

# CLUE: Crossmodal disambiguation via Language-vision Understanding with attEntion

Mouad Abrini and Mohamed Chetouani

*Institut des Systèmes Intelligents et de Robotique (ISIR), Sorbonne Université  
Paris, France*

{mouad.abrini, mohamed.chetouani}@isir.upmc.fr

**Abstract**—With the increasing integration of robots into daily life, human-robot interaction has become more complex and multifaceted. A critical component of this interaction is Interactive Visual Grounding (IVG), through which robots must interpret human intentions and resolve ambiguity. Existing IVG models generally lack a mechanism to determine when to ask clarification questions, as they implicitly rely on their learned representations. CLUE addresses this gap by converting the VLM’s cross-modal attention into an explicit, spatially grounded signal for deciding when to ask. We extract text to image attention maps and pass them to a lightweight CNN to detect referential ambiguity, while a LoRA fine-tuned decoder conducts the dialog and emits grounding location tokens. We train on a real-world interactive dataset for IVG, and a mixed ambiguity set for the detector. With InViG-only supervision, our model surpasses a state-of-the-art method while using parameter-efficient fine-tuning. Similarly, the ambiguity detector outperforms prior baselines. Overall, CLUE turns the internal cross-modal attention of a VLM into an explicit, spatially grounded signal for deciding when to ask. The data and code are publicly available at: [mouadabrini.github.io/clue/](https://mouadabrini.github.io/clue/)

## I. INTRODUCTION

Robots operating in human spaces must interpret underspecified language instructions and act safely. A central capability is interactive visual grounding (IVG): given a natural-language instruction (e.g., “Pick up the apple”), the robot must identify the intended referent and, when necessary, ask for clarification (see Fig 1). While recent IVG systems can localize objects from referring expressions [1], [2], [3], [4] and even engage in dialog [5], [6], [7], [8], an open question is how to detect that the current instruction is ambiguous in the specific visual scene. Classical referring expression comprehension (REC) methods, from modular attention networks to modern transformer detectors, largely assume uniqueness, optimizing for a single best match [1], [2]. Interactive approaches in HRI typically decide to ask using heuristics (e.g., multiple plausible candidates) or token-level uncertainty (e.g., entropy/confidence from a language or policy model), as seen in planners that query based on grounding confidence [9], and LLM-driven ambiguity classifiers like CLARA [10]. Recent dialog-trained systems further learn when to ask from demonstration, e.g., ClearVQA [11] and uncertainty-timed clarification [12]. These signals can be effective, but are policy-level (answer vs. ask) and only indirectly grounded in the spatial structure of the scene: they




Fig. 1. Problem illustration: when an instruction is underspecified, the robot should detect it and ask for clarification (AI generated, then edited)

rarely indicate where confusion arises or which distractors compete for attention. We argue that ambiguity in grounded instructions arises in the joint image-language representation, where textual tokens align with multiple plausible visual candidates. In such cross-modal spaces, the decoder’s text queries attend to visual patches that encode scene structure and attributes. When several objects fit the instruction, the alignment becomes non-selective, distributing mass across competing regions. We exploit this property by extracting the model’s cross-modal self-attention weights and training a convolutional network on these maps to classify ambiguity. This turns a latent alignment pattern into a spatial signal that both (i) indicates when the instruction is underspecified and (ii) localizes where the confusion lies for better interpretability. The contributions are structured as follows:

- Spatial ambiguity detection: a CNN over cross-modal attention maps that explicitly detects and localizes referential ambiguity.
- A synthetic dataset generated using the Isaac Sim Simulator [13] for multimodal ambiguity detection.
- End-to-end disambiguation (IVG): a fine-tuned VLM on InViG only [14] (real-world dataset) yielding state-of-the-art IVG performance, outperforming TiO [15] trained on InViG-only.

## II. RELATED WORK

### A. Interactive Visual Grounding and Clarification Dialogues

Ambiguity in referring expressions has long been addressed by interactive systems that can detect unclear references and ask users for clarification. Early human-robot

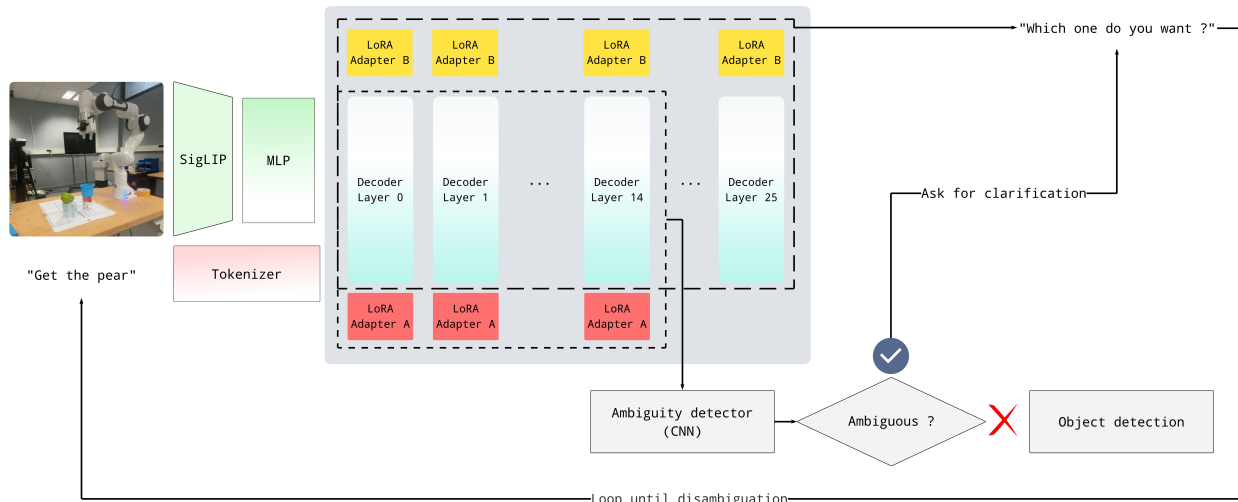


Fig. 2. Overall CLUE architecture. An RGB image is encoded by SigLIP and projected by an MLP. The text prefix is tokenized and passed with the image tokens into a Gemma2 decoder equipped with LoRA adapters. The decoder both (i) autoregressively generates clarification questions and (ii) exposes cross-modal attention maps from mid-layers to a lightweight CNN ambiguity detect (Fig. 4). If the detector predicts ambiguous, the model asks a follow-up and updates the dialog context; otherwise it emits location tokens and triggers object detection. The loop repeats until disambiguation.

interaction work integrated dialog or gestures to resolve ambiguity. For example, [5] combined user pointing gestures with a dialog system to disambiguate spoken commands, and [6] used a referring expression model to find candidate objects and then engaged in conversation to identify the correct target [7]. Similarly, [16] proposed generating spatial heatmaps and follow-up questions to locate the intended object, and [8] showed that clarification questions from a robot improved the grounding of object references. In essence, if a user says “Bring me that cup,” an interactive agent must recognize the referential ambiguity (e.g. multiple cups are present) and ask a targeted question to clarify which cup is meant. Modern systems explicitly predict such ambiguity and generate clarification questions. For instance, [9] presents an attribute-guided POMDP planner that uses vision-language grounding confidence to decide when to ask a question to differentiate between look-alike objects. Recent vision-language models are explicitly trained to ask when a question is ambiguous: [11] introduces the ClearVQA benchmark and methods that detect ambiguity, generate targeted clarification questions, and then answer after interaction. Likewise, [12] analyzed human dialogs and model uncertainty to determine the “right time” to ask for clarification, aligning the agent’s questions with moments of referential uncertainty. [14] introduce InViG, a large-scale benchmark with  $\sim 500K$  human-robot interactions specifically targeting ambiguity in open-ended scenes. It evaluates how an agent should respond with clarifying questions in a natural manner rather than relying on fixed templates. [15] builds on that line by proposing TiO, an end-to-end model for interactive visual grounding in the wild and learns to gather information actively during the interaction. The main limitation of these kinds of methods is that they rely heavily on the statistical presence of the dialog rounds. For example, the InViG dataset contains at least one

dialog round, so the model will always ask a clarification question, no matter what, even if the initial instruction is unambiguous. All these approaches typically rely on model confidence scores or predicted likelihoods to trigger clarifications. In contrast, our method enables triggering clarifications using a spatial, cross-modal attention signal from a VLM, rather than token likelihoods or confidence heuristics.

### B. Ambiguity Detection in Grounded Instructions

To detect ambiguity, models typically trigger clarification via heuristics (e.g., multiple candidate objects) or learned uncertainty (token/policy entropy, confidence). For example, the model proposed in [10] classifies commands as clear/ambiguous/infeasible using an LLM-based confidence signal and asks for clarification when needed. Interactive grounding also uses dialog/gestures or heatmap/question strategies to resolve reference ambiguity [5], [6], [7], [16], [8], and recent work learns when to ask from data; for instance, by planning over grounding confidence [9] or using model uncertainty to guide questioning [12]. To evaluate such capabilities, benchmarks like ClearVQA [11] have been introduced.

In an earlier work [17] for HRI, the authors used an existing object detector (DETR [18]) connected to a dense image captioning model (DenseCAP [19]), on which they applied GradCAM [20]. Although this method is quite promising, it is limited by the performance of the sub-components themselves and is not specifically tailored for ambiguity detection. Furthermore, even if we wanted to finetune (fully supervised) this further using the saliency maps, we would be limited by the computational complexity of GradCAM. In fact, if we define a loss that depends on the saliency map and backpropagate, we would be computing second order gradients, which is slow and memory hungry.

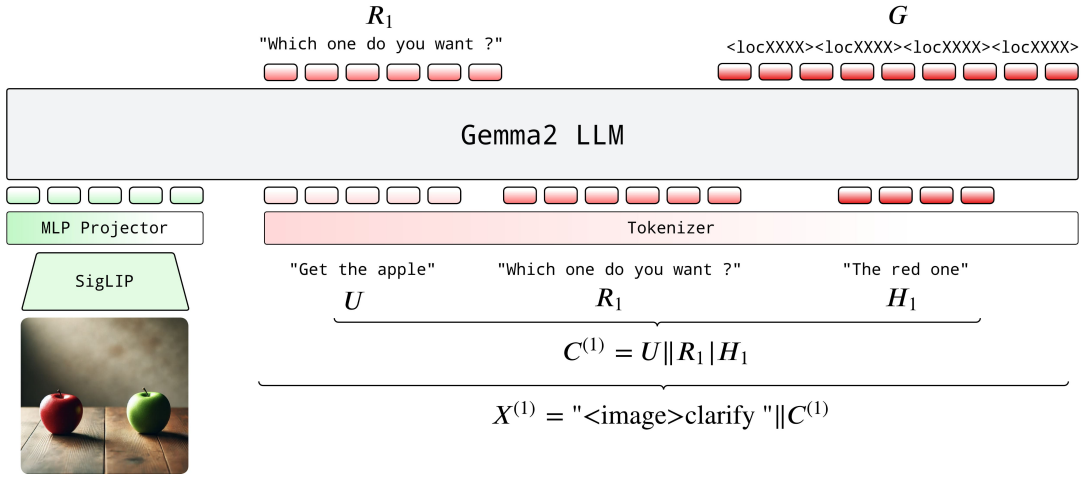


Fig. 3. IVG with notation. An RGB image (AI generated in this example) is encoded by SigLIP and projected to the decoder via an MLP, the text prefix is tokenized and concatenated with image tokens. The Gemma2 LLM takes  $X^{(1)} = \text{"<image>clarify " || } C^{(1)}$ , where the running context after the first turn is  $C^{(1)} = U || R_1 || H_1$  (initial user request  $U$ , assistant question  $R_1$ , human reply  $H_1$ ). The model autoregressively generates either a clarification segment (left stream, labeled  $R_1$ ) or location tokens  $G = \text{"<locXXXX> ... (right stream) that decode to a bounding box. Green blocks denote image features, pink blocks denote text tokens.$

Beyond that, recent work [21] fine-tuned an end-to-end foundation model (Molmo 7B [22]) to predict a yes/no response for ambiguity detection. However, for the balanced class ratio dataset they used, the performance was not very high (around 0.68 F1).

[23] introduced a benchmark for ambiguity detection, named “Focus Ambiguity”, which is precisely the type of ambiguity we are tackling in this paper. They showed that even large models do not achieve high performance when it comes to ambiguity classification.

### III. METHODS

Whenever applicable to fine-tuning the VLM’s language model, we use Low-Rank Adaptation (LoRA) [24]. Concretely, we freeze the visual encoder, the multimodal projector, and the tokenizer, and we insert LoRA adapters into the decoder’s projection layers (attention q/k/v/o and MLP up/down/gate). We use two different adapters for each task: Adapter A for ambiguity detection and adapter B for IVG, as represented in Fig 2.

#### A. CLUE architecture for IVG

1) **Model:** We implement IVG as a transformer decoder based on PaliGemma mix 2 (paligemma2-3b-mix-448) [25]. This model was pre-trained to handle different tasks by using a different initial token. For example, if one needs to do REC, the text instruction should be “DETECT the red cup”. If we want to do image captioning, we should use “CAPTION EN”. Since it was also trained for open-ended object detection, its vocabulary is extended with discrete tokens  $\{\text{"<lock>"}\}_{k=0}^{1024}$  that encode normalized box coordinates in  $[0, 1024)$ . In this line, we reserve a special conditioning token “CLARIFY” for our task: we attend to image tokens and the running text dialog context and autoregressively emit either a clarification question or a sequence of grounding location tokens. Figure 3 shows an illustration of how the model is used.

Let  $C$  be the running dialog context (the text prefix). Concretely,  $C^{(k)}$  is the concatenation (noted  $||$ ) of the initial user request  $U$  plus all turns (assistant/robot questions  $R_i$  and human answers  $H_i$ ) up to turn  $k$ . At token step  $t$ , conditioned on image  $I$ , the linearized prefix  $X$ , and the previously generated suffix tokens  $y_{<t}$ , the decoder predicts the next token  $y_t$ .

The interactive exchange is a sequence of assistant/robot (model) prompts and user replies. A typical dialog is defined as:

$$D = \{(R_1, H_1), (R_2, H_2), \dots, (R_K, H_K)\}$$

For training, We construct a supervised dataset  $S$  of prefix  $\rightarrow$  target pairs by cumulatively revealing pairs and alternating the prediction target:

$$\begin{aligned} (U \rightarrow R_1), \quad (U || R_1 H_1 \rightarrow R_2), \\ \dots, \\ (U || R_1 H_1 \dots R_{K-1} H_{K-1} \rightarrow R_K) \end{aligned}$$

The final grounding target would be:

$$(U || R_1 H_1 \dots R_K H_K \rightarrow G)$$

Each training example is linearized as input

$$X = (\text{"<image>clarify " || prefix})$$

and target suffix  $Y$  (either a textual response  $R_t$  or localization tokens  $G$ ). We apply token-level cross-entropy only over the suffix tokens:

$$\mathcal{L}_{IVG}(\theta) = \frac{1}{T} \sum_{t=1}^T [-\log p_{\theta}(y_t | I, X, y_{<t})]$$

Where  $T$  is the number of tokens in the target suffix  $Y$  for this example.

At test time, given the initial user instruction  $U$  and image  $I$ , we set  $C_0 = U$  and iterate. Algorithm 1 shows the pseudocode.

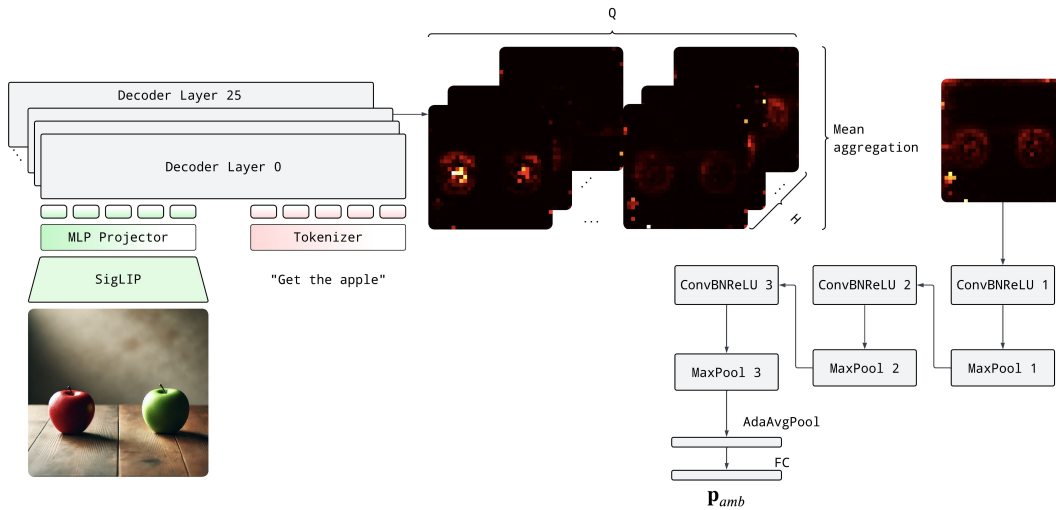


Fig. 4. Ambiguity detector. The image is encoded by SigLIP and projected with an MLP. The text prefix (e.g. "Get the apple") is tokenized and passed to the Gemma2 decoder. From the 14th layer, we read text to image cross-attention for the selected prefix query tokens  $Q$  and keep the first 1024 keys (image tokens), yielding per-token  $32 \times 32$  maps that are mean-aggregated into a single spatial map. This map is passed to a lightweight CNN to produce the ambiguity probability  $p_{amb}$ .

### Algorithm 1 Interactive Visual Grounding (IVG) Inference

**Require:** Image  $I$ , initial user request  $U$

**Ensure:** Bounding box decoded from a valid  $\langle loc \rangle$  sequence

- 1:  $C^{(0)} \leftarrow U$   $\triangleright$  running dialogue context / prefix
- 2: **for**  $k = 0, 1, 2, \dots$  **do**
- 3:  $X^{(k)} \leftarrow \langle \text{image} \rangle \text{clarify } \langle \rangle \parallel C^{(k)}$
- 4:  $\hat{Y}^{(k)} \leftarrow \text{Generate}(X^{(k)})$   $\triangleright$  multiple tokens
- 5: **if**  $\hat{Y}^{(k)}$  contains a valid  $\langle loc \rangle$  sequence **then**
- 6:     **return** DecodeBox( $\hat{Y}^{(k)}$ )
- 7: **else**
- 8:      $\hat{R}_k \leftarrow \text{ExtractQuestion}(\hat{Y}^{(k)})$
- 9:      $\hat{H}_k \leftarrow \text{GetHumanReply}(\hat{R}_k)$
- 10:      $C^{(k+1)} \leftarrow C^{(k)} \parallel \langle \text{assistant: } \rangle \parallel \hat{R}_k \parallel \langle \text{user: } \rangle \parallel \hat{H}_k$
- 11:     **end if**
- 12: **end for**

2) **Dataset:** We train and evaluate on the InViG-21K [14] subset of the InViG benchmark for interactive visual grounding. InViG is a large-scale corpus for open-ended clarification and grounding, built in two stages: (i) 21K human-human, visually grounded dialogs collected via an online chat tool, and (ii) 500K auto-generated human-robot disambiguation dialogs derived from the 21K seed to scale data and linguistic diversity. We use only the human-human InViG-21K subset.

### B. CLUE architecture for Ambiguity detection

1) **Model:** By inspecting the pretrained detector's (same base model) cross-modal attention, we observed that certain layers zero-shot produce sparse attention maps (Fig. 5). We hypothesize that this stems from the base model's multi-object localization training: exposure to naturally ambiguous references encourages an implicit ambiguity signal that

manifests in attention. On its own, the signal is insufficient for ambiguity detection because pretraining provides little explicit supervision for ambiguity. The idea is to use it as a starting point and fine-tune a model that is fully supervised on ambiguity detection. The architecture is represented in Fig. 4. Let the multimodal decoder backbone be a stack of  $L \in \mathbb{N}$  transformer blocks  $B_k$ . We define a half-depth variant  $f^{1/2}$  by selecting a subset of blocks  $\mathcal{S} = \{0, 1, \dots, 14\}$  (This choice is justified in Section IV-A.1). This gives,

$$f^{(1/2)} = B_{14} \circ \dots \circ B_1 \circ B_0$$

Given an image  $I$  and instruction  $x$ , we run a forward pass  $f^{1/2}$  and extract the attention tensor at the chosen block:  $A \in \mathbb{R}^{H \times Q \times K}$ , where  $H$  is the number of heads in the Grouped Query self-attention [26], [27],  $Q$  is the query length (text tokens), and  $K$  is the key length (image+text tokens). We create a mask over query positions  $\mathcal{Q}_x = \{q \in 1, \dots, Q : \text{token}(q) \notin \mathcal{E}\}$ , excluding special tokens  $\mathcal{E}$  (conditioning token, eos, pad). To prevent heads with high raw attention magnitudes from dominating the map, we perform per-head L1 renormalization over the image attention. Let  $L_{img}$  be the number of image tokens (patches). For each head  $h$  and query  $q \in \mathcal{Q}_x$ , we restrict keys to the image region  $k \in \{1, \dots, L_{img}\}$  and normalize:

$$\bar{A}_{h,q,k} = \frac{A_{h,q,k}}{\sum_{j=1}^{L_{img}} A_{h,q,j} + \epsilon} \quad (1)$$

This ensures that the attention distribution sums to 1 over the image for every head. We then mean-aggregate these normalized maps:

$$v = \frac{1}{|\mathcal{Q}_x|H} \sum_{h=1}^H \sum_{q \in \mathcal{Q}_x} \bar{A}_{h,q,:} \in \mathbb{R}^{L_{img}} \quad (2)$$




Fig. 5. Text→ image attention maps from layer 14 ( $32 \times 32$  patches) for different instructions. (a) Input image. (b) “detect the apple” yields two peaks over both apples (ambiguous). (c) “detect the red apple” concentrates on the left apple. (d) “detect apple on the right” concentrates on the right apple. Colors show min-max normalized attention weight.

We use mean aggregation for query-token fusion because PaliGemma applies bidirectional self-attention over the prefix (instruction) and a causal mask only on the suffix (See Fig. 2 in [28]). Thus, the instruction tokens (e.g., apple, left) fully attend to one another, yielding contextualized token embeddings by the time we read cross modal attention to image patches. Averaging these queries post Softmax is therefore a simple, stable and differentiable way to capture the composed semantics that already emerge from their mutual interaction.

Let  $v_{img} \in \mathbb{R}_{img}^K$  be the prefix of  $v$  corresponding to images keys (patch tokens). We reshape to a spatial grid,  $32 \times 32$  in our case. After per-sample min-max normalization, a lightweight probe  $g_\theta : [0, 1]^{H_M \times W_M} \rightarrow [0, 1]$  maps the attention map  $M$  to the probability of ambiguity  $p_{amb}$  (See Fig. 4). Given labels  $y \in \{0, 1\}$  (1=ambiguous, 0=unambiguous), we minimize binary cross-entropy.

2) **Dataset:** For this task, we generated synthetic data using the Isaac Sim simulator [13], consisting of 2000 tabletop scenes populated with randomized YCB objects. To construct the corresponding textual descriptions, we enforced a constraint during randomization such that each scene contains at least two visually similar objects. In this setting, ambiguous instructions simply take the form “GET THE <DUPLICATED\_OBJECT>”, whereas unambiguous instructions correspond to expressions such as “GET THE <UNIQUE\_OBJECT>”. Moreover, we combined it with another dataset (IT2P) from [16], which has 477 images, totaling 1847 image-text samples. In total, the training dataset (Dataset 1) includes around 4000 image–instruction pairs. We further set aside 100 manually labeled real-world images from InVIG [14] for evaluation. Each image has both an ambiguous and an unambiguous label, giving a total of 200 image-instruction samples (Dataset 2). The models were not trained on this dataset and we consider it Out-of-distribution (OOD).

## IV. RESULTS AND DISCUSSION

### A. Ambiguity detection

1) **Ablation study:** To understand how the choice of the decoder layer affects the attention signal and to justify the choice of layer 14, we freeze all parts of the model except



Fig. 6. Samples from the ambiguity detection datasets. (Left) Synthetic data from Isaac Sim, (Middle) real-world data (Dataset 2 (OOD), (Right) IT2P: dataset from [16]. Dataset 1 is a combination of samples from (Left) and (Right)

the CNN head and then train the model while changing the decoder layer used to extract that signal. Figure 7 shows the F1 score across the gradually increasing number of decoder layers. As can be seen, the F1 score improves steadily from the early stages, peaking significantly around Layer 14 (0.726), where the model achieves its optimal balance of Precision and Recall. However, a performance decrease is observed at the final stage, Layer 25, where the F1 score drops to 0.596.

2) **Setup:** We evaluated the ambiguity detector on the synthetic dataset of instruction-image pairs, in which the scenes contain visually similar distractors. An example is labeled ambiguous if the instruction under-specifies a unique referent; otherwise, it is unambiguous. We use a 60/15/25 split and identical preprocessing for all methods. We used the AdamW optimizer with a learning rate of  $5e-6$  for the LoRA adapters and  $1e-4$  for the CNN head. We applied a weight decay of  $1e-4$ . Training was done on a single A100 GPU.

3) **Baselines:** Because the pretraining for REC was conditioned on the token “detect”, using it gave the most interpretable signal (see Fig. 5). We therefore used “detect” for the CNN-based model and some autoregressive variants (where the decoder has to autoregressively generate “yes” or “no”), and introduced a new token, “disambiguate”, unused during pretraining, for the remaining variants. In detail, we performed a comparative analysis of several model variants. To ensure clarity, we adopt a naming convention denoting

TABLE I  
PERFORMANCE COMPARISON OF DIFFERENT MODELS ON DATASET 1 AND DATASET 2.

Model	Dataset 1 (Test)				Dataset 2 (OOD)			
	Acc	Prec	Rec	F1	Acc	Prec	Rec	F1
<i>Detect Instruction (Half-Depth)</i>								
Half-Last Detect (CNN)	0.850	0.847	0.846	0.846	0.732	0.680	0.874	0.765
Half-Last Detect (AR)	0.729	0.795	0.599	0.683	0.674	0.636	0.811	0.713
<i>Disambiguate Instruction (Half-Depth)</i>								
Half-Last Disambig. (AR)	0.874	0.839	0.917	0.876	0.763	0.695	<b>0.937</b>	0.798
Half-Full Disambig. (AR)	<b>0.912</b>	<b>0.934</b>	0.881	<b>0.907</b>	<b>0.826</b>	<b>0.793</b>	0.884	<b>0.836</b>
<i>Disambiguate Instruction (Full-Depth)</i>								
Full-Last Disambig. (AR)	0.701	0.629	0.940	0.754	0.632	0.634	0.621	0.628
Full-Full Disambig. (AR)	0.893	0.893	0.887	0.890	0.647	0.608	0.832	0.702
<i>Zero-Shot Baselines</i>								
Gemma3-4B-it	0.487	0.487	<b>1.000</b>	0.655	0.500	0.500	<b>1.000</b>	0.667
Gemma3-27B-it	0.609	0.560	0.931	0.699	0.479	0.486	0.747	0.589

**Model depth** (Half vs. Full decoder layers), **Finetuning strategy** (Last layer vs. Full network), and **Instruction type**.

- **Half-Last Detect (CNN)**: This follows the model in Fig. 4. A CNN classifier trained on attention maps extracted from around the first half of the decoder layers (“Half”), and only the last decoder layer is finetuned along with the CNN. It uses the instruction starting with “detect” (seen during pretraining).
- **Half-Last Detect (AR)**: An autoregressive approach using only the first half of the decoder layers. Only the last layer is finetuned, using the standard “detect” instruction.
- **Half-Last Disambig. (AR)**: Uses the first half of the decoder layers with only the last layer finetuned. However, it uses a new instruction starting with “disambiguate” (unseen during pretraining) to test generalization.
- **Half-Full Disambig. (AR)**: Uses the first half of the decoder layers, but unlike the “Last” variants, all parameters in these layers are LoRA finetuned using the “disambiguate” instruction.
- **Full-Last Disambig. (AR)**: Uses the complete set of decoder layers (“Full” depth). Only the very last layer is finetuned (“Last”) using the “disambiguate” instruction.
- **Full-Full Disambig. (AR)**: The full decoder depth with

all layers LoRA finetuned (“Full” finetuning) with the “disambiguate” instruction.

- **Gemma-it Baselines**: We include zero-shot evaluations of Gemma3-4B-it and Gemma3-27B-it [29] to provide a reference for standard instruction-tuned performance on this task.

The goal of this comparison is to see if using the internal representation of VLM provides a useful signal for ambiguity detection and up to which level.

4) **Results**: In the finetuned setting, we observe that under the same architecture and “Detect” instruction strategy, the CNN-based model significantly outperforms the classic Auto-Regressive (AR) baseline. As shown in Table I, the Half-Last Detect (CNN) achieves an F1-score of 0.846, compared to just 0.683 for the autoregressive approach. Furthermore, our results confirm that using around half the layers is actually superior to using the full layers; for instance, the AR “Disambiguate” task drops from 0.876 F1 at half-depth (Half-Last) to 0.754 at full-depth (Full-Last). Moreover, although **Full-Full Disambig. (AR)** attains almost the same in-domain performance as **Half-Full Disambig. (AR)** (0.89 vs. 0.907), its score drops sharply to 0.702 on the real-world OOD dataset, compared to 0.836 at half layers, indicating that middle layers exhibit greater robustness to domain shift than later layers.

Furthermore, while the generative “Disambiguate” task at half-layers slightly outperforms the **Half-Last Detect (CNN)** detection (0.876 vs. 0.846), this comes with a cost. Since the “disambiguate” conditioning token was never seen during pretraining, with finetuning, the model shapes its attention to achieve the task without producing a meaningful, explainable signal. This is illustrated in Figure 8.

5) **Generalization**: We observe that overall performance slightly drops on Dataset 2 (real-world OOD); for example, the **Half-Last Detect (CNN)** drops from 0.846 to 0.765.




Fig. 7. F1-score and accuracy as a function of the number of the decoder layer used



TABLE II  
PERFORMANCE COMPARISON BETWEEN OUR MODEL AND AMBRESVLM [21].

Model	Simulation				Real World			
	Acc	Prec	Rec	F1	Acc	Prec	Rec	F1
AmbResVLM [21]	–	0.52	0.97	0.68	–	0.48	0.78	0.59
<b>Half-Full Disambig. (AR)</b>	<b>0.995</b>	<b>1.000</b>	<b>0.989</b>	<b>0.994</b>	<b>0.795</b>	<b>0.898</b>	<b>0.644</b>	<b>0.750</b>

Model	LoRA config	Acc@0.5
<b>TiO [15]</b>	None	71.2%
<b>CLUE (non-mix backbone)</b>	$\alpha = 8, r = 16$	63.19%
<b>CLUE (mix)</b>	$\alpha = 32, r = 16$	75.55%
<b>CLUE (mix)</b>	$\alpha = 16, r = 8$	<b>75.66%</b>

TABLE III  
COMPARISON BETWEEN IVG MODELS TRAINED ON INVIG ONLY.

head with the IVG decoder using an auxiliary loss, so that clarification behaviour and detection can co-adapt. All in all, we see attention-based ambiguity signals as an, interpretable bridge between black-box VLMs and HRI: they are relatively cheap to compute, naturally grounded in the scene.

#### ACKNOWLEDGMENTS

This research was partially funded by the French National Research Agency (ANR) under the OSTENSIVE (ANR-24-CE33-6907-01), ANITA (ANR-22-CE38-0012-01), and euROBIN, the European Robotics and AI Network (Grant agreement 101070596), funded by the European Commission.

#### REFERENCES

- [1] L. Yu *et al.*, “Mattnet: Modular attention network for referring expression comprehension,” 2018. [Online]. Available: <https://arxiv.org/abs/1801.08186>
- [2] A. Kamath *et al.*, “Mdetr – modulated detection for end-to-end multi-modal understanding,” 2021. [Online]. Available: <https://arxiv.org/abs/2104.12763>
- [3] L. Jin *et al.*, “Refclip: A universal teacher for weakly supervised referring expression comprehension,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 2681–2690.
- [4] Z. Han *et al.*, “Zero-shot referring expression comprehension via structural similarity between images and captions,” 2024. [Online]. Available: <https://arxiv.org/abs/2311.17048>
- [5] M. Shridhar and D. Hsu, “Interactive Visual Grounding of Referring Expressions for Human-Robot Interaction,” Jun. 2018, arXiv:1806.03831 [cs]. [Online]. Available: <http://arxiv.org/abs/1806.03831>
- [6] J. Hatori *et al.*, “Interactively picking real-world objects with unconstrained spoken language instructions,” *CoRR*, vol. abs/1710.06280, 2017. [Online]. Available: <http://arxiv.org/abs/1710.06280>
- [7] J. Mi *et al.*, “Interactive Natural Language Grounding via Referring Expression Comprehension and Scene Graph Parsing,” *Frontiers in Neurorobotics*, vol. 14, Jun. 2020, publisher: Frontiers. [Online]. Available: <https://www.frontiersin.org/journals/neurorobotics/articles/10.3389/fnbot.2020.00043/full>
- [8] J. Thomason *et al.*, “Improving Grounded Natural Language Understanding through Human-Robot Dialog,” in *2019 International Conference on Robotics and Automation (ICRA)*, May 2019, pp. 6934–6941, arXiv:1903.00122 [cs]. [Online]. Available: <http://arxiv.org/abs/1903.00122>

- [9] Y. Yang, X. Lou, and C. Choi, “Interactive Robotic Grasping with Attribute-Guided Disambiguation,” Mar. 2022, arXiv:2203.08037 [cs]. [Online]. Available: <http://arxiv.org/abs/2203.08037>
- [10] J. Park *et al.*, “Clara: Classifying and disambiguating user commands for reliable interactive robotic agents,” 2024. [Online]. Available: <https://arxiv.org/abs/2306.10376>
- [11] P. Jian *et al.*, “Teaching vision-language models to ask: Resolving ambiguity in visual questions,” in *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, W. Che *et al.*, Eds. Vienna, Austria: Association for Computational Linguistics, Jul. 2025, pp. 3619–3638. [Online]. Available: <https://aclanthology.org/2025.acl-long.182/>
- [12] A. Testoni and R. Fernández, “Asking the right question at the right time: Human and model uncertainty guidance to ask clarification questions,” 2024. [Online]. Available: <https://arxiv.org/abs/2402.06509>
- [13] NVIDIA, “Isaac Sim.” [Online]. Available: <https://github.com/isaac-sim/IsaacSim>
- [14] H. Zhang *et al.*, “Invig: Benchmarking interactive visual grounding with 500k human-robot interactions,” 2023. [Online]. Available: <https://arxiv.org/abs/2310.12147>
- [15] J. Xu *et al.*, “Towards unified interactive visual grounding in the wild,” 2024. [Online]. Available: <https://arxiv.org/abs/2401.16699>
- [16] H. Ahn *et al.*, “Interactive Text2Pickup Network for Natural Language based Human-Robot Collaboration,” May 2018, arXiv:1805.10799 [cs]. [Online]. Available: <http://arxiv.org/abs/1805.10799>
- [17] F. I. Doğan, I. Torre, and I. Leite, “Asking Follow-Up Clarifications to Resolve Ambiguities in Human-Robot Conversation,” in *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI ’22. Sapporo, Hokkaido, Japan: IEEE Press, Mar. 2022, pp. 461–469.
- [18] N. Carion *et al.*, “End-to-end object detection with transformers,” 2020. [Online]. Available: <https://arxiv.org/abs/2005.12872>
- [19] J. Johnson, A. Karpathy, and L. Fei-Fei, “Densecap: Fully convolutional localization networks for dense captioning,” 2015. [Online]. Available: <https://arxiv.org/abs/1511.07571>
- [20] R. R. Selvaraju *et al.*, “Grad-cam: Visual explanations from deep networks via gradient-based localization,” in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 618–626.
- [21] E. Chisari *et al.*, “Robotic task ambiguity resolution via natural language interaction,” 2025. [Online]. Available: <https://arxiv.org/abs/2504.17748>
- [22] M. Deitke *et al.*, “Molmo and pixmo: Open weights and open data for state-of-the-art vision-language models,” 2024. [Online]. Available: <https://arxiv.org/abs/2409.17146>
- [23] C. Chen *et al.*, “Acknowledging focus ambiguity in visual questions,” 2025. [Online]. Available: <https://arxiv.org/abs/2501.02201>
- [24] E. J. Hu *et al.*, “Lora: Low-rank adaptation of large language models,” 2021. [Online]. Available: <https://arxiv.org/abs/2106.09685>
- [25] A. Steiner *et al.*, “Paligemma 2: A family of versatile vlm for transfer,” 2024. [Online]. Available: <https://arxiv.org/abs/2412.03555>
- [26] A. Vaswani *et al.*, “Attention is all you need,” 2023. [Online]. Available: <https://arxiv.org/abs/1706.03762>
- [27] J. Ainslie *et al.*, “Gqa: Training generalized multi-query transformer models from multi-head checkpoints,” 2023. [Online]. Available: <https://arxiv.org/abs/2305.13245>
- [28] L. Beyer *et al.*, “Paligemma: A versatile 3b vlm for transfer,” 2024. [Online]. Available: <https://arxiv.org/abs/2407.07726>
- [29] G. Team *et al.*, “Gemma 3 technical report,” 2025. [Online]. Available: <https://arxiv.org/abs/2503.19786>