

# CAPE: A CLIP-Aware Pointing Ensemble of Complementary Heatmap Cues for Embodied Reference Understanding

Fevziye Irem Eyiokur<sup>1,4</sup> Dogucan Yaman<sup>1,4</sup> Hazım Kemal Ekenel<sup>2</sup> Alexander Waibel<sup>1,3,4</sup>  
<sup>1</sup>Karlsruhe Institute of Technology, <sup>2</sup>Istanbul Technical University, <sup>3</sup>Carnegie Mellon University  
<sup>4</sup>KIT Campus Transfer GmbH (KCT)  
 irem.eyiokur@kit.edu

## Abstract

We address *Embodied Reference Understanding*, the task of predicting the object a person in the scene refers to through pointing gesture and language. This requires multimodal reasoning over text, visual pointing cues, and scene context, yet existing methods often fail to fully exploit visual disambiguation signals. We also observe that while the referent often aligns with the head-to-fingertip direction, in many cases it aligns more closely with the wrist-to-fingertip direction, making a single-line assumption overly limiting. To address this, we propose a dual-model framework, where one model learns from the head-to-fingertip direction and the other from the wrist-to-fingertip direction. We introduce a Gaussian ray heatmap representation of these lines and use them as input to provide a strong supervisory signal that encourages the model to better attend to pointing cues. To fuse their complementary strengths, we present the CLIP-Aware Pointing Ensemble module, which performs a hybrid ensemble guided by CLIP features. We further incorporate an auxiliary object center prediction head to enhance referent localization. We validate our approach on *YouRefIt*, achieving 75.0 mAP at 0.25 IoU, alongside state-of-the-art CLIP and  $C_D$  scores, and demonstrate its generality on unseen CAESAR and ISL Pointing, showing robust performance across benchmarks.

## 1. Introduction

Embodied Reference Understanding (ERU) [13] is the task of identifying a specific object in a visual scene based on language instructions and pointing cues in the image. It plays a key role in real-world applications like human-robot interaction, assistive robotics, and augmented reality, where systems must understand which object a person refers to.

While grounding models [39] and LMMs [3, 4, 6, 56] have made significant progress in detecting objects mentioned in natural language, they often fall short in ERU task,

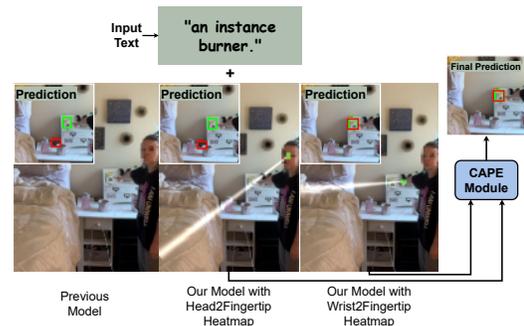


Figure 1. An example where previous models fail. Our  $M_H$  model also fails due to distractions from other objects along the head-to-fingertip pointing line. In contrast, our  $M_W$  model correctly identifies the target object, and CAPE selects this as the final prediction. This highlights the flexibility and robustness of our approach.

particularly in ambiguous scenes. When multiple instances of the same object type are present, these models tend to detect all matching candidates without the ability to disambiguate and identify the specific target intended by the user. Moreover, when the textual instruction itself is vague or ambiguous, grounding models struggle even further, often failing to identify the correct target object or producing no confident prediction at all due to the inherent ambiguity. This limitation highlights the need for additional disambiguation cues that can help resolve referential ambiguity and enable accurate identification of the intended object. Therefore, both embodied gesture signals (pointing) and language reference are crucial to identify the referent, aligning with early multimodal interaction studies [61, 72].

The ERU task was first defined in [13] with the *YouRefIt* dataset, and a method combining textual and visual inputs with PAF [8] and predicted saliency maps [36] was proposed. While effective, the model often fails to follow pointing instructions accurately. The *Touch-Line Transformer* [38] addresses this by predicting a Virtual Touch-Line along the head-to-fingertip direction. However, it still

faces limitations as the pointing line is not explicitly provided as input, despite its success [84], and the head-to-fingertip direction does not always align with the target. In some cases, accurate pointing is better represented from the hand alone (e.g., wrist-to-fingertip, Fig. 1), highlighting the need for more reliable pointing guidance.

To tackle these challenges, we propose a dual-model framework that leverages complementary pointing cues. Both models take as input a referring expression and an image with a pointing gesture, along with an additional heatmap that encodes the pointing direction. One model utilizes a head-to-fingertip heatmap, while the other uses a wrist-to-fingertip heatmap, capturing different aspects of the pointing behavior. Since these two models provide complementary information depending on the scenario, we introduce the CLIP-Aware Pointing Ensemble (CAPE) module to effectively combine the strengths of both models. In this module, we leverage the CLIP model to compute semantic similarity scores between the input text and candidate object images. We use CLIP because it was trained on a large-scale image-text dataset, learning a joint embedding space in which semantically matching images and text are closely aligned. Additionally, we introduce an auxiliary object center prediction head, providing a supervisory signal that guides the model to more accurately localize the pointed object.

Our contributions are: (1) We propose using a pointing heatmap as an additional modality to guide our model to focus more effectively on pointing cues. Specifically, we use detected head, wrist, and fingertip points to construct a Gaussian Ray Heatmap that highlights the approximate pointing area. The embedding of this pointing heatmap is then extracted via a heatmap encoder and provided to the model. (2) We propose two parallel complementary models to focus on head-to-fingertip and wrist-to-fingertip pointing lines, to individually alleviate inherent challenges of the task. (3) We present CAPE module to effectively ensemble two complementary models as ensembling has shown strong performance in other fields [19, 27, 78, 82, 83]. (4) We introduce an object center prediction head to improve prediction accuracy by emphasizing object localization without requiring precise bounding boxes.

## 2. Related Work

**Referring Expression Comprehension (REC)** In REC [26, 34, 44, 47, 48, 77], the goal is to identify a specific region in an image based on a given referring expression. In contrast to traditional object detection, REC interprets free-form text and can locate objects from any category, including previously unseen ones. Similarly, OV-DETR [79] integrates image and text embeddings from a CLIP model as queries within the DETR [11] to generate category-specific bounding boxes. ViLD [22]

distills knowledge from a CLIP teacher model into an R-CNN-like detector [21], enabling region embeddings to capture semantic information from language. GLIP [20] formulates object detection as a grounding task, using additional grounding datasets to align visual regions with textual phrases. DetCLIP [76] enriches its knowledge base using generated pseudo labels. While YOLO-World [14] extends traditional YOLO object detection [51] to an open-world setting, GroundingDINO [39] aligns visual and textual features to detect arbitrary objects described by free-form text. Recently, LMMs [1, 3, 6] get attention for their superior performance in visual grounding [12].

**Nonverbal Communication for Referent** Some studies [7, 23, 24, 28, 32, 46, 52, 54, 57–60, 62–68, 70, 71, 73] use gaze target detection as a nonverbal cue, however, gaze alone is unreliable due to distractions and the absence of a clear pointing moment. It is typically localized frame by frame and often used to support conversation [5, 10]. In [45], a multimodal exophora resolution method is proposed to disambiguate demonstrative expressions like “that one” by integrating object categories, pointing gestures, and prior environmental knowledge. [16] introduced an interactive robot dialogue system that uses multimodal interaction and pointing line estimation to accurately identify referent with an iterative correction process. Their follow-up work [17] extended this by recognizing unseen objects using a Region Proposal Network and VL-T5 multimodal network [15], moving beyond general object classes. On the other hand, several datasets have been collected to incorporate nonverbal cues, but most remain in simulated domains due to accessibility and controllability [2, 29–31].

**Embodied Reference Understanding (ERU)** ERU is a recent advancement of REC that considers the subject’s position while pointing at a reference object. Chen et al. [13] introduced the task and benchmarks using both verbal (text) and nonverbal (pointing gesture) cues. Their model leverages predicted saliency maps and PAF [8] as gestural features to better perceive pointing direction and detect referents. To address ambiguities caused by camera perspective, [53] mapped scenes to 3D coordinates using depth estimation and subject position for spatial attention, but this only yielded modest improvements. Li et al. [38] further improved ERU using an MDETR-based [33] model, introducing a “virtual-touch-line” from eye to fingertip and predicting its vector alongside bounding boxes. In this work, we also use a transformer-based multimodal object detector similar to MDETR [33]. However, instead of predicting a single line vector, we extend the visual-touch-line concept to a pointing heatmap representing the focus of attention. To capture different pointing variations, we use two different heatmaps processed in parallel networks. Some recent works [41, 42] have extended ERU task to 3D embodied settings. ScanERU [41] introduced the first 3D-ERU bench-

mark by inserting human avatars into existing 3D datasets, but relied on manual placement. In contrast, Ges3ViG [42] improves realism and scalability by automating avatar insertion, generating gesture-aware instructions, and incorporating human localization into the grounding pipeline.

### 3. Methodology

In ERU, given an RGB image  $x_{img} \in \mathbb{R}^{3 \times H \times W}$  and a text input  $x_{text} \in \mathbb{N}^L$ , we predict a bounding box  $x_{bbox} \in \mathbb{R}^4$  corresponding to an object referenced by the text and indicated by a pointing gesture in the image. We propose using an additional pointing heatmap  $x_{phm} \in \mathbb{R}^{1 \times H \times W}$  to supervise the model better focus on the visual pointing reference.

#### 3.1. Model Architecture

Fig. 2 illustrates our overall approach, which consists of two parallel models ( $M_H$  and  $M_W$ ; only one is shown since the sole difference is the heatmap encoder input). After processing the text, image, and heatmap inputs, the outputs of both models are passed to the CLIP-Aware Pointing Ensemble (CAPE) module for the final decision. CAPE combines the outputs of these two complementary models to enhance overall performance. Specifically,  $M_H$  and  $M_W$  each comprise a large-scale text encoder, image encoder, and heatmap encoder. The outputs of these encoders are concatenated and fed into a transformer [69] encoder block. A subsequent transformer decoder then processes the encoder output to generate predictions via task-specific heads.

**Encoders** We use a pretrained ROBERTA [40] which is a robustly optimized version of BERT [18] as the text encoder to obtain textual embeddings from the input text. To obtain image embeddings,  $F_I \in \mathbb{R}^{2048 \times 8 \times 8}$ , we utilize a pretrained ResNet-101 [25]. Since the pointing heatmap ( $x_{phm}$ ) is a sparse representation, we choose a lightweight encoder, ResNet-18 [25], to embed it:  $F_{phm} \in \mathbb{R}^{256 \times 8 \times 8}$ . This choice not only improves the computational efficiency of the network but also helps prevent redundant information in the final concatenated feature representation, which could otherwise harm the model’s learning capacity.

**Transformer encoder and decoder** After obtaining the embeddings from three encoders, we first apply a  $1 \times 1$  convolution to each embedding individually to project their channel dimensions to 256. Next, we flatten the spatial dimensions of the embedding into a single dimension, converting them into sequences of tokens. For example, for the heatmap embedding:  $F_{phm}^S = F(C_{1 \times 1}(F_{phm})) \in \mathbb{R}^{64 \times 256}$ , where  $F$  is flatten and  $C$  is convolution. We then concatenate the three embeddings at the sequence level to form the final representation, which is fed into the transformer encoder. It processes this concatenated input to learn multimodal representations. We feed the multimodal representation output from the transformer encoder into the

transformer decoder. Additionally, we provide a set of learnable object queries and gestural keypoint queries. The transformer decoder is responsible for generating object output embeddings and gestural output embeddings.

**Prediction head** The object and gestural output embeddings produced by the transformer decoder serve as inputs to our prediction heads. These heads are responsible for predicting candidate bounding boxes for the referent, gestural keypoints (pointing direction prediction head in Fig. 2), which predicts eye, fingertip, wrist coordinates as well as arm class, and center points of the referent. Finally, we select one bounding box prediction, one center point and a pair of gestural keypoints with the highest confidence scores as the final prediction. Please note that we use multi-layer perceptrons (MLPs) as the prediction heads.

#### 3.2. Pointing Heatmap Modality Learning

Incorporating a heatmap that represents the pointing direction provides a valuable spatial priors for the network, especially considering the fact that the existing models in this field tend to give less attention to the visual pointing. This explicit encoding pointing direction helps localize the target region by guiding the model’s attention toward direction and areas that are more likely to be referenced. This is particularly beneficial in real-world, cluttered scenes where gesture interpretation is inherently ambiguous.

**How should we create heatmaps?** While prior work such as VTL [38] shows that head-to-fingertip cues can often capture the intended pointing direction, our goal is to design a heatmap representation that remains reliable across a wider range of real-world conditions. Human pointing gestures are guided by visual attention, and the head-to-fingertip line naturally encodes this gaze-gesture alignment. However, relying on this cue alone can be problematic when the performer is looking elsewhere (e.g., toward a robot or camera) or when the head pose becomes ambiguous due to extreme angles or occlusions. Moreover, close proximity to the object or incomplete arm extension can also break alignment. To address these limitations, our approach incorporates both the head-to-fingertip and wrist-to-fingertip lines. The wrist-to-fingertip line provides a complementary, locally grounded cue that remains informative even when head orientation is unreliable. By combining these two directional sources, we aim to obtain a more robust and consistent representation of pointing intent. Based on this hypothesis, we generate Gaussian Ray Heatmaps for both cues rather than relying solely on explicit line overlap. Both heatmaps extend a ray toward the image boundary through the fingertip, originating at the eyes for  $x_{phm}^H$  and at the wrist for  $x_{phm}^W$ . These heatmaps serve as complementary spatial supervision signals generated through our heatmap process (see Appendix A), capturing not only overlap with a single pointing line but also

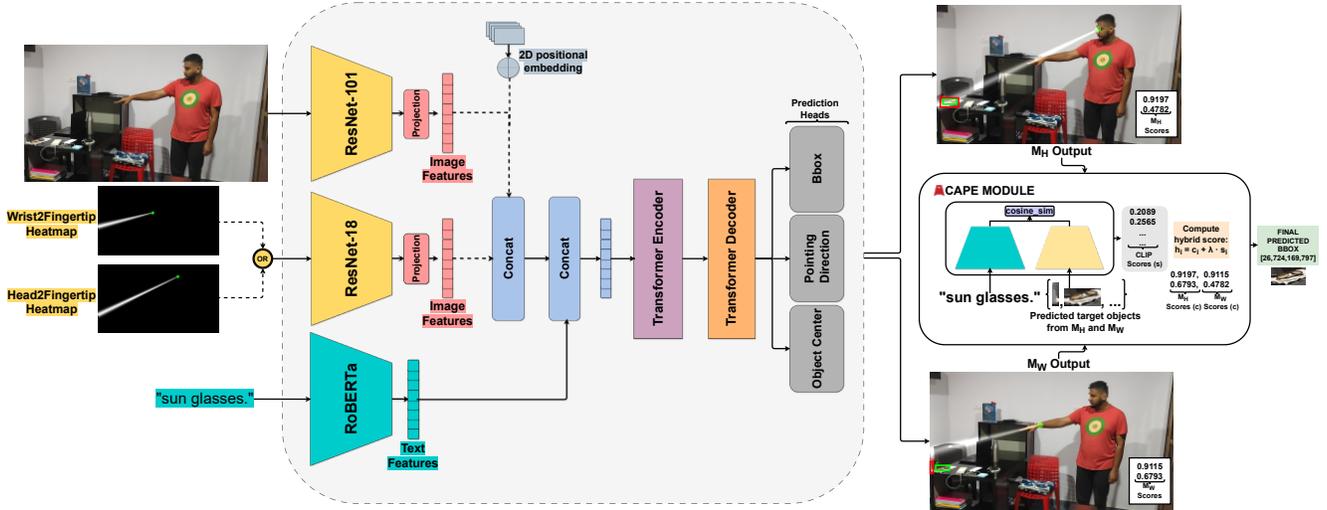


Figure 2. Overview of our approach.  $M_H$  and  $M_W$  share the same architectural design; therefore, only one model is illustrated. For visualization, we show the two different heatmaps:  $M_H$  uses a head-to-fingertip heatmap, while  $M_W$  uses a wrist-to-fingertip heatmap.

a broader region that reflects the performer’s focus of attention.

**How should we integrate these heatmaps?** We investigate two strategies for integrating the two heatmaps. The first merges them directly ( $x_{phm}^H + x_{phm}^W$ ) and feeds the combined input into a single network, testing whether the model can implicitly resolve the two cues and infer the correct pointing direction from pose context. The second trains two identical networks separately, one with the head-to-fingertip heatmap ( $M_H$ ) and one with the wrist-to-fingertip heatmap ( $M_W$ ), and ensembles their predictions. Empirically, the second strategy works substantially better. It yields two complementary models, with  $M_H$  providing the strongest gains. In contrast, the merged-heatmap approach offers only marginal improvement. The model behaves similarly to the head-only variant, and the extra wrist-to-fingertip signal introduces conflicting directional cues that ultimately reduce accuracy compared to using the head-to-fingertip heatmap alone (see Tab. 4, Setup G). We find that the most effective way to leverage heatmaps is to process them with a dedicated encoder whose embeddings are concatenated with the image encoder’s output. This design preserves the pretrained visual backbone while allowing the heatmap encoder to capture richer, modality-specific spatial cues.

### 3.3. CLIP-Aware Pointing Ensemble (CAPE)

We introduce an inference time ensemble module to increase the performance by benefiting from two complementary models,  $M_H$  and  $M_W$ , since we find that they demonstrate strength in different scenarios. CAPE is an adaptive scoring method that combines model confidence and CLIP-based similarity in a size-aware manner effectively. We choose the CLIP model because it was trained to understand

and evaluate image-text semantic similarity, which aligns well with our task. (1) For each prediction (top-2 predictions from both models), we compute a normalized CLIP score [49]. We sum pointing models’ confidence scores with the CLIP scores to obtain the fused score (CLIP Fusion in Tab. 6) since the models’ confidence scores also provide reliable insights. (2) We calculate CLIP scores for the top-1 predictions of both models. If the confidence scores of the second-highest predictions exceed a threshold, we also compute their CLIP scores. Finally, we select the prediction with the highest CLIP score from among these candidate boxes (CLIP-Only Top-2 + Threshold=0.95 in Tab. 6). CAPE applies strategy (1) when the referent is a small object (defined as occupying less than 0.48% of the image area, following [13]) to avoid relying solely on the CLIP prediction, as CLIP becomes less reliable for smaller objects. For all other objects, strategy (2) is used (see Appendix B.5 for details). This hybrid approach leverages the strengths of both signals while adapting to object scale. All thresholds and selection rules were tuned on the validation set to avoid test-time overfitting.

### 3.4. Gestural Signal Learning

**Referent alignment loss** Following [38], we incorporate a referent alignment loss ( $L_{RA}$ ) to enforce consistency between the predicted referent and the VTL. The core idea is that a correct referent should be geometrically aligned with the pointing direction. To quantify this alignment, we compute the cosine similarity between the eye-to-fingertip vector and the eye-to-object vector, defined as:

$$CS_p = \omega((x_f - x_e, y_f - y_e), (x_o - x_e, y_o - y_e)) \quad (1)$$

where  $(x_e, y_e)$  and  $(x_f, y_f)$  are the eye and the fingertip coordinates.  $(x_o, y_o)$  are the center of the predicted bounding box.  $\omega$  is cosine similarity. We compute the same similarity using GT bounding box to obtain  $CS_{GT}$ . The referent alignment loss then penalizes deviations between the predicted and GT alignment:

$$L_{RA} = ReLU(CS_p - CS_{GT}) \quad (2)$$

Please note that for both  $CS_p$  and  $CS_{GT}$ , we use GT eye and fingertip coordinates to ensure accurate directional representation. We apply this loss during the training of the model  $M_H$ , which uses the head-to-fingertip heatmap as reference. For training the  $M_W$  model, which utilizes the wrist-to-fingertip heatmap, we replace the eye coordinates with wrist coordinates. The rest remains unchanged, as we similarly aim to maximize the directional correlation.

**Object center prediction** In our task, identifying the correct referent and accurately localizing it are two distinct but equally important challenges. To enhance the localization capability of the model, we decouple the prediction of the object center from the rest of the bounding box regression. Specifically, we introduce an additional prediction head implemented as a MLP at the end of the network. This head is dedicated to predicting the  $(x, y)$  coordinates of the referent object’s center, independently from the standard bounding box regression head. During training, we supervise this center prediction using the L1 loss between the predicted and GT center coordinates:

$$L_{center} = \|(x_o^{GT}, y_o^{GT}) - (x_o^p, y_o^p)\|_1 \quad (3)$$

This explicit supervision encourages the model to focus on spatial alignment and leads to more accurate object localization, especially in scenes with dense or overlapping objects. Moreover, this center prediction serves as an auxiliary task, which improves the representational capacity of the model by encouraging the visual backbone to learn richer geometric and spatial features. We apply this auxiliary head and its corresponding loss to both models. Since this head is used only during training, it introduces no additional computational overhead during inference.

**Gesture prediction** In addition to the bounding box and object center predictions, our model incorporates several auxiliary heads to predict task-specific features, including eye (wrist in  $M_W$ ) and fingertip coordinates, and arm classification (See prediction head in Sec. 3.1). For each prediction, we apply a standard L1 loss between the predicted and GT data, which encourages spatial attention to task-relevant regions such as the eyes and fingertips, enhancing the model’s understanding of referential gestures. Moreover, we apply cross-entropy loss to classify whether the predicted gestural keypoints are correct.

**Total loss** We define our objective as follows:

$$L = \lambda_1 L_b + \lambda_2 L_{RA} + \lambda_3 L_{center} + \lambda_4 L_g + \lambda_5 L_t + \lambda_6 L_c \quad (4)$$

where  $L_b$  is bounding box loss and it is the combination of L1 and GIoU losses as in DETR-like methods [11, 37, 39, 43, 81, 85].  $L_{RA}$  denotes referent alignment loss,  $L_{center}$  indicates object center prediction loss, and  $L_g$  is gesture prediction loss.  $L_t$  and  $L_c$  are soft token loss and contrastive loss respectively to help visual and textural signals alignment as in [33]. We empirically determine the optimal coefficients on the validation set (See Appendix D).

## 4. Experimental Results

**Dataset** For both training and testing, we use the YouRefIt dataset [13], which contains 2,950 training and 1,245 test images. Text instructions and the annotations of bounding boxes, as well as pointing line coordinates, are provided by [13, 38]. Additionally, we evaluate our approach on the unseen and more challenging ISL Pointing dataset [16], and a simulation dataset named CAESAR [29] (for results see appx. Tab. E.13) to assess its generalization capability.

**Evaluation** For fair evaluation and comparison, we follow prior work [13] and report mean Average Precision (mAP) under three Intersection-over-Union (IoU) thresholds: 0.25, 0.50, and 0.75. Additionally, mAP scores are reported with respect to object size, categorized as Small (S), Medium (M), and Large (L) which based on the ratio of the object area to the image area using two thresholds: 0.48% and 1.76%. In addition to this standard evaluation protocol, we introduce two metrics for assessing referential grounding performance: CLIP score and center coordinate distance ( $C_D$ ). To compute the CLIP score, we extract features from the text input and the cropped region corresponding to the predicted bounding box using CLIP’s text and image encoders. We then calculate the cosine similarity between these feature vectors to quantify the semantic alignment between the predicted object and the text. For  $C_D$ , we compute the L1 distance between the center coordinates of the predicted and the GT bounding boxes, providing a direct measure of spatial alignment.

**Implementation** We use the AMSGrad optimizer [35, 50] during training and train our models for 30 epochs. Both the transformer encoder and decoder consist of 6 layers, each with 8 attention heads and an MLP dimension of 2048. We apply dropout [55] with  $p = 0.1$  in every layer of both the transformer encoder and decoder. The learning rate is set to  $1e - 4$  for the text encoder and  $5e - 5$  for the remaining components. All experiments are conducted on a single NVIDIA RTX A6000 GPU with a batch size of 4.

### 4.1. Results

Tab. 1 compares the performance of our method with existing approaches on the YouRefIt benchmark. Our model

IoU Threshold for mAP	0.25				0.50				0.75			
Object Sizes	All	S	M	L	All	S	M	L	All	S	M	L
PaliGemma2 [56]	58.8	29.0	53.5	75.8	46.9	22.1	50.8	68.0	31.7	6.2	34.1	54.8
Qwen2.5vl [4]	38.9	17.0	41.8	58.0	31.0	11.1	33.6	48.1	20.0	5.7	19.8	34.5
Grounding DINO [39]	57.9	38.0	60.9	74.9	54.9	35.7	59.3	69.6	<b>42.3</b>	<b>22.7</b>	<b>45.9</b>	<b>58.4</b>
FAOA [74]	44.5	30.6	48.6	54.1	30.4	15.8	36.5	39.3	8.5	1.4	9.6	14.4
ReSC [75]	49.2	32.3	54.7	60.1	34.9	14.1	42.5	47.7	10.5	0.2	10.6	20.1
YourRefit PAF [13]	52.6	35.9	60.5	61.4	37.6	14.6	49.1	49.1	12.7	1.0	16.5	20.5
YourRefit Full [13]	54.7	38.5	64.1	61.6	40.5	16.3	54.4	51.1	14.0	1.2	17.2	23.3
REP [53]	58.8	44.7	68.9	63.2	45.7	25.4	57.7	54.3	18.8	3.8	22.2	29.9
Touch-Line-EWL [38]	69.5	<u>56.6</u>	71.7	80.0	60.7	44.4	66.2	71.2	35.5	11.8	38.9	55.0
Touch-Line-VTL [38]	<u>71.1</u>	55.9	<u>75.5</u>	<u>81.7</u>	<u>63.5</u>	<u>47.0</u>	<u>70.2</u>	<b>73.1</b>	<u>39.0</u>	<u>13.4</u>	<u>45.2</u>	<u>57.8</u>
Ours (CAPE)	<b>75.0</b>	<b>63.2</b>	<b>80.2</b>	<b>81.8</b>	<b>65.4</b>	<b>49.5</b>	<b>74.3</b>	<u>72.7</u>	35.7	<u>13.4</u>	40.1	53.5

Table 1. Comparison of our model with prior works, SOTA LMMs, and Grounding-DINO in terms of mean Average Precision (mAP) at different IoU thresholds, across various object sizes, on the YouRefIt dataset [13].

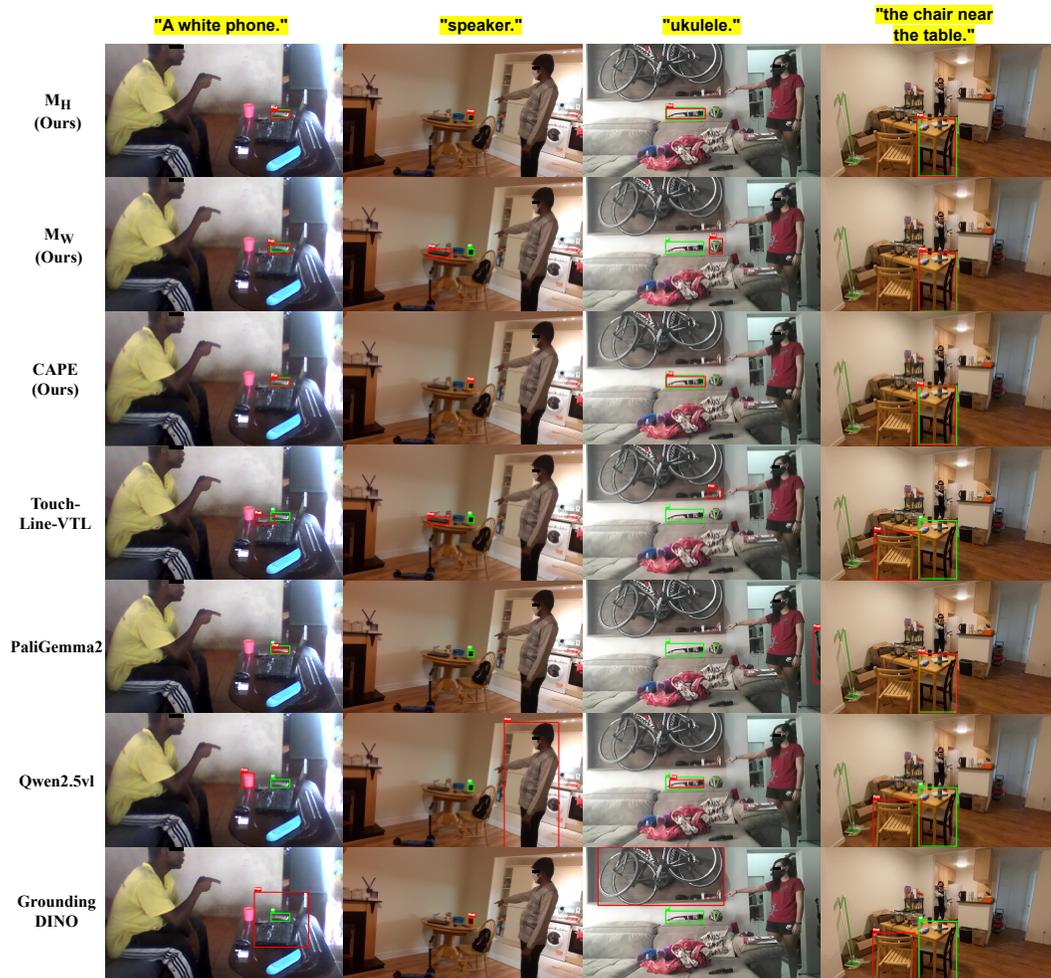


Figure 3. Qualitative comparison of our models with SOTA Touch-Line Transformer [38].

achieves SOTA results in most cases, despite being trained for only 30 epochs on a single GPU with a batch size of 4. In contrast, Touch-Line models [38] were trained for

200 epochs with a batch size of 56. This clearly demonstrates the efficiency and effectiveness of our approach. However, our optimized training setup affects bounding box

Objet Sizes	CLIP Score $\uparrow$				$C_D \downarrow$			
	All	S	M	L	All	S	M	L
PaliGemma2	0.2488	0.2302	0.2446	0.2659	0.4684	0.5784	0.4307	0.4175
Qwen2.5vl-3b	0.2408	0.2261	0.2388	0.2568	0.7288	0.7074	0.6893	0.7891
Grounding DINO	0.2437	0.2267	0.2413	0.2621	0.3496	0.2946	0.4521	0.3104
Touch-Line-EWL	0.2456	0.2312	0.2435	0.2615	0.3168	0.3006	0.2903	0.3564
Touch-Line-VTL	0.2456	0.2308	0.2440	0.2615	0.2843	0.2809	0.2276	0.3393
Ours (CAPE)	<b>0.2661</b>	<b>0.2642</b>	<b>0.2671</b>	<b>0.2670</b>	<b>0.2476</b>	<b>0.2137</b>	<b>0.2241</b>	<b>0.3023</b>

Table 2. Clip and  $C_D$  scores on the YouRefIt dataset.

Method	IoU=0.25	IoU=0.5	IoU=0.75	CLIP $\uparrow$	$C_D \downarrow$
PaliGemma2	47.2	39.5	<b>31.6</b>	0.2465	0.7449
Qwen2.5vl-3b	32.5	32.1	29.2	0.2264	0.8418
Grounding DINO	27.5	27.5	26.3	0.2082	0.7956
Touch-Line-EWL	45.0	35.8	22.0	0.2436	0.5160
Touch-Line-VTL	47.7	36.7	17.4	0.2473	0.5147
Ours, $M_W$	48.2	33.6	20.0	0.2453	0.4735
Ours, $M_H$	42.7	30.0	17.3	0.2449	0.6795
Ours, CAPE	<b>54.5</b>	<b>42.7</b>	25.5	<b>0.2642</b>	<b>0.4403</b>

Table 3. Quantitative results on unseen ISL pointing dataset [16].

precision, contributing to the suboptimal performance at  $IoU = 0.75$ , alongside limitations related to the simpler backbone. At  $IoU = 0.75$ , Grounding DINO with a more complex backbone achieves the best performance.

Further, we evaluate recent LMMs on this task, and they fall significantly short of our model’s performance, despite being among the largest models and trained for visual grounding purposes. These results show using complementary pointing heatmap modalities in addition to image and text is essential. Besides, despite having two separate models, our approach shows better inference-time performance (2.44 fps) compared to Paligemma2 (1.24fps) and Qwen2.5vl (0.5fps). Tab. 2 presents further comparison using two new metrics in this field: CLIP score and  $C_D$ . Higher CLIP scores indicate that our predicted bounding boxes have stronger semantic alignment with the input text, suggesting better object identification and improved bounding box precision. Similarly, our method achieves the best performance under the  $C_D$  metric. This confirms that our model produces more spatially accurate predictions. In Tab. 3, the experimental results on unseen ISL dataset show that our CAPE model surpasses the SOTA Touch-Line model as well as the evaluated LMMs, except for the  $IoU=0.75$  setup. This demonstrates the strong generalization capability of our method. In Fig. 3, we present sample images with predicted and GT bounding boxes. It is clearly seen from LMMs’ predictions, ambiguous text without pointing cues results in poor object localization outside of the pointing area and confusion with similar objects. For instance, Paligemma2 in third example and Qwen2.5vl in second and fourth examples detect wrong objects. On the other hand, the LMMs are strong at providing precise bounding box predictions due to extensive training and a robust backbone. Furthermore, even if Touch-Line model use pointing cues since it follows a pointing line instead of

Setup	Method	IoU=0.25	IoU=0.5	IoU=0.75	CLIP $\uparrow$	$C_D \downarrow$
A	Baseline	71.2	60.1	32.8	0.2469	0.2662
B	A + object center prediction	70.8	61.6	34.6	0.2446	0.2707
C	A + W2F heatmap	68.9	59.7	32.4	0.2451	0.2984
D	A + H2F heatmap	71.9	62.8	33.8	0.2458	0.2694
E	B + W2F heatmap ( $M_W$ )	69.6	60.7	31.5	0.2448	0.2770
F	B + H2F heatmap ( $M_H$ )	72.9	62.3	35.1	0.2457	0.2490
G	A + W2F heatmap + H2F heatmap	70.2	60.2	33.8	0.2448	0.2744
H	Full model - Ensemble of E and F	<b>75.0</b>	<b>65.4</b>	<b>35.7</b>	<b>0.2661</b>	<b>0.2476</b>

Table 4. Ablation study on YouRefIt dataset for our contributions.

heatmaps, this let model to detect other objects in the pointing direction. In contrast, as shown in the examples, CAPE leverages explicit pointing heatmaps, enabling it to more accurately localize the intended referent. Also, it enables to choose final prediction from which of the more accurate model.

## 4.2. Analysis

In this section, we provide a detailed analysis of the optimality of our design choices, including heatmap generation and the ensemble method. The baseline in Tab. 4 (Setup A) refers to a straightforward training of our model without incorporating any of the proposed contributions.

**Object center prediction** The object center prediction improves mAP at  $IoU = 0.5$  and  $IoU = 0.75$  but reduces performance at  $IoU = 0.25$ , as well as CLIP Score and  $C_D$ . Notably, when combined with heatmap (Setup E and F in Tab. 4), it boosts most metrics. The object center prediction alone enhances geometric precision more than semantic alignment, while pairing it with spatially dense cues makes it a more effective refining signal.

	IoU=0.25	IoU=0.5	IoU=0.75	CLIP $\uparrow$	$C_D \downarrow$
Channel-wise input	68.7	58.5	33.4	0.2449	0.2959
Channel-wise feature	72.6	58.4	27.9	0.2457	0.2598
Embedding feature	<b>72.9</b>	<b>62.2</b>	<b>35.1</b>	<b>0.2458</b>	<b>0.2490</b>

Table 5. Ablation of heatmap injection methods using Setup F.

**Where to inject heatmap** We compare several heatmap injection strategies in Tab. 5. The *channel-wise input* method, which feeds the heatmap as a fourth image channel, slightly hurts performance, likely because the heatmap is sparsely represented compared to the image content. Inspired by the head location prompting strategy in [52], the *channel-wise feature* approach adds learned heatmap embeddings to the visual tokens, which is lightweight but provides only limited gains. In contrast, the *embedding feature* setup uses a dedicated CNN to encode the heatmap before fusing it with image and text embeddings, and delivers the strongest improvements across all metrics. The experiments show that effective heatmap injection requires the model to interpret the heatmap as a structured spatial signal rather than as a raw image channel or a light additive bias on visual

Ensemble method	IoU=0.25	IoU=0.5	IoU=0.75	CLIP $\uparrow$	$C_D \downarrow$
Confidence-Only	73.4	64.1	34.1	0.2578	0.2470
CLIP-Only (Top-1)	73.5	64.2	35.4	0.2644	0.2391
CLIP-Only (Top-2 + Threshold=0.95)	73.8	64.1	35.5	0.2646	<b>0.2377</b>
CLIP Fusion	73.6	64.4	34.5	0.2459	0.2422
CAPE	<b>75.0</b>	<b>65.4</b>	<b>35.7</b>	<b>0.2661</b>	0.2476

Table 6. Ablation of different ensemble methods. See Appendix B.

tokens. Besides, the improvements seen in the *embedding feature* method support the gains observed in the main ablations (Setups D and F), where the H2F heatmap consistently strengthens localization and reduces  $C_D$ .

**CLIP-Aware Pointing Ensemble (CAPE)** We evaluate several ensemble strategies for selecting the final referent box from the two pointing models, as shown in Tab. 6. (1) *Confidence-Only* selects the Top-1 prediction with the highest model confidence. (2) *CLIP-Only (Top-1)* chooses the Top-1 box with the highest CLIP similarity. (3) *CLIP-Only (Top-2 + Threshold= 0.95)* additionally considers each model’s top-2 predictions when their confidence scores exceed 0.95. (4) *CLIP Fusion* combines normalized CLIP and confidence scores from the top-2 predictions of both models and selects the box with the highest total score. (5) *CAPE* adaptively switches between (3) and (4): it uses (4) for small objects, where CLIP alone is unreliable, and (3) otherwise. The confidence threshold in (3) and the small-object rule in CAPE were tuned on the validation set to avoid test-time bias. Across these strategies, the hybrid design in CAPE aligns with observations from the main ablations: the two models provide complementary strengths, CLIP-based scoring refines semantic alignment for most objects, and confidence-based cues are especially important for small or ambiguous targets. This contributes to the strong gains seen in the full model (Setup H).

**Failures** In Fig. 4, we present examples of failure cases. In the first example (column), although our model  $M_H$  correctly detects the target object, CAPE selects the prediction from  $M_W$ , which is incorrect in this case. In the second example, both of our models, as well as the Touch-Line models, produce incorrect predictions. However, guided by the pointing information, the predictions from  $M_H$  and  $M_W$  lie on the pointing line and are very close to the GT object. In contrast, the Touch-Line predictions are entirely off the pointing line and significantly distant from the GT object. In the third column, the bottles are extremely difficult to distinguish, even with explicit pointing cues. Nevertheless, the predicted bounding boxes cover other bottles adjacent to the GT one, which is acceptable, as they overlap with the pointing line and are semantically consistent with the text.

## 5. Conclusion

We address Embodied Reference Understanding by overcoming the limitations of relying on a single pointing

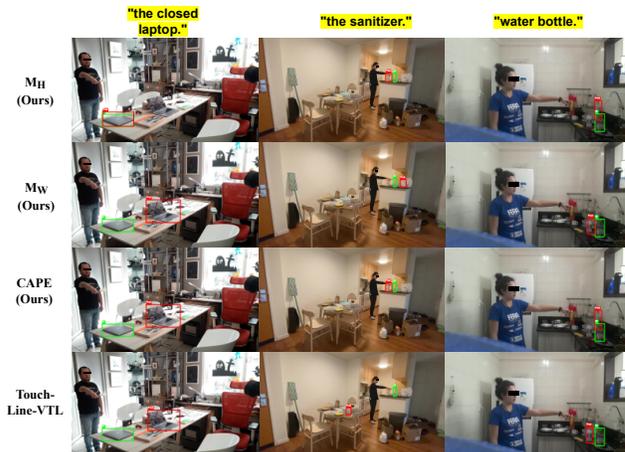


Figure 4. Failure cases from different models. The sample images are from YouRefIt dataset.

line, using our dual-model framework to leverage complementary Gaussian ray heatmaps for richer supervision and improved target detection. Combined through the CLIP-Aware Pointing Ensemble (CAPE) and enhancing the model with an auxiliary object center prediction, our approach achieves strong improvements on the YouRefIt, ISL pointing, and CAESAR datasets. Our results demonstrate that combining multiple pointing cues and in general, multimodal signals, leads to more accurate and robust referent understanding in complex visual scenes. Beyond empirical gains, our work demonstrates a general principle of multimodal cue integration: fusing complementary spatial and semantic signals resolves ambiguities that single modalities cannot, providing a principle for robust referent understanding in complex visual scenes.

**Limitations and Future Work** The model occasionally selects incorrect objects along the pointing direction due to the absence of depth information, and the CNN backbone limits bounding box precision. Future work could incorporate depth or 3D cues and adopt advanced backbones such as Swin Transformers. Additionally, formalizing a theoretical framework for multimodal cue integration and exploring adaptive weighting or attention mechanisms could further improve flexibility and robustness, allowing the system to optimally leverage multiple cues across diverse scenarios.

**Acknowledgment** This work was supported in part by the European Union’s Horizon research and innovation program, project Meetween (101135798) and project DVPS (Diversibus Viis Plurima Solvo) (101213369), and KIT Campus Transfer GmbH (KCT) under contract from Carnegie-AI LLC.

## References

- [1] Marah Abidin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J Hewett, Mojan Javaheripi, Piero Kauffmann, et al. Phi-4 technical report. *arXiv preprint arXiv:2412.08905*, 2024. 2
- [2] Nada H Alalyani. *Embodied Multimodal Referring Expressions Generation*. PhD thesis, Colorado State University, 2024. 2
- [3] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 1(2):3, 2023. 1, 2
- [4] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibong Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 1, 6
- [5] Loris Bazzani, Marco Cristani, Diego Tosato, Michela Farenzena, Giulia Paggetti, Gloria Menegaz, and Vittorio Murino. Social interactions by visual focus of attention in a three-dimensional environment. *Expert Systems*, 30(2):115–127, 2013. 2
- [6] Lucas Beyer, Andreas Steiner, André Susano Pinto, Alexander Kolesnikov, Xiao Wang, Daniel Salz, Maxim Neumann, Ibrahim Alabdulmohsin, Michael Tschannen, Emanuele Bugliarello, et al. Paligemma: A versatile 3b vlm for transfer. *arXiv preprint arXiv:2407.07726*, 2024. 1, 2
- [7] Brice Burger, Isabelle Ferrané, Frédéric Lerasle, and Guillaume Infantes. Two-handed gesture recognition and fusion with speech to command a robot. *Autonomous Robots*, 32(2):129–147, 2012. 2
- [8] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7291–7299, 2017. 1, 2
- [9] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE transactions on pattern analysis and machine intelligence*, 43(1):172–186, 2019. 1
- [10] Francesca Capozzi, Cigdem Beyan, Antonio Pierro, Atesh Koul, Vittorio Murino, Stefano Livi, Andrew P Bayliss, Jelena Ristic, and Cristina Becchio. Tracking the leader: Gaze behavior in group interactions. *Iscience*, 16:242–249, 2019. 2
- [11] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 2, 5, 6
- [12] Jierun Chen, Fangyun Wei, Jinjing Zhao, Sizhe Song, Bohuai Wu, Zhuoxuan Peng, S-H Gary Chan, and Hongyang Zhang. Revisiting referring expression comprehension evaluation in the era of large multimodal models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 513–524, 2025. 2
- [13] Yixin Chen, Qing Li, Deqian Kong, Yik Lun Kei, Song-Chun Zhu, Tao Gao, Yixin Zhu, and Siyuan Huang. Yourefit: Embodied reference understanding with language and gesture. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1385–1395, 2021. 1, 2, 4, 5, 6, 3
- [14] Tianheng Cheng, Lin Song, Yixiao Ge, Wenyu Liu, Xinggang Wang, and Ying Shan. Yolo-world: Real-time open-vocabulary object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16901–16911, 2024. 2
- [15] Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. Unifying vision-and-language tasks via text generation. In *International Conference on Machine Learning*, pages 1931–1942. PMLR, 2021. 2
- [16] Stefan Constantin, Fevziye Irem Eyiokur, Dogucan Yaman, Leonard Bärmann, and Alex Waibel. Interactive multimodal robot dialog using pointing gesture recognition. In *European conference on computer vision*, pages 640–657. Springer, 2022. 2, 5, 7, 10
- [17] Stefan Constantin, Fevziye Irem Eyiokur, Dogucan Yaman, Leonard Bärmann, and Alex Waibel. Multimodal error correction with natural language and pointing gestures. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1976–1986, 2023. 2
- [18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186, 2019. 3
- [19] Mudasir A Ganaie, Minghui Hu, Ashwani Kumar Malik, Muhammad Tanveer, and Ponnuthurai N Suganthan. Ensemble deep learning: A review. *Engineering Applications of Artificial Intelligence*, 115:105151, 2022. 2
- [20] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *International Journal of Computer Vision*, 132(2): 581–595, 2024. 2
- [21] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014. 2
- [22] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*, 2021. 2
- [23] Anshul Gupta, Samy Tafasca, and Jean-Marc Odobez. A modular multimodal architecture for gaze target prediction: Application to privacy-sensitive settings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5041–5050, 2022. 2
- [24] Anshul Gupta, Pierre Vuillecard, Arya Farkhondeh, and Jean-Marc Odobez. Exploring the zero-shot capabilities of vision-language models for improving gaze following. In

- Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 615–624, 2024. 2
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3
- [26] Shuting He, Henghui Ding, Chang Liu, and Xudong Jiang. Grec: Generalized referring expression comprehension. *arXiv preprint arXiv:2308.16182*, 2023. 2
- [27] Hieu Hoang, Huda Khayrallah, and Marcin Junczys-Dowmunt. On-the-fly fusion of large language models and machine translation. *arXiv preprint arXiv:2311.08306*, 2023. 2
- [28] Hartwig Holzapfel, Kai Nickel, and Rainer Stiefelwagen. Implementation and evaluation of a constraint-based multimodal fusion system for speech and 3d pointing gestures. In *Proceedings of the 6th international conference on Multimodal interfaces*, pages 175–182, 2004. 2
- [29] Md Mofijul Islam, Reza Mirzaiee, Alexi Gladstone, Haley Green, and Tariq Iqbal. Caesar: An embodied simulator for generating multimodal referring expression datasets. *Advances in Neural Information Processing Systems*, 35: 21001–21015, 2022. 2, 5, 10
- [30] Md Mofijul Islam, Alexi Gladstone, Riashat Islam, and Tariq Iqbal. Eqa-mx: Embodied question answering using multimodal expression. In *The Twelfth International Conference on Learning Representations*, 2023.
- [31] Aman Jain, Teruhisa Misu, Kentaro Yamada, and Hitomi Yanaka. Gesnavi: gesture-guided outdoor vision-and-language navigation. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 290–295, 2024. 2
- [32] Kritika Johari, Christopher Tay Zi Tong, Vigneshwaran Subbaraju, Jung-Jae Kim, and U-Xuan Tan. Gaze assisted visual grounding. In *International Conference on Social Robotics*, pages 191–202. Springer, 2021. 2
- [33] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr-modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1780–1790, 2021. 2, 5
- [34] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 787–798, 2014. 2
- [35] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [36] Alexander Kroner, Mario Senden, Kurt Driessens, and Rainer Goebel. Contextual encoder–decoder network for visual saliency prediction. *Neural Networks*, 129:261–270, 2020. 1
- [37] Feng Li, Hao Zhang, Shilong Liu, Jian Guo, Lionel M Ni, and Lei Zhang. Dn-detr: Accelerate detr training by introducing query denoising. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13619–13627, 2022. 5, 6
- [38] Yang Li, Xiaoxue Chen, Hao Zhao, Jiangtao Gong, Guyue Zhou, Federico Rossano, and Yixin Zhu. Understanding embodied reference with touch-line transformer. In *The Eleventh International Conference on Learning Representations*, 2023. 1, 2, 3, 4, 5, 6, 7
- [39] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European conference on computer vision*, pages 38–55. Springer, 2024. 1, 2, 5, 6
- [40] 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. 3
- [41] Ziyang Lu, Yunqiang Pei, Guoqing Wang, Peiwei Li, Yang Yang, Yinjie Lei, and Heng Tao Shen. Scanner: Interactive 3d visual grounding based on embodied reference understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3936–3944, 2024. 2
- [42] Atharv Mahesh Mane, Dulanga Weerakoon, Vigneshwaran Subbaraju, Sougata Sen, Sanjay E Sarma, and Archan Misra. Ges3vig: Incorporating pointing gestures into language-based 3d visual grounding for embodied reference understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 9017–9026, 2025. 2, 3
- [43] Depu Meng, Xiaokang Chen, Zejia Fan, Gang Zeng, Houqiang Li, Yuhui Yuan, Lei Sun, and Jingdong Wang. Conditional detr for fast training convergence. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3651–3660, 2021. 5, 6
- [44] Varun K Nagaraja, Vlad I Morariu, and Larry S Davis. Modeling context between objects for referring expression understanding. In *European Conference on Computer Vision*, pages 792–807. Springer, 2016. 2
- [45] Akira Oyama, Shoichi Hasegawa, Hikaru Nakagawa, Akira Taniguchi, Yoshinobu Hagiwara, and Tadahiro Taniguchi. Exophora resolution of linguistic instructions with a demonstrative based on real-world multimodal information. In *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 2617–2623. IEEE, 2023. 2
- [46] Maria Pateraki, Haris Baltzakis, and Panos Trahanias. Visual estimation of pointed targets for robot guidance via fusion of face pose and hand orientation. *Computer Vision and Image Understanding*, 120:1–13, 2014. 2
- [47] Renjie Pi, Lewei Yao, Jiahui Gao, Jipeng Zhang, and Tong Zhang. Perceptiongpt: Effectively fusing visual perception into llm. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 27124–27133, 2024. 2
- [48] Yanyuan Qiao, Chaorui Deng, and Qi Wu. Referring expression comprehension: A survey of methods and datasets. *IEEE Transactions on Multimedia*, 23:4426–4440, 2020. 2

- [49] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 4
- [50] Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. *arXiv preprint arXiv:1904.09237*, 2019. 5
- [51] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. 2
- [52] Fiona Ryan, Ajay Bati, Sangmin Lee, Daniel Bolya, Judy Hoffman, and James M Rehg. Gaze-llc: Gaze target estimation via large-scale learned encoders. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 28874–28884, 2025. 2, 7
- [53] Cheng Shi and Sibe Yang. Spatial and visual perspective-taking via view rotation and relation reasoning for embodied reference understanding. In *European Conference on Computer Vision*, pages 201–218. Springer, 2022. 2, 6
- [54] Kevin Smith, Sileye O Ba, Jean-Marc Odobez, and Daniel Gatica-Perez. Tracking the visual focus of attention for a varying number of wandering people. *IEEE transactions on pattern analysis and machine intelligence*, 30(7):1212–1229, 2008. 2
- [55] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014. 5
- [56] Andreas Steiner, André Susano Pinto, Michael Tschannen, Daniel Keysers, Xiao Wang, Yonatan Bitton, Alexey Gritsenko, Matthias Minderer, Anthony Sherbondy, Shangbang Long, et al. Paligemma 2: A family of versatile vlms for transfer. *arXiv preprint arXiv:2412.03555*, 2024. 1, 6
- [57] Rainer Stiefelhagen, Michael Finke, Jie Yang, and Alex Waibel. From gaze to focus of attention. In *International Conference on Advances in Visual Information Systems*, pages 765–772. Springer, 1999. 2
- [58] Rainer Stiefelhagen, Jie Yang, and Alex Waibel. Modeling focus of attention for meeting indexing. In *Proceedings of the seventh ACM international conference on Multimedia (Part 1)*, pages 3–10, 1999.
- [59] Rainer Stiefelhagen, Jie Yang, and Alex Waibel. Estimating focus of attention based on gaze and sound. In *Proceedings of the 2001 workshop on Perceptive user interfaces*, pages 1–9, 2001.
- [60] Rainer Stiefelhagen, Christian Fugen, R Gieselmann, Hartwig Holzapfel, Kai Nickel, and Alex Waibel. Natural human-robot interaction using speech, head pose and gestures. In *2004 IEEE/RSJ international conference on intelligent robots and systems (IROS)(IEEE cat. no. 04CH37566)*, pages 2422–2427. IEEE, 2004. 2
- [61] Bernhard Suhm, Brad Myers, and Alex Waibel. Model-based and empirical evaluation of multimodal interactive error correction. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 584–591, 1999. 1
- [62] Samy Tafasca, Anshul Gupta, and Jean-Marc Odobez. Child-play: A new benchmark for understanding children’s gaze behaviour. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20935–20946, 2023. 2
- [63] Samy Tafasca, Anshul Gupta, Victor Bros, and Jean-Marc Odobez. Toward semantic gaze target detection. *Advances in neural information processing systems*, 37:121422–121448, 2024.
- [64] Samy Tafasca, Anshul Gupta, and Jean-Marc Odobez. Sharingan: A transformer architecture for multi-person gaze following. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2008–2017, 2024.
- [65] Kosei Tanada, Shigemichi Matsuzaki, Kazuhito Tanaka, Shintaro Nakaoka, Yuki Kondo, and Yuto Mori. Pointing gesture understanding via visual prompting and visual question answering for interactive robot navigation. In *First Workshop on Vision-Language Models for Navigation and Manipulation at ICRA 2024*, 2024.
- [66] Francesco Tonini, Cigdem Beyan, and Elisa Ricci. Multimodal across domains gaze target detection. In *Proceedings of the 2022 International Conference on Multimodal Interaction*, pages 420–431, 2022.
- [67] Francesco Tonini, Nicola Dall’Asen, Cigdem Beyan, and Elisa Ricci. Object-aware gaze target detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 21860–21869, 2023.
- [68] Roberto Valenti, Nicu Sebe, and Theo Gevers. Combining head pose and eye location information for gaze estimation. *IEEE Transactions on Image Processing*, 21(2):802–815, 2011. 2
- [69] 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*, 30, 2017. 3
- [70] Yang Xiao, Zhijun Zhang, Aryel Beck, Junsong Yuan, and Daniel Thalmann. Human–robot interaction by understanding upper body gestures. *Presence*, 23(2):133–154, 2014. 2
- [71] Hee-Deok Yang, A-Yeon Park, and Seong-Whan Lee. Gesture spotting and recognition for human–robot interaction. *IEEE Transactions on robotics*, 23(2):256–270, 2007. 2
- [72] Jie Yang, Rainer Stiefelhagen, Uwe Meier, and Alex Waibel. Visual tracking for multimodal human computer interaction. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 140–147, 1998. 1
- [73] Yaokun Yang and Feng Lu. Gaze target detection based on head-local-global coordination. In *European Conference on Computer Vision*, pages 305–322. Springer, 2024. 2
- [74] Zhengyuan Yang, Boqing Gong, Liwei Wang, Wenbing Huang, Dong Yu, and Jiebo Luo. A fast and accurate one-stage approach to visual grounding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4683–4693, 2019. 6
- [75] Zhengyuan Yang, Tianlang Chen, Liwei Wang, and Jiebo Luo. Improving one-stage visual grounding by recursive subquery construction. In *European conference on computer vision*, pages 387–404. Springer, 2020. 6

- [76] Lewei Yao, Jianhua Han, Youpeng Wen, Xiaodan Liang, Dan Xu, Wei Zhang, Zhenguo Li, Chunjing Xu, and Hang Xu. Detclip: Dictionary-enriched visual-concept paralleled pre-training for open-world detection. *Advances in Neural Information Processing Systems*, 35:9125–9138, 2022. [2](#)
- [77] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. Mattnet: Modular attention network for referring expression comprehension. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1307–1315, 2018. [2](#)
- [78] Shoubin Yu, Jaehong Yoon, and Mohit Bansal. Crema: Generalizable and efficient video-language reasoning via multimodal modular fusion. *arXiv preprint arXiv:2402.05889*, 2024. [2](#)
- [79] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14393–14402, 2021. [2](#)
- [80] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11975–11986, 2023. [4](#)
- [81] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*, 2022. [5](#), [6](#)
- [82] Yuzhen Zhang, Jingjing Liu, and Wenjuan Shen. A review of ensemble learning algorithms used in remote sensing applications. *Applied Sciences*, 12(17):8654, 2022. [2](#)
- [83] Zongzheng Zhang, Xuchong Qiu, Boran Zhang, Guantian Zheng, Xunjiang Gu, Guoxuan Chi, Huan-ang Gao, Leichen Wang, Ziming Liu, Xinrun Li, et al. Delving into mapping uncertainty for mapless trajectory prediction. *arXiv preprint arXiv:2507.18498*, 2025. [2](#)
- [84] Hao Zhao, Ming Lu, Anbang Yao, Yurong Chen, and Li Zhang. Learning to draw sight lines. *International Journal of Computer Vision*, 128(5):1076–1100, 2020. [2](#)
- [85] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. [5](#), [6](#)