ItTakesTwo: Leveraging Peer Representations for Semi-supervised LiDAR Semantic Segmentation

Yuyuan Liu^{*1(\square)}, Yuanhong Chen^{*1}, Hu Wang^{2,1}, Vasileios Belagiannis³, Ian Reid^{2,1}, and Gustavo Carneiro^{4,1}

¹ Australian Institute for Machine Learning (AIML), University of Adelaide, AU

² Mohamed bin Zayed University of Artificial Intelligence (MBZUAI), UAE

³ University of Erlangen–Nuremberg, Germany

⁴ Centre for Vision, Speech and Signal Processing (CVSSP), University of Surrey, UK

Abstract. The costly and time-consuming annotation process to produce large training sets for modelling semantic LiDAR segmentation methods has motivated the development of semi-supervised learning (SSL) methods. However, such SSL approaches often concentrate on employing consistency learning only for individual LiDAR representations. This narrow focus results in limited perturbations that generally fail to enable effective consistency learning. Additionally, these SSL approaches employ contrastive learning based on the sampling from a limited set of positive and negative embedding samples. This paper introduces a novel semisupervised LiDAR semantic segmentation framework called ItTakesTwo (IT2). IT2 is designed to ensure consistent predictions from peer Li-DAR representations, thereby improving the perturbation effectiveness in consistency learning. Furthermore, our contrastive learning employs informative samples drawn from a distribution of positive and negative embeddings learned from the entire training set. Results on public benchmarks show that our approach achieves remarkable improvements over the previous state-of-the-art (SOTA) methods in the field. The code is available at: https://github.com/yyliu01/IT2.

Keywords: Semi-supervised Learning · LiDAR Semantic Segmentation

1 Introduction

Outdoor LiDAR semantic segmentation is emerging as an important area of research given its application to self-driving vehicles [16, 17]. Although LIDAR data has advantages relative to photometric camera data, it also presents unique challenges, such as data sparsity [66] and varying point density [27, 45]. Most existing works alleviate these issues by transforming the LiDAR point-clouds to other representations, including range-based image [13, 26, 45, 55], voxel grids [15, 21, 66], or fusing them with other modalities [38, 41] and views [58]. Despite their success, those methods rely on supervised learning that requires large labelled training sets, which can be problematic given the costly and time-consuming labelling process [39, 47].

^{*} denotes equal contribution. (>) Email: yuyuan.liu@adelaide.edu.au.



Fig. 1: The **left** figure represents the 3D point cloud scan (ground truth of four classes – Car, Pole, Terrain, and Sidewalk) while the **right** figures denote the prediction results (into the same four classes), provided by our model, from **range** (top-right) and voxel (bottom-right) representations, where the incorrect predictions are marked in gray. Please notice that: 1) the mapping of the four classes from the 3D point cloud scan to each representation demonstrates the consistent distribution of semantic classes, and 2) the histograms of the predictions for each representation exhibit distinct differences.

These issues have motivated semi-supervised learning (SSL) innovations for LiDAR semantic segmentation [23,27,31], which require a small labelled dataset used in conjunction with a large unlabelled dataset. Consistency Learning, one of the main strategies explored in SSL for image semantic segmentation, penalises inconsistent predictions from perturbed inputs, where these perturbations take various forms, including input augmentation [62, 69], adversarial noise [22, 39], network perturbations [6] and dropout [61]. The incorporation of strong and varied perturbations is crucial for enhancing generalisation [14, 39].

Current research [27,31] applies such consistency learning in SSL for LiDAR points using many strategies, such as laser-beam augmentation [27] or Mean Teacher (MT) [31] framework. Despite their success, these approaches solely rely on a single LiDAR representation (e.g., voxel grid) for training, which could potentially limit the model's ability to generalize effectively due to the restricted perturbations. Such an issue can be mitigated with the use of multiple LiDAR representations for consistency learning that explores the fundamental Clustering assumption⁵, where the results from different representations naturally indicate the same underlying semantic information from the 3D LiDAR points. For example, as illustrated in Fig. 1, the class Terrain (\triangle) in the top-right pixels and bottom-right voxels should be "clustered" according to the (left) original 3D point data. Thus, the peer LiDAR representations can be regarded as a novel form similar to perturbation, but yet under-explored in the field. Other disadvantages of focusing on only one of the representations when using small labelled datasets are: the range view suffers from the information loss during the projection process, particularly for the tail categories, such as Pole (\square), which has

⁵ Cluster hypothesis (or assumption) [42]: the data of various forms that behave similarly with respect to information relevance should be clustered together.

a low mIoU of 52.02; and the voxel grid tends to show lower accuracy in dense areas that are far away from the LiDAR source, such as Sidewalk (\checkmark), which has a low mIoU of 69.50. Such limited pseudo label results will inevitably introduce confirmation bias (i.e., the problem of relying on and reinforcing wrong predictions during training), reducing even further the efficacy of consistency learning based on single representation [23, 27, 31].

Another effective technique for semi-supervised LiDAR segmentation is contrastive learning [7,20,23,31], which consists of attracting positive and repelling negative embedding pairs at instance [46] or category [23,31] levels, thereby enhancing the latent space point distribution. However, to generate these pairs, methods typically rely on a limited number of embeddings extracted from randomly selected training samples due to the high computational cost associated with dense contrastive learning tasks [23,31]. Unfortunately, this limited sample size hinders a comprehensive understanding of the distribution across the entire embedding space [5], resulting in suboptimal convergence of LiDAR segmentation in SSL. This limitation becomes particularly pronounced when dealing with embeddings from different models in multiple views, where potential noisy data [44, 51] or model bias [59] from one network can adversely affect others, thereby impeding the optimization process.

In this work, we propose a novel semi-supervised LiDAR semantic segmentation framework, named ItTakesTwo (IT2). IT2 employs consistency learning based on peer-representations perturbation, which enforces the consistent predictions across two different representations. Such perturbation contributes with a more effective SSL generalisation and alleviates issues related to pseudo label qualities from a single representation. Furthermore, we propose cross-distribution contrastive learning, where instead of randomly selecting training samples in the embedding space, we learn informative embedding distributions (also known as prototype) via Gaussian Mixture Models from both representations. Then the embedding samples are drawn from such prototypical distributions to produce representative features that enhance the understanding of the latent space in peer representation to both LiDAR representations [27], we empirically show that applying different augmentations to distinct representations provides better generalisation. To summarise, our work makes the following key contributions:

- 1. A novel SSL LiDAR segmentation framework (IT2) that leverages peerrepresentation perturbation in consistency learning, encouraging the same predictions for the points with the same semantic meaning in different Li-DAR representations;
- 2. A new contrastive learning where informative positive and negative pairs of embeddings are sampled from prototype distributions (modeled via Gaussian Mixture Models) of both representations; and
- 3. An innovative representation-specific data augmentation that enables IT2 achieve better generalisation.

Our IT2 has state-of-the-art (SOTA) performance on nuScenes [4], SemanticKitti [2] and ScribbleKitti [52], and we also achieve best performance with different

sampling strategies to form the labelled and unlabelled subsets. Notably, in nuScenes [4] dataset, we outperform the previous SOTA [27] by around 4.6% in the limited labelled data situations.

2 Related Work

LiDAR Semantic Segmentation is an essential task for autonomous driving [17]. To more effectively extract the feature information from the unordered LiDAR point data, current approaches [45, 66, 67] transform such point data into different representations. Range representation methods [26, 45, 65] transfer the LiDAR points to the spherical coordinates [45] and project them to the image space [26, 65]. Voxel representation methods [15, 28, 66] voxelise the LiDAR points and project them into different sub-types of geometry views to alleviate the sparse data challenges [10,15], encompassing cartesian [53], spherical [28] and cylindrical [66] coordinate systems. Furthermore, some methods [34, 38] take inputs from a combination of LiDAR representations [48.67] or incorporate inputs from camera modalities [34, 40], fusing them in the intermediate layer to enhance feature understanding. Although these supervised methods show promising results [58, 65, 66], there has been a longstanding interest in minimising the extensive labeling efforts, with approaches such as scribble annotation [52], box supervision [37], active learning [35,60], and SSL [23,27]. In this paper, we focus on SSL to alleviate such annotation costs.

Semi-supervised LiDAR Semantic Segmentation leverages unlabelled data to enhance the model's generalisation beyond the small labelled dataset. The current literature on LiDAR SSL primarily focuses on indoor [7, 12, 32, 60] or object [25,49] point cloud scenarios, which have uniformly distributed and high density points in the scan. Such scenario is markedly different from the representation of outdoor LiDAR points that pose unique challenges, such as complex structure [27] and low density points [66]. Previous semi-supervised LiDAR segmentation methods [23,31,46] mainly rely on contrastive learning to improve the latent space understanding. For instance, GPC [23] utilises confidence-threshold to constrain potentially noisy pseudo labels, and Lim3D [31] uses a memory bank to save negative samples. To increase the variation of the LiDAR data, Laser-Mix [27] proposes a novel augmentation that mixes laser beams in various spatial positions and holds the consistency learning under the Mean Teacher (MT) architecture [50]. However, these methods use only single LiDAR representations which, as explained in Sec. 1, imposes many limitations, such as limited perturbation effectiveness and poor performance for specific categories. Our framework introduces a novel peer-representation perturbation that employs various LiDAR representations during training, resulting in improved generalisation.

Contrastive Learning [5, 24, 54, 56] enforces latent embeddings to be closer to instances with the same semantic meaning (i.e. positive) and to be far away from instances with different semantic meanings (i.e. negative). This approach has been explored in LiDAR segmentation [7, 12, 23, 31, 46], but the reliance on actual training samples leads to sub-optimal training as only a subset of em-



Fig. 2: Illustration of our approach with unlabeled data. The point scans in the batch (\mathbf{x}_1 and \mathbf{x}_2 on the left) are transformed into range (and processed by $f_{\phi_{\mathbf{FCN}}}(.)$) and voxel (and processed by $f_{\theta_{\mathbf{EXR}}}(.)$) representations. Subsequently, the results from $f_{\phi_{\mathbf{SEG}}}(f_{\phi_{\mathbf{FCN}}}(.))$ and $f_{\theta_{\mathbf{SEG}}}(f_{\theta_{\mathbf{EXR}}}(.))$ are transformed to the other representation with Eq. (2) and cross-supervised with Eq. (4). Also, the embedding $\{ \underline{*}, \underline{\bullet} \}$ selected by the pseudo labels are utilised to model the GMMs with Eq. (6) and penalised by the informative samples (generated from GMMs) based on $\ell_{cross}(.)$ in Eq. (8). Please note the arrow \cdots indicates that the operation does not require gradient computation.

bedding pairs [23, 31] can be selected for computing the contrastive loss. This becomes more challenging when working with embeddings obtained from different representations [36, 51, 59], as potentially noisy predictions [44, 59] from each representation can result in confirmation bias during training [1]. Prototypical contrastive learning [5, 18, 29, 30, 33] explores the use of a set of representative embeddings for the contrastive learning via memory bank [18, 54], optimal transport [5], or deep clustering [29, 30]. Nevertheless, such prototypes represent the modes of a more complex underlying distribution of prototypes that has not been explored before for SSL LiDAR segmentation. In our framework, we employ Gaussian Mixture Models (GMMs) to represent class-specific embedding distributions from both LiDAR representations and to generate informative embeddings for all classes, enhancing the effectiveness of the contrastive learning.

3 Methodology

Our approach is based on a labelled point dataset $\mathcal{P}_L = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{|\mathcal{P}_L|}$ consisting of $|\mathcal{P}_L|$ scans and an unlabelled set $\mathcal{P}_U = \{(\mathbf{x}_i)\}_{i=1}^{|\mathcal{P}_U|}$ with $|\mathcal{P}_U|$ scans. Each scan $\mathbf{x}_i \in \mathcal{P} \subset \mathbb{R}^{3 \times C \times N}$ has **3** spatial coordinates, with *C* denoting the number of features (intensity, elongation, etc.), *N* representing the number of points, and $\mathbf{y}_i \in \mathcal{Y} \subset \{0, 1\}^{Y \times N}$ represents the *one-hot* label with *Y* being the number of classes in the labelled set. This point cloud dataset is transformed into the range and voxel representations to build the two peer representation datasets that is required by our method. These transformations are defined by [58] as:

- $\begin{array}{l} -\psi_{p\to v}: \mathcal{P} \to \mathcal{V} \text{ is a point to voxel projection function that forms } \mathcal{V}_L = \\ \{(\mathbf{x}_i^v, \mathbf{y}_i^v) | (\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{P}_L\}_{i=1}^{|\mathcal{P}_L|} \text{ and } \mathcal{V}_U = \{\mathbf{x}_i^v | \mathbf{x}_i \in \mathcal{P}_U\}_{i=1}^{|\mathcal{P}_U|}, \text{ with each volume} \\ \mathbf{x}_i^v \in \mathcal{V} \subset \mathbb{R}^{C_v \times H \times W \times L} \text{ having size } H \times W \times L \text{ and } C_v \text{ features, and the} \\ \text{projected label being } \mathbf{y}_i^v \in \{0, 1\}^{Y \times H \times W \times L} \text{ in one-hot format.} \end{array}$
- $\begin{array}{l} -\psi_{p\rightarrow r}:\mathcal{P}\rightarrow\mathcal{R} \text{ is a function to project point to the range image}^6. \text{ With this,} \\ \text{we build } \mathcal{R}_L = \{(\mathbf{x}_i^r,\mathbf{y}_i^r)|(\mathbf{x}_i,\mathbf{y}_i)\in\mathcal{P}_L\}_{i=1}^{|\mathcal{P}_L|} \text{ and } \mathcal{R}_U = \{\mathbf{x}_i^r|\mathbf{x}_i\in\mathcal{P}_U\}_{i=1}^{|\mathcal{P}_U|}, \\ \text{with each range image } \mathbf{x}_i^r\in\mathcal{R}\subset\mathbb{R}^{C_r\times U\times V} \text{ having size } U\times V \text{ with } C_r \\ \text{features, and the projected label being } \mathbf{y}_i^r\in\{0,1\}^{Y\times U\times V}. \end{array}$

The inverse of the transformations above are respectively denoted by $\psi_{v\to p}$ and $\psi_{r\to p}$. Based on the transformed range and voxel representations from \mathcal{P}_L and \mathcal{P}_U , we define the labelled set $\mathcal{D}_L = \{(\mathbf{x}_i^v, \mathbf{y}_i^v), (\mathbf{x}_i^r, \mathbf{y}_i^r) | (\mathbf{x}_i^v, \mathbf{y}_i^v) \in \mathcal{V}_L, (\mathbf{x}_i^r, \mathbf{y}_i^r) \in \mathcal{R}_L, \}_{i=1}^{|\mathcal{D}_L|}$ and the larger unlabelled set $\mathcal{D}_U = \{(\mathbf{x}_i^v, \mathbf{x}_i^r) | \mathbf{x}_i^v \in \mathcal{V}_U, \mathbf{x}_i^r \in \mathcal{R}_U\}_{i=1}^{|\mathcal{D}_U|}$ where $|\mathcal{D}_L| \leq |\mathcal{D}_U|$. These datasets are used to train our IT2 methodology in Sec. 3.1, while Sec. 3.2 describes our Gaussian Mixture Models (GMMs) based prototypical contrastive learning. Lastly, we introduce the details of our augmentation method in Sec. 3.3.

3.1 ItTakesTwo (IT2) Framework

As illustrated in Fig. 2, our approach enforces prediction consistency between two peer representations to explore a potentially more effective perturbation. Our method can use various LiDAR representations, but we follow [27] and employ range and voxel networks for a fair comparison in the benchmark. **The range image network** consists of an FCN Network (e.g., ResNet [19]) that maps the input to a latent dense feature space with $f_{\phi_{\mathbf{FCN}}} : \mathcal{R} \to \mathcal{I}$, where $\mathcal{I} \subset \mathbb{R}^{U \times V \times D_r}$ with D_r denoting feature dimensionality. Then a segmentation head will map \mathcal{I} to the categorical segmentation map with $f_{\phi_{\mathbf{FCN}}} : \mathcal{I} \to \mathbb{R}^{U \times V \times Y}$. Similarly, **the voxel network** extracts features with $f_{\theta_{\mathbf{EXR}}} : \mathcal{V} \to \mathcal{K}$, where $\mathcal{K} \in \mathbb{R}^{H \times W \times L \times D_v}$ with D_v denoting the feature dimensionality, followed by a segmentation head $f_{\theta_{\mathbf{SEG}}} : \mathcal{K} \to \mathbb{R}^{H \times W \times L \times Y}$. The **peer representation consistency training** depends on the following categorical distributions for both views:

$$\hat{\mathbf{y}}^{r}(\omega^{r}) = \operatorname{softmax}^{(\omega^{r})}(f_{\phi_{\text{SEG}}}(f_{\phi_{\text{FCN}}}(\mathbf{x}^{r}))), \\
\hat{\mathbf{y}}^{v}(\omega^{v}) = \operatorname{softmax}^{(\omega^{v})}(f_{\theta_{\text{SEG}}}(f_{\theta_{\text{EXR}}}(\mathbf{x}^{v}))),$$
(1)

where $\hat{\mathbf{y}}^r$ and $\hat{\mathbf{y}}^v$ are the categorical distributions for the range image (in ω^r in the $U \times V$ lattice Ω^r) and voxel (in ω^v in the $H \times W \times L$ lattice Ω^v) representations, and softmax^(ω)(.) computes the softmax activation at ω . Based on these outputs, we perform the following cross-view transformations:

$$\begin{split} \tilde{\mathbf{y}}^{r}(\omega^{r}) &= \arg\max \Psi_{v \to r}^{(\omega^{r})}(\hat{\mathbf{y}}^{v}), \quad \mathbf{c}^{r}(\omega^{r}) = \max \Psi_{v \to r}^{(\omega^{r})}(\hat{\mathbf{y}}^{v}), \\ \tilde{\mathbf{y}}^{v}(\omega^{v}) &= \arg\max \Psi_{r \to v}^{(\omega^{v})}(\hat{\mathbf{y}}^{r}), \quad \mathbf{c}^{v}(\omega^{v}) = \max \Psi_{r \to v}^{(\omega^{v})}(\hat{\mathbf{y}}^{r}), \end{split}$$
(2)

⁶ Range-to-point projection can lead to information loss [45], which is typically mitigated by post-processing with K-Nearest Neighbors (KNN). In pursuit of efficiency, our IT2 does not use any post-processing during training.

where $\tilde{\mathbf{y}}^r$ and $\tilde{\mathbf{y}}^v$ represent the transformed hard pseudo label maps, \mathbf{c}^r and \mathbf{c}^v denote the transformed confidence maps for range and voxel representations, respectively, $\Psi_{v \to r} = \psi_{p \to r} \circ \psi_{v \to p}$, $\Psi_{r \to v} = \psi_{p \to v} \circ \psi_{r \to p}$, and the superscript ω indicates that we only consider the output at ω . We use the cross-view pseudo labels from Eq. (2) to formulate the loss function to train IT2, as follows:

 $\ell_{\mathrm{IT2}}(\mathcal{D}_L, \mathcal{D}_U, f_{\phi_{\mathbf{FCN}}}, f_{\phi_{\mathbf{EXR}}}, f_{\phi_{\mathbf{SEG}}}) = \ell_{\mathrm{range}}(\mathcal{D}_L, \mathcal{D}_U, f_{\phi_{\mathbf{FCN}}}, f_{\phi_{\mathbf{SEG}}}) + \ell_{\mathrm{voxel}}(\mathcal{D}_L, \mathcal{D}_U, f_{\theta_{\mathbf{EXR}}}, f_{\theta_{\mathbf{SEG}}}),$ (3)

where we formulate the loss for each view as:

$$\ell_{\text{range}}(\mathcal{D}_{L}, \mathcal{D}_{U}, f_{\phi_{\text{FCN}}}, f_{\phi_{\text{SEG}}}) = \frac{1}{|\mathcal{D}_{L}| \times |\Omega^{r}|} \sum_{i=1}^{|\mathcal{D}_{L}|} \sum_{\omega^{r} \in \Omega^{r}} \ell(\mathbf{y}_{i}^{r}(\omega^{r}), f_{\phi_{\text{SEG}}}^{(\omega^{r})}(f_{\phi_{\text{FCN}}}(\mathbf{x}_{i}^{r})) + \frac{1}{|\mathcal{D}_{U}| \times |\Omega^{r}|} \sum_{i=1}^{|\mathcal{D}_{U}|} \sum_{\omega^{r} \in \Omega^{r}} \ell(\mathbf{\tilde{y}}_{i}^{r}(\omega^{r}), f_{\phi_{\text{SEG}}}^{(\omega^{r})}(f_{\phi_{\text{FCN}}}(\mathbf{x}_{i}^{r})), \ell_{\text{voxel}}(\mathcal{D}_{L}, \mathcal{D}_{U}, f_{\theta_{\text{EXR}}}, f_{\theta_{\text{SEG}}}) = \frac{1}{|\mathcal{D}_{L}| \times |\Omega^{v}|} \sum_{i=1}^{|\mathcal{D}_{L}|} \sum_{\omega^{v} \in \Omega^{v}} \ell(\mathbf{\tilde{y}}_{i}^{v}(\omega^{v}), f_{\theta_{\text{SEG}}}^{(\omega^{v})}(f_{\theta_{\text{EXR}}}(\mathbf{x}_{i}^{v}))) + \frac{1}{|\mathcal{D}_{U}| \times |\Omega^{v}|} \sum_{i=1}^{|\mathcal{D}_{U}|} \sum_{\omega^{v} \in \Omega^{v}} \ell(\mathbf{\tilde{y}}_{i}^{v}(\omega^{v}), f_{\theta_{\text{SEG}}}^{(\omega^{v})}(f_{\theta_{\text{EXR}}}(\mathbf{x}_{i}^{v})),$$

with the loss function $\ell(.)$ denoting the combination of cross-entropy loss and Lovasz-softmax [3] loss, following [27, 31, 66].

3.2 Cross-distribution Contrastive Learning

We model the cross distribution of the embedding space using a multi-layer perceptron to generate latent embeddings in parallel with the segmentation head. Specifically, we obtain the range image embedding space via $f_{\phi_{\mathbf{PRJ}}}^{\tau}$: $\mathcal{I} \to \mathbb{R}^{U \times V \times Z}$ with Z denoting the pixel-wise number of dimensions in the projected space. Similarly, the voxel embedding space is obtained via $f_{\theta_{\mathbf{PRJ}}}^{v}$: $\mathcal{K} \to \mathbb{R}^{H \times W \times L \times Z}$, where \mathcal{I} and \mathcal{K} are intermediate feature spaces introduced in Sec. 3.1. We formulate the observed embedding results with:

$$\mathbf{z}^{r}(\omega^{r}) = f_{\phi_{\mathbf{PRJ}}}^{(\omega^{r})}(f_{\phi_{\mathbf{FCN}}}(\mathbf{x}^{r})), \quad \mathbf{z}^{v}(\omega^{v}) = f_{\theta_{\mathbf{PRJ}}}^{(\omega^{v})}(f_{\theta_{\mathbf{EXR}}}(\mathbf{x}^{v})), \tag{5}$$

where \mathbf{z}^r and \mathbf{z}^v represent the latent point-wise embeddings of the range and voxel representations, respectively.

As mentioned in Sec. 2, contrastive learning methods [7, 12, 23, 31, 46] only use a limited subset of actual training samples present in a mini batch, potentially leading to suboptimal convergence with confirmation bias. We propose a novel approach to represent the range and voxel embeddings with class-specific Gaussian Mixture Models (GMMs). Such approach enables the sampling of the most informative embeddings, fostering a comprehensive understanding of the distribution across the feature space of different representations. In practice, we concatenate the embeddings from both range and voxel based on their categories, where such embedding set is defined by $\mathcal{G}^y = \{(\mathbf{z}, \mathbf{c}, y) | \mathbf{z} = \mathbf{z}_i^r(\omega^r), \mathbf{c} =$ $\mathbf{c}_{i}^{r}(\omega^{r}), y = \tilde{\mathbf{y}}_{i}^{r}(\omega^{r}) \} \bigcup \{ (\mathbf{z}, \mathbf{c}, y) | \mathbf{z} = \mathbf{z}_{i}^{v}(\omega^{v}), \mathbf{c} = \mathbf{c}_{i}^{v}(\omega^{v}), y = \tilde{\mathbf{y}}_{i}^{v}(\omega^{v}) \}$ for the unlabelled data indexed by $i \in \{1, ..., |\mathcal{D}_{U}|\}$, with the related confidence (i.e., \mathbf{c}^{r} and \mathbf{c}^{v}) and pseudo labels (i.e., $\tilde{\mathbf{y}}^{r}$ and $\tilde{\mathbf{y}}^{v}$) from Eq. (2) at positions $\omega^{r} \in \Omega^{r}$ and $\omega^{v} \in \Omega^{v}$. The class-specific GMMs are defined for both representations by

$$\mathbf{P}_{\Gamma^{y}}(\mathbf{z}|y) = \sum_{m=1}^{M} \mathbf{P}_{\Gamma^{y}}(m|y) \mathbf{P}_{\Gamma^{y}}(\mathbf{z}|m,y) = \sum_{m=1}^{M} \pi^{y}_{m} \mathcal{N}(\mathbf{z};\boldsymbol{\mu}^{y}_{m},\boldsymbol{\Sigma}^{y}_{m}),$$
(6)

where $\pi_m^y = \mathbf{P}_{\Gamma^y}(m|y)$ denotes the prior probability for the *m*-th GMM component of class $y \in \{0,1\}^Y$, $\mu_m^y \in \mathbb{R}^Z$ and $\Sigma_m^y \in \mathbb{R}^{Z \times Z}$ represent the mean and covariance of the corresponding GMM component and we fix M = 5 for all experiments. The GMM parameter optimisation in Eq. (6) relies on the EM [11] algorithm using the training set \mathcal{G}^y defined above with the following steps:

$$\mathbf{E} \text{ step:} \quad \mathbf{q}_{n,m}^{y} = \frac{\boldsymbol{\pi}_{m}^{y} \mathcal{N}(\mathbf{z}_{n} \mid \boldsymbol{\mu}_{m}^{y}, \boldsymbol{\Sigma}_{m}^{y})}{\sum_{m}^{M} \boldsymbol{\pi}_{m}^{y} \mathcal{N}(\mathbf{z}_{n} \mid \boldsymbol{\mu}_{m}^{y}, \boldsymbol{\Sigma}_{m}^{y})}, \text{ for } (\mathbf{z}_{n}, \mathbf{c}_{n}, y_{n}) \in \mathcal{G}^{y}$$

$$\mathbf{M} \text{ step:} \quad \boldsymbol{\mu}_{m}^{y} = \frac{1}{|\mathcal{G}^{y}|} \sum_{(\mathbf{z}_{n}, \mathbf{c}_{n}, y_{n}) \in \mathcal{G}^{y}} \mathbf{c}_{n} \mathbf{q}_{n,m}^{y_{n}} \mathbf{z}_{n}$$

$$\boldsymbol{\Sigma}_{m}^{y} = \frac{1}{|\mathcal{G}^{y}|} \sum_{(\mathbf{z}_{n}, \mathbf{c}_{n}, y_{n}) \in \mathcal{G}^{y}} \mathbf{c}_{n} \mathbf{q}_{n,m}^{y_{n}} (\mathbf{z}_{n} - \boldsymbol{\mu}_{m}^{y_{n}}) (\mathbf{z}_{n} - \boldsymbol{\mu}_{m}^{y_{n}})^{\mathsf{T}},$$

$$(7)$$

where \mathcal{G}^y represents the set of embedding samples that share the same label y. The posterior $\mathbf{q}_{n,m}^y$ estimates the responsibility of the n^{th} embedding \mathbf{z} from class y in \mathcal{G}^y being assigned to the m-th component. The confidence \mathbf{c}_n represents the potential correctness of the pseudo label from Eq. (2), which is used as a weighting factor to reduce the effect of low-confidence segmentation results [8,9] in the estimation of the GMM parameters. Note that we maintain a uniform weight of $\pi_m^y = 1/M$ for each component [43, 57] to prevent one component from dominating others.

GMM-based Prototypical Contrastive Learning. The class-specific GMMs from Eq. (7) are used to generate a set of highly representative embeddings for the contrastive learning process explained next. These representative class-specific embeddings is denoted by $S^y = \{\mathbf{s} | \mathbf{s} \sim \mathbf{P}_{\Gamma^y}(\mathbf{z} | y)\}$, with $\mathbf{P}_{\Gamma^y}(\mathbf{z} | y)$ being the class-specific GMM defined in Eq. (6). Subsequently, these generated virtual prototypes are utilised to guide the contrastive learning for the incoming embedding features from both range and voxel representations [54], without joining the backpropagation computation graph. We define the set of incoming features with the anchor set $\mathcal{A} = \{(\mathbf{z}, y)\}$ that is formed similarly to how we form the set \mathcal{G}^y without including the confidence value. The samples in \mathcal{A} are selectively chosen for backpropagation [54], but the informative samples in \mathcal{S}^y contribute to a comprehensive understanding of the cross-distribution for both representations. The contrastive learning relies on the minimisation of the following loss:

$$\ell_{\rm cross}(\mathcal{D}_L, \mathcal{D}_U, f_{\phi_{\rm PRJ}}, f_{\phi_{\rm FCN}}, f_{\theta_{\rm PRJ}}, f_{\theta_{\rm EXR}}) = \sum_{(\mathbf{z}, y) \in \mathcal{A}} \sum_{\mathbf{s} \in \mathcal{S}^y} \frac{\exp(\mathbf{z} \cdot \mathbf{s}/\tau)}{\exp(\mathbf{z} \cdot \mathbf{s}/\tau) + \sum_{\mathbf{s}^- \in \bar{\mathcal{S}}^y} \exp(\mathbf{z} \cdot \mathbf{s}^-/\tau)},$$
(8)

where \bar{S}^y is the set of virtual negative embeddings for class y.

3.3 Representation Specific Data Augmentation

Following [39, 61], we apply spatial data augmentations to the pseudo-labeling strategy in Eq. (2) to increase the robustness of our approach to complex scenes. After experimenting with classic data augmentation processes (e.g., CutMix [63] and LaserMix [27]), we notice that the application of distinct spatial augmentations to different viewpoints in our IT2 framework results in more effective generalization. We have empirically found the multi-boxes CutMix [63] augmentation for the range image and single-inclination LaserMix [27] augmentation for the voxel representation yield the most robust results. Although this representation-specific data augmentation is not our main contribution, it is shown empirically to produce relevant improvements. Please see details in supplementary Sec. 2.

3.4 Overall Training

The overall training of our IT2 minimises the following loss function:

$$\mathcal{L}(\mathcal{D}_L, \mathcal{D}_U, \mathcal{D}_L, \mathcal{D}_U, f_{\phi_{\mathbf{FCN}}}, f_{\phi_{\mathbf{EXR}}}, f_{\phi_{\mathbf{SEG}}} f_{\phi_{\mathbf{PRJ}}}, f_{\phi_{\mathbf{PRJ}}}) = \ell_{\mathrm{IT2}}(\mathcal{D}_L, \mathcal{D}_U, f_{\phi_{\mathbf{FCN}}}, f_{\phi_{\mathbf{EXR}}}, f_{\phi_{\mathbf{SEG}}}) + \ell_{\mathrm{cross}}(\mathcal{D}_L, \mathcal{D}_U, f_{\phi_{\mathbf{PRJ}}}, f_{\phi_{\mathbf{FCN}}}, f_{\theta_{\mathbf{PRJ}}}, f_{\theta_{\mathbf{EXR}}}),$$
(9)

where the entire training pipeline is performed in an end-to-end manner, where in each iteration, we minimise the loss from $\ell_{\rm IT2}(.)$ and $\ell_{\rm cross}(.)$ simultaneously.

4 Experiment

4.1 Experimental Setup

We evaluate our approach on three benchmarks, including nuScenes [4], SemanticKITTI [2] and ScribbleKITTI [52]. The **nuScenes** [4] dataset has 29, 130 training scans and 6, 019 validation scans. Each LiDAR scan contains 32 beams with $[-30^{\circ}, 10^{\circ}]$ vertical field of views (FOV). The dataset has 23 categories, but following previous works [27, 66], we merge them into 16 for training and validation. The **SemanticKITTI** [2] and its extension **ScribbleKITTI** [52] are popular laser-based segmentation benchmarks, where the latter one is based on line-scribbles labelling to reduce the annotation effort. The SemanticKITTI [2] dataset has 19 categories for semantic segmentation, where every LiDAR scan contains 64 beams with $[3^{\circ}, -25^{\circ}]$ vertical FOV. There are 22 sequences of Li-DAR scans in the dataset, where sequences 0 to 9 are used for training and validation. Following [23, 27], we use 19, 130 scans (sequences 0 to 7 and sequence 9) for training and 4, 071 scans from sequence 8 for validation.

Sampling strategies. Previous works [23, 27, 31] propose different sampling strategies to select the labelled data for the semi-supervised learning setting, including: 1) uniform [27] sampling of the labelled data across all the training

Table 1:	Uniform sa	mpling re	sults ϵ	across	various be	encl	nmarks, when	re we	foll	ow	[27]
to utilise	half feature	dimension	and lo	ower	resolution	in	Cylinder3D	[66]	for	all	the
experime	nts. Best resu	ilts are high	hlighte	d in r	ed.						

Dopr	Mathad	n	uSce	nes [4	1]	SemanticKITTI [2]			ScribbleKITTI [52]				
nepr.	Method	1%	10%	20%	50%	1%	10%	20%	50%	1%	10%	20%	50%
	sup.	38.3	57.5	62.7	67.6	36.2	52.2	55.9	57.2	33.1	47.7	49.9	52.5
Q	MT [50]	42.1	60.4	65.4	69.4	37.5	53.1	56.1	57.4	34.2	49.8	51.6	53.3
bu	CBST [68]	40.9	60.5	64.3	69.3	39.9	53.4	56.1	56.9	35.7	50.7	52.7	54.6
Ra	CPS [6]	40.7	60.8	64.9	68.0	36.5	52.3	56.3	57.4	33.7	50.0	52.8	54.6
	LaserMix [27]	49.5	68.2	70.6	73.0	43.4	58.8	59.4	61.4	38.3	54.4	55.6	58.7
	IT2 (Ours)	56.5	71.3	73.4	74.0	51.9	60.3	61.7	62.1	46.6	57.1	57.3	58.6
Voxel	sup.	50.9	65.9	66.6	71.2	45.4	56.1	57.8	58.7	39.2	48.0	52.1	53.8
	MT [50]	51.6	66.0	67.1	71.7	45.4	57.1	59.2	60.0	41.0	50.1	52.8	53.9
	CBST [68]	53.0	66.5	69.6	71.6	48.8	58.3	59.4	59.7	41.5	50.6	53.3	54.5
	CPS [6]	52.9	66.3	70.0	72.5	46.7	58.7	59.6	60.5	41.4	51.8	53.9	54.8
	LaserMix [27]	55.3	69.9	71.8	73.2	50.6	60.0	61.9	62.3	44.2	53.7	55.1	56.8
	IT2 (Ours)	57.5	72.1	73.6	74.1	52.0	61.4	62.1	62.5	47.9	56.7	57.5	58.3

sequences, 2) partial [23] sampling of the labelled data only from certain sequences, and 3) sampling of significant [31] labelled frames from the sequences. In the experiment, we examine our approach under all these sampling protocols. **Metrics.** We measure performance with mean intersection over union (mIoU). In each benchmark, we use the sampling strategies above to label n% of the training set, with the rest (100 - n)% being unlabelled.

Implementation details. In the range view, we use the FIDNet architecture [65] with ResNet34 [19] backbone. The input range image has size 32×1920 for nuScenes [4] and 64×2048 for SemanticKITTI [2] and ScirbbleKITTI [52], following [27]. In the voxel representation, we use the Cylinder3D architecture [66], where we fix the intermediate layer at 16 dimensions with smaller resolution as specified in [27] for the uniform sampling experiments to enable a fair comparison. For contrastive learning, we build a 3-layer projector [24,54]. The projector outputs an embedding map with 64 channels, where we apply the temperature factor $\tau = .1$ in Eq. (8). See supplementary Sec. 3 and Sec. 7 for more details.

4.2 Results on Different Partition Protocols

Uniform sampling results. We first compare our method using the uniform sampling benchmark [27] in Tab. 1, with the experimental setup from [27] (i.e., architecture, split ID and the resolution of input representations). Our framework outperforms the previous SOTA Lasermix [27] by a large margin for almost all partition protocols. For example, we observe $\approx 3.4\%$ and $\approx 2.0\%$ improvements on average for range and voxel representations, respectively. Notably, we highlight that our approach yields significant improvements for the 1% partition protocol in the line-scribble dataset [52], providing 8.28% improvement in the

Table 2: Partial & significant sampling results in the SemanticKITTI [2] benchmark. We compare with previous SOTA based on Cylinder3D [66] architecture, where <u>TTA</u> denotes test time augmentation. † indicates that we re-implement results based on the supported checkpoints from their GitHub – please see Supplementary Sec. 1 for more details. Best results in each evaluation protocol are highlighted in red.

TTA	Method	SemanticKITTI(par.)				Mothod	SemanticKITTI(sig.)				
		5%	10%	20%	40%	Method	5%	10%	20%	40%	
x	GPC^{\dagger} [23]	39.9	45.9	55.7	57.1	lim3D [31]	56.8	59.6	60.5	61.4	
^	$IT2 \ ({\rm Ours})$	43.4	48.1	56.5	60.9	$IT2 \ ({\rm Ours})$	60.3	63.3	64.0	64.8	
 Image: A second s	GPC^{\dagger} [23]	40.2	46.3	56.5	57.6	lim3D [31]	59.5	62.2	63.1	63.3	
	$IT2 \ ({\rm Ours})$	43.8	49.2	58.8	62.4	$IT2 \ ({\rm Ours})$	61.7	64.4	64.7	65.8	

range image and 3.66% in the voxel representation.

Partial and significant sampling results. In Tab. 2, we compare our results with GPC [23] and lim3D [31] in partial sampling (par.) and significant sampling (sig.), based on both original performance and test time augmentation (TTA). In the **partial sampling** protocol ⁷, our approach yields superior results across all the labelled data situations for both with and without TTA. Notably, our results surpass GPC [23] by 4.5% and 3.6% in the limited 5% labelled data condition, demonstrating its good generalisation capability. In the **significant sampling** results, our method also achieves better performance than lim3D [31] for all the partition protocols. For example, our results are better than the previous SOTA by about 2.2% and 1.6% in 10% and 20% labelled data with TTA, respectively.

4.3 Ablation Study

Improvements from our contributions. Table 3 presents our ablation study on the nuScenes [4] and ScribbleKITTI [52] datasets across various partition protocols. We investigate the roles of the IT2 architecture, cross-distribution contrastive learning (Ctrs.), and representation-related data augmentation (Aug.). The baseline utilises network perturbation through single representations [6] from [27]. Our IT2 architecture results in improvements of approximately 4.1% and 2.5% for the voxel representation in 10% labelled data for nuScenes [4] and ScribbleKITTI [52], respectively. The Ctrs. further improves by $\approx 2.4\%$ and $\approx 1.1\%$ for range and voxel representations in both datasets in 10% labelled data. Such enhancements show the effectiveness of our cross-distribution prototypes and contrastive learning strategy. Furthermore, Aug. results in improvements of 2.2% and 1.4% for range and voxel representations in the ScribbleKITTI [52] dataset for 10% labelled data, demonstrating the better generalisation brought by the representation-specific data augmentation.

⁷ The original GPC paper [23] calculates the mIoU differently compared to common practices [27, 31, 66]. The reported performance is based on the re-implementation from the official supported checkpoints of their GitHub, where more details are in Supplementary Sec. 1.

12 Y. Liu et al.

Table 3: Ablation study in the nuScenes [4] and ScribbleKITTI [52] datasets via uniform sampling. IT2 represents our architecture, Ctrs. denotes our cross-distribution contrastive learning and Aug. is our representation-specific data augmentation. The vellow rows denote our single representation baseline with CPS [6] from [27].

Ropr	ITO	Ctra	Aug	n	uScenes [4]	ScribbleKITTI [52]			
		Ours.	Aug.	10% $20%$ $50%$		50%	10%	20%	50%	
	X	X	X	60.8	64.9	68.0	50.0	52.8	54.6	
nge		×	x	$67.3 \scriptstyle{(6.5\uparrow)}$	$70.8 \scriptstyle{(5.9\uparrow)}$	$72.6 \tiny (4.6\uparrow)$	$53.2 \hspace{0.1 cm} {}_{(3.2\uparrow)}$	55.0 (2.27)	$56.2 (1.6\uparrow)$	
Ra^{j}	1	1	X	$70.3 \scriptstyle{(9.5\uparrow)}$	$72.6 \ \text{(7.7\uparrow)}$	$73.5 \scriptscriptstyle{(5.5\uparrow)}$	$54.9 \tiny (4.9\uparrow)$	$56.1 \tiny (3.3\uparrow)$	$57.2 (2.6\uparrow)$	
	1	1	1	$71.3 \scriptscriptstyle{(10.5\uparrow)}$	$73.4 \scriptscriptstyle{(8.5\uparrow)}$	$74.0 \scriptstyle{~(6.0\uparrow)}$	$57.1 \scriptstyle{(7.1\uparrow)}$	$57.3 \tiny (4.5\uparrow)$	$58.6 \tiny (4.0\uparrow)$	
	X	X	X	66.3	70.0	72.5	51.8	53.9	54.8	
xel		×	x	$70.4 (4.1\uparrow)$	$71.9 \scriptstyle{(1.9\uparrow)}$	$73.1 \scriptstyle{(0.6\uparrow)}$	$54.3 \scriptscriptstyle{(2.5\uparrow)}$	$55.6 (2.7\uparrow)$	$56.6 \tiny (1.8\uparrow)$	
V_{0}	1	1	X	$71.6 \scriptscriptstyle ~(5.3\uparrow)$	$72.9 \scriptscriptstyle (2.9\uparrow)$	$73.9 \scriptscriptstyle (1.4\uparrow)$	$55.3 \scriptscriptstyle{(3.5\uparrow)}$	$56.5 \tiny (3.6\uparrow)$	$57.6 \scriptscriptstyle (2.8\uparrow)$	
	1	1	1	$72.1 \scriptstyle{(5.8\uparrow)}$	$73.6 \scriptscriptstyle (3.6\uparrow)$	$74.1 \scriptscriptstyle (1.6\uparrow)$	$56.7 \tiny (4.9\uparrow)$	$57.5 \tiny (4.6\uparrow)$	$58.3 \scriptscriptstyle{(3.5\uparrow)}$	



Table 4: Contrastive Learning study in nuScenes [4], where ContrasSeg [54] is reimplemented from the SOTA image (pixellevel) semantic segmentation.

Mothod	Bopr	nuScenes [4]				
Method	nepr.	10%	20%			
ContracSog [54]	range	70.3	72.3			
Contrasseg [54]	voxel	71.2	72.8			
Quira	range	$71.3 \scriptscriptstyle{(1.0\uparrow)}$	$73.4 \scriptscriptstyle (1.1\uparrow)$			
Juis	voxel	$72.1 \scriptstyle{(0.9\uparrow)}$	$73.6 \scriptstyle (0.8\uparrow)$			

Fig. 3: Ablation study of temperature value τ from Eq. (8) in nuScenes [4] dataset under 10% labelled partition protocol.

Cross-distribution contrastive learning. We investigate the impact of the hyperparameter τ from Eq. (8) in Fig. 3 using the 10% labelled protocol on the nuScenes [4] dataset. Our observations indicate that $\tau = 0.10$ yields the best results, so we fix this value for all experiments. In Tab. 4, we compare our approach with our re-implementation of the label-guided pixel-wise contrastive learning [54] based on our IT2 architecture with different ratio of labelled data, where our method shows $\approx 1\%$ improvement for both representations, demonstrating the effectiveness of our proposed cross-distribution contrastive learning. **Representation-specific augmentation.** In Fig. 4, we have demonstrated the augmentation results for each representation in the ScribbleKITTI [52] dataset, where 'multi' denotes multi-box [63], and '1-inc' represents single inclination setups [27]. In our framework, one representation's results rely on the other's generalisation. For instance, '(multi) CutMix' reaches 57.1% mIoU with '(1-inc) LaserMix,' improving by 0.5% compared to the results with 'LaserMix.' Meanwhile, '(multi) CutMix' combined with '(1-inc) LaserMix' improves by 1.3% in the range and 0.7% in voxel representations compared to LaserMix [27], highlighting the improved generalisation benefits from our augmentation.

Ablation study for other representations. In Tab. 5, we conduct abla-



Fig. 4: Ablation study for the repr.-specific augmentations in ScribbleKITTI [52] with 10% labelled data, where 'multi' is multiboxes and '1-inc' is single inclination.

Table 5: Ablation study for the *bird's eye* view (BEV) experiment in the nuScenes [4], where **results** are from PolarNet [64]. The baseline rely on single view [6], followed by BEV with range image or voxel grid.

Bopr	BEV (PolarNet [64])							
nepr.	10%	20%	50%					
sup.	58.5	63.9	68.4					
$\mathbf{w}/\ \mathbf{bev}$	$62.4 \scriptscriptstyle (3.9\uparrow)$	$66.7 \scriptscriptstyle (2.8\uparrow)$	$70.0 \tiny (1.6\uparrow)$					
w/ range	64.8 (6.3 [†])	$67.9 (4.0\uparrow)$	$70.6 \tiny (2.2\uparrow)$					
w/ voxel	$66.3 \scriptscriptstyle{(7.8\uparrow)}$	$69.1 \scriptscriptstyle{(5.2\uparrow)}$	$71.6 \scriptscriptstyle (3.2\uparrow)$					



Fig. 5: Error maps visualised from LiDAR points (top) and *range view* (bottom) in 2 examples of the nuScenes [4] dataset. The correct predictions are in green and the mistakes are highlighted in red.

tion studies with an additional representation, the Bird's Eye View (BEV), on the nuScenes [4] dataset, where <u>all the results</u> in the table are based on PolarNet [64]. We re-implement the supervised results and single representation supervision [6] in the first and second rows, respectively. We emphasize that employing peer representations can significantly enhance performance compared to the single-representation baseline [6]. For instance, under the 10% labelled partition protocol, employing BEV (based on PolarNet [64]) with the range image (based on FidNet [65]) and with the voxel grid (based on Cylinder3D [66]) lead to improvements of 2.4% and 3.9% in mIoU, respectively.



Fig. 6: T-SNE visualisation for the pixel- (and voxel-) wise latent embedding features in the nuScenes [4] dataset. The features \bigcirc and \bigcirc indicate range and voxel representations and they are colored according to the category label.

4.4 Visualisation

Visualisation of prediction error maps. We compare the error maps from predictions in Fig. 5, including Mean Teacher [50] (2^{nd} column), LaserMix [27] (3^{rd} column) and ours (last column). The visualisations show the bird's eye view's output in the first row, followed by the range view results in the second row. Notably, our method's prediction demonstrates a superior qualitative result, exhibiting fewer mistaken predictions (highlighted in red) and more correct predictions (marked in green) for both views.

T-SNE visualisation of the latent features. Figure 6 displays the latent embedding distribution from the nuScenes [4] dataset. Fig. 6a shows the results without contrastive learning, Fig. 6b displays the results based on the label-guided contrastive learning [54], and our prototypical contrastive learning is in Fig. 6c. The samples denoted by $\underline{\bullet}$ and \bigcirc indicate the latent embeddings from range and voxel models, respectively, where different colours represent the classes. Fig. 6a displays an unstructured latent representation due to the absence of contrastive learning, but Fig. 6b successfully clusters class embeddings from to the same modalities (e.g., $\{\underline{\bullet}, \underline{\bullet}\}$ and $\{\bigcirc, \bigcirc\}$), yet it fails to capture the semantic relationships between different modalities within the same category. Our method, depicted in Fig. 6c, successfully clusters embeddings from the same class and different modalities $\{\underline{\bullet}, \bigcirc\}$, enabling a more effective distribution of embeddings for subsequent segmentation.

5 Conclusion

In this study, we introduced the novel semi-supervised LiDAR semantic segmentation method ItTakesTwo (IT2). IT2 improves generalisation with consistent predictions from peer LiDAR representations, and it also introduces a more effective contrastive learning using informative embeddings sampled from a learned distribution of positives and negatives. Finally, through experimental studies, we identify an efficient way to apply data augmentation to each representation. Our approach achieves best results across various benchmarks under all partition protocols, demonstrating its effectiveness. Acknowledgements. The project is supported by the Australian Research Council (ARC) through grant FT190100525.

References

- Arazo, E., Ortego, D., Albert, P., O'Connor, N.E., McGuinness, K.: Pseudolabeling and confirmation bias in deep semi-supervised learning. In: 2020 International Joint Conference on Neural Networks (IJCNN). pp. 1–8. IEEE (2020)
- Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., Gall, J.: Semantickitti: A dataset for semantic scene understanding of lidar sequences. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9297–9307 (2019)
- Berman, M., Triki, A.R., Blaschko, M.B.: The lovász-softmax loss: A tractable surrogate for the optimization of the intersection-over-union measure in neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4413–4421 (2018)
- Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11621–11631 (2020)
- Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. Advances in neural information processing systems 33, 9912–9924 (2020)
- Chen, X., Yuan, Y., Zeng, G., Wang, J.: Semi-supervised semantic segmentation with cross pseudo supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2613–2622 (2021)
- Cheng, M., Hui, L., Xie, J., Yang, J.: Sspc-net: Semi-supervised semantic 3d point cloud segmentation network. In: Proceedings of the AAAI conference on artificial intelligence. vol. 35, pp. 1140–1147 (2021)
- Cheung, Y.m.: A rival penalized em algorithm towards maximizing weighted likelihood for density mixture clustering with automatic model selection. In: Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004. vol. 4, pp. 633–636. IEEE (2004)
- Cheung, Y.m.: Maximum weighted likelihood via rival penalized em for density mixture clustering with automatic model selection. IEEE Transactions on Knowledge and Data Engineering 17(6), 750–761 (2005)
- Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: Medical Image Computing and Computer-Assisted Intervention-MICCAI 2016: 19th International Conference, Athens, Greece, October 17-21, 2016, Proceedings, Part II 19. pp. 424– 432. Springer (2016)
- Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the en algorithm. Journal of the royal statistical society: series B (methodological) 39(1), 1–22 (1977)
- Deng, S., Dong, Q., Liu, B., Hu, Z.: Superpoint-guided semi-supervised semantic segmentation of 3d point clouds. In: 2022 International conference on robotics and automation (ICRA). pp. 9214–9220. IEEE (2022)
- Fan, L., Xiong, X., Wang, F., Wang, N., Zhang, Z.: Rangedet: In defense of range view for lidar-based 3d object detection. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 2918–2927 (2021)

- 14. French, G., Aila, T., Laine, S., Mackiewicz, M., Finlayson, G.: Semi-supervised semantic segmentation needs strong, high-dimensional perturbations (2019)
- Graham, B., Engelcke, M., Van Der Maaten, L.: 3d semantic segmentation with submanifold sparse convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 9224–9232 (2018)
- Grigorescu, S., Trasnea, B., Cocias, T., Macesanu, G.: A survey of deep learning techniques for autonomous driving. Journal of Field Robotics 37(3), 362–386 (2020)
- Guo, Y., Wang, H., Hu, Q., Liu, H., Liu, L., Bennamoun, M.: Deep learning for 3d point clouds: A survey. IEEE transactions on pattern analysis and machine intelligence 43(12), 4338–4364 (2020)
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9729–9738 (2020)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- Hou, J., Graham, B., Nießner, M., Xie, S.: Exploring data-efficient 3d scene understanding with contrastive scene contexts. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15587–15597 (2021)
- Hou, Y., Zhu, X., Ma, Y., Loy, C.C., Li, Y.: Point-to-voxel knowledge distillation for lidar semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8479–8488 (2022)
- Hung, W.C., Tsai, Y.H., Liou, Y.T., Lin, Y.Y., Yang, M.H.: Adversarial learning for semi-supervised semantic segmentation. arXiv preprint arXiv:1802.07934 (2018)
- Jiang, L., Shi, S., Tian, Z., Lai, X., Liu, S., Fu, C.W., Jia, J.: Guided point contrastive learning for semi-supervised point cloud semantic segmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6423– 6432 (2021)
- Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. Advances in Neural Information Processing Systems 33, 18661–18673 (2020)
- Kohli, A.P.S., Sitzmann, V., Wetzstein, G.: Semantic implicit neural scene representations with semi-supervised training. In: 2020 International Conference on 3D Vision (3DV). pp. 423–433. IEEE (2020)
- Kong, L., Liu, Y., Chen, R., Ma, Y., Zhu, X., Li, Y., Hou, Y., Qiao, Y., Liu, Z.: Rethinking range view representation for lidar segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 228–240 (2023)
- Kong, L., Ren, J., Pan, L., Liu, Z.: Lasermix for semi-supervised lidar semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21705–21715 (2023)
- Lai, X., Chen, Y., Lu, F., Liu, J., Jia, J.: Spherical transformer for lidar-based 3d recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17545–17555 (2023)
- Le-Khac, P.H., Healy, G., Smeaton, A.F.: Contrastive representation learning: A framework and review. Ieee Access 8, 193907–193934 (2020)
- Li, J., Zhou, P., Xiong, C., Hoi, S.C.: Prototypical contrastive learning of unsupervised representations. arXiv preprint arXiv:2005.04966 (2020)
- Li, L., Shum, H.P., Breckon, T.P.: Less is more: Reducing task and model complexity for 3d point cloud semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9361–9371 (2023)

- 32. Li, M., Xie, Y., Shen, Y., Ke, B., Qiao, R., Ren, B., Lin, S., Ma, L.: Hybridcr: Weakly-supervised 3d point cloud semantic segmentation via hybrid contrastive regularization. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14930–14939 (2022)
- Liang, C., Wang, W., Miao, J., Yang, Y.: Gmmseg: Gaussian mixture based generative semantic segmentation models. Advances in Neural Information Processing Systems 35, 31360–31375 (2022)
- Liang, T., Xie, H., Yu, K., Xia, Z., Lin, Z., Wang, Y., Tang, T., Wang, B., Tang, Z.: Bevfusion: A simple and robust lidar-camera fusion framework. Advances in Neural Information Processing Systems 35, 10421–10434 (2022)
- Liu, M., Zhou, Y., Qi, C.R., Gong, B., Su, H., Anguelov, D.: Less: Label-efficient semantic segmentation for lidar point clouds. In: European Conference on Computer Vision. pp. 70–89. Springer (2022)
- Liu, W., Yue, X., Chen, Y., Denoeux, T.: Trusted multi-view deep learning with opinion aggregation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 7585–7593 (2022)
- 37. Liu, Y., Hu, Q., Lei, Y., Xu, K., Li, J., Guo, Y.: Box2seg: Learning semantics of 3d point clouds with box-level supervision. arXiv preprint arXiv:2201.02963 (2022)
- Liu, Y., Kong, L., Cen, J., Chen, R., Zhang, W., Pan, L., Chen, K., Liu, Z.: Segment any point cloud sequences by distilling vision foundation models. arXiv preprint arXiv:2306.09347 (2023)
- Liu, Y., Tian, Y., Chen, Y., Liu, F., Belagiannis, V., Carneiro, G.: Perturbed and strict mean teachers for semi-supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4258–4267 (2022)
- Liu, Z., Tang, H., Amini, A., Yang, X., Mao, H., Rus, D., Han, S.: Bevfusion: Multi-task multi-sensor fusion with unified bird's-eye view representation. In: IEEE International Conference on Robotics and Automation (ICRA) (2023)
- Liu, Z., Tang, H., Amini, A., Yang, X., Mao, H., Rus, D.L., Han, S.: Bevfusion: Multi-task multi-sensor fusion with unified bird's-eye view representation. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 2774– 2781. IEEE (2023)
- Manning, C.D.: An introduction to information retrieval. Cambridge university press (2009)
- Mena, G., Nejatbakhsh, A., Varol, E., Niles-Weed, J.: Sinkhorn em: an expectationmaximization algorithm based on entropic optimal transport. arXiv preprint arXiv:2006.16548 (2020)
- Miech, A., Alayrac, J.B., Smaira, L., Laptev, I., Sivic, J., Zisserman, A.: End-to-end learning of visual representations from uncurated instructional videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9879–9889 (2020)
- Milioto, A., Vizzo, I., Behley, J., Stachniss, C.: Rangenet++: Fast and accurate lidar semantic segmentation. In: 2019 IEEE/RSJ international conference on intelligent robots and systems (IROS). pp. 4213–4220. IEEE (2019)
- Nunes, L., Marcuzzi, R., Chen, X., Behley, J., Stachniss, C.: Segcontrast: 3d point cloud feature representation learning through self-supervised segment discrimination. IEEE Robotics and Automation Letters 7(2), 2116–2123 (2022)
- Ouali, Y., Hudelot, C., Tami, M.: Semi-supervised semantic segmentation with cross-consistency training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12674–12684 (2020)

- Reichardt, L., Ebert, N., Wasenmüller, O.: 360deg from a single camera: A few-shot approach for lidar segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1075–1083 (2023)
- Sun, C.Y., Yang, Y.Q., Guo, H.X., Wang, P.S., Tong, X., Liu, Y., Shum, H.Y.: Semi-supervised 3d shape segmentation with multilevel consistency and part substitution. Computational Visual Media 9(2), 229–247 (2023)
- Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. arXiv preprint arXiv:1703.01780 (2017)
- Tian, Y., Krishnan, D., Isola, P.: Contrastive multiview coding. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16. pp. 776–794. Springer (2020)
- Unal, O., Dai, D., Van Gool, L.: Scribble-supervised lidar semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2697–2707 (2022)
- Vora, S., Radwan, N., Greff, K., Meyer, H., Genova, K., Sajjadi, M.S., Pot, E., Tagliasacchi, A., Duckworth, D.: Nesf: Neural semantic fields for generalizable semantic segmentation of 3d scenes. arXiv preprint arXiv:2111.13260 (2021)
- Wang, W., Zhou, T., Yu, F., Dai, J., Konukoglu, E., Van Gool, L.: Exploring crossimage pixel contrast for semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7303–7313 (2021)
- Wu, B., Wan, A., Yue, X., Keutzer, K.: Squeezeseg: Convolutional neural nets with recurrent crf for real-time road-object segmentation from 3d lidar point cloud. In: 2018 IEEE international conference on robotics and automation (ICRA). pp. 1887– 1893. IEEE (2018)
- Xie, S., Gu, J., Guo, D., Qi, C.R., Guibas, L., Litany, O.: Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16. pp. 574–591. Springer (2020)
- Xu, J., Hsu, D.J., Maleki, A.: Benefits of over-parameterization with em. Advances in Neural Information Processing Systems **31** (2018)
- Xu, J., Zhang, R., Dou, J., Zhu, Y., Sun, J., Pu, S.: Rpvnet: A deep and efficient range-point-voxel fusion network for lidar point cloud segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16024–16033 (2021)
- Xu, J., Tang, H., Ren, Y., Peng, L., Zhu, X., He, L.: Multi-level feature learning for contrastive multi-view clustering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16051–16060 (2022)
- Xu, Z., Yuan, B., Zhao, S., Zhang, Q., Gao, X.: Hierarchical point-based active learning for semi-supervised point cloud semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 18098–18108 (2023)
- Yang, L., Qi, L., Feng, L., Zhang, W., Shi, Y.: Revisiting weak-to-strong consistency in semi-supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7236–7246 (2023)
- Yang, L., Zhuo, W., Qi, L., Shi, Y., Gao, Y.: St++: Make self-training work better for semi-supervised semantic segmentation. arXiv preprint arXiv:2106.05095 (2021)
- 63. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6023–6032 (2019)

- Zhang, Y., Zhou, Z., David, P., Yue, X., Xi, Z., Gong, B., Foroosh, H.: Polarnet: An improved grid representation for online lidar point clouds semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9601–9610 (2020)
- Zhao, Y., Bai, L., Huang, X.: Fidnet: Lidar point cloud semantic segmentation with fully interpolation decoding. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 4453–4458. IEEE (2021)
- Zhu, X., Zhou, H., Wang, T., Hong, F., Ma, Y., Li, W., Li, H., Lin, D.: Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9939–9948 (2021)
- Zhuang, Z., Li, R., Jia, K., Wang, Q., Li, Y., Tan, M.: Perception-aware multisensor fusion for 3d lidar semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16280–16290 (2021)
- Zou, Y., Yu, Z., Kumar, B., Wang, J.: Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In: Proceedings of the European conference on computer vision (ECCV). pp. 289–305 (2018)
- Zou, Y., Zhang, Z., Zhang, H., Li, C.L., Bian, X., Huang, J.B., Pfister, T.: Pseudoseg: Designing pseudo labels for semantic segmentation. arXiv preprint arXiv:2010.09713 (2020)