

Bottom Up Top Down Detection Transformers for Language Grounding in Images and Point Clouds

Ayush Jain^{†1}, Nikolaos Gkanatsios^{†1}, Ishita Mediratta^{§2}, and Katerina Fragkiadaki¹

¹Carnegie Mellon University, ²Meta AI

Abstract. Most models tasked to ground referential utterances in 2D and 3D scenes learn to select the referred object from a pool of object proposals provided by a pre-trained detector. This is limiting because an utterance may refer to visual entities at various levels of granularity, such as the chair, the leg of the chair, or the tip of the front leg of the chair, which may be missed by the detector. We propose a language grounding model that attends on the referential utterance and on the object proposal pool computed from a pre-trained detector to decode referenced objects with a detection head, without selecting them from the pool. In this way, it is helped by powerful pre-trained object detectors without being restricted by their misses. We call our model Bottom Up Top Down DEtection TRansformers (BUTD-DETR) because it uses both language guidance (top down) and objectness guidance (bottom-up) to ground referential utterances in images and point clouds. Moreover, BUTD-DETR casts object detection as referential grounding and uses object labels as language prompts to be grounded in the visual scene, augmenting supervision for the referential grounding task in this way. The proposed model sets a new state-of-the-art across popular 3D language grounding benchmarks with significant performance gains over previous 3D approaches (12.6% on SR3D, 11.6% on NR3D and 6.3% on ScanRefer). When applied in 2D images, it performs on par with the previous state of the art. We ablate the design choices of our model and quantify their contribution to performance. Our code and checkpoints can be found at the project website <https://butd-detr.github.io>.

1 Introduction

Language-directed attention helps us localize objects that our “bottom-up”, task-agnostic perception may miss. Consider Fig. 1. The utterance “*bottle on top of the bathroom vanity*” suffices to direct our attention to the reference object, even though it is far from salient. Language-directed perception adapts the

[†]Equal contribution, order decided by `np.random.rand`

[§]Work done during an internship at CMU




Fig. 1: **Language-modulated 3D (top) and 2D (bottom) detection with BUTD-DETR.** *Middle:* State-of-the-art object detectors often fail to localize small, occluded or rare objects (here they miss the clock on the shelf and the bottle on the cabinet). *Right:* Language-driven and objectness-driven attention in BUTD-DETR modulates the visual processing depending on the referential expression while taking into account salient, bottom-up detected objects, and correctly localizes all referenced objects.

visual processing of the input scene according to the utterance. Object detectors instead apply the same computation in each scene, which can miss task-relevant objects.

Most existing language grounding models use object proposal bottlenecks: they select the referenced object from a pool of object proposals provided by the pre-trained object detector [9,20,22,11,17]. This means they cannot recover objects or parts that a bottom-up detector misses. This is limiting since small, occluded, or rare objects are hard to detect without task-driven guidance. For example, in Figure 1 middle, state-of-the-art 2D [39] and 3D [30] detectors miss the clock on the shelf and the bottle on the bathroom vanity, respectively.

Recently, Kamath et al. [21] introduced MDETR, a language grounding model for 2D images that decodes object boxes using a DETR [3] detection head and aligns them to the relevant spans in the input utterance, it does not select the answer from a box proposal pool. The visual computation is modulated based on the input utterance through several layers of self-attention on a concatenation of language and visual features. MDETR achieves big leaps in performance in 2D language grounding over previous box-bottlenecked methods.

We propose a model for grounding referential utterances in 3D and 2D visual scenes that builds upon MDETR, which we call BUTD-DETR (pronounced

Beauty-DETR), as it uses both box proposals, obtained by a pre-trained detector “bottom-up” and “top-down” guidance from the language utterance, to localize the relevant objects in the scene. BUTD-DETR uses box proposals obtained by a pre-trained detector as an additional input stream to attend on; however, it is not box-bottlenecked and still decodes objects with a detection head, instead of selecting them from the input box stream. Current object detectors provide a noisy tokenization of the input visual scene that, as our experiments show, is a useful cue to attend on for multimodal reasoning. Second, BUTD-DETR augments grounding annotations by configuring annotations for object detection as detection prompts to be grounded in visual scenes. A detection prompt is a list of object category labels, e.g., “*Chair. Door. Person. Bed.*”. We train the model to ground detection prompts by localizing the labels that are present in the image and learn to discard labels that are mentioned but do not correspond to any objects in the scene. Third, BUTD-DETR considers improved bounding box - word span alignment losses that reduce noise during alignment of object boxes to noun phrases in the referential utterance.

We test BUTD-DETR on the 3D benchmarks of [2,4] and 2D benchmarks of [23,47]. In 3D point clouds, we set new state-of-the-art in the two benchmarks of Referit3D [2] and ScanRefer [4] and report significant performance boosts over all prior methods (12.6% on SR3D, 11.6% on NR3D and 6.3% on ScanRefer), as well as over a direct MDETR-3D implementation of ours that does not use a box proposal stream or detection prompts during training. In 2D images, our model obtains competitive performance with MDETR on RefCOCO, RefCOCO+ and Flickr30k, and requires less than half of the GPU training time due to the cheaper deformable attention in the visual stream. We ablate each of the design choices of the model to quantify their contribution to performance.

In summary, our contributions are: (i) A model with SOTA performance across both 2D and 3D scenes with minor changes showing that modulated detection in 2D images can also work in 3D point clouds with appropriate visual encoder and decoder modifications. (ii) Augmenting supervision with detection prompts, attention on an additional input box stream and improved bounding box - word span alignment losses. (iii) Extensive ablations to quantify the contribution of different components of our model. We make our code publicly available at <https://butd-detr.github.io>.

2 Related work

Object detection with transformers Object detectors are trained to localize all instances of a closed set of object category labels in images and 3D point-clouds. While earlier architectures pool features within proposed boxes to decode objects and classify them into categories [14,28,38], recent methods pioneered by DETR [3] use transformer architectures where a set of object query vectors attend to the scene and among themselves to decode object boxes and their labels. DETR suffers from the quadratic cost of within image features self attention. D(eformable)-DETR [50] proposes deformable attention, a locally adaptive ker-

nel that is predicted directly in each pixel location without attention to other pixel locations, thus saving the quadratic cost of pixel-to-pixel attention. Our model builds upon deformable attention for feature extraction from RGB images. [30,34] extend detection transformers to 3D point cloud input.

2D referential language grounding Referential language grounding [23] is the task of localizing the object(s) referenced in a language utterance. Most 2D language grounding models obtain sets of object proposals using pre-trained object detectors and the original image is discarded upon extraction of the object proposals [9,20,22,11,17]. Many of these approaches use multiple layers of attention to fuse information across both, the extracted boxes and language utterance [31,6,46]. Recently, a few approaches directly regress the target bounding box without using pre-trained object proposals. In [5] language and visual features cross-attend and are concatenated to predict the box of the referential object. Yang et al. [45] extends the YOLO detector [38] to referential grounding by channel-wise concatenating language, visual and spatial feature maps and then regressing a single box using the YOLO box prediction head. [42] performs a fusion similar to [45], then selects a single box from a set of anchor boxes and predicts a deformation of it, much like the Faster-RCNN object detector [39]. While previous approaches encode the whole text input into a single feature vector, [44] further improves performance by recursively attending on different parts of the referential utterance. Lastly, [8] encodes the image and utterance with within- and cross-modality transformers, and a special learnable token regresses a single box. In contrast to our method, all these works predict a single bounding box per image-utterance pair. Our work builds upon MDETR of Kamath et al. [21] that modulates visual processing through attention to the input language utterance and decodes objects from queries similar to DETR, without selecting from a pool of proposals. Both our method and MDETR can predict multiple instances being referred to, as well as ground intermediate noun phrases. Concurrent to our work, GLIP [26] shows that adding supervision from detection annotations can improve 2D referential grounding. Our work independently confirms this hypothesis in 2D and also shows its applicability on the 3D domain.

3D referential language grounding has only recently gained popularity [4,2]. To the best of our knowledge, all related approaches are box-bottlenecked: they extract 3D object proposals and select one as their answer. Their pipeline can be decomposed into three main steps: i) Representation of object boxes as point features [46], segmentation masks [48] or pure spatial/categorical features [41]. ii) Encoding of language utterance using word embeddings [46,41] and/or scene graphs [10]. iii) Fusion of the two modalities and scoring of each proposal using graph networks [18] or Transformers [46]. Most of these works also employ domain-specific design choices by explicitly encoding pairwise relationships [18,13,48] or by relying on heuristics, such as restricting attention to be local [49,48] and ignoring input modalities [41]. Such design prevents those architectures from being applicable to both the 3D and 2D domains simultaneously.

Due to the inferior performance of 3D object detectors in comparison to their 2D counterparts, popular benchmarks for 3D language grounding, such as Referit3D [2] provide access to ground-truth object boxes at test time. The proposed BUTD-DETR is the first 3D language grounding model that is evaluated on this benchmark without access to oracle 3D object boxes.

3 Method

We first describe MDETR [21] in Section 3.1. Then, we present BUTD-DETR’s architecture in Section 3.2, supervision augmentation with detection prompts in Section 3.3 and its training objectives in Section 3.4.

3.1 Background: MDETR

MDETR is a 2D language grounding model that takes a referential utterance and an RGB image as input and localises in the image all objects mentioned in the utterance. MDETR encodes the image with a convolutional network [15] and the language utterance with a RoBERTa encoder [29]. It then fuses information across the language and visual features through multiple layers of self-attention on the concatenated visual and language feature sequences. In MDETR’s decoder, a set of query vectors iteratively attend to the contextualized visual features and self-attend to one another, similar to the DETR’s [3] decoder. Finally, each query decodes a bounding box and a confidence score over each word in the input utterance, which associates the box to a text span.

The predicted boxes are assigned to ground-truth ones using a Hungarian matching, similar to [3]. Upon matching, the following losses are computed:

- A bounding box loss between predicted boxes and the corresponding ground-truth ones. This is a combination of L1 and generalized IoU [40] losses.
- A soft token prediction loss. A query matched to a ground-truth box is trained to decode a uniform distribution over the language token positions that refer to that object. Queries not matched to ground-truth targets are trained to predict a no-object label.
- Two contrastive losses between query and language token features. The first one, called *object contrastive loss*, pulls an object query’s features closer to the features of the corresponding ground-truth span’s word tokens, and further than all other tokens. The second one, called *token contrastive loss*, pulls the features of a ground-truth span’s token closer to the corresponding object query features, and further than all other queries.

3.2 Bottom-up Top-down DETR (BUTD-DETR)

The architecture of BUTD-DETR is illustrated in Figure 2. Given a referential language utterance, e.g., “find the plant that is on top of the end table” and a visual scene, which can be a 3D point cloud or a 2D image, BUTD-DETR

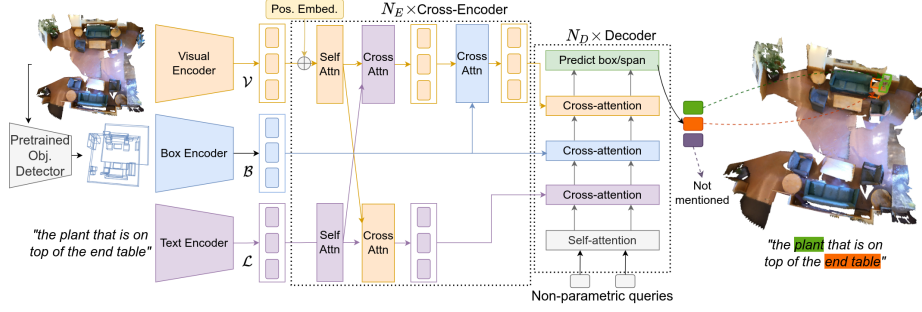


Fig. 2: **BUTD-DETR architecture.** Given a visual scene and a referential utterance, the model localizes all object instances mentioned in the utterance. A pre-trained object detector extracts object box proposals. The visual scene features, the language utterance and the labelled box proposals are encoded into corresponding sequences of visual, word and box tokens using visual, language and box encoders, respectively. The three streams cross-attend and finally decode boxes and corresponding spans in the language utterance that each decoded box refers to. We visualize here the model operating on a 3D point cloud; an analogous architecture is used for 2D image grounding.

is trained to localize all objects mentioned in the utterance. In the previous example, we expect one box for the “plant” and one for the “end table”. The model attends across image/point cloud, language and box proposal streams, then decodes the relevant objects and aligns them to input language spans.

Within-modality encoder In 2D, we encode an RGB image using a pre-trained ResNet101 backbone [16]. The 2D appearance visual features are added to 2D Fourier positional encodings, same as in [50, 19]. In 3D, we encode a 3D point cloud using a PointNet++ backbone [37]. The 3D point visual features are added to learnable 3D positional encodings, same as in [30]: we pass the coordinates of the points through a small multilayer perceptron (MLP). Let $\mathcal{V} \in \mathbb{R}^{n_v \times c_v}$ denote the visual token sequence, where n_v is the number of visual tokens and c_v is the number of visual feature channels.

The words of the input utterance are encoded using a pre-trained RoBERTa [29] backbone. Let $\mathcal{L} \in \mathbb{R}^{n_\ell \times c_\ell}$ denote the word token sequence.

A pre-trained detector is used to obtain 2D or 3D object box proposals. Following prior literature, we use Faster-RCNN [39] for RGB images, pre-trained on 1601 object categories of Visual Genome [24], and Group-Free detector [30] for 3D point clouds pre-trained on a vocabulary of 485 object categories on ScanNet [7]. The detected box proposals that surpass a confidence threshold are encoded using a box proposal encoder, by mapping their spatial coordinates and categorical class information to an embedding vector each, and concatenating them to form an object proposal token. We use a pre-trained and frozen RoBERTa [29]

backbone to encode the semantic categories of proposed boxes. Let $\mathcal{O} \in \mathbb{R}^{n_o \times c_o}$ denote the object token sequence.

The 3D detector is trained on ScanNet and all 3D benchmarks we use are also ScanNet-based. This creates a discrepancy in the quality of the detector’s predictions between train and test time, as it is far more accurate on the training set. As a result, we find that BUTD-DETR tends to rely on the detector at training time and generalizes less at test time, where the detector’s predictions are much noisier. To mitigate this, we randomly replace 30% of the detected boxes at training time with random ones. This augmentation leads to stronger generalization when the detector fails to locate the target object. Note that this is not the case in 2D, where the detector is trained on a different dataset.

All visual, word and box proposal tokens are mapped using (different per modality) MLPs to same-length feature vectors.

Cross-modality Encoder The visual, language and box proposals, interact through a sequence of N_E cross-attention layers. In each encoding layer, visual and language tokens cross-attend to one another and are updated using standard key-value attention. Then, the resulting language-conditioned visual tokens attend to the box proposal tokens. We use standard attention for both streams in 3D and deformable attention [50] for the visual stream in 2D.

In contrast to MDETR, BUTD-DETR keeps visual, language and box stream separate in the encoder instead of concatenating them. This enables us to employ deformable attention [50] in self and cross attention layers involving the visual stream in 2D domain. Deformable attention involves computing bilinearly interpolated features which is expensive and non-robust in discontinuous and sparse modalities like pointclouds, hence we use vanilla attention in 3D (for more details see supplementary). In our experiments, we show that concatenation versus keeping separate streams performs similarly in 3D referential grounding.

Decoder BUTD-DETR decodes objects from contextualized features using non-parametric queries in both 2D and 3D, similar to [50,30]. Non-parametric queries are predicted by visual tokens from the current scene, in contrast to parametric queries used in DETR [3] and MDETR [21] that correspond to a learned set of vectors shared across all scenes. Specifically, the contextualized visual tokens from the last multi-modality encoding layer predict confidence scores, one per visual token. The top- K highest scoring tokens are each fed into an MLP to predict a vector which stands for an *object query*, i.e., a vector that will decode a box center and size relative to the location of the corresponding visual token, similar to D-DETR [50]. The query vectors are updated in a residual manner through N_D decoder layers. In each decoder layer, we employ four types of attention operations. First, the queries self-attend to one another to contextually refine their estimates. Second, they attend to the contextualized word embeddings to condition on the language utterance. Next, they attend to the box proposal tokens and then in the image or point visual tokens. At the end of each decoding layer, there is a prediction head that predicts a box center

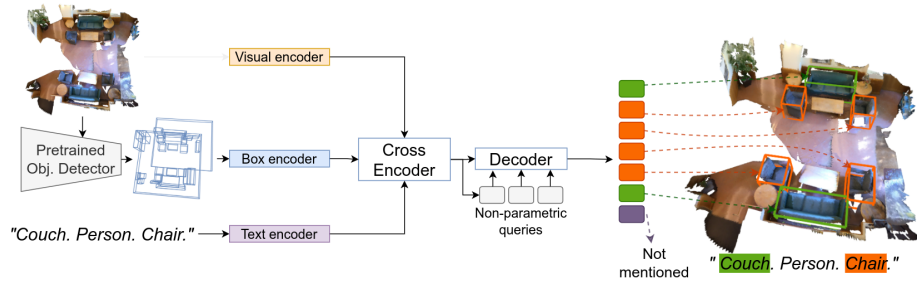


Fig. 3: **Augmenting referential grounding supervision with detection prompts.** A detection prompt is constructed by sequencing sampled object category labels (here *couch*, *person* and *chair*). The task is to localize all instances of mentioned objects and associate them with the correct span in the prompt. 50% of the sampled labels are negative, i.e., they have no corresponding object instance in the scene. The model learns not to associate these spans with predicted boxes.

displacement, height and width vector, and a token span for each object query that localizes the corresponding object box and aligns it with the language input. We refer the reader to our supplementary file for more implementation details.

3.3 Augmenting supervision with detection prompts

Object detection is an instance of referential language grounding in which the utterance is a single word, namely, the object category label. Language grounding models have effectively combined supervision across referential grounding, caption description and question answering tasks [31,32], which is an important factor for their success. Object detection annotations have not been considered so far as candidates for such co-training.

We cast object detection as grounding of detection prompts, namely, referential utterances comprised of a list of object category labels, as shown in Figure 3. Specifically, given the detector’s vocabulary of object category labels, we randomly sample a fixed number of them—some appear in the visual scene and some do not—and generate synthetic utterances by sequencing the sampled labels, e.g., “*Couch. Person. Chair. Fridge.*”, we call them detection prompts. We treat these prompts as referential utterances to be grounded: the task is to localize *all* object instances of the category labels mentioned in the prompt if they appear in the scene. The sampling of negative category labels (labels for which there are no object instances present) operates as negative training: the model is trained to not match any boxes to the negative category labels. Further details on this negative training can be found in the supplementary.

3.4 Supervision objectives

We supervise the outputs of all prediction heads in each layer of the decoder. We follow MDETR [21] in using Hungarian matching to assign a subset of object queries to the ground-truth object boxes and then compute the bounding box, soft token prediction and contrastive losses. Our bounding box and soft token prediction losses are identical to MDETR’s. However, we notice that MDETR’s contrastive losses do not compare all object queries and word tokens symmetrically. Specifically, the object contrastive loss supervises only the object queries that are matched to a ground-truth object box. On the other hand, the token contrastive loss includes only the tokens that belong to positive spans, namely, noun phrases with corresponding object instances in the scene. As a result, object queries not matched to any ground-truth object box are not pulled far from non-ground-truth text spans, which means at inference object queries can be close to negative spans. We find this asymmetry to hurt performance, as we show in our experiments.

To address this, we propose a symmetric alternative where the similarities between all object queries and language tokens are considered. We append the span “not-mentioned” to all input utterances. This acts as the ground-truth text span for all object queries that are not assigned to any of the ground-truth objects. The object contrastive loss now supervises all queries and considers the similarities with all tokens. We empirically find that gathering unmatched queries to “not mentioned” is beneficial. This is similar in principle to the soft token prediction loss, where unmatched queries have to predict “no object”. In fact, we find that this symmetric contrastive loss is sufficient for our model’s supervision, but we observe that co-optimizing for soft token prediction results in faster convergence.

4 Experiments

We test BUTD-DETR on grounding referential utterances in 3D point clouds and 2D images. Our experiments aim to answer the following questions:

1. How does BUTD-DETR perform compared to the state-of-the-art in 3D and 2D language grounding?
2. How does BUTD-DETR perform compared to a straightforward extension of the 2D state-of-the-art MDETR [21] model in 3D?
3. How much, if at all, attending to a bottom-up box proposal stream helps performance?
4. How much, if at all, co-training for grounding detection prompts helps performance?
5. How much, if at all, the proposed contrastive loss variant helps performance?

4.1 Language grounding in 3D point clouds

We test BUTD-DETR on SR3D, NR3D [2] and ScanRefer [4] benchmarks. All three benchmarks contain pairs of 3D point clouds of indoor scenes from ScanNet [7] and corresponding referential utterances, and the task is to localize the

objects referenced in the utterance. The utterances in SR3D are short and synthetic, e.g., “*choose the couch that is underneath the picture*”, while utterances in NR3D and ScanRefer are longer and more natural, e.g. “*from the set of chairs against the wall, the chair farthest from the red wall, in the group of chairs that is closer to the red wall*”. For fair comparison against previous methods, we train BUTD-DETR separately on each of SR3D, NR3D and ScanRefer. We augment supervision in each of the three datasets with ScanNet detection prompts. SR3D provides annotations for all objects mentioned in the utterance, so during training we supervise localization of all objects mentioned. In NR3D and ScanRefer, we use supervision for grounding only the referenced object.

All existing models that have been tested in SR3D or NR3D benchmarks are box-bottlenecked, namely, they are trained to select the answer from a pool of box proposals. They all use **ground-truth 3D object boxes (without category labels)** as the set of boxes to select from. We thus consider two evaluation setups:

1. **det**: where we re-train previous models using their publicly available code and provide the same 3D box proposals we use in BUTD-DETR, obtained by the Group-Free 3D object detector [30] trained to detect 485 object categories in ScanNet (Section **det** in Table 1).
2. **GT**, where we use ground-truth 3D object boxes for our model and baseline (Section **GT** in Table 1).

Alongside previous models, we also compare our model against our implementation of the MDETR model in 3D. This is similar to our model but without attention on a box stream, without co-training with detection prompts and with the original contrastive losses proposed by MDETR. We also replace MDETR’s parametric object queries with non-parametric one—similar to our model—since they have been shown to be crucial for good performance in 3D [30,34]. We call this model MDETR-3D. For the sake of completeness, we do have a 3D version of MDETR that uses parametric queries in Table 2 and, as expected, it is significantly worse. MDETR does not use a pool of box proposals in any way and hence we cannot report results of MDETR-3D under **GT**.

We show quantitative results of our models against previous works in Table 1. We use top-1 accuracy metric, which measures the percentage of times we can find the target box with an IoU higher than the threshold. We report results with IoU@0.25 on SR3D and NR3D; and with both IoU@0.25 and IoU@0.5 on ScanRefer. Please refer to supplementary for more detailed results.

BUTD-DETR outperforms existing approaches as well as MDETR-3D by a large margin under both evaluation setups, **det** and **GT**. It also outperforms the recent SAT-2D [46] that uses additional 2D RGB image features during training. BUTD-DETR does not use 2D image features, but it can be easily extended to do so. We show qualitative results in Figure 4. For more qualitative results, please check the supplementary file.

Table 1: **Results on language grounding in 3D point clouds.** We evaluate top-1 accuracy using ground-truth (GT) or detected (det) boxes. * denotes method uses extra 2D image features. † denotes evaluation with detected boxes using the authors’ code and checkpoints. ‡ denotes re-training using the authors’ code. For [49], we compare against their 3D-only version.

Method	SR3D		NR3D		ScanRefer (Val. Set)
	Acc@0.25(det)	Acc.(GT)	Acc@0.25(det)	Acc@0.25(det)	Acc@0.5(det)
ReferIt3DNet [2]	27.7 [†]	39.8	24.0 [†]	26.4	16.9
ScanRefer [4]	-	-	-	35.5	22.4
TGNN [18]	-	45.0	-	37.4	29.7
3DRefTransformer [1]	-	47.0	-	-	-
InstanceRefer [48]	31.5 [‡]	48.0	29.9 [‡]	40.2	32.9
FFL-3DOG [10]	-	-	-	41.3	34.0
LanguageRefer [41]	39.5 [†]	56.0	28.6 [†]	-	-
3DVG-Transformer [49]	-	51.4	-	45.9	34.5
TransRefer3D [13]	-	57.4	-	-	-
SAT-2D [46]*	35.4 [†]	57.9	31.7 [†]	44.5	30.1
MDETR-[21]-3D (our impl.)	45.4	-	31.5	47.2	31.9
BUTD-DETR (ours)	52.1	67.0	43.3	52.2	39.8

Table 2: **Ablation of design choices for BUTD-DETR on SR3D.**

Model	Accuracy
BUTD-DETR	52.1
w/o visual tokens	41.9
w/o detection prompts	47.9
w/o box stream	51.0
with MDETR’s [21] contrastive loss	49.6
w/o detection prompts; w/o box stream; (MDETR [21]-3D)	45.4
with parametric queries; w/o detection prompts; w/o box stream; (MDETR [21]-3D-Param)	33.8
with concatenated Visual, Language and Object Streams	51.3

Ablative analysis We ablate all our design choices for 3D BUTD-DETR on SR3D benchmark [2] in Table 2. We compare BUTD-DETR against the following variants:

- w/o visual tokens: an object-bottlenecked variant, which only attends to the language and box proposal streams and selects one box out of the proposals.
- w/o detection prompts: BUTD-DETR trained solely on SR3D grounding utterances.
- w/o box stream: BUTD-DETR without attention on the box stream.
- w/ MDETR’s contrastive loss: BUTD-DETR where we replace our modified contrastive loss with MDETR’s.
- w/o detection prompts, w/o box stream, w/ MDETR’s contrastive loss: an MDETR [21]-3D implementation.
- w/ parametric queries, w/o detection prompts, w/o box stream, w/ MDETR’s contrastive loss: an MDETR-3D implementation that uses parametric object queries, as in original MDETR.
- w/ concatenated visual, language and box streams: instead of attending to each modality separately, we concatenate the different streams along their sequence dimension.

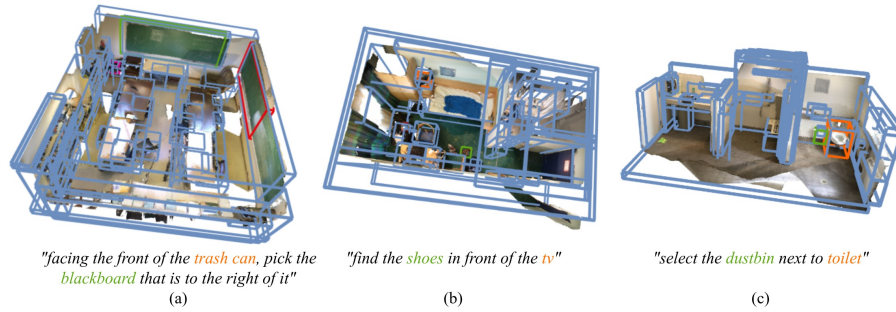


Fig. 4: **Qualitative results of BUTD-DETR in the SR3D benchmark.** Predictions for the target are shown in green and for other mentioned objects in orange. The detected proposals appear in blue. (a) The variant without box stream (red box) fails to exploit the information given by the detector, but BUTD-DETR succeeds. (b) The detector misses the “shoes” and any box-bottlenecked variant fails. (c) The detector is successful in finding the “dustbin”, still BUTD-DETR refines the box to get a more accurate bounding box.

The conclusions are as follows:

1. **Box bottlenecks hurt:** Models such as BUTD-DETR and MDETR-3D that decode object boxes instead of selecting them from a pool of given object proposals significantly outperform box-bottlenecked variants. BUTD-DETR outperforms by 10.2% an object-bottlenecked variant, that does not attend to 3D point features and does not decode boxes.
2. **BUTD-DETR outperforms MDETR-3D by 6.7%:**
3. **Attention on a box proposal stream helps:** Removing attention on the box stream causes an absolute 1.1% drop in accuracy.
4. **Co-training with detection prompts helps:** Co-training with detection prompts contributes 4.2% in performance (from 47.9% to 52.1%).
5. **BUTD-DETR’s contrastive loss helps:** Replacing our contrastive loss with MDETR’s results in drop of 2.5% in absolute accuracy.
6. **Concatenating Visual, Language and Object Streams performs worse than a model that has separate streams for each modality** Our motivation is to keep separate streams in 3D cross-modality encoder and decoder to be consistent with 2D BUTD-DETR as explained in Section 3.2. We additionally find that having separate streams gives a boost of 0.8%.

4.2 Language grounding in 2D images

We test BUTD-DETR on the referential grounding datasets of RefCOCO [23], RefCOCO+ [47] and Flickr30k entities dataset [36]. We follow the pretrain-then-finetune protocol of MDETR and first pre-train on combined grounding annotations from Flickr30k [36], referring expression datasets [23,47,33], Visual

Table 3: **Results on language grounding in 2D RefCOCO and RefCOCO+ Datasets on Top-1 accuracy metric using standard splits.** All training times are computed using same V100 GPUs. Training epochs are written as $x + y$ where x = number of pre-training epochs and y = number of fine-tuning epochs. All reported results use ResNet101 backbone.

Method	RefCOCO			RefCOCO+			Training Epochs	Training GPU Hours
	val	testA	testB	val	testA	testB		
UNITER-L [6]	81.4	87.0	74.2	75.9	81.5	66.7	-	-
VILLA-L [12]	82.4	87.5	74.8	76.2	81.5	66.8	-	-
MDETR [21]	86.8	89.6	81.4	79.5	84.1	70.6	40 + 5	5560
BUTD-DETR (ours)	85.9	88.5	81.5	78.2	82.8	70.0	12 + 5	2748

Table 4: **Results on language grounding in Flickr30k 2D images.** We use Recall@k metric. All training times are computed using same V100 GPUs.

Method	Val			Test			Training Epochs	Training GPU hours
	R@1	R@5	R@10	R@1	R@5	R@10		
VisualBERT [25]	70.4	84.5	86.3	71.3	85.0	86.5	-	-
MDETR [21]	82.5	92.9	94.9	83.4	93.5	95.3	40	5480
BUTD-DETR (ours)	81.2	90.9	92.8	81.0	91.6	93.2	12	2688

Genome [24]. During pre-training the task is to detect all instances of objects mentioned in the utterance. Different than MDETR, we augment this supervision with detection prompts from the MS-COCO dataset [27]. Following MDETR, we directly evaluate our pre-trained model on Flickr30k without any further fine-tuning and fine-tune for 5 epochs on RefCOCO and RefCOCO+.

We report top-1 accuracy on the standard splits of RefCOCO and RefCOCO+ in Table 3 and Recall metric with ANY-BOX protocol [25] on Flickr30k in Table 4. Our model and MDETR use the same 200k image-language pairs from COCO [27], Flickr30k [36] and Visual Genome [24]. VisualBERT [25] is trained on COCO captions. UNITER [6] and VILLA [12] use a larger dataset of 4.4M pairs from COCO, Visual Genome, Conceptual-Captions [43], and SBU Captions [35]. In addition, we augment our training set with detection prompts from COCO. BUTD-DETR trains two times faster than MDETR while getting comparable performance. This computational gain comes mostly from deformable attention which is much cheaper than original visual self-attention that scales quadratically with the number of visual tokens, as already reported in [50]. For qualitative results, please see the supplementary file.

Ablative analysis We ablate our model in RefCOCO without pre-training in Table 5, since pre-training is computationally expensive due to the size of the combined datasets. Consistent with 3D, removing detection prompts results in an accuracy drop of 2.4%. Additionally removing attention to the box proposal stream results in a drop of 3.1% in accuracy. When replacing our contrastive

Table 5: **Ablation for BUTD-DETR on the RefCOCO validation set.**

Model	Accuracy
BUTD-DETR	79.4
w/o det prompts	77.0
w/o box stream w/o det prompts	76.3
w/o box stream w/o det prompts w/ MDETR’s [21] contrastive	74.2

loss with MDETR’s, the model achieves 74.2%, resulting in an additional drop of 2.1% accuracy.

4.3 Limitations

Our work relies on language-image alignment and does not address how to ground language better and more robustly through abstraction of the visual features, e.g., the fact that *left* and *right* reverse when we change the user’s viewpoint, the fact that numbers requires precise counting, or the fact that the “*chair furthest away from the door*” requires to satisfy a logical constraint which our model can totally violate when presented with out-of-distribution visual input. This limitation is a direct avenue for future work.

5 Conclusion

We present BUTD-DETR, a model for referential grounding in 3D and 2D scenes, that attends to language, visual and box proposal streams to decode objects mentioned in the referential utterance and align them to corresponding spans in the input. BUTD-DETR builds upon MDETR [21] and outperforms its straightforward MDETR-3D equivalent by a significant margin thanks to attention on labelled bottom-up box proposals, co-training with detection prompts and improved contrastive losses, setting a new state-of-the-art in two 3D language grounding benchmarks. BUTD-DETR is also the first model in 3D referential grounding that operates on the realistic setup of not having access to oracle object boxes, but rather detects them from the input 3D point cloud.

Acknowledgement This material is based upon work supported by an Air Force grant FA95502010423, an AFOSR Young Investigator Award, DARPA Machine Common Sense, and an NSF CAREER award provided in affiliation to Carnegie Mellon University. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Army or the United States Air Force. The authors would also like to thank Adam W. Harley, Leonid Keselman, Muyang Li and Rohan Choudhary for their helpful feedback.

References

1. Abdelreheem, A., Upadhyay, U., Skorokhodov, I., Yahya, R.A., Chen, J., Elhoseiny, M.: 3DRefTransformer: Fine-Grained Object Identification in Real-World Scenes Using Natural Language. In: Proc. WACV (2022) [11](#)
2. Achlioptas, P., Abdelreheem, A., Xia, F., Elhoseiny, M., Guibas, L.: ReferIt3D: Neural Listeners for Fine-Grained 3D Object Identification in Real-World Scenes. In: Proc. ECCV (2020) [3](#), [4](#), [5](#), [9](#), [11](#)
3. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-End Object Detection with Transformers. In: Proc. ECCV (2020) [2](#), [3](#), [5](#), [7](#)
4. Chen, D.Z., Chang, A., Nießner, M.: ScanRefer: 3D Object Localization in RGB-D Scans using Natural Language. In: Proc. ECCV (2020) [3](#), [4](#), [9](#), [11](#)
5. Chen, X., Ma, L., Chen, J., Jie, Z., Liu, W., Luo, J.: Real-time referring expression comprehension by single-stage grounding network. ArXiv [abs/1812.03426](#) (2018) [4](#)
6. Chen, Y.C., Li, L., Yu, L., Kholy, A.E., Ahmed, F., Gan, Z., Cheng, Y., Liu, J.: UNITER: UNiversal Image-TEXT Representation Learning. In: Proc. ECCV (2020) [4](#), [13](#)
7. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T.A., Nießner, M.: ScanNet: Richly-Annotated 3D Reconstructions of Indoor Scenes. In: Proc. CVPR (2017) [6](#), [9](#)
8. Deng, J., Yang, Z., Chen, T., Zhou, W., Li, H.: Transvg: End-to-end visual grounding with transformers. In: Proc. ICCV (2021) [4](#)
9. Fang, H., Gupta, S., Iandola, F., Srivastava, R.K., Deng, L., Dollár, P., Gao, J., He, X., Mitchell, M., Platt, J.C., et al.: From Captions to Visual Concepts and Back. In: Proc. CVPR (2015) [2](#), [4](#)
10. Feng, M., Li, Z., Li, Q., Zhang, L., Zhang, X., Zhu, G., Zhang, H., Wang, Y., Mian, A.: Free-form Description Guided 3D Visual Graph Network for Object Grounding in Point Cloud. In: Proc. ICCV (2021) [4](#), [11](#)
11. Fukui, A., Park, D.H., Yang, D., Rohrbach, A., Darrell, T., Rohrbach, M.: Multimodal Compact Bilinear Pooling for Visual Question Answering and Visual Grounding. In: Proc. EMNLP (2016) [2](#), [4](#)
12. Gan, Z., Chen, Y.C., Li, L., Zhu, C., Cheng, Y., Liu, J.: Large-Scale Adversarial Training for Vision-and-Language Representation Learning. In: Proc. NeurIPS (2020) [13](#)
13. He, D., Zhao, Y., Luo, J., Hui, T., Huang, S., Zhang, A., Liu, S.: TransRefer3D: Entity-and-Relation Aware Transformer for Fine-Grained 3D Visual Grounding. In: Proc. ACMMM (2021) [4](#), [11](#)
14. He, K., Gkioxari, G., Dollár, P., Girshick, R.B.: Mask R-CNN. In: Proc. ICCV (2017) [3](#)
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016) [5](#)
16. He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition. In: Proc. CVPR (2016) [6](#)
17. Hu, R., Rohrbach, M., Andreas, J., Darrell, T., Saenko, K.: Modeling Relationships in Referential Expressions with Compositional Modular Networks. In: Proc. CVPR (2017) [2](#), [4](#)
18. Huang, P.H., Lee, H.H., Chen, H.T., Liu, T.L.: Text-Guided Graph Neural Networks for Referring 3D Instance Segmentation. In: Proc. AAAI (2021) [4](#), [11](#)

19. Jaegle, A., Gimeno, F., Brock, A., Zisserman, A., Vinyals, O., Carreira, J.: Perceiver: General Perception with Iterative Attention. In: Proc. ICML (2021) [6](#)
20. Johnson, J., Karpathy, A., Fei-Fei, L.: DenseCap: Fully Convolutional Localization Networks for Dense Captioning. In: Proc. CVPR (2016) [2](#), [4](#)
21. Kamath, A., Singh, M., LeCun, Y.A., Misra, I., Synnaeve, G., Carion, N.: MDETR - Modulated Detection for End-to-End Multi-Modal Understanding. In: Proc. ICCV (2021) [2](#), [4](#), [5](#), [7](#), [9](#), [11](#), [13](#), [14](#)
22. Karpathy, A., Fei-Fei, L.: Deep Visual-semantic Alignments for Generating Image Descriptions. In: Proc. CVPR (2015) [2](#), [4](#)
23. Kazemzadeh, S., Ordonez, V., Matten, M.A., Berg, T.L.: ReferItGame: Referring to Objects in Photographs of Natural Scenes. In: Proc. EMNLP (2014) [3](#), [4](#), [12](#)
24. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., Bernstein, M.S., Fei-Fei, L.: Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. *International Journal of Computer Vision* **123** (2016) [6](#), [13](#)
25. Li, L.H., Yatskar, M., Yin, D., Hsieh, C.J., Chang, K.W.: VisualBERT: A Simple and Performant Baseline for Vision and Language. ArXiv [abs/1908.03557](#) (2019) [13](#)
26. Li, L.H., Zhang, P., Zhang, H., Yang, J., Li, C., Zhong, Y., Wang, L., Yuan, L., Zhang, L., Hwang, J.N., Chang, K.W., Gao, J.: Grounded Language-Image Pre-training. In: Proc. CVPR (2022) [4](#)
27. Lin, T.Y., Maire, M., Belongie, S.J., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common Objects in Context. In: Proc. ECCV (2014) [13](#)
28. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S.E., Fu, C., Berg, A.C.: SSD: Single Shot MultiBox Detector. In: Proc. ECCV (2016) [3](#)
29. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: RoBERTa: A Robustly Optimized BERT Pretraining Approach. ArXiv [abs/1907.11692](#) (2019) [5](#), [6](#)
30. Liu, Z., Zhang, Z., Cao, Y., Hu, H., Tong, X.: Group-Free 3D Object Detection via Transformers. In: Proc. ICCV (2021) [2](#), [4](#), [6](#), [7](#), [10](#)
31. Lu, J., Batra, D., Parikh, D., Lee, S.: ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks. In: Proc. NeurIPS (2019) [4](#), [8](#)
32. Lu, J., Goswami, V., Rohrbach, M., Parikh, D., Lee, S.: 12-in-1: Multi-Task Vision and Language Representation Learning. In: Proc. CVPR (2020) [8](#)
33. Mao, J., Huang, J., Toshev, A., Camburu, O.M., Yuille, A.L., Murphy, K.P.: Generation and Comprehension of Unambiguous Object Descriptions. In: Proc. CVPR (2016) [12](#)
34. Misra, I., Girdhar, R., Joulin, A.: An End-to-End Transformer Model for 3D Object Detection. In: Proc. ICCV (2021) [4](#), [10](#)
35. Ordonez, V., Kulkarni, G., Berg, T.: Im2text: Describing images using 1 million captioned photographs. In: Proc. NIPS (2011) [13](#)
36. Plummer, B.A., Wang, L., Cervantes, C.M., Caicedo, J.C., Hockenmaier, J., Lazebnik, S.: Flickr30k Entities: Collecting Region-to-Phrase Correspondences for Richer Image-to-Sentence Models. In: Proc. ICCV (2015) [12](#), [13](#)
37. Qi, C., Yi, L., Su, H., Guibas, L.J.: PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In: Proc. NIPS (2017) [6](#)
38. Redmon, J., Divvala, S.K., Girshick, R.B., Farhadi, A.: You Only Look Once: Unified, Real-Time Object Detection. In: Proc. CVPR (2016) [3](#), [4](#)

39. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards Real-time Object Detection with Region Proposal Networks. In: Proc. NIPS (2015) 2, 4, 6
40. Rezatofighi, S.H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I.D., Savarese, S.: Generalized Intersection Over Union: A Metric and a Loss for Bounding Box Regression. In: Proc. CVPR (2019) 5
41. Roh, J., Desingh, K., Farhadi, A., Fox, D.: LanguageRefer: Spatial-Language Model for 3D Visual Grounding. In: Proc. CoRL (2021) 4, 11
42. Sadhu, A., Chen, K., Nevatia, R.: Zero-shot grounding of objects from natural language queries. In: Proc. ICCV (2019) 4
43. Sharma, P., Ding, N., Goodman, S., Soricut, R.: Conceptual Captions: A Cleaned, Hypernymed, Image Alt-text Dataset for Automatic Image Captioning. In: Proc. ACL (2018) 13
44. Yang, Z., Chen, T., Wang, L., Luo, J.: Improving one-stage visual grounding by recursive sub-query construction. In: Proc. ECCV (2020) 4
45. Yang, Z., Gong, B., Wang, L., Huang, W., Yu, D., Luo, J.: A fast and accurate one-stage approach to visual grounding. In: Proc. ICCV (2019) 4
46. Yang, Z., Zhang, S., Wang, L., Luo, J.: SAT: 2D Semantics Assisted Training for 3D Visual Grounding. In: Proc. ICCV (2021) 4, 10, 11
47. Yu, L., Poirson, P., Yang, S., Berg, A.C., Berg, T.L.: Modeling Context in Referring Expressions. In: Proc. ECCV (2016) 3, 12
48. Yuan, Z., Yan, X., Liao, Y., Zhang, R., Li, Z., Cui, S.: InstanceRefer: Cooperative Holistic Understanding for Visual Grounding on Point Clouds through Instance Multi-level Contextual Referring. In: Proc. ICCV (2021) 4, 11
49. Zhao, L., Cai, D., Sheng, L., Xu, D.: 3DVG-Transformer: Relation Modeling for Visual Grounding on Point Clouds. In: Proc. ICCV (2021) 4, 11
50. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable DETR: Deformable Transformers for End-to-End Object Detection. In: Proc. ICLR (2021) 3, 6, 7, 13