

# 4D Panoptic Segmentation as Invariant and Equivariant Field Prediction

Minghan Zhu<sup>1,\*</sup>Shizhong Han<sup>2</sup>  
Maani Ghaffari<sup>1,\*</sup>Hong Cai<sup>2</sup>  
Fatih Porikli<sup>2</sup>Shubhankar Borse<sup>2</sup><sup>1</sup>University of Michigan, Ann Arbor<sup>2</sup>Qualcomm AI Research<sup>†</sup>

## Abstract

In this paper, we develop rotation-equivariant neural networks for 4D panoptic segmentation. 4D panoptic segmentation is a benchmark task for autonomous driving that requires recognizing semantic classes and object instances on the road based on LiDAR scans, as well as assigning temporally consistent IDs to instances across time. We observe that the driving scenario is symmetric to rotations on the ground plane. Therefore, rotation-equivariance could provide better generalization and more robust feature learning. Specifically, we review the object instance clustering strategies and restate the centerness-based approach and the offset-based approach as the prediction of invariant scalar fields and equivariant vector fields. Other sub-tasks are also unified from this perspective, and different invariant and equivariant layers are designed to facilitate their predictions. Through evaluation on the standard 4D panoptic segmentation benchmark of SemanticKITTI, we show that our equivariant models achieve higher accuracy with lower computational costs compared to their non-equivariant counterparts. Moreover, our method sets the new state-of-the-art performance and achieves 1st place on the SemanticKITTI 4D Panoptic Segmentation leaderboard.

## 1. Introduction

Perception with LiDAR point clouds is an important part of building real-world autonomous systems, for example, self-driving cars [38, 32, 50, 36]. As the computer vision community gradually builds more capable neural networks, the tasks also become more complex. 4D panoptic segmentation [1] is an emerging task that combines several previously independent tasks: semantic segmentation, instance segmentation, and object tracking, in a unified framework, given sequential LiDAR scans. As the output provides abundant useful information for understanding the dynamic

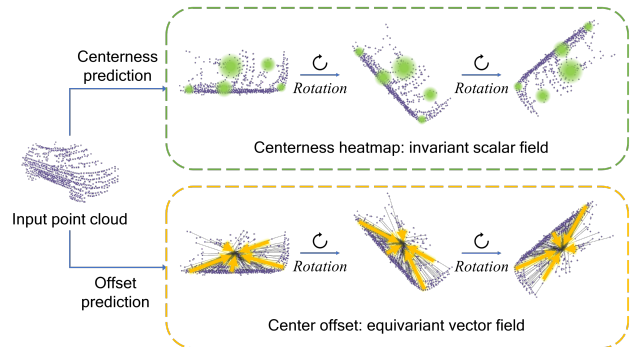


Figure 1. Predicting centerness or offset is a critical step in estimating object centers as part of the overall panoptic segmentation. A centerness heatmap can be viewed as an invariant field, while the offset can be viewed as an equivariant vector field. In our proposed approach, we devise different, corresponding methods to predict invariant and equivariant fields, respectively.

driving environment, solving this task has significant practical value.

While there are a few existing methods to solve 4D panoptic segmentation [1, 24], they ignore the rich, inherent symmetries present in this task. For example, for the point cloud of an object instance, its center is invariant to rotations, and the offset vector from any point on the object to the center is attached to and thus rotates along with the body frame. See Fig. 1 for a visual illustration.

As such, in this paper, we propose to develop equivariant neural networks to solve 4D panoptic segmentation. Equivariant networks [11] are deep learning models that are guaranteed to render outputs that respect the symmetries in data. For the 4D panoptic segmentation task, rotational equivariance can help the model perform consistently and generalize over rotations in the input data.

While equivariance is a nice property, equivariant models can be complex and incur high computational costs [52, 49, 45]. As a result, most existing equivariant models are only applied to small-scale problems, such as molecular analysis and single-object perception [17]. Recent works have looked into more efficient equivariant networks [52]

\*Work done at Qualcomm AI Research during an internship.

<sup>†</sup>Qualcomm AI Research, an initiative of Qualcomm Technologies, Inc.

and applications in larger problems [49], but significant performance improvement without extra computational cost has not been achieved in large-scale equivariant perception solutions.

In this work,  $SO(2)$ -equivariance is incorporated in the 4D segmentation model. We find that the equivariance brings consistent improvements in several performance metrics and that formulating the output as equivariant vector fields helps maximize the benefits of equivariant models, as compared to restricting to invariant scalar fields (see Fig. 1). Furthermore, our equivariant networks can improve the segmentation performance while reducing computational costs at the same time. With our proposed design, we outperform the non-equivariant models and, notably, achieve the **top-1** ranking position in the SemanticKITTI benchmark.

Our main contributions in this paper are as follows:

- We develop the first rotation-equivariant model for 4D panoptic segmentation, bringing improvements in both performance and efficiency.
- We investigate different strategies and designs to construct the equivariant architecture. Specifically, we discover the advantage of formulating prediction targets as equivariant vector fields, as compared to only invariant fields.
- Evaluated on the SemanticKITTI benchmark, our equivariant models significantly outperform existing methods, validating the value of equivariance in this large-scale perception task.
- Our code is open-sourced at <https://eq-4d-panoptic.github.io/>.

## 2. Related work

### 2.1. LiDAR 3D and 4D Panoptic segmentation

**3D Panoptic segmentation** The task of panoptic segmentation is first proposed in the image domain [23] and later extended to LiDAR point clouds with the release of a large-scale outdoor LiDAR point cloud dataset with panoptic labels, SemanticKITTI [3]. Similar to the semantic segmentation [8, 4, 21, 5, 6] and panoptic segmentation techniques in the image domain [31, 48, 7], their 3D counterparts can be classified into proposal-based and proposal-free methods. Proposal-based methods [3, 37] require a detection module to locate the objects first and then predict the instance mask for each bounding box and conduct semantic segmentation on the background pixels. This strategy needs to deal with potential conflicts among the segmentations. On the other hand, proposal-free methods conduct semantic segmentation first and then cluster the points belonging to different instances. The clustering strategy impacts overall efficiency

and performance. Offset prediction and centerness prediction are two major approaches for this. *Offset prediction* [19, 47] means that each point predicts the offset vector to the instance center, and the clustering is conducted on the predicted centers. *Centerness prediction* [1] is to regress a heatmap of the closeness to the instance center at any spatial location. Then the local maximums on the heatmap are used to cluster the points nearby. These two strategies can also be combined [51, 26]. Other clustering methods also exist. For example, [18] proposes an end-to-end clustering layer. [33] uses a graph network to cluster over-segmented points into instances.

**4D Panoptic segmentation** The 4D task is to provide temporally consistent object IDs on top of 3D panoptic segmentation. MOPT [22] is an early attempt to provide tracking ID for panoptic outputs. 4D-PLS [1] proposes evaluation metrics and a strong baseline method for this task. It accumulates point clouds in sequential timestamps to a common frame, applies segmentation on the aggregated point clouds, and clusters the instances using the centerness prediction. 4D-DS-Net [20] and 4D-StOP [24] follow this pipeline but use offset prediction to cluster the instances. CA-Net [28] uses an off-the-shelf 3D panoptic segmentation network and learns the temporal instances association through contrastive learning.

### 2.2. Equivariant Learning

**Equivariant neural networks** The equivariance to translations enables CNNs to generalize over translations of image content with much fewer parameters compared with fully connected networks. Equivariant networks extend the symmetries to rotations, reflections, permutations, etc. G-CNN [11] enables equivariance to 90-degree rotations of images. Steerable CNNs extend the symmetry to continuous rotations [42]. The input is also extended from 2D images to spherical [12, 15] and 3D data [41]. To deal with infinite groups (e.g., continuous rotations), generalized Fourier transforms and irreducible representations are adopted to formulate convolutions in the frequency domain [44, 40]. Equivariant graph networks [35] and transformers [17] are also proposed as non-convolutional equivariant layers. Equivariant models have applications in various areas such as physics, chemistry, and bioimaging [40, 17], where symmetries play an important role. They also attract research interest in robotic applications. For example, equivariant networks with  $SO(3)$ - and  $SE(3)$ -equivariance are developed to process 3D data such as point clouds [14, 10], meshes [13], and voxels [41]. However, due to the added complexity, most equivariant networks for 3D perception are restricted to relatively simple tasks with small-scale inputs, such as object-wise classification, registration, part segmentation, and reconstruction [34, 53, 9]. In the fol-

lowing, we will review recent progress in extending equivariance to large-scale outdoor 3D perception tasks.

**Equivariant networks in LiDAR perception** LiDAR perception demands two main considerations in equivariant models. Firstly, outdoor scene perception typically needs a sophisticated network design, and incorporating equivariant layers that are compatible with conventional network layers can build on existing successful designs. Secondly, due to the large sizes of LiDAR point clouds, it’s essential to create expressive equivariant models that fit within the memory constraints of standard GPUs.

Existing work mainly follows two strategies. With the first strategy, inputs are projected to rotation invariant features using local reference frames [27, 46]. In this way, the changes are mainly at the first layer of the network, causing limited memory overhead. The main drawback is that the invariant feature could cause information loss and limit performance. The second strategy adopts group convolution, by augmenting the domain of feature maps to include rotations [49, 45]. While achieving improved performance with the help of equivariance, they consume twice [45] or four times [49] of memory as their non-equivariant counterparts due to the augmented dimension of the feature maps and convolutions. With the two strategies, equivariant networks have been applied in 3D object detection [49, 45, 46] and semantic segmentation [27].

### 3. Equivariance for 4D Panoptic Segmentation

In this section, we provide preliminaries on equivariance and formulate the dense prediction task from the perspective of equivariant learning. We discuss how to learn equivariant features and fit equivariant prediction target fields, which leads to the design choices in our proposed architecture.

#### 3.1. Preliminaries on Equivariance

**Feature maps as fields** A feature map  $f_0$  is a *field*, assigning a value to each point in some space. In the context of point cloud perception, we have  $f_0 : \mathbb{R}^3 \rightarrow V$ , where  $V$  is some vector space. This map can represent the geometry of the point cloud, i.e.,  $f_0(x) = 1$  for every point  $x$  in the point cloud and  $f_0(x) = 0$  otherwise. It can also represent point properties or arbitrary learned point features. For example, the feature map  $f_0 : x \mapsto SC(x)$ , where  $SC(x)$  is the semantic class label of the point  $x$ , represents the semantic segmentation of a point cloud. We denote the space of all such feature maps as  $\mathcal{F}_0$ .

**Invariant and equivariant fields** For the *group* of transformations  $G$ , we use the rotation group  $SO(3)$  for the formulation, for which  $SO(2)$  is also valid.

A rotation  $R \in SO(3)$  can be applied to a point and to a feature map. A *rotated feature map*  $[Rf_0]$  is simply the feature map of the rotated point cloud. A point at  $x$  goes to  $Rx$  after the rotation, thus  $f_0(x)$  and  $[Rf_0](Rx)$  are the features of the same point before and after rotating the point cloud. The relation between them depends on the property of  $V$ . For instance, if  $V = SC$ , then we know

$$[Rf_0](Rx) = f_0(x), \quad \forall R \in SO(3), \quad (1)$$

i.e., the rotation does not change the semantic class.

As another example, if  $V$  represents the surface normal vector, then we have

$$[Rf_0](Rx) = Rf_0(x), \quad \forall R \in SO(3), \quad (2)$$

where on the right-hand side,  $R$  is applied to  $f_0(x) \in \mathbb{R}^3$ , meaning that the normal vector of a given point rotates along with the point cloud.

In both cases, the feature map of a rotated point cloud,  $[Rf_0]$ , can be generated from  $f_0$ , the feature map of the original point cloud. This property is called *equivariance*. In the case of Eq. (1), we call  $f_0$  an *invariant scalar field*. In the case of Eq. (2), we call  $f_0$  an *equivariant vector field*.

**Learning equivariant features** A dense prediction task on a point cloud is to reproduce a target field (e.g.,  $x \mapsto SC(x)$ ) using a feature map realized by a neural network. Naturally, it would be helpful to equip the network with the same invariant and/or equivariant properties as the target field. However, a general feature map learned by a network is typically neither invariant nor equivariant to rotations. To fix this, we can augment the space of feature maps to  $\mathcal{F} = \{f : \mathbb{R}^3 \times SO(3) \rightarrow V\}$ , defined by

$$f(x, R) := [R^{-1}f_0](R^{-1}x) \quad (3)$$

for some  $f_0 \in \mathcal{F}_0$ , i.e., the augmented feature map at rotation  $R$  equals the original feature map rotated by  $R^{-1}$ . In this way, we have,  $\forall R, R' \in SO(3)$ ,

$$[Rf](Rx, R') = [R'^{-1}Rf_0](R'^{-1}Rx) = f(x, R^{-1}R'), \quad (4)$$

which means that the augmented feature map of a rotated point cloud,  $[Rf]$ , can be generated by the augmented feature map of the original point cloud  $f$ , indicating that  $f$  is equivariant. Equivariant feature maps satisfying Eq. (3) can be constructed using group convolutions [11, 10, 52].

#### 3.2. Fitting Equivariant Targets

Now we can use the learned equivariant feature map  $f$  to fit the target field  $f_{gt} \in \mathcal{F}_0$ , which is invariant or equivariant. Suppose that we can fit  $f_{gt}$  using  $f$ , then  $[Rf]$ , the feature map of a rotated point cloud, should automatically fit  $[Rf_{gt}]$  which is the target field of the rotated point cloud. In other words, equivariance enables the *generalization* over rotations. Next, we introduce two strategies to achieve this.

**Rotational coordinate selection** Assuming that we can fit the target field of a point cloud without rotation:  $f(x, I) = f_{gt}(x)$ , where  $I$  is identity rotation. We want to show that the fitting generalizes to the rotated point clouds.

If  $f_{gt}$  is an invariant scalar field, from Eq. (4), we have

$$[Rf](Rx, R) = f(x, I) = f_{gt}(x) = [Rf_{gt}](Rx), \quad (5)$$

meaning that the feature map  $[Rf]$  at the rotational coordinate  $R$ ,  $[Rf](\cdot, R)$ , fits the target  $[Rf_{gt}]$  for the rotated point cloud.

If  $f_{gt}$  is an equivariant vector field, we have

$$[Rf](Rx, R) = f(x, I) = f_{gt}(x) = R^{-1}[Rf_{gt}](Rx), \quad (6)$$

which means that we need to apply a rotational matrix multiplication on features in  $[Rf](\cdot, R)$  to fit  $[Rf_{gt}]$ , i.e.,  $[Rf_{gt}](Rx) = R[Rf](Rx, R), \forall x \in \mathbb{R}^3$ .

The analysis above shows that the fitting of  $f \in \mathcal{F}$  to  $f_{gt} \in \mathcal{F}_0$  generalizes over rotations, if the rotational coordinate (i.e., the second argument in  $f$ ) is the same as the actual rotation of the point cloud. However, the actual rotation  $R$  is usually *unknown* during inference; thus needs to be learned. In practice, this is formulated as a *rotation classification* task to select the best rotational channel in a feature map  $f$ . We will discuss how we instantiate this in the architecture in Sec. 4.3.

**Invariant pooling layer** When the target field  $f_{gt}$  is invariant, there is another fitting strategy. For equivariant feature map  $f$  that satisfies Eq. (3), we can build a rotation-invariant layer by marginalizing over the second argument of  $f$ , i.e.,

$$f_{inv}(x) = \prod_R f(x, R), \quad (7)$$

where  $\prod$  denotes any summarizing operator symmetric to all its arguments, e.g., sum, average, and max.

If we can fit  $f_{inv}(x) = f_{gt}(x)$ , then based on Eq. (4), we have

$$\begin{aligned} [Rf_{inv}](Rx) &= \prod_{R'} [Rf](Rx, R') = \prod_{R'} f(x, R^{-1}R') \\ &= f_{inv}(x) = f_{gt}(x) = [Rf_{gt}](Rx), \end{aligned} \quad (8)$$

which indicates that generalization over rotations holds. In Sec. 4.3, we discuss different options of implementing the invariant layer in the network.

### 3.3. Equivariant Instance Segmentation

For instance segmentation, the prediction targets can be modeled as invariant or equivariant fields. As discussed in Sec. 2.1, centerness regression and offset regression are two predominately used approaches to estimate the object centers. Specifically, we can see that *centerness is an invariant*

*scalar field* and *offset is an equivariant vector field*. Accordingly, we propose different prediction layers in Sec. 4.3, resulting in different performances shown in Sec. 5.2.

## 4. Proposed Network Architecture

### 4.1. Network Overview

**Discretized SO(2)-equivariance** The model’s group equivariance should match the data’s actual transformations. An overly large group may lead to high computational costs with minimal performance improvement. For the outdoor driving scenario, the SO(2) group is chosen to represent planar rotations around the gravity axis.

We discretize SO(2) into a finite group and create a group CNN that is equivariant to the discretized rotations rather than the continuous ones. It allows us to leverage simpler structures akin to conventional deep learning models, allowing integration with existing state-of-the-art networks. Despite discretization, the network is expected to interpolate the equivariance gaps through training with rotational data augmentation.

Specifically, SO(2) is discretized into cyclic groups  $C_n$ , where  $n$  indicates the number of discretized rotations, such as  $C_3$  for 120-degree rotations. These are referred to as the *rotation anchors*.

**Network structure** We utilize the point-convolution style equivariant network E2PN [52] as our backbone, which is an equivariant version of KPConv [39]. We describe necessary adaptations to E2PN in Sec. 4.2, which allow us to build equivariant models on top of SOTA 4D panoptic segmentation network 4D-PLS [1] and 4D-StOP [24], both of which are based on KPConv. We refer to our equivariant models as *Eq-4D-PLS* and *Eq-4D-StOP*, respectively.

On a high level, both models first stack the point clouds from several sequential time instances within a common reference frame so that the temporal association becomes part of the instance segmentation. Each network consists of an encoder, a decoder, and prediction heads. The encoders and decoders are very similar for the two models, while the main differences lie in their prediction heads, which formulate the targets as invariant scalar fields and equivariant vector fields, respectively, as discussed in Sec. 4.3. An overview of our network structure is shown in Fig. 2.

### 4.2. Equivariant Encoder and Decoder

**Equivariant encoder** The encoder of the 4D panoptic segmentation networks can be made equivariant by simply swapping the KPConv [39] layers with E2PN [52] convolution layers. However, E2PN is originally designed for SO(3) equivariance, and uses quotient representations and efficient feature gathering to improve the efficiency. Using it for SO(2) equivariance requires two adaptations.

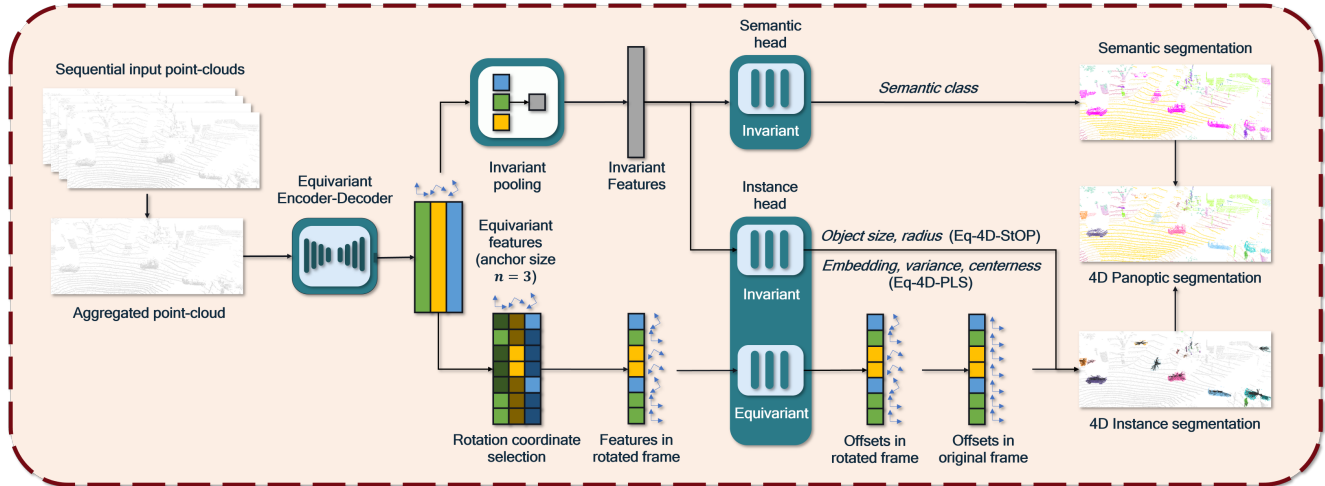


Figure 2. Overview of the network structure. All prediction targets can be classified as equivariant (e.g., offsets to object centers) and invariant (e.g., semantic classes). The gray color represents rotation-invariant features, which is used for invariant predictions. The green, yellow, and blue colors represent features at different rotational coordinates depicted by the small frames. Equivariant predictions use features from the rotational coordinate selection. Whether a rotation is selected is illustrated using the light and dark shades of colors.

First, we use the regular representation instead of the quotient representation in [52]. This is because  $SO(2)$  is abelian, in which case quotient representations cause loss of information (see appendix for details).

Second, to apply the efficient feature gathering in E2PN, the spatial position of the convolution kernel needs to be symmetric to the rotation anchors. It allows the rotation of kernel points to be implemented as a permutation of their indices. This implies that different  $n$ 's in  $C_n$  impose different constraints on the number and distribution of the kernel points. For example, the default KPConv kernel with 15 points is symmetric to 60-degree rotations and thus can be used to realize  $C_2$ ,  $C_3$ , and  $C_6$  equivariance. However,  $C_4$  requires a different kernel (for symmetry to 90-degree rotations), for which the 19-point KPConv kernel works.

**Equivariant decoder** The original E2PN only provides an encoder to predict a single output for an input point cloud. We need to devise an equivariant decoder for the dense 4D panoptic segmentation. Similar to conventional non-equivariant networks, our decoder adopts upsampling layers and 1-by-1 convolution layers. The upsampling layers simply assign the features of the coarser point cloud to the finer point cloud via nearest neighbor interpolation. The 1-by-1 convolution processes the feature at each point and each rotation independently. It is straightforward to prove the equivariance of such a decoder (see the appendix).

### 4.3. Equivariant Prediction Heads

Our baseline models 4D-PLS [1] and 4D-StOP [24] have different prediction heads. Their semantic segmentation heads are similar, but they employ different clustering ap-

proaches for instance segmentation. Correspondingly, we propose different equivariant prediction designs for them.

#### 4.3.1 Eq-4D-PLS: Segmentation as Invariant Scalar Field Prediction

In 4D-PLS [1], instance segmentation is done by clustering the point embeddings, assuming a Gaussian distribution for the embeddings of each instance. It also predicts a point-wise centerness score, measuring the closeness of a point to its instance center, which is used to initialize the cluster centers. Both the point embeddings and the centerness scores can be viewed as *invariant scalar fields*. Note that while these targets can appear like a vector, they are actually a stack of scalars invariant to rotations.

As discussed in Sec. 3.2, there are two strategies to fit invariant targets, i.e., rotation coordinate selection and invariant pooling. For invariant pooling layers, max pooling and average pooling over the rotational dimension are two obvious choices. For the rotational coordinate selection strategy, a key challenge is that the ground truth for the rotational coordinate  $R$  is unavailable or even undefined (as there is no canonical orientation for a LiDAR scan). In addition, existing dataset, e.g., SemanticKITTI, does not provide object bounding box annotations. Hence the object orientations are also unknown. We use an unsupervised strategy to address this issue. Instead of picking the best rotational dimension, we perform a weighted sum of all rotational dimensions, which is differentiable and allows the model to learn the weight for different rotations. This is equivalent to the group attentive pooling layer in EPN [10].

In summary, we study three designs, i.e., max pooling, average pooling, and group attentive pooling, for the pre-

diction of the invariant targets, including semantic classes, point embeddings, embedding variances, and centerness scores.

### 4.3.2 Eq-4D-StOP: Segmentation as Equivariant and Invariant Field Prediction

4D-StOP [24] uses an offset-based clustering method for instance segmentation. The network predicts an offset vector to the instance center for each input point, creating an *equivariant vector field*. The predicted center locations are clustered into instances, and features from points in the same instance are combined to predict instance properties.

As discussed in Sec. 3.2, we fit equivariant vector fields through rotational coordinate selection. While we do not have ground-truth rotations, the vector field of offsets naturally defines orientations. Denote an offset vector at point  $x$  as  $v(x) = x_{ctr} - x \in V = \mathbb{R}^3$ , where  $x_{ctr}$  is the center of the instance that  $x$  belongs to. We can define its rotation in  $SO(2)$  as  $\theta(v) = \text{atan2}(v_Y, v_X)$ , where  $X, Y$  are axes in the horizontal plane and  $Z$  is the vertical axis.  $\theta$  can be assigned to a discretized rotation coordinate (anchor)  $R_i^{gt}$  by nearest neighbor. In this way, each point has a rotation label based on its relative position to the instance center, which is well-defined and equivariant to the point cloud rotations.

Given the rotation label, we train the network to predict it as a classification task. Given the feature map at each rotational coordinate, a rotation score is predicted:

$$s_x(R_i) = \phi(f(x, R_i)), \quad \forall R_i \in SO(2)', \quad (9)$$

where  $x \in \mathbb{R}^3$ ,  $SO(2)'$  is the set of rotation anchors, i.e.,  $SO(2)' \cong C_n$ ,  $i = 1, \dots, n$ , and  $\phi$  is a scoring function. We concatenate  $s_x(R_i)$ 's as  $S_x = [s_x(R_1), \dots, s_x(R_n)]$  and apply a cross-entropy loss function on  $S_x$  with label  $R_i^{gt}$ .

The semantic segmentation, object size, and radius regression in 4D-StOP are invariant fields. However, note that a key difference from the prediction in Eq-4D-PLS (c.f. Sec. 4.3.1) is that we now have the rotation labels. As such, the rotational coordinate selection strategy can be applied to predicting the invariant targets. Therefore, for Eq-4D-StOP, we study four options for the invariant target prediction, including max pooling, average pooling, group attentive pooling, and rotational coordinate selection.

## 5. Experiments

### 5.1. Experimental setup

**Dataset:** We conduct our experiments primarily on the SemanticKITTI dataset [2]. The SemanticKITTI dataset establishes a benchmark for LiDAR panoptic segmentation [3]. It consists of 22 sequences from the KITTI dataset. 10 sequences are used for training, 1 for validation, and 11 for testing. In total, there are 43,552 frames. The dataset

Method	$LSTQ$	$S_{assoc}$	$S_{cls}$	$IOU^{St}$	$IOU^{Th}$
RangeNet++[29]+PP+MOT	43.8	36.3	52.8	60.5	42.2
KPConv[39]+PP+MOT	46.3	37.6	57.0	64.2	54.1
RangeNet++[29]+PP+SFP	43.4	35.7	52.8	60.5	42.2
KPConv[39]+PP+SFP	46.0	37.1	57.0	64.2	54.1
4D-DS-Net[20]	68.0	71.3	<b>64.8</b>	64.5	65.3
4D-PLS[1]	62.7	65.1	60.5	65.4	61.3
4D-StOP[24]	67.0	74.4	60.3	65.3	60.9
Eq-4D-PLS (ours)	65.0	67.7	62.3	66.4	64.6
Eq-4D-StOP (ours)	<b>70.1</b>	<b>77.6</b>	63.4	<b>66.4</b>	<b>67.1</b>

Table 1. SemanticKITTI validation set result. PP: PointPillars [25]. MOT: tracking-by-detection by [43]. SFP: tracking-by-detection with scene flow [30]. The best is highlighted in **bold**.

includes 28 annotated semantic classes, which are reorganized into 19 classes for the panoptic segmentation task. Among the 19 classes, 8 are classified as *things*, while the remaining 11 are categorized as *stuff*. Each point in the dataset is assigned a semantic label, and for points belonging to things, a temporally consistent instance ID.

**Metrics:** The core metric for 4D panoptic segmentation is  $LSTQ = \sqrt{S_{cls} \times S_{assoc}}$ , which is the geometric mean of the semantic segmentation metric  $S_{cls}$  and the instance segmentation and tracking metric  $S_{assoc}$ .  $S_{cls} = \frac{1}{C} \sum_{c=1}^C IoU(c)$  is the segmentation IoU averaged over all semantic classes. The average IoU for points belonging to things and stuff are denoted  $IoU^{Th}$  and  $IoU^{St}$ , respectively.  $S_{assoc}$  measures the spatial and temporal accuracy of segmenting object instances. See [1] for more details.

**Architecture:** For our Eq-4D-PLS and Eq-4D-StOP models, we keep the architectures unchanged from their 4D-PLS [1] and 4D-StOP [24] baselines except for the added rotational coordinate selection and invariant pooling layers necessary for equivariant and invariant field predictions. The input size, the batch size, and the learning rate also follow the baselines, respectively.

As the efficiency (especially the memory consumption) is a pain point for equivariant learning in large-scale LiDAR perception tasks, we specify the number of channels (network width)  $c$  and the rotation anchor size  $n$  in our analysis. The size of an equivariant feature map  $f$  is  $|f| = mcn$  for a point cloud with  $m$  points. Non-equivariant networks can be viewed as  $n = 1$ . We use the width of the first layer to denote the network width  $c$ , since the width of the following layers scales with the first layer proportionally. As feature maps play a major role in memory consumption,  $c \times n$  gives a rough idea of the memory cost of a model.

### 5.2. Quantitative Results

The evaluation results of our equivariant models on the SemanticKITTI validation set are shown in Tab. 1. Compared with 4D-PLS, our equivariant model improves by **2.3** points on  $LSTQ$ . Our Eq-4D-StOP model outperforms its non-equivariant baseline by **3.1**  $LSTQ$  points. The Eq-4D-StOP model achieves state-of-the-art performance among

Method	$LSTQ$	$S_{assoc}$	$S_{cls}$	$IoU^{St}$	$IoU^{Th}$
RangeNet++[29]+PP+MOT	35.5	24.1	52.4	64.5	35.8
KPConv[39]+PP+MOT	38.0	25.9	55.9	66.9	47.7
RangeNet++[29]+PP+SFP	34.9	23.3	52.4	64.5	35.8
KPConv[39]+PP+SFP	38.5	26.6	55.9	66.9	47.7
4D-PLS[1]	56.9	56.4	57.4	66.9	51.6
4D-DS-Net[20]	62.3	65.8	58.9	65.6	49.8
CA-Net[28]	63.1	65.7	60.6	66.9	52.0
4D-StOP[24]	63.9	69.5	58.8	67.7	53.8
Eq-4D-StOP ( <i>ours</i> )	<b>67.8</b>	<b>72.3</b>	<b>63.5</b>	<b>70.4</b>	<b>61.9</b>

Table 2. SemanticKITTI test set result.

published methods. From these experiments, we have the following observations:

- The improvements on  $IoU^{St}$  are similar for both Eq-4D-StOP and Eq-4D-PLS compared with their non-equivariant baselines.
- The improvements on  $IoU^{Th}$  and  $S_{assoc}$  are larger than on  $IoU^{St}$ .
- Eq-4D-StOP gains larger improvements on  $IoU^{Th}$  and  $S_{assoc}$  than Eq-4D-PLS does over their non-equivariant baselines.

The observations indicate the following. First, the introduction of equivariance brings improvements to both models across all metrics consistently. Second, objects (the *things* classes) enjoy more benefits from the equivariant models compared with the background (the *stuff* classes). We hypothesize that this is because objects present more rotational symmetry as compared to background classes. Third, the fact that more significant improvements are observed in Eq-4D-StOP shows the benefit of formulating the equivariant vector field regression problem induced from the offset-based clustering strategy. Improved clustering directly benefits the instance segmentation, i.e.,  $S_{assoc}$ , and it also improves the semantic segmentation of objects ( $IoU^{Th}$ ), since 4D-StOP unifies the semantic class prediction of all points belonging to a single instance by majority voting.

The evaluation results on the SemanticKITTI test set are shown in Tab. 2. Eq-4D-StOP achieves **3.9** points improvement in  $LSTQ$  over the non-equivariant model and ranks 1st in the leaderboard at the time of submission. We only test our best model on the test set, thus the result of Eq-4D-PLS is not available.

We use max-pooling in Eq-4D-PLS and average pooling in Eq-4D-StOP as the invariant pooling layer. The ablation study is in Sec. 5.3. The Eq-4D-PLS model is with  $c = 128, n = 6$ , and the Eq-4D-StOP model is with  $c = 128, n = 4$ . These parameters are selected based on the network scaling analysis in Sec. 5.5.

### 5.3. Ablation Study

**Invariant pooling and RCS** In Tab. 3, we show the segmentation performance with different invariant pooling lay-

Target type	Layer type	$LSTQ$	
		Eq-4D-StOP	Eq-4D-PLS
Invariant	Max	68.2	<b>63.7</b>
	Average	<b>69.2</b>	61.4
	Attentive	68.5	63.1
	RCS	68.4	n/a
Equivariant	RCS	<b>69.2</b>	n/a
	Average	61.6	n/a

Table 3. Ablation study on the invariant pooling layers and rotational coordinate selection (RCS). All comparisons use  $c = 128, n = 3$ . Some options are n/a for Eq-4D-PLS (see Sec. 4.3.1).

Method	$LSTQ$	$S_{assoc}$	$S_{cls}$	$IOU^{St}$	$IOU^{Th}$
4D-StOP	60.5	62.5	58.6	74.4	54.9
Eq-4D-StOP ( <i>ours</i> )	<b>67.3</b>	<b>73.7</b>	<b>61.5</b>	<b>76.4</b>	<b>58.7</b>

Table 4. 4D panoptic segmentation on nuScenes. The models are trained on the training split and evaluated on the validation split.

Anchor size $n$	1	2	3	4	6
Performance ( $LSTQ$ )	67.1	67.3	69.3	69.8	69.0
Inference speed (fps)	0.73	1.06	1.14	1.27	1.39

Table 5. Running speed analysis of Eq-4D-StOP, given constant feature map size  $c \times n = 256$ . Note that  $n = 1$  corresponds to the non-equivariant baseline.

ers and the effect of rotational coordinate selection (RCS). In terms of the invariant pooling layers, average pooling is optimal in Eq-4D-StOP, while max pooling is best in Eq-4D-PLS. Eq-4D-PLS favors max pooling, possibly because the point embeddings averaged over all rotational directions may be less discriminative, hindering instance clustering. In contrast, Eq-4D-StOP benefits from average pooling gathering information from all orientations, as the pooled features are used for object property prediction instead of clustering. For the equivariant field (offset) prediction, we compare the rotational coordinate selection with average pooling that wrongly treats the offsets as rotation-invariant targets. The performance drastically decreases when RCS is replaced with average pooling, showing that it is of vital importance to respect the equivariant nature of the targets.

### 5.4. Generalization on the nuScenes dataset

We extend experiments to the nuScenes [16] dataset. The hyperparameters follow the experiments on SemanticKITTI to show the generalizability of our proposed method. As shown in Tab. 4, our model (with  $c = 128, n = 4$ ) largely outperforms the baseline (with  $c = 512$ ), further validating the value of our equivariant approach.

### 5.5. Network Scaling and Computational Cost

We experiment on different network widths  $c$  and anchor sizes  $n$  to investigate the effect of network sizes on the model performance and efficiency. The  $LSTQ$  is evaluated

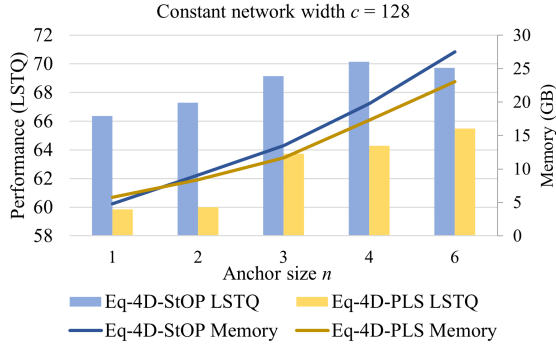


Figure 3. Effect of anchor size  $n$  on performance and memory usage in training, given constant network width  $c$ . Note that  $n = 1$  corresponds to the non-equivariant baselines.

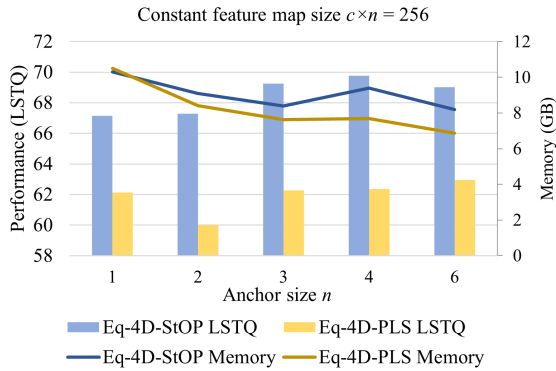


Figure 4. Effect of anchor size  $n$  on performance and memory usage in training, given a constant size of feature maps  $c \times n$ .

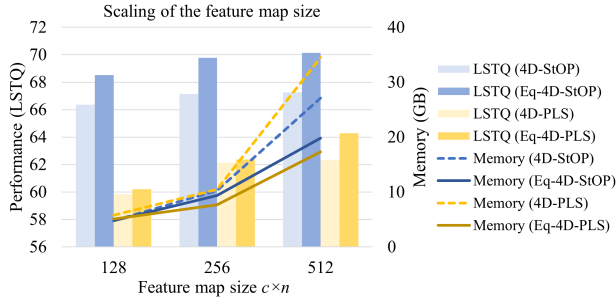


Figure 5. Comparison between equivariant and non-equivariant models at different network sizes.

on the SemanticKITTI validation set.

In Fig. 3, with  $c = 128$ , we test various anchor sizes  $n$ . A larger  $n$  more closely approximates continuous  $SO(2)$  equivariance but also proportionally expands feature maps and memory usage. The performance of both Eq-4D-PLS and Eq-4D-StOP enhances with increasing  $n$ , confirming the efficacy of equivariance in this task.

To rule out the factor of varying sizes of feature maps, we keep  $c \times n = 256$  with different combinations of  $c$  and  $n$  (for  $n = 3$ , we use  $c = 85$  as an approximation). As

shown in Fig. 4, Eq-4D-StOP significantly outperforms the non-equivariant version ( $n = 1$ ) when  $n \geq 3$ . The memory consumption even decreases when we increase  $n$ . There are two reasons. First, the size of the weight matrix in convolution layers gets smaller. Denote the number of kernel points as  $k$ , for a convolution layer with input and output width  $c$ , the size of the weight matrix is  $knc^2$ . Given constant  $c \times n$ , this number decreases with larger  $n$ . Second, the feature maps after the invariant pooling layer are smaller (of size  $m \times c$ ). The small bump-up of memory usage at  $n = 4$  is due to the larger convolution kernel ( $k = 19$  v.s.  $k = 15$  as discussed in Sec. 4.2), but it is still lower than the memory usage of the non-equivariant baseline. With constant feature map size  $c \times n$ , there is a trade-off between the number of rotation anchors  $n$  and the feature channels per rotation anchor, which explains the slight performance decline at  $n = 6$ . A similar trend can also be observed in Eq-4D-PLS, but with a smaller performance margin.

We also report the running time of our models in Tab. 5. The equivariant models run faster with higher accuracy.

We further investigate whether the advantage of equivariant models only occurs at a specific network size. In Fig. 5, we scale the feature map size of the networks. We use  $n = 4$  for equivariant models in this comparison. The equivariant models outperform the non-equivariant models at all network sizes. The memory consumption increases faster for the non-equivariant models, because the sizes of convolution kernel weight matrices grow quadratically with the network width  $c$ , while equivariant models have a smaller  $c$  given the same feature map size.

In summary, we found that the equivariant models have better performance than their non-equivariant counterparts with lower computational costs at different network sizes.

## 6. Conclusion

In this paper, we use equivariant learning to tackle a complicated large-scale perception problem, the 4D panoptic segmentation of sequential point clouds. While equivariant models were generally perceived as more expensive and complex than conventional non-equivariant models, we show that our method can bring performance improvements and lower computational costs at the same time. We also show that the advantage of equivariant models can be better leveraged if we formulate the learning targets as equivariant vector fields, compared with invariant scalar fields. A limitation of this work is that we did not propose drastically new designs on the overall structure of the 4D panoptic segmentation network under the equivariant setup, but it also allows us to conduct apple-to-apple comparisons regarding equivariance on this task so that our contribution is orthogonal to the improvements in the specific network design. We hope our work could inspire wider incorporation of equivariant networks in practical robotic perception problems.



## References

- [1] Mehmet Aygun, Aljosa Osep, Mark Weber, Maxim Maximov, Cyrill Stachniss, Jens Behley, and Laura Leal-Taixé. 4d panoptic lidar segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5527–5537, 2021. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#)
- [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 *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9297–9307, 2019. [6](#)
- [3] Jens Behley, Andres Milioto, and Cyrill Stachniss. A benchmark for lidar-based panoptic segmentation based on kitti. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13596–13603. IEEE, 2021. [2](#), [6](#)
- [4] Shubhankar Borse, Hong Cai, Yizhe Zhang, and Fatih Porikli. Hs3: Learning with proper task complexity in hierarchically supervised semantic segmentation. *arXiv preprint arXiv:2111.02333*, 2021. [2](#)
- [5] Shubhankar Borse, Debasmit Das, Hyojin Park, Hong Cai, Risheek Garrepalli, and Fatih Porikli. Dejavu: Conditional regenerative learning to enhance dense prediction. *arXiv preprint arXiv:2303.01573*, 2023. [2](#)
- [6] Shubhankar Borse, Marvin Klingner, Varun Ravi Kumar, Hong Cai, Abdulaziz Almuzairee, Senthil Yogamani, and Fatih Porikli. X-align: Cross-modal cross-view alignment for bird’s-eye-view segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3287–3297, 2023. [2](#)
- [7] Shubhankar Borse, Hyojin Park, Hong Cai, Debasmit Das, Risheek Garrepalli, and Fatih Porikli. Panoptic, instance and semantic relations: A relational context encoder to enhance panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1269–1279, 2022. [2](#)
- [8] Shubhankar Borse, Ying Wang, Yizhe Zhang, and Fatih Porikli. Inverseform: A loss function for structured boundary-aware segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5901–5911, 2021. [2](#)
- [9] Evangelos Chatzipantazis, Stefanos Pertigkiozoglou, Edgar Dobriban, and Kostas Daniilidis. Se (3)-equivariant attention networks for shape reconstruction in function space. *arXiv preprint arXiv:2204.02394*, 2022. [2](#)
- [10] Haiwei Chen, Shichen Liu, Weikai Chen, Hao Li, and Randall Hill. Equivariant point network for 3d point cloud analysis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14514–14523, 2021. [2](#), [3](#), [5](#)
- [11] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999. PMLR, 2016. [1](#), [2](#), [3](#)
- [12] Taco S Cohen, Mario Geiger, Jonas Köhler, and Max Welling. Spherical cnns. *arXiv preprint arXiv:1801.10130*, 2018. [2](#)
- [13] Pim De Haan, Maurice Weiler, Taco Cohen, and Max Welling. Gauge equivariant mesh cnns: Anisotropic convolutions on geometric graphs. *arXiv preprint arXiv:2003.05425*, 2020. [2](#)
- [14] Congyue Deng, Or Litany, Yueqi Duan, Adrien Poulenard, Andrea Tagliasacchi, and Leonidas J Guibas. Vector neurons: A general framework for so (3)-equivariant networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12200–12209, 2021. [2](#)
- [15] Carlos Esteves, Christine Allen-Blanchette, Ameesh Makadia, and Kostas Daniilidis. Learning so (3) equivariant representations with spherical cnns. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 52–68, 2018. [2](#)
- [16] Whye Kit Fong, Rohit Mohan, Juana Valeria Hurtado, Lubing Zhou, Holger Caesar, Oscar Beijbom, and Abhinav Valada. Panoptic nusenes: A large-scale benchmark for lidar panoptic segmentation and tracking. *IEEE Robotics and Automation Letters*, 7(2):3795–3802, 2022. [7](#)
- [17] Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. Se (3)-transformers: 3d roto-translation equivariant attention networks. *Advances in Neural Information Processing Systems*, 33:1970–1981, 2020. [1](#), [2](#)
- [18] Stefano Gasperini, Mohammad-Ali Nikouei Mahani, Alvaro Marcos-Ramiro, Nassir Navab, and Federico Tombari. Panoster: End-to-end panoptic segmentation of lidar point clouds. *IEEE Robotics and Automation Letters*, 6(2):3216–3223, 2021. [2](#)
- [19] Fangzhou Hong, Hui Zhou, Xinge Zhu, Hongsheng Li, and Ziwei Liu. Lidar-based panoptic segmentation via dynamic shifting network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13090–13099, 2021. [2](#)
- [20] Fangzhou Hong, Hui Zhou, Xinge Zhu, Hongsheng Li, and Ziwei Liu. Lidar-based 4d panoptic segmentation via dynamic shifting network. *arXiv preprint arXiv:2203.07186*, 2022. [2](#), [6](#), [7](#)
- [21] Hanzhe Hu, Yinbo Chen, Jiarui Xu, Shubhankar Borse, Hong Cai, Fatih Porikli, and Xiaolong Wang. Learning implicit feature alignment function for semantic segmentation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIX*, pages 487–505. Springer, 2022. [2](#)
- [22] Juana Valeria Hurtado, Rohit Mohan, Wolfram Burgard, and Abhinav Valada. Mopt: Multi-object panoptic tracking. *arXiv preprint arXiv:2004.08189*, 2020. [2](#)
- [23] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9404–9413, 2019. [2](#)
- [24] 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 *Computer Vision–ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part I*, pages 537–553. Springer, 2023. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#)

- [25] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12697–12705, 2019. 6
- [26] Jinke Li, Xiao He, Yang Wen, Yuan Gao, Xiaoqiang Cheng, and Dan Zhang. Panoptic-phnet: Towards real-time and high-precision lidar panoptic segmentation via clustering pseudo heatmap. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11809–11818, 2022. 2
- [27] Xingyi Li, Wenxuan Wu, Xiaoli Z Fern, and Li Fuxin. Improving the robustness of point convolution on k-nearest neighbor neighborhoods with a viewpoint-invariant coordinate transform. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1287–1297, 2023. 3
- [28] Rodrigo Marcuzzi, Lucas Nunes, Louis Wiesmann, Ignacio Vizzo, Jens Behley, and Cyrill Stachniss. Contrastive instance association for 4d panoptic segmentation using sequences of 3d lidar scans. *IEEE Robotics and Automation Letters*, 7(2):1550–1557, 2022. 2, 7
- [29] Andres Milioto, Ignacio Vizzo, Jens Behley, and Cyrill Stachniss. Rangenet++: Fast and accurate lidar semantic segmentation. In *2019 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 4213–4220. IEEE, 2019. 6, 7
- [30] Himangi Mittal, Brian Okorn, and David Held. Just go with the flow: Self-supervised scene flow estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11177–11185, 2020. 6
- [31] Rohit Mohan and Abhinav Valada. Efficienttps: Efficient panoptic segmentation. *International Journal of Computer Vision*, 129(5):1551–1579, 2021. 2
- [32] Rui Qian, Xin Lai, and Xirong Li. 3d object detection for autonomous driving: A survey. *Pattern Recognition*, 130:108796, 2022. 1
- [33] Ryan Razani, Ran Cheng, Enxu Li, Ehsan Taghavi, Yuan Ren, and Liu Bingbing. Gp-s3net: Graph-based panoptic sparse semantic segmentation network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16076–16085, 2021. 2
- [34] Rahul Sajani, Adrien Poulernard, Jivitesh Jain, Radhika Dua, Leonidas J Guibas, and Srinath Sridhar. Condor: Self-supervised canonicalization of 3d pose for partial shapes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16969–16979, 2022. 2
- [35] Victor Garcia Satorras, Emiel Hooeboom, and Max Welling. E (n) equivariant graph neural networks. In *International conference on machine learning*, pages 9323–9332. PMLR, 2021. 2
- [36] Yunxiao Shi, Haoyu Fang, Jing Zhu, and Yi Fang. Pairwise attention encoding for point cloud feature learning. In *2019 International Conference on 3D Vision (3DV)*, pages 135–144, 2019. 1
- [37] Kshitij Sirohi, Rohit Mohan, Daniel Büscher, Wolfram Burgard, and Abhinav Valada. Efficienttps: Efficient lidar panoptic segmentation. *IEEE Transactions on Robotics*, 38(3):1894–1914, 2021. 2
- [38] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2446–2454, 2020. 1
- [39] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6411–6420, 2019. 4, 6, 7
- [40] Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*, 2018. 2
- [41] Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco S Cohen. 3d steerable cnns: Learning rotationally equivariant features in volumetric data. *Advances in Neural Information Processing Systems*, 31, 2018. 2
- [42] Maurice Weiler, Fred A Hamprecht, and Martin Storath. Learning steerable filters for rotation equivariant cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 849–858, 2018. 2
- [43] Xinshuo Weng, Jianren Wang, David Held, and Kris Kitani. 3d multi-object tracking: A baseline and new evaluation metrics. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10359–10366. IEEE, 2020. 6
- [44] Daniel E Worrall, Stephan J Garbin, Daniyar Turmukhambetov, and Gabriel J Brostow. Harmonic networks: Deep translation and rotation equivariance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5028–5037, 2017. 2
- [45] Hai Wu, Chenglu Wen, Wei Li, Xin Li, Ruigang Yang, and Cheng Wang. Transformation-equivariant 3d object detection for autonomous driving. *arXiv preprint arXiv:2211.11962*, 2022. 1, 3
- [46] Liang Xie, Yibo Yang, Wenxiao Wang, Binbin Lin, Deng Cai, and Xiaofei He. General rotation invariance learning for point clouds via weight-feature alignment. *arXiv preprint arXiv:2302.09907*, 2023. 3
- [47] Shuangjie Xu, Rui Wan, Maosheng Ye, Xiaoyi Zou, and Tongyi Cao. Sparse cross-scale attention network for efficient lidar panoptic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2920–2928, 2022. 2
- [48] Tien-Ju Yang, Maxwell D Collins, Yukun Zhu, Jyh-Jing Hwang, Ting Liu, Xiao Zhang, Vivienne Sze, George Papandreou, and Liang-Chieh Chen. Deeperlab: Single-shot image parser. *arXiv preprint arXiv:1902.05093*, 2019. 2
- [49] Hong-Xing Yu, Jiajun Wu, and Li Yi. Rotationally equivariant 3d object detection. In *Proceedings of the IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition*, pages 1456–1464, 2022. [1](#), [2](#), [3](#)
- [50] Yuanxin Zhong, Minghan Zhu, and Huei Peng. Vin: Voxel-based implicit network for joint3d object detection and segmentation for lidars. In *32nd British Machine Vision Conference 2021, BMVC 2021, Online, November 22-25, 2021*, page 284. BMVA Press, 2021. [1](#)
- [51] Zixiang Zhou, Yang Zhang, and Hassan Foroosh. Panoptic-polarnet: Proposal-free lidar point cloud panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13194–13203, 2021. [2](#)
- [52] Minghan Zhu, Maani Ghaffari, William A Clark, and Huei Peng. E2pn: Efficient se (3)-equivariant point network. *arXiv preprint arXiv:2206.05398*, 2022. [1](#), [3](#), [4](#), [5](#)
- [53] Minghan Zhu, Maani Ghaffari, and Huei Peng. Correspondence-free point cloud registration with so (3)-equivariant implicit shape representations. In *Conference on Robot Learning*, pages 1412–1422. PMLR, 2022. [2](#)