Multi-Task Domain Adaptation for Language Grounding with 3D Objects

Penglei Sun¹*[®], Yaoxian Song²*[®], Xinglin Pan¹[®], Peijie Dong¹[®],

Xiaofei Yang^{1,3}[●], Qiang Wang^{⊠4}[●], Zhixu Li⁵[●],

Tiefeng Li^{2} , and Xiaowen $Chu^{1\boxtimes}$

¹ The Hong Kong University of Science and Technology (Guangzhou), China {psun012,xpan413,pdong212}@connect.hkust-gz.edu.cn, xwchu@ust.hk

² Zhejiang University, China {songyaoxian,litiefeng}@zju.edu.cn

³ Guangzhou University, China xiaofeiyang@gzhu.edu.cn

⁴ Harbin Institute of Technology (Shenzhen), China qiang.wang@hit.edu.cn ⁵ Fudan University, China zhixuli@fudan.edu.cn

Abstract. The existing works on object-level language grounding with 3D objects mostly focus on improving performance by utilizing the offthe-shelf pre-trained models to capture features, such as viewpoint selection or geometric priors. However, they have failed to consider exploring the cross-modal representation of language-vision alignment in the cross-domain field. To answer this problem, we propose a novel method called Domain Adaptation for Language Grounding (DA4LG) with 3D objects. Specifically, the proposed DA4LG consists of a visual adapter module with multi-task learning to realize vision-language alignment by comprehensive multimodal feature representation. Experimental results demonstrate that DA4LG competitively performs across visual and nonvisual language descriptions, independent of the completeness of observation. DA4LG achieves **state-of-the-art** performance in the single-view setting and multi-view setting with the accuracy of 83.8% and 86.8% respectively in the language grounding benchmark SNARE. The simulation experiments show the well-practical and generalized performance of DA4LG compared to the existing methods. Our project is available at https://sites.google.com/view/da4lg.

Keywords: Visual Language Grounding \cdot Multimodal Learning \cdot Domain Adaptation

1 Introduction

Visual language grounding aims to identify the region or object within visual content described by natural language [7,21]. It serves as an essential bridge for current embodied agents to connect symbolic concepts with the perceptible real world, enabling the evolution of an agent's intelligence from perceptual to cognitive decision-making [5, 16]. For example, an agent could make a cup of coffee

^{*} Equal Contribution.



Fig. 1: The comparison between existing works and our model. Existing works focus on (a) multi-view perception and (b) external prior. (c) We approach language grounding from domain adaptation.

following a series of primitive instructions including detailed descriptions of the target object from a Large Language Model (LLM) planner like GPT-4 [30]. Within the process, visual language grounding plays a key role in linking each step instruction with the physically observed object [3,36]. Therefore, visual language grounding with 3D objects is the indispensable means of enabling an agent to interact with the real world. A limited scale of high-quality vision-language paired data hinders the development of visual language grounding technology, especially 3D vision language grounding. To address this problem, existing work attempts [9,28,38,42] to use multi-view perception or external priors, which need extra data cost and existing domain gap caused by pre-trained feature encoders in fixed settings. In this paper, we give an exploration from a domain adaptation perspective for the language grounding task inspired by domain adaptation in parameter-efficient tuning of large language model [13, 14, 19, 22, 27, 39, 41].

As shown in Fig. 1, existing research in language grounding mainly focuses on two lines including the multi-view perception-enhanced method (Fig. 1 (a)) and the external prior-injected method (Fig. 1 (b)). For the former, Thomason et al. [42] and Mitra et al. [28] propose a view-based method to improve the prediction accuracy. Thomason et al. [42] design an auxiliary task of viewpoint angle estimation to enhance 3D object understanding. Mitra et al. [28] design a multi-view transformer to fuse visual features and textual features in shared space. For the latter, Corona et al. [9] propose a voxel-informed method by using pre-trained 3D volumetric generative model LegoFormer based on the view of 3D objects [45]. Song et al. [38] construct explicit scene-driven multimodal knowledge graph ManipMob-MMKG to design a knowledge-enhanced method. In all, current methods still depend heavily on viewpoints or external priors.

In the representation learning aspect, existing works usually adopt visionlanguage feature encoders using freezing parameter mode pre-trained on a source domain, which cannot work well for 3D language grounding tasks in multimodal alignment due to domain gap. Based on these findings, we propose a novel multimodal domain adaptation method **D**omain **A**daptation for **L**anguage Grounding named DA4LG, to improve 3D object-level understanding and multimodal alignment, which does not require extra visual or textual data, as shown in Fig. 1(c). Given the language similarity between the source domain (e.g., the WebImageText domain, where CLIP is pre-trained [31]) and the target domain (e.g., language grounding domain), coupled with the generalization capabilities of pre-trained language models [12, 20, 44], DA4LG concentrates on domain adaptation within the vision feature. Specifically, we design the pseudo-siamese visual encoder network [17] to realize domain adaptation, in which one visual encoder subnetwork is used to learn domain-specific 3D visual representation, named **Domain-specific Encoder**, while another one is freezing to encode visual representations associated with the source domain. For training the model, we design two auxiliary tasks with the main language grounding task (LGR) to learn cross-modal representation. The first task is to distinguish different objects using vision and language contrastive learning, while the second involves regenerating the input text from multi-modal fused features. We evaluate DA4LG via the language grounding dataset SNARE proposed by Thomason et al. [42], which discriminates natural language descriptions with 3D ShapeNet [6] objects. DA4LG achieves state-of-the-art (SOTA) performance in both single-view and multi-view settings. Additionally, through simulation experiments, DA4LG demonstrates generalization and robustness compared to existing models. Compared to multi-view perception-enhanced methods such as Thomason et al. [42] and Mitral et al. [28], our DA4LG is immune to the number of viewpoints or the selection of viewpoints. Compared to external prior-injected methods such as Coronal et al. [9] and Song et al. [38], our proposed method requires only adapting the visual encoder as a cloned module with limited parameter training, without the need for external prior injection. This approach presents a clear advantage in terms of reducing the model's parameter size while simultaneously enhancing its reliability.

Our main contributions can be summarized as follows:

- 1. We introduce a novel domain adaptation method (DA4LG) using multi-task learning to reduce visual domain gap in vision-language aligned representation for language grounding with 3D objects.
- DA4LG demonstrates SOTA performance, achieving 83.8% accuracy in a single-view setting and 86.8% accuracy in a multi-view setting on the language grounding benchmark SNARE [42].
- 3. We conduct the simulation 3D object grounding experiment including a grounding environment and a test set extended from Lang-SHAPE [37] referred to as **Simulation-SNARE**. The results indicate that DA4LG has an obvious advantage in the robustness and generalization of applications compared to existing models.

2 Related Work

Language Grounding. Language grounding with 3D can be classified into two categories: language grounding for navigation and language grounding for

4 P. Sun et al.



Fig. 2: The framework of DA4LG. DA4LG is comprised of Encoder Layer, Embedding Reweighting Layer and Embedding Fusion Layer. Encoder Layer contains three encoders: Language Encoder (L. Encoder), Vision Encoder, and Domain-specific Encoder. The snowflake and fire denote the freezing and unfreezing respectively.

interaction. Language grounding for navigation focuses on training an agent to follow a set of natural language instructions to navigate towards a target object in an environment, as discussed by Chen et al [8]. This task is relevant to scene understanding including object localization [1, 23, 33, 46], visual language navigation [18, 29, 34, 35].

Recently, research has emerged in the language grounding for interaction with 3D objects. ShapeGlot [2] investigates how linguistic expressions capture detailed differences in the shapes of common objects based on their 3D representations. Akula et al. [4] propose the use of neural module networks that share weights and exploit associations between similar textual contexts, such as "*dark cube on the left*" and "*black cube on the left*". SNARE [42] presents a challenge by requiring the identification of referent objects that are highly similar to distractors. This task encompasses various perspectives of objects in 3D space from ShapeNet [6], thereby increasing the complexity of the problem in language referencing.

Current research encompasses two main approaches: the multi-view perceptionenhanced method and the external prior-injected method. The former employs a multi-view framework with 3D objects to enhance prediction accuracy such as MAGiC [28], LAGOR [42], among others. The latter approach integrates exMulti-Task Domain Adaptation for Language Grounding with 3D Objects

ternal priors to augment task performance, including VLG [9], LOCKET [38], among others.

Domain Adaptation. Recent research has placed particular emphasis on adaptors designed for domain adaptation in the natural language processing (NLP). Hu et al. [22] propose Low-Rank Adaptation(LoRA), a novel approach involving freezing pre-trained model weights and introducing trainable rank decomposition matrices as adapter into each layer of the Transformer architecture. Effective adaptation strategies based on this are introduced to bridge the domain gap in NLP tasks [13, 14, 19, 22, 27, 39, 41]. Inspired by the existing research, we incorporate domain adapter into multimodal tasks and explore its utility in adapting language grounding with 3D objects, extending the application of it beyond traditional NLP tasks.

3 Proposed Method

Domain Adaptation. Given a pre-trained network model $\mathcal{F}_s(\cdot)$ on a source task \mathcal{T}_s (e.g., CLIP pre-trained model [31]) in the source domain \mathcal{D}_s (e.g., the WebImageText domain where CLIP is pre-trained on [31]), and a set of training examples with associated labels in the target domain \mathcal{D}_t for target task \mathcal{T}_t (i.e. the language grounding task). Our goal is to create an adaption strategy to promote the performance of target predictive function $\mathcal{F}_t(\cdot)$ in \mathcal{D}_t for \mathcal{T}_t through leveraging the $\mathcal{F}_s(\cdot)$ in \mathcal{D}_s and \mathcal{T}_s [10].

Task Definition. In the language grounding task with a 3D object, given a language description l as input, our objective is to identify the best matching object \hat{o} in a set of candidate objects $\mathcal{O} = \{o_i \mid i \in \{1,2\}\}$ [42]:

$$\hat{o} = \underset{o_i \in \mathcal{O}}{\arg\max} \ p(o_i|l), \tag{1}$$

where $p(o_i|l)$ denotes the conditional probability of the target object given the language description. We expect our output \hat{o} to closely align with the ground truth o^* as much as possible.

3.1 Domain Adaptation for Language Grounding

Overall of Network Structure. As shown in Fig. 2 (a), DA4LG comprises three encoders: a Vision Encoder $(E_{vision}(\cdot))$, a Language Encoder $(E_{lang}(\cdot))$, and a Domain-specific Encoder $(E_{domain}(\cdot))$. Additionally, it includes an Embedding Reweighting Layer $(RW(\cdot))$ and an Embedding Fusion Layer $(MLP(\cdot))$. Therefore Eq. 1 can be rewritten as:

$$f_{j}^{v} = E_{vision}(v_{j}),$$

$$f_{j}^{d} = E_{domain}(v_{j}),$$

$$f^{l} = E_{lang}(l),$$

$$Score = MLP(RW(f_{j}^{d}), f_{j}^{v}, f^{l}),$$

$$\hat{o} = \underset{o_{i} \in \mathcal{O}}{\operatorname{arg\,max}} Score_{o_{i}},$$

$$(2)$$

where the object o is projected into a set of images $\mathcal{V} = \{v_1, ..., v_j\}$ from discrete viewpoints. E_{vision} and E_{lang} are based on the pre-trained encoder in the freezing state. We compute the $Score_{o_i}$ to determine the output label \hat{o} .

Domain-specific Encoder. To reduce the domain gap between \mathcal{D}_s and \mathcal{D}_t , we design a Domain Vision Transformer named **Domain-specific Encoder** to encode the \mathcal{V} , as shown in Fig. 2 (c). Domain-specific Encoder is pre-trained on \mathcal{D}_s (e.g., the WebImageText domain [31]). Compared to vanilla Vision Transformer [15], we incorporate low-rank matrices W_*^A, W_*^B as domain adapters in the W_q, W_k, W_v of multi-head attention layer MHA. Different from the adapters employed for parameter-efficient tuning in NLP, the domain adapters in DA4LG are designed to capture the domain-specific representation. In Domain-specific Encoder, all other parameters are freezing except the W_*^A, W_*^B while training. v_j is input into the MHA to build domain features f_i^d ,

$$Q = W_q \cdot v_j + W_q^A \cdot W_q^B \cdot v_j,$$

$$K = W_k \cdot v_j + W_k^A \cdot W_k^B \cdot v_j,$$

$$V = W_v \cdot v_j + W_v^A \cdot W_v^B \cdot v_j,$$

$$f_j^d = MHA(Q, K, V).$$
(3)

Embedding Reweighting Layer. As shown in Fig. 2 (b), we use $RW(\cdot)$ to adjust the f_j^d and reduce the impact of those features that are irrelated to f^l . Specifically, given the f_j^d corresponding to the *j*-th view v_j and the description f^l , we compute the cosine similarity with $sim(f_j^d, f^l)$ and get the weighted combination of domain features $f_j^d = f_j^d \cdot softmax(sim(f_j^d, f^l))$. **Embedding Fusion Layer.** To enhance multimodal alignment and construct

Embedding Fusion Layer. To enhance multimodal alignment and construct joint features, we employ the aggregate operation agg in Embedding Fusion Layer to build the vision features $f^v = agg(f_j^v)$ and domain features $f^d = agg(f_j^d)$, where agg is max-pooling. We concatenate the features $f = [f^l, f^v, f^d]$ which is fed into a multi-layer perceptron $MLP(\cdot)$ to compute *Score*.

3.2 Multi-task Learning

As shown in Fig. 2 (a), the DA4LG framework incorporates three different tasks: the Language Grounding (LGR) Task, the Vision-Language Contrastive (VLC) Task, and the Vision Grounding Caption (VGC) Task in multi-task learning. The LGR Task is designed as the primary task following the existing works [42]. The VLC and VGC tasks serve as auxiliary tasks to optimize the training objectives inspired by BLIP-2 [24].

Language Grounding Task (LGR Task). The primary task is the LGR task, which involves predicting the target. We feed *Score* to predict the ground truth label o^* and apply the binary cross-entropy loss for optimization:

$$\mathcal{L}_{LGR} = -\mathbb{E}[o^* log(Score) + (1 - o^*) log(1 - Score)].$$
(4)

Vision-Language Contrastive Task (VLC Task). We propose the VLC task to learn an embedding that distinguishes samples from two different distributions. At each step, we sample some positive or negative pairs (l, o) during

Table 1: Performance of our model and existing work on the SNARE dataset. Absence of parentheses signifies that the original paper does not report standard deviations. Standard deviations over five seeds are shown in parentheses (the same as below).

Model	Aodel Method		Validation	L	Test		
		Visual	Blind	All	Visual	Blind	All
Vilbert [26]	External Prior	89.5	76.6	83.1	80.2	73	76.6
CLIP [31]	External Prior	83.7	65.2	74.5	80.0	61.4	70.9
MATCH [42]	External Prior	89.2 (0.9)	75.2(0.7)	82.2(0.4)	83.9(0.5)	68.7(0.9)	76.5(0.5)
LAGOR [42]	Multi-view Perception	89.8 (0.4)	75.3(0.7)	82.6(0.4)	84.3(0.4)	69.4(0.5)	77.0 (0.5)
VLG [9]	External Prior	91.2 (0.4)	78.4(0.7)	84.9(0.3)	86.0	71.7	79.0
BLIP2 [24]	External Prior	51.2	50.9	51.5	-	-	-
LOCKET [38]	External Prior	90.9	78.4	84.7	86.1	71.5	79.0
MAGiC [28]	Multi-view Perception	92.1(0.4)	$81.3\ (0.9)$	86.8 (0.5)	87.7	75.4	81.7
$\overline{\mathbf{DA4LG}(\mathrm{ours})}$	Domain Adaptation	91.8 (0.3)	81.8 (0.6)	86.8 (0.5)	88.5	75.0	81.9

training. Specifically, samples from matched pairs are termed positives, whereas those from unmatched pairs are termed negatives. We use cosine similarity, denoted as $s(f^l, f^o)$, to measure the alignment between the language features f^l and the object features f^o , where $f^o = agg([f_i^v, f_i^d])$. We optimize this function to correctly select a single positive description sample l with Φ negative object samples and calculate the contrastive loss for object description:

$$\mathcal{L}_{con}^{o \to l} = -\mathbb{E}[\log \frac{s(f^l, f^{o_p})}{s(f^l, f^{o_p}) + \sum_{\varphi=1}^{\Phi} s(f^l, f^{o_p^{\varphi}})}],\tag{5}$$

where f^{o_p} , $f^{o_n^{\varphi}}$ are the features of the positive and negatives sample for 3D objects. Similarly, we can obtain the contrastive loss for the description-object pairs:

$$\mathcal{L}_{con}^{l \to o} = -\mathbb{E}[\log \frac{s_j(f^{l_p}, f^o)}{s(f^{l_p}, f^o) + \sum_{\varphi=1}^{\Phi} s(f^{l_n^{\varphi}}, f^o)}].$$
(6)

And the VLC loss is denoted as $\mathcal{L}_{VLC} = \mathcal{L}_{con}^{l \to o} + \mathcal{L}_{con}^{o \to l}$. Vision Grounding Caption Task (VGC Task). Given the f^o and l, we design the VGC task to generate textual descriptions based on the freezing GPT-2 [32]. Here, l represents a sequence of tokens denoted as $\{c_i, i = [1, ..., N]\}$. Our training objective is to predict the caption tokens conditioned on the output token in auto-regressive. The training loss for the VGC task is formulated as:

$$\mathcal{L}_{VGC} = -\mathbb{E}[\sum_{i=1}^{N} \log p(c_i | f^o, c_1, ..., c_{i-1})].$$
(7)

During the multi-task learning phrase, our target loss function is

$$\mathcal{L} = \mathcal{L}_{LGR} + \mathcal{L}_{VLC} + \mathcal{L}_{VGC}.$$
(8)

4 Experiment Design

To evaluate our proposed method, we investigate four key research questions (RQs). The RQ1 concerns the DA4LG superior to the baselines. The RQ2

and **RQ3** focus on how to adopt DA4LG efficiently. The **RQ4** concerns the generalization of DA4LG in the downstream simulation environment.

• **RQ1**: What are the advantages of our DA4LG compared to other methods for the language grounding task?

• **RQ2**: Which training policy in Domain-specific Encoder can obtain the better performance with efficient parameters?

• **RQ3**: How do different learning tasks in DA4LG affect the language grounding performance?

• **RQ4**: Can our DA4LG perform more effectively deployed in a downstream task by a simulation environment compared to other methods?

4.1 Baseline Models

We conduct a comparative analysis between DA4LG and various public baselines, as summarized in Table 1. Two primary approaches are utilized in current research: the multi-view perception-enhanced method, exemplified by MAGiC [28] and LAGOR [42], and the external prior-injected method, represented by ViL-BERT [26], MATCH [42], VLG [9], CLIP [31], LOCKET [38], and BLIP2 [24]. We list these baselines below:

• **LAGOR** adopts a multi-task learning approach by predicting the canonical viewing angle for individual view images.

• **MAGiC** performs joint reasoning over candidate referent objects, considering each object from multiple possible perspectives.

• **MATCH** and **ViLBERT** uses CLIP-ViT and ViLBERT respectively to encode the views of each object. An MLP is trained to assign scores based on the encoded views and language description embedding

• **CLIP** uses the cosine distance in CLIP embedding between visual and language features to pick the object in the lowest distance.

• **BLIP-2** is the multimodal LLM-based method in zero-shot setting.

• VLG leverages implicit 3D prior information from predicted volumetric voxel maps to improve language grounding performance by LegoFormer [45].

• **LOCKET** is a knowledge enhancement method that employs a graph convolutional network to encode a multimodal knowledge graph.

4.2 Implementation Details

Training and Inference Details. We employ the Adam optimizer with weight decay $5e^{-4}$. The batchsize is 64, training epoch is 60, and learning rate is $5e^{-3}$. Experiments are implemented with CUDA 11.2 and PyTorch 1.7.1 and run on one NVIDIA RTX4090. In DA4LG, we employ the vision and language encoders from CLIP ViT-B/32 [31] as the Vision Encoder and Language Encoder, respectively. Domain-specific Encoder is initialized from the vision encoder in CLIP ViT-B/32.

Benchmark Datasets. We train and evaluate our proposed method on SNARE dataset [42] which is split into training, validation, and test sets following existing works [42]. SNARE is a benchmark for choosing the correct object with small differences in multiple views given a language description. Each data in the set has the label of visual or blind. The visual label means a comprehensive understanding of the object, providing relevant visual cues to guide the grounding process (e.g., "classic armchair with white seat"). The blind labels predominantly focus on the object's shape and specific distinguishing attributes, intentionally omitting color and other visual characteristics (e.g., "oval back and vertical legs"). The training set consists of 207 categories, 6, 153 objects, and 39, 104 references. The validation set contains 7 categories, 371 objects, and 2, 304 references. The test set consists of 48 categories, 1, 357 objects, and 8, 751 references.

Simulation Details. We build a new simulation benchmark by sampling the objects from the existing 3D datasets Lang-SHAPE [37] and annotating them with the original data annotation process utilized in SNARE, which we refer to as Simulation-SNARE. Specifically, the objects and domains in Lang-SHAPE are identical to those in SNARE, both of which are derived from ShapeNet [6]. We follow the original data annotation process utilized in SNARE. Simulation-SNARE consists of 327 objects and 2,876 references, where 634 classified as visual and 2,242 as blind, following the configuration of SNARE. We deploy the Simulation-SNARE into the simulation environment MuJoCo [43]. We replicate the existing methods CLIP, ViLBERT, MATCH, VLG, and LAGOR with open-source code, deploying these methods along with DA4LG in the simulation world.

Metric. The metric for this task is the accuracy (%) of the predictions in correctly identifying the object referred to by the language description from two candidate objects. We calculate the accuracy across all sets, the visual subset and the blind subset respectively. Additionally, we assess the model's parameter efficiency by calculating its parameter size.

5 Result Analysis

5.1 Benchmark Comparisons

Comparisons with Existing methods. To answer RQ1, we conduct a comparative analysis in multi-view and single-view settings and our proposed model DA4LG outperforms in all settings compared to baselines. In the multi-view setting, Table 1 demonstrates that DA4LG achieves the best improvements in validation performance compared to the External Prior methods utilizing the same backbone CLIP ViT-B/32. Specifically, DA4LG exhibits enhancements of 12.3% (74.5% \rightarrow 86.8%), 4.6% (82.2% \rightarrow 86.8%), 2.1% (84.7% \rightarrow 86.8%), and 1.9% (84.9% \rightarrow 86.8%) compared to CLIP, MATCH, LOCKET, and VLG methods, respectively. Specifically, the VLG model requires 168 million training parameters, while our method requires only less than half of the VLG

Table 2: Performance on the model with single-view. *We obtain the reported results from the curve [28], which is below but close to 82.0%. **The results come from our reproduction.

Models	Method	Validation				
		Visual	Blind	All		
CLIP [31]	External Prior	79.0	63.0	71.1		
MATCH [42]	External Prior	88.4 (0.4)	73.3(0.6)	80.9(0.4)		
VLG [9]**	External Prior	89.3 (0.7)	74.1(0.9)	81.8(0.6)		
MAGiC [28]	Multi-view Perception	-	-	82.0^{\star}		
DA4LG (ours)	Domain Adaptation	$90.1 \ (0.5)$	$77.1 \ (0.8)$	83.8 (0.5)		

model's parameters (79.5 million). This demonstrates the superior parameter efficiency of DA4LG. Contrary to the LOCKET model, which demands a knowledge graph for extensive data, our method eliminates the necessity for such a structure. When employing ViLBERT as the backbone, our model achieves a $3.7\%(83.1\% \rightarrow 86.8\%)$ improvement in validation accuracy. Furthermore, when utilizing a multimodal LLM in a zero-shot approach, our model exhibits an increase in performance exceeding 30%. The **Multi-view Perception** method includes LAGOR and MAGiC. DA4LG demonstrates improvements, with a $4.2\%(82.6\% \rightarrow 86.8\%)$ improvement in validation compared to LAGOR. DA4LG achieves a validation score of 86.8% on the validation set and 81.9% on the test set, tying with the existing SOTA model MAGiC in the validation and surpassing it by 0.2% in the test.

The performance of models in the **single-view setting** are shown in Table 2. DA4LG achieves the best result among all methods, surpassing the existing SOTA model MAGiC by over $1.8\%(82.0\% \rightarrow 83.8\%)$ in validation. DA4LG in the single-view setting achieves better performance than several models in the multi-view setting. Specifically, DA4LG exhibits improvements of $9.3\%(74.5\% \rightarrow 83.8\%)$, $1.6\%(82.2\% \rightarrow 83.8\%)$, $1.2\%(82.6\% \rightarrow 83.8\%)$, and $0.7\%(83.1\% \rightarrow 83.8\%)$ over the CLIP, MATCH, LAGOR, and ViLBERT models in the multi-view setting can enhance the overall performance of the model in the language grounding task. We believe that the observed improvement in performance can be attributed to the multimodal alignment within the target domain. In the following subsection, we provide ablation studies and visualizations to further demonstrate this result.

Training Policy in Domain-specific Encoder. To answer **RQ2**, we explore the different training mode and the source domains in Domain-specific Encoder. For training modes, we conduct a comparative analysis of following four scenarios where the Domain-specific Encoder is initialized using the same pre-trained parameters:

- Freezing-Param: The parameters of Domain-specific Encoder is fixed.
- Full-Param: The full parameter in the Domain-specific Encoder is trainable.

Table 3: Performance and training parameters across training mode in Domain-specific Encoder. M means one million and #Param means the training parameter size (the same as below).

Strategies		$ _{\#Param}$		
	Visual	Blind	All	
Freezing-Param	89.9 (0.5)	79.8(0.4)	84.9(0.4)	79.2M
Full-Param	86.4 (0.7)	75.0(0.2)	80.8(0.6)	254M
Partial-Param	89.9 (0.4)	77.7(0.6)	83.8(0.5)	107M
Domain-Adapter	91.8(0.3)	81.8 (0.6)	86.8 (0.5)	79.5M

Table 4: Comparative performance of Domain-Specific Encoders initialized with models from various source domains.

Models	Source Domain	#Param	Visual	Blind	All
Scratch-B	-	87M	86.4 (0.7)	76.7 (0.2)	81.6 (0.6)
BLIP-B	BLIP [25]	86M	+2.8(0.4)	+4.8(0.5)	+3.0 (0.6)
ViT-B	ImageNet21K [11]	87M	+3.4(0.5)	+4.8(0.7)	+4.0(0.4)
CLIP-B	WebImageText [31]	87M	+5.4(0.3)	+5.1 (0.6)	$+5.2\ (0.5)$
Scratch-L	-	303M	86.9 (0.5)	74.2(0.8)	80.6 (0.5)
BLIP-L	BLIP [25]	303M	+3.0(0.8)	+3.9(0.6)	+3.3 (0.5)
ViT-L	ImageNet21K [11]	304M	+4.3(0.6)	+3.8(0.8)	+4.1(0.5)
CLIP-L	WebImageText [31]	303M	+3.8(0.6)	+5.1 (0.7)	+4.4(0.4)

• **Partial-Param**: Inspired by Sun et al. [40], only the last two layers of the Domain-specific Encoder are updated during the training stage.

• **Domain-Adapter**: The Domain-specific Encoder is training with domain adaptor mentioned in Section 3.1 following Low-Rank Adaptation (LoRA) [22].

As presented in Table 3, Domain-Adapter achieves the best performance compared to other training modes, with an accuracy of 86.8% on the validation set with the limited training parameters. The Full-Param and Partial-Param strategies achieved performance of 80.8% and 83.8% in the validation set respectively. However, they come with the drawback of increased training parameters compared to the Freezing-Param and Domain-Adapter methods.

We select four source domains to initialize Domain-specific Encoder for **RQ2**. Table 4 demonstrates that a Domain-specific Encoder pre-trained on source domains exhibits better performance compared to training from scratch. Additionally, employing a larger-parameter Vision Transformer as the domain-specific encoder does not significantly improve the performance of the DA4LG. Domain-specific Encoder using CLIP-B (pre-trained on a specific domain WebImage-Text) demonstrates a significant improvement from 81.6% to 86.8% compared to Scratch-B. The performance of ViT-B and BLIP-B, which pre-trained on ImageNet domain and BLIP domain, demonstrates an enhancement by 3.0% and 4.0%, respectively. The performance for CLIP-L and Scratch-L are 85.0% and 80.6%, respectively. These figures do not show a significant improvement compared to the metrics for CLIP-B and Scratch-B, which are 86.8% and 81.6%, respectively. Therefore, initializing with CLIP-B is the preferred choice for domain-specific applications.

Input Encoders	Tas	Validation						
Language Vision Domain	LGR VL	C VGC	Vis	sual	Bl	ind	А	.11
\checkmark \checkmark \checkmark	√ ×	×	89.8	(1.1)	74.0	(0.4)	81.9	(0.8)
\checkmark \checkmark \checkmark	V V	×	91.0	(0.5)	79.8	(0.4)	85.4	(0.1)
\checkmark \checkmark \checkmark	√ ×	\checkmark	90.9	(0.4)	78.9	(0.7)	85.0	(0.5)
✓ ✓ ×	√ √	\checkmark	89.3	(0.8)	76.3	(0.6)	82.9	(0.7)
\checkmark \times \checkmark	\checkmark \checkmark	\checkmark	89.0	(0.6)	73.9	(0.7)	81.5	(0.6)
\checkmark \checkmark \checkmark	 ✓ ✓ 	\checkmark	91.8	(0.3)	81.8	(0.6)	86.8	(0.5)
the chair with black le	rgs			a cuj	v with c	offee in i		
	d a black bag]	ctangul	ar handl	,	J
						3		
the object with a round b	ack			the ch	air with	a yellow	headres	t

Table 5: The ablation study of multi-task learning and encoders. The terms Language, Vision, and Domain refer to the Language Encoder, Vision Encoder, and Domain-specific Encoder, respectively.

Fig. 3: Visualization of examples: Original images of the objects are displayed on the left. In the middle, attention score maps are visualized, and on the right, attention score maps are enhanced using a domain adapter in a Domain-specific Encoder. Warmer colors, such as red, indicate higher attention scores, while cooler colors, such as blue, represent lower attention scores.

Ablation Study. For RQ3, we perform an ablation study analysis to investigate the influence of different tasks and encoders in DA4LG. For different tasks, Table 5 demonstrates the effectiveness of different tasks and encoders within DA4LG. The DA4LG with all tasks and encoders achieves the best performance score of 86.6%. When DA4LG employs the LGR task exclusively, it yields a performance metric of 81.9% in validation. The incorporation of VLC and VGC enhances DA4LG performance, achieving scores of 85.4% and 85.0%, respectively. For different encoders, DA4LG without Vision Encoder achieves 81.5% in validation and DA4LG without Domain-specific Encoder achieves 82.9%. The DA4LG equipped with a Vision Encoder and Domain-specific Encoder demonstrates an enhanced performance of 86.8%, compared to its performance with only a single encoder.

Case Study and Visualization As depicted in Fig. 3, we randomly select cases for visualization to show the language vision alignment in the target domain. The

Models Method		Multi-view			Single-view			
		Visual	Blind	All	Visual	Blind	All	
CLIP [31]	External Prior	56.8	55.0	55.4	55.0	54.8	54.9	
MATCH [42]*	External Prior	61.5(0.8)	60.2(0.7)	60.5(0.6)	59.9 (0.5)	59.6(0.6)	59.7(0.5)	
LAGOR [42]*	Multi-view Perception	62.5(0.5)	61.0(0.6)	61.3(0.5)	60.7(0.5)	60.1(0.8)	60.2(0.6)	
VLG [9]*	External Prior	62.1(0.7)	61.2(0.4)	61.4(0.4)	61.4(0.6)	60.3(0.8)	60.5(0.7)	
DA4LG (ours)	Domain Adaptation	64.0(0.5)	63.8 (0.6)	63.9 (0.8)	63.1(0.7)	62.4(0.8)	62.6(0.5)	

Table 6: Results of simulation experiments comparing existing methods with DA4LG.

 *Results are obtained through our reproduction.

left column denotes the raw image of the target objects. The middle and right images illustrate the attention score map generated by the Domain-specific Encoder without domain adapter and with domain adapter respectively. For the descriptions "a cup with coffee in it" and "the chair with a yellow headrest", the attention map without the domain adapter fails to capture the descriptive elements such as "coffee" and "headrest" completely. The integration of the domain adapter results in an enhancement in the attention directed towards "coffee" and "headrest".

Additionally, the attention map extends its focus beyond the intended regions in the absence of the domain adapter. For the descriptions like "the chair with black legs", "the chair with wooden legs and a black back", "rectangular handle", and "the object with a round back", the attention is dispersed across the primary structure of the objects. However, the incorporation of the domain adapter refines the attention map, directing focus toward the detailed parts.

5.2 Simulation Results

To answer **RQ4**, we conduct a comparative analysis of the performance between existing methods and DA4LG. All models are trained on the SNARE dataset and subsequently deployed in a zero-shot setting to Simulation-SNARE. For our experiments, we provide multiple object observations including {bird, front, left, right, side} view images in a simulation environment, of which the birdview image is the input of a single-view experiment. Compared to SNARE, Simulation-SNARE has the following characteristics:

Viewpoint Diversity. Unlike the fixed viewpoint observed in the SNARE dataset, objects in the simulation environment exhibit a diverse viewpoint.
Physical Scene. The presence of physical scenes including the background and the worktop in the simulation environment is more similar to the real-world setting.

• **Texture Quality.** In the simulated environment, there is a richness in geometric details with a deficiency in textural rendering.

Table 6 shows the DA4LG architecture exhibits superior generalization and robustness compared to other existing models in the simulation environment.



Fig. 4: Examples illustrating instances where existing methods demonstrate success (\checkmark) in the SNARE dataset but failure (\times) in the Simulation-SNARE dataset, in contrast to DA4LG, which maintains robust performance across both datasets (\checkmark) . We visualize the language description (left), Simulation-SNARE (middle), and SNARE (right). For the Simulation-SNARE examples, we showcase the front, bird, and side views. For the SNARE examples, we showcase all eight views.

Specifically, DA4LG achieves the highest scores of 63.9% and 62.6% in multiview and single-view settings, respectively, outperforming the suboptimal model VLG by 2.5% and 2.1% in the respective settings.

As illustrated in Figure 4, We randomly select three cases where other methods (CLIP, MATCH, LAGOR, VLG) succeed in the SNARE dataset but fail in the Simulation-SNARE dataset. For instance, objects such as a "chair with a small metal seat" and a "circular canopy with a grip" are depicted in a fallen state, adding complexity to the task. The object "a circle and narrow legs" fails to capture the transparent top of the table in Simulation-SNARE. The DA4LG model exhibits strong performance in both the SNARE and Simulation-SNARE, indicating its enhanced suitability for simulation environments and operational scenes in the real scene.

6 Conclusion

In this work, we attempt to address language grounding tasks from the perspective of domain adaptation and introduce a novel method named **DA4LG**. A Domain-specific Encoder and a multi-task learning framework are proposed to improve language-based 3D object understanding, in which aligned cross-modal representation and domain information is encoded effectively. Evaluations on the benchmark demonstrate that DA4LG achieves SOTA performance of 83.8% and 86.8% in single-view and multi-view settings, respectively. The experiment results show the generalization and robustness of our proposed model compared to existing works. Our model reduces the domain gap for cross-modal aligned representation. We also reveal the improvement space in existing methods for domain gap research and underscore the domain adaptation in language grounding with 3D objects. Acknowledgments. This work is supported by the Postdoctoral Fellowship Program of CPSF under Grant Number GZC20232292.

References

- Achlioptas, P., Abdelreheem, A., Xia, F., Elhoseiny, M., Guibas, L.: Referit3d: Neural listeners for fine-grained 3d object identification in real-world scenes. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16. pp. 422–440. Springer (2020)
- Achlioptas, P., Fan, J., Hawkins, R., Goodman, N., Guibas, L.J.: Shapeglot: Learning language for shape differentiation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8938–8947 (2019)
- Ahn, M., Brohan, A., Brown, N., Chebotar, Y., Cortes, O., David, B., Finn, C., Fu, C., Gopalakrishnan, K., Hausman, K., et al.: Do as i can, not as i say: Grounding language in robotic affordances. arXiv preprint arXiv:2204.01691 (2022)
- Akula, A., Gella, S., Wang, K., Zhu, S.c., Reddy, S.: Mind the context: the impact of contextualization in neural module networks for grounding visual referring expressions. In: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing. pp. 6398–6416 (2021)
- Bisk, Y., Holtzman, A., Thomason, J., Andreas, J., Bengio, Y., Chai, J., Lapata, M., Lazaridou, A., May, J., Nisnevich, A., et al.: Experience grounds language. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP). pp. 8718–8735 (2020)
- Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
- Chen, D.Z., Chang, A.X., Nießner, M.: Scanrefer: 3d object localization in rgb-d scans using natural language. In: European Conference on Computer Vision. pp. 202–221. Springer (2020)
- Chen, S., Guhur, P.L., Tapaswi, M., Schmid, C., Laptev, I.: Language conditioned spatial relation reasoning for 3d object grounding. Advances in Neural Information Processing Systems 35, 20522–20535 (2022)
- Corona, R., Zhu, S., Klein, D., Darrell, T.: Voxel-informed language grounding. In: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers). pp. 54–60 (2022)
- Csurka, G.: Domain adaptation for visual applications: A comprehensive survey. arXiv preprint arXiv:1702.05374 (2017)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A largescale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
- Devillers, B., Choksi, B., Bielawski, R., Vanrullen, R.: Does language help generalization in vision models? In: Proceedings of the 25th Conference on Computational Natural Language Learning. pp. 171–182 (2021)
- Diao, S., Xu, R., Su, H., Jiang, Y., Song, Y., Zhang, T.: Taming pre-trained language models with n-gram representations for low-resource domain adaptation. In: Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers). pp. 3336–3349 (2021)

- 16 P. Sun et al.
- 14. Diao, S., Xu, T., Xu, R., Wang, J., Zhang, T.: Mixture-of-domain-adapters: Decoupling and injecting domain knowledge to pre-trained language models' memories. In: Rogers, A., Boyd-Graber, J., Okazaki, N. (eds.) Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 5113-5129. Association for Computational Linguistics, Toronto, Canada (Jul 2023). https://doi.org/10.18653/v1/2023.acl-long.280, https://aclanthology.org/2023.acl-long.280
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2020)
- F, J.: Language models and linguistic theories beyond words. Nature Machine Intelligence 5(7), 677-678 (Jul 2023). https://doi.org/10.1038/s42256-023-00703-8, https://doi.org/10.1038/s42256-023-00703-8
- 17. Gong, Y., Yue, Y., Ji, W., Zhou, G.: Cross-domain few-shot learning based on pseudo-siamese neural network. Scientific Reports **13**(1), 1427 (2023)
- Guo, Z., Tang, Y., Zhang, R., Wang, D., Wang, Z., Zhao, B., Li, X.: Viewrefer: Grasp the multi-view knowledge for 3d visual grounding. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15372–15383 (2023)
- Gururangan, S., Marasović, A., Swayamdipta, S., Lo, K., Beltagy, I., Downey, D., Smith, N.A.: Don't stop pretraining: Adapt language models to domains and tasks. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. pp. 8342–8360 (2020)
- Hao, Y., Dong, L., Wei, F., Xu, K.: Visualizing and understanding the effectiveness of bert. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP). pp. 4143–4152 (2019)
- Harnad, S.: The symbol grounding problem. Physica D: Nonlinear Phenomena 42(1-3), 335–346 (1990)
- Hu, E.J., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: Lora: Low-rank adaptation of large language models. In: International Conference on Learning Representations (2021)
- Huang, S., Chen, Y., Jia, J., Wang, L.: Multi-view transformer for 3d visual grounding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15524–15533 (2022)
- 24. Li, J., Li, D., Savarese, S., Hoi, S.: BLIP-2: bootstrapping language-image pretraining with frozen image encoders and large language models. In: ICML (2023)
- Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: International Conference on Machine Learning. pp. 12888–12900. PMLR (2022)
- Lu, J., Batra, D., Parikh, D., Lee, S.: Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. Advances in neural information processing systems **32** (2019)
- Malik, B., Kashyap, A.R., Kan, M.Y., Poria, S.: Udapter-efficient domain adaptation using adapters. In: Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics. pp. 2241–2255 (2023)
- Mitra, C., Anwar, A., Corona, R., Klein, D., Thomason, J.: Comparative multiview language grounding. arXiv preprint arXiv:2311.06694 (2023)
- Miyanishi, T., Azuma, D., Kurita, S., Kawanabe, M.: Cross3dvg: Baseline and dataset for cross-dataset 3d visual grounding on different rgb-d scans. arXiv preprint arXiv:2305.13876 (2023)

Multi-Task Domain Adaptation for Language Grounding with 3D Objects

17

- 30. OpenAI: Gpt-4 technical report (2023)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. OpenAI Blog (2019)
- Roh, J., Desingh, K., Farhadi, A., Fox, D.: Languagerefer: Spatial-language model for 3d visual grounding. In: Conference on Robot Learning. pp. 1046–1056. PMLR (2022)
- 34. Schumann, R., Riezler, S.: Analyzing generalization of vision and language navigation to unseen outdoor areas. In: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 7519–7532 (2022)
- Shrivastava, A., Gopalakrishnan, K., Liu, Y., Piramuthu, R., Tür, G., Parikh, D., Hakkani-Tur, D.: Visitron: Visual semantics-aligned interactively trained objectnavigator. In: Findings of the Association for Computational Linguistics: ACL 2022. pp. 1984–1994 (2022)
- Song, Y., Sun, P., Fang, P., Yang, L., Xiao, Y., Zhang, Y.: Human-in-the-loop robotic grasping using bert scene representation. In: Proceedings of the 29th International Conference on Computational Linguistics. pp. 2992–3006 (2022)
- 37. Song, Y., Sun, P., Jin, P., Ren, Y., Zheng, Y., Li, Z., Chu, X., Zhang, Y., Li, T., Gu, J.: Learning 6-dof fine-grained grasp detection based on part affordance grounding (2024), https://arxiv.org/abs/2301.11564
- Song, Y., Sun, P., Liu, H., Li, Z., Song, W., Xiao, Y., Zhou, X.: Scene-driven multimodal knowledge graph construction for embodied ai (2023)
- 39. Štefánik, M., Novotný, V., Groverová, N., Sojka, P.: Adaptor: Objective-centric adaptation framework for language models. In: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations. pp. 261–269 (2022)
- 40. Sun, W., Khan, H., Guenon des Mesnards, N., Rubino, M., Arkoudas, K.: Unfreeze with care: Space-efficient fine-tuning of semantic parsing models. In: Proceedings of the ACM Web Conference 2022. pp. 999–1007 (2022)
- Tai, W., Kung, H., Dong, X.L., Comiter, M., Kuo, C.F.: exbert: Extending pretrained models with domain-specific vocabulary under constrained training resources. In: Findings of the Association for Computational Linguistics: EMNLP 2020. pp. 1433–1439 (2020)
- Thomason, J., Shridhar, M., Bisk, Y., Paxton, C., Zettlemoyer, L.: Language grounding with 3d objects. In: Conference on Robot Learning. pp. 1691–1701. PMLR (2022)
- Todorov, E., Erez, T., Tassa, Y.: Mujoco: A physics engine for model-based control. In: 2012 IEEE/RSJ international conference on intelligent robots and systems. pp. 5026–5033. IEEE (2012)
- Wang, Z., Liang, J., He, R., Xu, N., Wang, Z., Tan, T.: Improving zero-shot generalization for clip with synthesized prompts. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3032–3042 (2023)
- Yagubbayli, F., Wang, Y., Tonioni, A., Tombari, F.: Legoformer: Transformers for block-by-block multi-view 3d reconstruction. arXiv preprint arXiv:2106.12102 (2021)

- 18 P. Sun et al.
- Zhang, Y., Gong, Z., Chang, A.X.: Multi3drefer: Grounding text description to multiple 3d objects. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15225–15236 (2023)