Federated Learning-Based Cross-Domain LiDAR Point Cloud Processing And Protection Of Sensitive Geographic Information

Yan Li*

Lyuliang University, Lyuliang, 033000, China

** Corresponding author. E-mail: liyan033000@126.com*

Received: Jul. 21, 2026; Accepted: Aug. 08, 2026

Cross-domain LiDAR point clouds have become an important data foundation for smart cities, high-precision map updating, road inspection, industrial-park digital twins, and critical infrastructure monitoring, yet they are usually collected in a decentralized manner by different institutions, regions, and sensing platforms, so direct centralized sharing may increase the risk of exposing sensitive geographic coordinates, facility outlines, and road topology. Because existing centralized point-cloud learning methods still struggle to balance cross-domain generalization, privacy protection, and system scalability, this paper proposes a scalable federated learning system for cross-domain LiDAR point cloud processing. The system adopts a three-layer architecture of edge point-cloud clients, regional coordination nodes, and a cloud-based secure aggregation server, in which point-cloud sampling, coordinate anonymization, local geometric feature extraction, and semantic segmentation are performed locally. Collaborative modeling without raw point-cloud upload is enabled through cross-domain category prototype alignment, client-adaptive weight aggregation, differential privacy, secure aggregation, Top-k sparsification, and 8-bit quantization. In simulations across four cross-domain scenarios, the method achieves an average mIoU of 67.2%, exceeding FedAvg (58.8%) and FedBN (62.0%), and retains 61.4% under severe Non-IID conditions. At a privacy budget of $\varepsilon = 2$, the attack success rate falls from 63.5% to 24.8% while mIoU remains 67.2%, and with 100 clients single-round communication drops from 21.96 GB to 7.92 GB. The system thus improves semantic performance, reduces sensitive-information leakage, and maintains communication efficiency and training stability as the number of clients grows.

Keywords: federated learning; LiDAR point cloud; cross-domain learning; sensitive geographic information protection; scalable information systems; differential privacy; edge-cloud collaboration.

http://dx.doi.org/10.6180/jase.202612_35.002

1. Introduction

LiDAR point clouds represent real spatial environments through three-dimensional coordinates, reflection intensity, and temporal scanning information. They can accurately depict roads, buildings, vehicles, poles, vegetation, industrial facilities, and terrain variations. In recent years, deep learning-based point-cloud processing has evolved from single models to segmentation, detection, and panoptic understanding of large-scale scenes, and point-cloud semantic segmentation has developed a relatively mature research foundation and task system [1]. With the development of intelligent transportation, urban digital twins, infrastructure inspection, and high-precision map updating, point-cloud data are no longer merely a single surveying product; they have gradually become a core data resource in scalable geospatial information systems.

In a centralized model, point-cloud data must be uploaded to a unified server for training and analysis. However, cross-domain LiDAR data typically come from dif-

ferent cities, parks, roads, industrial scenarios, and sensor platforms, leading to substantial differences in point density, scanning perspective, category distribution, and spatial scale. Although recent point-cloud models (reviewed in Section 1.1) have improved spatial feature representation, most still assume that training data can be used centrally, and in real multi-agency operations, raw point clouds often cannot be shared across domains, which calls for distributed learning architectures.

Beyond raw accuracy, intelligent point-cloud processing must also consider edge-computing capacity, scene variation, and annotation cost, as discussed further in Section 1.1. When road inspection units, park management units, industrial operation and maintenance teams, and urban surveying agencies each hold their own point-cloud data, centralized training faces constraints related to data scale, network bandwidth, compliance review, and leakage of sensitive geographic information. Absolute coordinates, road topology, building outlines, facility locations, and internal spatial relationships in raw point clouds may be re-registered with external maps or public images, creating a risk of spatial re-identification.

Federated learning provides a feasible way to address these problems by keeping data local while allowing model updates to participate in collaborative training, as reviewed in Section 1.2. However, cross-domain LiDAR scenarios are not a direct application of standard horizontal federated learning. Point-cloud data are irregular, sparse, spatially correlated, and long-tailed across categories, and Non-IID characteristics across clients appear not only as different category ratios but also as differences in point density, occlusion patterns, reflection intensity, and coordinate ranges, which can significantly reduce the convergence quality of the global model.

In terms of privacy protection, existing privacy-preserving federated learning techniques (reviewed in Section 1.2) can reduce the leakage risk of model updates, but for LiDAR point clouds, protecting gradients alone is not sufficient to protect sensitive geospatial structures; communication compression, device heterogeneity, regional aggregation, and throughput must also be considered in real-world deployments. Therefore, this paper treats scalable information systems, cross-domain point-cloud processing, and geographic privacy protection as an integrated problem.

The main contributions of this paper are fourfold. First, it constructs a three-layer edge-cloud collaborative federated architecture for cross-domain LiDAR point-cloud processing, enabling multiple institutions to train collaboratively without uploading raw point clouds. Second, it proposes cross-domain category prototype alignment and client-adaptive aggregation strategies to alleviate performance degradation caused by Non-IID point-cloud distributions. Third, it designs a joint protection mechanism that combines coordinate de-identification, differential privacy, and secure aggregation to reduce the leakage risk of absolute coordinates and model updates. Fourth, it introduces Top-k sparsification, 8-bit quantization, and hierarchical aggregation to improve communication efficiency as the number of clients increases.

1.1. LiDAR Point Cloud Processing

LiDAR point-cloud processing has evolved from artificial geometric features, voxel grids, and projected images to point-level deep networks. Existing reviews indicate that point-cloud classification, object detection, semantic segmentation, and panoptic segmentation now form a relatively complete technical system [1]. In large-scale environments, multi-scale point-wise convolutional networks can enhance the segmentation of objects at different spatial scales [2]. Transformer structures help capture long-range dependencies in point clouds [3]. These studies provide the model foundation for the point-cloud processing module in this paper.

However, existing point-cloud models are mainly designed for centralized training or single-domain evaluation. For point-cloud information systems spanning cities, parks, and industrial scenarios, models must be transferred across heterogeneous sensors and spatial structures. Research on real-time point-cloud segmentation at the edge emphasizes the needs of embedded devices, unmanned systems, and real-time perception [4], while studies of point clouds under adverse weather show that scene disturbances can significantly affect segmentation stability [5]. These characteristics indicate that cross-domain LiDAR systems require not only high-precision models but also training mechanisms that can adapt to data silos and edge deployments.

1.2. Federated Learning and Scalable Information Systems

The core idea of federated learning systems is to perform collaborative modeling without centralizing raw data. From a systems perspective, Li et al. summarized the participants, communication protocols, aggregation mechanisms, privacy protection methods, and evaluation frameworks of federated learning [6]. Kairouz et al. pointed out, from both theoretical and engineering perspectives, that federated learning still faces open issues involving communication efficiency, statistical heterogeneity, security attacks, and fairness [7]. For Non-IID data, Zhu et al. showed that

class skew, quantity skew, and domain skew can cause local model drift, making it difficult for standard FedAvg to obtain a stable global model [8]. These conclusions are particularly important for point-cloud scenarios because different clients may contain only some categories of roads, buildings, industrial facilities, or vegetation.

Research on privacy-preserving federated learning suggests that differential privacy, secure aggregation, and cryptographic computation can reduce leakage risks at the model-update level [9]. At the same time, studies of edge computing and blockchain-based federated learning emphasize trusted collaboration, edge-node management, and distributed scheduling [10]. Communication-efficient federated learning reduces model-update overhead through pruning, quantization, and sparse uploading [11]. Reviews of differential privacy further show that privacy mechanisms inevitably trade off against model performance [12]. Therefore, this paper does not treat privacy protection as an add-on module; instead, it incorporates privacy into the optimization objective of a cross-domain federated point-cloud system.

1.3. Geospatial Privacy and Point Cloud Information Systems

Geospatial data can be re-identified through spatial patterns. Even if explicit place names are removed, coordinates, road shapes, and building outlines can still be matched again. Differential privacy studies for geographic analysis show that spatial statistics and location data protection must consider both utility preservation and reconstruction risk [13]. Studies on large-scale point-cloud servers show that point-cloud management, indexing, compression, and multi-user access are important components of geospatial information systems [14]. In remote sensing and airborne laser scanning point clouds, semantic segmentation models need to use both local geometry and global context [15].

In summary, existing research has established a foundation in point-cloud models, federated systems, and privacy mechanisms, but three limitations remain. First, federated learning system designs for cross-domain LiDAR point clouds are still limited and do not fully account for heterogeneity across regions, devices, and institutions. Second, privacy protection often remains at the model-update level and lacks joint modeling of sensitive geographic coordinates and spatial-structure leakage. Third, experimental evaluations tend to focus on accuracy metrics and provide insufficient system-level verification of client scale, communication cost, aggregation latency, point-cloud visualization, and response heatmaps. This paper addresses these limitations.

2. Methods

2.1. Cross-Domain Federated LiDAR Scenarios

This paper considers a cross-domain point-cloud information system consisting of multiple LiDAR clients, regional coordination nodes, and a cloud-based secure aggregation server. Each client stores a local point-cloud dataset that includes point coordinates, reflection intensity, local timestamps, and semantic labels. Different clients correspond to different regions or organizations, such as urban road inspection units, park digital twin platforms, industrial infrastructure operation and maintenance units, and suburban traffic sensing nodes. During training, clients upload only model updates that have been pruned, perturbed, and compressed; they do not upload raw point clouds. Fig. 1 shows the overall system architecture.

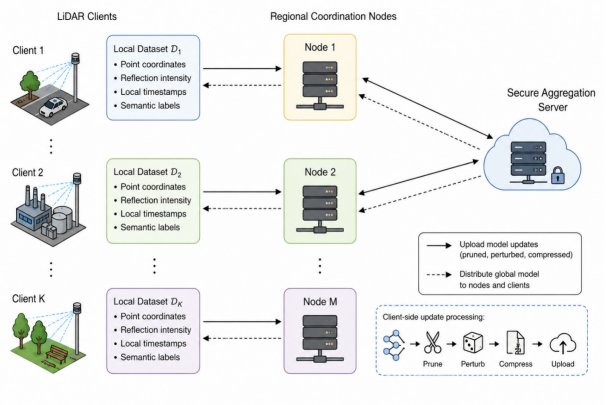

Fig. 1. Architecture of the scalable federated LiDAR point-cloud processing system

2.2. Sensitive Geographic Information and Threat Model

This paper defines sensitive geographic information as five categories: absolute geographic coordinates, critical infrastructure outlines, high-precision map structures, road topology, and internal institutional spatial assets. Attackers cannot directly obtain raw point clouds, but they may observe model updates, aggregated models, and partial public-map priors during communication. Attack targets include recovering local point-cloud morphology through gradient inversion, determining whether a specific area participated in training through membership inference, recovering the collection area's location through coordinate reconstruction, and identifying sensitive facilities through contour registration.

2.3. Optimization Objective and Mathematical Definition

This paper integrates the global point-cloud segmentation model, client data, sample size, local features, category prototypes, and cross-domain constraints into a mathematical formulation. The overall goal of federated training is to minimize the weighted empirical risk of each client while constraining privacy leakage and communication cost. The local empirical loss is defined in Eq. (1).

$$\mathcal{L}_k(\theta) = \frac{1}{n_k} \sum_{i=1}^{n_k} \ell \left(f_\theta(P_{k,i}), y_{k,i} \right) + \lambda \| z_{k,i} - c_{y_i} \|_2^2 \quad (1)$$

In Eq. (1), the variables represent the client point-cloud samples, point-level semantic labels, local point-cloud network, point-cloud local features, category prototypes, and cross-domain prototype constraint weights, respectively. This loss function incorporates both semantic segmentation error and category prototype distance into the optimization, enabling the local model to learn point-level labels while maintaining consistent geometric representations across domains. The global federated optimization objective is shown in Eq. (2).

$$\begin{aligned} \min_\theta F(\theta) = & \sum_{k=1}^{K} \frac{n_k}{N} \mathcal{L}_k(\theta) + \alpha \Omega_{\text{domain}}(\theta) \\ & + \beta \Omega_{\text{privacy}}(\theta) + \gamma \Omega_{\text{comm}}(\theta) \end{aligned} \quad (2)$$

In Eq. (2), the cross-domain alignment regularization term, privacy protection cost, and communication compression cost constrain cross-domain generalization, sensitive information protection, and communication efficiency, respectively. The three trade-off coefficients control the balance among accuracy, privacy, and communication cost.

2.4. Privacy-Preserving Federated Point-Cloud Processing Workflow

The proposed method uses an edge point-cloud client as the basic training unit. At the start of each training round, the cloud distributes the global model. The client first performs coordinate centering, relative coordinate transformation, and grid-based anonymization, followed by point-cloud sampling, neighborhood construction, and local feature learning. After local training is completed, the client applies gradient pruning, differential privacy perturbation, Top-k sparsification, and 8-bit quantization, and then uploads the model update to the regional coordination node. After the regional node completes local aggregation, the cloud performs secure aggregation and updates the global model.

2.5. Cross-Domain Category Prototype Alignment Mechanism

The main challenge in cross-domain point-cloud analysis is that the same category may have different point densities, boundary shapes, and occlusion patterns across domains, while different categories may share similar local geometry. This paper maintains a category prototype for each client and forms a cross-domain global category prototype at the regional coordination layer. During local client updates, the model minimizes not only segmentation loss but also the distance between features of the same point-cloud category and the corresponding prototype. As a result, categories such as roads, buildings, poles, and vehicles form a more consistent feature space across domains. Category prototype aggregation is shown in Eq. (3).

$$c_m^{t+1} = \sum_{k \in \mathcal{S}_t} \frac{n_{k,m}}{\sum_{j \in \mathcal{S}_t} n_{j,m}} c_{k,m}^t \quad (3)$$

In Eq. (3), the set of clients participating in training and the number of valid points in each client category jointly determine the global prototype update weight. This weighting method reduces the unstable influence of a small number of clients within a category while preventing large clients from fully dominating the global representation.

2.6. Sensitive Geographic Information Protection Mechanism

To protect sensitive geographic information, this paper adopts a joint mechanism that combines coordinate de-identification, secure aggregation, and differential privacy. Coordinate de-identification first converts absolute coordinates into local relative coordinates and then applies lightweight masking on the spatial grid, making it difficult for external attackers to register the real area directly through absolute coordinates or contours. At the model-update level, each client performs norm clipping on the gradient and adds Gaussian noise that satisfies differential privacy, as shown in Eq. (4).

$$\begin{aligned} \tilde{g}_k = & \text{clip}(g_k, C) + \mathcal{N}\left(0, \sigma^2 C^2 I\right), \\ & \text{clip}(g_k, C) = g_k \cdot \min\left(1, \frac{C}{\|g_k\|_2}\right) \end{aligned} \quad (4)$$

In Eq. (4), the local gradient, pruning threshold, noise multiplier, and identity matrix jointly determine the strength of the differential privacy perturbation. Secure aggregation ensures that the server can obtain only the aggregated result from multiple client updates and cannot directly read the model update of any individual client.

Table 1. Dataset domains and cross-domain client settings

Domain	Scene	Clients	Main classes	Distribution characteristics
A	Urban road	10-25	roads, vehicles, pedestrians, poles	dense traffic and clear road boundaries
B	Campus/park	10-25	buildings, vegetation, roads, fences	mixed buildings and vegetation
C	Industrial area	10-25	equipment, pipelines, buildings, ground	complex occlusions and sparse targets
D	Suburban road	10-25	roads, signs, barriers, vehicles	long-range sensing and low density

Table 2. Evaluation metrics used in the simulations

Category	Metrics	Description
Point cloud processing	mIoU, OA, Precision, Recall, F1	semantic segmentation performance
Cross-domain learning	target-domain mIoU, class-wise IoU, domain-gap reduction	generalization across domains
Privacy protection	privacy budget, attack success rate, coordinate reconstruction error	leakage risk under the threat model
System scalability	communication cost, latency, convergence rounds, throughput	scalable information-system performance

2.7. Communication-Efficient and Scalable Aggregation

Point-cloud networks have large parameter counts, and multi-client training can easily create communication bottlenecks. This paper combines Top-k sparsification with 8-bit quantization to reduce both the number of uploaded parameters and the bit width of each parameter representation. The upload update and global model update rules are shown in Eq. (5).

$$\tilde{\Delta}_k = Q_8\left(\mathrm{Top}_k\left(\tilde{g}_k\right)\right), \quad \theta^{t+1} = \theta^t + \eta \sum_{k \in \mathcal{S}_t} w_k \tilde{\Delta}_k \tag{5}$$

In Eq. (5), the client aggregation weight comprehensively considers sample size, domain differences, communication quality, and historical contribution. The system uses regional hierarchical aggregation to reduce cloud-side pressure and allows lightweight asynchronous updates to reduce waiting time caused by slow clients.

2.8. Dataset and Cross-Domain Client Partitioning

The simulation uses publicly available LiDAR point-cloud scenes and simulated institutional domains to build cross-domain clients. The four domains correspond to urban roads, industrial parks or campuses, industrial infrastructure, and suburban roads. Each domain contains multiple clients, with differences in category distribution, point density, scan radius, and noise level across clients. Table 1 shows the data domains and client settings.

2.9. Evaluation Metrics

To reflect the evaluation requirements of scalable information systems research, this paper uses not only segmentation accuracy but also evaluates the system from four perspectives: cross-domain generalization, privacy protection, communication efficiency, and system scalability. Table 2 lists the main metrics.

2.10. Comparison Methods and Parameter Settings

This paper compares eight methods: Local Only, Centralized, FedAvg, FedProx, FedBN, FedAvg+DP, FedAvg+Compression, and Proposed. Local Only trains each client individually. Centralized training is used as a theoretical upper bound. FedAvg applies standard federated averaging. FedProx mitigates local drift through proximal constraints. FedBN preserves local batch normalization to accommodate domain differences. FedAvg+DP is used to analyze the effect of privacy perturbation, and FedAvg+Compression is used to analyze the effect of communication compression. Proposed denotes the complete

Table 3. Overall performance comparison of federated LiDAR point cloud processing methods

Method	Urban mIoU	Campus mIoU	Industrial mIoU	Suburban mIoU	Avg. mIoU	OA	F1
Local Only	58.7	52.4	49.8	51.3	53.1	78.6	70.2
FedAvg	63.5	58.9	55.6	57.1	58.8	82.1	74.8
FedProx	65.1	60.7	57.4	58.8	60.5	83.4	76.1
FedBN	66.3	62.1	59.2	60.5	62.0	84.2	77.3
FedAvg+DP	61.8	56.7	53.9	55.2	56.9	81.0	73.2
Proposed	70.6	67.4	64.8	65.9	67.2	87.6	81.5

system presented in this paper. Unless otherwise specified, the experiments use 20 clients, 150 training rounds, 5 local epochs, a batch size of 64, an Adam learning rate of 0.001, a default privacy budget ε of 2, a gradient pruning threshold C of 1.0, and a default compression rate of 50%.

3. Results and discussion

3.1. Overall Point Cloud Processing Performance Comparison

Table 3 presents the overall performance of different methods in four cross-domain scenarios. As shown in Table 3, Local Only achieves the lowest performance under cross-domain conditions, indicating that single-client models struggle to learn shared features across scenarios. FedAvg improves over Local Only, but it remains strongly affected by domain shift in the industrial and suburban scenarios. FedProx and FedBN mitigate local drift and normalization differences, achieving average mIoUs of 60.5% and 62.0%, respectively. FedAvg+DP, because of the added gradient noise, obtains an average mIoU of only 56.9%, indicating that privacy protection weakens the identification of small targets and boundaries when it is not combined with point-cloud structure preservation. The proposed method achieves an average mIoU of 67.2% and an F1-score of 81.5%, giving the best performance in all four domains.

3.2. Algorithm Performance Comparison

Fig. 2 shows that conventional federated methods degrade more noticeably in the industrial infrastructure and suburban road domains because of complex occlusion, long-distance sparse points, and class imbalance. The proposed method outperforms the baselines in all four domains, demonstrating that cross-domain prototype alignment and adaptive aggregation can improve the identification of both large-area and sparse classes.

3.3. Robustness to Non-IID Strength

Fig. 3 illustrates performance changes under different Non-IID intensities. As the data distribution shifts gradually from IID to severe Non-IID, all methods decline, with FedAvg decreasing from 66.8% to 50.9%, a drop of 15.9 percentage points. FedProx and FedBN reach 54.6% and 56.1%, respectively, under severe Non-IID conditions, but they still show substantial degradation. The proposed method maintains 61.4% under severe Non-IID conditions, an improvement of 10.5 percentage points over FedAvg, indicating that cross-domain prototype alignment can mitigate the effect of client-side local optima on the global model.

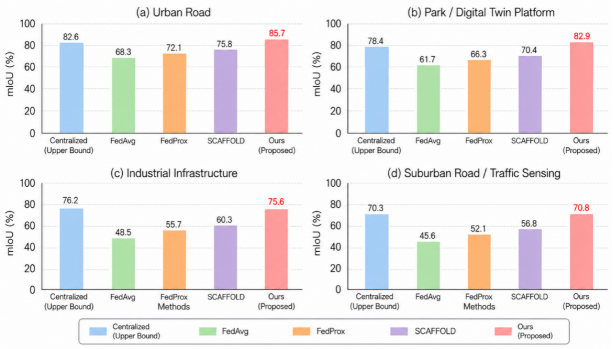

Fig. 2. Performance comparison under cross-domain LiDAR point-cloud scenarios

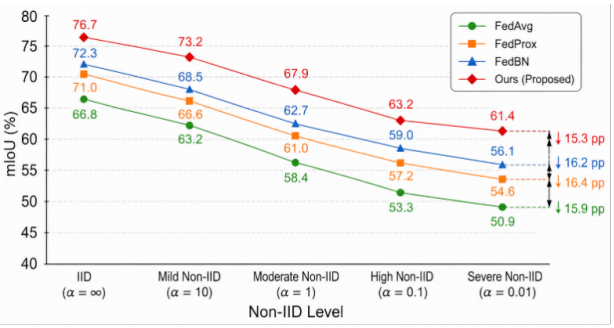

Fig. 3. Performance degradation under different Non-IID levels

3.4. Privacy Budget and Utility Trade-off

Fig. 4 illustrates changes in mIoU and attack success rate under different differential privacy budgets. Without differential privacy protection, point-cloud processing accuracy is the highest, but the attack success rate reaches 63.5%, indicating a high risk of sensitive geographic information leakage. When the privacy budget is 4, the mIoU is 68.9%, only 1.9 percentage points lower than without privacy protection, while the attack success rate drops to 36.4%. When the privacy budget is 2, the attack success rate further decreases to 24.8%, while the mIoU remains at 67.2%, showing a better balance. Although a privacy budget of 0.5 provides the strongest privacy protection, the mIoU drops to 59.6%, making it unsuitable for point-cloud applications that require high boundary and small-target accuracy.

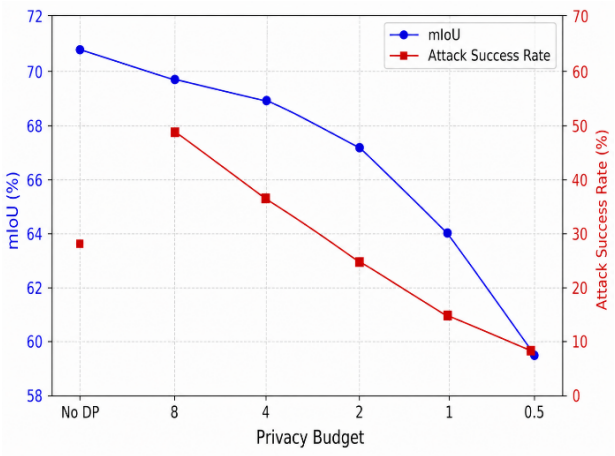

Fig. 4. Privacy-utility trade-off under different differential privacy budgets

3.5. Visualization of Sensitive Geographic Coordinate De-identification

Fig. 5 shows point-cloud visualization before and after coordinate anonymization. The original point cloud directly reflects road locations, building boundaries, and facility spatial relationships. If registered with a public map, it poses a high risk of location leakage. Coordinate centering can hide absolute locations while preserving scene shape. Grid-based masking can reduce spatial accuracy, but excessively coarse grids may cause boundary misalignment. The hybrid anonymization method balances relative coordinates, grid masking, and minimal noise, preserving the basic local structure of roads, buildings, and poles while increasing the difficulty of coordinate reconstruction.

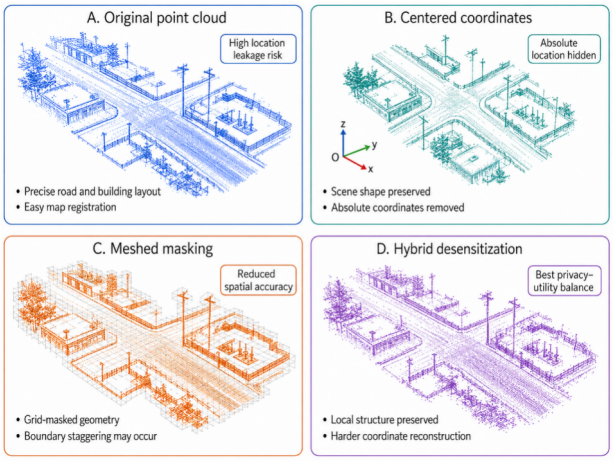

Fig. 5. Visualization of sensitive geospatial coordinate de-identification in LiDAR point clouds

3.6. Communication Efficiency and Scalability

Fig. 6 illustrates communication volume and model performance as the number of clients changes. Under FedAvg, the single-round communication volume increases approximately linearly with the number of clients, reaching 21.96 GB with 100 clients. By combining Top-k sparsification, 8-bit quantization, and region-level hierarchical aggregation, the proposed method reduces the single-round communication volume to 7.92 GB with 100 clients, a reduction of approximately 63.9%. Meanwhile, the mIoU reaches 70.3% around 50 clients and remains at 69.8% with 100 clients, indicating that system scaling does not cause substantial performance loss.

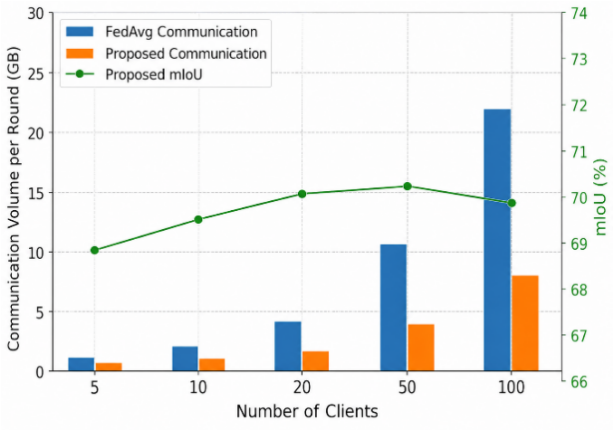

Fig. 6. Scalability evaluation with different numbers of federated LiDAR clients

3.7. Attention Heatmap Analysis

Fig. 7 illustrates the attention differences between the conventional federated model and the proposed method on key spatial structures. It shows that the proposed method produces higher attention responses for road boundaries, building facades, vehicle outlines, and pole-like regions, while producing lower responses for large flat backgrounds and noisy points. Compared with the conventional federated model, the proposed method focuses more on cross-domain stable structures, indicating that category prototype alignment guides the model toward discriminative geometric regions.

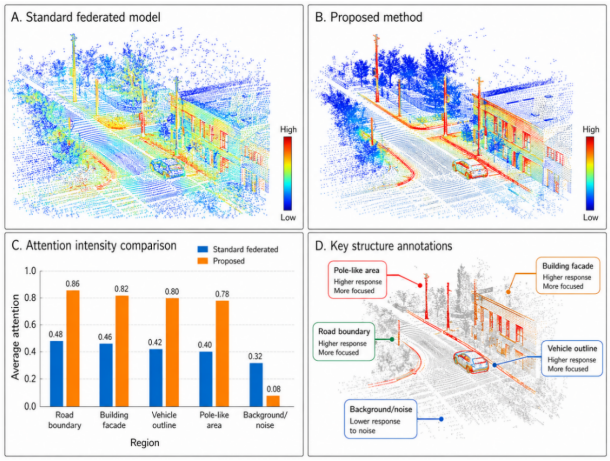

Fig. 7. Attention heatmap of key spatial structures in cross-domain LiDAR scenes

3.8. Response Heatmap Analysis

The response heatmap is used to observe whether the model's response to key regions remains stable after the privacy mechanism is added. Fig. 8 compares the response distributions of the model without privacy protection, with differential privacy (DP) only, with coordinate masking only, and with the proposed joint method. The model without privacy protection has the strongest response but also the highest privacy risk. When only DP is used, responses to some small targets and boundary regions are significantly weakened. When only coordinate masking is used, geometric continuity is affected to some extent. After coordinate de-identification and privacy perturbation are added, the proposed method still maintains strong responses at road boundaries, building outlines, and pole-like regions. This indicates that the proposed privacy protection mechanism does not simply add noise; it balances privacy risk and structural response.

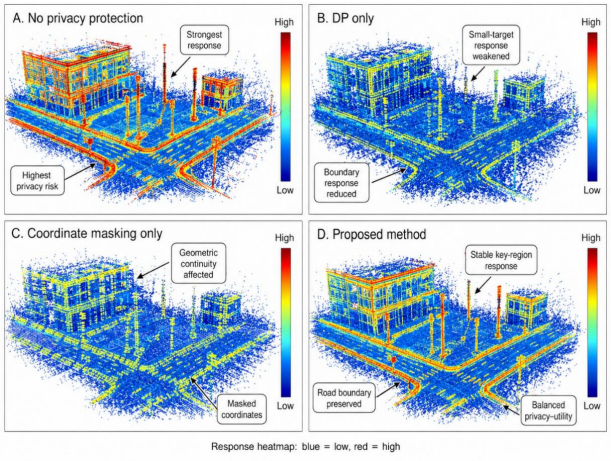

Fig. 8. Response heatmap under privacy-preserving federated LiDAR processing

3.9. Category-Level IoU Analysis

Table 4 presents the category-level IoU comparison. The results show that the most significant cross-domain degradation occurs not in large-area categories such as roads and buildings but in sparse categories such as poles, traffic signs, vehicle boundaries, and industrial equipment. The proposed method enhances the consistent representation of minority categories through category prototype constraints. The IoU for poles and traffic signs reaches 52.8%, an improvement of 10.2 percentage points over FedAvg, and the IoU for industrial equipment reaches 59.6%, an improvement of 11.1 percentage points over FedAvg.

Table 4. Class-wise IoU comparison under cross-domain scenarios

Category	FedAvg	FedBN	Proposed	Gain over FedAvg
Road	78.2	80.5	84.7	+6.5
Building	69.4	72.1	76.8	+7.4
Vehicle	63.7	65.9	71.5	+7.8
Pole / traffic sign	42.6	45.3	52.8	+10.2
Vegetation	66.1	68.4	73.2	+7.1
Industrial equipment	48.5	52.4	59.6	+11.1

3.10. Module Ablation Study

Table 5 presents the ablation results for key modules. Removing cross-domain alignment reduces the mIoU to 62.9%, indicating that this module contributes the most to cross-scenario generalization. Removing privacy protec-

tion increases the mIoU to 70.8%, but the attack success rate rises to 63.5%, failing to meet the requirements for sensitive geographic information protection. Removing communication compression increases the single-round communication volume to 21.12 GB. Removing hierarchical aggregation does not significantly change communication volume, but it increases aggregation latency to 161.2 s.

Table 5. Ablation study of key modules

Ablation variant	mIoU	Attack success	Communication cost/round	Latency
Without domain alignment	62.9	25.7%	7.95 GB	108.4 s
Without privacy protection	70.8	63.5%	7.88 GB	106.5 s
Without compression	68.1	24.6%	21.12 GB	119.6 s
Without hierarchical aggregation	67.6	24.9%	8.03 GB	161.2 s
Full model	67.2	24.8%	7.92 GB	107.8 s

3.11. Cross-Domain Transfer Matrix Analysis

Table 6 presents transfer results from different source domains to the target domain. Transfers from A to C and from B to C are the most difficult because the industrial infrastructure domain contains numerous pipelines, equipment boundaries, and complex occlusions, and its geometry differs substantially from urban roads and industrial park roads. The proposed method achieves an mIoU of 69.5% under the Multi-source→Target setting, demonstrating that complementary geometry from multiple source clients can be converted into effective generalization capability through prototype alignment.

3.12. Robustness of Heterogeneous Devices and Network Bandwidth

Table 7 presents robustness results under heterogeneous computing power and bandwidth conditions. When 60% of clients are low-computing-power nodes, FedAvg training time increases by a factor of 1.86, while the mIoU decreases to 54.7%. Under the same setting, the proposed method has a training time ratio of 1.18, and the mIoU remains at 65.1%. When 60% of clients have limited bandwidth, FedAvg training time increases by a factor of 2.14, while the proposed method increases by a factor of 1.31, indicating that hierarchical aggregation and adaptive compression can reduce system tail latency.

Table 6. Cross-domain transfer matrix under different source-target settings

Transfer setting	Local	FedAvg	FedBN	Proposed
A→B	54.8	59.6	62.4	67.1
A→C	48.2	54.1	57.3	63.5
A→D	50.6	56.7	59.1	64.2
B→C	49.5	55.2	58.0	64.1
C→D	51.1	57.0	60.4	65.3
Multi-source →target	58.9	63.8	66.2	69.5

Table 7. Robustness under heterogeneous computing and bandwidth settings

Scenario	FedAvg time ratio	FedAvg mIoU	Proposed time ratio	Proposed mIoU
20% of low-power clients	1.22	58.0	1.07	66.8
40% of low-power clients	1.51	56.2	1.12	65.9
60% of low-power clients	1.86	54.7	1.18	65.1
60% of clients bandwidth-limited	2.14	53.9	1.31	64.6

3.13. Attack Simulation and Sensitive Information Exposure Analysis

Table 8 presents simulations of three attack types: gradient inversion, membership inference, and sensitive-region exposure. Secure aggregation alone can reduce single-client update leakage, but it offers limited protection for the coordinate structure itself. DP alone can significantly reduce the success rate of gradient inversion, but it still leaves residual exposure of sensitive-region contours. Coordinate masking alone can reduce geographic location exposure, but it cannot adequately protect model updates. The joint mechanism presented in this paper shows the most stable performance across all three attack metrics.

Table 8. Attack simulation and sensitive geospatial exposure analysis

Defense setting	Gradient inversion	Membership inference	Sensitive-region exposure	Coordinate error
No protection	61.8%	58.4%	52.7%	3.2 m
Secure aggregation only	42.5%	44.8%	49.1%	4.8 m
DP only	18.7%	21.3%	38.6%	9.4 m
Coordinate masking only	50.2%	51.6%	23.8%	14.2 m
Proposed	14.9%	18.5%	12.6%	16.5 m

3.14. Discussion

From the perspective of scalable information systems, the core value of this paper lies not in a single segmentation network but in embedding cross-domain LiDAR point-cloud processing into a real-world system scenario characterized by distributed architecture, privacy protection, and limited communication. Fig. 1 shows edge clients, regional coordination nodes, and a cloud-based secure aggregation server, corresponding to data ownership and computing-resource levels in a real multi-agency point-cloud service. Experimental results demonstrate that the system maintains relatively stable performance even as the number of clients increases, data become more Non-IID, and network conditions deteriorate.

From a privacy perspective, the distinctive feature of point-cloud data is that the spatial structure itself can be sensitive information. Traditional differential privacy primarily protects model updates but cannot directly protect absolute coordinates and facility outlines. Simple coordinate perturbation may disrupt local neighborhoods, thereby reducing point-cloud semantic segmentation performance. This paper proposes a joint modeling approach based on coordinate centering, grid masking, differential privacy, and secure aggregation, and validates it using coordinate reconstruction error, attack success rate, and sensitive-region exposure rate. This approach better meets the need to protect sensitive geographic information.

From an interpretability perspective, attention heatmaps and response heatmaps can help operations personnel determine whether the model truly focuses on critical structures such as road boundaries, building outlines, poles, vehicle outlines, and facility edges. If the model mainly focuses on background noise or large ground areas after cross-domain transfer, a high overall mIoU may still conceal engineering safety risks. The method presented in this paper shows a more stable response to critical structures in the heatmaps, indicating that cross-domain prototype alignment improves not only metrics but also the model's decision regions.

It should be noted that the simulation results in this paper are used to support method design and feasibility analysis. Formal experiments should still be conducted using real datasets for retraining and statistical analysis. If publicly available datasets are used instead of measured data from sensitive areas, the term "sensitive geographic coordinates" should be clearly described as a simulated anonymization setting to avoid presenting public point clouds as classified data. The model cannot eliminate all leakage risks; it can only substantially reduce the success rate of reconstruction and inference under the specified threat model.

4. Conclusions

This paper proposes a scalable federated learning system for cross-domain LiDAR point-cloud processing and sensitive geographic information protection. The system enables multi-agency collaborative point-cloud training through a three-layer edge-cloud architecture, improves generalization in Non-IID environments through cross-domain category prototype alignment, reduces the leakage risk of sensitive geographic information through coordinate de-identification, differential privacy, and secure aggregation, and lowers communication costs through Top-k sparsification, 8-bit quantization, and hierarchical aggregation. Experimental results show that the proposed method achieves an average mIoU of 67.2%, which is 8.4 percentage points higher than FedAvg. In severe Non-IID scenarios, it exceeds FedAvg by 10.5 percentage points. When $\varepsilon = 2$, the attack success rate drops to 24.8%. With 100 clients, the single-round communication volume is reduced by approximately 63.9% compared with FedAvg. Attention heatmaps and response heatmaps further validate the advantages of the proposed method in key geometric structure identification and response preservation after privacy perturbation. Overall, this paper aligns with the journal Scalable Information Systems' focus on distributed data mining, information security, and system scalability evaluation.

Acknowledgements

This work is supported by the Innovative Teaching Reform Project for Higher Education Institutions in Shanxi Province (J20250275); Science and Technology Department of Shanxi Province, Youth Scientific Research Project of Shanxi Provincial Basic Research Program (202303021222252).

References

[1] Y. Guo, H. Wang, Q. Hu, H. Liu, L. Liu, and M. Bennamoun, (2021) *"Deep Learning for 3D Point Clouds: A Survey"* **IEEE Transactions on Pattern Analysis and Machine Intelligence *43***(12): 4338–4364. DOI: 10.1109/TPAMI.2020.3005434.

[2] L. Ma, Y. Li, J. Li, W. Tan, Y. Yu, and M. A. Chapman, (2021) *"Multi-Scale Point-Wise Convolutional Neural Networks for 3D Object Segmentation From LiDAR Point Clouds in Large-Scale Environments"* **IEEE Transactions on Intelligent Transportation Systems *22***(2): 821–836. DOI: 10.1109/TITS.2019.2961060.

[3] H.-X. Cheng, X.-F. Han, and G.-Q. Xiao, (2023) *"TransRVNet: LiDAR Semantic Segmentation With Transformer"* **IEEE Transactions on Intelligent Transportation Systems *24***(6): 5895–5907. DOI: 10.1109/TITS.2023.3248117.

[4] F. Wang, Z. Wu, Y. Yang, W. Li, Y. Liu, and Y. Zhuang, (2023) *"Real-Time Semantic Segmentation of LiDAR Point Clouds on Edge Devices for Unmanned Systems"* **IEEE Transactions on Instrumentation and Measurement *72***: 1–11. DOI: 10.1109/TIM.2023.3292948.

[5] A. Piroli, V. Dallabetta, J. Kopp, D. Meissner, M. Walessa, and K. Dietmayer, (2024) *"Label-Efficient Semantic Segmentation of LiDAR Point Clouds in Adverse Weather Conditions"* **IEEE Robotics and Automation Letters *9***(6): 5575–5582. DOI: 10.1109/LRA.2024.3396099.

[6] Q. Li, Z. Wen, Z. Wu, S. Hu, N. Wang, Y. Li, X. Liu, and B. He, (2023) *"A Survey on Federated Learning Systems: Vision, Hype and Reality for Data Privacy and Protection"* **IEEE Transactions on Knowledge and Data Engineering *35***(4): 3347–3366. DOI: 10.1109/TKDE.2021.3124599.

[7] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al., (2021) *"Advances and Open Problems in Federated Learning"* **Foundations and Trends in Machine Learning *14***(1–2): 1–210. DOI: 10.1561/2200000083.

[8] H. Zhu, J. Xu, S. Liu, and Y. Jin, (2021) *"Federated Learning on Non-IID Data: A Survey"* **Neurocomputing *465***: 371–390. DOI: 10.1016/j.neucom.2021.07.098.

[9] X. Yin, Y. Zhu, and J. Hu, (2021) *"A Comprehensive Survey of Privacy-Preserving Federated Learning: A Taxonomy, Review, and Future Directions"* **ACM Computing Surveys *54***(6): Article 131, 1–36. DOI: 10.1145/3460427.

[10] D. C. Nguyen, M. Ding, Q.-V. Pham, P. N. Pathirana, L. B. Le, A. Seneviratne, J. Li, D. Niyato, and H. V. Poor, (2021) *"Federated Learning Meets Blockchain in Edge Computing: Opportunities and Challenges"* **IEEE Internet of Things Journal *8***(16): 12806–12825. DOI: 10.1109/JIOT.2021.3072611.

[11] P. Prakash, J. Ding, R. Chen, X. Qin, Q. Cui, Y. Guo, and M. Pan, (2022) *"IoT Device Friendly and Communication-Efficient Federated Learning via Joint Model Pruning and Quantization"* **IEEE Internet of Things Journal *9***(15): 13638–13650. DOI: 10.1109/JIOT.2022.3145865.

[12] A. El Ouadrhiri and A. Abdelhadi, (2022) *"Differential Privacy for Deep and Federated Learning: A Survey"* **IEEE Access *10***: 22359–22380. DOI: 10.1109/ACCESS.2022.3151670.

[13] D. R. Harris. "Leveraging Differential Privacy in Geospatial Analyses of Standardized Healthcare Data". In: *2020 IEEE International Conference on Big Data (Big Data)*. 2020, 3119–3122. DOI: 10.1109/BigData50022.2020.9378390.

[14] R. Cura, J. Perret, and N. Paparoditis, (2017) *"A Scalable and Multi-Purpose Point Cloud Server (PCS) for Easier and Faster Point Cloud Data Management and Processing"* **ISPRS Journal of Photogrammetry and Remote Sensing *127***: 39–56. DOI: 10.1016/j.isprsjprs.2016.06.012.

[15] Y. Lin, G. Vosselman, Y. Cao, and M. Y. Yang, (2021) *"Local and Global Encoder Network for Semantic Segmentation of Airborne Laser Scanning Point Clouds"* **ISPRS Journal of Photogrammetry and Remote Sensing *176***: 151–168. DOI: 10.1016/j.isprsjprs.2021.04.016.