

Towards Federated Clustering: A Client-wise Private Graph Aggregation Framework

Guanxiong He, Zheng Wang*, Jie Wang, Liaoyuan Tang, Rong Wang, Feiping Nie

School of Artificial Intelligence, Optics and Electronics (iOPEN), Northwestern Polytechnical University
Xi'an Shaanxi, 710072, P.R.China

heguanx@mail.nwpu.edu.cn, zhengwangml@gmail.com, jiewang.dl@mail.nwpu.edu.cn, tangly@mail.nwpu.edu.cn,
wangrong07@tsinghua.org.cn, feipingnie@gmail.com

Abstract

Federated clustering addresses the critical challenge of extracting patterns from decentralized, unlabeled data. However, it is hampered by the flaw that current approaches are forced to accept a compromise between performance and privacy: *transmitting embedding representations risks sensitive data leakage, while sharing only abstract cluster prototypes leads to diminished model accuracy*. To resolve this dilemma, we propose Structural Privacy-Preserving Federated Graph Clustering (SPP-FGC), a novel algorithm that innovatively leverages local structural graphs as the primary medium for privacy-preserving knowledge sharing, thus moving beyond the limitations of conventional techniques. Our framework operates on a clear client-server logic; on the client-side, each participant constructs a private structural graph that captures intrinsic data relationships, which the server then securely aggregates and aligns to form a comprehensive global graph from which a unified clustering structure is derived. The framework offers two distinct modes to suit different needs. SPP-FGC is designed as an efficient one-shot method that completes its task in a single communication round, ideal for rapid analysis. For more complex, unstructured data like images, SPP-FGC+ employs an iterative process where clients and the server collaboratively refine feature representations to achieve superior downstream performance. Extensive experiments demonstrate that our framework achieves state-of-the-art performance, improving clustering accuracy by up to 10% (NMI) over federated baselines while maintaining provable privacy guarantees.

Extended version — <https://arxiv.org/abs/2511.10915>

Introduction

Federated Learning (FL) offers a promising approach to balancing data privacy with utilizing large-scale datasets for machine learning [McMahan et al. 2017, 2016, Yang et al. 2019]. FL methods are generally divided into three transmission strategies: *Model Averaging*, *Embedding Sharing*, and *Prototype Aggregation*, each offering different trade-offs among performance, communication cost, and privacy. Model Averaging preserves privacy effectively by aggregating parameters from locally trained models but incurs high communication overhead due to the complexity of model

structures [McMahan et al. 2017, Reddi et al. 2021]. Embedding Sharing reduces communication by transmitting intermediate representations but risks leaking sensitive information [Tan et al. 2022b, Wang et al. 2024]. Prototype Aggregation further minimizes communication and privacy risks by sharing only representative prototypes, although this limits the ability to fully capture complex data structures [Tan et al. 2022a, Huang et al. 2023]. Thus, choosing the appropriate strategy is essential to effectively balancing privacy protection, communication efficiency, and performance.

Most FL research has emphasized supervised tasks, presuming access to labeled data, which is not always realistic [Wu et al. 2023, Hu et al. 2024]. Clients may lack labels for their data, making supervised FL methods ineffective. This increased interest in unsupervised learning, particularly clustering, which groups data based on inherent similarities without needing labels [Dennis, Li, and Smith 2021, Xie et al. 2023]. Federated Clustering (FC) [Pan et al. 2022, Nardi, Valerio, and Passarella 2024] addresses this need by enabling decentralized clustering while preserving data privacy, providing significant advantages over traditional centralized approaches [Basagni 1999, Chen et al. 2016]. FC is especially valuable in privacy-sensitive sectors like healthcare and finance [Li et al. 2022], with practical applications including Human Activity Recognition [Presotto, Civitarese, and Bettini 2022] and Electricity Consumption Pattern Extraction [Wang et al. 2022]. By facilitating collaboration among isolated datasets, FC uncovers hidden patterns and relationships, becoming an important tool for unsupervised learning in industries where data sharing is restricted.

As shown in Fig 1, current FC methods fall into three main paradigms. In Model Averaging, clients share locally trained model weights, achieving strong performance but incurring high communication costs [Mashhadi, Sterner, and Murray 2021, Li et al. 2023]. The Embedding Sharing paradigm has clients send feature embeddings to the server, which allows knowledge sharing but risks leaking private information [Chen et al. 2023, Qiao, Ding, and Fan 2024]. In Prototype Aggregation, clients share only local cluster centers, a method that reduces communication and improves privacy but can compromise clustering quality and consistency [Dennis, Li, and Smith 2021, Pedrycz 2021]. Each paradigm offers a different trade-off between performance, privacy, and complexity. We aim to strike a better balance among these factors

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

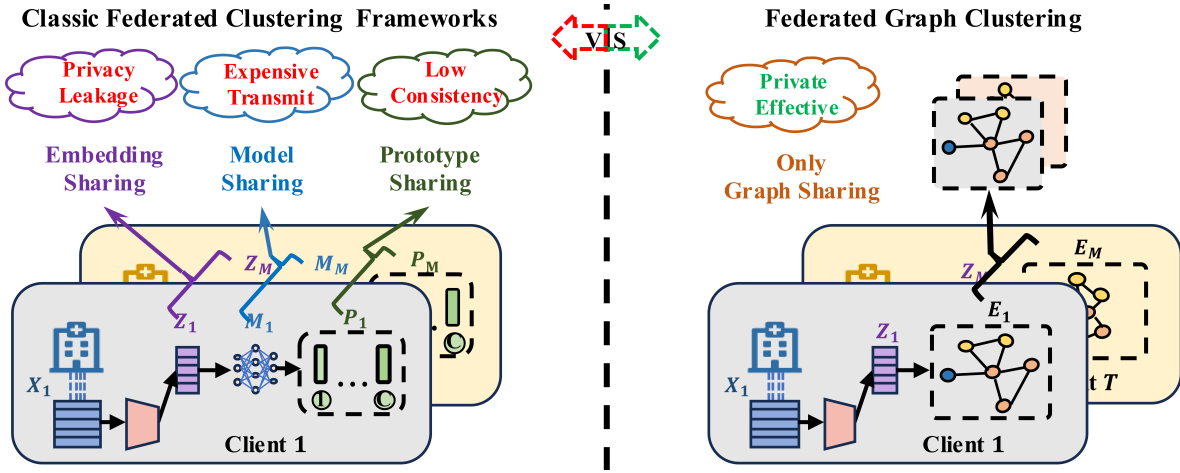


Figure 1: Classic FC paradigms and the proposed Federated Graph Clustering framework. The left shows the paradigms of Model Averaging, Embedding Sharing, and Prototype Aggregation, while the right highlights our graph-based approach.

while addressing the limitations of existing methods.

To achieve this balance, we propose a new method, SPP-FGC. Clients build structural graphs from their local data and send them, rather than raw data or embeddings, to the server, preserving privacy. During communication, only similarity graphs and prototype distributions are shared, and these are further protected using Differential Privacy (DP) techniques [Dwork 2006, Zhao and Chen 2022] to reduce leakage risk. The server aggregates these similarity graphs and uses the Constrained Laplacian Rank method [Nie et al. 2016] to refine the global structure. It then identifies clusters and sends global assignments back to clients. To better handle complex data, we extend the method to SPP-FGC+, which includes feature learning and iterative refinement. In this version, clients use Variational Autoencoders (VAE) [Doersch 2016, Xie, Girshick, and Farhadi 2016] with Deep Embedded Clustering (DEC) loss for self-supervised feature extraction. As the global structure updates through graph aggregation, new cluster prototypes are sent to guide further learning. This iterative process improves feature representation and clustering accuracy over time, making the method more adaptable to complex and dynamic datasets. The key contributions of our proposed methods are highlighted from the following three distinct perspectives.

- We introduce a novel method that leverages local structural graphs as the knowledge-sharing medium, effectively overcoming the privacy risks of sharing embeddings and the performance loss of sharing only prototypes.
- Our framework ensures security by abstracting client data into local private structures and employing differential privacy to facilitate their secure aggregation on the server.
- Our method is highly adaptable to diverse scenarios, offering an efficient one-shot mode (SPP-FGC) for simple datasets and an iterative version (SPP-FGC+) that achieves superior performance on complex, unstructured data through collaborative learning.

Related Works

Federated Clustering FC follows three main FL paradigms. Under Model Averaging, clients train local clustering models and send their parameters to a central server to build a global model [McMahan et al. 2017]. This approach protects privacy and achieves strong performance but incurs high communication costs. For example, F-DEC [Mashhadi, Sterner, and Murray 2021] integrates deep embedding clustering into Fed-Avg by sharing weighted local models and updating via gradients; FednamN [Hasan et al. 2025] adds Adam and Nadam optimizers to speed convergence; and PPFC-GAN [Yan et al. 2023] uses GANs [Goodfellow et al. 2020] to generate private data samples for global clustering. In Embedding Sharing, clients send feature embeddings instead of full models—FedDMVC [Chen et al. 2023] employs contrastive learning for robust embeddings, and FedCA [Zhang et al. 2023] uses contrastive averaging with a dictionary and alignment strategy. Finally, Prototype Aggregation has clients share cluster centers as prototypes: k-FED [Dennis, Li, and Smith 2021] extends k-means by aggregating centroids from each client, and FFCM [Pedrycz 2021] adds a fuzzy extension for greater robustness. Different from the above methods, the graph-based framework of SPP-FGC exchanges only private graphs, enhanced with differential privacy, to capture deeper relationships and deliver more consistent global clusters without excessive overhead.

Differential Privacy As a sophisticated mathematical framework designed to protect individual privacy within large datasets while still allowing meaningful aggregate analysis, DP often takes Gaussian and Laplace mechanisms into practice [Dwork 2006, Abadi et al. 2016]. It is widely used in FL to protect individual users’ data by adding noise to model updates, ensuring that sensitive information cannot be inferred from the aggregated results [Wei et al. 2020, Hu et al. 2020, Triastcyn and Faltings 2019].

Problem Formulation

FL is a decentralized machine learning mechanism that allows multiple clients to collaboratively train a global model without sharing their local data. It addresses privacy and data ownership concerns by keeping data on the client side and only exchanging non-sensitive information. In the scenario of M clients, each one has their own dataset D_1, \dots, D_M where $D_i = \{(v_j, y_j)\}_{j=1}^{N_i}$, with $v \in \mathbb{R}^d$ representing d -dimensional features and $y \in \{1, 2, \dots, C\}$ representing labels across C classes. The loss function $\ell(v, y, \theta)$ measures the loss for a sample (v, y) with model parameters θ . The combined dataset of all clients is $\mathbf{D} = \bigcup_{i=1}^M D_i$ and the goal of FL is to learn a model θ^* that minimizes:

$$\arg \min_{\theta} \sum_{i=1}^M \frac{|D_i|}{|\mathbf{D}|} \mathbb{E}_{(v,y) \sim D_i} \ell(v, y, \theta). \quad (1)$$

FC follows a similar objective on unlabeled data $\bar{D}_i = \{v_j\}_{j=1}^{N_i}$. It aims to optimize data clustering within each client by utilizing information from all clients, thereby enhancing clustering quality and consistency without compromising privacy. The optimization goal in FC is to minimize the clustering loss ℓ_c across all clients:

$$\arg \min_{\theta_1, \dots, \theta_M} \sum_{i=1}^M \frac{|\bar{D}_i|}{|\mathbf{D}|} \mathbb{E}_{v \sim \bar{D}_i} \ell_c(v, \theta). \quad (2)$$

The main challenge in FC is achieving high-quality, collaborative clustering while maintaining strict privacy standards. Current methods often have to sacrifice either performance or privacy to address this balance. Therefore, we aim to find an optimal balance by transmitting similarity graphs instead of raw data, ensuring both effective clustering and robust privacy protection.

Methodology

To address the challenge of achieving high-quality federated clustering while ensuring strong privacy protection, we introduce Structural Privacy-Preserving Federated Graph Clustering (SPP-FGC). This graph-based approach enhances data privacy by utilizing similarity graphs instead of raw data features, along with the DP mechanism. Additionally, we extend this framework with the SPP-FGC+ variant, which integrates deep feature extraction and iterative optimization. This enhancement enables adaptive and powerful feature discovery, making SPP-FGC+ suitable for complex, real-world data.

In our SPP-FGC method, we draw from traditional graph clustering, which typically involves two steps: building a similarity graph from the data and using it for clustering. This approach fits well within the FL framework, where data stays decentralized on clients, and only the similarity graphs are shared with the central server. The goal of SPP-FGC is to combine the clients' ability to capture local data relationships with the server's role in aggregating this information for better global clustering. By integrating data from all clients, SPP-FGC strengthens cluster connections and corrects local errors. Our method tackles three key challenges: **accurately**

mining structure on each client, **securely transmitting** information to the server, and **effectively merging** data from all clients to enhance overall clustering.

Private Structural Graph Construction

To uncover the underlying structure in each client's data, we construct a private graph $G^p = (V^p, E^p)$ that captures the neighborhood relationships of the samples. In this graph, similar samples are placed closer together, and dissimilar ones farther apart. This approach enhances the representation of data relationships, leading to better clustering results. To create a more informative graph, we define the following optimization objective:

$$\min_{E^p} \sum_{i,j=1}^N (\|v_i^p - v_j^p\|_2^2 E_{ij}^p + \gamma (E_{ij}^p)^2).$$

$$\text{s.t. } \forall i, (E_i^p)^T \mathbf{1} = 1, 0 \leq E_{ij}^p \leq 1, \text{rank}(L_{E^p}) = N - C \quad (3)$$

Here, E^p represents the relationship matrix for specified clients, each element E_{ij}^p measures the closeness of the relationship between the i -th and j -th private samples vectors v_i^p, v_j^p . Meanwhile, E_i^p represents the relationships between a given sample and all others. Additionally, L_{E^p} is the Laplacian matrix of E^p , defined as:

$$L_{E^p} = D_{E^p} - \frac{(E^p)^T + E^p}{2}, \quad (4)$$

where D_{E^p} denotes the degree matrix of E^p . Here are two constraints: 1) The rank constraint on L_{E^p} forces the graph to have exactly C connected components. The objective can be understood as obtaining the optimal private structural matrix E^p by minimizing the combined effect of the Euclidean distance between feature vectors. 2) The L2 regularization term $\gamma (E_{ij}^p)^2$ is taken to avoid the trivial solution where each data point connects only to its single nearest neighbor. It encourages a smoother distribution of connection weights across multiple neighbors.

To make the rank constraint in Eq.3 more tractable, we reformulate the problem by leveraging Ky Fan's Theorem [Fan 1949]. This involves introducing a continuous cluster indicator matrix $F \in \mathbb{R}^{N \times C}$. This matrix serves as a relaxed representation of cluster assignments, where its columns are the eigenvectors corresponding to the C smallest eigenvalues of L_{E^p} . The problem can be reformulated as:

$$\min_{E^p, F} \sum_{i,j=1}^N (\|v_i^p - v_j^p\|_2^2 E_{ij}^p + \gamma (E_{ij}^p)^2) + 2\lambda \text{Tr}(F^T L_{E^p} F). \quad (5)$$

$$\text{s.t. } \forall i, (E_i^p)^T \mathbf{1} = 1, 0 \leq E_{ij}^p \leq 1, F^T F = I$$

To solve the simplified problem, the parameters F and E^p can be optimized alternately. First, fix E^p , the optimal solution of the indicating matrix F is formed by the C eigenvectors of L_{E^p} corresponding to the C smallest eigenvalues.

Then, with F fixed, using the Lagrangian function and Karush-Kuhn-Tucker (KKT) conditions, the optimal solution for E_i^p is given by:

$$E_{ij}^p = -\frac{d_{ij}}{2\gamma_i} + \eta, \quad (6)$$

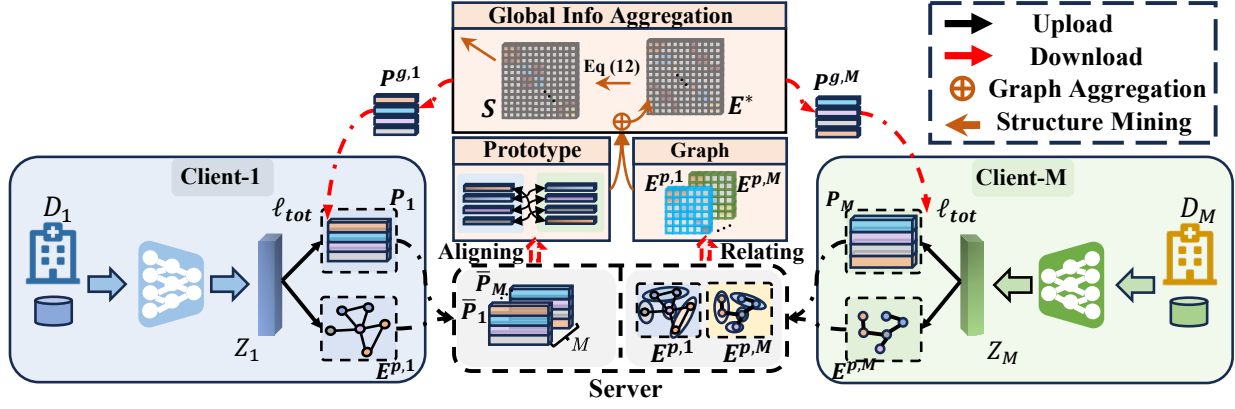


Figure 2: Overview of SPP-FGC+. Clients generate private structural graphs from learned features and upload them to the server. The server fuses these inputs into a global graph, derives new cluster prototypes, and sends them back as feedback.

where $d_{ij} = d_{ij}^x + d_{ij}^f$, with $d_{ij}^x = \|v_i^p - v_j^p\|_2^2$ and $d_{ij}^f = \|f_i - f_j\|_2^2$. Here, f_i denotes the row vector of F , also the indicating vector. The detailed process of the private structural graph optimization is provided in Appendix A.

For a specified client k , the private structural matrix $E^{p,k}$ is constructed through the above optimization. To capture the underlying data distribution, especially with heterogeneous data across clients, prototypes are generated for each cluster using a Gaussian Mixture Model (GMM). Specifically, for each cluster c within client k , its prototype is defined as $\mathcal{N}(\mu_{k,c}, \Sigma_{k,c})$, where $\mu_{k,c}$ is the mean centroid of cluster c , and $\Sigma_{k,c}$ is the covariance matrix that captures the variance and correlations within the cluster. These prototypes summarize the statistical properties of each cluster, aiding in more effective aggregation at the central server.

Private Information Transmission

To ensure data privacy and secure the transmission of sensitive information in the SPP-FGC framework, we use DP techniques, specifically the Laplace mechanism, to add calibrated noise to the transmitted data. This approach ensures that the exchange of the prototype set $P_k = \{(\mu_{k,c}, \Sigma_{k,c})\}_{c=1}^C$ adheres to ϵ -DP, providing strong privacy protection against potential data leaks.

For prototypes on client k , the sensitivities Δ^k are:

$$\begin{aligned} \Delta_{\mu}^k &= \max_{k,c} \|\mu_{k,c} - \mu'_{k,c}\|_1, \\ \Delta_{\Sigma}^k &= \max_{k,c} \|\Sigma_{k,c} - \Sigma'_{k,c}\|_1, \end{aligned} \quad (7)$$

where $\mu'_{k,c}$ and $\Sigma'_{k,c}$ are the mean and covariance matrices of the modified dataset D' . Laplace noise is then added to the parameters to ensure differential privacy:

$$\begin{aligned} \bar{\mu}_{i,c} &= \mu_{i,c} + \text{Lap}\left(\frac{\Delta_{\mu}}{\epsilon}\right), \\ \bar{\Sigma}_{i,c} &= \Sigma_{i,c} + \text{Lap}\left(\frac{\Delta_{\Sigma}}{\epsilon}\right). \end{aligned} \quad (8)$$

Here, $\text{Lap}(\cdot)$ represents the Laplace distribution with scale parameter $b = \Delta/\epsilon$, where Δ is the sensitivity and ϵ is the

privacy budget. Finally, the k -th private prototypes $\bar{P}_k = \{\bar{\mu}_{k,c}, \bar{\Sigma}_{k,c}\}_{c=1}^C$ are securely transmitted to the central server. Details about privacy sensitivity can be seen in Appendix C.

Global Information Aggregation

Upon receiving the Laplace-encrypted local prototypes \bar{P}_k and private structural matrix $E^{p,k}$ from all M clients, the central server faces the crucial task of mining consistent clustering structures from the aggregated information. This step is essential to ensure that the global clustering reflects the collective insights derived from all clients' data while preserving data privacy.

To achieve this, the central server constructs a global structural matrix E^* , which encapsulates both the local structural information and the inter-client similarities. The construction process begins by organizing E^* as a block matrix, where each diagonal block corresponds to a client's private structural matrix $E^{p,k}$, thus retaining the localized clustering information of each client. Formally, E^* is represented as:

$$E^* = \begin{bmatrix} E^{p,1} & E_2^{g,1} & \cdots & E_M^{g,1} \\ E_1^{g,2} & E^{p,2} & \cdots & E_M^{g,2} \\ \vdots & \vdots & \ddots & \vdots \\ E_1^{g,M} & E_2^{g,M} & \cdots & E^{p,M} \end{bmatrix}. \quad (9)$$

The off-diagonal blocks, $E_j^{g,i}$, represent the inter-client relation between clients i and j . These relationships are computed by measuring the affinity between their respective cluster prototypes. Specifically, the similarity between any sample from client i and any sample from client j is determined by the Kullback-Leibler (KL) divergence between their assigned prototypes. Let $\bar{P}_i^{c(m)}$ be the prototype for the cluster assigned to sample m in client i , and $\bar{P}_j^{c(n)}$ be the prototype for the cluster assigned to sample n in client j . The KL divergence $\text{KL}_{i,j}(\bar{P}_i^{c(m)}, \bar{P}_j^{c(n)})$ quantifies the dissimilarity between their underlying distributions.

To convert this divergence into a similarity score, we apply an exponential kernel, ensuring that greater divergence results in a lower similarity value. The inter-client relation

Algorithm 1: SPP-FGC

```

1: Input:  $M$  Clients,  $C$  Clusters, Privacy budget  $\epsilon$ 
2:
3: Client-side Procedures:
4: for each client  $k = 1$  to  $M$  do
5:   Construct  $E^{p,k} \leftarrow$  Eq. 5
6:   Compute  $P_k \leftarrow$  GMM( $D_k$ )
7:   Apply  $\bar{P}_k \leftarrow$  Lap( $\epsilon, P_k$ )
8:   Upload  $E^{p,k}$  and  $\bar{P}_k$  to Server
9: end for
10:
11: Server-side Procedures:
12: Compute  $\{E_j^{g,i}\}_{i,j=1}^M \leftarrow$  Eq. 10
13: Construct  $E^* \leftarrow$  Aggregate( $E^g, E^p$ )
14: Global aggregation  $S \leftarrow$  Eq. 12
15: Cluster assignments  $\mathbf{A} \leftarrow$  K-Means( $S$ )
16:
17: Output: Global cluster assignments  $\mathbf{A}$ 

```

graph $E_j^{g,i}$ is thus an $N_i \times N_j$ matrix where each element is computed as:

$$E_j^{g,i} = [\exp(-\text{KL}_{i,j}(\bar{P}_i, \bar{P}_j))]. \quad (10)$$

This process effectively captures the pairwise alignment of cluster distributions across the federation. By assembling the private structural matrices $E^{p,k}$ and the inter-client relation graphs $E_j^{g,i}$ into the global matrix E^* , the server fuses localized knowledge with global relational insights, creating a robust foundation for identifying consistent, overarching cluster structures.

To further learn the global matrix E^* and extract a decision similarity matrix S with clear clustering structures, we impose a Laplacian rank restriction to enhance the structural features in S . The desired matrix has a block diagonal structure, where each block corresponds to a cluster. This is achieved by solving the following optimization problem:

$$\begin{aligned} \min_S \|S - E^*\|_F^2. \\ \text{s.t. } \text{rank}(L_S) = N - C \end{aligned} \quad (11)$$

We apply the Lagrangian function and Ky Fan's Theorem to transform problem into the following optimization problem:

$$\begin{aligned} \min_S \|S - E^*\|_F^2 + 2\lambda \text{Tr}(F^T L_S F). \\ \text{s.t. } \sum_j S_{ij} = 1, S_{ij} \geq 0, F \in \mathbb{R}^{N \times C}, F^T F = I \end{aligned} \quad (12)$$

To solve this problem, we use the alternating optimization approach. Specifically, we iteratively optimize F and S by fixing one and solving for the other.

When S is fixed, the optimal solution to the optimization problem is obtained by selecting the C eigenvectors of L_S corresponding to the C smallest eigenvalues.

Using the KKT conditions, the optimal solution for S , when F is fixed, is given by:

$$S_{ij} = -\frac{d_{ij}}{2\gamma} + \eta. \quad (13)$$

Algorithm 2: SPP-FGC+

```

1: Input:  $T$  rounds,  $M$  clients,  $C$  clusters, privacy budget  $\epsilon$ 
2: Initialize VAE-DEC, obtain global prototypes  $P_g^0$ 
3:
4: for  $t = 1$  to  $T$  do
5:   Client-side Procedures:
6:   for each client  $k = 1$  to  $M$  do
7:     Train VAE-DEC with  $P_g^{t-1}$ 
8:     Construct  $E^{p,k} \leftarrow$  Eq. 5 on  $Z_k$ 
9:     Extract  $P_k \leftarrow$  DEC
10:    Apply  $\bar{P}_k \leftarrow$  Lap( $\epsilon, P_k$ )
11:    Upload  $E^{p,k}$  and  $\bar{P}_k$  to server
12:  end for
13:
14:  Server-side Procedures:
15:  Compute  $\{E_j^{g,i}\}_{i,j=1}^M \leftarrow$  Eq. 10
16:  Construct  $E^* \leftarrow$  Aggregate( $E^g, E^p$ )
17:  Global aggregation  $S \leftarrow$  Eq. 12
18:  Update  $P_g^t \leftarrow$  K-Means( $S$ )
19: end for
20:
21: Output: Global cluster assignments  $\mathbf{A}$  from  $P_g^T$ 

```

In summary, the primary framework of SPP-FGC is outlined in Algorithm 1. The detailed proof can be seen in Appendix B.

SPP-FGC with Self-Supervised Feature Learning

High-dimensional data often contains intricate structures that can impede effective clustering [Xie et al. 2023]. To enhance clustering performance, we integrate deep representation techniques that extract low-dimensional embeddings from high-dimensional data, thereby simplifying the data structure and improving cluster quality.

We propose the variant algorithm SPP-FGC+, which begins by using a VAE combined with a DEC loss to create effective embeddings for clustering. Specifically, the VAE is first pre-trained to learn robust initial embeddings and then fine-tuned with the DEC loss to optimize the latent space for clustering. The VAE encodes input features X into latent embeddings Z , which are used to construct private structural graphs on each client. The total loss function \mathcal{L}_{total} of the VAE-DEC framework is defined as:

$$\mathcal{L}_{total} = \mathcal{L}_{vae} + \lambda \cdot \mathcal{L}_{dec}(P), \quad (14)$$

where \mathcal{L}_{vae} is the standard VAE reconstruction loss, and $\mathcal{L}_{dec}(P)$ is the clustering loss.

Each client builds a private structural matrix $E^{p,k}$ from the learned embeddings Z and adds noise through DP to protect data during transmission. Protected prototypes are then sent to the central server, which combines them into a global similarity matrix and discover global structure as SPP-FGC. Finally, the global prototypes P_g are sent back to all clients. These prototypes act as a common target for the next round of local learning. By optimizing its local VAE model to minimize the distance between its data embeddings and these common prototypes, each client is implicitly guided

Dataset	M	Federated Methods					Our Methods	
		K-FED	FFCM	DSC	Fed-SC	PPFC-GAN	SPP-FGC	SPP-FGC+
Moon	A	0.743±0.010	0.742±0.010	0.687±0.127	0.998±0.002	0.761±0.039	<u>0.987±0.039</u>	0.999±0.001
	N	0.178±0.015	0.177±0.016	0.265±0.025	0.930±0.118	0.215±0.062	<u>0.900±0.013</u>	0.990±0.046
RING	A	0.371±0.004	0.370±0.003	0.929±0.006	0.956±0.046	0.376±0.008	<u>0.971±0.018</u>	0.990±0.011
	N	0.433±0.002	0.433±0.002	0.858±0.049	0.830±0.000	0.432±0.003	<u>0.936±0.009</u>	0.964±0.024
Letters	A	0.149±0.011	0.250±0.003	0.281±0.000	0.319±0.019	0.260±0.009	<u>0.623±0.023</u>	0.644±0.042
	N	0.221±0.010	0.335±0.003	0.377±0.002	0.432±0.010	0.374±0.009	<u>0.563±0.015</u>	0.614±0.013
MNIST	A	0.559±0.559	0.490±0.004	0.602±0.001	0.612±0.027	0.518±0.013	<u>0.631±0.017</u>	0.650±0.055
	N	0.504±0.002	0.428±0.004	0.577±0.009	0.657±0.007	0.497±0.002	<u>0.719±0.011</u>	0.740±0.011
Fashion	A	0.518±0.005	0.601±0.001	0.503±0.006	0.600±0.013	0.568±0.002	<u>0.686±0.028</u>	0.714±0.063
	N	0.491±0.003	0.592±0.002	0.508±0.002	0.602±0.012	0.575±0.002	<u>0.670±0.036</u>	0.709±0.027
CIFAR	A	0.725±0.000	0.656±0.003	0.610±0.001	0.705±0.016	0.726±0.005	<u>0.765±0.000</u>	0.803±0.022
	N	0.610±0.001	0.571±0.001	0.609±0.004	0.627±0.017	0.611±0.005	<u>0.646±0.001</u>	0.683±0.031
STL-10	A	0.850±0.006	0.879±0.002	0.832±0.001	0.878±0.013	0.881±0.003	<u>0.905±0.021</u>	0.923±0.014
	N	0.775±0.002	0.775±0.003	0.726±0.044	0.779±0.004	0.786±0.001	<u>0.810±0.033</u>	0.865±0.015
M-Image	A	0.733±0.010	0.438±0.010	0.399±0.028	0.369±0.002	0.448±0.009	<u>0.839±0.249</u>	0.861±0.027
	N	0.766±0.013	0.652±0.006	0.512±0.008	0.508±0.007	0.706±0.001	<u>0.810±0.037</u>	0.850±0.017

Table 1: Clustering Performance Comparison on Real-World Datasets. The performance of SPP-FGC and SPP-FGC+ is evaluated against federated baselines using Clustering Accuracy (A) and Normalized Mutual Information (NMI). Each cell shows the mean \pm standard deviation. Bold values denote the highest score in each row, while underlined values are the second-best.

to align its private feature space with the global consensus established by the server.

This approach leverages the strengths of both self-supervised learning and FC to provide a robust solution for privacy-preserving, efficient clustering of complex data. The enhanced algorithm is summarized in Algorithm 2.

Experimental Results

We evaluated the SPP-FGC and SPP-FGC+ algorithms on both synthetic and real-world datasets with varying complexity. The synthetic datasets included Moon (two interleaving half circles) and Ring (a high-dimensional ring). Real-world datasets comprised simple tabular data (Iris, Breast Cancer, Bank, COLI20, USPS, and Letters) and more complex image data (MNIST, Fashion-MNIST, CIFAR-10, STL-10, and Mini-Imagenet), as summarized in Table 4 of Appendix F. For tabular data, all algorithms processed the raw input directly, while for images, SPP-FGC+ used a four-layer Variational Autoencoder (VAE), and other algorithms used 512-dimensional embeddings extracted via pretrained ResNet-18. All experiments were run in a CUDA 12.6 environment on an RTX 3090 GPU.

Clustering Results Comparison

We compared the performance of our SPP-FGC and SPP-FGC+ algorithms with other FC methods, including k-FED, FFCM, DSC, Fed-SC, and PPFC-GAN, to demonstrate their effectiveness in preserving data privacy and handling complex data scenarios. The experiments, conducted in a federated setting, revealed that our methods outperformed the baseline algorithms, in terms of Clustering Accuracy (Acc) and Normalized Mutual Information (NMI). In detail, clustering Acc is calculated after aligning predicted cluster labels with ground truth labels using the Hungarian algorithm to

find the optimal one-to-one mapping. Notably, SPP-FGC+ achieved the highest NMI across most datasets, highlighting its superior ability to enhance clustering quality. Our methods excelled across eight selected datasets (with details in Appendix F), significantly outperforming other federated techniques. While SPP-FGC showed strong improvements on MNIST and Fashion-MNIST, its performance on CIFAR-10 and STL-10 suggested the potential for optimization in feature learning. SPP-FGC+’s adaptive feature embedding framework, which refines feature representations progressively, enabled more accurate clustering of high-dimensional and complex image datasets. In addition, the process of aggregation is shown in Figure.5 of Appendix F.I to visualize the effectiveness of clustering methods.

Data Heterogeneous Experiment

To evaluate the robustness of our framework against the critical challenge of data heterogeneity, we conducted an experiment on the MNIST dataset where the class imbalance across clients was systematically increased. As detailed in Table 2, SPP-FGC consistently outperforms baseline methods like Fed-Kmeans and Fed-SC across all levels of imbalance. While the performance of all algorithms naturally degrades as heterogeneity intensifies, the decline for SPP-FGC is far less severe, as it maintains significantly higher NMI and ACC scores. This demonstrates a remarkable resilience to the non-IID data distributions typical of practical federated applications, a conclusion further supported by visualizations in Appendix F.II.

Convergence Experiment

To validate the effectiveness of SPP-FGC+, we conducted two experiments on the MNIST dataset (shown in Fig 3). The first tested performance scaling with the number of clients,

Het	K-FED		Fed-SC		SPP-FGC	
	NMI	ACC	NMI	ACC	NMI	ACC
0.2	0.4811	0.5474	0.6351	0.6113	0.6788	0.6224
0.4	0.4639	0.4707	0.5858	0.5109	0.5956	0.5403
0.6	0.4495	0.4515	0.4751	0.4171	0.5476	0.5430
0.8	0.4108	0.4415	0.2815	0.2203	0.5127	0.5073
0.95	0.3864	0.3689	0.0279	0.1557	0.4849	0.4376

Table 2: Robustness to Heterogeneity on the MNIST Dataset. The ‘‘Het’’ column indicates the heterogeneity ratio, where a higher value signifies greater class imbalance across clients.




Figure 3: Clustering performance measured by ACC and NMI. The Left: changes over the course of iterations, The Right: changes as the number of clients increases.

showing improvement up to 10 clients, after which performance stabilized, demonstrating the method’s scalability. The second examined convergence over training iterations, with performance steadily improving and plateauing, highlighting the benefits of iterative feature learning. These results confirm SPP-FGC+’s robustness and adaptability in large-scale federated environments and complex data scenarios. Detailed experiment setting and analysis are shown in Appendix F.III.

Ablation Experiment

We conducted an ablation study on the SPP-FGC algorithm, testing five configurations: the original version, without Laplace Noise for privacy, without Private Structural Graph (PSG), without Global Structural Graph (GSG), and without both PSG and GSG. The results (Table 3) show that while removing Laplace Noise slightly improved clustering performance, it compromised privacy, which is crucial for FL. Omitting either PSG or GSG caused noticeable performance drops, and removing both had the most significant impact. These findings confirm that PSG and GSG are essential for effective clustering, while Laplace Noise introduces a minor trade-off between privacy and performance.

Privacy Preservation Experiment

To demonstrate the privacy benefits of our method under DP, we compared prototype visualizations from k-FED, PPFC-GAN, and our SPP-FGC on the MNIST dataset (on Fig 4). While the prototypes from k-FED and PPFC-GAN closely resemble the original digits, potentially exposing sensitive data, SPP-FGC prototypes appear blurred, indicating stronger privacy protection. This highlights the ability of SPP-FGC to

DP	PSG	GSG	NMI	ACC	ARI
✓	✓	✓	0.7193	0.6305	0.6722
✗	✓	✓	0.7228	0.6845	0.6931
✓	✗	✓	0.5847	0.6256	0.5599
✓	✓	✗	0.6793	0.6196	0.6417
✓	✗	✗	0.5793	0.4720	0.4315

Table 3: Ablation Study of SPP-FGC. It assess the impact on NMI, ACC, and ARI when removing key components.




Figure 4: Visualization of transmitted prototypes from k-FED (left column), PPFC-GAN (middle column), and SPP-FGC (right column). This figure illustrates the enhanced privacy preservation achieved by our proposed SPP-FGC algorithm.

safeguard individual data details while maintaining effective clustering, making it a reliable choice for privacy-sensitive unsupervised applications.

Conclusion

In this paper, we propose SPP-FGC and SPP-FGC+, which enhance privacy and communication efficiency in FL through private structure mining, global relation aggregation, and DP. These approaches improve clustering performance while protecting sensitive data, making them suitable for privacy-sensitive applications. Our experiments demonstrated their effectiveness and scalability. Future work will focus on reducing costs with sparse graph representations and exploring adaptive privacy techniques to support diverse data distributions, further advancing FL’s security and applicability.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant 62406250, 62236001, 62576277 and the Fundamental Research Funds for the Central Universities.

References

Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H. B.; Mironov, I.; Talwar, K.; and Zhang, L. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 308–318.

- Basagni, S. 1999. Distributed clustering for ad hoc networks. In *Proceedings Fourth International Symposium on Parallel Architectures, Algorithms, and Networks (I-SPAN'99)*, 310–315. IEEE.
- Chen, J.; Sun, H.; Woodruff, D.; and Zhang, Q. 2016. Communication-optimal distributed clustering. *Advances in Neural Information Processing Systems*, 29.
- Chen, X.; Xu, J.; Ren, Y.; Pu, X.; Zhu, C.; Zhu, X.; Hao, Z.; and He, L. 2023. Federated deep multi-view clustering with global self-supervision. In *Proceedings of the 31st ACM International Conference on Multimedia*, 3498–3506.
- Dennis, D. K.; Li, T.; and Smith, V. 2021. Heterogeneity for the win: One-shot federated clustering. In *International Conference on Machine Learning*, 2611–2620. PMLR.
- Doersch, C. 2016. Tutorial on variational autoencoders.
- Dwork, C. 2006. Differential privacy. In *International colloquium on automata, languages, and programming*, 1–12. Springer.
- Fan, K. 1949. On a theorem of Weyl concerning eigenvalues of linear transformations I. *Proceedings of the National Academy of Sciences*, 35(11): 652–655.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Commun. ACM*, 63(11): 139–144.
- Hasan, N.; Alam, M. G. R.; Ripon, S. H.; Pham, P. H.; and Hassan, M. M. 2025. An autoencoder-based confederated clustering leveraging a robust model fusion strategy for federated unsupervised learning. *Information Fusion*, 115: 102751.
- Hu, M.; Cao, Y.; Li, A.; Li, Z.; Liu, C.; Li, T.; Chen, M.; and Liu, Y. 2024. FedMut: Generalized Federated Learning via Stochastic Mutation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 12528–12537.
- Hu, R.; Guo, Y.; Li, H.; Pei, Q.; and Gong, Y. 2020. Personalized federated learning with differential privacy. *IEEE Internet of Things Journal*, 7(10): 9530–9539.
- Huang, W.; Ye, M.; Shi, Z.; Li, H.; and Du, B. 2023. Rethinking federated learning with domain shift: A prototype view. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16312–16322. IEEE.
- Li, S.; Hou, S.; Buyukates, B.; and Avestimehr, S. 2022. Secure federated clustering.
- Li, Y.; Wang, S.; Chi, C.-Y.; and Quek, T. Q. 2023. Differentially private federated clustering over non-IID data. *IEEE Internet of Things Journal*.
- Mashhadi, A.; Sterner, J.; and Murray, J. 2021. Deep embedded clustering of urban communities using federated learning. In *2021 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282. PMLR.
- McMahan, H. B.; Yu, F.; Richtarik, P.; Suresh, A.; Bacon, D.; et al. 2016. Federated learning: Strategies for improving communication efficiency. In *Proceedings of the 29th Conference on Neural Information Processing Systems (NIPS), Barcelona, Spain*, 5–10.
- Nardi, M.; Valerio, L.; and Passarella, A. 2024. Federated Clustering: An Unsupervised Cluster-Wise Training for Decentralized Data Distributions. *arXiv preprint arXiv:2408.10664*.
- Nie, F.; Wang, X.; Jordan, M.; and Huang, H. 2016. The constrained laplacian rank algorithm for graph-based clustering. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30.
- Pan, C.; Sima, J.; Prakash, S.; Rana, V.; and Milenkovic, O. 2022. Machine unlearning of federated clusters. *arXiv preprint arXiv:2210.16424*.
- Pedrycz, W. 2021. Federated FCM: Clustering under privacy requirements. *IEEE Transactions on Fuzzy Systems*, 30(8): 3384–3388.
- Presotto, R.; Civitaresse, G.; and Bettini, C. 2022. Fedclar: Federated clustering for personalized sensor-based human activity recognition. In *2022 IEEE international conference on pervasive computing and communications (PerCom)*, 227–236. IEEE.
- Qiao, D.; Ding, C.; and Fan, J. 2024. Federated spectral clustering via secure similarity reconstruction. volume 36.
- Reddi, S. J.; Charles, Z.; Zaheer, M.; Garrett, Z.; Rush, K.; Konečný, J.; Kumar, S.; and McMahan, H. B. 2021. Adaptive Federated Optimization. In *International Conference on Learning Representations*.
- Tan, Y.; Long, G.; Liu, L.; Zhou, T.; Lu, Q.; Jiang, J.; and Zhang, C. 2022a. Fedproto: Federated prototype learning across heterogeneous clients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 8432–8440.
- Tan, Y.; Long, G.; Ma, J.; Liu, L.; Zhou, T.; and Jiang, J. 2022b. Federated learning from pre-trained models: A contrastive learning approach. volume 35, 19332–19344.
- Triastcyn, A.; and Faltings, B. 2019. Federated learning with bayesian differential privacy. In *2019 IEEE International Conference on Big Data (Big Data)*, 2587–2596. IEEE.
- Wang, M.; Bodonhelyi, A.; Bozkir, E.; and Kasneci, E. 2024. TurboSVM-FL: Boosting Federated Learning through SVM Aggregation for Lazy Clients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 15546–15554.
- Wang, Y.; Jia, M.; Gao, N.; Von Krannichfeldt, L.; Sun, M.; and Hug, G. 2022. Federated clustering for electricity consumption pattern extraction. *IEEE Transactions on Smart Grid*, 13(3): 2425–2439.
- Wei, K.; Li, J.; Ding, M.; Ma, C.; Yang, H. H.; Farokhi, F.; Jin, S.; Quek, T. Q.; and Poor, H. V. 2020. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE transactions on information forensics and security*, 15: 3454–3469.
- Wu, X.; Huang, F.; Hu, Z.; and Huang, H. 2023. Faster adaptive federated learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 10379–10387.

- Xie, J.; Girshick, R.; and Farhadi, A. 2016. Unsupervised deep embedding for clustering analysis. In *International conference on machine learning*, 478–487. PMLR.
- Xie, S.; Wu, Y.; Liao, K.; Chen, L.; Liu, C.; Shen, H.; Tang, M.; and Sun, L. 2023. Fed-SC: One-Shot Federated Subspace Clustering over High-Dimensional Data. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, 2905–2918. IEEE.
- Yan, J.; Liu, J.; Qi, J.; and Zhang, Z.-Y. 2023. Privacy-Preserving Federated Deep Clustering based on GAN. arXiv:2211.16965.
- Yang, Q.; Liu, Y.; Chen, T.; and Tong, Y. 2019. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2): 1–19.
- Zhang, F.; Kuang, K.; Chen, L.; You, Z.; Shen, T.; Xiao, J.; Zhang, Y.; Wu, C.; Wu, F.; Zhuang, Y.; et al. 2023. Federated unsupervised representation learning. *Frontiers of Information Technology & Electronic Engineering*, 24(8): 1181–1193.
- Zhao, Y.; and Chen, J. 2022. A survey on differential privacy for unstructured data content. *ACM Computing Surveys (CSUR)*, 54(10s): 1–28.