






SFPNet: Sparse Focal Point Network for Semantic Segmentation on General LiDAR Point Clouds

Yanbo Wang^{1,2}, Wentao Zhao^{1,2}, Chuan Cao^{1,2}, Tianchen Deng^{1,2},
Jingchuan Wang^{1,2}, and Weidong Chen^{1,2,†}

¹ Institute of Medical Robotics and Department of Automation, Shanghai Jiao Tong University, China

² Key Laboratory of System Control and Information Processing, Ministry of Education, Shanghai 200240, China

{yanbowang319, wentaozhao, alex008, dengtianchen, jchwang, wdchen}@sjtu.edu.cn

Abstract. Although LiDAR semantic segmentation advances rapidly, state-of-the-art methods often incorporate specifically designed inductive bias derived from benchmarks originating from mechanical spinning LiDAR. This can limit model generalizability to other kinds of LiDAR technologies and make hyperparameter tuning more complex. To tackle these issues, we propose a generalized framework to accommodate various types of LiDAR prevalent in the market by replacing window-attention with our sparse focal point modulation. Our SFPNet is capable of extracting multi-level contexts and dynamically aggregating them using a gate mechanism. By implementing a channel-wise information query, features that incorporate both local and global contexts are encoded. We also introduce a novel large-scale hybrid-solid LiDAR semantic segmentation dataset for robotic applications. SFPNet demonstrates competitive performance on conventional benchmarks derived from mechanical spinning LiDAR, while achieving state-of-the-art results on benchmark derived from solid-state LiDAR. Additionally, it outperforms existing methods on our novel dataset sourced from hybrid-solid LiDAR. Code and dataset are available at <https://github.com/Cavendish518/SFPNet> and <https://www.semanticindustry.top>.

Keywords: Semantic Segmentation · Sparse Focal Point Network · LiDAR Point Clouds · Inductive Bias

1 Introduction

Various types of 3D LiDAR sensors (as shown in Fig. 1) have become popular choices in autonomous vehicles and robotics [5, 7, 15, 52] due to their accurate distance detection capabilities across diverse environments, including low-light

[†] Corresponding Author.

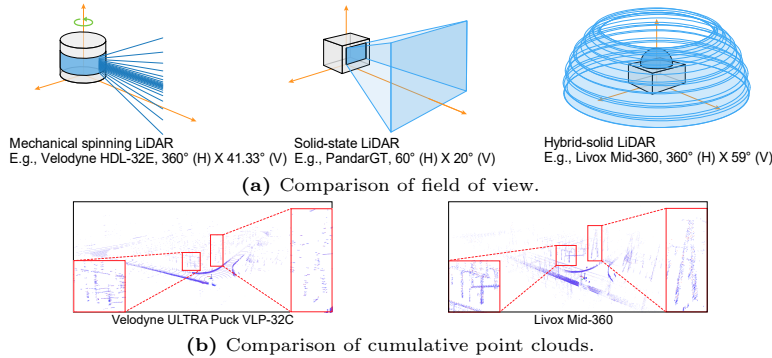


Fig. 1: Comparison of different types of LiDAR. Fig. 1 (a) compares three mainstream types of LiDAR technologies. Unlike camera, various types of LiDAR data have extremely different point distributions. Therefore, the generalizability of networks designed specifically for a particular LiDAR type is poor. Fig. 1 (b) contrasts the **cumulative** 1-second point clouds of Mid-360 (employed in our dataset) and commonly used VLP-32C. The non-repetitive scanning mode of Mid-360 covers a broader range of scenes, making it more suitable for industrial robots involving scene understanding tasks. Meanwhile, VLP-32C gathers more detailed road surface information.

conditions. The point clouds generated by LiDAR can accurately represent real-world scenes, facilitating direct 3D scene understanding through semantic segmentation. These advantages enable more effective support for subsequent tasks such as localization and planning compared to segmentation based on 2D images.

Despite the convenience afforded by LiDAR sensors, semantic segmentation based on LiDAR point clouds also encounters several challenges. These challenges primarily stem from characteristics inherent to LiDAR data. In general, the key features of all kinds of LiDAR data include sparsity, large scale, and non-uniform changes in point cloud density.

Most of the pioneering work [25, 34, 37, 61] did not take into account all the characteristics of the LiDAR data, or the model capacity was insufficient, result-

Table 1: Comparison of representative frameworks. For each input format, the second example improves the performance by introducing inductive bias for mechanical spinning LiDAR. FPS[†] and RS[‡] are the abbreviations for farthest point sampling and random sampling, respectively.

Input Format	Example	Core inductive bias
Point clouds	PointNet++ [37]	Neighbor consistency with FPS [†]
	RandLA-Net [21]	Neighbor consistency with RS [‡]
Range image	RangeNet++ [34]	Square locality
	RangeFormer [23]	Implicit square locality
Sparse voxel grid and points	SPVNAS [40]	Cubic locality
	Cylinder3D [61]	Cylindrical and asymmetrical locality
Sparse voxel grid with point property	VoTr [33]	Implicit cubic locality
	SphereFormer [25]	Implicit radial and cubic locality

ing in unsatisfactory performance. Cutting-edge works [21, 23, 25, 61] adapt to the distribution of mechanical spinning LiDAR data through specially designed inductive bias as shown in Tab. 1.

However, there are various types of LiDAR with distinct characteristics as illustrated in Fig. 1. According to the no free lunch theorem [44], state-of-the-art (SOTA) methods, which incorporate specifically designed inductive bias (*e.g.*, cylindrical partition [61] and radial window [25] as shown in Fig. 2), risk constraining model generalizability and complicating hyperparameter tuning when applied to other types of LiDAR technologies.

Motivated by the analysis in Fig. 1, we aim to propose a generalized framework capable of addressing the common characteristics of various types of LiDAR data prevalent in the market. Our goal is to ensure competitiveness on traditional benchmarks and demonstrate generality across other types of LiDAR data without introducing special inductive bias.

Inspired by focal attention [58] and focal modulation [57], we propose sparse focal point modulation (SFPM) by first extract features at different focal levels around each point. Then multi-level contexts are adaptively aggregated through a gate mechanism. Finally, a channel-wise information query is implemented to acquire the encoded features with both local and long-range information. Similar to window-attention [25], SFPM serves as a plugin module, thereby capable of seamlessly replacing window-attention in mainstream backbones.

To enable training and evaluation of LiDAR semantic segmentation based on hybrid-solid LiDAR, we build a new dataset, **SeMantic InDustry** (S.MID). We are the first to develop a hybrid-solid LiDAR semantic segmentation dataset based on the Livox Mid-360. To furnish robotic application scene for LiDAR semantic segmentation community, we used an industrial robot to collect a total of 38904 frames of LiDAR data in different substations. We annotated 25 categories under professional guidance and merged them into 14 classes for single-frame segmentation task.

We evaluate our sparse focal point network (SFPNet) on two mechanical spinning LiDAR datasets, nuScenes [6] and SemanticKITTI [4]. Our method achieves competitive results. We evaluate our SFPNet performance on solid-state LiDAR dataset through PandaSet [49] and performance on hybrid-solid LiDAR dataset through our dataset S.MID. Our method achieves better results compared with the SOTA works. Experimental results show that our frameworks have strong generalization capability and interpretability.

We summarize our contributions as below:

- We introduce SFPNet for feature encoding of sparse point clouds obtained from various types of LiDAR sensors. SFPNet effectively avoids introducing special inductive bias while ensuring expansive receptive fields. Additionally, it offers enhanced interpretability for semantic segmentation tasks.
- In our sparse focal point modulation, the introduced multi-level feature extraction and gated aggregation can adaptively learn local and global features from various LiDAR point cloud data with different distribution patterns and non-uniform density variations.

- A novel dataset for LiDAR semantic segmentation has been collected. S.MID is built with a novel hybrid-solid LiDAR in the substations. It fills the gap of public dataset in industrial outdoor scenes for robotic application.
- Our proposed method achieves the cutting-edge results on both nuScenes and SemanticKITTI, which are based on mechanical spinning LiDAR. More importantly, our new framework demonstrates its ability to generalize across different LiDAR technologies, as evidenced by its leading performance on PandaSet (solid-state LiDAR) and S.MID (hybrid-solid LiDAR).

2 Related Works

From advancements in 3D point cloud recognition to the development of LiDAR-based semantic segmentation, a variety of interesting techniques have been introduced. In accordance with the input format, methods are typically classified into point-based, projection-based, voxel-based, and multi-modality-based approaches. Previous high-performance methods usually have designed networks with tailored inductive bias to effectively address the characteristics of LiDAR point cloud data, including its large scale, low data volume, overall sparsity, and non-uniform density variations. In this section, we aim to summarize previous works from a novel inductive bias view.

2.1 Explicit Locality Assumption Methods

The explicit locality assumptions of mainstream frameworks are usually derived from the inherent properties of K-nearest neighbor (KNN), 2D CNN, 3D CNN and submanifold sparse convolutional networks (SSCN) [17].

Explicit 2D Locality. Explicit 2D locality assumption methods usually employ a 2D backbone to extract features from a range view [1, 4, 34, 39, 45, 46, 50], a bird-eye view (BEV) [59], or a combination of three planes [36] projected from LiDAR point clouds. In addition to the information loss caused during the projection process, the inductive bias of traditional 2D CNN also suffers from a certain degree of failure. Not only does the 2D CNN process irrelevant points in faraway locations, rendering locality invalid, but the deformation produced in the projection space also does not satisfy the spatial invariance assumption. Although the structure relying on explicit 2D locality is efficient, the above problems make the model capacity of this type of method easily reach the upper limit.

Explicit 3D Locality. Explicit 3D locality assumption methods usually adopt a PointNet++ [37] like network or SSCN [17] to extract features. The former kinds of methods [21, 38, 41, 42, 56] rely on neighbor consistency in 3D point clouds to extract and aggregate multi-scale features. The latter kinds of methods [10–12, 22, 28, 40, 54] depend on the inductive bias from 3D CNN or SSCN.

However, non-uniform density variations of LiDAR point clouds limit the effectiveness of SSCN. Further improvements have been made by modifying the output region [8] or enlarging the receptive field [9, 32]. Cylinder3D [61] proposes a cylindrical partition by investigating point distribution obtained from mechanical spinning LiDAR to make the data more consistent with the locality assumption of SSCN. In addition, the introduced asymmetrical convolution design [61] is more consistent with the density changes of point clouds in each direction. We have noticed that the key to the success of these improved methods enable SSCN to operate in a way that is more suitable for data distribution to obtain favorable properties such as better 3D locality assumptions.

2.2 Implicit Locality Assumption Methods

Transformer [43] has shown strong capability in modeling long-range dependency with the self-attention mechanism. However, it requires large dataset and longer training schedules to learn the implicit locality. Leveraging the attributes of mechanical spinning LiDAR, various architectures, tokenization strategies, window forms and positional encoding for transformer have been proposed.

Implicit 2D Locality. Recently, RangeViT [2] and RangeFormer [23] have been proposed to overcome the problem of insufficient model capacity of the explicit 2D locality assumption methods through the power of the transformer architecture. However, the performance still lags behind the latest 3D methods, requiring reliance on pre-processing and post-processing to achieve SOTA results.

Implicit 3D Locality. Several pioneering transformer-based works [14, 26, 33, 47] have been proposed to solve 3D point cloud perception. These works take 3D point cloud sequences or sparse voxel with point properties as input and learn 3D locality implicitly. These methods are more suitable for indoor datasets than LiDAR datasets. SphereFormer [25] introduces locality in the radial direction by calculating self-attention within the radial window to improve the perception of distant LiDAR point clouds. We have noticed that the key to the success of these methods is the introduction of long-range dependency while implicitly learning locality through specially designed positional encoding.

2.3 Multi-modality Assumption Methods

Fundamental assumptions in a single modality have drawbacks as mentioned above. Introducing additional modalities is regarded as a solution to address the limitation in model capacity arising from the assumption of a singular modality. RPVNet [51] takes advantage of multi-view fusion. This work leverages an inductive bias towards arbitrary relations between views, facilitating the flow of information across different views at the feature level. Moreover, [16, 29, 55, 62] introduce additional image information, which inherits the inherent inductive bias of the fundamental backbones. Although our work is dedicated to feature

encoding of general LiDAR point clouds, our approach still achieved superior performance than some of these methods.

Through summarizing the structure designed by cutting-edge frameworks, we propose our network architecture under explicit locality assumption with adaptive hierarchical contextualization aggregation to capture long-range dependence and global information. We take sparse voxel with point properties as our only input. Please note that our work focuses on the representational capabilities of the network design itself, rather than special data augmentation [23, 25, 48, 51, 53], training skills [24], post-processing [2, 34] and distillation [20, 28, 55]. Therefore, related powerful training techniques are not summarized here and have not been used in the experiments.

3 Methods

In this section, we first summarize the sparse submanifold convolution and self-attention mechanism employed in mainstream backbones and compare them with the sparse focal point modulation introduced in our research, as outlined in Sec. 3.1. Subsequently, we provide details for the implementation of sparse focal point modulation in Sec. 3.2. Then, we give our overall framework and objective function in the supplementary materials.

3.1 Rethinking Submanifold Sparse Convolution and Window-Attention

In general, considering a sorted LiDAR feature sequence $X \in \mathbb{R}^{N \times C_{in}}$ (feature in sparse tensor or feature inherited from original point sequence) as input, the output feature $y_i \in \mathbb{R}^{C_{out}}$ is encoded from token (or voxel) $x_i \in \mathbb{R}^{C_{in}}$ through neighborhood interaction ξ and contextual aggregation κ .

Submanifold Sparse Convolution. The submanifold sparse convolution without pooling operation can be formulated as:

$$y_i = \xi_{SubMconv}(i, X), \quad (1)$$

where neighbourhood features at location i are captured with well preserved sparsity [18] through efficient interaction $\xi_{SubMconv}$.

Window-Attention. The window-attention [30] can be formulated as:

$$y_i = \kappa_{attn}(x_i, \xi_{attn}(x_i, X)), \quad (2)$$

where the computation of attention scores between the query and key via the interaction ξ_{attn} occurs before the aggregation κ_{attn} within the window.

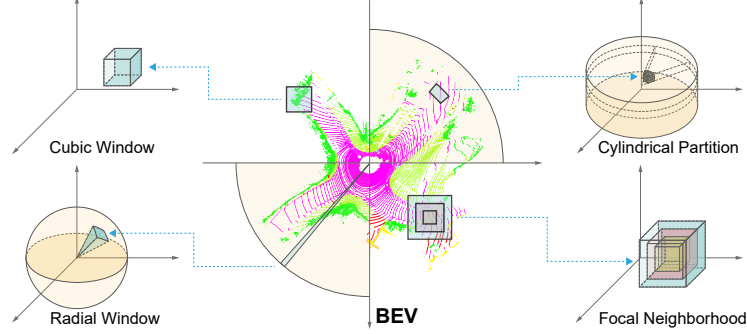


Fig. 2: Heuristic comparison with mainstream design. Based on point distribution of mechanical spinning LiDAR, Cylindrical partition [61] and radial window [25] are proposed to extract the features of distant points. Focal neighborhood adapts to this problem by aggregating multi-level contexts. Since no special inductive bias is introduced, it can be elegantly applied to various kinds of LiDAR as shown in Fig. 1.

Sparse Focal Point Modulation. Followed by focal modulation [57] paradigm, sparse focal point modulation (SFPM) can be formulated as:

$$y_i = \xi_{focal}(x_i, \kappa_{focal}(i, X)), \quad (3)$$

where the interaction operation ξ_{focal} performs the query task subsequent to aggregated contexts of the focal neighbors using κ_{focal} at each location i .

Comparison with Mainstream Design. As an extension of [57], we believe that SFPM combines the advantages of submanifold sparse convolution and window-attention by comparing Eqs. (1) to (3).

- *Explicit locality with contextual learning.* κ_{focal} extracts the contexts from location i . ξ_{focal} preserves channel information for each token x_i . Therefore, SFPM simultaneously possesses spatial-specific and channel-specific properties while exhibiting decoupled feature granularity.
- *Translation invariance.* SFPM is invariant to translation of sorted input LiDAR feature sequence X , since operation ξ_{focal} and κ_{focal} are always centered at location i . This also eliminates the need for positional encoding.

A heuristic comparison with mainstream design between our methods with Cylinder3D [61] and SphereFormer [25] can be found in Fig. 2. With the above advantages, it does not rely on special point distribution assumption, and has a large receptive field.

3.2 Sparse Focal Point Modulation

Fig. 3 shows the structure of SFPM. Based on Eq. (3), κ_{focal} is designed through **multi-level context extraction** and **adaptive feature aggregation**, and ξ_{focal} is designed into **channel-wise information query**.

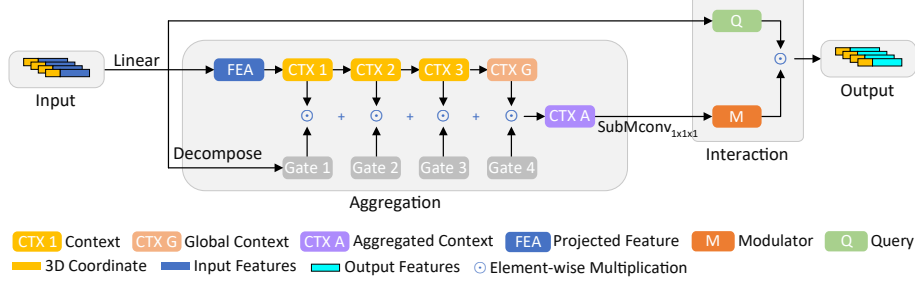


Fig. 3: Illustration of sparse focal point modulation.

Multi-level Context Extraction. Previous works have proved that both local features and long-range contexts are important for LiDAR segmentation [9, 25, 32]. Therefore, we hope to obtain features of different levels in the first step to acquire hierarchical information.

Given a sorted LiDAR feature sequence $X \in \mathbb{R}^{N \times C_{in}}$, we first employ a linear layer $S^0 = f_X^S(X) \triangleq MLP(X)$ to project it into a new feature space while keeping the same channel numbers C_{in} . Then the sparse multi-level contexts can be obtained through a sequence of 3D submanifold sparse convolution with $C_{in} = C_{out}$ and a kernel size of k^l at level l .

At focal level $l \in \{1, \dots, L\}$, the output S^l is derived by:

$$S^l = f_{l-1}^l(S^{l-1}) \triangleq LN(GeLU(SubMconv_{3d}^l(S^{l-1}))), \quad (4)$$

where LN represents for layer normalization [3], and GeLU is Gaussian error linear unit [19]. The kernel size at each focal level is increasing through $k^l = k^{l-1} + 2$. The effective receptive field at each level is $RF^l = 1 + \sum_{i=1}^l (k^i - 1)$, which leads to long-range dependence learning.

Finally, global average pooling $S^{L+1} = Avgpool(S^L)$ is performed on the L -th level features to obtain global context.

Adaptive Feature Aggregation. Through the above steps, multi-level contexts $\{S^1, \dots, S^L, S^{L+1}\}$ have been extracted. However, not all the contexts are equally important. Even point clouds from the same type of LiDAR belonging to the same type of objects will produce different multi-level features due to various point cloud distributions. Therefore, an adaptive feature aggregation is achieved through gated mechanism. Following the gated aggregation in [57], the spatial-aware gating weights for each level are calculated through $G = f_X^G(x) \triangleq MLP(X)$ with $L + 1$ channels. And the cross channel result F^{out} after gate aggregation is obtained through:

$$S^{out} = \sum_{l=1}^{L+1} G^l \odot S^l, \quad (5)$$

$$F^{out} = h(S^{out}) \triangleq SubMconv_{3d}^{1 \times 1 \times 1}(S^{out}), \quad (6)$$

where \odot is element-wise multiplication and $h(\cdot)$ represents the cross channel aggregation.

The whole aggregation step κ_{focal} can adaptively learn hard tokens through multi-level contexts, and can also avoid introducing too much invalid information to easy tokens. This also enables SFPM to accommodate various distribution patterns of point clouds without special inductive bias. Model interpretation through visualization experiments can be found in Sec. 4.3.

Channel-wise Information Query. Following [30, 43, 57, 58], ξ_{focal} is simply implemented through a query projection $q(x_i) \triangleq MLP(X)$ with the same channel number $C_{in} = C_{out}$. Derived from Eqs. (3) to (6), the features from SFPM is encoded via:

$$y_i = q(x_i) \odot h\left(\sum_{l=1}^{L+1} g_i^l \cdot s_i^l\right), \quad (7)$$

where g_i^l and s_i^l are the gate weights and sparse contexts at location i and level l , respectively. Through this lightweight element-wise multiplication, channel-wise information is well-preserved.

4 Experiments

In this section, we first introduce our experiment setups in Sec. 4.1. Then we show segmentation results across different datasets in Sec. 4.2. We further analyze the network design and interpretability in Sec. 4.3. Finally, ablation study is shown in Sec. 4.4.

4.1 Experiment Setups

Datasets.

- *Mechanical Spinning LiDAR.* Two large-scale driving-scene benchmarks, nuScenes (Velodyne-HDL-32E, 16 classes) [6] and SemanticKITTI [4] (Velodyne-HDL-64E, 19 classes) are employed to verify the performance for point clouds obtained from traditional mechanical spinning LiDAR.
- *Solid-State LiDAR.* We extract labeled solid-state LiDAR data from PandaSet (PandarGT) [49]. We merge and select 13 classes for evaluation. The data split is in an 4:1 ratio for training and validation.
- *Hybrid-Solid LiDAR.* We develop a novel dataset S.MID. We hope to setup a large-scale robotic application benchmark for LiDAR semantic segmentation task. We collect a total of 38904 frames of LiDAR data in different substations through an industrial robot equipped with Livox Mid-360, a novel hybrid-solid LiDAR. We carefully split the dataset as follows. 13,101 frames are collected from one complete substation for training. Validation and test sets are sourced from different substations, comprising 5,000 frames and 20,803 frames, respectively. 25 categories are annotated under professional

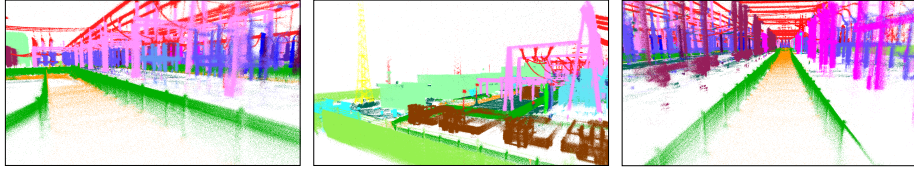


Fig. 4: Labeled cumulative point clouds in our novel dataset **S.MID**. We provide dense semantic annotations for each frame.

guidance as shown in Fig. 4. 14 classes are remained for single-frame segmentation task after merging classes with collective name and ignore classes with very few points. More details (sensors, scenes, annotation process, label distributions, *etc.*) about datasets can be found in the supplementary materials.

Implementation Details. We utilize two GeForce RTX 3090 GPUs for training, with the exception of SemanticKITTI, where four GPUs are employed. We train our models for 70 epochs using a batch size of 8, employing the AdamW optimizer [31] with a learning rate of 0.0008 and a polynomial learning rate scheduler.

4.2 Segmentation Results

As mentioned in Sec. 2.3, we focus on the representational capabilities of the network design itself and have not employed powerful training techniques in the experiments. Additionally, we select open source SOTA methods [25, 61] as our main comparison methods, since these methods play an important role as backbones or one of the modalities in cutting-edge works.

Mechanical Spinning LiDAR. The semantic segmentation results for the nuScenes and SemanticKITTI validation and test sets are displayed in Tab. 2, Tab. 3, and supplementary materials. The results demonstrate that our approach can achieve competitive performance even without inductive bias specifically tailored for mechanical spinning LiDAR. We attribute this achievement to the robust adaptability of our newly designed SFPM module. As shown in Tab. 3, our method outperforms all existing LiDAR-based methods on the nuScenes validation set. Compared to the methods incorporating explicit 2D locality assumptions [13, 27, 34, 59], ours yields a substantial performance improvement, ranging from 4.0% ~ 14.6% in terms of mIoU. Moreover, thanks to multi-level context aggregation, our method surpasses the models relying on explicit 3D locality assumptions [10, 35, 60, 61] by 4.0% ~ 17.9% mIoU. Our method also achieves 0.6% ~ 4.9% mIoU performance gain compared to methods with implicit 2D/3D locality assumptions [2, 23, 25]. It is notable that our LiDAR-based method even outperforms several approaches utilizing additional 2D information [16, 55, 62].

Table 2: Comparison with SOTA backbone type of works on four datasets. Note that all results are obtained from the literature[†] or from open source codes[‡] with carefully chosen parameters. No powerful training techniques are employed during our model training.

Methods	nuScenes [†]		SemanticKITTI [†]		PandaSet [‡]	S.MID [‡]	
	Val	Test	Val	Test	Val	Val	Test
Cylinder3D [61]	76.1	77.2	65.9	68.9	55.0	68.8	68.1
SphereFormer [25]	79.5	81.9	69.0	74.8	63.5	67.8	68.3
Ours	80.1	80.2	69.2	70.3	64.0	71.9	70.9

Table 3: Results of our proposed method and SOTA LiDAR Segmentation methods on nuScenes val set. Note that all results are obtained from the literature. All of the LiDAR-based methods published before March 7, 2024 have been listed. L and C represent for LiDAR and camera, respectively. Top 3 for each class are marked in **blue**.

Methods (year)	Input	mIoU	barrier	bicycle	bus	car	construction	motorcycle	pedestrian	traffic cone	trailer	truck	driveable	other flat	sidewalk	terrain	manmade	vegetation
(AF) ² -S3Net [10] (*21)	L	62.2	60.3	12.6	82.3	80.0	20.1	62.0	59.0	49.0	42.2	67.4	94.2	68.0	64.1	68.6	82.9	82.4
RangeNet53++ [34] (*19)	L	65.5	66.0	21.3	77.2	80.9	30.2	66.8	69.6	52.1	54.2	72.3	94.1	66.6	63.5	70.1	83.1	79.8
PolarNet [59] (*22)	L	71.0	74.7	28.2	85.3	90.9	35.1	77.5	71.3	58.8	57.4	76.1	96.5	71.1	74.7	74.0	87.3	85.7
PCSCNet [35] (*22)	L	72.0	73.3	42.2	87.8	86.1	44.9	82.2	76.1	62.9	49.3	77.3	95.2	66.9	69.5	72.3	83.7	82.5
SalsaNext [13] (*20)	L	72.2	74.8	34.1	85.9	88.4	42.2	72.4	72.2	63.1	61.3	76.5	96.0	70.8	71.2	71.5	86.7	84.4
SVASeg [60] (*22)	L	74.7	73.1	44.5	88.4	86.6	48.2	80.5	77.7	65.6	57.5	82.1	96.5	70.5	74.7	74.6	87.3	86.9
RangeViT-CS [2] (*23)	L	75.2	75.5	40.7	88.3	90.1	49.3	79.3	77.2	66.3	65.2	80.0	96.4	71.4	73.8	73.8	89.9	87.2
AMVNet [27] (*20)	L	76.1	79.8	32.4	82.2	86.4	62.5	81.9	75.3	72.3	83.5	65.1	97.4	67.0	78.8	74.6	90.8	87.9
Cylinder3D [61] (*21)	L	76.1	76.4	40.3	91.2	93.8	51.3	78.0	78.9	64.9	62.1	84.4	96.8	71.6	76.4	75.4	90.5	87.4
PVKD [20] (*22)	L	76.0	76.2	40.0	90.2	94.0	50.9	77.4	78.8	64.7	62.0	84.1	96.6	71.4	76.4	76.3	90.3	86.9
PMF [62] (*21)	L+C	76.9	74.1	46.6	89.8	92.7	57.0	77.7	80.9	70.9	64.6	82.9	95.5	73.3	73.6	74.8	89.4	87.7
RPVNet [51] (*21)	L	77.6	78.2	43.4	92.7	93.2	49.0	85.7	80.5	66.0	66.9	84.0	96.9	73.5	75.9	76.0	90.6	88.9
RangeFormer [23] (*23)	L	78.1	78.0	45.2	94.0	92.9	58.7	83.9	77.9	69.1	63.7	85.6	96.7	74.5	75.1	75.3	89.1	87.5
2D3DNet [16] (*21)	L+C	79.0	78.3	55.1	95.4	87.7	59.4	79.3	80.7	70.2	68.2	86.6	96.1	74.9	75.7	75.1	91.4	89.9
2DPASS [55] (*22)	L(C)	79.4	78.8	49.6	95.6	93.6	60.0	84.1	82.2	67.5	72.6	88.1	96.8	72.8	76.2	76.5	89.4	87.2
SphereFormer [25] (*23)	L	79.5	78.7	46.7	95.2	93.7	54.0	88.9	81.1	68.0	74.2	86.2	97.2	74.3	76.3	75.8	91.4	89.7
Ours	L	80.1	78.8	49.7	95.3	93.5	63.1	86.4	82.9	68.6	72.8	86.7	97.0	74.7	76.0	75.3	91.2	89.5

Solid-State LiDAR. The results on PandaSet val set are shown in Tab. 4. In comparison to mechanical spinning LiDAR, solid-state LiDAR exhibits a smaller horizontal field of view, greater detection range, and much finer vertical resolution. Our model surpasses the Cylinder3D [61], achieving a notable margin of 9.0% mIoU. Even when compared to the SphereFormer [25] which is proposed specifically for enhancing long-range dependencies fit for greater detection range, our approach still maintains a lead of 0.5% mIoU.

Hybrid-Solid LiDAR. In Tab. 2 and Tab. 5, we compare the results of our proposed method with open source SOTA LiDAR segmentation methods on our dataset S.MID. The point cloud distribution of each frame in hybrid-solid LiDAR exhibits **significant randomness** compared to mechanical spinning LiDAR. Previous methods demonstrate only marginal improvement compared to the common baseline SSCN [17] (cubic and radial window attention with +0.2%

Table 4: Results of our proposed method and SOTA LiDAR Segmentation methods on PandaSet validation set. Note that all results are obtained from open source code with carefully chosen parameters.

Methods	mIoU	reflection	vegetation	ground	road	lane line	stop Line	sidewalk	car	truck	motorcycle	pedestrian	signs	building
Cylinder3D [61]	55.0	21.0	88.2	65.3	94.9	42.1	14.4	59.0	90.3	50.7	0.0	43.5	57.8	87.7
SphereFormer [25]	63.5	23.4	92.5	77.2	97.3	57.0	33.4	73.9	93.3	56.5	0.0	59.9	69.6	91.9
Ours	64.0	23.6	92.1	77.9	97.6	58.8	41.7	75.1	94.1	63.4	0.0	55.4	61.5	91.3

Table 5: Results of our proposed method and SOTA LiDAR Segmentation methods on S.MID val set. Note that all results are obtained from open source code with carefully chosen parameters.

Methods	mIoU	knife switch	main xfmr	arrester	voltage xfmr	busbar	switch	current xfmr	scaffold	sup column	road	other-ground	fence	fire shelter	wall
SSCN [17] (baseline)	67.6	83.4	77.9	76.9	46.9	80.7	94.0	96.1	86.8	25.3	58.3	35.3	80.4	18.0	86.3
Cylinder3D [61]	68.8 (+1.2)	83.7	77.8	74.8	41.5	79.6	93.5	94.2	83.6	48.9	54.5	35.8	81.2	29.1	84.7
SphereFormer [25]	67.8 (+0.2)	84.2	68.7	76.8	45.1	80.7	96.2	96.5	86.7	39.0	61.8	38.8	74.5	18.2	82.4
Ours	71.9 (+4.3)	89.1	91.9	85.0	43.0	77.6	96.6	96.3	88.2	52.4	60.3	33.4	85.1	19.9	88.7

mIoU and cylindrical partition with +1.2% mIoU), due to the failure of specially designed inductive bias. However, our method still outperforms 4.3% mIoU. This demonstrates that when the distribution pattern of point clouds changes, our method is much more effective than previous approaches.

Visual Comparison. In Fig. 5, we visually compare the results from SphereFormer [25] and ours on S.MID. It visually indicates that our approach demonstrates superior performance in segmenting objects with similar geometric structures and distinguishing adjacent object boundaries.

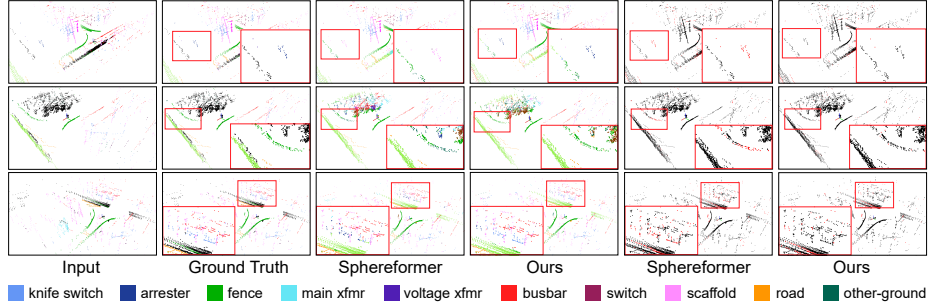


Fig. 5: Visual comparison between SphereFormer [25] and ours on S.MID val set. Details have been zoomed with red box. Difference maps are shown in the last two columns. More examples are given in the supplementary materials.

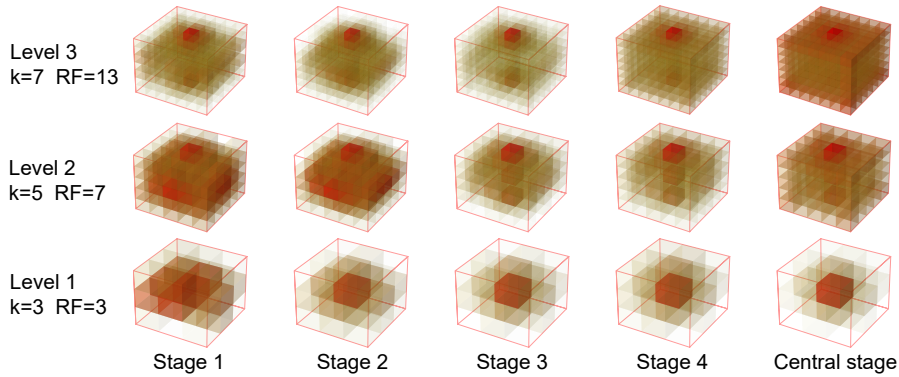


Fig. 6: Visualization of parameters of $SubMconv_{3d}^l$ in Eq. (4) at three focal levels in four down stages and central stage in SFPNet on nuScenes. For general pattern display, we take the average of all channels. The small cubes in the image are depicted as being redder and more opaque, representing higher weights.

4.3 Network Analysis

Patterns for Multi-level Context Extraction. Fig. 6 shows the learned $SubMconv_{3d}^l$ kernels in our SFPNet on nuScenes. On the one hand, it demonstrates the adaptive feature capturing ability of our network. The models prioritize low focal levels to capture local features during early stages. As the stages progress, greater reliance is placed on long-range contexts. On the other hand, it shows the ability to accommodate point distributions. The horizontal resolution of mechanical spinning LiDAR is much higher than the vertical resolution. Therefore, the distributions in vertical direction and BEV are not equivalent in the representation space. The changes in weights also coincide with this property.

Interpretability. Fig. 7 illustrates the correlation between location i and the feature sequence X in high-dimensional space obtained from Eq. (6). The visualization results simply elucidate the interpretability of our network. It is also noteworthy that even without utilizing a radial window or cylindrical partition, our model enables distant points to attend to a broader effective neighborhood.

4.4 Ablation Study

In order to evaluate the performance of each element within our networks, we carry out a series of ablation experiments utilizing the nuScenes dataset as shown in Tab. 6. We take SSCN [17] as our basic blocks. By removing SFPM, a decrease of 4.7% mIoU over the optimal design has been observed, which underscores the efficacy of SFPM.

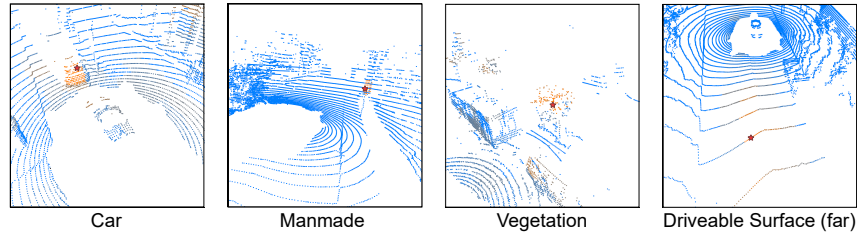


Fig. 7: Visualization of correlations between a certain location and feature sequence in Eq. (6) at central stage on nuScenes val set. The results are restored through the correspondence with the input point clouds. Location i has been marked as red star.

Table 6: Ablation study.

	Basic blocks	Focal level = 2	Focal level = 3	Global Avg Pooling	mIoU
Optimal design	✓		✓	✓	80.1
Ablation 1	✓	✓		✓	78.2 (-1.9)
Ablation 2	✓		✓		78.1 (-2.0)
Ablation 3	✓				75.4 (-4.7)

Fewer Focal Levels. We remove the last focal layer in SFPNet and it hurts 1.9% mIoU performance. This indicates that during multi-level context extraction, longer-range context can offer a larger effective receptive field, thereby compensating for the deficiency of local features faced by hard tokens due to sparsity.

No Global Average Pooling. Removing the global average pooling results in a 2.0% mIoU disparity compared to the optimal design. This demonstrates the significance of global information despite the presence of long-range context.

5 Conclusion

Our work proposes a generalized framework SFPNet to accommodate various types of LiDAR prevalent in the market. Our approach integrates multi-level context extraction and a gate mechanism to effectively aggregate both local and global features. Furthermore, we employ a lightweight element-wise multiplication operation with query to ensure the preservation of channel-wise information. A large-scale hybrid-solid LiDAR segmentation dataset for industrial robot applications has been collected and annotated under professional guidance. Our proposed network demonstrates outstanding performance on datasets derived from various types of LiDAR and exhibits excellent interpretability. We anticipate a diversification of LiDAR technologies for future industrial applications. We believe that our proposed network can better adapt to various types of LiDAR sensors or even their fused data. The limitations of our work are discussed in the supplementary materials.

Acknowledgements

This work is supported by the National Natural Science Foundation of China under Grant U2013203 and the Science and Technology Commission of Shanghai Municipality under Grant 20DZ2220400.

References

1. Aksoy, E.E., Baci, S., Cavdar, S.: Salsanet: Fast road and vehicle segmentation in lidar point clouds for autonomous driving. In: 2020 IEEE intelligent vehicles symposium (IV). pp. 926–932. IEEE (2020)
2. Ando, A., Gidaris, S., Bursuc, A., Puy, G., Boulch, A., Marlet, R.: Rangevit: Towards vision transformers for 3d semantic segmentation in autonomous driving. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 5240–5250 (2023)
3. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)
4. Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., Gall, J.: Semantickitti: A dataset for semantic scene understanding of lidar sequences. In: Int. Conf. Comput. Vis. pp. 9297–9307 (2019)
5. Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., Reid, I., Leonard, J.J.: Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. IEEE Transactions on robotics **32**(6), 1309–1332 (2016)
6. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 11621–11631 (2020)
7. Chen, X., Milioto, A., Palazzolo, E., Giguere, P., Behley, J., Stachniss, C.: Suma++: Efficient lidar-based semantic slam. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 4530–4537. IEEE (2019)
8. Chen, Y., Li, Y., Zhang, X., Sun, J., Jia, J.: Focal sparse convolutional networks for 3d object detection. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 5428–5437 (2022)
9. Chen, Y., Liu, J., Zhang, X., Qi, X., Jia, J.: Largekernel3d: Scaling up kernels in 3d sparse cnns. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 13488–13498 (2023)
10. Cheng, R., Razani, R., Taghavi, E., Li, E., Liu, B.: 2-s3net: Attentive feature fusion with adaptive feature selection for sparse semantic segmentation network. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 12547–12556 (2021)
11. Choy, C., Gwak, J., Savarese, S.: 4d spatio-temporal convnets: Minkowski convolutional neural networks. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 3075–3084 (2019)
12. Cohen, T.S., Geiger, M., Köhler, J., Welling, M.: Spherical cnns. arXiv preprint arXiv:1801.10130 (2018)
13. Cortinhal, T., Tzelepis, G., Erdal Aksoy, E.: Salsanext: Fast, uncertainty-aware semantic segmentation of lidar point clouds. In: Advances in Visual Computing: 15th International Symposium, ISVC 2020, San Diego, CA, USA, October 5–7, 2020, Proceedings, Part II 15. pp. 207–222. Springer (2020)
14. Fan, L., Pang, Z., Zhang, T., Wang, Y.X., Zhao, H., Wang, F., Wang, N., Zhang, Z.: Embracing single stride 3d object detector with sparse transformer. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 8458–8468 (2022)

15. Gao, B., Pan, Y., Li, C., Geng, S., Zhao, H.: Are we hungry for 3d lidar data for semantic segmentation? a survey of datasets and methods. *IEEE Transactions on Intelligent Transportation Systems* **23**(7), 6063–6081 (2021)
16. Genova, K., Yin, X., Kundu, A., Pantofaru, C., Cole, F., Sud, A., Brewington, B., Shucker, B., Funkhouser, T.: Learning 3d semantic segmentation with only 2d image supervision. pp. 361–372 (12 2021). <https://doi.org/10.1109/3DV53792.2021.00046>
17. Graham, B., Engelcke, M., Van Der Maaten, L.: 3d semantic segmentation with submanifold sparse convolutional networks. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 9224–9232 (2018)
18. Graham, B., Van der Maaten, L.: Submanifold sparse convolutional networks. *arXiv preprint arXiv:1706.01307* (2017)
19. Hendrycks, D., Gimpel, K.: Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415* (2016)
20. Hou, Y., Zhu, X., Ma, Y., Loy, C.C., Li, Y.: Point-to-voxel knowledge distillation for lidar semantic segmentation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 8479–8488 (2022)
21. Hu, Q., Yang, B., Xie, L., Rosa, S., Guo, Y., Wang, Z., Trigoni, N., Markham, A.: Randla-net: Efficient semantic segmentation of large-scale point clouds. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 11108–11117 (2020)
22. Jiang, L., Shi, S., Tian, Z., Lai, X., Liu, S., Fu, C.W., Jia, J.: Guided point contrastive learning for semi-supervised point cloud semantic segmentation. In: *Int. Conf. Comput. Vis.* pp. 6423–6432 (2021)
23. Kong, L., Liu, Y., Chen, R., Ma, Y., Zhu, X., Li, Y., Hou, Y., Qiao, Y., Liu, Z.: Rethinking range view representation for lidar segmentation. In: *Int. Conf. Comput. Vis.* pp. 228–240 (2023)
24. Kong, L., Ren, J., Pan, L., Liu, Z.: Lasermix for semi-supervised lidar semantic segmentation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 21705–21715 (2023)
25. Lai, X., Chen, Y., Lu, F., Liu, J., Jia, J.: Spherical transformer for lidar-based 3d recognition. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 17545–17555 (2023)
26. Lai, X., Liu, J., Jiang, L., Wang, L., Zhao, H., Liu, S., Qi, X., Jia, J.: Stratified transformer for 3d point cloud segmentation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 8500–8509 (2022)
27. Liong, V.E., Nguyen, T.N.T., Widjaja, S.A., Sharma, D., Chong, Z.J.: Amvnet: Assertion-based multi-view fusion network for lidar semantic segmentation. *ArXiv abs/2012.04934* (2020), <https://api.semanticscholar.org/CorpusID:228063957>
28. Liu, M., Zhou, Y., Qi, C.R., Gong, B., Su, H., Anguelov, D.: Less: Label-efficient semantic segmentation for lidar point clouds. In: *Eur. Conf. Comput. Vis.* pp. 70–89. Springer (2022)
29. Liu, Y., Chen, R., Li, X., Kong, L., Yang, Y., Xia, Z., Bai, Y., Zhu, X., Ma, Y., Li, Y., et al.: Uniseg: A unified multi-modal lidar segmentation network and the openpcseg codebase. In: *Int. Conf. Comput. Vis.* pp. 21662–21673 (2023)
30. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: *Int. Conf. Comput. Vis.* pp. 10012–10022 (2021)
31. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017)
32. Lu, T., Ding, X., Liu, H., Wu, G., Wang, L.: Link: Linear kernel for lidar-based 3d perception. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 1105–1115 (2023)

33. Mao, J., Xue, Y., Niu, M., Bai, H., Feng, J., Liang, X., Xu, H., Xu, C.: Voxel transformer for 3d object detection. In: *Int. Conf. Comput. Vis.* pp. 3164–3173 (2021)
34. Milioto, A., Vizzo, I., Behley, J., Stachniss, C.: Rangenet++: Fast and accurate lidar semantic segmentation. In: *2019 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. pp. 4213–4220. IEEE (2019)
35. Park, J., Kim, C., Kim, S., Jo, K.: Pscnet: Fast 3d semantic segmentation of lidar point cloud for autonomous car using point convolution and sparse convolution network. *Expert Systems with Applications* **212**, 118815 (2023)
36. Puy, G., Boulch, A., Marlet, R.: Using a waffle iron for automotive point cloud semantic segmentation. In: *Int. Conf. Comput. Vis.* pp. 3379–3389 (2023)
37. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: *Adv. Neural Inform. Process. Syst.* vol. 30 (2017)
38. Qian, G., Li, Y., Peng, H., Mai, J., Hammoud, H., Elhoseiny, M., Ghanem, B.: Pointnext: Revisiting pointnet++ with improved training and scaling strategies. In: *Adv. Neural Inform. Process. Syst.* vol. 35, pp. 23192–23204 (2022)
39. Razani, R., Cheng, R., Taghavi, E., Bingbing, L.: Lite-hdseg: Lidar semantic segmentation using lite harmonic dense convolutions. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 9550–9556. IEEE (2021)
40. Tang, H., Liu, Z., Zhao, S., Lin, Y., Lin, J., Wang, H., Han, S.: Searching efficient 3d architectures with sparse point-voxel convolution. In: *Eur. Conf. Comput. Vis.* pp. 685–702. Springer (2020)
41. Tatarchenko, M., Park, J., Koltun, V., Zhou, Q.Y.: Tangent convolutions for dense prediction in 3d. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 3887–3896 (2018)
42. Thomas, H., Qi, C.R., Deschaud, J.E., Marcotegui, B., Goulette, F., Guibas, L.J.: Kpconv: Flexible and deformable convolution for point clouds. In: *Int. Conf. Comput. Vis.* pp. 6411–6420 (2019)
43. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: *Adv. Neural Inform. Process. Syst.* vol. 30 (2017)
44. Wolpert, D.H., Macready, W.G.: No free lunch theorems for optimization. *IEEE transactions on evolutionary computation* **1**(1), 67–82 (1997)
45. Wu, B., Wan, A., Yue, X., Keutzer, K.: Squeezeseg: Convolutional neural nets with recurrent crf for real-time road-object segmentation from 3d lidar point cloud. In: *2018 IEEE international conference on robotics and automation (ICRA)*. pp. 1887–1893. IEEE (2018)
46. Wu, B., Zhou, X., Zhao, S., Yue, X., Keutzer, K.: Squeezesegv2: Improved model structure and unsupervised domain adaptation for road-object segmentation from a lidar point cloud. In: *2019 international conference on robotics and automation (ICRA)*. pp. 4376–4382. IEEE (2019)
47. Wu, X., Lao, Y., Jiang, L., Liu, X., Zhao, H.: Point transformer v2: Grouped vector attention and partition-based pooling. In: *Adv. Neural Inform. Process. Syst.* vol. 35, pp. 33330–33342 (2022)
48. Xiao, A., Huang, J., Guan, D., Cui, K., Lu, S., Shao, L.: Polarmix: A general data augmentation technique for lidar point clouds. In: *Adv. Neural Inform. Process. Syst.* vol. 35, pp. 11035–11048 (2022)
49. Xiao, P., Shao, Z., Hao, S., Zhang, Z., Chai, X., Jiao, J., Li, Z., Wu, J., Sun, K., Jiang, K., et al.: Pandaset: Advanced sensor suite dataset for autonomous driving. In: *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. pp. 3095–3101. IEEE (2021)

50. Xu, C., Wu, B., Wang, Z., Zhan, W., Vajda, P., Keutzer, K., Tomizuka, M.: Squeezesegv3: Spatially-adaptive convolution for efficient point-cloud segmentation. In: *Eur. Conf. Comput. Vis.* pp. 1–19. Springer (2020)
51. Xu, J., Zhang, R., Dou, J., Zhu, Y., Sun, J., Pu, S.: Rpvnet: A deep and efficient range-point-voxel fusion network for lidar point cloud segmentation. In: *Int. Conf. Comput. Vis.* pp. 16024–16033 (2021)
52. Xu, W., Cai, Y., He, D., Lin, J., Zhang, F.: Fast-lio2: Fast direct lidar-inertial odometry. *IEEE Transactions on Robotics* **38**(4), 2053–2073 (2022)
53. Xu, X., Kong, L., Shuai, H., Liu, Q.: Frnet: Frustum-range networks for scalable lidar segmentation. *arXiv preprint arXiv:2312.04484* (2023)
54. Yan, X., Gao, J., Li, J., Zhang, R., Li, Z., Huang, R., Cui, S.: Sparse single sweep lidar point cloud segmentation via learning contextual shape priors from scene completion. In: *AAAI*. vol. 35, pp. 3101–3109 (2021)
55. Yan, X., Gao, J., Zheng, C., Zheng, C., Zhang, R., Cui, S., Li, Z.: 2dpass: 2d priors assisted semantic segmentation on lidar point clouds. In: *Eur. Conf. Comput. Vis.* pp. 677–695. Springer (2022)
56. Yan, X., Zheng, C., Li, Z., Wang, S., Cui, S.: Pointasnl: Robust point clouds processing using nonlocal neural networks with adaptive sampling. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 5589–5598 (2020)
57. Yang, J., Li, C., Dai, X., Gao, J.: Focal modulation networks. In: *Adv. Neural Inform. Process. Syst.* vol. 35, pp. 4203–4217 (2022)
58. Yang, J., Li, C., Zhang, P., Dai, X., Xiao, B., Yuan, L., Gao, J.: Focal self-attention for local-global interactions in vision transformers. *arXiv preprint arXiv:2107.00641* (2021)
59. Zhang, Y., Zhou, Z., David, P., Yue, X., Xi, Z., Gong, B., Foroosh, H.: Polarnet: An improved grid representation for online lidar point clouds semantic segmentation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 9601–9610 (2020)
60. Zhao, L., Xu, S., Liu, L., Ming, D., Tao, W.: Svaseg: Sparse voxel-based attention for 3d lidar point cloud semantic segmentation. *Remote Sensing* **14**(18), 4471 (2022)
61. Zhu, X., Zhou, H., Wang, T., Hong, F., Ma, Y., Li, W., Li, H., Lin, D.: Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 9939–9948 (2021)
62. Zhuang, Z., Li, R., Jia, K., Wang, Q., Li, Y., Tan, M.: Perception-aware multi-sensor fusion for 3d lidar semantic segmentation. pp. 16260–16270 (10 2021). <https://doi.org/10.1109/ICCV48922.2021.01597>