Bayesian Self-Training for Semi-Supervised 3D Segmentation

Ozan Unal^{1,2}, Christos Sakaridis¹, and Luc Van Gool^{1,3,4}

¹ETH Zurich, ²Huawei Technologies, ³KU Leuven, ⁴INSAIT {ozan.unal, csakarid, vangool}@vision.ee.ethz.ch

Abstract. 3D segmentation is a core problem in computer vision and, similarly to many other dense prediction tasks, it requires large amounts of annotated data for adequate training. However, densely labeling 3D point clouds to employ fully-supervised training remains too labor intensive and expensive. Semi-supervised training provides a more practical alternative, where only a small set of labeled data is given, accompanied by a larger unlabeled set. This area thus studies the effective use of unlabeled data to reduce the performance gap that arises due to the lack of annotations. In this work, inspired by Bayesian deep learning, we first propose a Bayesian self-training framework for semi-supervised 3D semantic segmentation. Employing stochastic inference, we generate an initial set of pseudo-labels and then filter these based on estimated point-wise uncertainty. By constructing a heuristic n-partite matching algorithm, we extend the method to semi-supervised 3D instance segmentation, and finally, with the same building blocks, to dense 3D visual grounding. We demonstrate state-of-the-art results for our semisupervised method on SemanticKITTI and ScribbleKITTI for 3D semantic segmentation and on ScanNet and S3DIS for 3D instance segmentation. We further achieve substantial improvements in dense 3D visual grounding over supervised-only baselines on ScanRefer. Our project page is available at ouenal.github.io/bst/.

Keywords: Semi-supervised · Self-training · 3D semantic segmentation · 3D instance segmentation · 3D visual grounding

1 Introduction

Dense 3D semantic perception at the level of points is pivotal for several applications that are based on fine-grained processing and reasoning in 3D, e.g. autonomous cars, robotics, and augmented reality. Such perception requires a segmentation of the unstructured input 3D point cloud into different semantic classes, i.e. 3D semantic segmentation [60], or distinct instances of these classes, i.e., 3D instance segmentation [30]. The input point cloud is often coupled with human utterances that refer to individual objects in the former in cases where a visual AI system interacts with human users in a common 3D environment, necessitating a 3D instance-level segmentation of these objects based on the aforementioned referrals, a task which is known as dense 3D visual grounding or



Fig. 1: Illustration of our Bayesian pseudo-labeling pipeline for semi-supervised a) 3D semantic segmentation; b) 3D instance segmentation; and c) dense 3D visual grounding. With only slight adjustments, using the same building blocks, our method can be adapted to each of these tasks to achieve SOTA results.

referral-based 3D instance segmentation [48, 59]. In the majority of works that address these diverse 3D segmentation tasks, it is assumed that the training data come with full 3D semantic and/or verbal annotations, which creates the pressing need for large-scale labeled 3D datasets [1, 2, 6, 11] so that large end-to-end models can be adequately fitted on these data.

However, the creation of such fine-level 3D annotations and verbal referrals for 3D scans requires human labor to be fairly reliable. Such labor is time-consuming and expensive, especially due to the high dimensionality of the scanned point clouds, which makes it hard for human annotators to navigate through them, visually inspect them, and verbally describe them. These practical difficulties render learning frameworks for 3D segmentation that operate with partial labels essential, semi-supervised learning being a primary case. Answering this need, several recent works address semi-supervised 3D semantic segmentation [9, 25, 29], most of which focus on effective pseudo-labeling techniques for self-training on the larger, unlabeled portion of the training data. However, the generation of pseudo-labels in these works typically focuses on some grouping mechanism over 3D points, such as superpoints, to refine and propagate pseudo-labels and otherwise relies on naïve core techniques such as confidence thresholding [66] for generating the *initial* point-level pseudo-labels that are then post-processed by their proposed modules. Adding to this, there is a relative scarcity in works on semi-supervised 3D instance segmentation [10, 33], which we attribute to the relative difficulty in pseudo-labeling instances compared to creating plain semantic pseudo-labels. Last but not least, to our knowledge, no work has considered the setting of semi-supervised dense 3D visual grounding.

We jointly address these shortcomings of previous works by proposing a common self-training framework for semi-supervised 3D segmentation which is based on Bayesian deep learning and applies to all three aforementioned segmentation tasks. At the core of our self-training framework lies a variational Bayesian approximation of the uncertainty of the respective network's predictions via dropout-based Monte Carlo integration [15] over multiple augmented versions of the original 3D input. This approximation delivers a more informed uncertainty estimate than simple confidence thresholding techniques [66], which is crucial for eliminating confirmation bias from the pseudo-labeling process that we adopt for self-training. In particular, we utilize both the estimated Shannon entropy from the sampled predictions and the predicted label from each sample to filter out uncertain or inconsistent points from the generated set of pseudo-labels.

While this formulation is sufficient for the task of 3D semantic segmentation, it does not work out of the box for instance-level segmentation, as Monte Carlo integration cannot be applied directly to instance predictions. To account for this issue, we propose a novel and efficient, heuristic matching strategy across instances coming from different stochastic dropout forward passes based on linear sum assignment. This strategy enables the application of the same filtering steps as described above for semantic pseudo-labels to *instance* pseudo-labels and thus extends the scope of our approach to 3D instance segmentation. Going one step further, we observe that the same matching strategy that we devise for instance segmentation applies to the selection stage of dense 3D grounding, in which the correct instance from the set of segmented candidate instances needs to be selected based on the input verbal description. Thus, for dense 3D grounding, we apply our heuristic linear sum assignment twice, first for the purely 3Dbased candidate instance proposal generation and afterward for the verbo-visual instance selection. Our complete self-training pipeline for each of the three 3D segmentation tasks to which our method applies is summarized in Fig. 1.

We test our method on leading benchmarks across a wide range of domains and datasets. Specifically, our Bayesian self-training framework achieves state-of-the-art results for semi-supervised 3D semantic segmentation on SemanticKITTI [3] and ScribbleKITTI [47] and semi-supervised 3D instance segmentation on ScanNet [11] and S3DIS [2]. Furthermore, we are the first to explore a semi-supervised setting on ScanRefer [6] for 3D visual grounding and obtain consistent significant improvements over supervised-only baselines.

In summary, our contributions are as follows:

- 1. Following established uncertainty estimation techniques from active learning, we construct a Bayesian self-training framework for semi-supervised 3D semantic segmentation by generating reliable pseudo-labels through stochastic inference and entropy-based uncertainty filtering.
- 2. We introduce a novel heuristic matching algorithm to extend the proposed framework to 3D instance segmentation.
- 3. We further extend the method to dense 3D visual grounding utilizing the same building blocks as in the above tasks and are the first to explore this task in a semi-supervised setting.
- We obtain state-of-the-art results for semi-supervised 3D segmentation on SemanticKITTI [3], ScribbleKITTI [47], ScanNet [11] and S3DIS [2].

4 O. Unal et al.

2 Related Work

Semi-supervised 3D semantic segmentation has recently grown in popularity in the area of 3D understanding along with weakly supervised 3D semantic segmentation [47, 52, 57, 61], due to the difficulty in large-scale 3D labeling that is necessary for the supervised setup [60]. In this work, we tackle the task of semi-supervised segmentation where a large portion of data is completely unlabeled. As opposed to methods developed for weakly-supervised settings, a semisupervised method cannot leverage the availability of partial true label information within each data sample, e.g. via the ArcPoint loss [31] or a PLS descriptor [47]. A common technique for leveraging unlabeled data under this restrictive scheme is explored under pseudo-labeling. Semantic pseudo-labels are utilized by Jiang et al. [25] to guide a contrastive loss and in SSPC [9] to propagate labels on a superpoint graph. Superpoints are also leveraged by Deng *et al.* [13] to constrain pseudo-label propagation. In the 2D counterpart task, Chen et al. [8] enforces consistency of semantic pseudo-labels across perturbed versions of the 2D segmentation network. By contrast, we perturb both the network weights and the 3D inputs across different forward passes, and we consider both the entropy and the consistency among predictions in our pseudo-labels. Another related line of 2D semi-supervised works generates pseudo-labels based on direct thresholding of max-softmax scores to reduce confirmation bias [19, 66], while we threshold the *entropy* of the distribution of softmax scores. Active learning is combined with a 3D semi-supervised approach in DiAL [46] to streamline practical 3D segmentation annotation pipelines. LaserMix [29] introduces variability through cylindrical and range-view partitioning and mixing of 3D LiDAR scans, encouraging consistency across different mixes.

Semi-supervised 3D instance segmentation is still in its infancy compared to its instance-agnostic semantic segmentation counterpart reviewed above. In particular, most object-level semi-supervised 3D understanding works focus on the coarser task of 3D object detection [17,35,37,38,42,50,51,63], which outputs 3D bounding boxes that are not fine enough for precision-critical downstream tasks such as robotic grasping. For 3D instance/object segmentation, approaches that depart from the standard fully supervised setup have mostly considered the cases of unsupervised pre-training [56], weak supervision via one-point-oneinstance [20, 43], or unsupervised generic 3D object segmentation [32]. To our knowledge, the only two semi-supervised works that specifically consider 3D instance segmentation are TWIST [10] and WS3D [33]. The former work proposes a pseudo-label generation strategy tailored for the instance segmentation task, for which the technique of confidence thresholding that is ubiquitously used to the same end in semantic segmentation is not applicable. It thus introduces pseudo offset vectors, besides semantic proposal-level pseudo-labels, for performing self-training. The latter work proposes a region-level contrastive loss, which performs unsupervised instance discrimination, and leverages boundary information via a region-level unsupervised energy-based loss. On the contrary, we propose a unified semi-supervised Bayesian self-training framework that spans multiple 3D segmentation tasks including 3D instance segmentation, and handles the additional matching required between instances generated in different forward passes through the network via an efficient heuristic optimization step.

3D visual grounding is a central 3D task in the area of vision and language [4, 14, 28, 34, 36, 62]. Relevant models accept a 3D point cloud as visual input and perform grounding of the accompanying verbal description to the referred object. While most previous works [7, 23, 24, 34, 39] focus on coarse 3D bounding box outputs, the task of *dense* 3D visual grounding or referral-based 3D instance segmentation has been explored far less in comparison, in which a fine-grained point-level mask is predicted [22, 48, 59]. Regardless, all aforementioned methods invariably assume a *supervised* setting, where full ground-truth instance annotations are available for all samples at training time. By contrast, the formulation of our method makes it fit for also handling the more practical *semi-supervised* setting, in which only a part of the training data is labeled.

Uncertainty estimation is crucial in learning setups that lack full supervision as ours, given that such estimates can provide information about the reliability of the predictions of the model on unlabeled samples [21]. In the more recent deep learning era. Bayesian neural networks form the core for uncertainty estimation by modeling network weights as stochastic distributions, where testtime-dropout [41] is utilized as a variational Bayesian approximation [15] and Monte Carlo integration is computed via multiple stochastic forward passes to approximate prediction uncertainty [16]. This Bayesian deep learning framework has been applied successfully to supervised 2D dense prediction tasks [26,27]. We focus on the aggregate uncertainty in the predictions of a general 3D segmentation model and build on the Bayesian deep learning approach of Gal et al. [16] which examines the approximate predictive entropy as an acquisition function for active learning. However, contrary to Gal et al. [16], in our selftraining setup, we use *low* rather than high entropy as a positive indication for considering a point as a source of supervision. To the best of our knowledge, our work is the first to introduce such a Bayesian self-training framework for general 3D point cloud segmentation, specifically including 3D instance segmentation and referral-based 3D instance segmentation.

Related uncertainty-based approaches to ours have been proposed for the semantic segmentation of 3D grid volume data in the field of medical image computing [12,54,58]. In particular, Yu *et al.* [58] also use Monte Carlo dropout to estimate predictive entropy but only applies dropout to a single network input, whereas we vary both network weights and input views. Moreover, that work leverages the estimated entropy merely as the weight in a consistency loss between a teacher and a student network, while we filter pseudo-labels based on our entropy estimates. Besides, Xia *et al.* [54] perform semi-supervised semantic segmentation with multi-view co-training and applies Bayesian deep learning across views for uncertainty-based weighting of pseudo-labels. We only train one network instead of multiple networks employed in Xia *et al.* [54] and we derive the uncertainty estimates by feeding all input views to that single network, which is more parameter-efficient. Both of these approaches only address 3D semantic segmentation, and cannot be utilized in instance-level dense prediction tasks.



Fig. 2: Illustration of Bayesian pseudo-labeling for semi-supervised instance segmentation. We initialize a seed prediction via a forward pass using a non-augmented input. Then our heuristic *n*-partite matching algorithm is employed to pair each seed instance with the best matching predicted instance from each of the K stochastic forward passes. For each aligned object, we compute the aggregated label through unanimous voting and filter it based on uncertainty to obtain the final pseudo mask.

3 Bayesian Self-Training

In Sec. 3.1 we start by formulating a self-training approach for semi-supervised 3D semantic segmentation based on established uncertainty estimation techniques from the field of active learning [15, 46]. We then extend our formulation to include instance labels in Sec. 3.2 and finally build up to also take natural language prompts as input, i.e., towards semi-supervised dense 3D visual grounding in Sec. 3.3.

3.1 Semi-Supervised 3D Semantic Segmentation

In a semi-supervised setting, a dataset D consists of a labeled set S and an unlabeled set U. The baseline approach when working with such data is to treat the problem similarly to any fully supervised task and employ a semantic loss \mathcal{L}_{sem} on the available labeled scenes S, typically stated as cross-entropy. Formally, the objective function over the model parameters ω is formulated as:

$$\min_{\omega} \sum_{i=1}^{|S|} \mathcal{L}_{sem}(\hat{\boldsymbol{y}}, \mathbf{y}).$$
(1)

Starting from this baseline formulation, a semi-supervised training strategy aims to leverage the unlabeled scenes U to further improve model performance. One such common technique is self-training, where the model is iteratively trained on its own predictions, i.e. on generated per-point one-hot pseudo-label vectors $\hat{\mathbf{y}} = [\hat{\mathbf{y}}^{(1)}, \dots, \hat{\mathbf{y}}^{(C)}]$, with C denoting the number of classes. However, the task of generating pseudo-labels is not trivial, as the model may suffer from *confirmation bias* that occurs when repeatedly overfitting to incorrect labels. A simple but effective solution is explored under the umbrella of *threshold-ing* [5, 19, 55, 65, 66]. In essence, instead of training on all model predictions, a pseudo-label set L is formed by predictions that exceed a confidence of a determined threshold τ to ensure a higher level of accuracy. Formally, the objective function of threshold-based self-training for semantic segmentation can be formulated as:

$$\min_{\boldsymbol{\omega}, \hat{\mathbf{y}}} \sum_{d=1}^{|D|} \left(\mathcal{L}_{sem}(\hat{\mathbf{y}}, \mathbf{y}) - (\log \hat{\mathbf{y}}_d - \tau) \cdot \hat{\mathbf{y}}_d \right)$$
(2)

with \hat{y}_d denoting the softmax model predictions. Eq. 2 is optimized in two steps that are repeated until convergence:

- 1. **Training:** First, $\hat{\mathbf{y}}$ is fixed, and the objective function is optimized with respect to ω .
- 2. **Pseudo-labeling:** Then, ω is fixed and the objective function is optimized with respect to $\hat{\mathbf{y}}$. Consequently, L is updated given $\hat{\mathbf{y}}$.

The process is initialized by setting the latent variable $\hat{\mathbf{y}} = \mathbf{0}$ for all points, i.e. by reverting to the naïve solution of Eq. 1 and only training on the labeled set S. The two steps can be repeated to take advantage of the improved representation capability of the model thanks to pseudo-labeling.

While naïve thresholding relies on the softmax scores of the model to determine confidence, these scores may be unreliable and often skewed when facing long-tailed distributions, resulting in overconfident predictions [18]. Inspired by uncertainty estimation techniques commonly employed in active learning (AL) pipelines [21,46], we introduce a Bayesian self-training strategy that leverages an uncertainty-aware filtering technique for generating more reliable pseudo-labels.

We first start by decomposing the goal of 3D semantic segmentation, which is to discover the dependency of the point-wise distribution over the labels $y \in Y$ on an input sample $x \in X$ via the model weights ω . Formally, the task can be reduced to per-point classification, where y can be of class $c \in \{1, ..., C\}$. The conditional probability is therefore defined by:

$$p(y=c|x,D) = \int p(y=c|a(x),\omega)p(\omega|D)p(a(x)|x)d\omega dx$$
(3)

for the dataset D = (X, Y), under randomly applied affine transformations modeled through $a(\cdot)$.

We set p(a(x)|x) as a constant, restate the intractable $p(\omega|D)$ via the variational approximation $q(\omega)$ and minimize $\text{KL}(q(\omega)|p(\omega|X,Y))$. This objective can be approximated through Monte Carlo integration by employing variational inference via Bayesian networks, i.e. by applying stochastic forward passes with weights ω_k :

$$p(y = c|x, D) \approx \int p(y = c|a(x), \omega)q(\omega)d\omega dx$$

$$\approx \frac{1}{K} \sum_{k=1}^{K} p(y = c|a(x), \omega_k).$$
(4)

8 O. Unal et al.

Compared to the common decomposition of BALD [21], the inclusion of the affine transformations via $a(\cdot)$ allows a higher deviation amongst stochastic forward passes which in return allows better estimation of predictive uncertainty [49]. We capture the uncertainty via the Shannon entropy [40]:

$$\mathbb{H}[y|x, D] = -\sum_{c=1}^{C} p(y=c|x, D) \log p(y=c|x, D)$$
(5)

and the accumulated predicted label \hat{y}^{\ast} is given by the unanimous voting principle:

$$\hat{\mathbf{y}}^{*(c)} = \begin{cases} 1, \text{ if } K = \sum_{k} \mathbb{1}(\operatorname{argmax}(p(y|a(x), \omega_k)) = c) \\ 0, \text{ otherwise,} \end{cases}$$
(6)

with 1 denoting the indicator function.

Unlike what is commonly seen in active learning settings, our goal is not to maximize the total entropy but rather to minimize it in order to ensure a high-quality pseudo-label set. We therefore redefine the self-training objective (Eq. 2) for stochastic pseudo-labeling as:

$$\min_{\boldsymbol{\omega}, \hat{\mathbf{y}}} \sum_{d=1}^{D} \left(\mathcal{L}_{sem}(\hat{\mathbf{y}}, \mathbf{y}) + (\mathbb{H}[y|x, D] - \tau)(\hat{\mathbf{y}}^* \cdot \hat{\mathbf{y}}) \right)$$
(7)

with $\tau \in (0, 1)$ determining the cut-off entropy threshold. To solve the nonlinear integer optimization task, we employ the following solver:

$$\hat{\mathbf{y}}^{(c)} = \begin{cases} 1, \text{ if } \mathbb{H}[y|x, D] < \tau \text{ and} \\ c = \operatorname{argmax} \hat{\mathbf{y}}^* \\ 0, \text{ otherwise,} \end{cases}$$
(8)

where labels $\hat{\mathbf{y}}^{(c)} = 0 \ \forall \ c$ are ignored during training.

In practice, following precedent [47,66], we set τ based on p_{τ} , with p_{τ} determining the percentage of pseudo-labels to be sampled amongst the accumulated predicted labels.

In summary, we first determine a set of labels that are agreed upon by all stochastic forward passes, i.e. by forward passes with active dropout [15, 41]. We then compute the Shannon entropy for all predictions, filtering those that exhibit high entropy. This surprisingly simple but effective approach allows us to reduce the incorrect targets that we introduce to our pseudo-label set and to avoid the overconfidence issue of naïve thresholding-based approaches.

3.2 Semi-Supervised 3D Instance Segmentation

The proposed Bayesian self-training framework provides a strong basis for any semi-supervised 3D dense prediction task, one such example being 3D instance segmentation. Yet, compared to semantic segmentation where accumulation of

Algorithm 1: Bayesian pseudo-label generation for instance segmentation.

Input: Point cloud $P \in \mathbb{R}^{N \times 3}$, trained model with dropout ω , number of passes K, threshold τ 1 seed_output = $\omega(P) \in \mathbb{R}^{M \times N}$ $\mathbf{2} \ \mathbf{i} = \text{seed output.copy}()$ $\mathbf{3} \text{ seed}_{\max} = \text{to}_{\max}(\mathbf{i}) \in \{0,1\}^{M \times N}$ 4 unanimous = seed mask.copy() **5** for k in 1:K do $ext{output} = \omega_k(a(P)) \in \mathbb{R}^{M' imes N}$ 6 $mask = to_mask(output) \in \{0,1\}^{M' \times N}$ 7 cost = -iou(seed mask, mask)8 row, col = LSA(cost)9 unanimous[col] &= mask[row] 10 i[col] += output[row] 11 12 $\mathbb{H} = (K - i) \cdot \log(K - i) + i \log i$ **13** pseudo label = $(\mathbb{H} < \tau)$ & unanimous **Return:** pseudo label $\in \{0, 1\}^{M \times N}$

labels can be easily handled across multiple stochastic forward passes, accumulation of predicted instance masks is not trivial and requires *n*-partite matching.

However, the *n*-partite matching problem is NP-hard with a complexity of $|I|!^{(K-1)}$ with K passes and |I| number of instances assuming full bijection. To match predicted instance masks across various stochastic forward passes, we propose a heuristic method based on the Hungarian algorithm. First, we generate an initial prediction via an unaugmented forward pass (Alg. 1 L1&3). The goal is to establish a seed that is likely of a higher quality with minimal perturbations of the data. We then iteratively assign predicted masks of each stochastic forward pass via linear sum assignment, with the cost matrix formed by the negative pairwise IoUs (Alg. 1 L5-10).

Having matched predicted binary instance masks from individual stochastic forward passes, each resulting set of masks can then be treated as a set of 2-class semantic segmentation outputs, and thus our proposed Bayesian pseudo-labeling approach from Eq. 8 can be trivially applied (Alg. 1 L11-15). Our pipeline for semi-supervised 3D instance segmentation is illustrated in Fig. 2.

3.3 Semi-Supervised Dense 3D Visual Grounding

Similar to both semi-supervised 3D semantic and 3D instance segmentation, our proposed Bayesian self-training formulation can easily be extended to include natural language prompts as inputs and thus tackle semi-supervised dense 3D visual grounding, i.e., referral-based 3D instance segmentation.

A common technique when tackling dense 3D visual grounding is to employ a grounding-by-selection strategy which can be summarised in two steps [6,48,62]:

- 10 O. Unal et al.
- 1. Visual: First, a 3D backbone computes object candidates O from the scene.
- 2. Verbo-visual: Then, a verbo-visual fusion module ω_{sel} selects the correct instance candidate from predicted instances I based on the natural language description W, i.e., it predicts the index:

$$i^* = \operatorname{argmax} G$$
, with $G = \omega_{sel}(O, W) \in \mathbb{R}^{|I|}$, (9)

To extend our Bayesian self-training approach to this task, we follow the above structure and individually adapt each of the two steps, with the adaptation of the visual step simply following Sec. 3.2.

For the verbal side of 3D visual grounding, i.e. referral-based instance selection, we observe that the pseudo-label generation can be reduced to the same formulation as the pseudo-label generation for 3D instance segmentation. Utilizing the results from the heuristic *n*-partite matching (Alg. 1 L9), vector G can be reordered to match across stochastic forward passes. Again, having matched predicted probability vectors, the method presented in Sec. 3.1 can be trivially applied (Alg. 1 L11-15).

4 Experiments

We evaluate our method on all three semi-supervised 3D semantic perception tasks: 3D semantic segmentation, 3D instance segmentation, and dense 3D visual grounding. The dataset and implementation details regarding each task can be found in the supplementary materials.

4.1 3D Semantic Segmentation

The comparison to the state of the art for semi-supervised 3D semantic segmentation is shown in Tab. 1. As seen, our Bayesian self-training strategy outperforms existing work in both SemanticKITTI [3] and ScribbleKITTI [47] on splits of 10%, 20% and 50%, while performing on par for 1%. Specifically, compared to the baseline Cylinder3D [64] model denoted as *Sup-only*, our strategy allows significant improvements ranging from 4.4% to 10.9% mIoU. In Fig. 3 (a-c) we showcase qualitative results on SemanticKITTI 10%.

4.2 3D Instance Segmentation

For semi-supervised 3D instance segmentation, we report results for ScanNet [11] in Tab. 2 and for S3DIS [2] in Tab. 3. As seen, we observe a similar behavior to semi-supervised 3D semantic segmentation, where at higher thresholds of 10%, 20% for ScanNet and 20%, 50% for S3DIS, our simple yet effective Bayesian self-training strategy significantly outperforms state-of-the-art methods that require highly engineered solutions. In contrast, at lower thresholds, we see a significant gap in performance (especially for ScanNet). Still, compared to the baseline *Suponly*, we see significant gains across the board for all semi-splits. Additionally in Fig. 3 (d-f), we show qualitative results on the ScanNet *val*-set.

Table 1: Comparison of state-of-the-art methods for semi-supervised LiDAR semantic segmentation (mIoU [%]). For fair comparison, all methods use the same backbone given by *Sup-only*, which at 100% labels achieves 64.3% mIoU on SemanticKITTI and 57.0% mIoU on ScribbleKITTI.

	Labeled	1%	10%	20%	50%
~	Sup-only [64]	45.4	56.1	57.8	58.7
	DiAL [46]	45.4	57.1	59.2	60.0
E	CBST [66]	48.8	58.3	59.4	59.7
IX	CPS [8]	46.7	58.7	59.6	60.5
lich	LaserMix [29]	50.6	60.0	61.9	62.3
ant	IGNet [45]	49.0	61.3	63.1	64.8
em	Ours	49.8	61.7	63.7	64.1
Ň	Δ Sup-only	+4.4	+5.6	+5.9	+5.4
	Sup-only [64]	39.2	48.0	52.1	53.8
[4]	DiAL [46]	41.0	50.1	52.8	53.9
II	CBST [66]	41.5	50.6	53.3	54.5
Ē	CPS [8]	41.4	51.8	53.9	54.8
eK	LaserMix [29]	44.2	53.7	55.1	56.8
ldc	IGNet [45]	44.4	57.7	59.6	60.8
cril	Ours	43.9	58.9	60.1	61.2
Ň	Δ Sup-only	+4.7	+10.9	+8.0	+7.4

Table 2: Comparison to state-of-the-art methods for semi-supervised 3D instance segmentation on the ScanNet Limited Reconstruction *valid*-set. For a fair comparison, all methods apart from that marked with * use the same backbone [20,53] as *Sup-only*.

Labeled		1%			5%			10%			20%	
Metric	mAP	AP50	AP25	mAP	AP50	AP25	mAP	AP50	AP25	mAP	AP50	AP25
Sup-only [10]	5.1	9.8	17.6	18.2	32.0	47.0	26.7	42.8	58.9	29.3	47.9	63.0
PointCont. [56]	7.2	12.5	20.3	19.4	35.4	48.5	27.0	43.9	59.5	30.2	49.5	63.6
CSC [20]	7.1	13.0	21.2	20.9	36.7	50.6	27.3	45.0	60.2	30.6	50.3	64.1
TWIST [10]	9.6	17.1	26.2	27.0	44.1	56.2	30.6	49.7	63.0	32.8	52.9	66.8
WS3D* [33]	-	32.5	-	-	45.6	-	-	49.2	-	-	51.3	-
Ours	7.2	14.2	23.0	24.2	39.3	51.4	32.7	51.9	65.5	37.4	56.1	68.9
Δ Sup-only	+2.1	+4.4	+5.4	+6.0	+7.3	+4.4	+6.0	+9.1	+6.6	+8.1	+8.2	+5.9

4.3 Dense 3D Visual Grounding

For semi-supervised 3D visual grounding, we report results on the ScanRefer [6] dataset. As seen in Tab. 4 (Vi), our self-training approach shows improvements across the board, with significant gains on the "unique" category.

While natural language prompts still require human effort to generate, they are considerably cheaper and easier to acquire than per-point high-quality panoptic labels. In Tab. 4 (ViVe), we show that the inclusion of such prompts in the unlabeled dataset further benefits the model's accuracy. As seen, we observe even larger gains over the baseline method with the overall accuracy (OA) showing +4.7% improvement with a 50% IoU threshold.

Furthermore, in Fig. 3 (g) we qualitatively compare our method to the baseline *Sup-only* on ScanRefer.



Fig. 3: Qualitative results from the SemanticKITTI (10%) *val*-set, comparing a) the ground truth, b) *Sup-only* and c) ours; on ScanNet (10%), comparing d) the ground-truth instance masks, e) *Sup-only* and f) ours; and finally g) ScanRefer (10%).

Table 3: Comparison to state-of-the-art methods for semi-supervised 3D instance segmentation on S3DIS Area 5 via AP50 [%]. For a fair comparison, all methods use the same backbone Sparse U-Net [20, 53] as *Sup-only*.

Labeled	5%	10%	20%	50%
Sup-only $[10]$	30.4	36.8	41.2	46.5
PointContrast [56]	33.6	38.7	43.1	48.9
CSC [20]	34.2	41.0	44.7	50.4
TWIST [10]	37.1	45.6	48.4	-
Ours	35.0	42.9	51.1	56.6
Δ Sup-only	+4.6	+6.1	+9.9	+10.1

4.4 Ablation Studies

We conduct ablation studies on 10% splits of SemanticKITTI [3], ScanNet [11] and ScanRefer [6] for semi-supervised 3D semantic segmentation, instance segmentation and visual grounding respectively.

Entropy Thresholding: In Tab. 5, we show an analysis of the effectiveness of our proposed uncertainty thresholding via Shannon entropy, comparing it to the baseline unanimous voting results as well as established thresholding techniques of softmax thresholding [55] and class-balanced thresholding [66]. Alongside the performance for each individual task, we also report the accuracy of the generated pseudo-labels. For instance segmentation, a pseudo-mask is considered correct if its intersection with the associated ground-truth mask is identical to the former. For visual grounding, we reduce the accuracy threshold to 0.5 following the benchmark metric used. As seen, our uncertainty-based filtering improves pseudo-label quality across the board. We also showcase the performance implications of uncertainty filtering, where it can be observed that entropy thresholding outperforms both unanimous voting and current thresholding strategies across all tasks. We further provide a qualitative analysis of the uncertainty met-

Table 4: Results for semi-supervised 3D dense visual grounding while assuming (Vi) that the model is given no natural language descriptions for the unlabeled split and (ViVe) when the model has access to natural language descriptions for the unlabeled split. Reported values are given as Acc@25/50 for the 10% split.

	Category	Unique	Multiple	Overall
	Sup-only [48]	$51.9 \ / \ 43.6$	$19.3 \ / \ 14.1$	$25.4 \ / \ 19.6$
. .	Ours	58.8 / 48.8	22.0 / 18.2	28.8 / 23.9
	Δ Sup-only	+6.9 / +5.2	+2.7 / +4.1	+3.4 / +4.3
Ve	Ours	$61.0 \ / \ 51.2$	$23.4 \ / \ 18.1$	$30.4 \ / \ 24.3$
21	Δ Sup-only	+9.1 / +7.6	+4.1 / +4.0	+5.0 / +4.7

Table 5: Comparison of the pseudo-label quality of current thresholding strategies and their effects on performance on SemanticKITTI 10%, ScanNet 10%, ScanRefer 10% (with visual and verbal) for each respective task. Pseudo-label thresholding is applied after unanimous voting for a fair comparison. The letters indicate each pseudolabel type: S- for semantic, I- for instance, and G- for grounding.

	Sema	antic		Instan	ce	V	isual (Ground	ing
Method	S-Acc	mIoU	S-Acc	I-Acc	AP@50	$S ext{-Acc}$	I-Acc	$G\operatorname{-Acc}$	OA@50
Sup-only [48]	-	56.1	-	-	42.8	-	-	-	19.6
Unanimous	96.4	60.0	79.3	75.3	47.0	72.9	69.9	61.5	16.9
Naïve [5, 55, 65]	98.3	60.8	82.9	77.7	50.2	82.0	73.0	68.6	21.9
Class-balanced [66]	96.6	61.1	81.2	70.1	44.4	80.9	69.0	60.6	16.2
Entropy-based	98.6	61.7	84.8	77.8	51.9	83.8	75.2	75.6	24.3
Δ Unanimous	+2.2	+1.7	+5.5	+2.5	+4.9	+5.8	+2.3	+5.7	+7.4



Fig. 4: Qualitative analysis of the uncertainty estimation. We compare softmax confidence to our MC-derived Shannon entropy.

ric in Fig. 4 to illustrate the better correlation of our MC-derived uncertainty with pseudo-label accuracy compared to softmax based uncertainty.

Labeling Percentage (p_{τ}) : We compare various set labeling percentages p_{τ} across all tasks in Tab. 6. While we obtain higher accuracy pseudo-labels at 50%, it can be seen that the reduced supervision negatively impacts performance.

5 Limitations and Discussion

The main advantage of the proposed Bayesian self-training is its simplicity and ease of implementation compared to existing SOTA approaches that utilize highly engineered and often costly modules. However, the method has a limitation. As observed from the results, the performance is comparably inferior to

14 O. Unal et al.

Table 6: Ablation study analyzing the effect of the hyperparameter p_{τ} for each task. We vary the labeling percentage and show the performance on SemanticKITTI 10%, ScanNet 10%, and ScanRefer 10% (incl. verbal input) for each respective task.

Threshold	mIoU	AP@50	OA@50
50%	61.4	50.0	22.8
75%	61.7	51.9	24.3
100%	60.0	47.0	16.9

Table 7: Ablation study showing the performance on SemanticKITTI 10%, ScanNet 10% and ScanRefer 10% (incl. verbal input) for varying label aggregation strategies from majority rules to unanimous voting over 9 stochastic forward passes.

Threshold	mIoU	AP@50	OA@50
5/9	60.7	47.9	20.1
7/9	61.0	51.0	23.9
9/9	61.7	51.9	24.3

SOTA semi-supervised approaches for highly sparse training splits. This behavior stems from the reliance on unanimous voting (Eq. 6). With low amounts of labeled data, the initial naïve model performance suffers, causing a reduction in the number of unanimously voted points. This effect is further pronounced when tackling instance segmentation, where fused instance masks may affect multiple predictions. In Tab. 7, we show that relaxation does not aid the performance as it further introduces incorrect predictions into the pseudo-label set.

While semi-supervised training methods that enhance the robustness of the pseudo-labels either via a teacher model (DiAL [44, 46]), mix-augmentations (LaserMix [29]) or multi-stage refinement (TWIST [10]) are presented as competing to ours, we argue that as a solution, they can in fact be utilized within our framework in place of the 3D model to improve the quality of the initial predictions. This would allow a larger set of unanimously voted samples and would thus re-enable the benefits of our proposed method on highly sparse splits.

6 Conclusion

In this work, we introduce a simple yet effective Bayesian self-training approach for semi-supervised 3D semantic segmentation by utilizing uncertainty estimation strategies from the field of active learning. Specifically, through multiple stochastic forward passes, we generate an initial pseudo-label set and then filter it based on estimated predictive entropy. We further extend our semi-supervised 3D method to newly include both instance segmentation and referral-based instance segmentation via an efficient, heuristic *n*-partite matching algorithm. We conduct extensive experiments on multiple datasets, covering both indoor and outdoor domains, showing that our method performs on par if not better than current domain- and task-specific methods.

Acknowledgments: This work is funded by Toyota Motor Europe via the research project TRACE-Zürich.

References

- Achlioptas, P., Abdelreheem, A., Xia, F., Elhoseiny, M., Guibas, L.: ReferIt3D: Neural listeners for fine-grained 3D object identification in real-world scenes. In: ECCV (2020)
- Armeni, I., Sener, O., Zamir, A.R., Jiang, H., Brilakis, I., Fischer, M., Savarese, S.: 3D semantic parsing of large-scale indoor spaces. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR) (2016)
- 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)
- 4. Cai, D., Zhao, L., Zhang, J., Sheng, L., Xu, D.: 3DJCG: A unified framework for joint dense captioning and visual grounding on 3D point clouds. In: CVPR (2022)
- Cascante-Bonilla, P., Tan, F., Qi, Y., Ordonez, V.: Curriculum labeling: Revisiting pseudo-labeling for semi-supervised learning. In: Proceedings of the AAAI conference on artificial intelligence. vol. 35, pp. 6912–6920 (2021)
- Chen, D.Z., Chang, A.X., Nießner, M.: ScanRefer: 3D object localization in RGB-D scans using natural language. In: ECCV (2020)
- Chen, J., Luo, W., Wei, X., Ma, L., Zhang, W.: HAM: Hierarchical attention model with high performance for 3D visual grounding. arXiv preprint arXiv:2210.12513 (2022)
- 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 (CVPR) (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 (2021)
- Chu, R., Ye, X., Liu, Z., Tan, X., Qi, X., Fu, C.W., Jia, J.: TWIST: Two-way interlabel self-training for semi-supervised 3D instance segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
- Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scan-Net: Richly-annotated 3D reconstructions of indoor scenes. In: ICCV (2017)
- Dai, W., Li, X., Cheng, K.T.: Semi-supervised deep regression with uncertainty consistency and variational model ensembling via bayesian neural networks. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 7304– 7313 (2023)
- Deng, S., Dong, Q., Liu, B., Hu, Z.: Superpoint-guided semi-supervised semantic segmentation of 3D point clouds. In: International Conference on Robotics and Automation (ICRA) (2022)
- Feng, M., Li, Z., Li, Q., Zhang, L., Zhang, X., Zhu, G., Zhang, H., Wang, Y., Mian, A.: Free-form description guided 3D visual graph network for object grounding in point cloud. In: ICCV (2021)
- Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In: International Conference on Machine Learning (2016)
- Gal, Y., Islam, R., Ghahramani, Z.: Deep Bayesian active learning with image data. In: International Conference on Machine Learning (2017)

- 16 O. Unal et al.
- Gao, H.a., Tian, B., Li, P., Zhao, H., Zhou, G.: DQS3D: Densely-matched quantization-aware semi-supervised 3D detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
- Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: International Conference on Machine Learning (2017)
- He, R., Yang, J., Qi, X.: Re-distributing biased pseudo labels for semi-supervised semantic segmentation: A baseline investigation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021)
- 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 (CVPR) (2021)
- Houlsby, N., Huszár, F., Ghahramani, Z., Lengyel, M.: Bayesian active learning for classification and preference learning. arXiv preprint arXiv:1112.5745 (2011)
- 22. Huang, P.H., Lee, H.H., Chen, H.T., Liu, T.L.: Text-guided graph neural networks for referring 3D instance segmentation. In: AAAI (2021)
- 23. Huang, S., Chen, Y., Jia, J., Wang, L.: Multi-view transformer for 3D visual grounding. In: CVPR (2022)
- Jain, A., Gkanatsios, N., Mediratta, I., Fragkiadaki, K.: Bottom up top down detection transformers for language grounding in images and point clouds. In: ECCV (2022)
- 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 (ICCV) (2021)
- Kendall, A., Badrinarayanan, V., Cipolla, R.: Bayesian SegNet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. In: Proceedings of the British Machine Vision Conference (BMVC) (2017)
- Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- Kong, C., Lin, D., Bansal, M., Urtasun, R., Fidler, S.: What are you talking about? text-to-image coreference. In: ICCV (2014)
- 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 (CVPR) (2023)
- Lai, X., Yuan, Y., Chu, R., Chen, Y., Hu, H., Jia, J.: Mask-attention-free transformer for 3D instance segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
- Lee, M.S., Yang, S.W., Han, S.W.: Gaia: Graphical information gain based attention network for weakly supervised point cloud semantic segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 582–591 (2023)
- Lei, J., Deng, C., Schmeckpeper, K., Guibas, L., Daniilidis, K.: EFEM: Equivariant neural field expectation maximization for 3D object segmentation without scene supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
- Liu, K., Zhao, Y., Nie, Q., Gao, Z., Chen, B.M.: Weakly supervised 3D scene segmentation with region-level boundary awareness and instance discrimination. In: European Conference on Computer Vision (ECCV) (2022)

- 34. Luo, J., Fu, J., Kong, X., Gao, C., Ren, H., Shen, H., Xia, H., Liu, S.: 3D-SPS: Single-stage 3D visual grounding via referred point progressive selection. In: CVPR (2022)
- 35. Meng, Q., Wang, W., Zhou, T., Shen, J., Jia, Y., Van Gool, L.: Towards a weakly supervised framework for 3D point cloud object detection and annotation. IEEE Transactions on Pattern Analysis and Machine Intelligence 44(8), 4454–4468 (2021)
- Prabhudesai, M., Tung, H.Y.F., Javed, S.A., Sieb, M., Harley, A.W., Fragkiadaki, K.: Embodied language grounding with 3D visual feature representations. In: CVPR (2020)
- Qin, Z., Wang, J., Lu, Y.: Weakly supervised 3D object detection from point clouds. In: Proceedings of the ACM International Conference on Multimedia (2020)
- Ren, Z., Misra, I., Schwing, A.G., Girdhar, R.: 3D spatial recognition without spatially labeled 3D. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
- Roh, J., Desingh, K., Farhadi, A., Fox, D.: LanguageRefer: Spatial-language model for 3D visual grounding. In: CoRL (2022)
- 40. Shannon, C.E.: A mathematical theory of communication. The Bell system technical journal **27**(3) (1948)
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. The Journal of Machine Learning Research 15(1), 1929–1958 (2014)
- Tang, Y.S., Lee, G.H.: Transferable semi-supervised 3D object detection from rgbd data. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2019)
- Tao, A., Duan, Y., Wei, Y., Lu, J., Zhou, J.: SegGroup: Seg-level supervision for 3D instance and semantic segmentation. IEEE Transactions on Image Processing 31, 4952–4965 (2022)
- 44. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. Advances in Neural Information Processing Systems (2017)
- Unal, O., Dai, D., Hoyer, L., Can, Y.B., Van Gool, L.: 2d feature distillation for weakly- and semi-supervised 3d semantic segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 7336–7345 (January 2024)
- 46. Unal, O., Dai, D., Unal, A.T., Van Gool, L.: Discwise active learning for LiDAR semantic segmentation. IEEE Robotics and Automation Letters (2023)
- 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 (CVPR) (2022)
- Unal, O., Sakaridis, C., Saha, S., Gool, L.V.: Four ways to improve verbo-visual fusion for dense 3d visual grounding (2024), https://arxiv.org/abs/2309.04561
- Wang, G., Li, W., Aertsen, M., Deprest, J., Ourselin, S., Vercauteren, T.: Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks. Neurocomputing 338, 34–45 (2019)
- Wang, H., Cong, Y., Litany, O., Gao, Y., Guibas, L.J.: 3DIoUMatch: Leveraging IoU prediction for semi-supervised 3D object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)

- 18 O. Unal et al.
- Wang, J., Gang, H., Ancha, S., Chen, Y.T., Held, D.: Semi-supervised 3D object detection via temporal graph neural networks. In: International Conference on 3D Vision (3DV) (2021)
- Wei, J., Lin, G., Yap, K.H., Hung, T.Y., Xie, L.: Multi-path region mining for weakly supervised 3D semantic segmentation on point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- 53. Wu, Y., Shi, M., Du, S., Lu, H., Cao, Z., Zhong, W.: 3d instances as 1d kernels. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIX. pp. 235–252. Springer (2022)
- 54. Xia, Y., Liu, F., Yang, D., Cai, J., Yu, L., Zhu, Z., Xu, D., Yuille, A., Roth, H.: 3D semi-supervised learning with uncertainty-aware multi-view co-training. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (March 2020)
- Xie, Q., Luong, M.T., Hovy, E., Le, Q.V.: Self-training with noisy student improves imagenet classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10687–10698 (2020)
- 56. Xie, S., Gu, J., Guo, D., Qi, C.R., Guibas, L., Litany, O.: Pointcontrast: Unsupervised pre-training for 3D point cloud understanding. In: ECCV (2020)
- 57. Xu, X., Lee, G.H.: Weakly supervised semantic point cloud segmentation: Towards 10x fewer labels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- Yu, L., Wang, S., Li, X., Fu, C.W., Heng, P.A.: Uncertainty-aware self-ensembling model for semi-supervised 3D left atrium segmentation. In: MICCAI (2019)
- 59. Yuan, Z., Yan, X., Liao, Y., Zhang, R., Wang, S., Li, Z., Cui, S.: InstanceRefer: Cooperative holistic understanding for visual grounding on point clouds through instance multi-level contextual referring. In: ICCV (2021)
- 60. Zhang, F., Fang, J., Wah, B., Torr, P.: Deep fusionnet for point cloud semantic segmentation. In: ECCV (2020)
- Zhang, Y., Li, Z., Xie, Y., Qu, Y., Li, C., Mei, T.: Weakly supervised semantic segmentation for large-scale point cloud. In: Proceedings of the AAAI Conference on Artificial Intelligence (2021)
- 62. Zhao, L., Cai, D., Sheng, L., Xu, D.: 3DVG-Transformer: Relation modeling for visual grounding on point clouds. In: ICCV (2021)
- Zhao, N., Chua, T.S., Lee, G.H.: SESS: Self-ensembling semi-supervised 3D object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- 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 (CVPR) (2021)
- Zoph, B., Ghiasi, G., Lin, T.Y., Cui, Y., Liu, H., Cubuk, E.D., Le, Q.V.: Rethinking pre-training and self-training. arXiv preprint arXiv:2006.06882 (2020)
- Zou, Y., Yu, Z., Vijaya Kumar, B., Wang, J.: Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In: The European Conference on Computer Vision (ECCV) (2018)