

Article

Instance Segmentation of LiDAR Point Clouds with Local Perception and Channel Similarity

Xinmiao Du * and Xihong Wu

School of Intelligence Science and Technology, Peking University, Beijing 100080, China; wxh@cis.pku.edu.cn

* Correspondence: 2001111389@stu.pku.edu.cn

Highlights

What are the main findings?

- The LPM fuses multi-scale/multi-neighborhood point cloud features on a uniform BEV/polar coordinate grid.
- The ICCM calculates the attention weight matrix between feature channels, which adaptively enhances effective channel features and removes redundant channel features.

What is the implication of the main finding?

- The LPM is more effective against distance sparsity and occlusion in LiDAR than conventional FPN/CBAM that only does channel weighting or in-scale fusion.
- The ICCM makes representative features of point cloud extracted more effectively.

Abstract

Lidar point clouds are crucial for autonomous driving, but their sparsity and scale variations pose challenges for instance segmentation. In this paper, we propose LCPSNet, a Light Detection and Ranging (LiDAR) channel-aware point segmentation network designed to handle distance-dependent sparsity and scale variation in point clouds. A top-down FPN is adopted, where high-level features are progressively upsampled and fused with shallow layers. The fused features at 1/16, 1/8, and 1/4 are further aligned to a common BEV/polar grid and processed by the Local Perception Module (LPM), which applies cross-scale, position-dependent weighting to enhance intra-object coherence and suppress interference. The Inter-Channel Correlation Module (ICCM) employs ball queries to model spatial and channel correlations, computing an inter-channel similarity matrix to reduce redundancy and highlight valid features. Experiments on SemanticKITTI and Waymo show that LPM and ICCM effectively improve local feature refinement and global semantic consistency. LCPSNet achieves 70.9 PQ and 77.1 mIoU on SemanticKITTI, surpassing mainstream methods and reaching state-of-the-art performance.

Keywords: LiDAR; point clouds; instance segmentation; channel attention mechanism; inter-channel similarity; local perception



Academic Editor: Kevin Tansey

Received: 31 July 2025

Revised: 7 September 2025

Accepted: 15 September 2025

Published: 19 September 2025

Citation: Du, X.; Wu, X. Instance Segmentation of LiDAR Point Clouds with Local Perception and Channel Similarity. *Remote Sens.* **2025**, *17*, 3239. <https://doi.org/10.3390/rs17183239>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

With the rapid development of 3D sensing technology, the application of LiDAR has made the acquisition of point cloud data increasingly convenient and efficient [1], showing great application potential in many fields, including autonomous driving, robot navigation, building information modeling, and virtual reality [2–4]. The core processing tasks of

point clouds include classification and segmentation of point clouds [5,6]. Among these challenging tasks, instance segmentation task for point clouds is notable because it not only requires the algorithm to distinguish different categories but also to correctly distinguish each independent instance within the same category [7].

2D instance segmentation methods have been well-developed, and following the top-down paradigm, they can effectively segment instances in images. However, due to complex interactions and occlusions between 3D instances, pixel-based 2D segmentation methods cannot be directly applied to process unstructured and sparse 3D point clouds.

In recent years, alongside the rapid advancement of deep learning approaches [8–11], a series of instance segmentation networks grounded in deep learning have been put forward one after another. These methods can be broadly categorized into top-down-based and bottom-up-based [12,13]. The top-down method first identifies the region where the instance is located based on object detection methods, then filters the image features through non-maximum suppression, and refines them through mask segmentation. The bottom-up method distinguishes different instances through clustering and metric learning based on points with the same semantics in the semantic segmentation results. However, both methods still have obvious drawbacks. The limitation of the former lies in its dependence on the quality of bounding boxes, and the uneven distribution of point clouds will lead to unstable quality of bounding boxes, thus affecting the quality of instance segmentation. The latter has the problem of loss of local feature information after voxelization, and excessive post-processing steps will increase the amount of computation, while it is difficult to improve the segmentation performance in complex scenes.

Methods for event camera calibration and trajectory reconstruction offer novel approaches for instance segmentation in lidar point clouds [14]. Specifically, key frontiers and boundary points are extracted from locally sparse points, followed by instance-level aggregation using geometric consistency and physical models. Among these, a high-precision event camera calibration method based on a collimator [15] was designed to address the challenge of geometric calibration for event cameras during long-range and high-precision measurements. This method employs a collimator featuring a twinkling star pattern to first achieve linear parameter solutions through a spherical motion model. It then combines nonlinear optimization to further refine calibration results, thereby significantly enhancing parameter estimation accuracy and stability while maintaining advantages in high dynamic range and high temporal resolution.

To address issues such as insufficient capability in extracting local features and redundant repeated features in existing technologies, this paper innovatively proposes a Light Detection and Ranging (LiDAR) Channel-Aware Point Segmentation Network named LCPSNet. The network takes local perception and channel similarity as its core design concepts, with innovations reflected in the following three aspects.

(1) The Local Perception Module (LPM) performs location-dependent local perception and fusion of multi-scale/multi-neighborhood point cloud features on a uniform BEV/polar coordinate grid. The multi-scale local features are first spliced in the group dimension, and then a lightweight convolutional branch generates the global spatial saliency map and group-related components. Position-by-position modulation of local features at each scale is followed by fusion. This mechanism is more effective against distance sparsity and occlusion in LiDAR than conventional FPN/CBAM that only does channel weighting or in-scale fusion.

(2) The Inter-Channel Correlation Module (ICCM) is proposed, which innovatively introduces the channel attention mechanism in the clustering stage. By calculating the attention weight matrix between feature channels, ICCM adaptively enhances effective

channel features and removes redundant channel features. Thus, the representative features of the target instances in the point cloud are extracted more effectively.

(3) Aiming at the distance-dependent sparsity and drastic scale changes in LiDAR, a position-aligned multi-scale fusion is designed based on the top-down path of FPN. It not only up-samples and sums the corresponding shallow layers element by element but also uniformly resamples the fused features of $1/16$, $1/8$, $1/4$, and other layers to the same BEV/polar coordinate grid. The fusion features of $1/16$, $1/8$, $1/4$, and other layers are uniformly resampled to the same BEV/polar coordinate grid, and then the LPM performs cross-scale position-dependent weighting and modulation at the same spatial location.

2. Related Work

In 3D point clouds, 3D instance segmentation can accurately tell apart distinct categories as well as separate instances within the same category, thus providing key technical support to fields including autonomous driving and robot navigation. Research on 3D instance segmentation-related technologies promotes the innovation and performance improvement of point cloud processing algorithms. Deep learning-based instance segmentation network methods are mainly divided into top-down-based and bottom-up-based, which is shown in Figure 1.




Figure 1. Deep learning-based 3D instance segmentation methods.

2.1. Top-Down-Based Methods

In the top-down based methods, the system first generates potential object proposals in 3D space, that is, predicts the bounding boxes where each entity may exist. As one of the early 3D instance segmentation methods, SGPN segments point clouds with the help of point group networks and completes instance clustering by learning the similarity between points, laying the foundation for the development of subsequent 3D instance segmentation technologies [16]. However, the SGPN method exhibits limited performance when handling high-complexity data. In contrast, the 3D SIS method, which applies the Mask R-CNN architecture to instance segmentation [17,18], enhances instance segmentation precision by extracting image features from both RGB images and 3D point clouds. But the detection network and segmentation network of 3D-SIS do not share networks and weights, and the data processing process is relatively complex.

3D-BoNet emphasizes the significance of network depth in capturing complex features of 3D data. It performs instance segmentation by directly regressing 3D bounding boxes and predicting point-level masks for each instance, thereby reducing reliance on post-processing steps [19]. However, as point clouds are distributed unevenly, bounding box quality tends to be unstable, which affects the quality of instance segmentation. OccuSeg [20] is mainly designed to address the sparsity of point cloud data. Nevertheless, this method still faces many technical challenges in determining the position of bounding boxes, fitting bounding

boxes to cover target entities, and handling incompleteness in point cloud data. These challenges often lead to inaccurate positioning of bounding boxes, which in turn affects the quality of instance segmentation. TD3D conducts fully convolutional end-to-end training in a data-driven manner without relying on information of detected objects, reducing the burden of manual parameter tuning [21].

However, all top-down methods depend on the quality of bounding boxes. Given the uneven distribution of point clouds, the quality of these bounding boxes tends to be unstable, thereby impairing the performance of instance segmentation.

2.2. Bottom-Up-Based Methods

These methods are based on the results of point cloud semantic segmentation and complete segmentation by aggregating points with the same semantics into individual instances. PointNet++, built on the pioneering PointNet, addresses issues such as sparsity, disorder, and permutation invariance in point clouds by independently extracting features from all points and aggregating global features [22,23]. It also introduces adaptive sampling and points cloud partitioning strategies, extracting features from each region using PointNet before aggregating them. PartNet uses PointNet++ as the backbone network to predict semantic labels for each point and incorporates a segment detection network to achieve point cloud instance segmentation [24]. SSTNet effectively handles dynamic scenes through innovative methods, enhancing the ability to extract spatiotemporal features, especially excelling in processing moving objects [25]. PointGroup processes instance segmentation tasks in point clouds using dual-set point grouping, which has advantages in speed and significantly improves runtime performance through simplified strategies [26]. However, the success of this strategy largely depends on accurate semantic predictions, and even minor errors can lead to significant deviations in instance segmentation results.

Some new technologies such as HAIS have achieved high accuracy in 3D instance segmentation tasks with hierarchical network design and effective feature fusion strategies. However, their complex network structure increases computational requirements and shows weak segmentation accuracy in complex scenes [27]. DyCo3D adopts dynamic convolution technology, without a predefined proposal framework, and can directly generate instance masks through a small number of simple convolution layers [28]. This design dynamically generates convolution filters by responding to the spatial distribution of data and semantic predictions, thereby enhancing the ability to distinguish different instances. However, it relies heavily on centroid offset prediction, and due to insufficient local feature extraction capabilities, it is prone to incorrect grouping in dense object scenarios. ISBNNet includes an instance farthest point sampling and box-aware mechanism, which can improve point recall while enhancing features through geometric clues [29]. As a 3D point cloud segmentation approach based on a recursive slicing network, TD3D technology realizes point cloud semantic and instance-wise segmentation through slicing processing and the BiLSTM [1]. However, this method is limited by the processing capacity of BiLSTM and may encounter gradient vanishing problems in complex environments, which in turn limits its performance.

Although these novel bottom-up methods have achieved good performance, new problems have emerged. This type of method will lose some local feature information after voxelization. DyCo3D is more dependent on centroid offset prediction. Some methods add excessive post-processing steps to improve accuracy, such as the segment detection network introduced by PartNet, the small Transformer introduced by DyCo3D, and the box-aware mechanism of ISBNNet. Through these post-processing steps, the network can extract more feature information. However, while increasing the amount of computation, these methods cannot improve the instance segmentation performance for complex scenes,

resulting in more repeated or merged invalid 3D information in scale, and the segmentation efficiency and accuracy are not ideal.

Traditionally, 3D semantic segmentation, instance segmentation, and pan-segmentation have employed distinct model architectures, resulting in high training costs and ineffective utilization of implicit task relationships. OneFormer3D [30] aims to address all three segmentation tasks with a unified model, reducing training costs while improving overall performance. It integrates a Transformer decoder, semantic and instance query mechanisms, flexible pooling, and decoupled matching into a unified framework. This unified modeling approach achieves state-of-the-art performance across multiple 3D point cloud benchmarks. Test-Time Adaptive (TTA) for 3D point cloud segmentation, HGL [31] model geometric information at the point, object, and temporal levels. It generates reliable pseudo labels through neighborhood aggregation, corrects global consistency via category prototypes, and maintains feature stability with cross-frame constraints. By combining adaptive loss and temporal consistency regularization, HGL mitigates domain shifts without source domain data and achieves state-of-the-art performance on synthetic-to-real, cross-sensor, and continuous learning tasks with reduced adaptation overhead. To address the degradation of lidar point cloud quality caused by hardware and environmental interference, Zhao et al. [32] proposed a robust multi-task learning network for point cloud preprocessing. This method employs a shared PointNet encoder and incorporates three branches: denoising, single-object segmentation, and completion. The denoising branch combines geometric projection with dual-driver modeling, while the segmentation branch introduces an attention mechanism to enhance object-level segmentation. The overall framework effectively improves point cloud quality, thereby supporting more accurate downstream tasks.

3. Methods

3.1. The Overall Structure of the LiDAR Channel-Aware Point Segmentation Network

In this paper, a LiDAR Channel-Aware Point Segmentation Network (LCPSNet) is proposed. It is an end-to-end instance segmentation network composed of five main parts including the backbone network, prediction module, Local Perception Module (LPM), Inter-Channel Correlation Module (ICCM) and instance decoder. LCPSNet takes the original point cloud 3D coordinates and RGB color vectors as input. First, the point cloud is first voxelized into a raster representation and fed into the backbone network, where semantic and geometric features are extracted at scales such as $1/2$, $1/4$, $1/8$, $1/16$, and $1/32$. The top-down path up-samples ($\text{up} \times N$) the high-level features and sums them with the corresponding shallow features element by element to realize FPN multi-scale fusion. Next, the fused multi-scale features enter the LPM. Location-dependent spatial weights are generated and weighted and fused to the local features at different scales to enhance the response within the object and suppress cross-instance interference at the boundaries. These features are then passed to the Prediction Module containing the Offset Prediction Linear Layer and the Semantic Prediction Linear Layer. The instance decoder then clusters/correlates the semantic masks with the offset fields to generate instance masks. The immediately following ICCM further refines the instance and semantic results to output the final instance segmentation. The network structure of LCPSNet is shown in Figure 2. Its overall mechanism is multiscale fusion through a feature pyramid in the backbone, followed by local perceptual weighting using LPM, and then joint decoding using semantics and offsets. Finally, ICCM further refines the instance features through ball query and a three-dimensional attention mechanism.




Figure 2. The network structure of LCPSNet.

The backbone in LCPSNet is a top-down multiscale fusion, as shown in Figure 2. The point cloud is first voxelized to a fixed-resolution grid, and the backbone extracts pyramidal features at steps $1/2$, $1/4$, $1/8$, $1/16$, and $1/32$. Deeper features have stronger semantics and larger sensory fields and are able to cover context at longer distances and large targets, but with coarser spatial localization. Shallow features, on the other hand, retain geometric details for close range and small targets. The top-down path is upsampled by a factor of 2 at each level, starting from the coarsest $1/32$, to match the resolution of the next level of features. At each level, the channel is aligned using a 1×1 convolution and then summed element-by-element with the shallow features at the same level, thus injecting the deep global semantics into the shallow fine geometry. This process is performed sequentially at $1/16$, $1/8$, and $1/4$ to obtain a fused feature map with both context and boundary information. The up-sampling implementation can use nearest-neighbor/bilinear up-sampling followed by 3×3 convolution (range-view) or back-convolution (BEV) to mitigate aliasing and stabilize training. During the unified resampling of multi-scale features onto BEV/polar coordinate grids, both the sparsity of point clouds and the presence of voids must be addressed. To this end, a two-step strategy is adopted. First, missing regions are filled using nearest-neighbor or bilinear interpolation during scale alignment, ensuring spatial continuity and smoothness of features. Second, a 3×3 convolutional smoothing is applied to the interpolated feature map to mitigate interpolation artifacts and instability, thereby enhancing training convergence and generalization. In extremely sparse or completely void regions, the network compensates for missing information through global contextual features. These global features, extracted by deep networks, provide robust semantic guidance for distant objects or weak point cloud regions.

In order to allow the subsequent Local Perception Module to perform location-dependent weighting on the same spatial grid, the network resamples each of the fused features at $1/16$, $1/8$, and $1/4$ scales to a uniform target resolution. $8\times$ magnification from $1/16$, $4\times$ magnification from $1/8$, and $2\times$ magnification from $1/4$, and feeds them into the LPM along with the $1/2$ features that are themselves at the target resolution. Doing so allows position-by-position alignment and modulation of responses from different receptive fields on the same polar coordinate unit. Distant or large objects rely on the contextual enhancement brought by the deep layers, while close or small objects benefit from the fine

boundaries preserved by the shallow layers. This improves the separability of neighboring instances in sparse, occluded, and scale-variable LiDAR scenes.

3.2. Local Perception Module

Considering the sparse nature of LiDAR point cloud data itself, local feature information is inevitably lost during systematic processing. This lack of information cannot fully reflect the local details of the LiDAR point cloud data, which in turn may adversely affect the subsequent processing steps and reduce the accuracy of segmentation. Therefore, in this paper, the Local Perception Module (LPM) is designed to follow the backbone. This module enriches the contextual information of each point by extracting the structural information between neighboring points through local and global attention mechanism. Subsequently, enhancement of the features contained in each point is achieved by combining the local features with the global features through a dynamic fusion module. This approach better maintains the integrity of the local features when dealing with LiDAR point cloud tasks. This enhances the generalization ability of the model and the perception of multi-scale features and ultimately improves the segmentation performance of the network. The overall structure of LPM is shown in Figure 3.




Figure 3. The overall structure of Local Perception Module.

The core of LPM is to solve the problem of local detail loss caused by the sparsity of LiDAR point clouds. In fields such as autonomous driving and robotics, LiDAR is a key sensor for acquiring three-dimensional information about the environment. However, due to factors such as object proximity, surface material, and scanning angle, the density distribution of LiDAR point clouds is very uneven, and the point clouds can be very sparse, especially in distant or occluded areas. Specifically, let $P = \{p_i\}_{i=1}^N$ represent a point cloud consisting of N points, where the features of the points are $F = \{F_{p_i}\}_{i=1}^N$.

Specifically, during the local feature extraction stage, we employ 3×3 convolutions (stride = 1) to capture neighborhood spatial features, ensuring comprehensive coverage of the local receptive field. Concurrently, in the global attention generation process, 1×1 convolutions are used to map and apply nonlinear transformations to the features after global pooling. The tanh activation function is employed in both local and global

paths to enhance the nonlinear expressive power of features. This approach improves the discriminative capability between features at different scales during the fusion stage. Input features undergo concatenation and convolution to extract local representations, followed by global max pooling to obtain global contextual features. Next, two layers of 1×1 convolutions with tanh activation are used to generate global attention weights. These weights are spatially expanded and then position-wise added to local features. Finally, element-wise multiplication performs weighted fusion to yield the final LPM output.

The LPM structure mentioned above is shown in Algorithm 1. First, the feature F is input, and then concatenated to obtain F' . Then, convolve F' to obtain F_1 and enter the global spatial attention. Then, F_1 is maxpooled to obtain F_2 . After two *ConvTanh* processes, F_2'' is expanded and added to F_2 . Next, F_3 goes through *ConvTanh* to obtain the final output. Finally, output is expanded and multiplied to obtain the fused feature map, achieving local perception and fusion of features.

Algorithm 1. Local Perception Module

```

1  Input  $F$ 
2   $F' = \text{Concat}(F)$ 
3   $F_1 = \text{Conv}(F')$ 
4  Global Spatial Attention:
5     $F_2 = \text{MaxPooling}(F_1)$ 
6     $F_2' = \text{ConvTanh}(F_2)$ 
7     $F_2'' = \text{ConvTanh}(F_2')$ 
8     $F_3 = \text{Expand}(F_2'') + F_2$ 
9     $\text{output} = \text{ConvTanh}(\text{ConvTanh}(F_3))$ 
10 Return output
11  $\text{expanded\_weights} = \text{Expand}(\text{output})$ 
12  $\text{fused} = \text{WiseProduct}(F', \text{expanded\_weights})$ 
13 Output fused

```

3.3. Inter-Channel Correlation Module Based on Channel Similarity

For traditional LiDAR point cloud clustering module works by first dividing the point clouds in the offset space into subsets based on the semantic labels of the LiDAR point clouds, which contain only point clouds with the same semantics. The subsets are further grouped using a clustering algorithm. Ultimately, each LiDAR point cloud subset obtained after clustering is considered as an instance. That is, the points in each LiDAR point cloud subset have the same instance label.

However, some post-processing operations aimed at improving the segmentation accuracy result in getting more duplicates or merging invalid 3D information when clustering. This results in increased network computation and poor segmentation network performance in complex scenarios. Therefore, a new clustering module, Inter-Channel Correlation Module (ICCM), is proposed in this paper. The structure of ICCM is shown in Figure 4. LiDAR point cloud data is usually high-dimensional, sparse, and noisy, which makes instance segmentation very challenging. The ICCM weights the temporal, spatial, and channel dimensions to help the LCPSNet model focus on key regions and features, thus improving the accuracy and robustness of segmentation.




Figure 4. The structure of the ICCM.

As shown in Figure 4, the spatial feature attention module is first utilized to process the LiDAR point cloud data because point cloud data naturally has a spatial structure. Each point in the laser point cloud has corresponding spatial coordinates, and by weighting these spatial features, it can help the model focus on the most important spatial regions. The sphere query algorithm is first used to identify points within a specified spherical neighborhood. Then, the acquired point set context information is encoded and processed and converted into preliminary instance features. The temporal attention module is then utilized to process the temporal data. The successive sweeps of the LiDAR consist of multiple time steps, and the temporal attention module helps the model to learn which time points are important at different time steps. Instance features that incorporate temporal information are output by calculating Q , K , and V to weight the temporal dimension of the processed point cloud data. Finally, the channel attention weights are assigned by calculating the channel similarity of the instance features to extract the effective channel features and release the redundant ones.

Specifically, for each point selected by the farthest point sampling, a sphere with a unit ball of R is defined with its center. Subsequently, the individual points of the original point cloud are filtered to cull out all points within that sphere that are not more than R away from the center point, which are considered as neighborhood points. If the number of point clouds within a cluster of neighboring points is more than a set threshold B , the closest B neighboring points are selected as local features in order of distance. If the number of neighboring points is less than B , repeat the filling process for these points to ensure that the number reaches B . These B neighboring points are collected as local point cloud features $F_{local} \in \mathbb{R}^{B \times D}$, where D denotes the number of channels. The relative coordinates between the sampled points and the neighboring points are also calculated and normalized by the neighborhood radius r to form the local coordinates P_{local} . The initial instance features (Feature map 3) are finalized by weighted averaging of local features and local coordinates, as shown in Equation (1).

$$Feature\ map\ 3 = Feature\ map\ 1 \times \frac{F_{local} + P_{local}}{2} \quad (1)$$

The temporal attention module computes the temporal attention of the input feature graph. It computes the *Query*, *Key* and *Value* tensor and computes the attention weights

using the dot product attention mechanism. The *Query*, *Key*, and *Value* tensor are shown in Equations (2)–(4).

$$Q = W_q \times \text{Feature map 3} \quad (2)$$

$$K = W_k \times \text{Feature map 3} \quad (3)$$

$$V = W_v \times \text{Feature map 3} \quad (4)$$

Then, the attention fraction is obtained as shown in Equation (5).

$$\text{Attention} = \text{softmax}(QK^T) \quad (5)$$

The values are weighted and summed using the attention score, as shown in Equation (6). The final feature map 4 to be input into the channel attention is the instance feature F_{ins} in the LiDAR point cloud instance segmentation task.

$$F_{ins} = \text{Feature map 4} = \gamma \times (V \times \text{Attention}^T) + \text{Feature map 3} \quad (6)$$

where W_q , W_k , and W_v are the weights of the convolutional layers used for query, key, and value computation, γ is the learnable parameter, and softmax computes softmax along the time dimension.

A transpose operation is performed on F_{ins} to change the shape. Two MLPs are then used to manipulate each channel's feature vector $C_i \in \mathbb{R}^N$ with $C_i = \mathbb{R}^N$, $F_{ins} = [C_1, C_2, \dots, C_D]$ to encode the contextual information of each channel feature vector. Thereby, the original B points are randomly replaced with fewer B' points, where $B' = B/r_{ratio}$; $r_{ratio} > 1$ is the scale parameter, and the value is taken as 4 here. Compared with obtaining all points or points extracted under other constraints, this operation can effectively reduce the number of features at the origin and fully preserve the original information, as shown in Equations (7) and (8).

$$q_i = \text{MLP}_q(C_i) \quad (7)$$

$$k_i = \text{MLP}_k(C_i) \quad (8)$$

where MLP_q and MLP_k are the two MLPs operating on the query matrix and the key matrix. Synthesizing the above yields Equations (9) and (10).

$$Q_{B' \times D} = [q_1, q_2, \dots, q_D] \quad (9)$$

$$K_{B' \times D} = [k_1, k_2, \dots, k_D] \quad (10)$$

The corresponding channel similarity matrix is subsequently computed by dot-producing the transpose of the query matrix with the key matrix, which can be expressed as Equation (11).

$$S_{D \times D} = Q^T K \quad (11)$$

where Q is the query matrix and K is the key matrix.

S_{ij} is the similarity between the i th channel and the j th channel of the F_{ins} feature map and an element of the channel similarity matrix. The obtained similarity matrix (S) between the channels is maximally pooled to obtain the column direction expansion, which results in a weight matrix of the same size as the original matrix. By subtracting the weight matrix from the similarity matrix, similar features are retained while redundant features are removed. The refined version of the weight matrix $D \times D$ is derived by multiplying any value with the value matrix via Softmax's activation function. This design ensures that

the weights optimized for the other channels are de-redundant between each channel. The accumulation of invalid information in the original features is effectively circumvented, thus enhancing the uniqueness and accuracy of the overall feature expression. The process of calculating the weight matrix can be expressed as Equation (12).

$$A_{D \times D} = \text{Softmax} \left\{ \underset{1 \rightarrow D}{\text{expand}} \left[\underset{1 \rightarrow D}{\text{Maxpool}}(S) \right] - S \right\} \quad (12)$$

Based on the weight matrix, the instance feature channel information is refined by performing a weighted sum over all channels, and another MLP operation is applied to obtain the weight matrix V . Specifically, this can be expressed as Equations (13) and (14).

$$V_{B \times D} = [v_1, v_2, \dots, v_D] \quad (13)$$

$$v_i = \text{MLP}(C_i); v_i \in R^B \quad (14)$$

Multiplying the value matrix V with the weight matrix achieves the same effect as refining the instance feature channel. In addition, the training is simplified by using residual concatenation with learnable weights α . The processed instance features are shown in Equation (15).

$$F'_{ins} = F_{ins} + \alpha \times VA \quad (15)$$

During the calculation of the channel similarity matrix in ICCM, numerical stability issues were specifically addressed. Given that Softmax is prone to numerical instability, vanishing gradients, or exploding gradients under high-dimensional features, this paper employs the following strategies to ensure stability. First, normalization is applied. Before computing the channel similarity matrix $S = Q^T K$, input features undergo L2 regularization. This ensures consistent vector norm lengths across channels, thereby preventing numerical overflow caused by excessively large inner product values. Then, during the Softmax calculation, a scaling factor of $\frac{1}{\sqrt{d}}$ is introduced, where d represents the channel feature dimension. This aligns with the Scaled Dot-Product Attention method commonly used in Transformers, effectively mitigating the issue of excessively large gradients in high-dimensional features. In the implementation, numerically stable Softmax is employed to avoid the sensitivity of exponential functions to large values, ensuring the stability of gradient computations. During the final feature update, a residual structure is employed, as shown in Equation (15). Residual connections provide a direct path during gradient propagation, further mitigating the vanishing gradient problem.

3.4. Cross Entropy Loss Function

In LiDAR point cloud instance segmentation, in order to supervise the point-by-point semantic prediction, a weighted multi-class cross entropy loss $\mathcal{L}_{CE}$ is used for the set of valid points, as shown in Equation (16). The predicted probability of each point is compared with its true semantic label and averaged as semantic loss. Invalid/filler points are ignored during training, and weights can be set by category frequency to mitigate category imbalance. If the labeling is noisy, light label smoothing can be added.

$$\mathcal{L}_{CE} = -\frac{1}{|\mathcal{V}|} \sum_{i \in \mathcal{V}} \sum_{c=1}^C w_c y_{i,c} \log p_{i,c} \quad (16)$$

where $y_{i,c}$ is the one-hot truth value, w_c is the optional category weight to mitigate the class imbalance, and $p_{i,c}$ is the probability of each point to the C semantic categories obtained by picking up MLP and softmax at the output F'_{ins} .

4. Experiments

4.1. Datasets and Metrics

The SemanticKITTI dataset [33] is a large-scale dataset constructed on the original KITTI dataset for 3D LiDAR point cloud segmentation. Its core contribution is to provide dense, point-by-point semantic and instance annotations for dynamic outdoor scenes for KITTI. Based on data from the KITTI Odometry Benchmark, it provides point-by-point annotations for all 22 of these sequences (sequences 00 to 21), totaling over 43,000 LiDAR scans in length. This means that it is a dynamic dataset containing temporal information. Unlike the sparse bounding box of KITTI, SemanticKITTI provides labels for all points in each scan, truly enabling a comprehensive understanding of the entire 3D space. It contains a total of 28 categories, 19 of which are used for evaluation, covering almost all elements in an autonomous driving scenario. The examples of SemanticKITTI dataset are shown in Figure 5.



Figure 5. The examples of SemanticKITTI dataset.

The Waymo dataset [34] is a large-scale LiDAR point-cloud benchmark for autonomous driving. Data were collected by a fleet equipped with five LiDAR sensors and five high-resolution RGB cameras, operating day and night in cities such as San Francisco. The dataset comprises 1150 scenes, each 20 s long, yielding 230,000 frames in total. Range images and RGB images are synchronized and released at 10 Hz. Among these scenes, 1000 are designated for training and 159 for testing. The examples of Waymo dataset are shown in Figure 6.

Instance segmentation is essentially a subset of panoramic segmentation. Using metrics for panoramic segmentation provides a more comprehensive view by evaluating not only the performance of instance segmentation but also the semantic segmentation capability that underlies it. Accordingly, we evaluate semantic segmentation using mean Intersection over Union (mIoU). This standard metric quantifies the overlap between predictions and ground truth. For each class we compute IoU as intersection over union and then average the classwise IoUs. As a point level measure, mIoU captures semantic classification accuracy and is suitable when the evaluation focuses only on semantic information.



Figure 6. The examples of Waymo dataset.

PQ metrics were initially used for image panorama segmentation evaluation and have been gradually introduced to the 3D point cloud domain in recent years. PQ integrates the recognition quality (RQ) and segmentation quality (SQ) of object instances and is able to evaluate the instance recognition ability and segmentation accuracy of the model simultaneously. Its definition is shown in Equation (17).

$$PQ = \frac{\sum_{TP} IoU}{|TP| + \frac{1}{2}|FN| + \frac{1}{2}|FP|} \quad (17)$$

where TP denotes the number of true positive cases, and FN and FP are the number of false negative and false positive cases, respectively. The numerator part denotes the sum of IoUs of all TPs, and the denominator is the weighted sum of detection and segmentation errors.

4.2. Comparison Experiments

To fully evaluate the performance of LCPSNet as described in this paper in the LiDAR point cloud semantic segmentation challenge, we compare it with a number of comparable techniques. The results are displayed in Table 1. Since top-down methods are not mainstream in outdoor panorama segmentation, they are difficult to handle backgrounds without fixed shapes and are not competitive. Therefore, in this paper, we only compare with bottom-up methods.

Table 1. Comparison results between LCPSNet and other methods on SemanticKITTI dataset.

Method	PQ	mIoU (%)	FPS	Params (M)
PointNet [21]	17.5	18.2	120	3.5
PointNet++ [20]	20.8	23.4	84.8	1.7
SSCNS [25]	35.2	37.9	25.2	16
PolarNet [35]	54.3	55.7	16.9	13.6
PointGroup [24]	41.7	42.5	12.4	22.8
Cylinder3D [36]	66.8	68.9	11.5	25.1
AF2S3Net [37]	64.9	69.7	9	31.4
RangeFormer [38]	64.1	73.6	10.2	28.2
SDSeg3D [39]	62.6	70.4	8.5	25
SpAtten [40]	70.5	76.8	6	40
LCPSNet (Ours)	70.9	77.1	13.5	18.2

Early methods such as PointNet and PointNet++ only achieve PQs of 17.5 and 20.8 and mIoUs of 18.2 and 23.4, respectively. This indicates that their structures are difficult to model effective spatial context relations when facing complex point cloud geometries, resulting in

severely limited segmentation performance. Some subsequent improved methods, such as SSCNS, PolarNet, and PointGroup, introduce sparse convolution, polar coordinate transformation, or clustering mechanisms in the feature extraction strategy, which leads to improved performance. For example, PolarNet achieves a PQ of 54.3 and an mIoU of 55.7, demonstrating progress in spatial structure alignment. However, PointGroup achieves a PQ of 41.7 with an mIoU of 42.5, reflecting its shortcomings in instance boundary modeling.

Mainstream high-performance methods such as Cylinder3D, AF2S3Net, RangeFormer, and SDSeg3D achieve high scores in both metrics. Among them, Cylinder3D reaches 66.8 in PQ and 68.9 in mIoU, which reflects the modeling advantages of columnar structures in dealing with dense voxels and ground targets. AF2S3Net has a PQ of 64.9 and an mIoU of 69.7, relying on its ability to model semantic spatial relations to achieve a balanced performance. RangeFormer, with the introduction of the Transformer architecture, achieves a higher result of 73.6 in mIoU, but with a PQ of 64.1, suggesting that its global semantic enhancement is effective but still needs to be optimized for instance, matching at the target level. The current best performer is SpAtten with a PQ of 70.5 and mIoU of 76.8.

The LCPSNet proposed in this paper achieves the current best level in both PQ and mIoU, which are 70.9 and 77.1, respectively. Compared with the existing methods, LCPSNet performs better in both semantic consistency and instance separation ability. In conclusion, LCPSNet achieves the simultaneous improvement of semantic accuracy and panoramic quality in the LiDAR point cloud semantic segmentation task, which verifies its excellent performance in the real scene perception task.

Additionally, as shown in Figure 1, lightweight methods such as PointNet and PointNet++ are significantly faster but lag far behind subsequent complex methods in performance. This indicates they sacrifice accuracy for real-time capability. Mainstream methods exhibit lower speeds, with FPS ranging from 8 to 17 for models like RangeFormer and PointGroup. This demonstrates that inference speed decreases significantly as model complexity increases. LCPSNet occupies a middle ground between lightweight and heavyweight models in terms of speed. While it does not match the PointNet series, it is over twice as fast as SpAtten while maintaining high accuracy. This indicates a successful balance between speed and precision.

4.3. Ablation Experiments

In this paper, ablation experiments are conducted on SemanticKITTI dataset and Waymo dataset, respectively, to verify the effectiveness of each module in LCPSNet. Since the LiDAR point cloud possesses natural defects such as sparsity, inhomogeneous density, and viewpoint dependence, it puts higher requirements on the model's local perception ability and global semantic modeling ability. Therefore, the design of the modules should not only improve the point-level semantic prediction accuracy but also enhance the instance differentiation and structural consistency modeling ability.

First, mIoU and PQ on SemanticKITTI dataset are shown in Table 2, and IoU on each category is shown in Table 3, where mIoU is a weighted average of the weights by category, which is not equal to the simple average.

Table 2. Ablation experiments of Different Module Combinations on SemanticKITTI Dataset.

Combination	LPM	ICCM	mIoU (%)	PQ
Baseline	×	×	69.9	61.4
LPM	✓	×	72.4	64.8
ICCM	×	✓	72.9	65.2
LCPSNet (Ours)	✓	✓	77.1	70.9

Table 3. Ablation experiments of different modules under different categories in the SemanticKITTI dataset.

Category	Baseline	LPM	ICCM	LCPSNet (Ours)
car	94.3	94.7	95	98.2
bicy	68.3	69.3	68.5	72.4
moto	70.8	72.3	72.8	75.7
truc	59.1	60.8	60.2	63.9
o.veh	69.4	71	71.8	74.5
ped	73.7	75.8	76.1	79.3
b.list	70.5	71.3	71.5	75.2
m.list	56.1	58.1	58.1	60.9
road	88.2	89.2	89.2	92.9
park	69.9	72	71.9	74.4
walk	75.6	76.4	76.8	79.7
o.gro	42.5	44	43.8	46.5
build	89.9	90.1	91	93.8
fenc	67.4	69.7	69.5	72.9
veg	83	84.8	85.3	87.8
trun	72.4	73	73.6	76.6
terr	68.1	70.3	69.6	73.4
pole	63.9	65.1	66.1	68.7
sign	64.9	65.6	65.9	68.9
mIoU	69.9	72.4	72.9	77.1

From Table 2, it can be seen that when no module is introduced, the overall performance of Baseline is relatively weak, the mIoU is only 69.9%, and the PQ value is 61.4%, which indicates that there is much room for improvement of the model in terms of semantic recognition and instance matching accuracy. After the introduction of the LPM, the mIoU increases to 72.4%, and the PQ increases to 64.8%. This shows that the module effectively enhances the semantic discrimination of local spatial regions in the LiDAR point cloud, and LPM especially improves the segmentation accuracy of dynamic object categories. In particular, LPM improves the segmentation accuracy of dynamic object categories, such as ped from 73.7 to 75.8%, walk from 75.6 to 76.4%, and truc from 59.1 to 60.8%. This validates its ability to recognize fine-grained semantic boundaries.

Furthermore, when the ICCM is introduced alone, the mIoU is 72.9% and the PQ reaches 65.2%, which achieves a significant improvement on several structural or low sample categories. ICCM enhances the model's ability to model structurally complex regions by capturing long-range dependencies and reinforcing semantic consistency among point clouds. For example, m.list increases from 56.1 to 58.1%, terr from 68.1 to 69.6%, and sign from 64.9 to 65.9%. This shows that ICCM has a significant advantage in modeling global contextual semantic consistency, which is especially suitable for improving the recognition accuracy of small targets and fuzzy boundary regions.

When the two modules are jointly introduced to construct the full LCPSNet model, the performance is optimized, with the mIoU reaching 77.1% and the PQ improving to 70.9%, which is 7.2% and 9.5% compared to the baseline, respectively. In Table 3, it can be seen that the model achieves the highest segmentation accuracy in most of the classes. For example, it is significantly ahead in the ground static categories such as build 93.8%, veg 87.8%, and trun 76.6%, as well as the complex m.list 60.9% and walk 79.7% categories. This indicates that the synergy of the two modules effectively mitigates the common problems of uneven density, semantic boundary blurring, and category imbalance in LiDAR point clouds.

It is worth emphasizing that the significant enhancement of the PQ metric, as an important indicator for panoramic segmentation evaluation, reflects the combined enhance-

ment of the model in terms of instance separation and prediction accuracy. Unlike mIoU, which focuses more on point-level classification, PQ also considers the match between real and predicted instances, including IoU and target overlap determination. Its significant enhancement thus indicates that LCPSNet not only improves the semantic recognition capability but also strengthens the consistent modeling of target boundary and instance structure when dealing with point cloud panorama segmentation tasks.

Both LPM and ICCM provide complementary capabilities in different dimensions of point clouds. The former focuses on local spatial detail enhancement, while the latter strengthens cross-regional semantic consistency. Together, they can significantly improve the segmentation robustness and generalization ability of the model in LiDAR point cloud scenarios with sparse, heterogeneous, and multi-scale structures.

The comparison results of different loss functions on SemanticKITTI are shown in Table 4. It can be observed that weighted cross-entropy (CE) achieves the best performance in both PQ at 70.9% and mIoU at 77.1%, indicating its robustness and effectiveness in point cloud instance segmentation. In contrast, while IoU loss optimizes boundaries, it leads to a decrease in overall accuracy, with PQ of 69.8% and mIoU of 76.2%. The mAcc loss prioritizes category balance but performs worst on long-tail data distributions, yielding PQ of 68.9% and mIoU of 75.6%. This indicates that CE is the most suitable loss function as the primary optimization objective, while IoU and mAcc are better suited as auxiliary losses for joint optimization.

Table 4. Comparison of Different Loss Functions on SemanticKITTI Dataset.

Loss Function	PQ (%)	mIoU (%)
Cross-Entropy (CE)	70.9	77.1
IoU	69.8	76.2
mAcc	68.9	75.6

To further validate the robustness and generalization ability of the proposed module under different LiDAR point cloud scenarios, we conduct the same ablation experiments on the Waymo Open Dataset with a larger field of view, higher density, and more complex traffic environments, and the results are shown in Table 5. Compared with SemanticKITTI, the Waymo dataset contains richer urban traffic scenarios, sampled point clouds under different time and weather conditions, and is more challenging. Therefore, this experiment better reflects the adaptability of the module under different point cloud distribution conditions.

Table 5. Ablation Study of Different Module Combinations on Waymo Dataset.

Combination	LPM	ICCM	mIoU (%)
Baseline	×	×	62.7
LPM	✓	×	66.9
ICCM	×	✓	67.5
LCPSNet (Ours)	✓	✓	70.4

From the results in the table, it can be seen that the mIoU of the base model on this dataset is 62.7%, which is slightly lower than that on SemanticKITTI, reflecting the challenge of target class distribution and environment complexity on the performance of the base model in the Waymo dataset. When the LPM is introduced, the mIoU increases to 66.9%, which is 4.2% higher than that of the baseline. This indicates that LPM can effectively enhance the local feature expression and improve the ability to discriminate the detailed targets when facing the point cloud data with both dense scenes and sparse structures.

Similarly, when the ICCM is introduced alone, mIoU improves to 67.5%, an increase of 4.8%. This reaffirms the effectiveness of ICCM in semantic consistency modeling, especially in the Waymo dataset. Since there are often fuzzy boundaries and semantic overlaps between categories such as lanes, buildings, etc., ICCM can effectively align the contextual semantics to improve prediction consistency. In deep learning, different feature channels often represent different information. The channel attention mechanism allows the model to dynamically learn and adjust the importance weights of different channels. By calculating the similarity between channels, the model can focus on those feature channels that are most useful for the task at hand, while suppressing those that are redundant or noisy. In the task of instance segmentation of LiDAR point clouds, this means that the model can focus more intelligently on the key features that distinguish different objects.

In our full model LCPSNet, LPM is used in conjunction with ICCM, and the mIoU is finally increased to 70.4%, which is 7.7% higher than the base model and further increased by 2.9–3.5% on top of the two modules. This shows that the two modules still have strong complementarity in the more challenging point cloud scenarios of Waymo, and the joint modeling can capture the local and global features of the point cloud more effectively, which significantly improves the overall performance of the model.

Overall, the experimental results on the Waymo dataset not only further validate the effectiveness of LPM and ICCM in different scenarios. It also shows that the method in this paper has good cross-dataset migration capability and generalization performance and is suitable for real LiDAR semantic sensing tasks in larger-scale and more complex environments.

4.4. Hyperparameter Sensitivity Analysis

Figure 7 shows the performance trends of the three LiDAR point cloud segmentation methods on the Waymo and SemanticKITTI datasets for different ball query radius settings. The ball query radius is an important parameter for constructing local neighborhoods in point cloud feature extraction, and its size directly affects the geometric perceptual range and semantic representation effect of the aggregated region. This experiment systematically analyzes the sensitivity and adaptability of OccuSeg, ISBNNet, and LCPSNet methods to this hyperparameter on the Waymo dataset by adjusting the value of the radius r .




Figure 7. Influence of ball query radius on mAP: (a) on the Waymo dataset; (b) on the SemanticKITTI dataset.

As can be seen in Figure 7 and Table 6, traditional clustering-based approaches, such as OccuSeg and ISBNNet, are highly dependent on aggregation size. When the radius is small, these methods can only capture local details, and it is difficult to obtain the complete structure and context information of the target, especially when dealing with large targets

or complex shapes. When the radius is too large, it is easy to introduce cross-instance interference, resulting in feature mixing and boundary blurring, which affects the accuracy of instance recognition. Therefore, the performance of such methods fluctuates greatly and lacks stability.

Table 6. Effect of different ball radius on mAP for different methods on Waymo dataset and SemanticKITTI dataset.

	r (cm)	1	2	3	4	5
Waymo	OccuSeg	77.79	85.02	87.81	79.37	76.93
	ISBNet	82.07	82.55	88.5	87.24	86.47
	LCPSNet	84.78	87.64	87.25	88.61	87.39
	r (cm)	1	2	3	4	5
SemanticKITTI	OccuSeg	87.29	90.26	92.82	89.26	89.89
	ISBNet	89.24	92.47	93.08	92.84	91.26
	LCPSNet	90.78	92.16	93.25	93.61	91.74

In contrast, LCPSNet has stronger scale adaptation capability. By introducing local feature enhancement and a semantic consistency modeling mechanism, the method is able to maintain robust expression ability under different ball query ranges, avoiding performance jitter due to radius changes. In the whole test range, the mAP of LCPSNet on the Waymo dataset is consistently above 85%, and the mAP of LCPSNet on the SemanticKITTI dataset is consistently above 91%. Meanwhile, the fluctuation of LCPSNet's mAP with sphere radius is minimized on both datasets, which fully proves its good structural generality and practical deployment stability.

Ball query radius, as a key structural parameter in the point cloud aggregation process, has a significant impact on segmentation performance. Traditional clustering-based methods are prone to unstable performance under the change in this parameter, indicating their dependence on the modeling of feature space structure. LCPSNet, on the other hand, effectively mitigates this problem through the multi-scale fusion and semantic guidance mechanism, reflecting stronger adaptability and better LiDAR point cloud understanding.

4.5. Visualization Analysis

Figures 8 and 9 show the visualization results on SemanticKITTI and Waymo dataset, respectively. In perception systems for autonomous driving, LiDAR is known for its ability to provide accurate 3D spatial information. However, the point clouds it generates are inherently sparse and lack color and texture details. These two characteristics aptly define the need for the complementary LPM and ICCM. In the SemanticKITTI dataset, LCPSNet accurately distinguishes road regions and their boundaries, maintaining stable contour recognition even at intersections and curves. This demonstrates LPM's advantage in capturing sparse point cloud details. Visualization results show that LCPSNet effectively differentiates between large vehicles and surrounding dynamic objects while preserving boundary continuity, reducing discontinuities caused by point cloud sparsity. Compared to traditional methods, the target contours are more complete, which is closely related to the introduction of ICCM. In sparse point cloud regions at a distance, LCPSNet still maintains high segmentation consistency. Particularly in the right-hand scene of Figure 8, distant buildings and road edges can still be correctly annotated, demonstrating the method's excellent performance in long-range perception tasks. Figure 9 demonstrates the visualization results of LCPSNet on Waymo dataset. Despite differences in scene resolution and point cloud density compared to SemanticKITTI, the model accurately distinguishes large vehicles from their surroundings. This indicates that LCPSNet exhibits strong robustness

and generalization capabilities across different datasets. Overall, LCPSNet demonstrates advantages in visualization results, including clear boundaries, strong semantic consistency, and low sensitivity to sparse point clouds. These effects further validate the effectiveness of the LPM and ICCM in enhancing semantic segmentation quality.




Figure 8. Visualization results for LCPSNet on SemanticKITTI dataset.




Figure 9. Visualization results for LCPSNet on Waymo dataset.

LPM provides the geometric skeleton for the network. It ensures that the underlying spatial structure of even poorly characterized objects due to distance or few scanning points can be better captured and understood by the network by enhancing the learning of local geometric features. For example, when a car whose color blends in with the environment (with blurred image features) but whose LiDAR point cloud clearly outlines the 3D structure, LPM becomes the key to identifying the object. It enhances the robustness of LCPSNet to the geometric properties of the LiDAR data itself.

ICCM, on the other hand, provides perceptual focus. It utilizes the rich texture and color information of the image to direct the network's attention to those details that are easily overlooked in the LiDAR point cloud. In particular, it directs the network's attention to key details that are easily overlooked in sparse point clouds, such as the contours of pedestrians or the patterns of traffic signs.

Therefore, LCPSNet is not just a simple multimodal fusion network but a deeply complementary system. It combines the robustness originated from LPM for LiDAR geometric information and the acuity originated from ICCM for image perception information. This enables the most accurate judgment to be made for optimal detection performance in difficult scenarios where either LiDAR data is dominant or image data is dominant.

5. Conclusions

Aiming at the problems of information loss due to voxelization processing and invalid feature redundancy due to post-processing operations after feature enhancement in the LiDAR point cloud instance segmentation network, this paper designs and implements a LiDAR point cloud instance segmentation method (LCPSNet) based on local perception and channel similarity. Based on the traditional FPN that only does prediction or simple cascading in each layer, we propose a location-consistent multi-scale fusion for LiDAR.

The Local Perception Module (LPM) and Global Spatial Attention (GSA) are applied after the backbone for LiDAR point cloud instance segmentation to address cross-scale aliasing and adjacent instance coalescence issues caused by distance sparsity and scale variations. The point cloud is first projected into the range view, and multi-scale features are extracted via FPN. GSA generates a global spatial saliency map and produces group-wise correlation weights, while LPM broadcasts these weights to each channel and fuses them with local features before feeding them into the semantic and offset heads. Both modules perform position-aware cross-scale modulation on a unified polar coordinate grid, effectively integrating deep semantic information with shallow boundary information to significantly enhance the separability of adjacent instances. Additionally, the network proposes an Inter-Channel Correlation Module (ICCM) based on channel similarity, simultaneously modeling the temporal and channel characteristics of point clouds while performing explicit redundancy removal. This module utilizes local geometry obtained through spherical queries to adaptively select key information across multiple frames via temporal attention. It then applies channel similarity-weighted weighting to suppress redundancy and highlight differences, combined with learnable channel compression and residual refinement. Compared to methods relying on post-processing clustering, ICCM achieves higher accuracy, robustness, and efficiency in sparse, noisy, and occlusion-prone scenarios.

Experimental results demonstrate that LCPSNet outperforms numerous existing classical and state-of-the-art methods on both the SemanticKITTI and Waymo datasets, particularly when handling dense and complex outdoor traffic scenes. On the SemanticKITTI dataset, it achieves a PQ of 70.9 and an mIoU of 77.1, attaining state-of-the-art performance. Furthermore, the proposed network exhibits greater stability and universality against parameter variations while enhancing local perception and improving segmentation performance.

Given the strong modularity and versatility of LPM's position-wise weighting and ICCM's channel-correlation modeling, they are expected to adapt well to lightweight backbone networks and demonstrate advantages in edge computing or real-time application scenarios. In future work, we will further explore the performance of LPM and ICCM on lightweight backbone networks and investigate how to reduce computational overhead and parameter size while maintaining high accuracy, thereby enhancing the universality and practicality of these methods in real-world applications. Additionally, given that the core design principles of LCPSNet lie in Local Perception Modules (LPM) and Inter-Channel Correlation Modeling (ICCM), the method itself exhibits strong modularity and transferability. Consequently, we anticipate this framework will also demonstrate advantages in tasks such as indoor point cloud segmentation on datasets like ScanNet

and S3DIS, and point cloud perception for robotic navigation, particularly when handling sparse points, heavily occluded environments, or complex structures. Moving forward, we plan to conduct further experiments in these scenarios to validate LCPSNet’s cross-scenario generalization capabilities and explore its potential applications in multimodal fusion, such as integrating event cameras with LiDAR, and in real-time robot navigation.

Author Contributions: Conceptualization, X.D.; Methodology, X.D.; Software, X.D.; Validation, X.D.; Formal analysis, X.W.; Investigation, X.W.; Resources, X.W.; Data curation, X.D.; Writing—original draft, X.D.; Writing—review & editing, X.W.; Supervision, X.W. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Data are contained within the article.

Conflicts of Interest: The authors declare no conflict of interest.

References

- Huang, X.S.; Mei, G.F.; Zhang, J.; Abbas, R. A comprehensive survey on point cloud registration. *arXiv* **2021**, arXiv:2103.02690. Available online: <https://arxiv.org/abs/2103.02690v2> (accessed on 12 December 2024).
- Zeng, Y.H.; Jiang, C.H.; Mao, J.G.; Han, J.H.; Ye, C.Q.; Huang, Q.Q. CLIP2: Contrastive language-image-point pretraining from real world point cloud data. In Proceedings of the 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Vancouver, BC, Canada, 17–24 June 2023; IEEE Press: New York, NY, USA, 2023; pp. 15244–15253.
- Marinos, V.; Farmakis, I.; Chatzitheodosiou, T.; Papouli, D.; Theodoropoulos, T.; Athanasoulis, D.; Kalavria, E. Engineering Geological Mapping for the Preservation of Ancient Underground Quarries via a VR Application. *Remote Sens.* **2025**, *17*, 544. [\[CrossRef\]](#)
- Qian, R.; Lai, X.; Li, X.R. 3D object detection for autonomous driving: A survey. *Pattern Recognit.* **2022**, *130*, 108796. [\[CrossRef\]](#)
- Lee, S.; Lim, H.; Myung, H. Patchwork: Fast and robust ground segmentation solving partial under-segmentation using 3D point cloud. In Proceedings of the 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Kyoto, Japan, 23–27 October 2022; IEEE Press: New York, NY, USA, 2022; pp. 13276–13283.
- Xiao, A.R.; Yang, X.F.; Lu, S.J.; Guan, D.; Huang, J. FPS-Net: A convolutional fusion network for large-scale LiDAR point cloud segmentation. *ISPRS J. Photogramm. Remote Sens.* **2021**, *176*, 237–249. [\[CrossRef\]](#)
- Hafiz, A.M.; Bhat, G.M. A survey on instance segmentation: State of the art. *Int. J. Multimed. Inf. Retr.* **2020**, *9*, 171–189. [\[CrossRef\]](#)
- Guo, Y.L.; Wang, H.Y.; Hu, Q.Y.; Liu, H.; Bennamoun, M. Deep learning for 3D point clouds: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **2021**, *43*, 4338–4364. [\[CrossRef\]](#) [\[PubMed\]](#)
- Lu, B.; Liu, Y.W.; Zhang, Y.H.; Yang, Z.Y. Point cloud segmentation algorithm based on density awareness and self-attention mechanism. *Laser Optoelectron. Prog.* **2024**, *61*, 0811004. [\[CrossRef\]](#)
- Ai, D.; Zhang, X.Y.; Xu, C.; Qin, S.Y.; Yuan, H. Advancements in semantic segmentation methods for large-scale point clouds based on deep learning. *Laser Optoelectron. Prog.* **2024**, *61*, 1200003. [\[CrossRef\]](#)
- Zhang, K.; Zhu, Y.W.; Wang, X.H.; Zhang, L.T.; Zhong, R.F. Three-dimensional point cloud semantic segmentation network based on spatial graph convolution network. *Laser Optoelectron. Prog.* **2023**, *60*, 0228007. [\[CrossRef\]](#)
- Xu, X. *Research on 3D Instance Segmentation Method for Indoor Scene*; Northeast Petroleum University: Daqing, China, 2023; pp. 12–13.
- Cui, L.Q.; Hao, S.Y.; Luan, W.Y. Lightweight 3D point cloud instance segmentation algorithm based on Mamba. *Comput. Eng. Appl.* **2025**, *61*, 194–203. [\[CrossRef\]](#)
- Lei, T.; Guan, B.; Liang, M.; Li, X.Y.; Liu, J.B.; Tao, J.; Shang, Y.; Yu, Q.F. Event-based multi-view photogrammetry for high-dynamic, high-velocity target measurement. *arXiv* **2025**, arXiv:2506.00578.
- Liu, Z.; Liang, S.; Guan, B.; Tan, D.; Shang, Y.; Yu, Q.F. Collimator-assisted high-precision calibration method for event cameras. *Opt. Lett.* **2025**, *50*, 4254–4257. [\[CrossRef\]](#) [\[PubMed\]](#)
- Wang, W.Y.; Yu, R.; Huang, Q.G. SGPN: Similarity group proposal network for 3D point cloud instance segmentation. *arXiv* **2017**, arXiv:1711.08588.
- Hou, J.; Dai, A.; Nießner, M. 3D-SIS: 3D semantic instance segmentation of RGB-D scans. In Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019; IEEE Press: New York, NY, USA, 2019; pp. 4416–4425.
- Lin, K.H.; Zhao, H.M.; Lv, J.J.; Li, C.Y.; Liu, X.Y.; Chen, R.J.; Zhao, R.Y. Face detection and segmentation based on improved mask R-CNN. *Discret. Dyn. Nat. Soc.* **2020**, *2020*, 9242917. [\[CrossRef\]](#)

19. Yang, B.; Wang, J.; Clark, R.; Hu, Q.; Wang, S.; Markham, A.; Trigoni, N. Learning object bounding boxes for 3D instance segmentation on point clouds. In Proceedings of the 33rd Conference on Neural Information Processing Systems, Vancouver, BC, Canada, 8–14 December 2019; pp. 1–10.
20. Han, L.; Zheng, T.; Lan, X.; Lu, F. OccuSeg: Occupancy aware 3D instance segmentation. In Proceedings of the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 14–19 June 2020; pp. 2937–2946.
21. Kolodiazny, M.; Vorontsova, A.; Konushin, A.; Rukhovich, D. Top-down beats bottom-up in 3D instance segmentation. In Proceedings of the 2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), Waikoloa, HI, USA, 3–8 January 2024; IEEE Press: New York, NY, USA, 2024; pp. 3554–3562.
22. Charles, R.Q.; Li, Y.; Hao, S.; Leonidas, J.G. PointNet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv* **2017**, arXiv:1706.02413. Available online: <https://arxiv.org/abs/1706.02413> (accessed on 12 December 2024).
23. Charles, R.Q.; Hao, S.; Mo, K.C.; Leonidas, J.G. PointNet: Deep learning on point sets for 3D classification and segmentation. In Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017; IEEE Press: New York, NY, USA, 2017; pp. 77–85.
24. Mo, K.C.; Zhu, S.L.; Chang, A.X.; Yi, L.; Tripathi, S.; Guibas, L.J. PartNet: A large-scale benchmark for fine-grained and hierarchical part level 3D object understanding. In Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019; IEEE Press: New York, NY, USA, 2019; pp. 909–918.
25. Graham, B.; Engelcke, M.; van der Maaten, L. 3D semantic segmentation with submanifold sparse convolutional networks. In Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–22 June 2018; pp. 9224–9232.
26. Jiang, L.; Zhao, H.; Shi, S.; Liu, S.; Fu, C.W.; Jia, J. PointGroup: Dual-set point grouping for 3D instance segmentation. In Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 14–19 June 2020; pp. 4866–4875.
27. Chen, S.Y.; Fang, J.M.; Zhang, Q.; Liu, W.Y.; Wang, X.G. Hierarchical aggregation for 3D instance segmentation. *arXiv* **2021**, arXiv:2108.02350. [[CrossRef](#)]
28. He, T.; Shen, C.; Hengel, V.D.A. DyCO3D: Robust instance segmentation of 3D point clouds through dynamic convolution. In Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 19–25 June 2021; pp. 354–363.
29. Ngo, T.D.; Hua, B.S.; Nguyen, K. ISBNet: A 3D point cloud instance segmentation network with instance-aware sampling and box-aware dynamic convolution. In Proceedings of the 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Vancouver, BC, Canada, 17–24 June 2023; IEEE Press: New York, NY, USA, 2023; pp. 13550–13559.
30. Kolodiazny, M.; Vorontsova, A.; Konushin, A.; Rukhovich, D. Oneformer3d: One transformer for unified point cloud segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 17–18 June 2024; pp. 20943–20953.
31. Zou, T.; Qu, S.; Li, Z.; Knoll, A.; He, L.H.; Chen, G.; Jiang, C.J. Hgl: Hierarchical geometry learning for test-time adaptation in 3d point cloud segmentation. In *European Conference on Computer Vision*; Springer Nature: Cham, Switzerland, 2024; pp. 19–36.
32. Zhao, L.; Hu, Y.; Yang, X.; Dou, Z.L.; Kang, L.S. Robust multi-task learning network for complex LiDAR point cloud data preprocessing. *Expert Syst. Appl.* **2024**, *237*, 121552. [[CrossRef](#)]
33. Behley, J.; Garbade, M.; Milioto, A.; Quenzel, J.; Behnke, S.; Stachniss, C. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Seoul, Republic of Korea, 27 October–2 November 2019; pp. 9297–9307.
34. Sun, P.; Kretschmar, H.; Dotiwala, X.; Chouard, A.; Patnaik, V.; Tsui, P. Scalability in perception for autonomous driving: Waymo open dataset. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*; IEEE Press: New York, NY, USA, 2020; pp. 2446–2454.
35. Zhang, Y.; Zhou, Z.; David, P.; Yue, X.; Xi, Z.; Gong, B. Polarnet: An improved grid representation for online lidar point clouds semantic segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 14–19 June 2020; pp. 9601–9610.
36. Zhou, H.; Zhu, X.; Song, X.; Ma, Y.; Wang, Z.; Li, H.; Lin, D. Cylinder3d: An effective 3d framework for driving-scene lidar semantic segmentation. *arXiv* **2020**, arXiv:2008.01550.
37. Cheng, R.; Razani, R.; Taghavi, E.; Li, E.; Liu, B. (AF)²-s3net: Attentive feature fusion with adaptive feature selection for sparse semantic segmentation network. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 19–25 June 2021; pp. 12547–12556.
38. Kong, L.; Liu, Y.; Chen, R.; Ma, Y.; Zhu, X.; Li, Y. Rethinking range view representation for lidar segmentation. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Paris, France, 2–6 October 2023; pp. 228–240.

39. Cardace, A.; Spezialetti, R.; Ramirez, P.Z.; Salti, S.; Stefano, L.D. Self-distillation for unsupervised 3d domain adaptation. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, Waikoloa, HI, USA, 2–7 January 2023; pp. 4166–4177.
40. Wei, Y.; Liu, H.; Xie, T.; Ke, Q.; Guo, Y. Spatial-temporal transformer for 3d point cloud sequences. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, Waikoloa, HI, USA, 3–8 January 2022; pp. 1171–1180.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.