

# Fine-Grained Spatial and Verbal Losses for 3D Visual Grounding

Sombit Dey<sup>1,2</sup> Ozan Unal<sup>1,3\*</sup> Christos Sakaridis<sup>1</sup> Luc Van Gool<sup>1,2,4</sup>  
<sup>1</sup>ETH Zurich, <sup>2</sup>INSAIT, <sup>3</sup>Huawei Technologies, <sup>4</sup>KU Leuven

## Abstract

3D visual grounding consists of identifying the instance in a 3D scene which is referred to by an accompanying language description. While several architectures have been proposed within the commonly employed grounding-by-selection framework, the utilized losses are comparatively under-explored. In particular, most methods rely on a basic supervised cross-entropy loss on the predicted distribution over candidate instances, which fails to model both spatial relations between instances and the internal fine-grained word-level structure of the verbal referral. Sparse attempts to additionally supervise verbal embeddings globally by learning the class of the referred instance from the description or employing verbo-visual contrast to better separate instance embeddings do not fundamentally lift the aforementioned limitations. Responding to these shortcomings, we introduce two novel losses for 3D visual grounding: a visual-level offset loss on regressed vector offsets from each instance to the ground-truth referred instance and a language-related span loss on predictions for the word-level span of the referred instance in the description. In addition, we equip the verbo-visual fusion module of our new 3D visual grounding architecture AsphaltNet with a top-down bidirectional attentive fusion block, which enables the supervisory signals from our two losses to propagate to the respective converse branches of the network and thus aid the latter to learn context-aware instance embeddings and grounding-aware verbal embeddings. AsphaltNet proposes novel auxiliary losses to aid 3D visual grounding with competitive results compared to the state-of-the-art on the ReferIt3D benchmark.

## 1. Introduction

With the ever-increasing popularity of the interaction- and 3D-heavy applications of augmented/virtual reality and human-support-oriented robotics in indoor home environments, bridging the gap between language understanding and 3D visual perception is becoming essential. This need has motivated a large body of recent works on the task of




Figure 1. We propose two auxiliary losses to aid the task of 3D visual grounding: i) offset loss  $\mathcal{L}_o$  between instances to promote better localization and improve object separability; ii) span loss  $\mathcal{L}_{sp}$  on the prompts to provide a higher granularity of supervision.

3D visual grounding [16,18,32,42,43], surveys papers [21], which consists in identifying and localizing the object in the input 3D scene which is referred by a language description.

3D visual grounding works are classified into two main categories: (i) referral-based 3D object identification methods and (ii) referral-based 3D object localization methods. The former variant requires as visual inputs both the raw 3D point cloud and the ground-truth 3D bounding boxes or instance masks of all objects in the input scene which belong to the class of the referred object [2]. On the contrary, the latter variant consists in the 3D localization of the referred object directly from the raw 3D point cloud [5]. While the two variants might appear fundamentally different at first sight, referral-based localization reduces to referral-based identification after candidate object boxes or masks have been generated. 3D grounding is typically performed in both cases in a grounding-by-selection fashion, using a verbo-visual fusion module which aims to classify the re-

\*Corresponding author: Ozan Unal, ozan.unal@vision.ee.ethz.ch

ferred object based on the language description.

The de-facto choice in the literature [5, 18, 36, 42, 43] for supervising the above verbo-visual fusion module is via a supervised cross-entropy loss on the module’s softmax classification outputs over the set of candidate objects. The limitations of this loss are that (i) it treats the candidate 3D objects as isolated entities that independently correlate with the language description, even though these objects co-exist in the 3D scene and participate in mutual spatial arrangements and relative configurations, and (ii) it views the language description as a monolithic entity, neglecting its internal structure. Some works [2, 17, 43] extend supervision via an additional, language-driven cross-entropy loss on a class distribution of the referred object predicted from a global language embedding. While this language-driven loss can increase the awareness of language features about the semantics of the referred object, it does not lift limitation (ii) above, as it still operates on a global language embedding, which obscures the attribution of gradients to individual word embeddings that are more closely associated with the referred object. Another line of works [18, 36] proposes contrastive losses between verbal and visual embeddings to inject a better structure in the joint verbo-visual embedding space. However, such cross-modal contrast improves separation between distinct instance embeddings only *indirectly* by pulling them to / pushing them away from sentence [36] or word [18] embeddings and it still treats each instance individually and not jointly with others, leaving limitation (i) above unaddressed.

In this work, we attempt to smooth the loss manifold for 3D visual grounding by proposing two novel losses to overcome the two aforementioned limitations of the basic supervised grounding-by-selection setup. First, we propose to make 3D instances aware of spatial visual grounding context by letting each of them regress a 3D offset vector pointing to the ground-truth referred instance. In this way, we encourage the instance embeddings to attend to each other in a way that is informed about the *spatial* location to which the semantics of the referral is grounded. The specific implementation we follow for the verbo-visual fusion module involves a basic attention-based building block. As we aim at injecting the aforementioned grounding awareness into embeddings from an early-on stage of the verbo-visual fusion module, we add offset vector regression heads after each such attention block and supervise each of them with a distinct offset regression loss. For the attention-based block, we build on the generic design proposed in [36] but reverse the bottom-up masked attention into a top-down counterpart. In this transformer architecture, earlier attention blocks are allowed to attend to a larger spatial context of instances in order to be able to already predict potentially large offsets if necessary, while later blocks are confined to a narrower attention mask, as they have already attended to

more distant parts of the scene in earlier stages.

Second, we propose to exploit the input language description at a finer level of granularity than the previously considered global sentence level by supervising the predicted word-level *span* of the ground-truth referred object in the description, which is produced by the final verbal embeddings of the verbo-visual fusion module. These spans are normally available in standard 3D visual grounding datasets [2], but they have been largely overlooked by virtually all existing works when designing their models. By applying our proposed span loss, we can precisely attribute the gradient of each of the terms of the loss back to the embedding of the corresponding word, treating words inside the span differently from words outside it. This makes our word embeddings aware of the verbally described characteristics of the referred object. To allow this verbal awareness to be combined with visual awareness and thus produce word embeddings that are discriminative for the referred instance, we additionally enable cross-attention in the direction from the visual branch to the verbal branch in the attention-based building blocks of our verbo-visual fusion module, in a bidirectional attention formulation similar to that in [18]. Put together, our two novel losses along with top-down bidirectional attentive fusion form our complete AsphaltNet architecture for 3D visual grounding. AsphaltNet is evaluated on the Nr3D and Sr3D benchmarks and compared to the state-of-the-art approaches.

## 2. Related Work

3D visual grounding is one of the core problems that lies in the intersection between computer vision and language, aiming to localize objects in 3D space that are referred to by natural language prompts. To date, its 2D counterpart has been extensively studied [15, 20, 24, 28]. In the literature, 3D visual grounding is explored in two ways: (i) by purely focusing on the verbo-visual fusion, where given all available ground truth objects a selection task is carried out to identify the best fitting candidate, i.e. referral-based object identification [2]; or (ii) in an end-to-end fashion, where given a language utterance, a referred object is localized within the 3D scene via a bounding box or instance mask, i.e. referral-based object localization [5]. While seemingly different, a common strategy when tackling referral-based object localization is to employ a grounding-by-selection approach, where after a visual backbone produces candidate objects, the task boils down to an identification problem [5, 36, 43]. In this work, we undertake referral-based object identification to focus on improving the training process for multi-modal training.

The seminal works of Chen *et al.* [5] and Achlioptas *et al.* [1] provided the first 3D visual grounding benchmarks of ScanRefer, Nr3D and Sr3D by providing grounding annotations for the popular indoor RGB-D dataset ScanNet [9]. In

ScanRefer, a baseline model was introduced that tackled the fusion of language and vision in a rudimentary way, where each modality is individually encoded to later be joined via concatenation and processed via a shallow MLP.

ConcreteNet [36] improved 3DVG-Transformer by (i) by utilizing bottom-up attentive fusion module that localized mask attention in a progressively increasing neighborhood, (ii) by implicitly promoting separation between instance candidates via a cross-modal contrastive loss, and (iii) by tackling view-dependent prompts via a novel learned global camera token. In our work, we utilize a similar mask attention block that progressively scales in size, but do so in a top-down manner. Furthermore, instead of a multiplex attention layer, we employ bidirectional cross-attention to not only pass information from the natural language description to the 3D scene, but vice versa.

3D grounding solely from 2D image inputs was addressed in [29] which a scene-graph based 3D grounding was explored in [12]. Joint multi-task architectures for 3D dense captioning and 3D visual grounding were presented in [4], in [6]. While most 3D grounding approaches decouple object proposal generation from verbo-visual fusion, [23] proposes a single-stage approach which does not involve vision-only prediction of object proposals but performs mid-level fusion of language with 3D visual features.

InstanceRefer [42] constructs attention based verbo-visual fusion between global sentence embedding and instance embeddings, and tackled dense 3D visual grounding, i.e. referral-based 3D instance segmentation. TransRefer3D [13] proposes an entity-aware attention module and a relation-aware attention module to conduct fine-grained cross-modal feature matching. LanguageRefer [32] uses combines spatial embeddings extracted from provided ground truth bounding boxes with fine-tuned language embeddings from DistilBert [33] to identify the referred object. Furthermore, they analyze view-dependent utterances and viewpoint annotations. SAT [41] utilizes 2D image semantics in the training stage to ease verbo-visual joint representation learning and assist 3D visual grounding. LAR [3] tackles a similar problem of assisting 3D feature extraction by incorporating 2D object information. However in contrast to SAT, LAR does not rely on the existence of available 2D data, but synthetically generates 2D clues from 3D point clouds. MVT [17] repeatedly augments the 3D scene to generate a multiple view input and simultaneously processes and then aggregates the results. MDETR [19] is a language grounding model for 2D images that decodes object boxes and aligns them with relevant spans in input utterances. BUTD-DETR [18] builds upon MDETR but relies on a 3D object detector to generate referential grounding results, which are then grounded in the visual scene, augmenting supervision for the referential grounding. BUTD-DETR [20], builds upon MDETR [19], utiliz-

ing a 3D object detector for referential grounding, leveraging bidirectional attention layers to integrate visual and verbal inputs. Similar to BUTD-DETR, our work employs bidirectional cross-attention to allow information flow between visual tokens and verbal tokens. However, we introduce a top-down spherical masking technique to localize attention, improving grounding performance in referrals that require distinguishing objects based on neighboring relations. ScanEnts3D [1] advances the field of referential grounding by providing ground truth anchors, enabling the study of the effects of anchors and distractors for 3D grounding. 3DVista [44] aims to tackle 3D visual grounding by utilizing aggregation of datasets and tasks with a unified pre-trained transformer for 3D visual and text alignment.

### 3. Method

In this section we introduce our model AsphaltNet. First, we individually encode the visual and verbal cues (Sec.3.1 and Sec.3.2 respectively), which form the input for the verbo-visual fusion module that processes the two modalities to select the referred instance (Sec. 3.3). Later we explore two auxiliary losses to aid the training of the model (Sec 3.4). An illustration can be found in Fig. 2.

#### 3.1. Encoding Visual Cues

Typically, segmentation models operate directly on available dense information, which can be effectively encoded to extract high-fidelity scene-level features and decoded to generate dense predictions [30, 31, 35, 37]. However, transformer-based verbo-visual fusion modules are heavily constrained by the number of the input tokens due to the considerable memory and computing requirements of self- and cross-attention operations [11, 38]. To this end, a common solution is to construct a sparse representation of the scene that can capture relevant information for visual grounding [18, 36, 43].

In referral-based 3D object identification, the visual input is given by a point cloud  $P \in \mathbb{R}^{N \times 6}$  with  $N$  points, each described by its 3D coordinates  $(x, y, z)$  and color values  $(r, g, b)$ , as well as the ground truth instance masks  $I \in \{0, 1\}^{K \times N}$  with  $K$  instances. Following set precedent [18, 36, 43], to avoid directly operating on the large point cloud data during verbo-visual fusion, we construct instance embeddings using the available mask information to act as visual tokens. Specifically, we extract high-level features from the input point cloud via a UNet backbone and mean pool pointwise features for each instance mask. To preserve spatial and color information, we concatenate the center coordinates of the axis aligned bounding box along with the mean color vector with the pooled UNet features.

Formally, the instance embeddings  $e_i$  that form the vi-

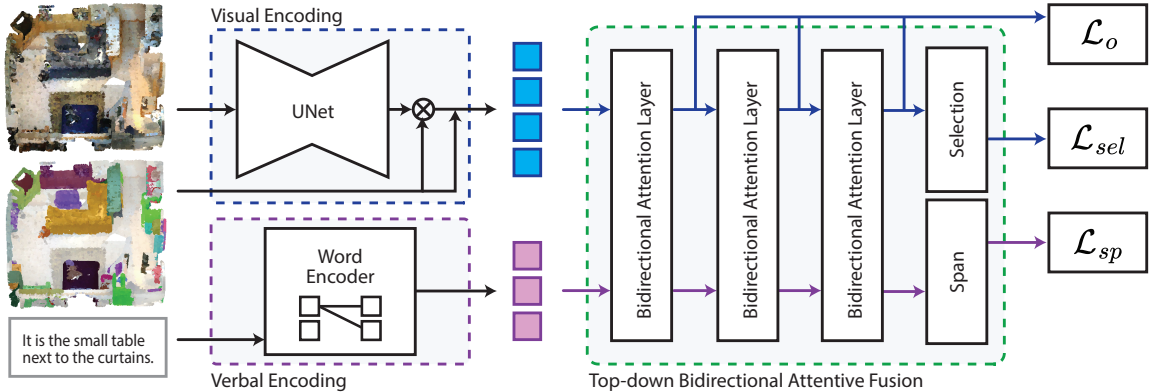


Figure 2. Illustration of the AsphaltNet pipeline. Via a UNet, we first compute per point features from the 3D point cloud and construct instance tokens via mean pooling on the provided instance masks and concatenation with the corresponding bounding box centroid and mean color vector. We encode the verbal input and generate word tokens using a pretrained LLM. A top-down bidirectional attentive fusion module consumes the instance and word tokens to ground the natural language description within the scene, i.e. predict the referred object. To aid the training process, we employ two auxiliary losses:  $\mathcal{L}_o$  that aids localization, and  $\mathcal{L}_{sp}$  to supervise the language encoding.

sual input for the verbo-visual fusion are given by:

$$\begin{aligned} e_i &= [f_i, c, a] \in \mathbb{R}^{d+6}, \text{ with} \\ f_i &= (I \times \phi_{UNet}(P)) \oslash (I \times [1, 1, 1, \dots]^T) \end{aligned} \quad (1)$$

$c = (c_x, c_y, c_z)$  and  $a = (r, g, b)$  denoting the center and mean color of the instance candidate respectively,  $\phi_{UNet}$  denoting the UNet encoding the produces a  $d$  dimensional output and  $\oslash$  the Hadamard division.

### 3.2. Encoding Verbal Cues

Large pretrained transformers have risen in popularity in recent years following their wide success in natural language processing [22, 34, 38]. This success has carried over to computer vision applications that require or benefit from verbal cues. Given that such models are trained with large-scale datasets they are capable of capturing context and intent within natural language prompts. While most early work in 3D visual grounding utilize a more traditional approach in text encoding, namely GloVe [27], recent efforts have moved on to transformer-based encoders such as BERT [10], RoBERTa [22] or MPNet [34].

We employ the pretrained BERT to initially encode word tokens. The features are then projected using a  $1 \times 1$  convolutional layer to match the dimension of the instance candidates  $e_w \in \mathbb{R}^{|W| \times (d+6)}$  which form the verbal input tokens of our fusion module.

### 3.3. Verbo-visual Fusion

In referral-based 3D object localization, verbal and visual features are processed, to output a likelihood vector over the input candidates to ground the natural language prompt. In the current literature, a common strategy to fuse

the two modalities is through the employment of attention:

$$g = \text{softmax}(qk^T)v \quad (2)$$

with residual attention features  $g$ , the queries  $q$  extracted from instance embeddings  $e_i$  and key-value pairs  $k$  and  $v$  extracted from word features  $e_w$  for cross-attention or from instance features  $e_i$  for self-attention. The cross-attention allows information extracted from the pretrained verbal model to be routed to each instance candidate. Subsequent self-attention layers that are formed between instance candidates routes this information, allowing the tokens to extract further relevant features for visual grounding.

**Top-down Bidirectional Attentive Fusion (TBA):** Often natural language prompts localize a repetitive object, i.e. an object that is semantically not unique within a scene, by establishing its relation to another object. Such relational cues are commonly constrained within local neighborhoods, e.g. “Chair next to the door”, or “The ottoman in front of the sofa”. While second nature to us, such localized information routing is difficult to learn for transformer-based architectures that rely on global information routing schemes as found in self-attention layers. To combat this, inspired by Mask2former [8], ConcreteNet [36] proposes using masked self-attention layers to explicitly limit information routing within local neighborhoods. Formally, Eq. 2 is restated as:

$$\begin{aligned} g &= \text{softmax}(M + qk^T)v, \text{ with} \\ M(i, j) &= \begin{cases} 0, & \text{if } \|c_i - c_j\| < r \\ -\infty, & \text{otherwise} \end{cases} \end{aligned} \quad (3)$$

with  $r$  denoting the radius of the spherical masking set per layer. Similarly, we construct a transformer-based verbo-visual fusion module by iteratively stacking masked self-attention layers followed by cross attention layers.

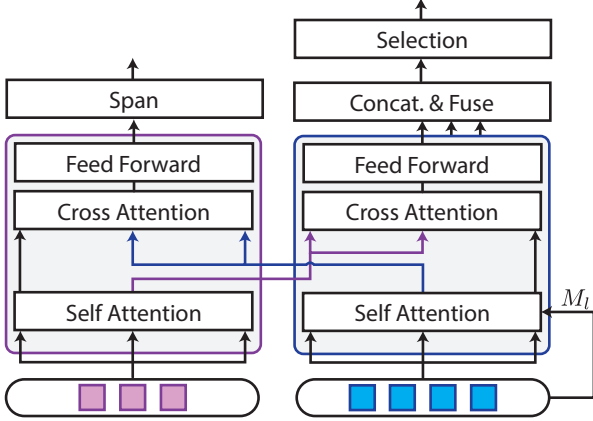


Figure 3. Illustration of top-down bidirectional attentive fusion.

Yet, our fusion module differs from that of ConcreteNet [36] in two key aspects: (i) we utilize bidirectional cross-attention, i.e. cross-attention not only to route information from verbal cues to the visual, but also to allow the verbal cues to learn about the 3D space. We argue that while the features extracted from the pretrained verbal model are rich in information, the instance features that describe the 3D scene can provide strong priors to the word tokens and help to resolve view-dependent relations. Bidirectional cross-attention has previously shown to strongly aid 3D grounding performance [18] but was never utilized in a local attention setting; (ii) next, we employ a top-down approach when masking compared to a bottom-up approach, i.e. after each bidirectional attention block, starting from global attention we reduce the radius  $r$  to further limit the spherical neighborhood. This is later discussed in Sec. 3.4 - Offset Loss. An illustration of TBA can be seen in Fig. 3.

**Selection Loss:** Following the TBA, a  $1 \times 1$  convolutional layer maps each embedding  $e_i$  to a confidence score in order to form the prediction  $u$ . The prediction is supervised by the ground truth target  $\hat{u}$  via a cross-entropy loss:

$$\mathcal{L}_{sel} = H(u, \hat{u}). \quad (4)$$

### 3.4. Auxiliary Losses for 3D Visual Grounding

The literature is currently focused on exploring new architectural setups to better ground natural language prompts to 3D scenes [18, 41, 43]. The respective training objectives remain highly under-explored. Similarly as Sec. 3.3, current methods are supervised only via a limited pool of losses such as (i) cross-entropy that aims to either maximize the softmax probability of the ground-truth object when applied on the instance tokens, or (ii) identify the semantic class of the target object when applied on the word tokens; and (iii) contrastive loss that aims to map matching word and instance tokens to close representations. We argue that while these losses have shown promising results, they remain in-

sufficient as they do not explicitly promote localization of the prompt within the 3D scene or simply remain too coarse for informative supervision.

**Offset Loss ( $\mathcal{L}_o$ ):** To remedy this, inspired by common practices in 3D instance segmentation [26, 40], we propose to define an offset loss at the level of objects, where each object in the scene is required to predict its offset to the referred object, in order to better ground the natural language in 3D space. Formally, we define the offset loss as:

$$\mathcal{L}_o = \sum_{i=1}^{|I|} L_2(o_i, (c_i - c_{gt})) \quad (5)$$

with  $o_i$ ,  $c_i$  denoting the predicted offset value and centroid for instance  $i$  and  $c_{gt}$  the referred object centroid.

In essence, in order for each instance to be able to predict the offset to the referred object, grounding-aware features are learned for all instances in the self-attention layers. With such an auxiliary loss, we not only achieve better localization through direct supervision, but also implicitly promote separation between the features of the referred object and those of spatially distant repetitive objects which may bear semantic similarity to the former.

We employ the offset loss after every bidirectional attentive fusion block to encourage the network the extract grounding-aware instance features at varying spherical neighborhoods. It is for this reason that a top-down approach is required, as in a bottom-up setting, where masking is carried out in a local-then-global fashion, early layers do not allow information routing across the entire scene, which limits the ability of an instance candidate to predict an accurate offset to a distant referred object. In Fig. 4 we showcase example offset predictions from AsphaltNet.

**Span Loss ( $\mathcal{L}_{sp}$ ):** To better exploit the input verbal utterances at a finer level of granularity than the aforementioned considerations that supervise the model at a sentence level through cross-entropy, we propose predicting and supervising on a word-level through the *span* of a description. The span  $\hat{l} \in \{0, 1\}^{|W|}$  is a vector with entry 1 for where the word or words describe the referred object, and 0 otherwise. Formally we define the span loss as:

$$\mathcal{L}_{sp} = H^1(l, \hat{l}) \quad (6)$$

with  $H^1$  denoting the binary cross entropy loss,  $l$  the ground truth span labels and  $\hat{l}$  the predicted span logits.

The span loss allows precisely gradient flow from each loss term back to the corresponding word embedding. Each words inside the span is treated differently from words outside it, increasing awareness towards the verbally described characteristics of the referred object.




Figure 4. Sample offset predictions of AsphaltNet from the Nr3D [2] *val*-set given the prompt “Facing the armchairs, the chair in the upper left”. We overlay the heatmap of the predicted offset values per instance to the target on top of the scene for all the three layers of the fusion module. The offset prediction is refined at each step as the distribution of predicted offsets tightens and converges on the target instance.

	Method	Input	Overall	Easy	Hard	V-Dep.	V-Ind.
Nr3D	ReferIt3D [2]	3D	35.6	43.6	27.9	32.5	37.1
	3DVG-Transformer [43]	3D	40.8	48.5	34.8	34.8	43.7
	LanguageRefer [32]	3D	43.9	51.0	36.6	41.7	45.0
	ConcreteNet [36]	3D	45.5	52.0	38.7	38.9	48.3
	SAT [41]	2D+3D	49.2	56.3	42.4	46.9	50.4
	BUTD-DETR [18]	3D	54.6	60.7	48.4	46.0	58.0
	MVT [17]	MV-3D	55.1	61.3	49.1	54.3	55.4
	3DVista (scratch)	3D	57.5	65.9	49.4	53.7	59.4
	AsphaltNet	3D	<b>58.9</b>	<b>64.3</b>	<b>53.9</b>	<b>57.4</b>	<b>59.8</b>
	ScanEnts [1]	3D	59.3	65.4	53.5	57.3	60.4
3DVista [44]	3D*	64.2	72.1	56.7	61.5	65.1	
Vil3DRel [7]	3D	64.4	70.2	57.4	62.0	64.5	
Sr3D	ReferIt3D [2]	3D	40.8	44.7	31.5	39.2	40.8
	3DVG-Transformer [43]	3D	51.4	54.2	44.9	44.6	51.7
	LanguageRefer [32]	3D	56.0	58.9	49.3	49.2	56.3
	SAT [41]	2D+3D	57.9	61.2	50.0	49.2	58.3
	NS3D [14]	3D	62.7	64.0	59.6	62.0	62.7
	MVT [17]	MV-3D	64.5	66.9	58.8	58.4	64.7
	BUTD-DETR [18]	3D	67.0	68.6	63.2	53.0	67.6
	3DVista (scratch)	3D	69.6	72.1	63.6	57.9	70.14
	AsphaltNet	3D	<b>69.7</b>	<b>71.9</b>	<b>64.5</b>	<b>67.2</b>	<b>70.0</b>
	Vil3DRel [7]	3D	72.8	74.9	67.9	63.8	73.2
3DVista [44]	3D*	76.4	78.8	71.3	58.9	77.3	

Table 1. 3D visual grounding results for referral-based object identification on the Nr3D and Sr3D datasets [2]. V-Dep. and V-Ind. refer to the view-dependent and independent subcategories. MV refers to multi-view. 3D\* refers to the use of additional 3D scenes. ScanEnts [1] is evaluated along with GT anchor proposals. Vil3Dref has a reduced evaluation protocol, excluding challenging longer prompts.

## 4. Experiments

### 4.1. Implementation details

**Implementation details:** For the visual encoder, we use the 3D Unet backbone of ISBNet [26] pretrained on ScanNet and frozen. We employ standard augmentation strategies such as rotation, translation and cropping. For the verbal backbone we use a pretrained frozen MPNet [34]. Following Zhao *et al.* [43], we mask the referred object nouns with a probability of 0.2. We also employ text augmentation by randomly replacing word with their synonyms with a prob-

ability of 0.3 [25]. For the fusion module, we use a dimension of 128 across all layers, using 2 transformer encoder layers for the language abstraction with sinusoidal positional encoding. TBA consists of 3 bidirectional attention layers, each with a set spherical radius of  $[\infty, 2.5m, 1m]$  following the inverted scheme of ConcreteNet [36]. We use a batch size of 1 with each batch consisting of up to 64 utterances. This allows us to reduce the total memory requirements while retaining clean gradients for grounding. We train for 300 epochs using AdamW with a learning rate of 0.0001 on a single Nvidia Geforce RTX 3090.

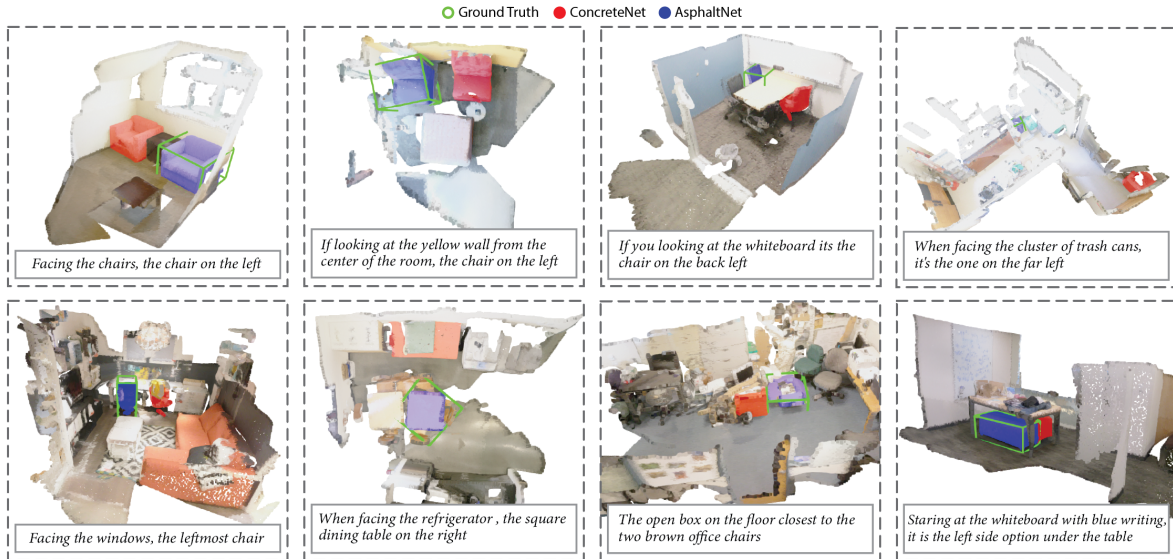


Figure 5. Qualitative results on the Nr3D *val*-set comparing AsphaltNet with the baseline ConcreteNet model. We illustrate the predicted instances via instance masks (red) for ConcreteNet, and (blue) for AsphaltNet. In green, we provide the axis aligned bounding box of the referred ground truth object.

$\mathcal{L}_o$	TBA	$\mathcal{L}_{sp}$	Overall	Easy	Hard	V-Dep	V-Ind
			45.5	52.0	38.7	38.9	48.3
✓			51.1	57.9	45.6	50.4	51.5
✓	✓		55.2	60.2	50.5	54.6	55.4
✓	✓	✓	<b>58.9</b>	<b>64.3</b>	<b>53.9</b>	<b>57.4</b>	<b>59.8</b>

Table 2. Ablation study showing the benefits on each introduced contribution. Starting with a baseline model using a top-down BAF module [36], we introduce our offset loss ( $\mathcal{L}_o$ ), top-down bidirectional attentive fusion module (TBA), and the span loss ( $\mathcal{L}_{sp}$ ). We report results on the NR3D *val*-set [%].

## 4.2. Results

In Tab. 1 we report the performance of our proposed AsphaltNet on the Nr3D and Sr3D *val*-sets and compare them to state-of-the-art methods. As seen, AsphaltNet is competitive with the existing work on both datasets, without relying on added supervision from the introduction of handcrafted detection prompts like BUTD-DETR (which account for +4.2% overall accuracy on Nr3D), or on boosting performance by averaging the multiple predictions from view augmented forward passes like MVT. 3DVista proposes a pretrained foundational model for 3D and text alignment, utilizing a rich and diverse dataset. We compare our results with 3DVista(scratch) which follows the same dataset and evaluation protocol as AsphaltNet, outperforming on both Nr3D (+1.4) and SR3D datasets.

## 4.3. Ablation Studies

**Effects of Network Components:** In Tab. 2, we perform an ablation study to investigate the effects of the different com-

Method	Overall	Failure Cases
Baseline	1.33	2.54
+ $\mathcal{L}_o$	0.93	2.02

Table 3. Mean L2 distance of the predicted instance mask center to the referred instance mask center with and without the offset loss [m]. We report the distance across the validation-set (overall) along with only on failure cases where the model incorrectly predicts the referred object.

Verbal Backbone	Overall Acc
BERT [10]	<b>58.9</b>
ROBERTA [22]	58.8
MPNet [34]	56.1

Table 4. Ablation Study showing the effect of changing the verbal backbone

Method	Overall	Easy	Hard	V-Dep	V-Ind
with $\mathcal{L}_{cls}$	56.3	62.6	50.6	55.5	56.9
with $\mathcal{L}_{sp}$	<b>58.9</b>	<b>64.3</b>	<b>53.9</b>	<b>57.4</b>	<b>59.8</b>

Table 5. Comparison of the standard cross-entropy based sentence-level supervision ( $\mathcal{L}_{cls}$ ) to our word-level span loss  $\mathcal{L}_{sp}$ .

ponents of our proposed AsphaltNet. We set up a baseline model using a top-down attentive fusion module from ConcreteNet [36]. We then systematically introduce our offset loss ( $\mathcal{L}_o$ ), the top-down bidirectional attention block (TBA), and finally the span loss ( $\mathcal{L}_{sp}$ ). As observed, the introduction of the offset loss drastically improves *hard* grounding

	Weight	Acc	Easy	Hard	V-Dep	V-Ind
$\mathcal{L}_o$	0.1	55.7	61.1	49.8	53.8	56.8
	1	<b>58.9</b>	64.3	53.9	57.4	59.8
	10	49.8	55.3	45.1	49.0	50.8
$\mathcal{L}_{sp}$	0.1	55.8	61.1	50.9	54.6	56.4
	1	<b>58.9</b>	64.3	53.9	57.4	59.8
	10	56.8	61.4	52.4	55.7	57.4

Table 6. Hyperparameter study on the effect of varying weights of the offset and span losses on 3D visual grounding performance.

performance by +6.9%. As mentioned above, the *hard* category consists of referred objects with more than 3 same class distractors within the same scene, which illustrates the effectiveness of the offset loss in candidate separation as a byproduct of grounding-aware feature routing.

Following the offset loss, we introduce the bidirectional attention layers into the top-down fusion module, allowing the word tokens to get updated based on instance-level features. As seen, the overall accuracy further improves by +4.1%. Unlike 2D images, since 3D scenes don’t inherently possess an orientation, view-dependent prompts in Nr3D are often constructed through object references, e.g. “If you looking at the whiteboard it’s the chair on the back left.” Such cases can be ill-posed when language tokens are unable to obtain information from the 3D scene. The inclusion of the top-down bidirectional attentive fusion block especially aids the *view-dependent* accuracy, promoting scene-aware word encoding.

Finally, we introduce supervision to the word tokens via the span loss ( $\mathcal{L}_{sp}$ ). As seen in Tab. 2, the loss shows benefits across all categories, with the overall accuracy increasing by +3.7. With all three components, AsphaltNet reaches 58.9% overall accuracy on Nr3D dataset.

**Impact of  $\mathcal{L}_o$  on Localization:** In Tab. 3 we conduct an ablation study to investigate the effect of the offset loss on referral-based object localization. Given the baseline model, we compute the average L2 distance of each model prediction to the referred target object with and without the inclusion of the offset loss during training. Furthermore, we report the same metric over only failure cases to eliminate the bias stemming from the increased number of accuracy (i.e. 0 distance) samples. As seen, in both overall and failure cases, the inclusion of the offset loss reduces the average distance of the predicted object to the target object. In other words, the offset loss enables better localization within the network, promoting predictions nearer to the target object even during failure cases.

Additionally in Fig. 4 we showcase a qualitative example of the evolution of the offset predictions within AsphaltNet. Specifically on each layer, for each instance, we obtain the predicted target object centroid and fit a 2D kernel density function [39] to the obtained set. As expected, the distribution of the centroid estimations converges on the referred

object with increased precision, i.e. decreased uncertainty.

**Impact of the Verbal Backbone:** In Tab. 4 we provide an ablation study with different verbal encoding choices. We show that our model is robust against this choice, with BERT [10] and RoBERTa [22] achieving 58.9% and 58.8% overall accuracy. We observe a decrease in performance when employing MPNet which has seen previous success in 3D visual grounding models [36]. We account this drop in performance towards the permuted language modeling of MPNet, which fails to capture the relational contexts in more structured ReferIt3D prompts adequately.

**Auxiliary Supervision on Language:** In Tab. 5 we compare our proposed span loss  $\mathcal{L}_{sp}$  to the standard approach of constructing an auxiliary classification problem using the language tokens and supervising it via a cross-entropy  $\mathcal{L}_{cls}$  [5, 43]. As seen, providing a higher granularity in the language supervision aids the model performance, boasting an additional +2.6% overall accuracy over the cross-entropy formulation, with the majority of the benefits coming from the *hard* category where utilization of language becomes critical to distinguish the referred object between a large number of semantically similar candidates.

**Hyperparameter Study:** Having introduced two additional losses to aid 3D referral-based object identification, we explore the effect of loss weights on the models performance. In Tab. 6, we train AsphaltNet with varying weights on either the offset loss  $\mathcal{L}_o$  or the span loss  $\mathcal{L}_{sp}$ . While the model is quite robust to the weight changes of the losses – showing comparable performance to the second-best performing method despite an order of magnitude change in any loss weight – the best configuration is obtained by setting each weight to 1.

## 5. Conclusion

In this work we tackle the problem of 3D visual grounding, where the goal is to identify the correct instance based on a natural language prompt. We note that while continual efforts are put into designing better architectures, the exploration of losses remains stagnant. This work proposes two auxiliary losses to improve grounding performance and increase training robustness. Firstly we construct an offset loss to aid localization where each object aims to predict its offset to the referred target. Next, we extend our baseline method with bidirectional attention to allow information flow from the 3D scene to the language branch. This allows us to construct a span loss, providing highly granular word-level supervision to the verbal encoding. Combining all three components, we construct AsphaltNet, a novel 3D visual grounding approach with impressive performance on ReferIt3D without relying on any additional data.

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



## References

- [1] Ahmed Abdelreheem, Kyle Olszewski, Hsin-Ying Lee, Peter Wonka, and Panos Achlioptas. Scanents3d: Exploiting phrase-to-3d-object correspondences for improved visiolinguistic models in 3d scenes. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3524–3534, 2024. 3, 6, 11
- [2] Panos Achlioptas, Ahmed Abdelreheem, Fei Xia, Mohamed Elhoseiny, and Leonidas J. Guibas. ReferIt3D: Neural listeners for fine-grained 3d object identification in real-world scenes. In *16th European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 6, 11, 12
- [3] Eslam Bakr, Yasmeeen Alsaedy, and Mohamed Elhoseiny. Look around and refer: 2d synthetic semantics knowledge distillation for 3d visual grounding. In *Advances in Neural Information Processing Systems*, 2022. 3
- [4] Daigang Cai, Lichen Zhao, Jing Zhang, Lu Sheng, and Dong Xu. 3DJCG: A unified framework for joint dense captioning and visual grounding on 3D point clouds. In *CVPR*, 2022. 3
- [5] Dave Zhenyu Chen, Angel X Chang, and Matthias Nießner. Scanrefer: 3d object localization in rgb-d scans using natural language. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 8, 11
- [6] Dave Zhenyu Chen, Qirui Wu, Matthias Nießner, and Angel X Chang. D3Net: a speaker-listener architecture for semi-supervised dense captioning and visual grounding in RGB-D scans. In *ECCV*, 2022. 3
- [7] Shizhe Chen, Pierre-Louis Guhur, Makarand Tapaswi, Cordelia Schmid, and Ivan Laptev. Language conditioned spatial relation reasoning for 3d object grounding. *Advances in neural information processing systems*, 35:20522–20535, 2022. 6
- [8] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022. 4
- [9] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2017. 2, 11
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 4, 7, 8
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3
- [12] Mingtao Feng, Zhen Li, Qi Li, Liang Zhang, XiangDong Zhang, Guangming Zhu, Hui Zhang, Yaonan Wang, and Ajmal Mian. Free-form description guided 3D visual graph network for object grounding in point cloud. In *ICCV*, 2021. 3
- [13] Dailan He, Yusheng Zhao, Junyu Luo, Tianrui Hui, Shaofei Huang, Aixi Zhang, and Si Liu. Transrefer3d: Entity-and-relation aware transformer for fine-grained 3d visual grounding. In *Proceedings of the 29th ACM International Conference on Multimedia*, MM '21. ACM, Oct. 2021. 3
- [14] Joy Hsu, Jiayuan Mao, and Jiajun Wu. Ns3d: Neuro-symbolic grounding of 3d objects and relations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 6
- [15] Ronghang Hu, Huazhe Xu, Marcus Rohrbach, Jiashi Feng, Kate Saenko, and Trevor Darrell. Natural language object retrieval. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4555–4564, 2016. 2
- [16] Pin-Hao Huang, Han-Hung Lee, Hwann-Tzong Chen, and Tyng-Luh Liu. Text-guided graph neural networks for referring 3D instance segmentation. In *AAAI*, 2021. 1
- [17] Shijia Huang, Yilun Chen, Jiaya Jia, and Liwei Wang. Multi-view transformer for 3D visual grounding. In *CVPR*, 2022. 2, 3, 6, 11
- [18] Ayush Jain, Nikolaos Gkanatsios, Ishita Mediratta, and Kateirina Fragkiadaki. Bottom up top down detection transformers for language grounding in images and point clouds. In *ECCV*, 2022. 1, 2, 3, 5, 6
- [19] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetrmultimodal detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 3
- [20] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. ReferItGame: Referring to objects in photographs of natural scenes. In Alessandro Moschitti, Bo Pang, and Walter Daelemans, editors, *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Oct. 2014. 2
- [21] Daizong Liu, Yang Liu, Wencan Huang, and Wei Hu. A survey on text-guided 3d visual grounding: Elements, recent advances, and future directions. *arXiv preprint arXiv:2406.05785*, 2024. 1
- [22] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 4, 7, 8
- [23] Junyu Luo, Jiahui Fu, Xianghao Kong, Chen Gao, Haibing Ren, Hao Shen, Huaxia Xia, and Si Liu. 3D-SPS: Single-stage 3D visual grounding via referred point progressive selection. In *CVPR*, 2022. 3
- [24] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20, 2016. 2
- [25] Vukosi Marivate and Tshephisho Sefara. Improving short text classification through global augmentation methods. In *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*, 2020. 6

- [26] Tuan Duc Ngo, Binh-Son Hua, and Khoi Nguyen. Isbnet: a 3d point cloud instance segmentation network with instance-aware sampling and box-aware dynamic convolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 5, 6
- [27] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *EMNLP*, volume 14, 2014. 4
- [28] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015. 2
- [29] Mihir Prabhudesai, Hsiao-Yu Fish Tung, Syed Ashar Javed, Maximilian Sieb, Adam W. Harley, and Katerina Fragkiadaki. Embodied language grounding with 3D visual feature representations. In *CVPR*, 2020. 3
- [30] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [31] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in Neural Information Processing Systems*, 2017. 3
- [32] Junha Roh, Karthik Desingh, Ali Farhadi, and Dieter Fox. LanguageRefer: Spatial-language model for 3D visual grounding. In *CoRL*, 2022. 1, 3, 6
- [33] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019. 3
- [34] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. MpNet: Masked and permuted pre-training for language understanding. *Advances in Neural Information Processing Systems*, 2020. 4, 6, 7
- [35] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019. 3
- [36] Ozan Unal, Christos Sakaridis, Suman Saha, and Luc Van Gool. Four ways to improve verbo-visual fusion for dense 3d visual grounding. In *European Conference on Computer Vision (ECCV)*, 2024. 2, 3, 4, 5, 6, 7, 8, 11
- [37] Ozan Unal, Luc Van Gool, and Dengxin Dai. Improving point cloud semantic segmentation by learning 3d object detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021. 3
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 3, 4
- [39] M. P. Wand and M. C. Jones. *Kernel Smoothing*. Chapman and Hall/CRC, 1st edition, 1994. 8
- [40] Yizheng Wu, Min Shi, Shuaiyuan Du, Hao Lu, Zhiguo Cao, and Weicai Zhong. 3d instances as 1d kernels. In *ECCV*, 2022. 5
- [41] Zhengyuan Yang, Songyang Zhang, Liwei Wang, and Jiebo Luo. SAT: 2D semantics assisted training for 3D visual grounding. In *ICCV*, 2021. 3, 5, 6
- [42] Zhihao Yuan, Xu Yan, Yinghong Liao, Ruimao Zhang, Sheng Wang, Zhen Li, and Shuguang Cui. Instancerefer: Cooperative holistic understanding for visual grounding on point clouds through instance multi-level contextual referring. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 1, 2, 3
- [43] Lichen Zhao, Daigang Cai, Lu Sheng, and Dong Xu. 3DVG-Transformer: Relation modeling for visual grounding on point clouds. In *ICCV*, 2021. 1, 2, 3, 5, 6, 8
- [44] Zhu Ziyu, Ma Xiaojian, Chen Yixin, Deng Zhidong, Huang Siyuan, and Li Qing. 3d-vista: Pre-trained transformer for 3d vision and text alignment. In *ICCV*, 2023. 3, 6




Figure 6. Plot visualizing the overall accuracy against the length of the natural language description #tokens.

Method	Overall
Top-down	58.9
Bottom-up	56.3

Table 7. Comparison of top-down and bottom-up attention masking for the verbo-visual fusion module

## A. Dataset

We evaluate our method on the Nr3D and Sr3D datasets [2] that provide human-annotated and synthetically generated utterances to identify instances on ScanNet [9]. Specifically, Nr3D consists of 41,503 natural language queries and Sr3D consists of 83,572 descriptions over 707 unique indoor scenes referring to 76 object classes. For evaluation, the dataset is split in two ways. First split considers the difficulty of the grounding task, with referred objects that have only 2 distractors considered *easy*, and objects that have up to 6 distractors are considered *hard*. Second form of splitting considered view-dependency. Based on distinct word mining, utterances are classified as either view-dependent or view-independent.

## B. Results and Evaluations

### B.1. Multiple Objects

Description referring to multiple objects is a compelling research problem that the ReferIt3D benchmark doesn’t address. We carry out further analysis of our approach using the ScanEnts dataset [1]. to evaluate the network performance with the number of anchors in the utterance 6. A drop in performance with more objects in the utterance is observed.

Method	Accuracy
Baseline	43.89
AsphaltNet	47.99

Table 8. Evaluation of OOD performance on ScanRefer

Method	Accuracy
Baseline	34.87
AsphaltNet	38.53

Table 9. Evaluation of OOD performance on NR3D

## C. Top-Down Masking

To explicitly limit information routing within local neighborhoods, ConcreteNet [36] employs a masking operation that hinders self-attention between spatially distant instance candidates. This masking operation is implemented in a bottom-up manner where the masking sphere grows as the layers increase. To be able to utilize the offset loss introduced in the main manuscript at various stages of the bidirectional attentive fusion module, we instead use a top-down approach, allowing all instances to be aware of one another even at the shallowest levels. In Tab. 7, we ablate this masking approach and show that a top-down method works better than a bottom-up one.

### C.1. Out of Distribution

We address the out-of-distribution ability of AsphaltNet. In table 8 we evaluate our network on the scanrefer dataset [5] with GT boxes provided in the input with network weights corresponding to NR3D dataset. The network trained without any of the contribution of AsphaltNet is used as baseline for the comparisons. We observe decent performance on the scanrefer dataset given the vast difference in the language domain of ReferIt3D and scanrefer datasets showcasing the OOD robustness of AsphaltNet. We further investigate the synthetic to real transferability of the network. AsphaltNet trained on SR3D dataset is tested on the NR3D dataset to test for synthetic - real domain adaptability. AsphaltNet yields relatively lower performance, table 9, when evaluated on NR3D dataset. This can be attributed to a lack of diversity in the synthetic dataset hence unable to transfer to the natural language settings.

### C.2. Multi-View Ensembling

MVT [17] proposes a multi-view feature aggregation approach that aggregates visual features from different viewpoints during both training and testing. While this acts as a form of regularization during training with the goal of smoothing the gradients, during test time, the multi-view aggregations act as a form of test-time-augmentation (TTA). AsphaltNet’s performance can be enhanced by utilizing TTA by employing multiple forward passes through affine-transformed point clouds. Specifically, we employ multi-view ensembling (MVE) following ConcreteNet [36] where we infer a scene  $N$  times, each with the point cloud rotated by an angle  $r \in [0, 2\pi]$ . The final prediction is de-

	Network	Overall	Easy	Hard	View-Dep	View-Ind
NR3D	AsphaltNet	58.9	64.3	53.9	57.4	59.8
	AsphaltNet + MVE	<b>59.3</b>	64.9	54.2	57.8	59.9
SR3D	AsphaltNet	69.7	71.9	64.5	67.2	70.0
	AsphaltNet + MVE	<b>72.7</b>	74.6	68.2	70.4	73.0

Table 10. Evaluation of AsphaltNet with multi-view ensembling (MVE)

terminated through majority voting across the  $N$  inferences (with  $N = 9$ ). As seen in table 10, MVE further boosts the performance of AsphaltNet for both NR3D and SR3D splits of ReferIt3D.




Figure 7. Plot visualizing the overall accuracy against the length of the natural language description #tokens.




Figure 8. Plot visualizing the overall accuracy against the total number of anchors in the description.

## D. Limitation and Discussion

While AsphaltNet achieves state-of-the-art results on both ReferIt3D [2] benchmarks, we still observe an important limitation when analyzing the predictions. In Fig. 7 we show the accuracy of our model on the Nr3D *val*-set against the number of tokens in the natural language prompt. As seen, we observe a linear drop in performance as the number of tokens increases. We speculate this is due to the increased difficulty of the span prediction task that limits our method’s generalization ability for longer text descriptions.

Furthermore, in Fig. 8, we show that while our method’s performance weakens as the number of tokens increases (verbal input), the same behavior is not observed when investigating the visual input changes, i.e. as the number of same-class distractors increases. This illustrates the benefits of the offset loss, where the increased robustness towards spatial localization reduces the negative effects of increased same-class instances within the scene.