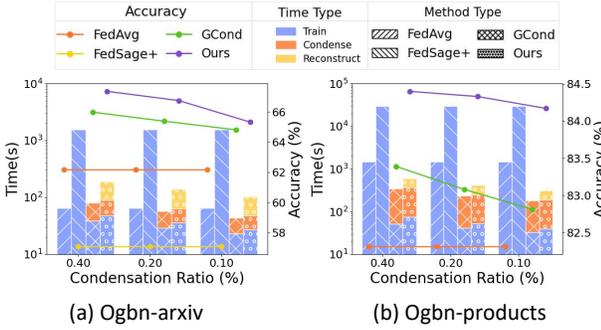


Figure 6: Efficiency evaluation on Arxiv and Products.

dataset and the heterogeneous *Roman-empire* [13] dataset to evaluate scalability and robustness of FedC4.

Baselines. We compare our proposed method to 12 baselines: (1) two FL methods (FedAvg, FedDC) [3, 12], (2) three Graph Reduction methods in the FL setting (Random, Herding, Coarsening) [11, 20], (3) three GC methods in the FL setting (GCond, SFGC, SGDD) [8, 26, 33], and (4) three C-C FGL methods (FedSage, FedDEP, FedGCN) [29, 31, 32]. Additionally, we include the best-performing S-C FGL method (FGSSL, FedGTA, FedTAD) on each dataset [7, 10, 35]. For Graph Reduction and GC baselines, we apply client-side condensation with three different ratios and train local models on the condensed graphs. The compression ratios are determined based on prior studies [8] and specific FGL settings.

Implementation Details. We simulate the subgraph systems using the Louvain community detection algorithm, specifically partitioning the original graph into 5 distinct communities. For local GNNs, we employ a 2-layer GCN with a hidden dimension of 64. All networks are optimized using SGD [17] with a weight decay of 5×10^{-4} . Each client is trained for 5 local epochs with 200 communication rounds. For the parameters of the NS and GR module, we set $\tau = 0.38$, $\alpha = 150$, and $\beta = 250$.

5.2 Performance Comparison

To answer Q1, the results in Table 1 demonstrate that FedC4 consistently outperforms most baseline methods across FL, FL+Graph Reduction, FL+GC, and FGL categories: **(1) FL methods:** Compared to FedAvg and FedDC, FedC4 achieves an average improvement of 12.03% by leveraging the structural information of graph data,

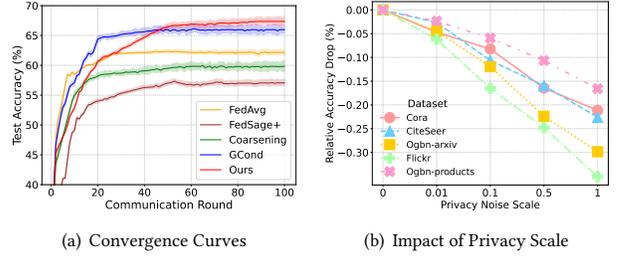


Figure 7: (a) Convergence curves. (b) Impact of privacy scale.

effectively addressing data heterogeneity and overcoming the limitations of vanilla aggregation. **(2) FL+Graph Reduction methods:** Compared to Herding and Coarsening, FedC4 not only maintains efficiency but also improves accuracy significantly by 7.88% and 7.76%. GC strategies consistently outperform other graph reduction techniques in both accuracy and scalability, making GC the preferred choice for our framework. **(3) FL+GC methods:** FL+GC methods such as GCond and SFGC, while achieving efficiency gains, suffer from inter-client heterogeneity. Compared to GCond, FedC4 achieves an accuracy improvement of 3.01%, and compared to SFGC, FedC4 improves by 2.69%. **(4) FGL methods:** In the FGL category, specialized methods like FedSage+ (C-C) and FGSSL (S-C) face scalability challenges, often resulting in OOM or OOT errors on large-scale datasets. Compared to FGL methods, FedC4 improves by an average of 7.52% over S-C methods and 3.58% over other C-C methods. These performance comparisons between S-C, C-C, and FedC4 validate the communication granularity issues in Sec. 1, with FedC4 overcoming the limitations of traditional S-C and C-C methods through more fine-grained personalized communication.

5.3 Ablation Study

To answer Q2, we conduct an ablation study by systematically removing Graph Rebuilder (-GR) and Node Selector (-NS) modules. Ogbn-arxiv, a moderate-size transductive dataset, and Flickr, an inductive dataset, were chosen for the ablation study to evaluate the performance of the modules in different learning settings. As shown in Fig. 3, combining all modules achieves the highest performance, validating the effectiveness of fine-grained communication.

In addition to the GR and NS modules, we also conducted an ablation study on the CM module to explore the impact of different broadcasting strategies on model performance. Specifically, we compared the effects of full broadcasting and selective broadcasting within the CM module. In terms of theoretical communication cost, full broadcasting incurs a cost of $O(C^2 N' d)$, whereas selective broadcasting significantly reduces this cost to $O(C \log C N' d)$. As shown in Fig. 4(a), selective broadcasting within the CM module achieves superior performance by reducing unnecessary communications and focusing on the most pertinent cross-client information exchanges. This result further confirms the effectiveness of fine-grained communication strategies in simultaneously enhancing performance and communication efficiency.

In addition to performance-level ablations, we also evaluate the structural impact of the GR module to understand how well it recovers the original graph topology after condensation in Appendix B.

5.4 Efficiency and Scalability Study

To answer Q3, we conduct four experiments: a Robustness Study to investigate the impact of client number variation on model performance, an Efficiency Study to evaluate efficiency and accuracy, a Hyper-parameter Study to analyze the impact of key parameters in NS and GR module, and a Convergence Study to compare the training stability of FedC4 with other methods.

Robustness study. To assess the robustness of our method to different client configurations, we conduct a study under varying numbers of clients (5, 10, and 15) on three datasets: Cora (8%), Products (0.4%), and Reddit (1%). As shown in Fig. 4(b), FedC4 consistently maintains high accuracy as the number of clients increases, indicating strong robustness to data partitioning granularity. This robustness study highlights the stability of our method across different federated settings (number of clients).

Efficiency study. To evaluate the efficiency of our proposed method, we conduct experiments on two datasets, Arxiv and Products, under varying condensation ratios (0.40%, 0.20%, 0.10%). As shown in Fig. 6(a) and (b), FedC4 achieves superior accuracy compared to baseline approaches (FedAvg, FedSage+ and GCond), while maintaining a competitive efficiency. (FedSage+ OOM on products)

Although the condensation and graph reconstruction processes introduce additional computation, they significantly reduce the overall training time by producing more compact and informative graph representations. This trade-off effectively speeds up model training, highlighting the practicality of integrating GC into C-C FGL for efficiency and scalability.

Hyper-parameter study. We analyze the performance under different values of τ (Node Selector) and α, β (Graph Rebuilder). These parameters control node selection and graph reconstruction during federated learning. Results are shown in Fig. 5, where Fig. 5(a) illustrates the impact of varying α and β with a fixed τ , highlighting how different feature weightings affect model performance. Fig. 5(b) shows the effect of varying τ while fixing α and β , indicating how NS module influences model accuracy.

Convergence study. Fig. 7(a) shows the average test accuracy curves during training across five random runs on Ogbn-arxiv. It can be observed that FedC4 consistently outperforms baselines in terms of average test accuracy and achieves faster convergence. It demonstrates that FedC4 not only performs well but also stabilizes quickly, highlighting its robustness and efficiency in FGL settings.

5.5 Privacy Protection

To answer Q4, we refer to the theoretical analysis in Privacy for Free[2], which proves that propagating condensed datasets is inherently secure. In Sec. 4.1, we extend this proof to graph data, demonstrating that transmitting node embeddings of condensed graphs also satisfies privacy guarantees.

Building on this theoretical foundation, we introduced varying levels of Laplace noise during the condensation process to simulate real-world privacy challenges. As shown in Fig. 7(b), results on real-world datasets show that FedC4 maintains high accuracy even under significant noise, confirming its robustness in privacy-preserving FGL and demonstrating its effectiveness in maintaining model performance while safeguarding privacy.

6 Conclusion

This paper presents the first integration of GC into the C-C paradigm in FGL, introducing three novel modules to address key challenges: fine-grained personalization, communication overhead, and privacy risks. The fine-grained personalization challenge, stemming from the need for tailored information exchange, is effectively addressed through three core modules: CM selectively shares global statistics tailored to client-specific contexts, NS identifies representative nodes for targeted knowledge transfer, and GR reconstructs personalized graph structures, enabling efficient and fine-grained collaboration. To tackle both communication overhead and privacy concerns, GC techniques are seamlessly integrated. Theoretical analysis further validates that GC significantly reduces communication costs while mitigating privacy risks. Experimental results demonstrate that FedC4 achieves superior performance and efficiency compared to competitive baselines.

Appendix

A Algorithm

A.1 Customizer Module

Algorithm 1: Customizer Module

```

1 Input: Local node embeddings  $\mathbf{H}_c = \{\mathbf{h}_i \mid i \in \mathcal{V}_c\}$ 
2 for each communication round  $t = 1, \dots, T$  do
3   for each client  $c$  do
4     Compute statistics & Normalization: Eq. 8, Eq.9, Eq.10
5     if  $t = 1$  then
6       Broadcast statistics to all clients
7     else
8       Use the clustering results from the previous round's
9       Node Selector:  $C_{\text{same}}$ 
10      Broadcast statistics to clients in  $C_{\text{same}}$ : Eq. 11
11    end if
12  end for
13 end for

```

A.2 Node Selector Module

Algorithm 2: Node Selector Module

```

1 Input: Embedding distributions  $\{\text{Dis}_c \mid c \in C\}$ , mean
2 embeddings  $\{\mu_c \mid c \in C\}$ 
3 for  $c \in C$  do
4   Compute Sliced Wasserstein Distances: Eq. 12
5   Cluster clients:
6      $C_c = \{c' \mid \text{SWD}_{c,c'} \leq \delta\}, \quad \forall c \in C$ 
7   for  $c' \in C_c$  do
8     Compute similarity: Eq. 13
9     Select nodes:
10     $S_c = \{\mathbf{h}_i \mid S(\mathbf{h}_i, \mu_{c'}) > \tau, \quad \forall c' \in C_c\}$ 
11  end for
12 end for
13 Return  $\{S_c \mid c \in C\}$ 

```

A.3 Graph Rebuilder Module

The Graph Rebuilder module reconstructs reliable graph structures for selected nodes to preserve connectivity and information flow after graph condensation. By adaptively restoring topology based on node relationships, it mitigates structural loss and helps the model capture key dependencies. Algorithm 3 outlines this process.

Algorithm 3: Graph Rebuilder Module

```

1 Input: Node features  $\{X_c\}$ , selected nodes  $\{S_c\}$ 
2 Output: Rebuilt graph  $G_c$ 
   1: for each node  $i$  in  $\{S_c\}$  do
   2:   for each neighbor  $j$  of node  $i$  do
   3:     Compute similarity: Eq.14
   4:   end for
   5: end for
6: Compute the reconstruction loss: Eq.15
7: Optimize  $Z$  to minimize  $\mathcal{L}_{rec}$ 
8: Return Rebuilt graph  $G_c = \{Z, X_c\}$ 

```

B Discussion of GR Module

To assess the GR module’s ability to restore topological information after condensation, we analyze three structural metrics: degree distribution KL divergence, graph density, and homophily.

Table 3: Topological metrics comparison among original, condensed, and reconstructed graphs.

Metric	Original	Condensed	Reconstructed
KL Divergence ↓	0.000	0.612	0.105
Graph Density ↓	0.007	0.855	0.059
Homophily ↑	0.731	0.422	0.694

As shown in Table 3, condensation leads to significant structural deviation. The KL divergence rises to 0.612, indicating notable changes in degree distribution. Graph density increases drastically from 0.007 to 0.855, suggesting an overly connected structure, while homophily drops from 0.731 to 0.422, weakening semantic consistency between connected nodes.

In contrast, the reconstructed graph significantly reduces divergence across all metrics: KL divergence decreases to 0.105, density is corrected to 0.059, and homophily rises to 0.694. These improvements show that GR effectively recovers both sparse connectivity and meaningful label relationships. Notably, although homophily is often emphasized in homophilic datasets, GR also proves effective on heterophilic graphs. For example, our accuracy on the *Roman-empire* dataset in Sec. 5.2 confirm GR’s adaptability under heterophily conditions.

C Training Configuration

C.1 Statistics of Datasets

We evaluate our proposed method on eight publicly available graph datasets widely used in federated and decentralized learning research, including citation networks (Cora, Citeseer, Ogbn-arxiv), co-purchase networks (Ogbn-products), social and community graphs

(Reddit, Flickr), physical systems (Physics), and heterophilic graphs (Roman-empire). The datasets vary significantly in terms of scale, feature dimensionality, and class distributions, thus providing a comprehensive benchmark for assessing the scalability, generalizability, and robustness of our approach. Each dataset is partitioned into training, validation, and testing subsets. Specifically, we use a 60%/20%/20% split for Cora, Citeseer, Arxiv, Physics, and Reddit, a 50%/25%/25% split for Flickr and Empire, and a 30%/10%/60% split for Products. Detailed statistics are summarized in Table 4.

Table 4: Statistics of Datasets.

Dataset	Nodes	Edges	Features	Classes
Cora	2,708	5,429	1,433	7
Citeseer	3,327	4,732	3,703	6
Arxiv	169,343	1,166,243	128	40
Physics	34,493	495,924	841	5
Flickr	89,250	899,756	500	7
Reddit	232,965	11,606,919	602	41
Products	2,449,029	61,859,140	100	47
Empire	22,678	342,334	50	18

C.2 Computation Resource

Experiments are conducted on a Linux server equipped with 2 Intel(R) Xeon(R) Gold 6240 CPUs @ 2.60GHz (36 cores per socket, 72 threads total), 251GB RAM (with approximately 216GB available), and 4 NVIDIA A100 GPUs with 40GB memory each. The software environment includes Python 3.11.10 and PyTorch 2.1.0 with CUDA 11.8. We adopt a multi-GPU training setup where each simulated client is assigned to a dedicated GPU during federated optimization. This design enables parallel execution across clients in the Client-Client (C-C) communication paradigm.

The 4 NVIDIA A100 GPUs provide substantial computational power and memory capacity, enabling efficient handling of large-scale datasets such as Ogbn-products. However, during experiments, OOM (Out of Memory) occurs when a model exceeds the 40GB memory capacity of a single A100 GPU, while OOT (Out of Time) is defined as operations exceeding 8 hours of computation time on a single A100 GPU. These constraints highlight the challenges in scaling graph learning tasks under limited hardware resources.

In contrast to existing FGL methods that frequently suffer from OOM or OOT issues on large-scale graphs, FedC4 demonstrates superior scalability and efficiency. By employing graph condensation and selective communication, FedC4 significantly reduces the memory footprint and computational load during training. This allows it to complete training on massive datasets such as Ogbn-products without encountering resource bottlenecks, thereby validating its practicality and robustness in real-world federated settings.

Acknowledgments

This work is supported by the NSFC Grants U2241211, 62427808 and U24A20255. Rong-Hua Li is the corresponding author of this paper.

GenAI Usage Disclosure

This article utilized ChatGPT-4o and Grok-3 for textual improvement and experimental code restructuring.

References

- [1] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. 2013. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203* (2013).
- [2] Tian Dong, Bo Zhao, and Lingjuan Lyu. 2022. Privacy for free: How does dataset condensation help privacy?. In *International Conference on Machine Learning*. PMLR, 5378–5396.
- [3] Liang Gao, Huazhu Fu, Li Li, Yingwen Chen, Ming Xu, and Cheng-Zhong Xu. 2022. Feddc: Federated learning with non-iid data via local drift decoupling and correction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10112–10121.
- [4] Xinyi Gao, Tong Chen, Yilong Zang, Wentao Zhang, Quoc Viet Hung Nguyen, Kai Zheng, and Hongzhi Yin. 2024. Graph condensation for inductive node representation learning. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 3056–3069.
- [5] Xinyi Gao, Tong Chen, Yilong Zang, Wentao Zhang, Quoc Viet Hung Nguyen, Kai Zheng, and Hongzhi Yin. 2024. Graph condensation for inductive node representation learning. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 3056–3069.
- [6] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems* 30 (2017).
- [7] Wenke Huang, Guancheng Wan, Mang Ye, and Bo Du. 2024. Federated graph semantic and structural learning. *arXiv preprint arXiv:2406.18937* (2024).
- [8] Wei Jin, Lingxiao Zhao, Shichang Zhang, Yozen Liu, Jiliang Tang, and Neil Shah. 2021. Graph condensation for graph neural networks. *arXiv preprint arXiv:2110.07580* (2021).
- [9] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. *arXiv:1609.02907 [cs.LG]* <https://arxiv.org/abs/1609.02907>
- [10] Xunkai Li, Zhengyu Wu, Wentao Zhang, Yinlin Zhu, Rong-Hua Li, and Guoren Wang. 2024. Fedgta: Topology-aware averaging for federated graph learning. *arXiv preprint arXiv:2401.11755* (2024).
- [11] Andreas Loukas. 2019. Graph reduction with spectral and cut guarantees. *Journal of Machine Learning Research* 20, 116 (2019), 1–42.
- [12] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*. PMLR, 1273–1282.
- [13] Oleg Platonov, Denis Kuznedelev, Michael Diskin, Artem Babenko, and Liudmila Prokhorenkova. 2023. A critical look at the evaluation of GNNs under heterophily: Are we really making progress? *arXiv preprint arXiv:2302.11640* (2023).
- [14] Yuxin Qiu. 2023. Default Risk Assessment of Internet Financial Enterprises Based on Graph Neural Network. In *2023 IEEE 6th Information Technology, Networking, Electronic and Automation Control Conference*, Vol. 6. IEEE, 592–596.
- [15] Zongshuai Qu, Tao Yao, Xinghui Liu, and Gang Wang. 2023. A Graph Convolutional Network Based on Univariate Neurodegeneration Biomarker for Alzheimer’s Disease Diagnosis. *IEEE Journal of Translational Engineering in Health and Medicine* (2023).
- [16] Maria Rigaki and Sebastian Garcia. 2023. A survey of privacy attacks in machine learning. *Comput. Surveys* 56, 4 (2023), 1–34.
- [17] Herbert Robbins and Sutton Monro. 1951. A stochastic approximation method. *The annals of mathematical statistics* (1951), 400–407.
- [18] Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann. 2018. Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868* (2018).
- [19] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, Yoshua Bengio, et al. 2017. Graph attention networks. *stat* 1050, 20 (2017), 10–48550.
- [20] Max Welling. 2009. Herding dynamical weights to learn. In *Proceedings of the 26th annual international conference on machine learning*. 1121–1128.
- [21] Shiwen Wu, Fei Sun, Wentao Zhang, Xu Xie, and Bin Cui. 2022. Graph neural networks in recommender systems: a survey. *Comput. Surveys* 55, 5 (2022), 1–37.
- [22] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* 32, 1 (2020), 4–24.
- [23] Han Xie, Jing Ma, Li Xiong, and Carl Yang. 2021. Federated graph classification over non-iid graphs. *Advances in neural information processing systems* 34 (2021), 18839–18852.
- [24] Han Xie, Li Xiong, and Carl Yang. 2023. Federated node classification over graphs with latent link-type heterogeneity. In *Proceedings of the ACM Web Conference 2023*. 556–566.
- [25] Bo Yan. 2024. Federated Graph Condensation with Information Bottleneck Principles. *arXiv preprint arXiv:2405.03911* (2024).
- [26] Beining Yang, Kai Wang, Qingyun Sun, Cheng Ji, Xingcheng Fu, Hao Tang, Yang You, and Jianxin Li. 2024. Does graph distillation see like vision dataset counterpart? *Advances in Neural Information Processing Systems* 36 (2024).
- [27] Zhilin Yang, William Cohen, and Ruslan Salakhudinov. 2016. Revisiting semi-supervised learning with graph embeddings. In *International conference on machine learning*. PMLR, 40–48.
- [28] Zhilin Yang, William Cohen, and Ruslan Salakhudinov. 2016. Revisiting semi-supervised learning with graph embeddings. In *International conference on machine learning*. PMLR, 40–48.
- [29] Yuhang Yao, Weizhao Jin, Srivatsan Ravi, and Carlee Joe-Wong. 2024. FedGCN: Convergence-communication tradeoffs in federated training of graph convolutional networks. *Advances in neural information processing systems* 36 (2024).
- [30] Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. 2019. Graphsaint: Graph sampling based inductive learning method. *arXiv preprint arXiv:1907.04931* (2019).
- [31] Ke Zhang, Lichao Sun, Bolin Ding, Siu Ming Yiu, and Carl Yang. 2024. Deep Efficient Private Neighbor Generation for Subgraph Federated Learning. In *Proceedings of the 2024 SIAM International Conference on Data Mining (SDM)*. SIAM, 806–814.
- [32] Ke Zhang, Carl Yang, Xiaoxiao Li, Lichao Sun, and Siu Ming Yiu. 2021. Sub-graph federated learning with missing neighbor generation. *Advances in Neural Information Processing Systems* 34 (2021), 6671–6682.
- [33] Xin Zheng, Miao Zhang, Chunyang Chen, Quoc Viet Hung Nguyen, Xingquan Zhu, and Shirui Pan. 2024. Structure-free graph condensation: From large-scale graphs to condensed graph-free data. *Advances in Neural Information Processing Systems* 36 (2024).
- [34] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2020. Graph neural networks: A review of methods and applications. *AI open* 1 (2020), 57–81.
- [35] Yinlin Zhu, Xunkai Li, Zhengyu Wu, Di Wu, Miao Hu, and Rong-Hua Li. 2024. FedTAD: Topology-aware Data-free Knowledge Distillation for Subgraph Federated Learning. *arXiv preprint arXiv:2404.14061* (2024).