

Mobile Manipulation Instruction Generation from Multiple Images with Automatic Metric Enhancement

Kei Katsumata, Motonari Kambara, Daichi Yashima, Ryosuke Korekata, and Komei Sugiura

Abstract—We consider the problem of generating free-form mobile manipulation instructions based on a target object image and receptacle image. Conventional image captioning models are not able to generate appropriate instructions because their architectures are typically optimized for single-image. In this study, we propose a model that handles both the target object and receptacle to generate free-form instruction sentences for mobile manipulation tasks. Moreover, we introduce a novel training method that effectively incorporates the scores from both learning-based and n -gram based automatic evaluation metrics as rewards. This method enables the model to learn the co-occurrence relationships between words and appropriate paraphrases. Results demonstrate that our proposed method outperforms baseline methods including representative multimodal large language models on standard automatic evaluation metrics. Moreover, physical experiments reveal that using our method to augment data on language instructions improves the performance of an existing multimodal language understanding model for mobile manipulation.

Index Terms—Deep Learning Methods, Deep Learning for Visual Perception, Mobile Manipulation

I. INTRODUCTION

THE advancement and deployment of service robots are essential in a variety of contexts such as elderly care facilities and daily support for disabilities. In particular, the integration of service robots in elderly care facilities significantly reduces the burden on caregivers and addresses the growing demand driven by the rise in the elderly population. Incorporating natural language understanding capabilities into robots would enhance their functionality and make them more user-friendly. To enhance their understanding of natural language instructions, training models with datasets containing high-quality instructions are essential. Nonetheless, the construction of such datasets presents a significant challenge: the required mobile manipulation instructions are labor-intensive to produce. This is because annotators are required to create sufficiently clear instructions based on multiple images. In fact, instructions for mobile manipulation are often complex and lengthy [1]. Consequently, the ability to automatically generate high-quality instructions would be highly beneficial.

In this study, we focus on the task of generating free-form mobile manipulation instructions based on a target object image and receptacle image. Fig. 1 shows a typical scene of this task. In this example, the images in Figs. 1 (a) and (b) are the given target object image and the receptacle image,

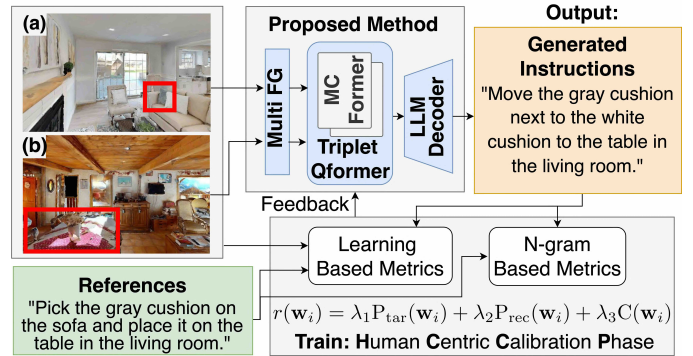


Fig. 1. Overview of our method. Our method generates a mobile manipulation instruction for a given target object image and receptacle image.

respectively. In addition, the target object is a gray cushion and the receptacle is a wooden table. We aim to generate a sentence such as “Move the gray cushion next to the white cushion to the table in the living room.” In this task, instructions need to be generated considering both the target object present in one image and the receptacle in the other image. Therefore, models are required to appropriately handle both images. Most existing image captioning models (e.g., [2], [3]) do not have architectures that can handle multiple images. Hence, these methods are inappropriate for generating mobile manipulation instructions based on multiple images.

We propose a model that generates mobile manipulation instructions using a target object image and a receptacle image. An overview of our method is also shown in Fig. 1. Our model identifies both the target object and receptacle, and it generates the corresponding instruction for the mobile manipulation task. We introduce Triplet Qformer, which enables the alignment of multiple visual features with respect to text features. Triplet Qformer embeds two types of visual features and one linguistic feature into the same space. This enables the grounding of visual features in natural language for both target objects and receptacles. Moreover, we introduce human centric calibration phase (HCCP), a training method that uses the human centric calibration training (HCCT) loss function in combination with learning-based and n -gram-based automatic evaluation metrics. Using HCCP, the model learns co-occurrence relationships between words and appropriate paraphrases of sentences.

The contributions of this study are as follows:

- We introduce Triplet Qformer. This architecture enables multiple types of visual features to be aligned individually with text features. The text features are used as anchors to align the visual features with each other.
- We introduce HCCP, a training method that uses the HCCT loss function with learning-based and n -gram-

Manuscript received: November 27, 2024; Accepted January 27, 2025. This paper was recommended for publication by Editor Vincze Markus upon evaluation of the Associate Editor and Reviewers’ comments. This work was partially supported by JSPS KAKENHI Grant Number 23K28168, JST Moonshot, NEDO, and JSPS Fellows Grant Number JP23KJ1917.

The authors are with Keio University, 3-14-1 Hiyoshi, Kohoku, Yokohama, Kanagawa 223-8522, Japan. {ke59ka77, motonari.k714, ydaichil207, rkorekata, komei.sugiura}@keio.jp

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

based automatic evaluation metrics.

Our code is available at this URL.¹

II. RELATED WORK

A. Multimodal Language Processing for Robotics

Research related to vision and language in robotics is widely conducted, and several survey papers exist [4], [5]. For example, many benchmarks for mobile manipulation tasks have been proposed [6], [7]. Most existing methods handle only template-based instructions or a limited number of fixed instruction types [6], [7]. On the other hand, some free-form instruction-based mobile manipulation methods [8]–[10] have also been proposed.

There are some standard datasets for the tasks with free-form natural language instruction (e.g., [1]). Matterport 3D (MP3D) [11] and Habitat-Matterport 3D (HM3D) [12] are representative benchmarks in research focused on indoor environments. These datasets are smaller in size compared to standard datasets used in other Vision & Language tasks. For example, the COCO dataset [13], a standard for image captioning, contains approximately 1.7M sentences, and the average sentence length is 10.47 words [14]. In contrast, the Remote Embodied Visual Referring Expression in Real Indoor Environments dataset [1], one of the representative datasets for mobile manipulation, contains only 21K sentences and the average sentence length is 18 words. Thus, compared with datasets for other tasks, those with natural language instructions for mobile manipulation are not adequate in size to enable effective training. Moreover, as can be inferred from the difference in average sentence length, instructions for mobile manipulation are often complex and lengthy. This complexity arises from the necessity to clarify tasks using referring expressions. Therefore, constructing such datasets poses a challenge because of the labor-intensive and time-consuming nature of annotation.

B. Data Augmentation

Data augmentation aims to create additional training data. Many methods have been proposed for the augmentation of captions [15], [16]. These approaches are often limited to word substitutions or paraphrasing. Nevertheless, text data can also be augmented using image captioning models. Wang et al. [17] generated large-scale data for VLN using photo-realistic environment. However, they found that the quality of captions in their work was limited.

C. Image Captioning

Image captioning is the task of generating sentences based on a given image. Many image captioning methods have been proposed [2], [3]. Recent developments in LLMs have enhanced their expressive capabilities [3], [18]. BLIP-2 [3] is a representative captioning model employing an LLM decoder. It leverages a two-stage pretraining strategy to bridge the modality gap between vision and language. However, these methods are not designed to handle two images simultaneously.

Many training methods have also been proposed for image captioning models. In particular, methods that use automatic evaluation metrics as reward functions are widely employed [2], [19]. Self-critical sequence training [19] aligns the training process with the evaluation metric, directly optimizing for CIDEr [20]. Cho et al. [21] also introduced a strategy that combines n -gram based metrics and learning-based metrics. Though this approach has been successful in retrieval tasks, approaches that do not combine these metrics have remained more effective in image captioning tasks. In this study, we introduce the novel training method HCCP, which effectively incorporates both learning-based and n -gram-based metrics into the loss function.

III. PROBLEM STATEMENTS

In this paper, we focus on the mobile manipulation instruction generation (MMIG) task. We define the OMIG task as follows: given two images, one of a target object and one of a receptacle, the model generates a free-form instruction sentence for fetching the target object and carrying it to the receptacle. In this task, appropriate mobile manipulation instructions that include the target object and the receptacle based on the input images are to be generated. Fig.1 shows a typical scene from the OMIG task. In the images in Fig.1 (a) and (b), the target object and the receptacle are a gray cushion and a table, respectively. The goal is to generate an instruction sentence, such as “Move the gray cushion next to the white cushion to the table in the living room.” The terms used in this paper are defined as follows: A ‘target object’ refers to an object needed in daily life, identified as the object to be grasped in the instruction, and a ‘receptacle’ refers to a piece of furniture identified as the destination in the instruction where the target is to be placed.

The input for the task consists of an image of the target object (target object image) and the receptacle (receptacle image). Bounding boxes for specific objects are not provided. The output is a free-form mobile manipulation instruction, which includes both the target object and the receptacle. In this study, we assume that the model is required to generate an instruction for a single pick-and-place task and path planning is not required for this task. In addition, we assume that reference instructions are executable; in other words, if we can generate instructions similar to the reference, then the feasibility of execution is ensured.

IV. PROPOSED METHOD

We propose a model that generates mobile manipulation instructions by identifying the target object and receptacle in each image, respectively. The proposed model consists of three main modules: Multi Image Feature Generator, Triplet Qformer, and LLM Decoder. Fig. 2 shows the structure of our method. The input to our model is defined as $x = \{\mathbf{X}_{\text{tar}}, \mathbf{X}_{\text{rec}}\}$. Here, \mathbf{X}_{tar} and \mathbf{X}_{rec} denote a target object image and a receptacle image, respectively.

A. Multi Image Feature Generator

Multi Image Feature Generator generates the region feature \mathbf{H}_R , which includes local features of objects, and the grid

¹<https://github.com/keio-smilab24/MMIG>

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

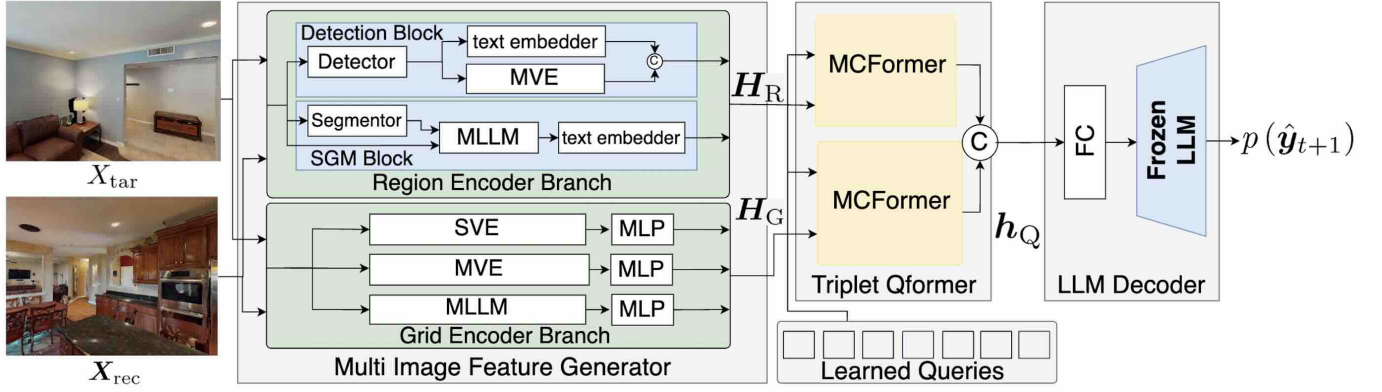


Fig. 2. Overall architecture of the proposed model. \odot , SVE, MVE, MLLM, and FC represent concatenation, single-modal visual encoder, multimodal visual encoder, multimodal large language model, and fully connected layer, respectively.

feature H_G , which include the global features of the entire image, from x . Relying solely on either type of feature can lead to a lack of information regarding the details of objects or spatial relationships between objects. Therefore, using both types of features enables effective image representations. This module consists of two branches: the region encoder branch and the grid encoder branch.

1) **Region encoder branch:** The region encoder branch computes the region features, which contain detailed information about objects. The effectiveness of region features has been reported [2]. Most existing methods use single-modal region features based solely on image features. Integrating linguistic information into region features can help to align text with image features more appropriately. Therefore, our method uses multimodal region features that integrate image features obtained by the object detector with text features derived from the predicted labels. Moreover, we leverage a multimodal LLM (MLLM) to obtain comprehensive textual descriptions of the objects in the image, including their spatial relationships and detailed characteristics. The region encoder branch consists of two blocks: the detection block and the spatial grounded marks (SGM) block.

The region encoder branch is applied to each image X_{tar} and X_{rec} independently. The following explanation focuses on X_{tar} , because the same process is applied to X_{rec} . In this branch, we obtain two types of region features: the detection image feature $h_{R,D}$ in the detection block and SGM feature $h_{R,S}$ in the SGM block from X_{tar} .

First, in the Detection block, we obtain the sets of detected objects D_k with predicted labels using an object detector from X_{tar} . Here, k denotes the number of detected objects. We provide the detector with a dictionary containing noun phrases extracted from the training data. The dictionary is used to map the detector's outputs to the corresponding noun phrases, ensuring more appropriate association between image features of detected objects and text features from their labels. Using a multimodal vision encoder (e.g., [22]), we extract image feature $v_{De,k}$ from D_k . Moreover, we extract text feature s_k from the predicted text labels of the detected objects. Then, we obtain the feature $h_{R,D} = [v_{De,k}; s_k]$.

In the SGM block, we first apply marks to the images with segmentation masks using segmentation models and obtain

descriptions using an MLLM from those images following a method similar to Set of Marks (SoM) [23]. Then, we acquire the SGM feature $h_{R,S}$ from the description. By executing the aforementioned procedure on X_{rec} , we obtain detection region feature $h'_{R,D}$ and the SGM feature $h'_{R,S}$. Finally, we obtain the region feature H_R as

$$H_R = \{[h_{R,D}; h'_{R,D}], [h_{R,S}; h'_{R,S}]\}. \quad (1)$$

2) **Grid encoder branch:** While the region features are helpful for capturing local object characteristics, they have a limited ability to provide a comprehensive representation of the image. In many cases, region features do not include information about the space between objects, making it challenging to obtain positional relationships. However, as previously mentioned, complex referring relationships, such as spatial relationships between objects, are crucial for identifying objects. To address these limitations, we introduce a grid encoder branch which computes grid features from the entire image. The grid features are capable of acquiring global information, including the spatial relationships between objects. The grid encoder branch generates three types of latent representations, and the grid features are obtained using three encoders: a single-modal visual encoder (e.g., DINOv2 [24]), a pre-trained multimodal visual encoders (e.g., CLIP), and an MLLM (e.g., LLaVA [25]). Each latent representation individually emphasizes different aspects: a single-modal visual encoder captures fine-grained visual details, a pre-trained multimodal visual encoder addresses structural and relational aspects, and the MLLM aligns these aspects with natural language for comprehensive understanding [26]. Combining these multimodal features is an effective way to generate mobile manipulation instructions using diverse expressions. Grid features $h_{G,D}$, $h_{G,C}$, and $h_{G,L}$ are obtained from X_{tar} as $(h_{G,D}, h_{G,C}, h_{G,L}) = (\text{MLP}(v_D), \text{MLP}(v_C), \text{MLP}(v_L))$, where v_D , v_C , v_L , and $\text{MLP}(\cdot)$ denote the visual features extracted by a single-modal visual encoder, the visual features extracted by a multimodal encoder, the visual features extracted by an MLLM, and a linear transformation, respectively. Similarly, we obtain image features $h'_{G,D}$, $h'_{G,C}$, and $h'_{G,L}$ from X_{rec} . Finally, we obtain a grid feature

$$H_G = \{[h_{G,D}; h'_{G,D}], [h_{G,C}; h'_{G,C}], [h_{G,L}; h'_{G,L}]\}. \quad (2)$$

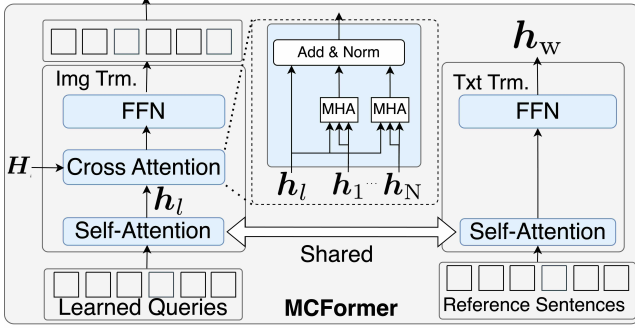


Fig. 3. The details of MCFormer. Img Trm., Txt Trm., FFN and MHA represent image transformer, text transformer, feed-forward network and multi-head attention, respectively. Here, we define \mathbf{H} as $\mathbf{H} = \{\mathbf{h}_1, \dots, \mathbf{h}_N\}$.

B. Triplet Qformer

Next, we introduce Triplet Qformer, which aligns the two kinds of visual features with text features from the ground-truth description, using text features as the anchor. This approach indirectly aligns the three types of features. By aligning the entirety of the reference sentences with the object and receptacle features, the model can learn co-occurrence relationships between pairs of target objects and receptacles. This module takes \mathbf{H}_G , \mathbf{H}_R , and learnable queries as inputs. During pretraining, the reference sentence \mathbf{y} is also used as input. The output is multimodal feature \mathbf{h}_Q . Triplet Qformer consists of two Multi Image Cross Attention Transformers (MCFormers), the structure of which is shown in Fig. 3. The first and second MCFormers use \mathbf{H}_G and \mathbf{H}_R as the image features, respectively. For readability, the following describes a block that uses \mathbf{H}_G as the image feature. MCFormer is an extension of Qformer [3], which was designed to integrate multiple visual features and align these visual features with text features.

As shown in Fig. 3, the text and image features are aligned through the shared self-attention layers in the two transformers: the image transformer and the text transformer. For the image transformer, each layer includes a self-attention layer, a cross-attention layer, and a feed-forward network layer. First, in the self-attention layer, we obtain \mathbf{h}_l by applying self-attention to the learnable query embedding. The parameters of the self-attention layer are shared with the self-attention layer in the text transformer. As shown in Fig. 3, the cross-attention layer consists of N multi-head attention (MHA) blocks. Here, N represents the number of elements in a set of image features. That is to say, $N = 3$ and $N = 2$ for \mathbf{H}_G and \mathbf{H}_R , respectively. Note that the parameters of each MHA block are not shared. The output of each MHA block is denoted as \mathbf{a}_n , where $n = 1, \dots, N$. We obtain an intermediate feature

$$\mathbf{h}_{Q,G} = \text{FFN} \left(\text{LN} \left(\sum_{n=1}^N \mathbf{a}_n + \mathbf{h}_l \right) \right), \quad (3)$$

where $\text{LN}(\cdot)$ and $\text{FFN}(\cdot)$ denote the layer normalization and feed-forward network, respectively. By applying the same procedure to \mathbf{H}_R , we also obtain the feature $\mathbf{h}_{Q,R}$. Finally, we obtain intermediate feature $\mathbf{h}_Q = [\mathbf{h}_{Q,G}; \mathbf{h}_{Q,R}]$.

For the text transformer, each layer includes a self-attention layer and feed-forward network layer. The output of the text

transformer is \mathbf{h}_w . This is used only during training, as discussed in Section IV-D.

C. LLM Decoder

Existing image captioning models do not generalize well to out-of-domain images containing novel scenes or objects. Therefore, we incorporate a pre-trained LLM as decoder. The LLM possesses highly expressive capabilities and an extensive vocabulary, which enhances its linguistic expressiveness. LLM Decoder generates an instruction sentence based on \mathbf{h}_Q . In this module, we first use a fully connected layer to linearly project \mathbf{h}_Q . Then, we obtain the probability for the next token $p(\mathbf{y}_{t+1} | \mathbf{x}, \mathbf{y}_{1:t})$ from a frozen LLM. Here, \mathbf{y}_{t+1} , and $\mathbf{y}_{1:t}$ represent the predicted token at time $t + 1$ and the predicted token sequence up to time t , respectively.

D. Training Stages

Our model is trained in three stages: Triplet Qformer pre-training phase (TQPP), probability distribution matching phase (PDMP), and human centric calibration phase (HCCP). In TQPP, we pre-train Triplet Qformer using a loss function similar to that used in Qformer [3]. Next, in PDMP, we train our entire model using the cross-entropy loss function. Finally, in HCCP, we train the model selected in PDMP using the human centric calibration training (HCCT) loss function. The HCCT loss function employs both n -gram and learning-based automatic evaluation metrics. Existing methods [2], [19] have proposed loss functions that utilize n -gram-based metrics. This enables them to learn the co-occurrence relationships between words present in the reference sentences. However, several studies have reported that the performance of n -gram-based metrics is insufficient [29], [30]. In particular, it is challenging to appropriately evaluate paraphrases of reference sentences. Some learning-based automatic evaluation metrics [29], [30] are reported to have a higher correlation coefficient with human evaluations than those n -gram-based metrics. The evaluation metrics based on the similarity between the generated sentences and images can appropriately evaluate paraphrases. Therefore, by incorporating both learning-based and n -gram-based metrics into the loss function, the frozen LLM decoder's expressiveness should be leveraged more effectively for generating instruction sentences.

We use different loss functions for each stage. In TQPP, we employ the loss function $\mathcal{L}_{\text{TQPP}}$. We delineate the loss function in the training of the block that handles \mathbf{h}_G . Here, $\mathcal{L}_{\text{TQPP}}$ is defined as

$$\mathcal{L}_{\text{TQPP}} = \mathcal{L}_{\text{itc}} + \mathcal{L}_{\text{CE}}(\mathbf{y}, \mathbf{h}_w) + \mathcal{L}_{\text{BCE}}(\mathbf{h}_{Q,G}), \quad (4)$$

where $\mathcal{L}_{\text{CE}}(\cdot)$, \mathbf{y} , \mathbf{h}_w , and $\mathcal{L}_{\text{BCE}}(\cdot)$ denote the cross-entropy function loss, a reference sentence, an output text feature of MCFormer, and the binary cross-entropy function loss, respectively. We use the image-text contrastive loss \mathcal{L}_{itc} defined in [3]. In PDMP, we employ the cross-entropy function. Furthermore, in HCCP, we use the HCCT loss function $\mathcal{L}_{\text{HCCT}}$. $\mathcal{L}_{\text{HCCT}}$ is defined as follows:

$$\mathcal{L}_{\text{HCCT}} = -\frac{1}{k} \sum_{i=1}^k (r(\mathbf{w}_i) - b) \log p(\mathbf{w}_i). \quad (5)$$

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE I

QUANTITATIVE COMPARISON BETWEEN THE PROPOSED METHOD AND BASELINE METHODS ON THE TEST SETS OF THE HM3D-FC DATASET. DENSE CAPTION REFERS TO THE METHOD DENSE CAPTION BASELINE. TARGET AND RECEPTACLE CORRESPOND TO THE SCORES FOR THE TARGET OBJECT IMAGES AND THE RECEPTACLE IMAGES, RESPECTIVELY. THE BEST SCORE FOR EACH METRIC IS IN BOLD.

Method	Polos \uparrow		PAC-S \uparrow		RefPAC-S \uparrow		SPICE \uparrow	CIDEr \uparrow	BLEU4 \uparrow
	Target	Receptacle	Target	Receptacle	Target	Receptacle			
Dense Caption	30.0 \pm 5.6	30.8 \pm 6.0	70.3 \pm 2.1	69.6 \pm 2.1	73.5 \pm 2.2	73.3 \pm 2.1	10.0 \pm 2.9	25.1 \pm 7.3	6.6 \pm 0.7
GRIT [2]	41.2 \pm 1.8	39.0 \pm 1.8	70.8 \pm 1.4	67.6 \pm 1.9	75.5 \pm 1.1	73.7 \pm 1.6	19.8 \pm 0.5	59.2 \pm 2.8	10.5 \pm 0.3
BLIP-2 [3]	42.4 \pm 1.4	40.1 \pm 1.5	72.3 \pm 0.8	67.4 \pm 0.9	74.7 \pm 0.6	73.6 \pm 0.7	16.9 \pm 0.4	36.6 \pm 4.9	8.5 \pm 1.2
Gemini [27]	29.5 \pm 0.2	29.4 \pm 0.2	69.7 \pm 0.5	68.1 \pm 0.2	73.3 \pm 0.5	72.7 \pm 0.2	11.2 \pm 0.6	26.2 \pm 1.6	5.2 \pm 0.4
GPT-4o [28]	34.5 \pm 0.2	35.5 \pm 0.06	72.1 \pm 1.4	70.7 \pm 0.5	75.4 \pm 0.8	75.0 \pm 0.7	15.2 \pm 0.5	33.4 \pm 1.5	7.5 \pm 0.2
Ours	50.9 \pm 0.8	50.7 \pm 0.7	74.3 \pm 0.6	72.1 \pm 0.7	79.8 \pm 0.4	78.5 \pm 0.5	22.7 \pm 0.4	64.8 \pm 4.0	11.5 \pm 1.1

Here, \mathbf{w}_i , $r(\mathbf{w}_i)$, b , and k denote the i -th generated sentence in the beam, reward function, reward baseline, and index of the sample in the batch, respectively. Moreover, $r(\mathbf{w}_i)$ and b are defined as follows:

$$r(\mathbf{w}_i) = \lambda_1 P_{\text{tar}}(\mathbf{w}_i) + \lambda_2 P_{\text{rec}}(\mathbf{w}_i) + \lambda_3 C(\mathbf{w}_i),$$

$$b = \frac{1}{k} \sum_{i=1}^k r(\mathbf{w}_i), \quad (6)$$

where $P_{\text{tar}}(\cdot)$, $P_{\text{rec}}(\cdot)$, and $C(\cdot)$ represent the Polos score for the target object image, the Polos score for the receptacle image, and the CIDEr [20] score, respectively. Here, λ_1 , λ_2 , and λ_3 are hyperparameters.

V. EXPERIMENTAL SETTINGS

In the experiments, we used the Habitat-Matterport 3D Dataset (HM3D) subset of the Learning-To-Rank in Real Indoor Environments for Fetch-and Carry (LTRRIE-FC) dataset (HM3D-FC [10]). The dataset for the OMIG task should include target object images, receptacle images, and instruction sentences corresponding to each image pair. Most standard datasets for the VLN task (e.g., [1]) include information related to path planning in instructions. Such datasets were not appropriate for this study because we did not address path planning. By contrast, the HM3D-FC dataset includes instruction sentences that are acceptable for our task. Thus, we used the HM3D-FC dataset.

In our experiments, the reward weights in the human centric calibration training loss function, we use 0.25, 0.25, and 0.5 for λ_1 , λ_2 , and λ_3 , respectively. We trained our model on a GeForce RTX4090 with 24 GB of GPU memory. It took approximately 16 hours in total to train our model. The inference time was approximately 92 ms/sample.

In the region feature branch in Multi Image Feature Generator, we employed the following pre-trained models. In the Detection block, we employed Detic [31], CLIP, and RoBERTa [32] as the detector, multimodal visual encoder for the detected object images, and text embedder for the labels of detected objects, respectively. Additionally, in the SGM block, we used SEEM [33], GPT-4V [28], and embedding-large-3 [34] as the segment model, MLLM, and text embedder, respectively. In the grid feature branch, DINOv2, CLIP, LLaVA were used for single-modal visual encoder, multimodal encoder, and MLLM, respectively. In LLM Decoder, we used OPT-2.7B [35] as the frozen LLM.

VI. EXPERIMENTAL RESULTS

A. Quantitative Results

Table I shows the quantitative results of the baseline and proposed methods on the HM3D-FC dataset. The values in the table show the average and standard deviation over five trials. We used Dense Caption Baseline, GRIT, BLIP-2, Gemini [27], and GPT-4o [28] as the baseline methods. Dense Caption Baseline was the method used to generate dense captions for target object and receptacle images with BLIP-2 and fused them into instructions using an LLM. We employed GRIT and BLIP-2 because they are representative methods for image captioning in the COCO [13] benchmark. Moreover, we also used Gemini and GPT-4o because they are representative MLLMs that have been reported to produce favorable results in many vision & language tasks. We trained GRIT and BLIP-2 on HM3D-FC dataset. For models that can only handle a single input image, only the target object image was used as input. Experiments with the MLLM models, Gemini and GPT-4o, were conducted in a few-shot setting.

In this study, we used Polos, PAC-S [30], RefPAC-S [30], SPICE [36], CIDEr, and BLEU4 as the evaluation metrics. Polos is a supervised automatic evaluation metric for image captioning models. Polos computes scores from multimodal inputs, using a parallel feature extraction mechanism that leverages embeddings trained through large-scale contrastive learning. PAC-S and Ref-PAC-S evaluate captions in an unsupervised manner by computing their similarity with embeddings derived from fine-tuned CLIP. Ref-PAC-S uses reference-based metrics in its evaluation. By contrast, PAC-S is a reference-free metric that calculates scores based solely on the similarities between image features and text features. We used Polos and RefPAC-S because scores in these metrics correlate strongly with human judgment. To evaluate the sentences generated by MLLMs appropriately, we also employed PAC-S. This is a representative reference-free evaluation metric that considers only the similarities between the generated sentences and images. These evaluation metrics only calculate the similarity between an instruction and a single image. They cannot handle multiple images as simultaneous inputs. Therefore, we independently calculated scores for the target object images and the receptacle images. In addition, we also used SPICE, CIDEr, and BLEU4 because these are standard metrics in image captioning.

These results indicate that the proposed method outperformed BLIP-2 (which had the highest Polos scores of the

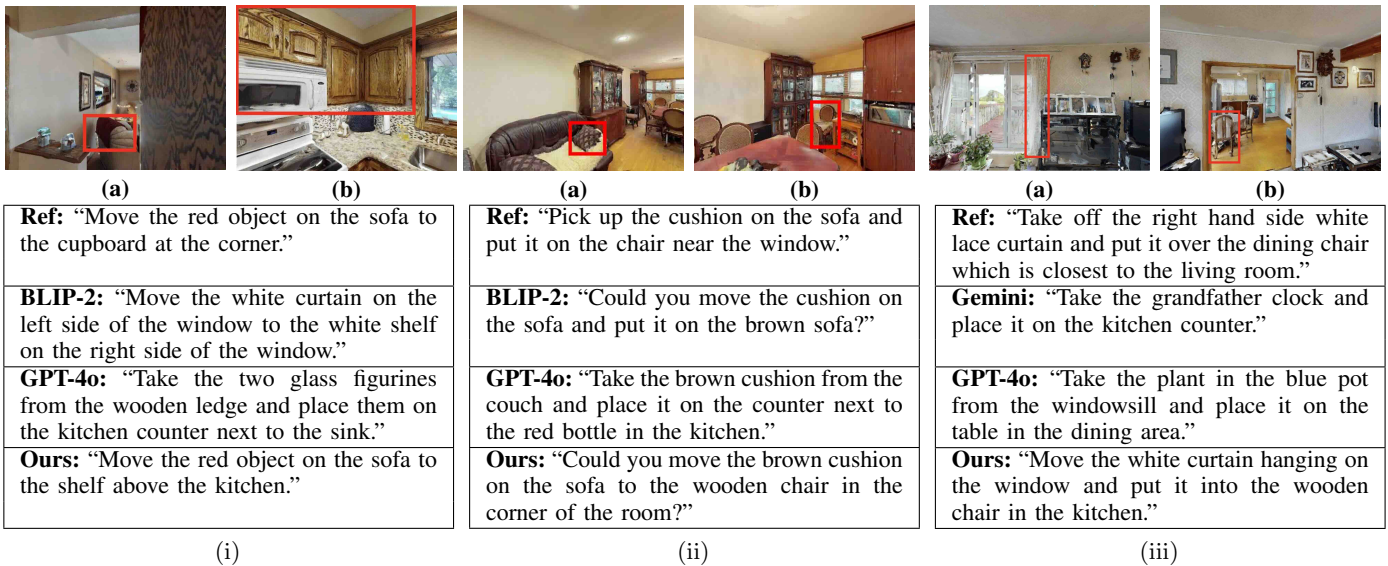


Fig. 4. Succeeded samples with the proposed method. (a) and (b) show the target object image and the receptacle image, respectively. The bounding box indicates the target object and receptacle that are included in the references.

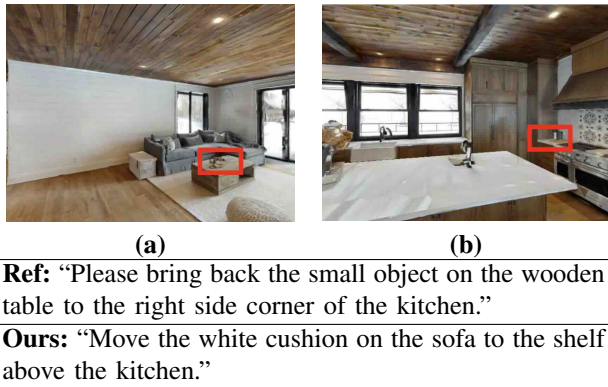


Fig. 5. A failed samples obtained by the proposed method. (a) Target object image and (b) receptacle image, respectively.

baseline methods) by 8.5 points for the target object images and 10.6 points for the receptacle images. Similarly, the proposed method outperformed the baseline methods with respect to PAC-S score for target object images and receptacle images, with improvements of 2.0, and 1.4 points compared with the best-performing baseline methods for each metric, respectively. It indicates that even with a representative reference-free evaluation metric, the proposed method outperformed all baseline methods. Moreover, the proposed method outperformed in RefPAC-S, SPICE, CIDEr, and BLEU4 compared with the highest baseline in each metric. There were statistically significant differences in the performances of our method and the baseline methods in terms of all evaluation metrics (p -value < 0.05).

B. Qualitative Results

Fig. 4 shows successful examples obtained by using the proposed method. In Fig. 4 (i), the target object and receptacle were a red object and cupboard, respectively. The reference was "Move the red object on the sofa to the cupboard at the corner." Our method generated the sentence "Move the red object on the sofa to the shelf above the kitchen." Thus, our method described the target object and the receptacle using appropriate colors and referring expressions. By contrast, BLIP-2 and GPT-4o output "Move the white curtain

on the left side of the window to the white shelf on the right side of the window." and "Take the two glass figurines from the wooden ledge and place them on the kitchen counter next to the sink.", respectively. In these generated sentences, the descriptions of both the target object and the receptacle object were inappropriate. In Fig. 4 (ii), the target object and receptacle were a cushion and a chair, respectively. The reference sentence for this example was "Pick up the cushion on the sofa and put it on the chair near the window." The proposed method generated the sentence "Could you move the brown cushion on the sofa to the wooden chair in the corner of the room?", appropriately included both the target object and the receptacle in each image. By contrast, BLIP-2 and GPT-4o output "Could you move the cushion on the sofa and put it on the brown sofa?" and "Take the brown cushion from the couch and place it on the counter next to the red bottle in the kitchen.", respectively. In the sentence generated by BLIP-2, the descriptions of both the target object and the receptacle object represented the wrong color. On the other hand, the generated sentence by GPT-4o refers to a receptacle with the nonexistence object 'red bottle'. In the case of Fig. 4 (iii), the target object and receptacle were a curtain and a chair, respectively. The reference was "Take off the right hand side white lace curtain and put it over the dining chair which is closest to the living room." Our method generated the sentence "Move the white curtain hanging on the window and put it into the wooden chair in the kitchen.", describing the positions and materials of the target object and the receptacle appropriately. By contrast, Gemini and GPT-4o generated the sentences "Take the grandfather clock and place it on the kitchen counter." and "Take the plant in the blue pot from the windowsill and place it on the table in the dining area.", respectively. GPT-4o incorrectly recognized the color of the pot, while Gemini misidentified the small clock-like objects as a 'grand-father clock'.

Fig. 5 shows a failed example of the proposed method. In Fig. 5, the target object and receptacle were a small object and right side corner of the kitchen, respectively. For this example,

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

the reference sentence was “Please bring back the small object on the wooden table to the right side corner of the kitchen.” The sentence generated by the proposed method was “Move the white cushion on the sofa to the shelf above the kitchen.” The proposed method incorrectly described the target object and the receptacle as different from those in the reference captions. Moreover, it incorrectly described the color of the target object (cushions) as white, despite the target object image showing gray cushions.

C. Ablation Studies

Table II shows the results of the ablation studies. The following two ablation conditions were set:

HCCP Ablation: The proposed method was trained using the three stages. The effectiveness of the human-centric calibration phase (HCCP) was then investigated by removing this stage. As shown in Table II, Model (i) achieved Polos scores of 44.7 and 44.4 for the target object images and receptacle images, respectively. Compared with Model (iii), the Polos scores obtained were 6.2 and 6.3 points lower. In fact, Model (iii) outperformed Model (i) in all evaluation metrics. These results indicate that training with the HCCP improved the quality of instruction sentences, allowing them to more closely resemble instructions given by humans.

Triplet Qformer Ablation: The effectiveness of Triplet Qformer was demonstrated by replacing it with a standard Qformer [3]. In this comparison, the concatenated features of H_G and H_R were used as the input to the Qformer in Model (ii). According to Table II, Model (ii) achieved Polos scores of 46.2 and 44.7 for the target object images and receptacle images, respectively. Compared with Model (iii), the Polos scores decreased by 4.7 and 6.0 points, respectively. In general, Model (iii) outperformed Model (ii) across all evaluation metrics. These results reveal that the structure of Triplet Qformer enables the two types of visual features to be appropriately aligned with text features.

D. Dataset Augmentation Experiments

To investigate the effectiveness of the instruction sentences generated by the proposed method for data augmentation, we conducted data augmentation experiments. In these experiments, we focused on the Image Retrieval-based Open-Vocabulary Fetch-and-Carry (IROV-FC) task [10]. For this task, a system in which a robot identifies and retrieves images of the specified target object and receptacle based on a free-form mobile manipulation instruction must be developed. This task was selected because of the applicability of the sentences generated by the proposed method. We used mean reciprocal rank (MRR) and recall@ K ($R@K$) because they are standard metrics used in machine learning problems that involve ranking [9], [10]. We used MultiRankIt [9] as an IROV-FC model because it is a representative IROV-FC model. We trained this model on the LTRRIE-FC dataset [10] which is a superset of the HM3D-FC dataset. Here, we trained MultiRankIt for target objects and receptacles separately as it cannot handle both modes within a single model. Therefore, we evaluated MultiRankIt for target objects and receptacles individually, and then, calculated the average.

We created the Aug-train dataset with instruction sentences generated using images in the training set of the HM3D-FC dataset. To build the Aug-train dataset, we attached a single instruction sentence generated by the proposed method to each image included in the training set of the HM3D-FC dataset. Here, the proposed method was trained on HM3D-FC. We set the following four conditions in the data augmentation experiment:

- (i) **half:** Half of the training set of the LTRRIE-FC dataset.
- (ii) **half + Aug:** Half of the training set of the LTRRIE-FC dataset (randomly chosen) plus the Aug-train dataset.
- (iii) **full:** All of the training samples in the training set of the LTRRIE-FC dataset.
- (iv) **full + Aug:** The full training set of the LTRRIE-FC and Aug-train datasets.

Table III shows the quantitative results. From the table, Models (iv-T) and (iv-R) achieved MRRs of 22.3% and 22.1%, whereas Models (iii-T) and (iii-R) achieved 20.5% and 19.8%, respectively. These results indicate that Models (iv-T) and (iv-R) outperformed Models (iii-T) and (iii-R) by 1.8 points and 2.3 points in terms of MRR on the HM3D-FC test set. Similarly, the models with augmented datasets outperformed those trained without augmented datasets in terms of $R@K$. We confirmed that performing data augmentation on language instructions with our model improved the performance of the existing IROV-FC model.

E. Physical Experiments

In this experiment, we confirmed that dataset augmentation using the proposed method is effective even in the physical experiments. The experimental environment and robot platform were set up according to [10]. We validated the effects of data augmentation using MultiRankIt trained on conditions (iii) and (iv), as described in Section VI-D. In the experiments, the robot executed mobile manipulation based on free-form instructions given by the users. The experiments were conducted over 20 episodes: four different environmental settings times five episodes in per setting. We performed the remaining settings according to the physical experiment in [10]. We used MRR, $R@10$ and task success rate (SR) as the evaluation metrics for the physical experiments. Table IV shows that Model (iv) achieved SR of 55%, which was 10 points higher than the unaugmented model. These results indicate that data augmentation on language instructions with our model contributed to improving the performance of MultiRankIt, leading to an increase in the SR of fetching and carrying actions.

VII. CONCLUSIONS

In this study, we focused on the mobile manipulation instruction generation task. The contributions of this study were as follows. We introduced Triplet Qformer, which enabled the alignment of multiple types of visual features individually with text features and then used the text features as an anchor to align the visual features with each other. We also introduced human centric calibration phase, a training method that uses human centric calibration training loss function with learning-based and n -gram-based automatic evaluation metrics. Our

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE II
QUANTITATIVE RESULTS OF THE ABLATION STUDIES. HERE, TQ REFERS TO TRIPLET QFORMER.

Model	Polos \uparrow		PAC-S \uparrow		RefPAC-S \uparrow		SPICE \uparrow	CIDEr \uparrow	BLEU4 \uparrow
	Target	Receptacle	Target	Receptacle	Target	Receptacle			
(i) w/o HCCP	44.7 \pm 0.1	44.4 \pm 0.2	72.4 \pm 0.2	70.6 \pm 0.2	77.8 \pm 0.1	76.7 \pm 0.1	20.9 \pm 0.5	53.3 \pm 1.6	9.0 \pm 0.5
(ii) w/o TQ	46.2 \pm 0.2	44.7 \pm 0.2	74.2 \pm 0.2	69.9 \pm 0.2	78.6 \pm 0.1	76.0 \pm 0.1	18.3 \pm 0.3	55.3 \pm 1.4	10.2 \pm 0.2
(iii) full	50.9 \pm 0.8	50.7 \pm 0.7	74.3 \pm 0.6	72.1 \pm 0.7	79.8 \pm 0.4	78.5 \pm 0.5	22.7 \pm 0.4	64.8 \pm 4.0	11.5 \pm 1.1

TABLE III

QUANTITATIVE RESULTS OF THE DATA AUGMENTATION EXPERIMENT. (*-T) AND (*-R) DENOTE THE MODELS FOR THE TARGET OBJECTS AND THE RECEPTACLES, RESPECTIVELY. THE NUMBERS IN BOLD INDICATE THE BEST SCORES FOR EACH IMAGE IN EACH METRIC.

[%]	Condition	HM3D-FC (unseen)			
		MRR \uparrow	R@5 \uparrow	R@10 \uparrow	R@20 \uparrow
(i-T)	half	14.9 \pm 2.0	19.0 \pm 3.7	37.1 \pm 3.9	64.3 \pm 4.9
(i-R)		17.9 \pm 1.1	25.3 \pm 3.5	43.2 \pm 3.4	70.3 \pm 1.7
(ii-T)	half + Aug	19.5 \pm 2.4	27.0 \pm 4.2	44.2 \pm 3.7	69.8 \pm 4.3
(ii-R)		18.7 \pm 2.0	26.8 \pm 4.3	44.5 \pm 4.4	70.1 \pm 4.0
(iii-T)	full	20.5 \pm 2.3	30.1 \pm 3.4	48.2 \pm 1.4	73.2 \pm 2.8
(iii-R)		19.8 \pm 1.1	27.1 \pm 3.2	49.1 \pm 5.9	74.6 \pm 3.1
(iv-T)	full + Aug	22.3 \pm 1.8	31.9 \pm 4.4	50.2 \pm 3.5	75.1 \pm 4.3
(iv-R)		22.1 \pm 1.2	32.0 \pm 3.8	52.3 \pm 1.9	77.3 \pm 3.4

TABLE IV

QUANTITATIVE RESULTS FOR THE PHYSICAL EXPERIMENTS.

[%]	Condition	MRR \uparrow	R@10 \uparrow	SR \uparrow
(iii)	full	21	51	45 (9/20)
(iv)	full + Aug	23	58	55 (11/20)

proposed method outperformed the baseline methods including representative MLLMs on automatic evaluation metrics on the HM3D-FC dataset. We confirmed that performing data augmentation on language instructions with our model improved the performance of an existing Image Retrieval-based Open-Vocabulary Fetch-and-Carry model in both the dataset augmentation and physical experiments.

There were still some samples with color-related errors in the generated sentences by the proposed method. Therefore, introducing color detectors that label the color of each detected object and use it as an additional feature is one task for future work.

REFERENCES

- Y. Qi *et al.*, "REVERIE: Remote Embodied Visual Referring Expression in Real Indoor Environments," in *CVPR*, 2020, pp. 9982–9991.
- V. Nguyen *et al.*, "GRIT: Faster and Better Image Captioning Transformer Using Dual Visual Features," in *ECCV*, 2022, pp. 167–184.
- J. Li, D. Li, S. Savarese, and S. Hoi, "BLIP-2: Bootstrapping Language-Image Pre-Training with Frozen Image Encoders and Large Language Models," in *ICML*, 2023, pp. 19730–19742.
- R. Firoozi *et al.*, "Foundation Models in Robotics: Applications, Challenges, and the Future," *arXiv preprint arXiv:2312.07843*, 2023.
- K. Kawaharazuka *et al.*, "Real-World Robot Applications of Foundation Models: A Review," *arXiv preprint arXiv:2402.05741*, 2024.
- L. Iocchi, D. Holz, J. Ruiz-del Solar, K. Sugiura, and T. Van Der Zant, "RoboCup@Home: Analysis and Results of Evolving Competitions for Domestic and Service Robots," *AIJ*, vol. 229, pp. 258–281, 2015.
- H. Okada, T. Inamura, and K. Wada, "What Competitions Were Conducted in the Service Categories of the World Robot Summit?" *AR*, vol. 33, no. 17, pp. 900–910, 2019.
- D. Driess, F. Xia, M. Sajjadi, C. Lynch, *et al.*, "PaLM-E: An Embodied Multimodal Language Model," in *ICML*, 2023, pp. 8469–8488.
- S. Kaneda, S. Nagashima, R. Korekata, M. Kambara, *et al.*, "Learning-To-Rank Approach for Identifying Everyday Objects Using a Physical-World Search Engine," *IEEE RA-L*, vol. 9, no. 3, pp. 2088–2095, 2024.
- R. Korekata, K. Kaneda, *et al.*, "DM2RM: Dual-Mode Multimodal Ranking for Target Objects and Receptacles Based on Open-Vocabulary Instructions," *arXiv preprint arXiv:2408.07910*, 2024.
- A. Chang, A. Dai, T. Funkhouser, *et al.*, "Matterport3D: Learning from RGB-D Data in Indoor Environments," in *3DV*, 2017, pp. 667–676.
- S. Ramakrishnan *et al.*, "Habitat-Matterport 3D Dataset (HM3D): 1000 Large-scale 3D Environments for Embodied AI," in *NeurIPS*, 2021.
- Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, *et al.*, "Microsoft COCO: Common Objects in Context," in *ECCV*, 2014, pp. 740–755.
- E. Hirsch and A. Tal, "CLID: Controlled-Length Image Descriptions with Limited Data," in *WACV*, 2024, pp. 5531–5541.
- C. Xiao *et al.*, "Multimodal Data Augmentation for Image Captioning using Diffusion Models," in *LGM3A workshop*, 2023, pp. 23–33.
- J. Yi, B. Uzcent, O. Ignat, Z. Li, *et al.*, "Augment the Pairs: Semantics-Preserving Image-Caption Pair Augmentation for Grounding-Based Vision and Language Models," in *WCCV*, 2024, pp. 5520–5530.
- Z. Wang, J. Li, and Y. Qiao, "Scaling Data Generation in Vision-and-Language Navigation," in *ICCV*, 2023, pp. 12009–12020.
- J. Li, M. Vo, A. Sugimoto, and H. Nakayama, "EVCap: Retrieval-Augmented Image Captioning with External Visual-Name Memory for Open-World Comprehension," in *CVPR*, 2024, pp. 13733–13742.
- S. Rennie, E. Marcheret, Y. Mroueh, *et al.*, "Self-Critical Sequence Training for Image Captioning," in *CVPR*, 2017, pp. 7008–7024.
- R. Vedantam, C. Zitnick, and D. Parikh, "CIDEr: Consensus-based Image Description Evaluation," in *CVPR*, 2015, pp. 4566–4575.
- J. Cho, S. Yoon, A. Kale, F. Deroncourt, *et al.*, "Fine-grained Image Captioning with CLIP Reward," in *NAACL*, 2022, pp. 517–527.
- A. Radford, J. Kim, *et al.*, "Learning Transferable Visual Models From Natural Language Supervision," in *ICML*, 2021, pp. 8748–8763.
- J. Yang *et al.*, "Set-of-Mark Prompting Unleashes Extraordinary Visual Grounding in GPT-4V," *arXiv preprint arXiv:2310.11441*, 2023.
- M. Oquab, T. Darcet, T. Moutakanni, H. Vo, *et al.*, "DINOv2: Learning Robust Visual Features without Supervision," *TMLR*, 2024.
- H. Liu, C. Li, Q. Wu, and J. Lee, "Visual Instruction Tuning," in *NeurIPS*, 2023, pp. 34892–34916.
- M. Goko *et al.*, "Task success prediction for open-vocabulary manipulation based on multi-level aligned representations," in *CoRL*, 2024.
- M. Reid, N. Savinov, D. Teplyashin, D. Lepikhin, T. Lillicrap, *et al.*, "Gemini 1.5: Unlocking Multimodal Understanding Across Millions of Tokens of Context," *arXiv preprint arXiv:2403.05530*, 2024.
- J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, *et al.*, "GPT-4 Technical Report," *arXiv preprint arXiv:2303.08774*, 2023.
- Y. Wada *et al.*, "Polos: Multimodal Metric Learning from Human Feedback for Image Captioning," in *CVPR*, 2024, pp. 13559–13568.
- S. Sarto *et al.*, "Positive-Augmented Contrastive Learning for Image and Video Captioning Evaluation," in *CVPR*, 2023, pp. 6914–6924.
- X. Zhou, R. Girdhar, A. Joulin, *et al.*, "Detecting Twenty-Thousand Classes using Image-Level Supervision," in *ECCV*, 2022, pp. 350–368.
- Y. Liu, M. Ott, N. Goyal, *et al.*, "RoBERTa: A Robustly Optimized Bert Pretraining Approach," *arXiv preprint arXiv:1907.11692*, 2019.
- X. Zou, J. Yang, H. Zhang, F. Li, *et al.*, "Segment Everything Everywhere All at Once," in *NeurIPS*, 2023, pp. 19769–19782.
- OpenAI, "text-embedding-3-large," <https://platform.openai.com/docs/models/embeddings>, 2024, accessed: July. 2024.
- S. Zhang, S. Roller, *et al.*, "OPT: Open Pre-Trained Transformer Language Models," *arXiv preprint arXiv:2205.01068*, 2022.
- P. Anderson, B. Fernando, *et al.*, "SPICE: Semantics Propositional Image Caption Evaluation," in *ECCV*, 2016, pp. 382–398.