

InViG: Open-Ended Interactive Visual Grounding in Human-Robot Interaction with 500K Dialogues

Hanbo Zhang^{*1}, Jie Xu^{*1}, Yuchen Mo¹, Tao Kong¹

Abstract—Ambiguity is ubiquitous in human communication. Previous approaches in Human-Robot Interaction (HRI) have often relied on predefined interaction templates, leading to reduced performance in realistic and open-ended scenarios. To address these issues, we present a large-scale dataset, InViG, for interactive visual grounding under language ambiguity. Our dataset comprises over 520K images accompanied by open-ended goal-oriented disambiguation dialogues, encompassing millions of object instances and corresponding question-answer pairs. Leveraging the InViG dataset, we conduct extensive studies and propose a set of baseline solutions for end-to-end interactive visual disambiguation and grounding, achieving a 45.6% success rate during validation. To the best of our knowledge, the InViG dataset is the first large-scale dataset for resolving open-ended interactive visual grounding, presenting a practical yet highly challenging benchmark for ambiguity-aware HRI. Codes and datasets are available at: <https://openivg.github.io>.

I. INTRODUCTION

Robots are gradually becoming more prevalent in our homes. Being an important part of our daily lives, it is crucial for robots to understand visual concepts and comprehend natural language instructions. Additionally, uncertainties and ambiguities are inherent in our daily communication. To effectively handle these uncertainties, robots must actively engage in interactions to gather more information. This process helps them develop more accurate models of the world, leading to better decision-making. Previous research has demonstrated the significance of interaction in reducing failures and improving user experience [1]–[7]. However, they often rely on predefined interaction templates, which can lead to confusing or problematic questions and hinder their ability to handle ambiguous language instructions in challenging scenarios.

Recent advancements in large-scale visual-language models (VLMs) [9]–[17] have shown promise in modeling visually-grounded HRI using end-to-end neural networks. These models, pre-trained on large amounts of image and text data, exhibit potential in understanding natural language [18]–[20], grounding symbolic concepts in visual inputs [21], [22], generating open-ended text or verbal responses [23], [24], and even reasoning from knowledge [25]. Despite these achievements, multi-turn HRI remains challenging, particularly when reasoning over historical information due to the scarcity of visually-grounded interaction data.

In this paper, we introduce the Interactive Visual Grounding dataset (InViG dataset) and benchmarks for this task. The InViG dataset is designed to train the neural networks to interact directly with users in natural language using raw



Fig. 1: InViG 500K dataset compared to previous works. (a) GuessWhat?! [8] dataset, which only contains questions with answers from “Yes”, “No”, and “N/A”; (b) INGRESS [4], which constrains the questions to follow predefined templates; (c) Our InViG 500K with free-form instructions, questions, and answers, designed to facilitate open-ended goal-oriented dialogues for human-robot disambiguation.

images and texts as inputs. For the first part, our dataset contains over 21K human-to-human visually-grounded dialogs collected based on our open-source online chat program. Furthermore, based on the InViG 21K dataset, we train models to generate 500K human-robot disambiguation dialogues automatically and leverage ChatGPT to enrich the diversity of languages. Our aim is to enable the model to learn text generation externally from ChatGPT while being trained visually. Based on the proposed benchmark, we show that our simple yet efficient baseline can archive a 35.6% success rate in the interactive visual grounding task when trained solely on InViG 21K dataset, and moreover, a 45.6% success rate after pre-training on InViG 500K dataset with the help of a state-of-the-art object detector [26].

Our contributions include three points:

- We present the first large-scale dataset including 520K images specifically designed for object-oriented open-ended interactive visual grounding and disambiguation.
- We propose a simple yet strong baseline solution to interactive visual grounding and provide a set of promis-

^{*} Equal Contribution. ¹ ByteDance Research



Fig. 2: Examples of InViG 21K dataset.

ing results.

- We conduct extensive empirical studies exploring different formulations, training strategies, and data combinations, providing comprehensive analysis to facilitate further research in interactive visual grounding for HRI.

II. RELATED WORK

a) Disambiguation in HRI: Disambiguation is a practical yet still challenging field in HRI due to the inevitable ambiguity and vagueness of languages. It has been shown that with a specific task setting and a predefined interaction corpus, verbal instruction generation is typically easier than recognition [27]. Therefore, early works usually rely on a predefined corpus to generate verbal expressions for disambiguation to, for example, navigate with the supervision of users [28], [29], collaborate with humans in complicated tasks [30], or recover from failures [1], [31]. With deep vision methods, it has been demonstrated promising to generate free-form expression with raw images as inputs [32]–[35], and hence facilitate the open-ended interaction in open-world scenarios [4]–[7]. Instead of following a fixed set of words, the agents may generate object-specific descriptions to make the interaction more natural and flexible. Yet, such methods are still limited by a set of templates, sometimes making the interaction confusing and even annoying.

b) Datasets for HRI: With progress in deep learning, end-to-end approaches have drawn great attention. As is notoriously known, data-driven approaches are always data-intensive. Therefore, some works have contributed large-scale datasets for HRI. Image Chat [36] is proposed to train agents for the emotional interaction. The robots are expected to respond differently given different emotional states. Visual Dialog [37] is designed following a similar setting of Visual Question Answering (VQA) [38], but the answer to questions



Fig. 3: Examples of InViG 500K dataset.

in it may depend on the dialog histories. GuessWhat?! [8] is inspired by the well-known interaction game that involves two persons at the same time, one for the Guesser and the other for the Oracle. The Guesser will be guessing the target of the Oracle by iteratively asking judgment questions, and the Oracle can only answer ‘Yes’ or ‘No’ to complete the game. Following these datasets, end-to-end approaches have shown great potential for language-based interaction [39]–[41]. More impressively, by training the interaction models end-to-end simply using supervised learning, it even learns to track the dialogue states during interaction [39]. These results show promise to achieve natural HRI through data. More datasets can be related to visual-language interaction like VQA [38], [42]–[46], Visual Reasoning [47], Visual Grounding [48], and Image Captioning [49]–[51]. Nevertheless, they can hardly be applied to data-driven HRI directly since they only include single-turn interactions. On the other hand, the most related work GuessWhat?! [8] simplified the problem too much by restricting the answers to be chosen from “Yes” and “No”, and hence is limited in natural interactions. In this paper, we contribute a dataset for interactive visual grounding to reduce the gap between current multi-modal models and practical application scenarios in robotics.

c) Deep Learning Models for HRI: Recently, large-scale language models (LLMs) [52]–[55] are shocking the world with their surprisingly emergent zero-shot and few-shot learning abilities [56]. By simply training the model with intensive text data, such models unite most natural language tasks by text generation. Though powerful and impressive in natural language tasks, current LLMs still lack efficiency in multi-modal scenes, especially ones that involve the alignment of visual and linguistic concepts. The most recent works have tried to either assemble or finetune

TABLE I: InViG Dataset Compared to Other Related Datasets

	#Image	#Dialogue	#Object	#Class	#Question	#Token	Object-Oriented?	Open-Ended?	Interactive?	Ambiguity?
VQA 2.0	123K	-	-	-	658K	13M				
RefCOCO	20K	-	197K	80	-	530K	✓	✓		
Image Chat	202K	202K	-	-	401K	4.7M		✓	✓	
Visual Dialog	125K	125K	-	-	1.3M	11M		✓	✓	
GuessWhat?!	67K	155K	1.3M	80	822K	5.8M	✓		✓	✓
InViG	526K	526K	43M	>500	1.9M	80M	✓	✓	✓	✓

from LLMs to keep the emergent abilities while learning to align visual and linguistic embeddings [17], [57], [58]. They, nevertheless, are still not capable enough to be deployed in specific task settings like interactive visual grounding due to the lack of data.

III. INVIG DATASET

A. Overview

We start with manually annotated interactive dialog collection using crowd-sourcing. In this phase, we sample and filter 21K images from OpenImages Dataset [59], which consists of massive images with object instances and class-level annotations. Based on the sampled images, we recruit annotators to label each image with one or more targets and human-to-human dialogues.

With 21K labeled data, we further develop an annotation system to automatically generate HRI data, based on which we further generate 500K goal-oriented disambiguation dialogues in extremely low costs. Therefore, in total, our InViG dataset contains more than 520K dialogues for interactive visual grounding.

We demonstrate the comparison between InViG Dataset and previous works in Table I. In summary, InViG dataset is proposed to solve the problem of object-oriented open-ended interactive ambiguity in HRI, which widely appears in daily communications between humans. Therefore, differentiated from all previous works, InViG dataset contains extensive interactive disambiguation data to facilitate the development of HRI systems.

B. Data Collection

We hope our dataset contains images commonly seen in daily life and vary in ambiguity. To do so, we sample images from OpenImages dataset [59] with a two-stage filter.

In the first stage, we remove those images with inappropriate contents, like dense insects and animals like *spiders*, *ladybugs*, *scorpions*, and *marine invertebrates*. Also, classes that are parts of one another, like *human body*, *human ear*, and *human teeth*, are also filtered out. Besides, we also filter out most images with low ambiguity, which is defined as the maximum number of objects belonging to the same class.

In the second stage, we propose a Bayesian Filter on the object classes to further identify if an image is suitable for our task. To be specific, we first randomly sample a small image set and manually classify these images by their suitability into a positive set I_{pos} and a negative set I_{neg} . For every image I that contains class c_1, \dots, c_n , we assume

that the appearance of each class c_i ($1 \leq i \leq n$) in each image is independent of each other. Therefore, we have

$$p(I|c_1, \dots, c_n) \propto \prod_{i=1}^n p(c_i|I) \quad (1)$$

We approximate $p(c_i|I \in I_{pos})$ and $p(c_i|I \in I_{neg})$ by simply counting the frequency that each object class appears in positive images and negative images, respectively. Then, we can apply Eq. 1 to estimate the probability of each image being positive. Finally, we sort the scores and keep only those images with a positive probability over 0.5, which means they probably portray a good InViG scenario.

Intuitively, the first stage cleans a large portion of undesired images efficiently, while the second stage refines the results by filtering out images that cannot be handled simply using heuristics.


C. Data Annotation

To collect interaction data, we developed an online interactive website that enables image-based chat between two users. To start the interaction, each user will be paired randomly with a partner. During the interaction, each user can choose to skip the current image or reset all labeled data if unexpected situations are encountered. We force the ‘‘User’’ who plays the role of the real user to start the conversation by telling the ‘‘Agent’’ what he/she is interested in. Due to ambiguity in languages, the ‘‘Agent’’ needs to ask questions to disambiguate. We force two annotators to take turn-talking. Hence, whether a conversation should stop is totally up to the ‘‘Agent’’. That is if the ‘‘Agent’’ gets enough information, and finally, locates the targets, it will stop the conversation actively. During labeling the targets, the ‘‘Agent’’ is forced to upload the results first, after which the results will be sent to the ‘‘User’’ for checking.

In total, we spent one month recruiting 230 annotators to label the disambiguation dialogues. Each annotator is required to participate in the label of at most 300 images, considering the diversity of languages with different backgrounds. To streamline the annotation process and ensure its effectiveness, we utilize specially developed image-based chat tools to collect all the labels. Finally, we collected more than 21K images and the corresponding targets with disambiguation dialogues. We have shown some examples of the dataset together with their labels in Fig. 2.

D. From 21K to 520K

Due to the time and labor-intensive nature of manual annotation, we developed an automatic annotation system



(a) Dialog rounds (b) Sentence lengths

Fig. 4: Statistics of dialogues in InViG 21K dataset.




Fig. 5: Word clouds of InViG 21K dataset.

to generate disambiguation dialogues only given unlabeled images. It enables us to efficiently collect a large-scale dataset with diverse visual and linguistic labels. By doing so, we collected 500K automatically images and generated dialogues, with diversified visual and linguistic labels.

In this phase, we choose the SAM dataset [60] as the data source considering the diversity of visual inputs. Besides, it also contains many more object segments in each image than OpenImages [59], making it too hard for human annotators yet suitable for automatic data generation and collection. The 500K images are randomly sampled from the Segment Anything dataset. However, we found that there are around 60% objects only occupying less than 5% area in the image, making it too small to be appropriate for our task. Therefore, we assign the weights for each object according to their occupied areas, meaning that larger objects will be more likely selected.

To label disambiguation dialogues automatically, we train models to interact with each other using the InViG 21K dataset together with other related public datasets [8], [37], [48], [50], [51], [61]–[64]. Our models follow the encoder-decoder transformer architectures for visual grounding and text generation [16], [65], [66]. Concretely, we first train an object captioner to generate a query expression for the selected target, which is the start of each dialogue. Based on the generated query, a well-trained questioner and oracle will speak alternatively to generate the dialogue until the questioner stops talking. Finally, a visual grounding model will output the bounding box according to the dialogue and the input image. To ensure the diversity of language, we then use ChatGPT [67] to rephrase the generated dialogues by our trained models. Concretely, we prompt ChatGPT to consider the diversity of personality, vocation, and language proficiency during the dialogue augmentation, so that the models can adapt to different kinds of users. Besides, we also set restrictions in the prompt, to make the output precise and



(a) Dialog rounds (b) Sentence lengths

Fig. 6: Statistics of dialogues in InViG 500K dataset.




Fig. 7: Word clouds of InViG 500K dataset.

concise. Some examples are shown in Fig. 3. We can see that the generated dialogues contain more diversified expressions compared to the manually labeled ones as shown in Fig. 2.

E. Data Statistics and Analysis

We have demonstrated the statistics of objects and dialogues in InViG dataset.

a) Images and Objects: As shown in Table I, InViG dataset totally contains more than 520K images with 43M segmented object instances. Each image is labeled with a multi-turn goal-oriented disambiguation dialogue. Among all the data, we have question-answer pairs, more than 4.4M sentences, and around 80M tokens. To the best of our knowledge, this is the first large-scale dataset for the task of open-ended interactive visual grounding and disambiguation.

b) Human-to-Human Dialogues: InViG includes 21K human-to-human dialogues from crowd-sourcing based on 311K object instances belonging to 510 different categories. The distribution of dialogue rounds and sentence lengths are shown in Fig. 4. We can see that more than 99% dialogues end in 6 interaction rounds, and simultaneously, more than 80% sentences include less than 10 words. We also demonstrate the word clouds of the dialogues in Fig. 5. The first column shows the overall word clouds for questions and answers, and the following ones illustrate word distributions by turns. We can see that people used to ask questions like “which one” in the first turn, and then move on to confirming questions. The end signals (e.g. “thank”) usually appear after the 3rd turn of dialogues. Moreover, interestingly, the ratio of spatial words decreases gradually with an increasing number of dialog turns, meaning that spatial relationships are preferred to answer the question “which one”. These observations accord with our experience, which might be considered in the design of interaction models.

c) Automatically Generated Dialogues: We illustrate the distribution of automatically generated dialogues in Fig. 6.

We can see that after the rephrase of ChatGPT, dialogues in InViG 500K are much more complex than human utterances. The maximum number of dialogue turns increases to more than 20 and the maximum length of each sentence is larger than 100 words. We also visualize the word clouds for InViG 500K in Fig. 7. Similar to InViG 21K, spatial relationships are used as answers mostly in the 1st turn while the confirming answers like “Yes” mostly appear after the 2nd dialogue turn.

IV. BASELINES AND EXPERIMENTAL RESULTS

Based on the collected human-to-human dialogues, we conduct a series of downstream tasks to set up a benchmark for interactive visual grounding. We aim to answer the following questions in this section:

- (1) What tasks can we do by taking advantage of InViG dataset? (Sect. IV-A)
- (2) How can InViG 21K and InViG 500K be used to perform better in disambiguation tasks? (Sect. IV-D)
- (3) What are the pros and cons of state-of-the-art visual-language models for disambiguation tasks? (Sect. IV-D)

A. Tasks

We design four downstream tasks based on our InViG dataset: a) Multi-Turn Visual Question Answering; b) Visual Question Generation; c) Multi-Turn Visual Grounding; d) Interactive Visual Grounding.

a) *Multi-turn Visual Question-Answering (M-VQA)*: requires the model to answer open-ended questions based on image observations and dialogue history, hence, the inputs are an image I , an initial referring expression e , and a group of questions and answers $\{(q_i, a_i)\}_{i=1}^{T-1}$, and the question to be answered q_T . The output is a sequence of words that form an answer a_T .

b) *Multi-turn Visual Question Generation (M-VQG)*: requires the model to generate a question for clarification in the next turn. Therefore, the input includes an image I , an initial referring expression e , and a group of questions and answers $\{(q_i, a_i)\}_{i=1}^T$. The output should be a sequence of words that form a question of the next turn.

c) *Multi-turn Visual Grounding (M-VG)*: is defined as grounding the target object $x^* = (x_{min}, y_{min}, x_{max}, y_{max})$ given an image I , an initial referring expression e , and a dialogue history $\{(q_i, a_i)\}_{i=1}^T$ that contains the complete information about the targets.

d) *Interactive Visual Grounding (I-VG)*: is a new task that involves both active interaction and interactive visual grounding. In this setting, the ground truth dialogue history is not available. Instead, the agent needs to actively interact with humans to collect the necessary information for the final grounding task. In previous works, humans are usually involved in evaluating the performance of a model [4], which is time-consuming and also expensive to evaluate a large dataset like our InViG. Instead, we introduce an Oracle by following [8] in this paper to conduct the automatic evaluation. The Oracle follows the same form as the multi-turn visual question-answering model.

Formally, the input in this task contains an image I and an initial referring expression e . The model needs to generate a set of questions $\{q_i\}_{i=1}^T$ to finally ground the target object $x^* = (x_{min}, y_{min}, x_{max}, y_{max})$. Meanwhile, the Oracle model should play the role of users, and give answers $\{a_i\}_{i=1}^T$ to the corresponding questions. High performance in this task requires seamless coordination of all the above three tasks.

B. Metrics

a) *M-VQA and M-VQG*: Since we do not restrict the question and answer space, we introduce two sets of metrics for the evaluation of these two tasks. Firstly, we follow previous works in natural language processing and measure the similarity between the predictions and the ground truths. In detail, we use BLEU-1 and BLEU-4 [68], CIDEr [69], METEOR [70], and ROUGE [71] as the metrics for text similarity. Nevertheless, these metrics assume that the generated text should be unique, which violates the intuition of open-ended dialogues. Hence, we also introduce the retrieval-based multi-choice metrics: *Recall@k* and *GT Rank*. To do so, we assign 30 candidates for each question and answer. To make the choices challenging, the candidates are sampled from the dialogues of the top 30 similar images using CLIP [15] embeddings. The model is required to rank all candidates together with the ground truth. *Recall@k* measures the fraction of the ground truth being in top- k candidates. *GT Rank* measures the mean rank of the ground truth among all candidates.

b) *M-VG and I-VG*: We follow the traditional settings [48] and use accuracy to evaluate the performance of Interactive Visual Grounding. Concretely, for each model, we measure the fraction of predicted bounding boxes that have the Intersection of Union (IoU) no smaller than m , given the dialog histories, where $m \in \{0.1, 0.2, \dots, 1.0\}$.

C. Baselines

We implement a simple yet efficient baseline algorithm named *InViG-Base*, based on the state-of-the-art vision-language foundation model X-VLM [65]. It includes three parts: an *Oracle* model which can answer open-ended questions, a *Questioner* model that interacts with the Oracle to collect information, and a *Guesser* model to guess targets with the input of ground truth bounding boxes of all objects or the ones detected using Detic [26]. The overall architecture is shown in Fig. 8. To validate different formulations for multi-turn dialogues, we also implement a multi-turn variant called *InViG-Base-MT*, following previous work [9]. It reformulates the Agent in a multi-turn mode. In each turn, the model takes as inputs the belief of each object from the last turn as well as the current textual and visual embeddings. The third one is similar to *InViG-Base*, which is also demonstrated in Fig. 8. In this version, the visual grounding head is replaced with a bounding box regressor rather than a classifier following the formulation from [18], [65]. It gets rid of object proposals and directly outputs the location of

TABLE II: M-VQA Performance

	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE	CIDEr	R@1	R@5	Rank
XVLM 21K	24.6	15.8	10.4	7.2	16.0	28.3	72.8	49.8	73.9	5.0
<i>InViG-Base</i> 21K	26.7	17.8	12.1	8.6	16.9	30.2	81.9	50.9	75.0	4.8
<i>InViG-Base</i> 500K	23.2	14.5	9.5	6.7	15.4	22.8	52.3	39.7	61.5	7.1
<i>InViG-Base</i> 500+21K	27.7	19.3	14.0	10.5	17.8	33.4	96.2	49.9	74.3	4.8

TABLE III: M-VQG Performance

	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE	CIDEr	R@1	R@5	Rank
<i>InViG-Base-MT</i> 21K	31.0	24.9	21.1	18.3	16.2	48.2	117.4	38.0	63.6	5.8
<i>InViG-Base-BBOX</i> 21K	34.2	27.1	22.8	19.5	17.3	49.2	120.5	49.9	68.2	5.4
<i>InViG-Base</i> 21K	34.7	27.6	23.1	19.9	17.4	49.7	122.2	49.9	67.9	5.4
<i>InViG-Base</i> 500K	24.0	15.8	10.9	7.9	13.9	38.7	36.8	33.1	58.1	6.9
<i>InViG-Base</i> 500+21K	34.2	27.7	23.5	20.4	17.6	50.3	128.2	52.6	74.7	4.5

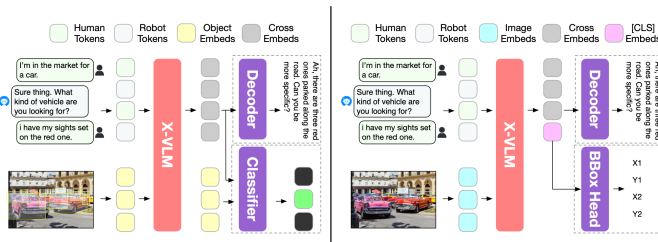


Fig. 8: Networks of baseline methods. Left: *InViG-Base*. Right: *InViG-Base* BBOX. Note that *InViG-Base* BBOX is conditioned on image embeddings and does not rely on an external bounding box detector.

the target. For M-VQA, we also compare to the original X-VLM with only the question to be asked as language inputs, which is trained on our *InViG* 21K dataset (XVLM 21K).

D. Results

We do experiments for all four tasks introduced in Sect. IV-A on *InViG* 21K and 500K datasets.

a) *Pre-training improves performance*: For the text generation tasks, specifically M-VQA and M-VQG (see Table II and Table III), training on a large amount of noisy data (*InViG-Base* 500K) yields inferior performance compared to training on a smaller dataset with less clean data (*InViG-Base* 21K). However, pre-training on a large amount of noisy data still leads to a significant improvement in performance. Notably, the pre-trained model *InViG-Base* 500+21K outperforms the model without pre-training across most metrics. For M-VG and I-VG (Fig. 9a), we also have similar conclusions.

b) *Dialog history is crucial*: In the case of M-VG, as shown in Fig. 9b, *Turn-m*, $m \in \{1, 2, 3\}$, means grounding performance without the last m dialogue turns. We can see that the grounding performance drops consistently with less information. Besides, dialog history also helps M-VQA. As shown in Table II, in the case of M-VQA, our *InViG-Base* consistently outperforms the original X-VLM, albeit by a small margin. This observation highlights two key findings: 1) Dialog histories contribute to enhanced performance in M-VQA, and 2) Most questions can be reasonably answered without relying on dialog history.




Fig. 9: Multi-Turn Visual Grounding (M-VG) and Interactive Visual Grounding (I-VG) performance.

c) *Target classification performs stably*: From Fig. 9a, we can conclude that when considering different visual grounding approaches, direct bounding box prediction demonstrates superior performance at lower IoU thresholds, but its performance declines significantly as the threshold increases. In contrast, the performance of *InViG-Base* based on the target classification is more stable. We can see that with ground truth object bounding boxes, *InViG-Base* 500+21K still achieves more than 60% success rate with ground truth dialogues. When paired with Detic, its performance is stable when the threshold increases. Besides, *InViG-Base-MT* exhibits the poorest performance, suggesting that the multi-turn dialogue formulation is not optimal for interactive visual grounding tasks.

V. CONCLUSIONS

In this paper, we present *InViG* dataset, the first large-scale dataset including more than 520K images and dialogues for interactive visual grounding, to resolve the challenge of language ambiguity in HRI. We have conducted extensive and comprehensive experiments and set up a suite of baseline solutions to resolve HRI ambiguity. Our results demonstrate that based on the *InViG* dataset, the robot can successfully disambiguate interaction with a success rate of 45.6% in challenging and realistic scenarios in validation. Future works include developing an automatic data-cleaning process for high-quality data collection, integration into more interactive robot systems, and validation of performance on more downstream interactive robot tasks.

REFERENCES

- [1] Stefanie Tellex, Ross Knepper, Adrian Li, Daniela Rus, and Nicholas Roy. Asking for help using inverse semantics. 2014.
- [2] David Whitney, Eric Rosen, James MacGlashan, Lawson LS Wong, and Stefanie Tellex. Reducing errors in object-fetching interactions through social feedback. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1006–1013. IEEE, 2017.
- [3] Jun Hatori, Yuta Kikuchi, Sosuke Kobayashi, Kuniyuki Takahashi, Yuta Tsuboi, Yuya Unno, Wilson Ko, and Jethro Tan. Interactively picking real-world objects with unconstrained spoken language instructions. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3774–3781. IEEE, 2018.
- [4] Mohit Shridhar, Dixant Mittal, and David Hsu. Ingress: Interactive visual grounding of referring expressions. *The International Journal of Robotics Research*, 39(2-3):217–232, 2020.
- [5] Hanbo Zhang, Yunfan Lu, Cunjun Yu, David Hsu, Xuguang Lan, and Nanning Zheng. INVIGORATE: Interactive Visual Grounding and Grasping in Clutter. In *Proceedings of Robotics: Science and Systems, Virtual*, July 2021.
- [6] Yang Yang, Xibai Lou, and Changhyun Choi. Interactive robotic grasping with attribute-guided disambiguation. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 8914–8920. IEEE, 2022.
- [7] Yuchen Mo, Hanbo Zhang, and Tao Kong. Towards open-world interactive disambiguation for robotic grasping. In *2023 International Conference on Robotics and Automation (ICRA)*, 2023.
- [8] Harm De Vries, Florian Strub, Sarath Chandar, Olivier Pietquin, Hugo Larochelle, and Aaron Courville. Guesswhat?! visual object discovery through multi-modal dialogue. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5503–5512, 2017.
- [9] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *arXiv preprint arXiv:1908.02265*, 2019.
- [10] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. In *International Conference on Learning Representations*, 2019.
- [11] Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 12-in-1: Multi-task vision and language representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10437–10446, 2020.
- [12] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *European Conference on Computer Vision*, pages 104–120. Springer, 2020.
- [13] Pengchuan Zhang, Xiujuan Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5579–5588, 2021.
- [14] Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. Unifying vision-and-language tasks via text generation. In *International Conference on Machine Learning*, pages 1931–1942. PMLR, 2021.
- [15] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [16] Yan Zeng, Xinsong Zhang, and Hang Li. Multi-grained vision language pre-training: Aligning texts with visual concepts. In *Proceedings of the 39th international conference on Machine learning*, 2022.
- [17] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. In *Advances in Neural Information Processing Systems*, 2022.
- [18] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr-modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1780–1790, 2021.
- [19] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10965–10975, 2022.
- [20] Yuying Ge, Yixiao Ge, Xihui Liu, Jinpeng Wang, Jianping Wu, Ying Shan, Xiaohu Qie, and Ping Luo. Miles: visual bert pre-training with injected language semantics for video-text retrieval. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXV*, pages 691–708. Springer, 2022.
- [21] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Cliport: What and where pathways for robotic manipulation. In *Conference on Robot Learning*, pages 894–906. PMLR, 2022.
- [22] Sanjay Subramanian, Will Merrill, Trevor Darrell, Matt Gardner, Sameer Singh, and Anna Rohrbach. Reclip: A strong zero-shot baseline for referring expression comprehension. *arXiv preprint arXiv:2204.05991*, 2022.
- [23] Xiaowei Hu, Zhe Gan, Jianfeng Wang, Zhengyuan Yang, Zicheng Liu, Yumao Lu, and Lijuan Wang. Scaling up vision-language pre-training for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17980–17989, 2022.
- [24] Jun Chen, Han Guo, Kai Yi, Boyang Li, and Mohamed Elhoseiny. Visualgpt: Data-efficient adaptation of pretrained language models for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18030–18040, 2022.
- [25] Fangyu Liu, Francesco Piccinno, Syrine Krichene, Chenxi Pang, Kenton Lee, Mandar Joshi, Yasemin Altun, Nigel Collier, and Julian Martin Eisenschlos. Matcha: Enhancing visual language pre-training with math reasoning and chart derendering. *arXiv preprint arXiv:2212.09662*, 2022.
- [26] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Philipp Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. In *European Conference on Computer Vision*, pages 350–368. Springer, 2022.
- [27] Andrea Thomaz, Guy Hoffman, and Maya Cakmak. Computational human-robot interaction. *Foundations and Trends in Robotics*, 4(2-3):105–223, 2016.
- [28] Geert-Jan M Kruijff, Hendrik Zender, Patric Jensfelt, and Henrik I Christensen. Clarification dialogues in human-augmented mapping. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*, pages 282–289, 2006.
- [29] Sachithra Hemachandra and Matthew R Walter. Information-theoretic dialog to improve spatial-semantic representations. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5115–5121. IEEE, 2015.
- [30] Stephanie Rosenthal, Joydeep Biswas, and Manuela M Veloso. An effective personal mobile robot agent through symbiotic human-robot interaction. In *AAMAS*, volume 10, pages 915–922, 2010.
- [31] Robin Deits, Stefanie Tellex, Pratiksha Thaker, Dimitar Simeonov, Thomas Kollar, and Nicholas Roy. Clarifying commands with information-theoretic human-robot dialog. *Journal of Human-Robot Interaction*, 2(2):58–79, 2013.
- [32] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Densecap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4565–4574, 2016.
- [33] Linjie Yang, Kevin Tang, Jianchao Yang, and Li-Jia Li. Dense captioning with joint inference and visual context. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2193–2202, 2017.
- [34] Jingyu Liu, Liang Wang, and Ming-Hsuan Yang. Referring expression generation and comprehension via attributes. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4856–4864, 2017.
- [35] Duo Zheng, Tao Kong, Ya Jing, Jiaan Wang, and Xiaojie Wang. Towards unifying reference expression generation and comprehension. *arXiv preprint arXiv:2210.13076*, 2022.
- [36] Kurt Shuster, Samuel Humeau, Antoine Bordes, and Jason Weston. Image chat: Engaging grounded conversations. *arXiv preprint arXiv:1811.00945*, 2018.
- [37] Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. Visual dialog. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 326–335, 2017.
- [38] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual ques-

- tion answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015.
- [39] Wei Pang and Xiaojie Wang. Visual dialogue state tracking for question generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11831–11838, 2020.
- [40] Tao Tu, Qing Ping, Govindarajan Thattai, Gokhan Tur, and Prem Natarajan. Learning better visual dialog agents with pretrained visual-linguistic representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5622–5631, 2021.
- [41] Shoya Matsumori, Kosuke Shingyouchi, Yuki Abe, Yosuke Fukuchi, Komei Sugiura, and Michita Imai. Unified questioner transformer for descriptive question generation in goal-oriented visual dialogue. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1898–1907, 2021.
- [42] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709, 2019.
- [43] Linjie Li, Jie Lei, Zhe Gan, and Jingjing Liu. Adversarial vqa: A new benchmark for evaluating the robustness of vqa models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2042–2051, 2021.
- [44] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017.
- [45] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, pages 3195–3204, 2019.
- [46] Tejas Gokhale, Pratyay Banerjee, Chitta Baral, and Yezhou Yang. Vqalol: Visual question answering under the lens of logic. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 379–396. Springer, 2020.
- [47] Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. A corpus for reasoning about natural language grounded in photographs. *arXiv preprint arXiv:1811.00491*, 2018.
- [48] Licheng Yu, Patrick Poisson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *European Conference on Computer Vision*, pages 69–85. Springer, 2016.
- [49] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017.
- [50] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015.
- [51] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018.
- [52] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [53] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.
- [54] Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*, 2022.
- [55] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
- [56] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022.
- [57] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023.
- [58] Shaohan Huang, Li Dong, Wenhui Wang, Yaru Hao, Saksham Singhal, Shuming Ma, Tengchao Lv, Lei Cui, Owais Khan Mohammed, Qiang Liu, et al. Language is not all you need: Aligning perception with language models. *arXiv preprint arXiv:2302.14045*, 2023.
- [59] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision*, 128(7):1956–1981, 2020.
- [60] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023.
- [61] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [62] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 787–798, 2014.
- [63] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023.
- [64] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023.
- [65] Yan Zeng, Xinsong Zhang, and Hang Li. Multi-grained vision language pre-training: Aligning texts with visual concepts. *arXiv preprint arXiv:2111.08276*, 2021.
- [66] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *International Conference on Machine Learning*, pages 23318–23340. PMLR, 2022.
- [67] OpenAI. Gpt-4 technical report. *arXiv*, pages 2303–08774, 2023.
- [68] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002.
- [69] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015.
- [70] Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005.
- [71] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.