

SQN: Weakly-Supervised Semantic Segmentation of Large-Scale 3D Point Clouds

Qingyong Hu¹, Bo Yang^{2*}, Guangchi Fang³, Yulan Guo³,
Aleš Leonardis⁴, Niki Trigoni¹, and Andrew Markham¹

¹ University of Oxford

² vLAR Group, The Hong Kong Polytechnic University

³ Sun Yat-sen University

⁴ Huawei Noah's Ark Lab

<https://github.com/QingyongHu/SQN>

qingyong.hu@cs.ox.ac.uk, bo.yang@polyu.edu.hk

Abstract. Labelling point clouds fully is highly time-consuming and costly. As larger point cloud datasets with billions of points become more common, we ask whether the full annotation is even necessary, demonstrating that existing baselines designed under a fully annotated assumption only degrade slightly even when faced with 1% random point annotations. However, beyond this point, *e.g.*, at 0.1% annotations, segmentation accuracy is unacceptably low. We observe that, as point clouds are samples of the 3D world, the distribution of points in a local neighbourhood is relatively homogeneous, exhibiting strong semantic similarity. Motivated by this, we propose a new weak supervision method to implicitly augment highly sparse supervision signals. Extensive experiments demonstrate the proposed Semantic Query Network (SQN) achieves promising performance on seven large-scale open datasets under weak supervision schemes, while requiring only 0.1% randomly annotated points for training, greatly reducing annotation cost and effort.

Keywords: Semantic Query, Weak Supervision, Large-Scale Point Clouds

1 Introduction

Learning precise semantic meanings of large-scale point clouds is crucial for intelligent machines to truly understand complex 3D scenes in the real world. This is a key enabler for autonomous vehicles, augmented reality devices, *etc.*, to quickly interpret the surrounding environment for better navigation and planning.

With the availability of large amounts of labeled 3D data for fully-supervised learning, the task of 3D semantic segmentation has made significant progress in the past four years. Following the seminal works PointNet [46] and SparseConv [16], a series of sophisticated neural architectures [47,34,11,24,66,38,103,10] have been proposed in the literature, greatly improving the accuracy and efficiency of semantic estimation on raw point clouds. The performance of these

* Corresponding author

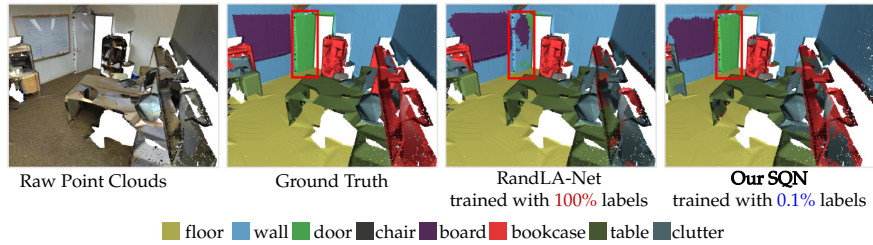


Fig. 1: Qualitative results of RandLA-Net [24] and our SQN on the S3DIS dataset. Trained with only 0.1% annotations, SQN achieves comparable or even better results than the fully-supervised RandLA-Net. Red bounding boxes highlight the superior segmentation accuracy of our SQN.

fully-supervised methods can be further boosted with the aid of self-supervised pre-training representation learning as seen in recent studies [85,36,73,7,96,64]. The success of these approaches primarily relies on densely annotated per-point semantic labels to train the deep neural networks. However, it is extremely costly to fully annotate 3D point clouds due to the unordered, unstructured, and non-uniform data format (*e.g.*, over 1700 person-hours to annotate a typical dataset [3] and around 22.3 minutes for a single indoor scene ($5\text{m} \times 5\text{m} \times 2\text{m}$) [14]). In fact, for very large-scale scenarios *e.g.*, an entire city, it becomes infeasible to manually label every point in practice.

Inspired by the success of weakly-supervised learning techniques in 2D images, a few recent works have started to tackle 3D semantic segmentation using fewer point labels to train neural networks. These methods can be generally divided into five categories: 1) Using 2D image labels for training as in [72,102]; 2) Using fewer 3D labels with gradient approximation/supervision propagation/perturbation consistency [87,94,75,79]; 3) Generating pseudo 3D labels from limited indirect annotations [60,78]; 4) Using superpoint annotations from over-segmentation [60,9,37], and 5) Contrastive pretraining followed by fine-tuning with fewer 3D labels [22,85,97]. Although they achieve encouraging results on multiple datasets, there are a number of limitations still to be resolved.

Firstly, existing approaches usually use custom methods to annotate different amounts of data (*e.g.*, 10%/5%/1% of raw points or superpoints) for training. It is thus unclear what proportion of raw points should be annotated and how, making fair comparison impossible. **Secondly**, to fully utilize the sparse annotations, existing weak-labelling pipelines usually involve multiple stages including careful data augmentation, self-pretraining, fine-tuning, and/or post-processing such as the use of dense CRF [28]. As a consequence, it tends to be more difficult to tune the parameters and deploy them in practical applications, compared with the standard end-to-end training scheme. **Thirdly**, these techniques do not adequately consider the strong local semantic homogeneity of point neighbors

in large-scale point clouds, or do so ineffectively, resulting in the limited, yet valuable, annotations being under-exploited.

Motivated by these issues, we propose a new paradigm for weakly-supervised semantic segmentation on large-scale point clouds, addressing the above shortcomings. In particular, we first explore weak-supervision schemes purely based on existing fully-supervised methods, and then introduce an effective approach to learn accurate semantics given extremely limited point annotations.

To explore weak supervision schemes, we take into account two key questions: 1) *whether, and how, do existing fully-supervised methods deteriorate given different amounts of annotated data for training?* 2) *given fewer and fewer labels, where the weakly supervised regime actually begins?* Fundamentally, by doing so, we aim to explore the limit of current fully-supervised methods. This allows us to draw insights about the use of mature architectures when addressing this challenging task, instead of naïvely borrowing off-the-shelf techniques developed in 2D images [61]. Surprisingly, we find that the accuracy of existing fully-supervised baselines drops only slightly when faced with 1% of random labelled points. However, beyond this point, *e.g.*, 0.1% of the full annotations, the performance degrades rapidly.

With this insight, we propose a novel yet simple **Semantic Query Network**, named **SQN**, for semantic segmentation given as few as 0.1% labeled points for training. Our SQN firstly encodes the entire raw point cloud into a set of hierarchical latent representations via an existing feature extractor, and then takes an arbitrary 3D point position as input to query a subset of latent representations within a local neighborhood. These queried representations are summarized into a compact vector and then fed into a series of multilayer perceptrons (MLPs) to predict the final semantic label. Fundamentally, our SQN explicitly and effectively considers the semantic similarity between neighboring 3D points, allowing the extremely sparse training signals to be back-propagated to a much wider spatial region, thereby achieving superior performance under weak supervision.

Overall, this paper takes a step to bridge the gap between the highly successful fully-supervised methods to the emerging weakly-supervised schemes, in an attempt to reduce the time and labour cost of point-cloud annotation. However, unlike the existing weak-supervision methods, our SQN does not require any self-supervised pretraining, hand-crafted constraints, or complicated post-processing steps, whilst obtaining close to fully-supervised accuracy using as few as 0.1% training labels on multiple large-scale open datasets. Remarkably, for similar accuracy, we find that labelling costs (time) can be reduced up to 98% according to our empirical evaluation in Appendix. Figure 1 shows the qualitative results of our method. Our key contributions are:

- We propose a new weakly supervised method that leverages a point neighbourhood query to fully utilize the sparse training signals.
- We observe that existing fully-supervised methods degrade slowly until 1% point annotations, showing that dense labelling is redundant and unnecessary.
- We demonstrate a significant improvement over baselines in our benchmark, and surpass the state-of-the-art weak-supervision methods by large margins.

2 Related Work

2.1 Learning with Full Supervision

End-to-End Full Supervision. With the availability of densely-annotated point cloud datasets [23,2,18,3,52,68,58], deep learning-based approaches have achieved unprecedented development in semantic segmentation in recent years. The majority of existing approaches follow the standard end-to-end training strategy. They can be roughly divided into three categories according to the representation of 3D point clouds [17]: **1) Voxel-based methods.** They [10,88,16,42] usually voxelize the irregular 3D point clouds into regular cubes [63,11], cylinders [103], or spheres [33]. **2) 2D Projection-based methods.** This pipeline projects the unstructured 3D points into 2D images through multi-view [4,29], bird-eye-view [1], or spherical projections [43,13,80,81,86], and then uses the mature 2D architectures [39,21] for semantic learning. **3) Point-based methods.** These methods [24,46,47,66,34,83,100] directly operate on raw point clouds using shared MLPs. Hybrid representations, such as point-voxel representation [59,38,49], 2D-3D representation [92,26], are also studied.

Self-supervised Pretraining + Full Finetuning. Inspired by the success of self-supervised pre-training representation learning in 2D images [7,20], several recent studies [85,36,73,96,64,53,27,8] apply contrastive techniques for 3D semantic segmentation. These methods usually pretrain the networks on additional 3D source datasets to learn initial per-point representations via self-supervised contrastive losses, after which the networks are carefully finetuned on the target datasets with full labels. This noticeably improves the overall accuracy.

Although these methods have achieved remarkable results on existing datasets, they rely on a large amount of labeled data for training, which is costly and prohibitive in real applications. By contrast, this paper aims to learn semantics from a small fraction of annotations, which is cheaper and more realistic in practice.

2.2 Unsupervised Learning

Saudar and Sievers [53] learn the point semantics by recovering the correct voxel position of every 3D point after the point cloud is randomly shuffled. Sun et al. propose Canonical Capsules [57] to decompose point clouds into object parts and elements via self-canonicalization and auto-encoding. Although they have obtained promising results, they are limited to simple objects and cannot process the complex large-scale point clouds.

2.3 Learning with Weak Supervision

Limited Indirect Annotations. Instead of having point-level semantic annotations, only sub-cloud level or seg-level labels are available. Wei et al. [78] firstly train a classifier with sub-cloud labels, and then generate point-level pseudo labels using class activation mapping technique [101]. Tao et al. [60] present a grouping network to learn semantic and instance segmentation of 3D point

clouds, with the seg-level labels generated by over-segmentation pre-processing. Ren et al. [48] present a multi-task learning framework for both semantic segmentation and 3D object detection with scene-level tags.

Limited Point Annotations. Given a small fraction of points with accurate semantic labels for training, Xu and Lee [87] propose a weakly supervised point cloud segmentation method by approximating gradients and using handcrafted spatial and color smoothness constraints. Zhang et al. [94] explicitly added a perturbed branch, and achieve weakly-supervised learning on 3D point clouds by enforcing predictive consistency. Shi et al. [55] further investigate label-efficient learning by introducing a super-point-based active learning strategy. In addition, self-supervised pre-training methods [54,85,22,96,36,97] are also flexible to fine-tune networks on limited annotations. Our SQN is designed for limited point annotations which we believe has greater potential in practical applications. It does not require any pre-training, post-processing, or active labelling strategies, while achieving similar or even higher performance than the fully-supervised counterpart with only 0.1% randomly annotated points for training.

Fair comparison with 1T1C [37]. In the interests of fair and reproducible comparison, we point out that a few published works claim state-of-the-art results yet make misleading assumptions. Specifically, 1T1C [37] reports impressive results in the paper. However, a deeper investigation of its official GitHub code-base reveals two serious issues:

- **Ground truth label leakage.** 1T1C [37] uses the ground truth instance segments as the super-voxel partition for training on ScanNet⁵. However, given the semantic label of 1 click on ground truth instance segments, the super-voxel semantic labels used by 1T1C are actually dense and full ground truth semantic labels, rather than weak labels.
- **Misleading (over-exaggerated) labeling ratios.** 1T1C calculates its labeling ratio by using the number of labeled instances divided by the total number of raw points, resulting in a fantastically low labeling ratio (*e.g.*, 0.02%)⁶. A fairer method, as used in prior art [87,97,93], is to use the total number of labelled points (*i.e.*, to keep consistency) divided by the total number of points.

For these reasons, 1T1C [37] and its follow-up work PointMatch [84] can be regarded as almost full supervision (all instances are fully annotated) methods on ScanNet. Therefore, our method cannot directly compare with them on ScanNet.

3 Exploring Weak Supervision

As weakly-supervised 3D semantic segmentation is still in its infancy, there is no consensus about what are the sensible formulations of weak training signals, and what approach should be used to sparsely annotate a dataset such that a direct

⁵ <https://github.com/liuzhengzhe/One-Thing-One-Click/issues/13>

⁶ <https://github.com/liuzhengzhe/One-Thing-One-Click/issues/8>

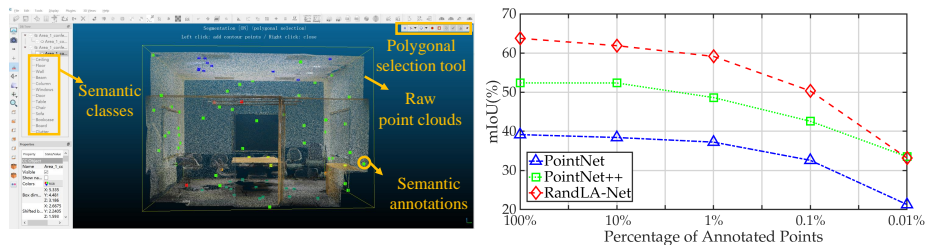


Fig. 2: Left: Illustration of the sparse annotation tool. Right: Degradation of three baselines in the *Area-5* of S3DIS [2] when decreasing proportions of points that are randomly annotated. (Logarithmic scale used in horizontal axis).

comparison is possible. We first explore this, then we investigate how existing fully supervised techniques perform under a weak labelling regime.

Weak Annotation Strategy: The fundamental objective of weakly-supervised segmentation is to obtain accurate estimations with as low as possible annotation cost, in terms of labeller time. However, it is non-trivial to compare the cost of different annotation methods in practice. Existing annotation options include 1) randomly annotating sparse point labels [87,94,93], 2) actively annotating sparse point labels [22,55] or region-wise labels [82], 3) annotating seg-level labels or superpoint labels [60,37,9] and 4) annotating sub-cloud labels [78]. All methods have merits. For the purpose of fair reproducibility, we opt for the random point annotation strategy, considering the practical simplicity of building such an annotation tool.

Annotation Tool: To verify the feasibility of random sparse annotations in practice, we develop a user-friendly labelling pipeline based on the off-the-shelf CloudCompare⁷ software. Specifically, we first import raw 3D point clouds to the software and randomly downsample them to 10%/1%/0.1% of the total points for sparse annotation. Considering the sparsity of the remaining points, we explicitly enlarge the size of selected points and take the original full point clouds as a reference. As illustrated in left part of Figure 2, we then use the standard labelling mode such as polygonal edition for point-wise annotating. (Details and video recordings of our annotation pipeline are supplied in the appendix).

Annotation Cost: With the developed annotation tool, it takes less than 2 minutes to annotate 0.1% of points of a standard room in the S3DIS dataset. For comparison, it requires more than 20 minutes to fully annotate all points for the same room. Note that, the sparse annotation scheme is particularly suitable for large-scale 3D point clouds with billions of points. As detailed in the appendix, it only takes about 18 hours to annotate 0.1% of the urban-scale SensatUrban dataset [23], while annotating all points requires more than 600 person-hours.

Experimental Settings: We choose the well-known S3DIS dataset [2] as the testbed. The Areas {1/2/3/4/6} are selected as the training point clouds,

⁷ <https://www.cloudcompare.org/>

the Area 5 is fully annotated for testing only. With the random sparse annotation strategy, we set up the following four groups of weak signals for training. Specifically, we only annotate the randomly selected 10%/1%/0.1%/0.01% of the 3D points in each room in all training areas.

Using Fully-supervised Methods as Baselines. We select the seminal works PointNet/PointNet++ [46,47] and the recent large-scale-point-cloud friendly RandLA-Net [24] as baselines. These methods are end-to-end trained on the four groups of weakly annotated data without using any additional modules. During training, only the labeled points are used to compute the loss for back-propagation. In total, 12 models (3 models/group \times 4 groups) are trained for evaluation on the full Area 5. Detailed results can be found in Appendix.

Results and Findings. Figure 2 shows the mIoU scores of all models for segmenting the total 13 classes. The results under full supervision (100% annotations for all training data) are included for comparison. It can be seen that:

- The performance of all baselines only decreases marginally (less than 4%) even though the proportion of point annotations drops significantly from 100% to 1%. This clearly shows that the dense annotations are actually unnecessary to obtain a comparable and favorable segmentation accuracy under the simple random annotation strategy.
- The performance of all baselines drops significantly once the annotated points are lower than 0.1%. This critical point indicates that keeping a certain amount of training signals is also essential for weak supervision.

Above all, we may conclude that for segmenting large-scale point clouds which are usually dominated by major classes and have numerous repeatable local patterns, it is desirable to develop weakly-supervised methods which have an excellent trade-off between annotation costs and estimation accuracy. With this motivation, we propose SQN which achieves close to fully-supervised accuracy using only 0.1% labels for training.

4 SQN

4.1 Overview

Given point clouds with sparse annotations, the fundamental challenge for weakly-supervised learning is how to fully utilize the sparse yet valuable training signals to update the network parameters, such that more geometrically meaningful local patterns can be learned. To resolve this, we design a simple SQN which consists of two major components: 1) a point local feature extractor to learn diverse visual patterns; 2) a flexible point feature query network to collect as many as possible relevant semantic features for weakly-supervised training. As shown in Figure 3, our two sub-networks are illustrated by the stacked blocks.

4.2 Point Local Feature Extractor

This component aims to extract local features for all points. As discussed in Section 2.1, there are many excellent backbone networks that are able to extract

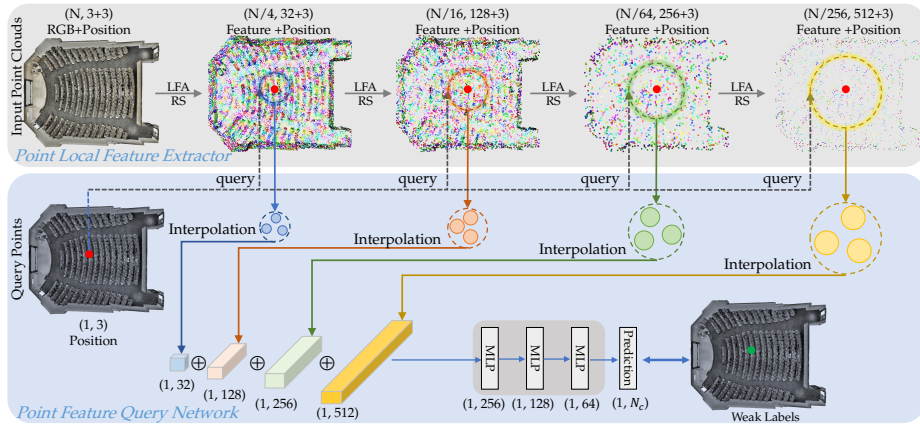


Fig. 3: The pipeline of our SQN at the training stage with weak supervision. We only show one query point for simplicity.

per-point features. In general, these networks stack multiple encoding layers together with downsampling operations to extract hierarchical local features. In this paper, we use the encoder of RandLA-Net [24] as our feature extractor thanks to its efficiency on large-scale point clouds. Note that SQN is not restricted to any particular backbone network *e.g.* as we demonstrate in the Appendix with MinkowskiNet [11].

As shown in the top block of Figure 3, the encoder includes four layers of Local Feature Aggregation (LFA) followed by a Random Sampling (RS) operation. Details refer to RandLA-Net [24]. Given an input point cloud \mathcal{P} with N points, four levels of hierarchical point features are extracted after each encoding layer, *i.e.*, 1) $\frac{N}{4} \times 32$, 2) $\frac{N}{16} \times 128$, 3) $\frac{N}{64} \times 256$, and 4) $\frac{N}{256} \times 512$. To facilitate the subsequent query network, the corresponding point location xyz are always preserved for each hierarchical feature vector.

4.3 Point Feature Query Network

Given the extracted point features, this query network is designed to collect as many relevant features, to be trained using the available sparse signals. In particular, as shown in the bottom block of Figure 3, it takes a specific 3D query point as input and then acquires a set of learned point features relevant to that point. Fundamentally, this is assumed that the query point shares similar semantic information with the collected point features, such that the training signals from the query points can be shared and back-propagated for the relevant points. The network consists of: 1) Searching Spatial Neighbouring Point Features, 2) Interpolating Query Point Features, 3) Inferring Query Point Semantics.

Searching Spatial Neighbouring Point Features. Given a 3D query point p with its location xyz , this module is to simply search the nearest K

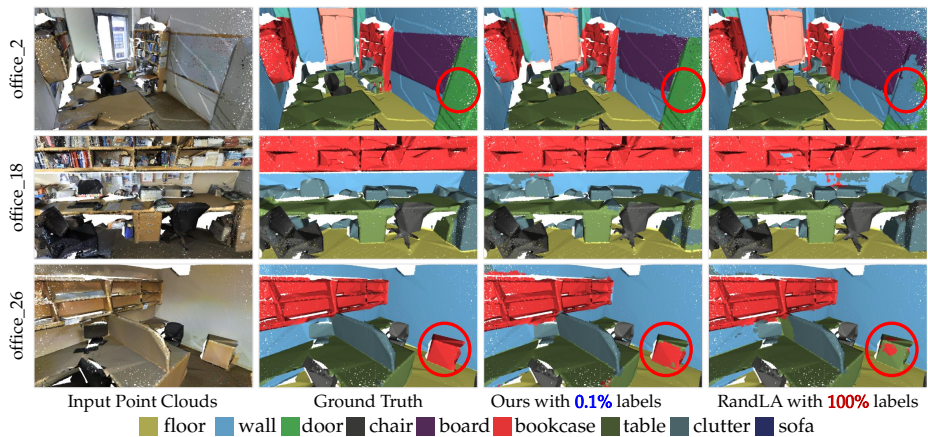


Fig. 4: Qualitative results achieved by our SQN and the fully-supervised RandLA-Net [24] on the *Area-5* of the S3DIS dataset.

points in each of the previous 4-level encoded features, according to the point-wise Euclidean distance. For example, as to the first level of extracted point features, the most relevant K points are selected, acquiring the raw features $\{F_p^1, \dots, F_p^K\}$.

Interpolating Query Point Features. For each level of features, the queried K vectors are compressed into a compact representation for the query point p . For simplicity, we apply the trilinear interpolation method to compute a feature vector for p , according to the Euclidean distance between p and each of K points. Eventually, four hierarchical feature vectors are concatenated together, representing all relevant point features from the entire 3D point cloud.

Inferring Query Point Semantics. After obtaining the unique and representative feature vector for the query point p , we feed it into a series of MLPs, directly inferring the point semantic category.

Overall, given a sparse number of annotated points, we query their neighbouring point features in parallel for training. This allows the valuable training signals to be back-propagated to a much wider spatial context. During testing, all 3D points are fed into the two sub-networks for semantic estimation. In fact, our simple query mechanism allows the network to infer the point semantic category from a significantly larger receptive field.

4.4 Implementation Details

The hyperparameter K is empirically set to 3 for semantic query in our framework and kept consistent for all experiments. Our SQN follows the dataset pre-processing used in RandLA-Net [24], and is trained end-to-end with 0.1% randomly annotated points. All experiments are conducted on a PC with an Intel

		Methods mIoU(%) ceiling floor wall beam column window door table chair sofa bookcase board clutter													
Full supervision	PointNet [46]	41.1	88.8	97.3	69.8	0.1	3.9	46.3	10.8	58.9	52.6	5.9	40.3	26.4	33.2
	PointCNN [34]	57.3	92.3	98.2	79.4	0.0	17.6	22.8	62.1	74.4	80.6	31.7	66.7	62.1	56.7
	SPGraph [31]	58.0	89.4	96.9	78.1	0.0	<u>42.8</u>	48.9	61.6	<u>84.7</u>	75.4	69.8	52.6	2.1	52.2
	SPH3D [33]	59.5	93.3	97.1	81.1	0.0	33.2	45.8	43.8	79.7	86.9	33.2	71.5	54.1	53.7
	PointWeb [99]	60.3	92.0	<u>98.5</u>	79.4	0.0	21.1	59.7	34.8	76.3	88.3	46.9	69.3	64.9	52.5
	RandLA-Net [24]	63.0	92.4	96.7	80.6	0.0	18.3	61.3	43.3	77.2	85.2	71.5	71.0	69.2	52.3
	KPCConv rigid [66]	65.4	92.6	97.3	81.4	0.0	16.5	54.5	69.5	80.2	90.1	66.4	74.6	63.7	58.1
Limited superpoint labels [†]	IT1C (0.02%) [37]	50.1	-	-	-	-	-	-	-	-	-	-	-	-	-
	SSPC-Net (0.01%) [9]	51.5	-	-	-	-	-	-	-	-	-	-	-	-	-
Limited point-wise labels	<i>H</i> Model (10%) [30]	46.3	91.8	97.1	73.8	0.0	5.1	42.0	19.6	67.2	66.7	47.9	19.1	30.6	41.3
	MT (10%) [61]	47.9	92.2	96.8	74.1	0.0	10.4	46.2	17.7	70.7	67.0	50.2	24.4	30.7	42.2
	Xu (10%) [87]	48.0	90.9	97.3	74.8	0.0	8.4	49.3	27.3	71.7	69.0	53.2	16.5	23.3	42.8
	Zhang et al. (1%) [93]	61.8	<u>91.5</u>	<u>96.9</u>	<u>80.6</u>	0.0	18.2	58.1	47.2	75.8	85.7	65.2	68.9	65.0	50.2
	PSD (1%) [94]	63.5	92.3	97.7	80.7	0.0	27.8	56.2	62.5	78.7	84.1	63.1	70.4	58.9	53.2
	<i>H</i> Model (0.2%) [30]	44.3	89.1	97.0	71.5	0.0	3.6	43.2	27.4	63.1	62.1	43.7	14.7	24.0	36.7
	MT (0.2%) [61]	44.4	88.9	96.8	70.1	0.1	3.0	44.3	28.8	63.7	63.6	47.7	15.5	23.0	35.8
	Xu (0.2%) [87]	44.5	90.1	97.1	71.9	0.0	1.9	47.2	29.3	64.0	62.9	42.2	15.9	18.9	37.5
	RandLA-Net (0.1%)	52.9	89.9	95.9	75.3	0.0	7.5	52.4	26.5	62.2	74.5	49.1	60.2	49.3	45.1
	Ours (0.1%)	61.4	91.7	95.6	78.7	0.0	24.2	55.9	63.1	70.5	83.1	60.7	67.8	56.1	50.6

Table 1: Quantitative results of different methods on the *Area-5* of S3DIS dataset. Mean IoU (mIoU, %), and per-class IoU (%) scores are reported. Bold represents the best result in weakly labelled settings and underlined represents the best under fully labelled settings. [†]As mentioned in Sec. 2.3, misleading labeling ratio is reported, and hence a direct comparison is not possible.

Core™ i9-10900X CPU and an NVIDIA RTX Titan GPU. Note that, the proposed SQN framework allows flexible use of different backbone networks such as voxel-based MinkowskiNet [11], please refer to the appendix for more details.

5 Experiments

5.1 Comparison with SOTA Approaches

We first evaluate the performance of our SQN on three commonly-used benchmarks including S3DIS [2], ScanNet [14] and Semantic3D [18]. Following [24], we use the Overall Accuracy (OA) and mean Intersection-over-Union (mIoU) as the main evaluation metrics.

Evaluation on S3DIS. Following [87], we report the results on Area-5 in Table 1. Note that, our SQN is compared with three groups of approaches: 1) Fully-supervised methods including SPGraph [31], KPCConv [66] and RandLA-Net with 100% training labels; 2) Weakly supervised approaches that learn from limited superpoint annotations including IT1C [37] and SSPC-Net [9]; 3) Weakly-supervised methods [87,61,30] that learning from limited annotations. We also list the proportion of annotations used for training.

Considering different backbones and different labelling ratios are used by existing methods, we focus on the comparison of our SQN and the baseline RandLA-Net, which under the same weakly-supervised settings. It can be seen that our SQN outperforms RandLA-Net by nearly 9% under the same 0.1% random sparse annotations. In particular, our SQN is also comparable to the fully-supervised RandLA-Net [24]. Figure 4 shows qualitative comparisons of RandLA-Net and our SQN.

Settings	Methods	mIoU(%)
Full supervision	PointNet++ [47]	33.9
	SPLATNet [56]	39.3
	TangentConv [62]	43.8
	PointCNN [34]	45.8
	PointConv [83]	55.6
	SPH3D-GCN [33]	61.0
	KPConv [66]	68.4
	RandLA-Net [24]	64.5
Weak supervision	MPRM* [78]	41.1
	Zhang <i>et al.</i> (1%) [93]	51.1
	PSD (1%) [94]	54.7
	Ours (0.1%)	56.9

Table 2: Quantitative results on ScanNet (online test set). *MPRM [78] takes sub-cloud labels as supervision signal.

	Methods	<i>Semantic8</i>		<i>Reduced8</i>	
		OA(%)	mIoU(%)	OA(%)	mIoU(%)
Full sup.	SnapNet [4]	91.0	67.4	88.6	59.1
	PointNet++ [47]	85.7	63.1	-	-
	ShellNet [98]	-	-	93.2	69.3
	GACNet [74]	-	-	91.9	70.8
	RGNet [67]	90.6	72.0	94.5	74.7
	SPG [31]	92.9	76.2	94.0	73.2
	KPConv [66]	-	-	92.9	74.6
	ConvPoint [6]	93.4	76.5	-	-
	WreathProdNet [76]	94.6	77.1	-	-
	RandLA-Net [24]	95.0	75.8	94.8	77.4
	Weak sup.	Zhang <i>et al.</i> (1%) [93]	-	-	-
PSD (1%) [94]		-	-	-	75.8
Ours (0.1%)		94.8	72.3	93.7	74.7
Ours (0.01%)		91.9	58.8	90.3	65.6

Table 3: Quantitative results on Semantic3D [18]. The scores are obtained from the recent publications.

Evaluation on ScanNet. We report the quantitative results achieved by different approaches on the hidden test set in Table 2. It can be seen that our SQN achieves higher mIoU scores with only 0.1% training labels, compared with MPRM [78] which is trained with sub-cloud labels, and Zhang *et al.* [93] and PSD [94] trained with 1% annotations. Considering that the actual training settings in the ScanNet Data-Efficient benchmark cannot be verified, hence we do not provide the comparison in this benchmark.

Evaluation on Semantic3D. Table 3 compares our SQN with a number of fully-supervised methods. It can be seen that our SQN trained with 0.1% labels achieves competitive performance with fully-supervised baselines on both *Semantic8* and *Reduced8* subsets. This clearly demonstrates the effectiveness of our semantic query framework, which takes full advantage of the limited annotations. Additionally, we also train our SQN with only 0.01% randomly annotated points, considering the extremely large amount of 3D points scanned. We can see that our SQN trained with 0.01% labels also achieves satisfactory accuracy, though there is space to be improved in the future.

5.2 Evaluation on Large-Scale 3D Benchmarks

To validate the versatility of our SQN, we further evaluate our SQN on four point cloud datasets with different density and quality, including SensatUrban [23], Toronto3D [58], DALES [68], and SemanticKITTI [3]. Note that, all existing weakly supervised approaches are only evaluated on the dataset with dense point clouds, and there are no results reported on these datasets. Therefore, we only compare our approach with existing fully-supervised methods in this section.

As shown in Table 4, the performance of our SQN is on par with the fully-supervised counterpart RandLA-Net on several datasets, whilst the model is only supplied with 0.1% labels for training. In particular, our SQN trained with 0.1% labels even outperforms the fully supervised RandLA-Net on the SensatUrban

Settings	Methods	DALES [68]		SensatUrban [23]			Toronto3D [58]		SemanticKITTI [3]
		OA(%)	mIoU(%)	OA(%)	mAcc (%)	mIoU(%)	OA(%)	mIoU(%)	mIoU(%)
Full supervision	PointNet [46]	-	-	80.8	30.3	23.7	-	-	14.6
	PointNet++ [47]	95.7	68.3	84.3	40.0	32.9	84.9	41.8	20.1
	PointCNN [34]	97.2	58.4	-	-	-	-	-	-
	TangentConv [62]	-	-	77.0	43.7	33.3	-	-	40.9
	ShellNet [98]	96.4	57.4	-	-	-	-	-	-
	DGCNN [77]	-	-	-	-	-	94.2	61.8	-
	SPG [31]	95.5	60.6	85.3	44.4	37.3	-	-	17.4
	SparseConv [16]	-	-	88.7	63.3	42.7	-	-	-
	KPConv [66]	97.8	81.1	93.2	63.8	57.6	95.4	69.1	58.1
	ConvPoint [5]	97.2	67.4	-	-	-	-	-	-
RandLA-Net [24]	97.1	80.0	89.8	69.6	52.7	92.9	77.7	53.9	
Weak supervision	Ours (0.1%)	97.0	72.0	91.0	70.9	54.0	96.7	77.7	50.8
	Ours (0.01%)	95.9	60.4	85.6	49.4	37.2	94.2	68.2	39.1

Table 4: Quantitative results of different approaches on the DALES [68], SensatUrban [23], Toronto3D [58] and SemanticKITTI [3].

dataset. This shows the great potential of our method, especially for extremely large-scale point clouds with billions of points, where the manual annotation is unrealistic and impractical. The detailed results can be found in Appendix.

5.3 Ablation Study

To evaluate the effectiveness of each module in our framework, we conduct the following ablation studies. All ablated networks are trained on Areas{1/2/3/4/6} with 0.1% labels, and tested on the *Area-5* of the S3DIS dataset.

(1) Varying Number of Queried Neighbors.

Intuitively, querying a larger neighborhood is more likely to achieve better results. However, an overly large neighborhood may include points with very different semantics, diminishing overall performance. To investigate the impact of the number of neighboring points used in our semantic query, we conduct experiments by varying the number of neighboring points from 1 to 25. As shown in Figure 5, the overall performance with differing numbers of neighboring points does not change significantly, showing that our simple query mechanism is robust to the size of the neighboring patch. Instead, the mixture of different feature levels plays a more important role (Table 5).

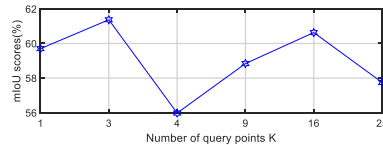


Fig. 5: The results of our SQN with different number of query points on the *Area-5* of the S3DIS dataset.

(2) Variants of Semantic Queries.

The hierarchical point feature query mechanism is the major component of our SQN. To evaluate this component, we perform semantic query at different encoding layers. In particular, we train four additional models, each of which has a different combination of queried neighbouring point features. From Table 5 we can see that the segmentation performance drops significantly if we only collect the relevant point features at a single layer (*e.g.*, the first or the last layer), whilst querying at the last layer can

Model	1st	2nd	3rd	4st	OA(%)	mIoU(%)
A	✓				48.66	22.89
B				✓	75.54	46.02
C	✓	✓			70.76	38.18
D	✓	✓	✓		82.37	54.21
E	✓	✓	✓	✓	86.15	61.41

Table 5: Ablations of different levels of semantic query.

	OA(%)	mIoU(%)
Trial1	86.15	61.41
Trial2	85.63	59.24
Trial3	86.39	60.93
Trial4	86.32	59.40
Trial5	86.40	61.56
Mean	86.25	60.42
STD	0.32	0.93

Table 6: Sensitivity analysis of the proposed SQN on S3DIS dataset (*Area 5*) over 5 runs.

achieve much better results than in the first layer. This is because the points in the last encoding layer are quite sparse but representative, aggregating a large number of neighboring points. Additionally, querying at different encoding layers and combining them is likely to achieve better segmentation results, mainly because it integrates different spatial levels of semantic content and considers more neighboring points.

(3) Varying Annotated Points. To verify the sensitivity of our SQN to different randomly annotated points, we train our models five times with exactly same architectures, *i.e.*, the only change is that different subsets of randomly selected 0.1% of points are labeled. The results are reported in Table 6. It can be seen that there are slight, but not significant, differences between different runs. This indicates that the proposed SQN is robust to the choice of randomly annotated points. We also notice that the major performance change lies in minor categories such as *door*, *sofa*, and *board*, showing that the underrepresented classes are more sensitive to weak annotation. Please refer to appendix for details.

(4) Varying Proportion of Annotated Points. We further examine the performance of SQN with differing amounts of annotated points. As shown in Table 7, the proposed SQN can achieve satisfactory segmentation performance when there are only 0.1% labels available, but the performance drops significantly when there are only 0.01% labeled points available, primarily because the supervision signal is too sparse and limited in this case. It is also interesting to see that our framework achieves slightly better mIoU performance when using 10% labels compared with full supervision. In particular, the performance on minority categories such as *column/window/door* has improved by 2%-5%. This implies that: 1) In a sense, the supervision signal is sufficient in this case; 2) Another way to address the critical issue of imbalanced class distribution may be to use a portion of training data (*i.e.*, weak supervision). This is an interesting direction for further research, and we leave it for future exploration.

(5) Extension to Region-wise Annotated Data. Beyond evaluating on randomly point-wise annotated datasets, we also extend our SQN on the region-wise sparsely labeled S3DIS dataset. Following [82], point clouds are firstly grouped into regions by unsupervised over-segmentation methods [45], and then a sparse

Settings	Methods	mIoU(%)	ceiling	floor	wall	beam	col.	win.	door	chair	table	book.	sofa	board	clutter
100%	SQN	63.73	92.76	96.92	81.84	0.00	25.93	50.53	65.88	79.52	85.31	55.66	72.51	65.78	55.85
10%	SQN	64.67	93.04	97.45	81.55	0.00	28.01	55.77	68.68	80.11	87.67	55.25	72.31	63.91	57.02
1%	SQN	63.65	92.03	96.41	81.32	0.00	21.42	53.71	73.17	77.80	85.95	56.72	69.91	66.57	52.49
0.1%	SQN	61.41	91.72	95.63	78.71	0.00	24.23	55.89	63.14	70.50	83.13	60.67	67.82	56.14	50.63
0.01%	SQN	45.30	89.16	93.49	71.28	0.00	4.14	34.67	41.02	54.88	66.85	25.68	55.37	12.80	39.57

Table 7: Quantitative results achieved by our SQN on *Area-5* of S3DIS under different amounts of labeled points.

	SPVCNN [59]	MinkowskiUnet [11]	SQN (Ours)
Random	49.61	46.15	60.19
Softmax Confidence [71]	51.05	45.45	57.24
Softmax Margin [71]	50.80	44.33	57.94
Softmax Entropy [71]	50.35	49.99	57.98
MC Dropout [15]	50.39	49.94	58.30
ReDAL [82]	50.89	47.88	54.24

Table 8: Quantitative results achieved by different methods on the region-wise labeled S3DIS dataset.

number of regions are manually annotated through various active learning strategies [71,82,15]. As shown in Table 8, our SQN can consistently achieve better results than vanilla SPVCNN [59] and MinkowskiNet [11] under the same supervision signal (10 iterations of active selection), regardless of the active learning strategy used. This is likely because the SparseConv based methods [59,11] usually have larger models and more trainable parameters compared with our point-based lightweight SQN, thus naturally exhibiting a stronger demand and dependence for more supervision signals. On the other hand, this result further validates the effectiveness and superiority of our SQN under weak supervision.

6 Conclusion

In this paper, we propose SQN, a conceptually simple and elegant framework to learn the semantics of large-scale point clouds, with as few as 0.1% supplied labels for training. We first point out the redundancy of dense 3D annotations through extensive experiments, and then propose an effective semantic query framework based on the assumption of semantic similarity of neighboring points in 3D space. The proposed SQN simply follows the concept of wider label propagation, but shows great potential for weakly-supervised semantic segmentation of large-scale point clouds. It would be interesting to extend this method for weakly-supervised instance segmentation, panoptic segmentation, and further integrate it into semantic surface reconstruction [70].

Acknowledgements. This work was partially supported by the National Natural Science Foundation of China (No. 61972435, U20A20185), China Scholarship Council (CSC) scholarship, and Huawei UK AI Fellowship. Qingyong Hu and Bo Yang were partially supported by Shenzhen Science and Technology Innovation Commission (JCYJ20210324120603011).

References

1. Aksoy, E.E., Baci, S., Cavdar, S.: Salsanet: Fast road and vehicle segmentation in LiDAR point clouds for autonomous driving. In: IV. pp. 926–932 (2019) [4](#)
2. Armeni, I., Sax, S., Zamir, A.R., Savarese, S.: Joint 2D-3D-semantic data for indoor scene understanding. In: ICCV (2017) [4](#), [6](#), [10](#), [16](#), [20](#)
3. 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: ICCV. pp. 9297–9307 (2019) [2](#), [4](#), [11](#), [12](#), [16](#), [24](#), [26](#)
4. Boulch, A., Saux, B.L., Audebert, N.: Unstructured point cloud semantic labeling using deep segmentation networks. In: 3DOR. pp. 17–24 (2017) [4](#), [11](#), [21](#), [22](#)
5. Boulch, A.: Generalizing discrete convolutions for unstructured point clouds. In: 3DOR. pp. 71–78 (2019) [12](#), [21](#), [24](#)
6. Boulch, A., Puy, G., Marlet, R.: Fkaconv: Feature-kernel alignment for point cloud convolution. In: ACCV (2020) [11](#), [21](#), [22](#)
7. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML. pp. 1597–1607 (2020) [2](#), [4](#)
8. Chen, Y., Liu, J., Ni, B., Wang, H., Yang, J., Liu, N., Li, T., Tian, Q.: Shape self-correction for unsupervised point cloud understanding. In: ICCV (2021) [4](#)
9. Cheng, M., Hui, L., Xie, J., Yang, J.: Sspc-net: Semi-supervised semantic 3d point cloud segmentation network. arXiv preprint arXiv:2104.07861 (2021) [2](#), [6](#), [10](#)
10. Cheng, R., Razani, R., Taghavi, E., Li, E., Liu, B.: 2-S3Net: Attentive feature fusion with adaptive feature selection for sparse semantic segmentation network. arXiv preprint arXiv:2102.04530 (2021) [1](#), [4](#)
11. Choy, C., Gwak, J., Savarese, S.: 4D spatio-temporal convnets: Minkowski convolutional neural networks. In: CVPR. pp. 3075–3084 (2019) [1](#), [4](#), [8](#), [10](#), [14](#), [17](#)
12. Contreras, J., Denzler, J.: Edge-convolution point net for semantic segmentation of large-scale point clouds. In: IGARSS. pp. 5236–5239 (2019) [21](#)
13. Cortinhal, T., Tzelepis, G., Aksoy, E.E.: SalsaNext: Fast semantic segmentation of LiDAR point clouds for autonomous driving. In: ISVC (2020) [4](#), [24](#)
14. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In: CVPR. pp. 5828–5839 (2017) [2](#), [10](#), [20](#), [21](#)
15. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: ICML (2016) [14](#)
16. Graham, B., Engelcke, M., van der Maaten, L.: 3D semantic segmentation with submanifold sparse convolutional networks. In: CVPR (2018) [1](#), [4](#), [12](#), [21](#), [23](#)
17. Guo, Y., Wang, H., Hu, Q., Liu, H., Liu, L., Bennamoun, M.: Deep learning for 3D point clouds: A survey. IEEE TPAMI (2020) [4](#)
18. Hackel, T., Savinov, N., Ladicky, L., Wegner, J.D., Schindler, K., Pollefeys, M.: Semantic3D.Net: A new large-scale point cloud classification benchmark. ISPRS (2017) [4](#), [10](#), [11](#), [16](#), [21](#), [22](#)
19. Hackel, T., Wegner, J.D., Schindler, K.: Fast semantic segmentation of 3D point clouds with strongly varying density. ISPRS **3**, 177–184 (2016) [21](#)
20. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR. pp. 9729–9738 (2020) [4](#)
21. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016) [4](#)

22. Hou, J., Graham, B., Nießner, M., Xie, S.: Exploring data-efficient 3D scene understanding with contrastive scene contexts. In: CVPR (2021) [2](#), [5](#), [6](#)
23. Hu, Q., Yang, B., Khalid, S., Xiao, W., Trigoni, N., Markham, A.: Towards semantic segmentation of urban-scale 3D point clouds: A dataset, benchmarks and challenges. In: CVPR (2021) [4](#), [6](#), [11](#), [12](#), [15](#), [16](#), [25](#)
24. Hu, Q., Yang, B., Xie, L., Rosa, S., Guo, Y., Wang, Z., Trigoni, N., Markham, A.: RandLA-Net: Efficient semantic segmentation of large-scale point clouds. In: CVPR (2020) [1](#), [2](#), [4](#), [7](#), [8](#), [9](#), [10](#), [11](#), [12](#), [15](#), [20](#), [21](#), [22](#), [23](#), [24](#)
25. Huang, Q., Wang, W., Neumann, U.: Recurrent slice networks for 3D segmentation of point clouds. In: ICCV (2018) [20](#)
26. Jaritz, M., Gu, J., Su, H.: Multi-view pointnet for 3D scene understanding. In: ICCVW. pp. 0–0 (2019) [4](#)
27. 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: ICCV (2021) [4](#)
28. Krähenbühl, P., Koltun, V.: Efficient inference in fully connected CRFs with gaussian edge potentials. In: NeurIPS. pp. 109–117 (2011) [2](#)
29. Kundu, A., Yin, X., Fathi, A., Ross, D., Brewington, B., Funkhouser, T., Pantofaru, C.: Virtual multi-view fusion for 3D semantic segmentation. In: ECCV. pp. 518–535 (2020) [4](#)
30. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. In: ICLR (2017) [10](#)
31. Landrieu, L., Simonovsky, M.: Large-scale point cloud semantic segmentation with superpoint graphs. In: CVPR. pp. 4558–4567 (2018) [10](#), [11](#), [12](#), [20](#), [21](#), [22](#), [23](#), [24](#)
32. Lei, H., Akhtar, N., Mian, A.: SegGCN: Efficient 3D point cloud segmentation with fuzzy spherical kernel. In: CVPR (2020) [21](#)
33. Lei, H., Akhtar, N., Mian, A.: Spherical kernel for efficient graph convolution on 3D point clouds. IEEE TPAMI (2020) [4](#), [10](#), [11](#), [21](#)
34. Li, Y., Bu, R., Sun, M., Wu, W., Di, X., Chen, B.: PointCNN: Convolution on X-transformed points. In: NeurIPS (2018) [1](#), [4](#), [10](#), [11](#), [12](#), [20](#), [21](#), [22](#), [24](#)
35. Li, Y., Ma, L., Zhong, Z., Cao, D., Li, J.: TGNNet: Geometric graph cnn on 3D point cloud segmentation. IEEE TGRS (2019) [23](#)
36. Liu, Y., Yi, L., Zhang, S., Fan, Q., Funkhouser, T., Dong, H.: P4contrast: Contrastive learning with pairs of point-pixel pairs for rgb-d scene understanding. arXiv preprint arXiv:2012.13089 (2020) [2](#), [4](#), [5](#)
37. Liu, Z., Qi, X., Fu, C.W.: One thing one click: A self-training approach for weakly supervised 3d semantic segmentation. In: CVPR. pp. 1726–1736 (2021) [2](#), [5](#), [6](#), [10](#), [15](#)
38. Liu, Z., Tang, H., Lin, Y., Han, S.: Point-voxel cnn for efficient 3D deep learning. In: NeurIPS (2019) [1](#), [4](#)
39. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR. pp. 3431–3440 (2015) [4](#)
40. Ma, L., Li, Y., Li, J., Tan, W., Yu, Y., Chapman, M.A.: Multi-scale point-wise convolutional neural networks for 3D object segmentation from LiDAR point clouds in large-scale environments. IEEE TITS (2019) [23](#)
41. Ma, Y., Guo, Y., Liu, H., Lei, Y., Wen, G.: Global context reasoning for semantic segmentation of 3D point clouds. WACV (2020) [21](#)
42. Meng, H.Y., Gao, L., Lai, Y.K., Manocha, D.: VV-Net: Voxel vae net with group convolutions for point cloud segmentation. In: ICCV (2019) [4](#)
43. Milioto, A., Vizzo, I., Behley, J., Stachniss, C.: Rangenet++: Fast and accurate LiDAR semantic segmentation. In: IROS. pp. 4213–4220 (2019) [4](#), [24](#)

44. Montoya-Zegarra, J.A., Wegner, J.D., Ladický, L., Schindler, K.: Mind the gap: modeling local and global context in (road) networks. In: GCPR (2014) [21](#)
45. Papon, J., Abramov, A., Schoeler, M., Worgotter, F.: Voxel cloud connectivity segmentation-supervoxels for point clouds. In: CVPR (2013) [13](#)
46. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: PointNet: Deep learning on point sets for 3D classification and segmentation. In: CVPR. pp. 652–660 (2017) [1](#), [4](#), [7](#), [10](#), [12](#), [20](#), [23](#), [24](#)
47. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: PointNet++: Deep hierarchical feature learning on point sets in a metric space. In: NeurIPS (2017) [1](#), [4](#), [7](#), [11](#), [12](#), [20](#), [21](#), [22](#), [23](#), [24](#)
48. 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. pp. 13204–13213 (2021) [5](#)
49. Rethage, D., Wald, J., Sturm, J., Navab, N., Tombari, F.: Fully-convolutional point networks for large-scale point clouds. In: ECCV (2018) [4](#)
50. Rosu, R.A., Schütt, P., Quenzel, J., Behnke, S.: LatticeNet: Fast point cloud segmentation using permutohedral lattices. In: RSS (2020) [24](#)
51. Roynard, X., Deschaud, J.E., Goulette, F.: Classification of point cloud for road scene understanding with multiscale voxel deep network. In: PPNIV (2018) [22](#)
52. Roynard, X., Deschaud, J.E., Goulette, F.: Paris-Lille-3D: A large and high-quality ground-truth urban point cloud dataset for automatic segmentation and classification. *IJRR* **37**(6), 545–557 (2018) [4](#)
53. Sauder, J., Sievers, B.: Self-supervised deep learning on point clouds by reconstructing space. In: NeurIPS. pp. 12962–12972 (2019) [4](#)
54. Sharma, C., Kaul, M.: Self-supervised few-shot learning on point clouds. In: NeurIPS (2020) [5](#)
55. Shi, X., Xu, X., Chen, K., Cai, L., Foo, C.S., Jia, K.: Label-efficient point cloud semantic segmentation: An active learning approach. arXiv preprint arXiv:2101.06931 (2021) [5](#), [6](#)
56. Su, H., Jampani, V., Sun, D., Maji, S., Kalogerakis, E., Yang, M.H., Kautz, J.: SPLATNet: sparse lattice networks for point cloud processing. In: CVPR. pp. 2530–2539 (2018) [11](#), [21](#), [24](#)
57. Sun, W., Tagliasacchi, A., Deng, B., Sabour, S., Yazdani, S., Hinton, G., Yi, K.M.: Canonical capsules: Unsupervised capsules in canonical pose. arXiv preprint arXiv:2012.04718 (2020) [4](#)
58. Tan, W., Qin, N., Ma, L., Li, Y., Du, J., Cai, G., Yang, K., Li, J.: Toronto-3D: A large-scale mobile LiDAR dataset for semantic segmentation of urban roadways. In: CVPRW. pp. 202–203 (2020) [4](#), [11](#), [12](#), [16](#), [23](#)
59. Tang, H., Liu, Z., Zhao, S., Lin, Y., Lin, J., Wang, H., Han, S.: Searching efficient 3D architectures with sparse point-voxel convolution. In: ECCV. pp. 685–702 (2020) [4](#), [14](#), [17](#)
60. Tao, A., Duan, Y., Wei, Y., Lu, J., Zhou, J.: Seggroup: Seg-level supervision for 3D instance and semantic segmentation. arXiv preprint arXiv:2012.10217 (2020) [2](#), [4](#), [6](#)
61. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: NeurIPS. pp. 1195–1204 (2017) [3](#), [10](#)
62. Tatarchenko, M., Park, J., Koltun, V., Zhou, Q.Y.: Tangent convolutions for dense prediction in 3D. In: CVPR. pp. 3887–3896 (2018) [11](#), [12](#), [21](#), [23](#), [24](#)
63. Tchapmi, L., Choy, C., Armeni, I., Gwak, J., Savarese, S.: Segcloud: Semantic segmentation of 3D point clouds. In: 3DV. pp. 537–547 (2017) [4](#), [22](#)

64. Thabet, A., Alwassel, H., Ghanem, B.: Self-supervised learning of local features in 3D point clouds. In: CVPRW. pp. 938–939 (2020) [2](#), [4](#)
65. Thomas, H., Goulette, F., Deschaud, J.E., Marcotegui, B., LeGall, Y.: Semantic classification of 3D point clouds with multiscale spherical neighborhoods. In: 3DV. pp. 390–398 (2018) [22](#)
66. Thomas, H., Qi, C.R., Deschaud, J.E., Marcotegui, B., Goulette, F., Guibas, L.J.: KPConv: Flexible and deformable convolution for point clouds. In: ICCV. pp. 6411–6420 (2019) [1](#), [4](#), [10](#), [11](#), [12](#), [15](#), [20](#), [21](#), [22](#), [23](#), [24](#)
67. Truong, G., Gilani, S.Z., Islam, S.M.S., Suter, D.: Fast point cloud registration using semantic segmentation. In: DICTA. pp. 1–8 (2019) [11](#), [21](#), [22](#)
68. Varney, N., Asari, V.K., Graehling, Q.: DALES: A large-scale aerial LiDAR data set for semantic segmentation. In: CVPRW. pp. 186–187 (2020) [4](#), [11](#), [12](#), [16](#), [17](#)
69. Varney, N., Asari, V.K., Graehling, Q.: Pyramid point: A multi-level focusing network for revisiting feature layers. arXiv preprint arXiv:2011.08692 (2020) [24](#)
70. Wang, B., Yu, Z., Yang, B., Qin, J., Breckon, T., Shao, L., Trigoni, N., Markham, A.: RangUDF: Semantic surface reconstruction from 3d point clouds. arXiv preprint arXiv:2204.09138 (2022) [14](#)
71. Wang, D., Shang, Y.: A new active labeling method for deep learning. In: IJCNN. pp. 112–119. IEEE (2014) [14](#)
72. Wang, H., Rong, X., Yang, L., Wang, S., Tian, Y.: Towards weakly supervised semantic segmentation in 3D graph-structured point clouds of wild scenes. In: BMVC. p. 284 (2019) [2](#)
73. Wang, H., Liu, Q., Yue, X., Lasenby, J., Kusner, M.J.: Pre-training by completing point clouds. arXiv preprint arXiv:2010.01089 (2020) [2](#), [4](#)
74. Wang, L., Huang, Y., Hou, Y., Zhang, S., Shan, J.: Graph attention convolution for point cloud semantic segmentation. In: CVPR (2019) [11](#), [22](#)
75. Wang, P., Yao, W.: A new weakly supervised approach for als point cloud semantic segmentation. arXiv preprint arXiv:2110.01462 (2021) [2](#)
76. Wang, R., Albooyeh, M., Ravanbakhsh, S.: Equivariant maps for hierarchical structures. arXiv preprint arXiv:2006.03627 (2020) [11](#), [21](#)
77. Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph cnn for learning on point clouds. ACM TOG **38**(5), 1–12 (2019) [12](#), [23](#)
78. 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: CVPR. pp. 4384–4393 (2020) [2](#), [4](#), [6](#), [11](#)
79. Wei, J., Lin, G., Yap, K.H., Liu, F., Hung, T.Y.: Dense supervision propagation for weakly supervised semantic segmentation on 3d point clouds. arXiv preprint arXiv:2107.11267 (2021) [2](#)
80. 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: ICRA. pp. 1887–1893 (2018) [4](#), [24](#)
81. Wu, B., Zhou, X., Zhao, S., Yue, X., Keutzer, K.: SqueezeSegV2: Improved model structure and unsupervised domain adaptation for road-object segmentation from a LiDAR point cloud. In: ICRA. pp. 4376–4382 (2019) [4](#), [24](#)
82. Wu, T.H., Liu, Y.C., Huang, Y.K., Lee, H.Y., Su, H.T., Huang, P.C., Hsu, W.H.: Redal: Region-based and diversity-aware active learning for point cloud semantic segmentation. In: ICCV. pp. 15510–15519 (2021) [6](#), [13](#), [14](#)
83. Wu, W., Qi, Z., Fuxin, L.: PointConv: Deep convolutional networks on 3D point clouds. In: CVPR. pp. 9621–9630 (2018) [4](#), [11](#), [21](#)

84. Wu, Y., Yan, Z., Cai, S., Li, G., Yu, Y., Han, X., Cui, S.: Pointmatch: A consistency training framework for weakly supervised semantic segmentation of 3d point clouds. arXiv preprint arXiv:2202.10705 (2022) [5](#)
85. Xie, S., Gu, J., Guo, D., Qi, C.R., Guibas, L., Litany, O.: PointContrast: Unsupervised pre-training for 3D point cloud understanding. In: ECCV. pp. 574–591 (2020) [2](#), [4](#), [5](#)
86. Xu, C., Wu, B., Wang, Z., Zhan, W., Vajda, P., Keutzer, K., Tomizuka, M.: SqueezeSegV3: Spatially-adaptive convolution for efficient point-cloud segmentation. In: ECCV. pp. 1–19 (2020) [4](#), [24](#)
87. Xu, X., Lee, G.H.: Weakly supervised semantic point cloud segmentation: Towards 10x fewer labels. In: CVPR. pp. 13706–13715 (2020) [2](#), [5](#), [6](#), [10](#), [16](#)
88. Yan, X., Gao, J., Li, J., Zhang, R., Li, Z., Huang, R., Cui, S.: Sparse single sweep lidar point cloud segmentation via learning contextual shape priors from scene completion. In: AAAI (2020) [4](#)
89. Yan, X., Zheng, C., Li, Z., Wang, S., Cui, S.: PointASNL: Robust point clouds processing using nonlocal neural networks with adaptive sampling. In: ICCV. pp. 5589–5598 (2020) [20](#)
90. Ye, X., Li, J., Huang, H., Du, L., Zhang, X.: 3D recurrent neural networks with context fusion for point cloud semantic segmentation. In: ECCV (2018) [20](#)
91. Zhang, B., Wang, Y., Hou, W., Wu, H., Wang, J., Okumura, M., Shinozaki, T.: Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Advances in Neural Information Processing Systems* **34** (2021) [17](#)
92. Zhang, F., Fang, J., Wah, B., Torr, P.: Deep fusionnet for point cloud semantic segmentation. In: ECCV. vol. 2, p. 6 (2020) [4](#)
93. Zhang, Y., Li, Z., Xie, Y., Qu, Y., Li, C., Mei, T.: Weakly supervised semantic segmentation for large-scale point cloud. In: AAAI (2021) [5](#), [6](#), [10](#), [11](#)
94. Zhang, Y., Qu, Y., Xie, Y., Li, Z., Zheng, S., Li, C.: Perturbed self-distillation: Weakly supervised large-scale point cloud semantic segmentation. In: ICCV. pp. 15520–15528 (2021) [2](#), [5](#), [6](#), [10](#), [11](#)
95. 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: CVPR. pp. 9601–9610 (2020) [24](#)
96. Zhang, Z., Girdhar, R., Joulin, A., Misra, I.: Self-supervised pretraining of 3D features on any point-cloud. arXiv preprint arXiv:2101.02691 (2021) [2](#), [4](#), [5](#)
97. Zhang, Z., Girdhar, R., Joulin, A., Misra, I.: Self-supervised pretraining of 3d features on any point-cloud. In: ICCV (2021) [2](#), [5](#)
98. Zhang, Z., Hua, B.S., Yeung, S.K.: ShellNet: Efficient point cloud convolutional neural networks using concentric shells statistics. In: ICCV. pp. 1607–1616 (2019) [11](#), [12](#), [20](#), [22](#), [24](#)
99. Zhao, H., Jiang, L., Fu, C.W., Jia, J.: PointWeb: Enhancing local neighborhood features for point cloud processing. In: CVPR (2019) [10](#), [20](#)
100. Zhao, H., Jiang, L., Jia, J., Torr, P., Koltun, V.: Point Transformer. arXiv preprint arXiv:2012.09164 (2020) [4](#)
101. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: CVPR. pp. 2921–2929 (2016) [4](#)
102. Zhu, X., Chen, J., Zeng, X., Liang, J., Li, C., Liu, S., Behpour, S., Xu, M.: Weakly supervised 3d semantic segmentation using cross-image consensus and inter-voxel affinity relations. In: ICCV (2021) [2](#)
103. 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: CVPR (2021) [1](#), [4](#)