

# SVQNet: Sparse Voxel-Adjacent Query Network for 4D Spatio-Temporal LiDAR Semantic Segmentation

Xuechao Chen<sup>1\*</sup> Shuangjie Xu<sup>2\*</sup> Xiaoyi Zou<sup>3</sup> Tongyi Cao<sup>3</sup> Dit-Yan Yeung<sup>2 †</sup> Lu Fang<sup>1 †</sup>  
<sup>1</sup>Tsinghua University <sup>2</sup>Hong Kong University of Science and Technology <sup>3</sup>Deeprouete.ai

## Abstract

LiDAR-based semantic perception tasks are critical yet challenging for autonomous driving. Due to the motion of objects and static/dynamic occlusion, temporal information plays an essential role in reinforcing perception by enhancing and completing single-frame knowledge. Previous approaches either directly stack historical frames to the current frame or build a 4D spatio-temporal neighborhood using KNN, which duplicates computation and hinders real-time performance. Based on our observation that stacking all the historical points would damage performance due to a large amount of redundant and misleading information, we propose the Sparse Voxel-Adjacent Query Network (SVQNet) for 4D LiDAR semantic segmentation. To take full advantage of the historical frames high-efficiently, we shunt the historical points into two groups with reference to the current points. One is the Voxel-Adjacent Neighborhood carrying local enhancing knowledge. The other is the Historical Context completing the global knowledge. Then we propose new modules to select and extract the instructive features from the two groups. Our SVQNet achieves state-of-the-art performance in LiDAR semantic segmentation of the SemanticKITTI benchmark and the nuScenes dataset.

## 1. Introduction

Serving as a robust 3D perception solution, LiDAR-based perception is under enthusiastic exploration by researchers, among which 3D LiDAR semantic segmentation, aiming at assigning a category label to each point in the whole LiDAR scene at the semantic level, is of great significance in autonomous driving and robotics. Recently, a large number of literature [33, 43, 45, 28, 13, 19] concentrates on semantic segmentation within a single frame. However, the information in a single frame is affected by multiple factors: 1) occlusion problems caused by obstacles or the movement of ego-car, leading to incomplete information of the

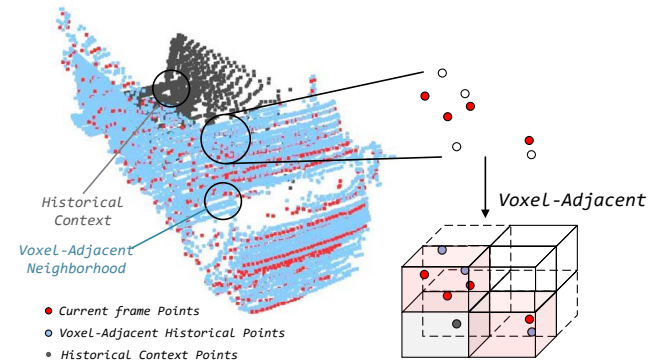


Figure 1. 4D spatio-temporal LiDAR points for a truck. We shunt the historical points into 1) *Voxel-Adjacent* points that lie in the voxel containing current frame points; 2) The remaining points named *Historical Context* points whose features will be adaptively fused to complete the missing features in the current frame.

occluded objects; and 2) ambiguity between similar point clusters, for example, the fence looks similar to one side of a big truck, which severely degrades the performance of single-frame based LiDAR semantic segmentation.

To eliminate the distortion within the single frame, the use of sequential knowledge has attracted widespread attention [24, 44, 42, 1, 39] as the LiDAR continuously transmits and receives sensory data. Therefore, 4D Spatio-temporal information (LiDAR video) is increasingly playing an essential role in reinforcing perception by enhancing and completing single-frame knowledge. Classic temporal methods [1, 48] directly stack frames in the last few timestamps by adding additional channel  $t$  to the coordinates  $xyz$  of each point, which is quite straightforward but superimposing all historical points without any selection brings redundancy, masking the useful temporal knowledge and weakens the benefits of the time series.

To model spatio-temporal relationship instead of stacking all frames, approaches based on KNN or radius neighbors query [21, 4, 31, 37] apply point-wise nearest neighbor search methods to extract instructive features across time and space. However, these approaches will not only fail when the target object is moving at high speed but also bear the high complexity of searching algorithms that lead to the

\*Equal contribution

†Corresponding authors

inability to adopt long time-series information. Other approaches based on RNN [15, 10, 18] or memory [9, 23] using a recurrent neural network or sequence-independent storage to memorize the instructive features from historical frames, can model the long sequence knowledge. Nevertheless, these approaches are unable to align recurrent features under sparse representation and thus adopt a range-image view [8, 38], which are impossible to gain from the sparse representation [12] of point clouds.

To efficiently extract valuable spatio-temporal features in 3D voxel representation, we propose a Sparse Voxel-Adjacent Query Network (*SVQNet*). Our *SVQNet* shunts the historical information into two groups based on the observation from Fig. 1: 1) *Voxel-Adjacent Neighborhood*: historical points around points in current frame can *enhance* the spatial semantic features from sparse to dense across time to disambiguate current frame semantics; 2) *Historical Context*: some occlusion can be *completed* from multiple frames in a learning manner, by activating valuable historical context according to current voxel features. Unlike previous work [31, 21] requiring the calculation and sorting of distances between current and historical points to find nearest neighbors, the search of *Voxel-Adjacent* is highly efficient, which performs sparse hash query from current points to historical points, under several scales from small to big, acquiring spatio-temporal neighbors from near to far. The sparse hash algorithm allows us to reduce the complexity from quadratic to linear, which further endows us with real-time performance. The proposed *SVQNet* achieves state-of-the-art performance on SemanticKITTI [2] and nuScenes [3] datasets while maintaining a real-time runtime. Our main contributions are as followed:

- The Spatio-temporal information is formulated as *enhancing* and *completing* in the first time, with a novel *Spatio-Temporal Information Shunt* module to efficiently shunt the stream of historical information.
- An efficient *Sparse Voxel-Adjacent Query* module is proposed to search instructive neighbors in 4D sparse voxel space, and extract knowledge from the *Voxel-Adjacent Neighborhood*.
- The learnable *Context Activator* is introduced to activate and extract historical *completing* information.
- We furthermore introduce a lightweight *Temporal Feature Inheritance* method to collect features of historical frames and reuse them in the current frame.

## 2. Related Work

### 2.1. LiDAR Semantic Segmentation

3D semantic segmentation is to classify each LiDAR point with a semantic label. Early work is mainly based on the indoor semantic dataset. Pointnet [27] treated point cloud data as one-dimensional and directly applied MLPs to

extract features. RandLA-Net [14] proposed a local feature aggregation method based on K Nearest Neighbor (KNN). KPConv [35] proposed a novel convolution on point clouds, which is called kernel point convolution. Others [7, 12] employed sparse convolution which significantly speeds up 3D convolution and improves the performance.

In recent years, with the advent of outdoor datasets [2, 3], more and more 3D segmentation methods for large scenes are proposed. JS3C-Net [41] utilized shape priors from the scene completion task to help semantic segmentation. SPVNAS [33] proposed a two-branch Point-Voxel convolution method to extract sparse features of point clouds and a method of automatically searching the best model construct, which is called Neural Architecture Search. DRINet [43] proposed a dual representation that contains Point-Voxel and Voxel-Point feature extraction. Some research transferred Cartesian coordinates to polar coordinates [46] and cylindrical coordinates [48]. However, these methods only use features from a single frame and lack context information in temporal space.

### 2.2. Temporal LiDAR Perception

Temporal information, namely 4D Spatio-temporal information, is usually considered useful information in LiDAR perception. Recently, lots of work made efforts on temporal LiDAR detection [15, 24, 44, 18], motion prediction [39, 42], scene flow estimation [29, 26, 20]. As for temporal LiDAR semantic segmentation, it aims to utilize 4D Spatio-temporal information to improve the performance of semantic segmentation. SpSequenceNet [31] employed KNN to gather 4D Spatio-temporal information globally and locally. 4D MinkNet [7] directly processed 4D spatio-temporal point clouds using high-dimensional convolutions. Some work [4, 23, 30] recurrently utilized sequential information. DeepTemporalSeg [8] employed dense blocks and utilized depth separable convolutions to explore the temporal consistency. LMNet [5] effectively performs scene segmentation by accurately distinguishing moving and static objects. 4DMOS [25] utilizes sparse 4D convolutions to extract spatio-temporal features and predict moving object scores for each point. Previous work can be summarized into three categories: directly frame stacking [1, 17], KNN based [31, 7], RNN/memory based [4, 23, 30]. Nevertheless, the methods mentioned above not only lack the mining of the effective part of the 4D information but also suffer from the high computation of 4D processing. In contrast, our method takes full advantage of the historical frames high-efficiently owing to the shunt of *Voxel-Adjacent Neighborhood* and *Historical Context*.

## 3. Method

**Architecture.** As shown in Fig. 1, we shunt the historical points into two phases: 1) By modeling *Voxel-Adjacent*

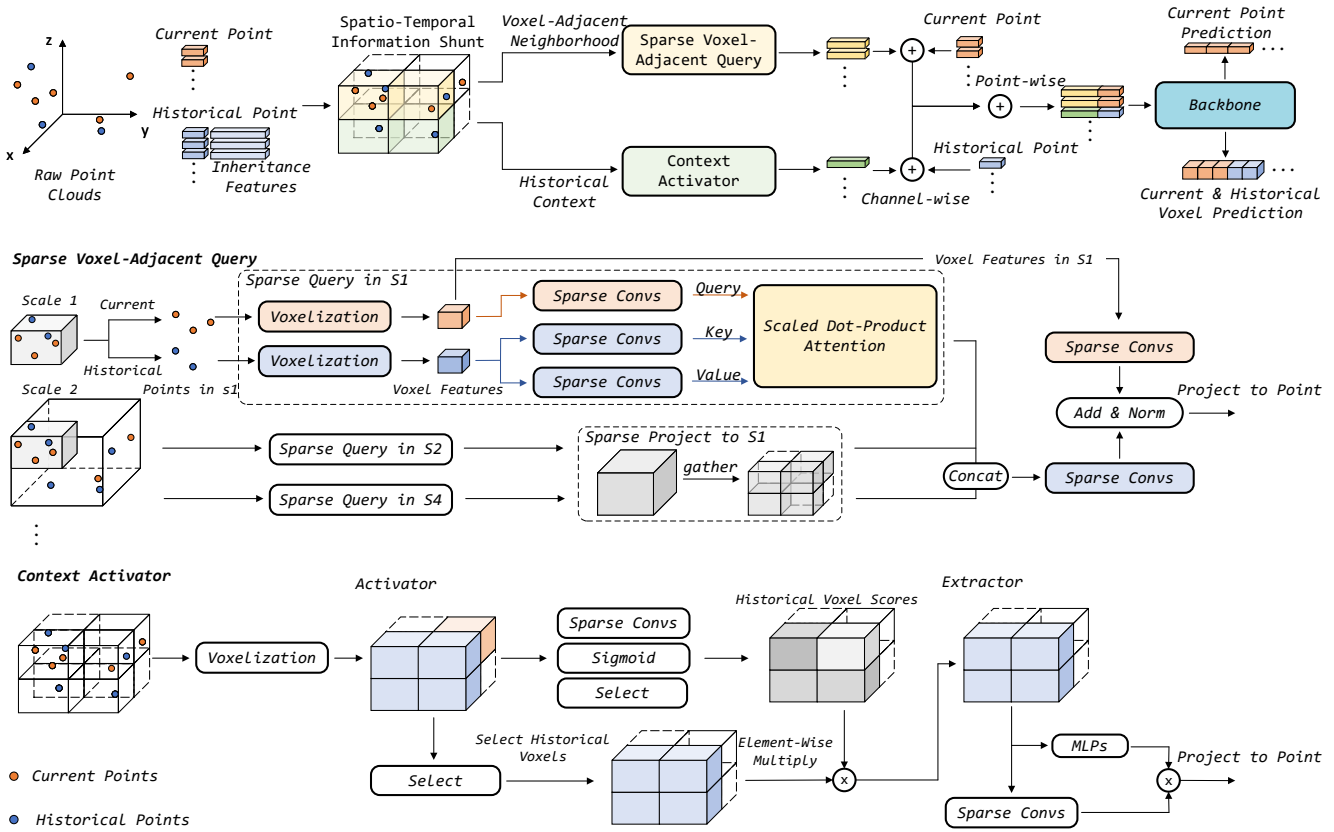


Figure 2. The architecture of our proposed *SVQNet*. *Spatio-Temporal Information Shunt (STIS)* shunts historical sequences into *Voxel-Adjacent Neighborhood* and *Historical Context* information high-efficiently. The shunted streams are then fed into *Sparse Voxel-Adjacent Query (SVAQ)* module and *Context Activator (CA)* module, where the former aggregates multi-scale voxel features from historical *Voxel-Adjacent Neighborhood* to enhance current features with self-attention, and the latter activates *Historical Context* features with a learnable scoring strategy to dynamic select the global context that is truly instructive. Finally, the output attentive features from the two modules are concatenated and fed into our backbone.

*Neighborhood*, we enhance the current points using adjacent historical points. 2) The non-adjacent historical points that we call *Historical Context* help complete the current scan. And we show the modeling processes of the two phases in Sec. 3.1. Then the divided two streams are fed into *Sparse Voxel-Adjacent Query* (detailed in Sec. 3.2) and *Context Activator* (detailed in Sec. 3.3) respectively, as pipeline demonstrated in the top row of Fig. 2.

**Plugin properties.** Shown as Fig. 2, the proposed approach is a kind of “Plugin” which can be used as a plugin to be inserted in mainstream backbone networks, making those networks able to take advantage of temporal information. The experiment on other methods is detailed in Sec. 4.4.

### 3.1. Spatio-Temporal Information Shunt

As shown in Fig. 2, the role of *Spatio-Temporal Information Shunt (STIS)* is to model the *Voxel-Adjacent Neighborhood* and *Historical Context* swiftly. Specifically, *STIS* takes in both current points  $P^c = \{p_i^c, i = 1, \dots, N\}$  where  $p_i^c \in \mathbb{R}^5$  and historical points  $P^h = \{p_j^h, j = 1, \dots, M\}$  where  $p_j^h \in \mathbb{R}^{5+d}$ . The original point features contain

$x, y, z$ , intensity, and timestamp.  $d$  represents the dimension of inheritance features from historical frames later detailed in Sec. 3.4. All historical coordinates  $x, y, z$  have been converted in the current coordinate system by translation and rotation according to the pose matrix of ego-motion.

**Voxelization** is to divide 3D space into voxels with size  $w, l, h$  and then assign each point to the voxel it lies in. In practice, the point clouds will be projected to voxels at different scales  $s$  (a positive number), which means that the voxel size will become  $s \times w, s \times l, s \times h$ . Therefore, the voxelized coordinate for a point under scale  $s$  is  $\left\lfloor \frac{x}{s \times w} \right\rfloor, \left\lfloor \frac{y}{s \times l} \right\rfloor, \left\lfloor \frac{z}{s \times h} \right\rfloor$ . Given point clouds  $P$ , we can apply voxelization to get voxelized coordinates of all points, and these unique voxelized coordinates form the voxel set  $V^s = \{v_k^s, k = 1, \dots, L\}$ . Each  $v_k^s$  contains a voxelized coordinate and the corresponding features for this voxel  $f_k^s$ , which is aggregated from point features inside the voxel by applying DynamicVFE [47]. To simplify symbols, we hide  $s$  if all variables in one formula are in the same scale.

**Voxel-Adjacent Neighborhood Modeling.** Previous attempts [4, 37, 31, 21] have proved that building a lo-

cal spatio-temporal neighborhood can enhance the performance in sequential perception. Unlike previous KNN or radius-based methods [31, 21], we model our proposed *Voxel-Adjacent Neighborhood* with real-time inference speed. *Voxel-Adjacent Neighborhood* search is to query historical voxel features whose voxelized coordinates also exist in the current frame. The physical meaning is to use the method of coincidence of voxelized coordinates to quickly find the nearest neighbor of the current points from the historical points. Given current point clouds  $P^c$  and historical point clouds  $P^h$  whose points have been projected into the coordinate system of  $P^c$ , we can get the corresponding current voxel set  $V^c$  and historical voxel set  $V^h$  by voxelization under scale  $s$ . Then the queried historical voxel set  $V^q$  can be obtained by the  $\Psi_{\text{query}}$  function:

$$\begin{aligned} V^q &= \Psi_{\text{query}}(V^h, V^c) \\ &= \text{HashQuery}(\mathcal{C}(V^c), \mathcal{C}(V^h)). \end{aligned} \quad (1)$$

$\mathcal{C}$  is the function to get the voxelized coordinate if input one voxel or a set of coordinates if input a set of voxels. Using the coordinates as hash keys, we can use the sparse HashQuery function to map  $V^h$  to  $V^c$ :

$$\text{HashQuery}(v_i^c, V^h) = \begin{cases} \emptyset, & \mathcal{C}(v_i^c) \notin \mathcal{C}(V^h), \\ v_j^h, & \mathcal{C}(v_i^c) = \mathcal{C}(v_j^h). \end{cases} \quad (2)$$

$\emptyset$  denotes if there's no voxel under the corresponding coordinate in  $V^h$ , a zero placeholder will be padded in the query results. Otherwise, HashQuery returns *Voxel-Adjacent Neighbor*  $v_j^h$  which has the same coordinate as the input  $v_i^c$ . The previous methods employ KNN or radius-based search to find 4D neighbors and the time complexity is  $O(NM)$  ( $M > N$ ). Ours benefits from the hash algorithm, whose time complexity is  $O(N)$ . TorchSparse [32] is used as the implementation of sparse hash query in this paper.

By the same procedures under multiple scales, we can get the *Voxel-Adjacent Neighborhood* from small scale to big, from small scope to large. Thus we model a multi-scale voxel-aware neighborhood search procedure under multiple scales to obtain the queried historical voxel sets  $\{V^{q_s} | s = s_1, s_2, s_4\}$ , where the  $s_i = i$  that is 1, 2, 4 respectively in our setting. Note that  $V^c$  has the same length as  $V^q$  and they are index-corresponding. Then the built *Voxel-Adjacent Neighborhood* carries local 4D Spatio-temporal information, which is extracted as enhancing features to the current frame by *SVAQ*. More details about *SVAQ* are described in Sec. 3.2.

**Historical Context Modeling.** With the motion of the ego vehicle, part of the scene may not be visible caused of occlusion from other dynamic agents or static obstacles, leading to the target object being too sparse to be recognized. In this case, points in historical frames can help to complete the lost information by aggregating *Historical Context*

from multiple frames. *Historical Context* denotes those historical voxel features that are not queried as *Voxel-Adjacent Neighborhood*. To activate valuable contexts selectively, we first select historical voxels that are not queried, which is  $V^n$ , and then provide these features to the *Context Activator* module for further *completing* features extraction. Similar to the modeling of *Voxel-Adjacent Neighborhood*, the unqueried voxels are named *Historical Context*, which can be modeled by a negative function to  $\Psi_{\text{query}}$ :

$$V^n = \Psi_{\text{unquery}}(V^h, V^c), \quad (3)$$

where  $V^h$  is the historical voxels and  $V^c$  is the current voxels.  $V^n$  is the remaining elements of historical voxel set  $V^h$  that are not queried by  $\Psi_{\text{query}}$ . The scale of *Historical Context*  $s_1 = 1$  in our setting. Subsequently, the formed *Historical Context* carrying global 4D Spatio-temporal information is fed to *Context Activator* module detailed in Sec. 3.3.

### 3.2. Sparse Voxel-Adjacent Query

The key to enhancing the current features is to extract the relative historical features with adaptive learning. Therefore, we propose a *Sparse Voxel-Adjacent Query (SVAQ)* module to benefit from *Voxel-Adjacent Neighborhood* with Transformer attention. As illustrated in Fig. 2, given voxel set  $\{V^c, V^q\}$  under scale  $s$ , voxel features from current frame  $V^c$  are fed into one Sparse-Convolution [7, 12] layer (SPC) to generate *Query* features. The queried historical voxels  $V^q$  are fed into two independent Sparse-Convolution layers to generate *Key* and *Value* features. Later the generated *Query*, *Key* of dimension  $d_k$  and *Value* are fed into Scaled Dot-Product Attention [36] to extract attentive local 4D spatio-temporal voxel features  $T$  under scale  $s$ , which can be written as:

$$T = \text{SPC}(V^q) \cdot \text{Softmax}\left(\frac{\text{SPC}(V^c) \cdot \text{SPC}(V^q)^T}{\sqrt{d_k}}\right). \quad (4)$$

$T$  contains local dependencies between current features  $V^c$  and historical *Voxel-Adjacent* features  $V^q$ , making the repetitive historical features more instructive to *enhance* the same voxel in current frame. Then the attentive features under three scales  $\{s_1, s_2, s_4\}$  extracted by Scaled Dot-Product Attention are concatenated at the feature channel and fused by another three SPC layers. Specially, the attentive features  $T^{s_2}, T^{s_4}$  should be projected back to scale 1. The calculation can be denoted as

$$T_o = \text{SPC}(T^{s_1} \oplus \text{Proj}(T^{s_2}) \oplus \text{Proj}(T^{s_4})), \quad (5)$$

where  $\oplus$  denotes channel-wise concatenation, Proj denotes sparse projection to scale  $s_1 = 1$ , which can be implemented as a HashQuery task by querying  $T^s$  for  $s > 1$  according to the scale-normlized coordinates of  $T^{s_1}$  to scatter features from low resolution to high resolution:

$$\text{Proj}(T^s) = \text{HashQuery}(\mathcal{C}(T^{s_1})/s, \mathcal{C}(T^s)). \quad (6)$$

In particular, we take  $V^c$  under scale  $s_1$  as a skip connection followed by three SPC layers. The final output features  $O_v$  of *SVAQ* can be represented as:

$$O_v = \text{Prop}(\text{Norm}(\text{SPC}(V^{cs_1}) + T_o)), \quad (7)$$

where  $O_v \in \mathbb{R}^{N \times N_C}$ ,  $N$  is current point number and  $N_C$  is the feature channel number.  $\text{Prop}$  denotes the projection from voxels to points,  $\text{Norm}$  denotes BatchNorm layer [16].  $\text{Prop}$  can be formulated as the inversion of *Voxelization*, which outputs point features by assigning the feature of each point with the voxel feature it lies in. The number of the projected point features equals the number of the input current points and they are index-corresponding. In summary, our *SVAQ* module acts in a multi-head attention way to encode *Voxel-Adjacent* features to current point features.

### 3.3. Context Activator

Since the information in the current frame is incomplete caused of dynamic or static occlusion and the sparsity of LiDAR beams, we utilize 4D Spatio-temporal information to *complete* current features with the automatically selected valuable *Historical Context*. Yet directly stacking historical points damages the efficiency and brings a lot of repetitive information, we propose learning-based *Context Activator* (CA) shown in Fig. 2 to flexibly activate and extract the global *Context* that are truly instructive. Firstly, the *Activator* generates voxel scores  $S$  for each context voxel of  $V^n$ , with the help of current voxels  $V^c$  as reference. It employs a three-layer SPC along with a Sigmoid layer Sigmoid to scoring  $V^n$  with a predicted score between 0 to 1. Then we multiply the *Context*  $V^n$  with corresponding voxel scores to perform an element-wise attentive selection with a defined threshold  $S_{th}$ . Note that with  $S_{th}$ , the number of historical voxels reserved after the selection process can be controlled to balance the performance and computation at inference time. The *Activator* can be represented as:

$$S = \text{Select}_{V^n}(\text{Sigmoid}(\text{SPC}(V^n \circ V^c))), \quad (8)$$

$$R = \text{Select}_{S > S_{th}}(V^n \otimes S), \quad (9)$$

where  $\otimes$  denotes element-wise multiplication,  $\circ$  denotes the union set of two voxel sets. The  $\text{Select}_{V^n}$  selects scores generated by  $V^n$ . The  $\text{Select}_{S > S_{th}}$  selects scores bigger than  $S_{th}$ , which is only activated at inference time. In the training process, the  $R$  is obtained by  $R = V^n \otimes S$  to keep samples in the low score as negative samples to balance the learning of *Activator*.

After the dynamic selection, the *Extractor* performs channel-wise self-attention on activated *Historical Context*. *Extractor* employs a three-layer MLP and a three-layer SPC to extract features from  $R$  respectively, where the former takes voxels as points to extract inner-voxel features

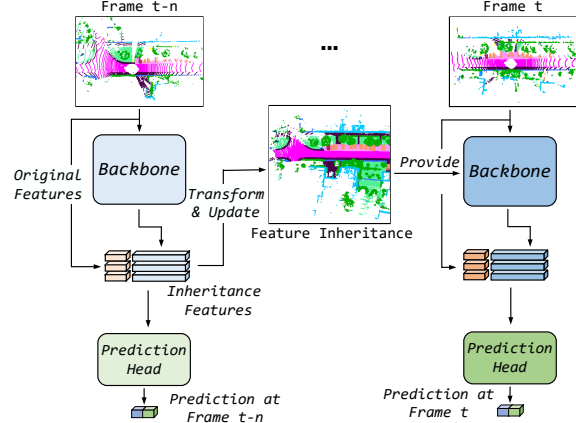


Figure 3. *Temporal Feature Inheritance*. We store the historical point features, namely inheritance features, along with the meta-information  $x, y, z$ , intensity, and timestamp. Then we reuse them by transforming and updating coordinates in meta-information into coordinates under the current frame with the pose matrix.

and the latter captures the inter-voxel relationship. Then the final output  $O_c$  is obtained by the product of two groups:

$$O_c = \text{Prop}(\text{MLP}(R) \cdot \text{SPC}(R)), \quad (10)$$

where  $O_c \in \mathbb{R}^{M' \times N_C}$ ,  $M'$  is the number of points in activated *Historical Context* voxels and  $N_C$  denotes channel number same as  $O_v$ . To sum up, our CA acts in a learnable way to activate *Historical Context* selectively.

### 3.4. Temporal Feature Inheritance

We find that it is a waste to discard the high-dimensional features of previous frames and relearn them during the process of the current frame. Therefore, as demonstrated in Fig. 3, we propose a simple but efficient method named *Temporal Feature Inheritance* (TFI) to inherit sparse features from historical sequences. Assume that *SVQNet* takes  $n$  historical frames as input, all historical point features for frame  $t$  (points from frame  $t-n$  to  $t-1$ ) are ready when the process of frame  $t-1$  is finished by *SVQNet*, which is like a sliding window on time series. Based on the above observation, we concatenate high-dimensional point features from historical frames at the end of *SVQNet* backbone with the corresponding meta-information including  $x, y, z$ , intensity, and timestamp, which is then stored in the buffer memory. The historical meta-information is mainly used for the transformation of the coordinates of the points from the coordinate system of the historical frame into the current frame. When processing the next frame  $t$ , we fetch historical point features from the memory and project  $x, y, z$  to the coordinate system of the current frame by pose transformation, and then the projected coordinates can be used for *SVAQ*. Furthermore, the stored absolute timestamp should be transformed into a relative timestamp by subtraction with

Methods	road	sidewalk	parking	other ground	building	car	truck	bicycle	motorcycle	other vehicle	vegetation	trunk	terrain	person	bicyclist	motorcyclist	fence	pole	traffic sign	mov. car	mov. bicyclist	mov. person	mov. motorcycle	mov. truck	mov. other veh.	mIoU (%)
TangentConv[34]	83.9	64.0	38.3	15.3	85.8	84.9	21.1	2.0	18.2	18.5	79.5	43.2	56.7	1.6	0.0	0.0	49.1	36.4	31.2	40.3	1.1	6.4	1.9	<b>42.2</b>	<b>30.1</b>	34.1
DarkNet53Seg[2]	91.6	75.3	64.9	27.5	85.2	84.1	20.0	30.4	32.9	20.7	78.4	50.7	64.8	7.5	0.0	0.0	56.5	38.1	53.3	61.5	14.1	15.2	0.2	37.8	28.9	41.6
SpSequenceNet[31]	90.1	73.9	57.6	27.1	91.2	88.5	29.2	24.0	26.2	22.7	84.0	66.0	65.7	6.3	0.0	0.0	66.8	50.8	48.7	53.2	41.2	26.2	36.2	0.1	2.3	43.1
TemporallidarSeg[9]	91.8	75.8	59.6	23.2	89.8	92.1	39.2	47.7	40.9	35.0	82.3	62.5	64.7	14.4	0.0	0.0	63.8	52.6	60.4	68.2	42.8	40.4	12.9	2.1	12.4	47.0
KPConv[35]	86.5	70.5	58.4	26.7	90.8	93.7	<b>42.5</b>	44.9	47.2	38.6	84.6	70.3	66.0	21.6	0.0	0.0	64.5	57.0	53.9	68.1	67.4	67.5	47.2	0.5	0.5	51.2
Cylinder3D[48]	90.4	74.9	66.3	32.1	92.4	93.8	41.2	<b>67.6</b>	<b>63.3</b>	37.6	85.4	72.8	68.1	12.9	<b>0.1</b>	<b>0.1</b>	65.8	62.6	61.3	68.1	60.0	63.1	0.4	0.0	0.1	51.5
<b>SVQNet(ours)</b>	<b>93.2</b>	<b>80.5</b>	<b>71.6</b>	<b>37.0</b>	<b>93.7</b>	<b>96.1</b>	40.4	64.4	60.3	<b>60.9</b>	<b>87.3</b>	<b>76.7</b>	<b>72.3</b>	<b>27.4</b>	0.0	0.0	<b>72.6</b>	<b>68.4</b>	<b>71.0</b>	<b>80.5</b>	<b>72.4</b>	<b>84.7</b>	<b>91.0</b>	3.9	7.5	<b>60.5</b>
Improvements $\Delta$	<b>+1.4</b>	<b>+4.7</b>	<b>+5.3</b>	<b>+4.9</b>	<b>+1.3</b>	<b>+2.3</b>	-2.1	-3.2	-3.0	<b>+22.3</b>	<b>+1.9</b>	<b>+3.9</b>	<b>+4.2</b>	<b>+5.8</b>	-0.1	-0.1	<b>+5.8</b>	<b>+5.8</b>	<b>+9.7</b>	<b>+12.3</b>	<b>+5.0</b>	<b>+17.2</b>	<b>+43.8</b>	-38.3	-22.6	<b>+9.0</b>

Table 1. The experiment results on the semantic segmentation of SemanticKITTI test set (multi-scan phase). All listed methods utilized 4D spatio-temporal information according to their paper.  $\Delta$ : comparing with the best previous results for each class. (mov. denotes moving.)

Methods	barrier	bicycle	bus	car	con. vehicle	motorcycle	pedestrian	traffic cone	trailer	truck	surface	other flat	sidewalk	terrain	manmade	vegetation	mIoU (%)
PolarNet[46]	72.2	16.8	77.0	86.5	51.1	69.7	64.8	54.1	69.7	63.5	96.6	67.1	77.7	72.1	87.1	84.5	69.4
JS3C-Net[41]	80.1	26.2	87.8	84.5	55.2	72.6	71.3	66.3	76.8	71.2	96.8	64.5	76.9	74.1	87.5	86.1	73.6
Cylinder3D[48]	82.8	29.8	84.3	89.4	63.0	79.3	77.2	73.4	<b>84.6</b>	69.1	<b>97.7</b>	70.2	80.3	75.5	90.4	87.6	77.2
SPVNAS[33]	80.0	30.0	91.9	90.8	64.7	79.0	75.6	70.9	81.0	74.6	97.4	69.2	80.0	76.1	89.3	87.1	77.4
AF2S3Net[6]	78.9	<b>52.2</b>	89.9	84.2	<b>77.4</b>	74.3	77.3	72.0	83.9	73.8	97.1	66.5	77.5	74.0	87.7	86.8	78.3
<b>SVQNet(ours)</b>	<b>84.5</b>	41.8	<b>93.3</b>	<b>92.5</b>	69.1	<b>85.5</b>	<b>83.7</b>	<b>78.3</b>	84.5	<b>77.5</b>	97.1	<b>70.3</b>	<b>81.6</b>	<b>77.9</b>	<b>91.8</b>	<b>90.1</b>	<b>81.2</b>

Table 2. Results on the nuScenes test set. All listed methods employed frames stacking strategy according to their implementation.

the absolute timestamp of the current frame. At the end of *SVQNet* at frame  $t$ , we update the memory with newly extracted features for the  $t + 1$  frame.

## 4. Experiments

### 4.1. Datasets and Metrics

**SemanticKITTI** [2] is a large-scale outdoor point clouds dataset for autonomous driving which was collected by a 64-beam LiDAR sensor. The train set contains 23201 sequential LiDAR scans and the test set contains 20351 sequential scans. The semantic segmentation task is officially divided into two phases. One is the single-scan phase containing 19 semantic classes without distinction of moving or static objects. The other is the multi-scan phase containing 25 semantic classes with the distinction between moving and static objects. To test and ablate our *SVQNet*, we run experiments on the multi-scan phase.

**nuScenes** [3, 11] is a large-scale outdoor multi-modal dataset for autonomous driving. Their point clouds data was collected by a 32-beam LiDAR sensor in sequence. It contains 1,000 scenes and 16 semantic classes with no distinction of moving or static objects. They only annotated LiDAR data one frame every ten frames so the default setting of most methods is multiple scans.

**mIoU** is the mean intersection over the union. The IoU is defined as  $\frac{TP}{TP+FP+FN}$ , where the  $TP$ ,  $FP$ ,  $FN$  represent the true positive, false positive, and false negative of the prediction. The IoU score is first calculated for each class and then the mIoU is obtained by averaging across classes.

Methods	Latency
RandLA-Net [14]	880 ms
SqueezeSegV3 [40]	238 ms
SPVNAS [33]	259 ms
Cylinder3D [49]	170 ms
<b>SVQNet(ours)</b>	<b>97 ms</b>

Table 3. Latency comparison with single-scan methods.

**Implementation details.** We implement DRINet [43] as our baseline, and plus an additional voxel-wise loss to predict a semantic label for each voxel with the majority category of points in the voxel as the ground truth, which is an auxiliary loss to make *CA* module to learn the activation for valuable historical voxels. Other settings are the same as DRINet reported. The point-wise loss is only applied to current points, but the voxel-wise loss is applied both to current voxels and historical voxels.

We adopt 1 current frame and 2 historical frames as the inputs without any down-sampling. All the experiments are conducted on a machine with 8 \* NVIDIA RTX 3090 GPU. The learning rate is set to  $2e - 3$  with an adamW optimizer [22]. The training epoch is set to 40. The dimension  $d$  of inheritance features,  $d_k$  in *SVAQ*, and  $N_C$  in *CA* are all set to 64. The  $S_{th}$  in *CA* is set as  $S_{th} = 0.1$  to reserve activated historical voxel number  $M' \approx 40k$  at inference time. In training, we set the  $S_{th} = 0.0$  to disable the selection.

### 4.2. Results on SemanticKITTI

The multi-scan phase of SemanticKITTI distinguishes moving and static objects. In this phase, we achieve state-of-the-art performance in terms of mIoU. As shown in Tab. 1, our proposed *SVQNet* surpasses the multi-scan

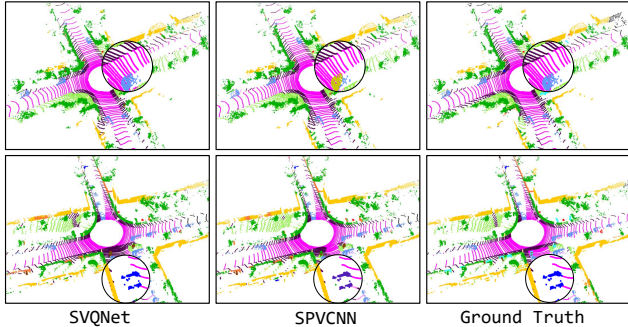


Figure 4. The qualitative visualization results on SemanticKITTI validation set, where the first column is the prediction from *SVQNet*, the second column is the prediction from multi-scan SPVCNN [33] and the third column is the ground truth. Enlarged circles display our good cases (cars and motorcycles).

Backbone+	CA	SVAQ	TFI	mIoU (%)	$\Delta$	Latency
✓				52.8	-	125 ms
✓	✓			53.9	+1.1	61 ms
✓	✓	✓		54.8	+0.9	92 ms
✓	✓	✓	✓	<b>55.3</b>	+0.5	97 ms

Table 4. Ablation study on our proposed modules *CA*, *SVAQ*, and *TFI*. Backbone+ denotes that the backbone network directly stacks 2 historical frames with the current frame. Note that the Backbone+ takes in the same quantity of information as our *SVQNet*.

$M'$	50k	45k	40k	35k	30k
mIoU (%)	55.0	55.2	<b>55.3</b>	55.1	54.9

Table 5. Ablation study on *Context Activator*. As the number of activated contextual voxels  $M'$  varies by threshold  $S_{th}$ , *SVQNet* achieves the best performance when  $M' = 40k$ .

Cylinder3D [48] and KPConv [35], which employed directly stacking strategy according to their implementation, by 9% and 9.3% in terms of mIoU separately. We also obtain a 17.4% performance gain in terms of mIoU than SpSequenceNet [31], which employed KNN to build the spatio-temporal relationship. Besides, our *SVQNet* gets a mIoU improvement of 13.5% compared with TemporalLidarSeg [9], which proposed Temporal Memory Alignment to utilize spatio-temporal information. The results show the excellent capability of our proposed *SVQNet* in distinguishing moving and static semantics. Due to the lack of reported runtime data of multi-scan methods listed in Tab. 1, we compare the latency of our method with some 3D single-scan methods in Tab. 3, showing that our multi-scan method is even faster than previous single-scan methods. Moreover, the qualitative visualization results are shown in Fig. 4.

### 4.3. Results on nuScenes

To further test our proposed *SVQNet*, we run experiments on the nuScenes dataset. As a result, our method ranks 2<sup>nd</sup> on the leaderboard of the semantic segmentation task, in terms of mIoU. As demonstrated in Tab. 2,




Figure 5. The visualization of the attention map of our proposed *Context Activator*, where orange denotes current points, blue denotes activated *Historical Context*, and green denotes deactivated *Historical Context*.

Distance (m)	[0,12)	[12,24)	[24,36)	[36,∞)	all
Backbone+	54.7	50.9	49.3	36.1	52.8
SVQNet (Ours)	56.5	54.5	53.5	42.3	55.3
$\Delta$	+1.8	+3.6	+4.2	+6.2	+2.5

Table 6. The distance-based mIoU results on SemanticKITTI validation set (seq 08, multi-scan phase). The mIoU(%) changes as the distance interval varies from [0, 12) to [36, ∞) (unit: meter). "all" shows the mIoU results of all points from 0 m to ∞. Note that only current points are considered in the evaluation.

compared to JS3C-Net [41], Cylinder3D [48] and SPVNAS [33], which employed directly stacking strategy according to their implementation, our method receives a performance gain of 7.6%, 4% and 3.8%, respectively, in terms of mIoU, which shows the superior performance of our proposed *SVQNet* on the semantic segmentation task.

### 4.4. Ablation Studies

In the ablation study, we conduct several experiments on the validation set (sequence 08, multi-scan phase) of the SemanticKITTI dataset.

**Ablation study on our proposed modules.** As shown in Tab. 4, getting started from our naive multi-scan backbone that stacks all points as input along with 2 historical frames, we get the baseline performance of 52.8% mIoU. By enabling the *CA* module, the network gains the most improvement of 1.1% mIoU, which proves that sequence data does have a lot of redundancy and the proposed *CA* module can activate the truly instructive information. Furthermore, if we activate the *SVAQ* module to enhance features of the current frame, the performance continues to grow 0.9% mIoU, demonstrating that historical knowledge is also crucial to enhance current voxel features. In the end, *TFI* method is applied to inherit previous features, and we get a gain of 0.5%, which proves the effectiveness to reuse previous computed features with proposed *TFI*.

**Ablation study on Context Activator.** To examine the effectiveness of the proposed *CA*, we control the number of activated contextual voxels  $M'$  by adjusting the threshold value  $S_{th}$  when applying selection on *Historical Context*.

Methods	road	sidewalk	parking	other ground	building	car	truck	bicycle	motorcycle	other vehicle	vegetation	trunk	terrain	person	bicyclist	motorcyclist	fence	pole	traffic sign	mov. car	mov. bicyclist	mov. person	mov. motorcycle	mov. truck	mov. other veh.	mIoU (%)	Latency(ms)	GPU Memory(MB)
Backbone+	94.1	82.0	56.8	1.9	90.8	95.4	66.7	40.0	63.3	71.1	88.7	69.1	79.9	28.5	0.0	0.0	60.1	63.0	53.0	63.4	84.2	59.8	0.0	0.0	0.2	52.8	125	8893
SVQNet(Ours)	94.9	83.7	57.0	0.3	88.6	97.3	90.0	47.1	75.8	78.4	89.4	66.9	79.1	32.0	0.0	0.0	47.8	65.5	55.5	74.7	92.5	67.3	0.0	0.0	0.0	55.3	97	6854
$\Delta$	+0.8	+1.7	+0.2	-1.6	-2.2	+1.9	+23.3	+7.1	+12.5	+7.3	+0.7	-2.2	-0.8	+3.5	0.0	0.0	-12.3	+2.5	+2.5	+11.3	+8.3	+7.5	0.0	0.0	-0.2	+2.5	-28	-2039

Table 7. The per-class results on SemanticKITTI validation set (seq 08, multi-scan phase). (mov. denotes moving)

With the growth of  $M'$ , the provided information is more but the inference time is larger. As illustrated in Tab. 5, we vary the number from  $30k$  to  $50k$ , finding that *SVQNet* achieves the highest mIoU 55.3% at  $M' = 40k$ . As  $M'$  increases to  $45k$ , the performance falls, further proving the importance of historical feature selection. In our experiments mentioned above, we fix  $M' = 40k$ .

**Visualization of attention map.** Additionally, we visualize the attention map to directly understand what our proposed *Context Activator* has learned. As shown in Fig. 5, we color the current points orange, activated *Historical Context* blue, and deactivated *Historical Context* green. In this visualization, we fix the number of activated *Historical Context* to about  $40k$  by altering the threshold  $S_{th}$  mentioned in our method. Based on the observation, we conclude that, 1) our proposed *Context Activator* tends to complement the current point clouds at the distant place, where the objects are too sparse to be recognized; 2) our proposed *Context Activator* tends to complete the intricate objects such as cars since the instructive shape priors are significant to semantic segmentation as JS3C-Net [41] proved.

**Comparisons with Backbone+.** To reveal the source of mIoU gain under the same data input, we conduct experiments and make comparisons with the Backbone+, which directly stacks two historical frames with the current frame as the input of the backbone network. We show distance-based mIoU results in Tab. 6. As the distance interval varies from  $[0, 12)$  to  $[36, \infty)$  (unit: meter), the mIoU results of Backbone+ and ours descend because the farther away the LiDAR scene gets, the more sparse the point clouds become. However, ours drops less than Backbone+ when distance increases. And we show the mIoU gap in  $\Delta$ , revealing our robustness withstanding the distance sparsity of LiDAR. In addition, Tab. 7 shows the per-class results of Backbone+ and ours, and also the gap. We can see that, under the input of the same data, ours achieves distinct improvement in intricate objects such as static/moving cars, static/moving people, static trucks, bicycles, motorcycles, and traffic signs. That reveals the source of our mIoU gain from different categories and the strong ability to distinguish static and moving objects. Last but not least, we compare ours with Backbone+ regarding latency and GPU memory. As Tab. 7 demonstrates, under the input of the

SPVCNN	stacking	+ CA&SVAQ&TFI
mIoU (%)	50.04	<b>52.36</b>

Table 8. The mIoU of SPVCNN (w/o or w/ our modules) on the SemanticKITTI dataset (validation set, multi-scan settings).

same data and the same experiment settings, our method is 22.4% faster than Backbone+, and we also use 22.9% less GPU memory, benefiting from the proposed information shunt module *STIS* and context selection module *CA* that avoid wasting computation on redundant information. That exhibits our device-friendly character and the potential for real-time application.

**Plugin properties.** To prove the plugin properties mentioned in Sec. 3, we perform experiments using SPVCNN [33] as our backbone. The network architecture is demonstrated in the top row of Fig. 2. Under the same input, which is two historical frames along with the current frame, equipping with our modules outperforms the SPVCNN baseline employing stacking strategy by 2.32% mIoU, as shown in Tab. 8.

## 5. Conclusion

We propose *Sparse Voxel-Adjacent Query Network* to concentrate on efficiently extracting 4D spatio-temporal features. The *Spatio-Temporal Information Shunt* module is proposed to high-efficiently shunt the 4D spatio-temporal information in two groups, *Voxel-Adjacent Neighborhood* and *Historical Context*. Further, we propose two novel modules, *Sparse Voxel-Adjacent Query* and *Context Activator*, benefiting from the locally *enhancing* information of *Voxel-Adjacent Neighborhood* and globally *completing* information of *Historical Context*. In addition, *Temporal Feature Inheritance* method is introduced to collect the features preserved in historical frames as an input in the current frame. The proposed *SVQNet* reaches state-of-the-art performance on nuScenes and SemanticKITTI leaderboards.

**Acknowledgement** This research has been made possible by funding support from 1) Deeprouete.ai, 2) the Research Grants Council of Hong Kong under the Research Impact Fund project R6003-21, and 3) Natural Science Foundation of China (NSFC) under contract No. 62125106, 61860206003 and 62088102; Ministry of Science and Technology of China under contract No. 2021ZD0109901.



## References

- [1] Mehmet Aygun, Aljosa Osep, Mark Weber, Maxim Maximov, Cyrill Stachniss, Jens Behley, and Laura Leal-Taixé. 4D Panoptic LiDAR Segmentation. In *CVPR*, 2021.
- [2] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. In *ICCV*, 2019.
- [3] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A Multi-modal Dataset for Autonomous Driving. In *CVPR*, 2020.
- [4] Hanwen Cao, Yongyi Lu, Cewu Lu, Bo Pang, Gongshen Liu, and Alan Yuille. ASAP-Net: Attention and Structure Aware Point Cloud Sequence Segmentation. *arXiv preprint arXiv:2008.05149*, 2020.
- [5] Xieyuanli Chen, Shijie Li, Benedikt Mersch, Louis Wiesmann, Jürgen Gall, Jens Behley, and Cyrill Stachniss. Moving Object Segmentation in 3D LiDAR Data: A Learning-based Approach Exploiting Sequential Data. *IEEE RA-L*, 2021.
- [6] Ran Cheng, Ryan Razani, Ehsan Taghavi, Enxu Li, and Bingbing Liu. 2-S3Net: Attentive Feature Fusion with Adaptive Feature Selection for Sparse Semantic Segmentation Network. In *CVPR*, 2021.
- [7] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks. In *CVPR*, 2019.
- [8] Ayush Dewan and Wolfram Burgard. DeepTemporalSeg: Temporally Consistent Semantic Segmentation of 3D LiDAR Scans. In *ICRA*, 2020.
- [9] Fabian Duerr, Mario Pfaller, Hendrik Weigel, and Jürgen Beyerer. LiDAR-based Recurrent 3D Semantic Segmentation with Temporal Memory Alignment. In *3DV*, 2020.
- [10] Emeç Erçelik, Ekim Yurtsever, and Alois Knoll. TempFrustrum Net: 3D Object Detection with Temporal Fusion. In *IEEE IV*, 2021.
- [11] Whye Kit Fong, Rohit Mohan, Juana Valeria Hurtado, Lubing Zhou, Holger Caesar, Oscar Beijbom, and Abhinav Valada. Panoptic nuScenes: A Large-Scale Benchmark for LiDAR Panoptic Segmentation and Tracking. *IEEE RA-L*, 2022.
- [12] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3D Semantic Segmentation with Submanifold Sparse Convolutional Networks. In *CVPR*, 2018.
- [13] Yuenan Hou, Xinge Zhu, Yuexin Ma, Chen Change Loy, and Yikang Li. Point-to-Voxel Knowledge Distillation for LiDAR Semantic Segmentation. In *CVPR*, 2022.
- [14] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. RandLA-Net: Efficient Semantic Segmentation of Large-Scale Point Clouds. In *CVPR*, 2020.
- [15] Rui Huang, Wanyue Zhang, Abhijit Kundu, Caroline Pantofaru, David A Ross, Thomas Funkhouser, and Alireza Fathi. An LSTM Approach to Temporal 3D Object Detection in LiDAR Point Clouds. In *ECCV*, 2020.
- [16] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *ICML*, 2015.
- [17] Lars Kreuzberg, Idil Esen Zulfikar, Sabarinath Mahadevan, Francis Engelmann, and Bastian Leibe. 4D-StOP: Panoptic Segmentation of 4D LiDAR Using Spatio-Temporal Object Proposal Generation and Aggregation. In *ECCV Workshops*, 2023.
- [18] K. Kumar and Samir Al-Stouhi. Real-time Spatial-temporal Context Approach for 3D Object Detection using LiDAR. In *VEHITS*, 2020.
- [19] Xiaoyan Li, Gang Zhang, Hongyu Pan, and Zhenhua Wang. CPGNet: Cascade Point-Grid Fusion Network for Real-Time LiDAR Semantic Segmentation. In *ICRA*, 2022.
- [20] Xingyu Liu, Charles R Qi, and Leonidas J Guibas. FlowNet3D: Learning Scene Flow in 3D Point Clouds. In *CVPR*, 2019.
- [21] Xingyu Liu, Mengyuan Yan, and Jeannette Bohg. MeteorNet: Deep Learning on Dynamic 3D Point Cloud Sequences. In *ICCV*, 2019.
- [22] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [23] Fan Lu, Guang Chen, Yinlong Liu, Zhijun Li, Sanqing Qu, and Tianpei Zou. MoNet: Motion-based Point Cloud Prediction Network. *arXiv preprint arXiv:2011.10812*, 2020.
- [24] Scott McCrae and Avidesh Zakhori. 3D Object Detection For Autonomous Driving Using Temporal Lidar Data. In *ICIP*, 2020.
- [25] Benedikt Mersch, Xieyuanli Chen, Ignacio Vizzo, Lucas Nunes, Jens Behley, and Cyrill Stachniss. Receding Moving Object Segmentation in 3D LiDAR Data Using Sparse 4D Convolutions. *IEEE RA-L*, 2022.
- [26] Gilles Puy, Alexandre Boulch, and Renaud Marlet. FLOT: Scene Flow on Point Clouds Guided by Optimal Transport. In *ECCV*, 2020.
- [27] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *CVPR*, 2017.
- [28] Shi Qiu, Saeed Anwar, and Nick Barnes. Semantic Segmentation for Real Point Cloud Scenes via Bilateral Augmentation and Adaptive Fusion. In *CVPR*, 2021.
- [29] Rishav Rishav, Ramy Battrawy, René Schuster, Oliver Wasenmüller, and Didier Stricker. DeepLiDARFlow: A Deep Learning Architecture For Scene Flow Estimation Using Monocular Camera and Sparse LiDAR. In *IROS*, 2020.
- [30] Peer Schutt, Radu Alexandru Rosu, and Sven Behnke. Abstract Flow for Temporal Semantic Segmentation on the Permutohedral Lattice. In *ICRA*, 2022.
- [31] Hanyu Shi, Guosheng Lin, Hao Wang, Tzu-Yi Hung, and Zhenhua Wang. SpSequenceNet: Semantic Segmentation Network on 4D Point Clouds. In *CVPR*, 2020.
- [32] Haotian Tang, Zhijian Liu, Xiuyu Li, Yujun Lin, and Song Han. TorchSparse: Efficient Point Cloud Inference Engine. In *MLSys*, 2022.
- [33] Haotian Tang, Zhijian Liu, Shengyu Zhao, Yujun Lin, Ji Lin, Hanrui Wang, and Song Han. Searching Efficient 3D Architectures with Sparse Point-Voxel Convolution. In *ECCV*, 2020.

- [34] Maxim Tatarchenko, Jaesik Park, Vladlen Koltun, and Qian-Yi Zhou. Tangent Convolutions for Dense Prediction in 3D. In *CVPR*, 2018.
- [35] Hugues Thomas, Charles R. Qi, Jean-Emmanuel Deschaud, Beatriz MarcoteGui, Francois Goulette, and Leonidas J. Guibas. KPConv: Flexible and Deformable Convolution for Point Clouds. In *ICCV*, 2019.
- [36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All You Need. In *NIPS*, 2017.
- [37] Guangming Wang, Hanwen Liu, Muyao Chen, Yehui Yang, Zhe Liu, and Hesheng Wang. Anchor-Based Spatio-Temporal Attention 3D Convolutional Networks for Dynamic 3D Point Cloud Sequences. *IEEE Transactions on Instrumentation and Measurement*, 2021.
- [38] Song Wang, Jianke Zhu, and Ruixiang Zhang. Meta-RangeSeg: LiDAR Sequence Semantic Segmentation Using Multiple Feature Aggregation. *IEEE RA-L*, 2022.
- [39] Pengxiang Wu, Siheng Chen, and Dimitris N Metaxas. MotionNet: Joint Perception and Motion Prediction for Autonomous Driving Based on Bird’s Eye View Maps. In *CVPR*, 2020.
- [40] Chenfeng Xu, Bichen Wu, Zining Wang, Wei Zhan, Peter Vajda, Kurt Keutzer, and Masayoshi Tomizuka. SqueezeSegV3: Spatially-Adaptive Convolution for Efficient Point-Cloud Segmentation. In *ECCV*, 2020.
- [41] Xu Yan, Jiantao Gao, Jie Li, Ruimao Zhang, Zhen Li, Rui Huang, and Shuguang Cui. Sparse Single Sweep LiDAR Point Cloud Segmentation via Learning Contextual Shape Priors from Scene Completion. In *AAAI*, 2021.
- [42] Maosheng Ye, Tongyi Cao, and Qifeng Chen. TPCN: Temporal Point Cloud Networks for Motion Forecasting. In *CVPR*, 2021.
- [43] Maosheng Ye, Shuangjie Xu, Tongyi Cao, and Qifeng Chen. DRINet: A Dual-Representation Iterative Learning Network for Point Cloud Segmentation. In *ICCV*, 2021.
- [44] Zhenxun Yuan, Xiao Song, Lei Bai, Zhe Wang, and Wanli Ouyang. Temporal-Channel Transformer for 3D Lidar-Based Video Object Detection for Autonomous Driving. *IEEE TCSVT*, 2021.
- [45] Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaoang Wang, Amrbrish Tyagi, and Amit Agrawal. Context Encoding for Semantic Segmentation. In *CVPR*, 2018.
- [46] Yang Zhang, Zixiang Zhou, Philip David, Xiangyu Yue, Zerong Xi, Boqing Gong, and Hassan Foroosh. PolarNet: An Improved Grid Representation for Online LiDAR Point Clouds Semantic Segmentation. In *CVPR*, 2020.
- [47] Yin Zhou, Pei Sun, Yu Zhang, Dragomir Anguelov, Jiyang Gao, Tom Ouyang, James Guo, Jiquan Ngiam, and Vijay Vasudevan. End-to-End Multi-View Fusion for 3D Object Detection in LiDAR Point Clouds. In *CoRL*, 2020.
- [48] Xinge Zhu, Hui Zhou, Tai Wang, Fangzhou Hong, Wei Li, Yuexin Ma, Hongsheng Li, Ruigang Yang, and Dahua Lin. Cylindrical and Asymmetrical 3D Convolution Networks for LiDAR-based Perception. *IEEE TPAMI*, 2021.
- [49] Xinge Zhu, Hui Zhou, Tai Wang, Fangzhou Hong, Yuexin Ma, Wei Li, Hongsheng Li, and Dahua Lin. Cylindrical and

Asymmetrical 3D Convolution Networks for LiDAR Segmentation. In *CVPR*, 2021.