

Mask3D: Mask Transformer for 3D Semantic Instance Segmentation

Jonas Schult¹, Francis Engelmann^{2,3}, Alexander Hermans¹, Or Litany⁴, Siyu Tang², Bastian Leibe¹

Abstract—Modern 3D semantic instance segmentation approaches predominantly rely on specialized voting mechanisms followed by carefully designed geometric clustering techniques. Building on the successes of recent Transformer-based methods for object detection and image segmentation, we propose the first Transformer-based approach for 3D semantic instance segmentation. We show that we can leverage generic Transformer building blocks to directly predict instance masks from 3D point clouds. In our model – called Mask3D – each object instance is represented as an *instance query*. Using Transformer decoders, the instance queries are learned by iteratively attending to point cloud features at multiple scales. Combined with point features, the instance queries directly yield all instance masks in parallel. Mask3D has several advantages over current state-of-the-art approaches, since it neither relies on (1) voting schemes which require hand-selected geometric properties (such as centers) nor (2) geometric grouping mechanisms requiring manually-tuned hyper-parameters (e.g. radii) and (3) enables a loss that directly optimizes instance masks. Mask3D sets a new state-of-the-art on ScanNet test (+6.2 mAP), S3DIS 6-fold (+10.1 mAP), STPLS3D (+11.2 mAP) and ScanNet200 test (+12.4 mAP).

I. INTRODUCTION

This work addresses the task of semantic instance segmentation of 3D scenes. That is, given a 3D point cloud, the desired output is a set of object instances represented as binary foreground masks (over all input points) with their corresponding semantic labels (e.g. ‘chair’, ‘table’, ‘window’).

Instance segmentation resides at the intersection of two problems: semantic segmentation and object detection. Therefore methods have opted to either first learn semantic point features followed by grouping them into separate instances (bottom-up) or detecting object instances followed by refining their semantic mask (top-down). Bottom-up approaches (ASIS [59], SGPN [58], 3D-BEVIS [12]) employ contrastive learning, mapping points to a high-dimensional feature space where features of the same instance are close together, and far apart otherwise. Top-down methods (3D-SIS [22], 3D-BoNet [61]) use an approach akin to Mask R-CNN [19]: First detect instances as bounding boxes and then perform mask segmentation on each box individually. While 3D-SIS [22] relies on predefined anchor boxes [19], 3D-BoNet [61] proposes an interesting variation that predicts bounding boxes from a global scene descriptor and optimizes an association loss based on bipartite matching [27]. A major step forward was sparked by powerful feature backbones [17, 53, 57] such as sparse convolutional networks [8, 16] that improve over existing PointNets [42, 44] and dense 3D CNNs [36, 43, 60].



Fig. 1: Mask3D. We train an end-to-end model for 3D semantic instance segmentation on point clouds. Given an input 3D point cloud (left), our Transformer-based model uses an attention mechanism to produce instance heatmaps across all points (center) and directly predicts all semantic object instances in parallel (right).

Well established 2D CNN architectures [20, 46] can now easily be adapted to sparse 3D data. These models can process large-scale 3D scenes in one pass, which is necessary to capture global scene context at multiple scales. As a result, bottom-up approaches which benefit from strong features (MTML [28], MASC [32]) experienced another performance boost. Soon after, inspired by Hough voting approaches [24, 30], VoteNet [41] proposed *center-voting* for 3D object detection. Instead of mapping points to an abstract high-dimensional feature space (as in bottom-up approaches), points now vote for their object center – votes from the same object are then closer to each other which enables geometric grouping into instance masks. This idea quickly influenced the 3D instance segmentation field, and by now, the vast majority of current state-of-the-art 3D instance segmentation methods [4, 13, 18, 26, 56] make use of both object center-voting and sparse feature backbones.

Although 3D instance segmentation has made impressive progress, current approaches have several major problems: typical state-of-the-art models are based on manually-tuned components, such as voting mechanisms that predict hand-selected geometric properties (e.g., centers [26], bounding boxes [7], occupancy [18]), and heuristics for clustering the votes (e.g., dual-set grouping [26], proposal aggregation [13], set aggregation/filtering [4]). Another limitation of these models is that they are not designed to directly predict instance masks. Instead, masks are obtained by grouping votes, and the model is trained using proxy-losses on the votes. A more elegant alternative consists of directly predicting and supervising instance masks, such as 3D-BoNet [61] or DyCo3D [21]. Recently, this idea gained popularity in 2D object detection (DETR [2]) and image segmentation (Mask2Former [5, 6]) but so far received less attention in 3D [21, 37, 61]. At the same time, in 2D image processing, we observe a strong shift from ubiquitous CNN architectures [19, 20, 45, 46] towards Transformer-based models [6, 11, 33]. In 3D, the move towards Transformers is less pronounced

¹ Computer Vision Group, RWTH Aachen University, Germany.

² Computer Vision and Learning Group, ETH Zürich, Switzerland.

³ ETH AI Center, Zürich, Switzerland.

⁴ NVIDIA, Santa Clara, USA

with only a few methods focusing on 3D object detection [34, 37, 39] or 3D semantic segmentation [29, 63, 63] and no methods for 3D instance segmentation. Overall, these approaches are still behind in terms of performance compared to current state-of-the-art methods [38, 48, 56, 57].

In this work, we propose the first Transformer-based model for 3D semantic instance segmentation of large-scale scenes that sets new state-of-the-art scores over a wide range of datasets, and addresses the aforementioned problems on hand-crafted model designs. The main challenge lies in directly predicting instance masks and their corresponding semantic labels. To this end, our model predicts *instance queries* that encode semantic and geometric information of each instance in the scene. Each instance query is then further decoded into a semantic class and an *instance feature*. The key idea (to directly generate masks) is to compute similarity scores between individual instance features and all point features in the point cloud [4, 6, 21]. This results in a heatmap over the point cloud, which (after normalization and thresholding) yields the final binary instance mask (*c.f.* Fig. 1). Our model, called Mask3D, builds on recent advances in both Transformers [5, 37] and 3D deep learning [8, 17, 57]: to compute strong point features, we leverage a sparse convolutional feature backbone [8] that efficiently processes full scenes and naturally provides multi-scale point features. To generate instance queries, we rely on stacked Transformer decoders [5, 6] that iteratively attend to learned point features in a coarse-to-fine fashion using non-parametric queries [37]. Unlike voting-based methods, directly predicting and supervising masks causes some challenges during training: before computing a mask loss, we first have to establish correspondences between predicted and annotated masks. A naïve solution would be to choose for each predicted mask the nearest ground truth mask [21]. However, this does not guarantee an optimal matching and any unmatched annotated mask would not contribute to the loss. Instead, we perform bipartite graph matching to obtain optimal associations between ground truth and predicted masks [2, 61]. We evaluate our model on four challenging 3D instance segmentation datasets, ScanNet v2 [9], ScanNet200 [47], S3DIS [1] and STPLS3D [3] and significantly outperform prior art, even surpassing architectures that are highly tuned towards specific datasets. Our experimental study compares various query types, different mask losses, and evaluates the number of queries as well as Transformer decoder steps.

Our contributions are as follows: (1) We propose the first competitive Transformer-based model for 3D semantic instance segmentation. (2) Our model named Mask3D builds on domain-agnostic components, avoiding center voting, non-maximum suppression, or grouping heuristics, and overall requires less hand-tuning. (3) Mask3D achieves state-of-the-art performance on ScanNet, ScanNet200, S3DIS and STPLS3D. To reach that level of performance with a Transformer-based approach, it is key to predict instance queries that encode the semantics and geometry of the scene and objects.

II. RELATED WORK

3D Instance Segmentation. Numerous methods have been proposed for 3D instance semantic segmentation, including bottom-up approaches [12, 28, 32, 58, 59], top-down approaches [22, 61, 61], and more recently, voting-based approaches [4, 13, 18, 26, 56]. MASC [32] uses a multi-scale hierarchical feature backbone, similar to ours, however, the multi-scale features are used to compute pairwise affinities followed by an offline clustering step. Such backbones are also successfully employed in other fields [5, 48]. Another influential work is DyCo3D [21], which is among the few approaches that directly predict instance masks without a subsequent clustering step. DyCo3D relies on *dynamic convolutions* [25, 54] which is similar in spirit to our mask prediction mechanism. However, it does not use optimal supervision assignment during training, resulting in subpar performance. Optimal assignment of the supervision signal was first implemented by 3D-BoNet [61] using Hungarian matching. Similar to ours, [61] directly predicts all instances in parallel. However, it uses only a single-scale scene descriptor which cannot encode object masks of diverse sizes.

Transformers. Initially proposed by Vaswani *et al.* [55] for NLP, Transformers have recently revolutionized the field of computer vision with successful models such as ViT [11] for image classification, DETR [2] for 2D object detection, or Mask2Former [5, 6] for 2D segmentation tasks. The success of Transformers has been less prominent in the 3D point cloud domain though and recent Transformer-based methods focus on either 3D object detection [34, 37, 39] or 3D semantic segmentation [29, 40, 63]. Most of these rely on specific attention modifications to deal with the quadratic complexity of the attention [29, 39, 40, 63]. Liu *et al.* [34] use vanilla Transformer decoder, but only to refine object proposals, whereas Misra *et al.* [37] are the first to show how to apply a vanilla Transformer to point clouds, still relying on an initial learned downsampling stage though. DyCo3D [21] also uses a Transformer, however at the bottleneck of the feature backbone to increase the receptive field size and is not related to our mechanism for 3D instance segmentation. In this work, we show how a vanilla Transformer decoder can be applied to the task of 3D semantic instance segmentation and achieve state-of-the-art performance.

III. METHOD

Fig. 2 illustrates our end-to-end 3D instance segmentation model Mask3D. As in Mask2Former [5], our model includes a feature backbone (□), a Transformer decoder (□) built from mask modules (□) and Transformer decoder layers used for query refinement (□). At the core of the model are *instance queries*, which each should represent one object instance in the scene and predict the corresponding point-level instance mask. To that end, the instance queries are iteratively refined by the Transformer decoder (Fig. 2, □) which allows the instance queries to cross-attend to point features extracted from the feature backbone and self-attend the other instance queries. This process is repeated for multiple iterations and

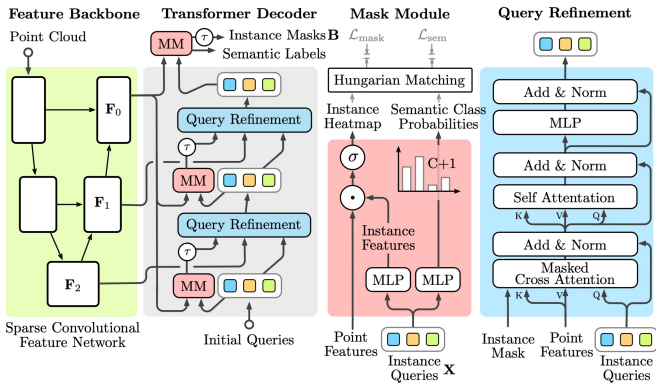


Fig. 2: Illustration of the Mask3D model. The feature backbone outputs multi-scale point features \mathbf{F} , while the Transformer decoder iteratively refines the instance queries \mathbf{X} . Given point features and instance queries, the mask module predicts for each query a semantic class and an instance heatmap, which (after thresholding) results in a binary instance mask \mathbf{B} . \odot applies a threshold of 0.5 and spatially rescales if required. \odot is the dot product. σ is the sigmoid function. We show a simplified model with fewer layers.

feature scales, yielding the final set of refined instance queries. A mask module consumes the refined instance queries together with the point features, and returns (for each query) a semantic class and a binary instance mask based on the dot product between point features and instance queries. Next, we describe each of these components in more detail.

Sparse Feature Backbone. (Fig. 2, \square) We use a sparse convolutional U-net backbone with a symmetrical encoder and decoder, based on the MinkowskiEngine [8]. Given a colored input point cloud $P \in \mathbb{R}^{N \times 6}$ of size N , it is first quantized into M_0 voxels $V \in \mathbb{R}^{M_0 \times 3}$, where each voxel is assigned the average RGB color of the points within that voxel as its initial feature. Next to the full-resolution output feature map $\mathbf{F}_0 \in \mathbb{R}^{M_0 \times D}$, we also extract a multi-resolution hierarchy of features from the backbone decoder before upsampling to the next finer feature map. At each of these resolutions $r \geq 0$ we can extract features for a set of M_r voxels, which we linearly project to a fixed and common dimension D , yielding feature matrices $\mathbf{F}_r \in \mathbb{R}^{M_r \times D}$. We let the queries attend to features from coarser feature maps of the backbone decoder, *i.e.* $r \geq 1$, and use the full-resolution feature map ($r = 0$) to compute the auxiliary and final per-voxel instance masks.

Mask Module. (Fig. 2, \square) Given the set of K instance queries $\mathbf{X} \in \mathbb{R}^{K \times D}$, we predict a binary mask for each instance and classify each of them as one of C classes or as being inactive. To create the binary mask, we map the instance queries through an MLP $f_{\text{mask}}(\cdot)$, to the same feature space as the backbone output features. We then compute the dot product between these *instance features* and the backbone features \mathbf{F}_0 . The resulting similarity scores are fed through a sigmoid and thresholded at 0.5, yielding the final binary mask $\mathbf{B} \in \{0, 1\}^{M \times K}$:

$$\mathbf{B} = \{b_{i,j} = [\sigma(\mathbf{F}_0 f_{\text{mask}}(\mathbf{X})^T)_{i,j} > 0.5]\}. \quad (1)$$

We apply the mask module to the refined queries \mathbf{X} at each Transformer layer using the full-resolution feature map \mathbf{F}_0 , to

create auxiliary binary masks for the masked cross-attention of the following refinement step. When this mask is used as input for the masked cross-attention, we reduce the resolution according to the voxel mask by average pooling. Next to the binary mask, we predict a single semantic class per instance. This step is done via a linear projection layer into $C + 1$ dimensions, followed by a softmax. While prior work [4, 13, 56] typically needs to obtain the semantic label of an instance via majority voting or grouping over per-point predicted semantics, this information is directly contained in the refined instance queries.

Query Refinement. (Fig. 2, \square) The Transformer decoder starts with K instance queries, and refines them through a stack of L Transformer decoder layers to a final set of accurate, scene specific instance queries by cross-attending to scene features, and reasoning at the instance-level through self-attention. We discuss different types of instance queries in Sec. III-A. Each layer attends to one of the feature maps from the feature backbone using standard cross-attention:

$$\mathbf{X} = \text{softmax}(\mathbf{Q}\mathbf{K}^T/\sqrt{D})\mathbf{V}. \quad (2)$$

To do so, the voxel features $\mathbf{F}_r \in \mathbb{R}^{M_r \times D}$ are first linearly projected to a set of keys and values of fixed dimensionality $\mathbf{K}, \mathbf{V} \in \mathbb{R}^{M_r \times D}$ and our K instance queries \mathbf{X} are linearly projected to the queries $\mathbf{Q} \in \mathbb{R}^{K \times D}$. This cross-attention thus allows the queries to extract information from the voxel features. The cross-attention is followed by a self-attention step between the queries, where the keys, values, and queries are all computed based on linear projections of the instance queries. Without such inter-query communications, the model could not avoid multiple instance queries latching onto the same object, resulting in duplicate instance masks. Similar to most Transformer-based approaches, we use positional encodings for our keys and queries. We use Fourier positional encodings [52] based on voxel positions. We add the resulting positional encodings to their respective keys before computing the cross-attention. All instance queries are also assigned a fixed (and potentially learned) positional embedding, that is not updated throughout the query refinement process. These positional encodings are added to the respective queries in the cross-attention, as well as to both the keys and queries in the self-attention. Instead of using the vanilla cross-attention (where each query attends to all voxel features in one resolution) we use a masked variant where each instance query only attends to the voxels within its corresponding intermediate instance mask \mathbf{B} predicted by the previous layer. This is realized by adding $-\infty$ to the attention matrix to all voxels for which the mask is 0. Eq. 2 then becomes:

$$\mathbf{X} = \text{softmax}(\mathbf{Q}\mathbf{K}^T/\sqrt{D} + \mathbf{B}')\mathbf{V} \text{ with } \mathbf{B}'_{i,j} = -\infty \cdot [\mathbf{B}_{i,j} = 0] \quad (3)$$

where $[\cdot]$ are Iverson brackets. In [5], masking out the context from the cross-attention improved segmentation. A likely reason is that the Transformer does not need to *learn* to focus on a specific instance instead of irrelevant context, but is *forced* to do so by design.

In practice, we attend to the 4 coarsest levels of the feature backbone, from coarse to fine, and do this a total of 3 times, resulting in $L = 12$ query refinement steps. The Transformer decoder layers share weights for all 3 iterations. Early experiments showed that this approach preserves the performance while keeping memory requirements in bound. **Sampled Cross-Attention.** Point clouds in a training batch typically have different point counts. While MinkowskiEngine can handle this, current Transformer implementations rely on a fixed number of points in each batch entry. In order to leverage well-tested Transformer implementations, in this work we propose to pad the voxel features and mask out the attention where needed. In case the number of voxels exceeds a certain threshold, we resort to *sampling* voxel features. To allow instances to have access to all voxel features during cross-attention, we resample the voxels in each Transformer decoder layer though, and use all voxels during inference. This can be seen as a form of dropout [50]. In practice, this procedure saves significant amounts of memory and is crucial for obtaining competitive performance. In particular, since the proposed sampled cross-attention requires less memory, it enables training on higher-resolution voxel grids which is necessary for achieving competitive results on common benchmarks (e.g., 2 cm voxel side-length on ScanNet [9]).

A. Training and Implementation Details

Correspondences. Given that there is no ordering to the set of instances in a scene and the set of predicted instances, we need to establish correspondences between the two sets during training. To that end, we use bipartite graph matching. While such a supervision approach is not new (e.g. [51, 61]), recently it has become more common in Transformer-based approaches [2, 5, 6]. We construct a cost matrix $\mathcal{C} \in \mathbb{R}^{K \times \hat{K}}$, where \hat{K} is the number of ground truth instances in a scene. The matching cost for a predicted instance k and a target instance \hat{k} is given by:

$$\mathcal{C}(k, \hat{k}) = \lambda_{\text{dice}} \mathcal{L}_{\text{dice}}(k, \hat{k}) + \lambda_{\text{BCE}} \mathcal{L}_{\text{BCE}_{\text{mask}}}(k, \hat{k}) + \lambda_{\text{cl}} \mathcal{L}_{\text{CE}_{\text{cl}}}(k, \hat{k}) \quad (4)$$

We set the weights to $\lambda_{\text{dice}} = \lambda_{\text{cl}} = 2.0$ and $\lambda_{\text{BCE}} = 5.0$. The optimal solution for this cost assignment problem is efficiently found using the Hungarian method [27]. After establishing the correspondences, we can directly optimize each predicted mask as follows:

$$\mathcal{L}_{\text{mask}} = \lambda_{\text{BCE}} \mathcal{L}_{\text{BCE}} + \lambda_{\text{dice}} \mathcal{L}_{\text{dice}}, \quad (5)$$

where \mathcal{L}_{BCE} is the binary cross-entropy loss (over the foreground and background of that mask) and $\mathcal{L}_{\text{dice}}$ is the Dice loss [10]. We use the default multi-class cross-entropy loss $\mathcal{L}_{\text{CE}_{\text{cl}}}$ to supervise the classification. If a mask is left unassigned, we seek to maximize the associated *no-object* class, for which the $\mathcal{L}_{\text{CE}_{\text{cl}}}$ loss is weighted by an additional $\lambda_{\text{no-obj.}} = 0.1$. The overall loss for all auxiliary instance predictions after each of the L layers is defined as:

$$\mathcal{L} = \sum_l^L \mathcal{L}_{\text{mask}}^l + \lambda_{\text{cl}} \mathcal{L}_{\text{CE}_{\text{cl}}}^l \quad (6)$$

Prediction Confidence Score. We seek to assign a confidence to each predicted instance. While other existing methods

require a dedicated ScoreNet [26] which is trained to estimate the intersection over union with the ground truth instances, we directly obtain the confidence scores from the refined query features and point features as in Mask2Former [5]. We first select the queries with a dominant semantic class, for which we obtain the class confidence based on the softmax output $c_{\text{cl}} \in [0, 1]$, which we additionally multiply with a mask based confidence:

$$c = c_{\text{cl}} \cdot (\sum_i^M m_i \cdot [m_i > 0.5]) / (\sum_i^M [m_i > 0.5]), \quad (7)$$

where $m_i \in [0, 1]$ is the instance mask confidence for the i^{th} voxel given a single query. In essence, this is the mean mask confidence of all voxels falling inside of the binarized mask [5]. For an instance prediction to have a high confidence, it needs both a confident classification among C classes, and a mask that predominantly consists of high-confidence voxels.

Query Types. Methods like DETR [2] or Mask2Former [5, 6] use parametric queries. During training both the instance query features and the corresponding positional encodings are learned. This thus means that during training the set of K instance queries has to be optimized in such a way that it can cover all instances present in a scene during inference.

Misra *et al.* [37] propose to initialize queries with sampled point coordinates from the input point cloud based on farthest point sampling. Since this initialization does not involve learned parameters, they are called *non-parametric* queries. Interestingly, the instance query features are initialized with zeros and only the 3D position of the sampled points is used to set the corresponding positional encoding. We also experiment with a variant where we use sampled point features as instance query features. Similar to [37], we observe improved performance when using non-parametric queries although less pronounced. The key advantage of non-parametric queries is that, during inference, we can sample a different number of queries than during training. This provides a trade-off between inference speed and performance, without the need to retrain the model when using more instance queries.

Training Details. The feature backbone is a Minkowski Res16UNet34C [8]. We train for 600 epochs using AdamW [35] and a one-cycle learning rate schedule [49] with a maximal learning rate of 10^{-4} . Longer training times (1000 epochs) did not further improve results. One training on 2 cm voxelization takes ~ 78 hours on an NVIDIA A40 GPU. We perform standard data augmentation: horizontal flipping, random rotations around the z-axis, elastic distortion [46] and random scaling. Color augmentations include jittering, brightness and contrast augmentations. During training on ScanNet, we reduce memory consumption by computing the dot product between instance queries and aggregated point features within segments (obtained from a graph-based segmentation [15], similar to OccuSeg [18] or Mix3D [38]). Wrongly merged instances can be separated using connected components [14] (Sec. IV-C).

IV. EXPERIMENTS

In this section, we compare Mask3D with prior state-of-the-art on four publicly available 3D indoor and outdoor datasets

TABLE I: 3D Instance Segmentation Scores on ScanNet v2 [9]. We report mean average precision (mAP) with different IoU threshold over 18 classes on the ScanNet validation and test set. The inference speed is averaged over the validation set and computed on a TITAN X GPU (*c.f.* [56]), excluding postprocessing. Test scores accessed on 13. September 2022.

Method	ScanNet Val		ScanNet Test		Runtime (in ms)
	mAP	mAP ₅₀	mAP	mAP ₅₀	
SGPN [58]	–	–	4.9	14.3	158439
GSPN [62]	19.3	37.8	–	30.6	12702
3D-SIS [22]	–	18.7	16.1	38.2	–
MASC [32]	–	–	25.4	44.7	–
3D-Bonet [61]	–	–	25.3	48.8	9202
MTML [28]	20.3	40.2	28.2	54.9	–
3D-MPA [13]	35.5	59.1	35.5	61.1	–
DyCo3D [21]	35.4	57.6	39.5	64.1	–
PointGroup [26]	34.8	56.7	40.7	63.6	452
MaskGroup [64]	42.0	63.3	43.4	66.4	–
OccuSeg [18]	44.2	60.7	48.6	67.2	1904
SSTNet [31]	49.4	64.3	50.6	69.8	428
HAIS [4]	43.5	64.1	45.7	69.9	339
SoftGroup [56]	46.0	67.6	50.4	76.1	345
Mask3D (Ours)	55.2	73.7	56.6	78.0	339

(Sec. IV-A). Then, we provide analysis experiments on the proposed model investigating query types and the impact of the number of query refinement steps as well as the number of queries during inference. (Sec. IV-B). Finally, we show qualitative results and discuss limitations (Sec. IV-C).

A. Comparing with State-of-the-Art Methods

Datasets and Metrics. We evaluate Mask3D on four publicly available 3D instance segmentation datasets.

ScanNet [9] is a richly-annotated dataset of 3D reconstructed indoor scenes. It contains hundreds of different rooms showing a large variety of room types such as hotels, libraries and offices. The provided splits contain 1202 training, 312 validation and 100 hidden test scenes. Each scene is annotated with semantic and instance segmentation labels covering 18 object categories. The benchmark evaluation metric is mean average precision (mAP). *ScanNet200* [47] extends the original ScanNet scenes with an order of magnitude more classes. ScanNet200 allows to test an algorithm’s performance under the natural imbalance of classes, particularly for challenging long-tail classes such as *coffee-kettle* and *potted-plant*. We keep the same train, validation and test splits as in the original ScanNet dataset. *S3DIS* [1] is a large-scale indoor dataset showing six different areas from three different campus buildings. It contains 272 scans and is also annotated with semantic instance masks over 13 different classes. We follow the common splits and evaluate on Area-5 and 6-fold cross validation. We report scores using the mAP metric from ScanNet and mean precision/recall at IoU threshold 50% (mPrec₅₀/mRec₅₀) as initially introduced by ASIS [59]. Unlike mAP, this metric does not consider confidence scores, therefore we filter out instance masks with a prediction confidence score below 80% to avoid excessive false positives. *STPLS3D* [3] is a synthetic outdoor dataset closely mimicking

TABLE II: 3D Instance Segmentation Scores on S3DIS [1]. We report mean average precision (mAP) with different IoU threshold (as in [9]) as well as mean precision (mPrec) and mean recall (mRec) with 50% IoU threshold (as in [59]) over 13 classes on S3DIS Area 5 and 6-fold cross validation. Scores in light gray are pre-trained on ScanNet [9] and fine-tuned on S3DIS [1].

Method	S3DIS Area 5				S3DIS 6-fold CV			
	AP	AP ₅₀	Prec ₅₀	Rec ₅₀	AP	AP ₅₀	Prec ₅₀	Rec ₅₀
SGPN [58]	–	–	36.0	28.7	–	–	38.2	31.2
ASIS [59]	–	–	55.3	42.4	–	–	63.6	47.5
3D-Bonet [61]	–	–	57.5	40.2	–	–	65.6	47.6
OccuSeg [18]	–	–	–	–	–	–	72.8	60.3
3D-MPA [13]	–	–	63.1	58.0	–	–	66.7	64.1
PointGroup [26]	–	57.8	61.9	62.1	–	64.0	69.6	69.2
DyCo3D [21]	–	–	64.3	64.2	–	–	–	–
MaskGroup [64]	–	65.0	62.9	64.7	–	69.9	66.6	69.6
SSTNet [31]	42.7	59.3	65.5	64.2	54.1	67.8	73.5	73.4
Mask3D (Ours)	56.6	68.4	68.7	66.3	64.5	75.5	72.8	74.5
HAIS [4]	–	–	71.1	65.0	–	–	73.2	69.4
SoftGroup [56]	51.6	66.1	73.6	66.6	54.4	68.9	75.3	69.8
Mask3D (Ours)	57.8	71.9	74.3	63.7	61.8	74.3	76.5	66.2

TABLE III: 3D Instance Segmentation Scores on ScanNet200 [47] and STPLS3D [3]. We report mean average precision (mAP) with different IoU threshold over 14 classes on the STPLS3D test set. Hidden test scores accessed on 13. September 2022.

Method	ScanNet 200			Method	STPLS3D	
	head	com	tail		mAP	mAP ₅₀
CSC [23]	22.3	8.2	4.6	PointGroup [26]	23.3	38.5
Mink34D [8]	24.6	8.3	4.3	HAIS [4]	35.1	46.7
LGround [47]	27.5	10.8	6.0	SoftGroup [56]	46.2	61.8
Mask3D (Ours)	38.3	26.3	16.8	Mask3D (Ours)	57.3	74.3

the data generation process of aerial photogrammetry point clouds. 25 urban scenes totalling 6 km² are densely annotated with 14 instance classes. We follow the common splits [3, 56].

Results are summarized in Tab. I (ScanNet), Tab. II (S3DIS), Tab. III (*left*, ScanNet200) and Tab. III (*right*, STPLS3D). Mask3D outperforms prior work by a large margin on the most challenging metric mAP by at least **6.2 mAP** on ScanNet, **6.2 mAP** on S3DIS, **10.8 mAP** on ScanNet200 and **11.2 mAP** on STPLS3D. As in [4, 56], we also report scores for models pre-trained on ScanNet and fine-tuned on S3DIS. For Mask3D, pre-training improves performance by 1.2 mAP on Area 5. Mask3D’s strong performance on indoor and outdoor datasets as well as its ability to work under challenging class imbalance settings *without* inherent modifications to the architecture or the training regime highlights its generality. Trained models are available at: <https://github.com/JonasSchult/Mask3D>

B. Analysis Experiments

Query Types. Mask3D iteratively refines instance queries by attending to voxel features (Fig. 2, □). We distinguish two types of query initialization prior to attending to voxel features: ① parametric and ②-③ non-parametric initial

TABLE IV: Ablations. a) We explore two variants for query positions and features. Parametric queries ① are learned during training. Non-parametric queries consist of FPS point positions ② and potentially their features ③, resembling scene-specific queries. b) We optimize the instance mask prediction using the binary cross-entropy loss \mathcal{L}_{CE} and the dice loss \mathcal{L}_{dice} . A weighted combination of dice and cross-entropy loss results in best performance.

Position Features			mAP	\mathcal{L}_{dice} \mathcal{L}_{BCE} mAP		
① Param.	Param.		39.7±0.7	✗	✓	27.0±0.6
② FPS	Zeros		40.6±0.3	✓	✗	38.0±0.3
③ FPS	Point Feat.		38.4±0.3	✓	✓	40.6±0.3




Fig. 3: Number of queries and decoder layers.

queries. Parametric refers to *learned* positions and features [2], while non-parametric refers to point positions sampled with *furthest point sampling* (FPS) [44]. When selecting query positions with FPS, we can either initialize the queries to zero (②, as in 3DETR [37]) or use the point features at the sampled position (③). Tab. IV (left) shows the effects of using parametric or non-parametric queries on ScanNet validation (5 cm). In line with [37], we see that non-parametric queries ② outperform parametric queries ①. Interestingly, ③ results in degraded performance compared to both parametric ① and position-only non-parametric queries ②.

Number of Queries and Decoders. We analyze the effect of varying numbers of queries K during inference on models trained with $K = 100$ and $K = 200$ non-parametric queries sampled with FPS. By increasing K from 100 to 200 during training, we observe a slight increase in performance (Fig. 3, left) at the cost of additional memory. When evaluating with fewer queries than trained with, we observe reduced performance but faster runtime. When evaluating with more queries than trained with, we observe slightly improved performance, typically less than 1 mAP. Our final model uses $K = 100$ due to memory constraint when using 2 cm voxels in the feature backbone. In this study, we report scores using 5 cm on ScanNet validation. We also analyse the mask quality that we obtain after each Transformer decoder layer in our trained model (Fig. 3, right). We see a rapid increase up to 4 layers, then the quality increases a bit slower.

Mask Loss. The mask module (Fig. 2, □) generates instance heatmaps for every instance query. After Hungarian matching, the corresponding ground truth mask is used to compute the mask loss \mathcal{L}_{mask} . The binary cross entropy loss \mathcal{L}_{BCE} is the obvious choice for binary segmentation tasks. However, it does not perform well under large class imbalance (few

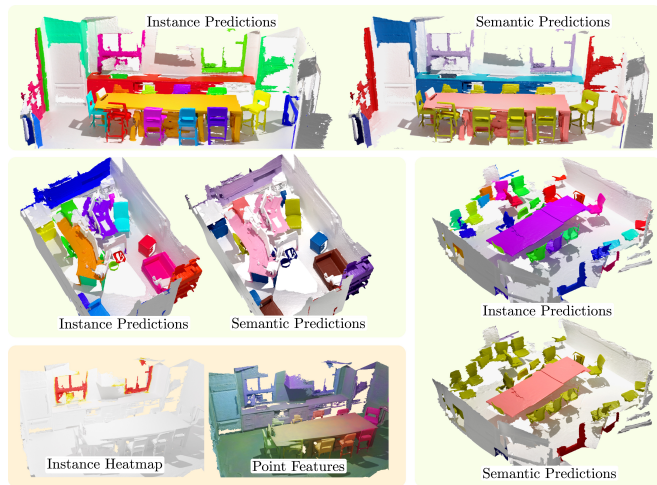


Fig. 4: Qualitative Results on ScanNet. We show pairs of predicted instance masks and predicted semantic labels. On the bottom left, we show the heatmap of a failure case of two windows that are wrongly assigned to a single instance. The corresponding point features are visualized as RGB after projecting them to 3D using PCA.

foreground mask points, many background points). The Dice loss \mathcal{L}_{dice} is specifically designed to address such data imbalance. Tab. IV (right) shows scores on ScanNet validation for combinations of both losses. While \mathcal{L}_{dice} improves over \mathcal{L}_{BCE} , we observe an additional improvement by training our model with a weighted sum of both losses (Eq. 5).

C. Qualitative Results and Limitations

Fig. 4 shows several representative examples of Mask3D instance segmentation results on ScanNet. The scenes are quite diverse and present a number of challenges, including clutter, scanning artifacts and numerous similar objects. Still, our model shows quite robust results. There are still limitations in our model though. A systematic mistake that we observed are merged instances that are far apart (see Fig. 4, bottom left). As the attention mechanism can attend to the full point cloud, it can happen that two objects with similar semantics and geometry expose similar learned point features and are therefore combined into one instance even if they are far apart in the scene. This is less likely to happen with methods that explicitly encode geometric priors.

V. CONCLUSION

In this work, we have introduced Mask3D, for 3D semantic instance segmentation. Mask3D is based on Transformer decoders, and learns instance queries that, combined with learned point features, directly predict semantic instance masks without the need for hand-selected voting schemes or hand-crafted grouping mechanisms. We think that Mask3D is an attractive alternative to current voting-based approaches and expect to see follow-up work along this line of research.

Acknowledgments: This work is supported by the ERC Consolidator Grant DeeViSe (ERC-2017-CoG-773161), SNF Grant 200021 204840, compute resources from RWTH Aachen University (rwth1238) and the ETH AI Center post-doctoral fellowship. We additionally thank Alexey Nekrasov, Ali Athar and István Sárárdi for helpful discussions and feedback.

REFERENCES

- [1] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3D Semantic Parsing of Large-Scale Indoor Spaces. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-End Object Detection with Transformers. In *European Conference on Computer Vision*, 2020.
- [3] Meida Chen, Qingyong Hu, Thomas Hugues, Andrew Feng, Yu Hou, Kyle McCullough, and Lucio Soibelman. STPLS3D: A Large-Scale Synthetic and Real Aerial Photogrammetry 3D Point Cloud Dataset. *arXiv:2203.09065*, 2022.
- [4] Shaoyu Chen, Jiemin Fang, Qian Zhang, Wenyu Liu, and Xinggang Wang. Hierarchical Aggregation for 3D Instance Segmentation. In *International Conference on Computer Vision*, 2021.
- [5] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention Mask Transformer for Universal Image Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [6] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-Pixel Classification is Not All You Need for Semantic Segmentation. In *Advances in Neural Information Processing Systems*, 2021.
- [7] Julian Chibane, Francis Engelmann, Tuan Anh Tran, and Gerard Pons-Moll. Box2Mask: Weakly Supervised 3D Semantic Instance Segmentation Using Bounding Boxes. In *European Conference on Computer Vision*, 2022.
- [8] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [9] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3D Reconstructions of Indoor Scenes. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [10] Ruoxi Deng, Chunhua Shen, Shengjun Liu, Huibing Wang, and Xinru Liu. Learning to Predict Crisp Boundaries. In *European Conference on Computer Vision*, 2018.
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*, 2021.
- [12] Cathrin Elich, Francis Engelmann, Theodora Kontogianni, and Bastian Leibe. 3D Bird's-eye-view Instance Segmentation. In *German Conference on Pattern Recognition*, 2019.
- [13] Francis Engelmann, Martin Bokeloh, Alireza Fathi, Bastian Leibe, and Matthias Nießner. 3D-MPA: Multi-Proposal Aggregation for 3D Semantic Instance Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [14] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1996.
- [15] Pedro F Felzenszwalb and Daniel P Huttenlocher. Efficient Graph-Based Image Segmentation. *International Journal of Computer Vision*, 59(2):167–181, 2004.
- [16] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3D Semantic Segmentation with Submanifold Sparse Convolutional Networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [17] Benjamin Graham and Laurens van der Maaten. Submanifold Sparse Convolutional Networks. *arXiv:1706.01307*, 2017.
- [18] Lei Han, Tian Zheng, Lan Xu, and Lu Fang. OccuSeg: Occupancy-aware 3D Instance Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [19] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *International Conference on Computer Vision*, 2017.
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [21] Tong He, Chunhua Shen, and Anton van den Hengel. DyCo3D: Robust Instance Segmentation of 3D Point Clouds through Dynamic Convolution. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [22] Ji Hou, Angela Dai, and Matthias Nießner. 3D-SIS: 3D Semantic Instance Segmentation of RGB-D Scans. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [23] Ji Hou, Benjamin Graham, Matthias Nießner, and Saining Xie. Exploring Data-Efficient 3D Scene Understanding with Contrastive Scene Contexts. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [24] Paul VC Hough. Machine Analysis of Bubble Chamber Pictures. In *International Conference on High Energy Accelerators and Instrumentation*, 1959.
- [25] Xu Jia, Bert De Brabandere, Tinne Tuytelaars, and Luc V Gool. Dynamic Filter Networks. *Neural Information Processing Systems*, 2016.
- [26] Li Jiang, Hengshuang Zhao, Shaoshuai Shi, Shu Liu, Chi-Wing Fu, and Jiaya Jia. PointGroup: Dual-Set Point Grouping for 3D Instance Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [27] Harold W Kuhn. The Hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- [28] Jean Lahoud, Bernard Ghanem, Marc Pollefeys, and Martin R Oswald. 3D Instance Segmentation via Multi-Task Metric Learning. In *International Conference on Computer Vision*, 2019.
- [29] Xin Lai, Jianhui Liu, Li Jiang, Liwei Wang, Hengshuang Zhao, Shu Liu, Xiaojuan Qi, and Jiaya Jia. Stratified Transformer for 3D Point Cloud Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [30] Bastian Leibe, Aleš Leonardis, and Bernt Schiele. Robust Object Detection with Interleaved Categorization and Segmentation. *International Journal of Computer Vision*, 2008.
- [31] Zhihao Liang, Zhihao Li, Songcen Xu, Mingkui Tan, and Kui Jia. Instance Segmentation in 3D Scenes using Semantic Superpoint Tree Networks. In *International Conference on Computer Vision*, 2021.
- [32] Chen Liu and Yasutaka Furukawa. MASC: Multi-Scale Affinity with Sparse Convolution for 3D Instance Segmentation. *arXiv:1902.04478*, 2019.
- [33] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In *International Conference on Computer Vision*, 2021.
- [34] Ze Liu, Zheng Zhang, Yue Cao, Han Hu, and Xin Tong. Group-Free 3D Object Detection via Transformers. In *International Conference on Computer Vision*, 2021.
- [35] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*, 2019.
- [36] Daniel Maturana and Sebastian Scherer. VoxNet: A 3D Convolutional Neural Network for Real-time Object Recognition. In *International Conference on Intelligent Robots and Systems*, 2015.
- [37] Ishan Misra, Rohit Girdhar, and Armand Joulin. An End-to-End Transformer Model for 3D Object Detection. In *International Conference on Computer Vision*, 2021.
- [38] Alexey Nekrasov, Jonas Schult, Or Litany, Bastian Leibe, and Francis Engelmann. Mix3D: Out-of-Context Data Augmentation for 3D Scenes. In *International Conference on 3D Vision*, 2021.
- [39] Xuran Pan, Zhuofan Xia, Shiji Song, Li Erran Li, and Gao Huang. 3D Object Detection With Pointformer. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [40] Chunhyun Park, Yoonwoo Jeong, Minsu Cho, and Jaesik Park. Fast Point Transformer. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [41] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep Hough Voting for 3D Object Detection in Point Clouds. In *International Conference on Computer Vision*, 2019.
- [42] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [43] Charles R Qi, Hao Su, Matthias Nießner, Angela Dai, Mengyuan Yan, and Leonidas J Guibas. Volumetric and Multi-View CNNs for Object Classification on 3D Data. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [44] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *Advances in Neural Information Processing Systems*, 2017.

- [45] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Neural Information Processing Systems*, 2015.
- [46] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2015.
- [47] David Rozenberszki, Or Litany, and Angela Dai. Language-Grounded Indoor 3D Semantic Segmentation in the Wild. In *European Conference on Computer Vision*, 2022.
- [48] Danila Rukhovich, Anna Vorontsova, and Anton Konushin. FCAF3D: Fully Convolutional Anchor-Free 3D Object Detection. *arXiv:2112.00322*, 2021.
- [49] Leslie N Smith and Nicholay Topin. Super-Convergence: Very Fast Training of Neural Networks Using Large Learning Rates. In *Artificial intelligence and machine learning for multi-domain operations applications*, 2019.
- [50] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [51] Russell Stewart, Mykhaylo Andriluka, and Andrew Y Ng. End-to-end people detection in crowded scenes. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [52] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains. In *Advances in Neural Information Processing Systems*, 2020.
- [53] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotequi, François Goulette, and Leonidas J Guibas. KPConv: Flexible and Deformable Convolution for Point Clouds. In *International Conference on Computer Vision*, 2019.
- [54] Zhi Tian, Chunhua Shen, and Hao Chen. Conditional Convolutions for Instance Segmentation. In *European Conference on Computer Vision*, 2020.
- [55] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. In *Advances in Neural Information Processing Systems*, 2017.
- [56] Thang Vu, Kookhoi Kim, Tung M Luu, Xuan Thanh Nguyen, and Chang D Yoo. SoftGroup for 3D Instance Segmentation on Point Clouds. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [57] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-CNN: Octree-based Convolutional Neural Networks for 3D Shape Analysis. *ACM Transactions on Graphics*, 2017.
- [58] Weiyue Wang, Ronald Yu, Qiangui Huang, and Ulrich Neumann. SGPN: Similarity Group Proposal Network for 3D Point Cloud Instance Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [59] Xinlong Wang, Shu Liu, Xiaoyong Shen, Chunhua Shen, and Jiaya Jia. Associatively Segmenting Instances and Semantics in Point Clouds. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [60] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3D ShapeNets: A Deep Representation for Volumetric Shapes. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [61] Bo Yang, Jianan Wang, Ronald Clark, Qingyong Hu, Sen Wang, Andrew Markham, and Niki Trigoni. Learning Object Bounding Boxes for 3D Instance Segmentation on Point Clouds. In *Advances in Neural Information Processing Systems*, 2019.
- [62] Li Yi, Wang Zhao, He Wang, Minhyuk Sung, and Leonidas J Guibas. GSPN: Generative Shape Proposal Network for 3D Instance Segmentation in Point Cloud. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [63] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point Transformer. In *International Conference on Computer Vision*, 2021.
- [64] Min Zhong, Xinghao Chen, Xiaokang Chen, Gang Zeng, and Yunhe Wang. MaskGroup: Hierarchical Point Grouping and Masking for 3D Instance Segmentation. *arXiv:2203.14662*, 2022.