

PROGrasp: Pragmatic Human-Robot Communication for Object Grasping

Gi-Cheon Kang^{1,3} Junghyun Kim^{1,3} Jaein Kim^{2,3} Byoung-Tak Zhang^{1,2,3}

Abstract—Interactive Object Grasping (IOG) is the task of identifying and grasping the desired object via human-robot natural language interaction. Current IOG systems assume that a human user initially specifies the target object’s category (e.g., bottle). Inspired by *pragmatics*, where humans often convey their intentions by relying on context to achieve goals, we introduce a new IOG task, Pragmatic-IOG, and the corresponding dataset, Intention-oriented Multi-modal Dialogue (IM-Dial). In our proposed task scenario, an intention-oriented utterance (e.g., “I am thirsty”) is initially given to the robot. The robot should then identify the target object by interacting with a human user. Based on the task setup, we propose a new robotic system that can interpret the user’s intention and pick up the target object, Pragmatic Object Grasping (PROGrasp). PROGrasp performs Pragmatic-IOG by incorporating modules for visual grounding, question asking, object grasping, and most importantly, answer interpretation for pragmatic inference. Experimental results show that PROGrasp is effective in offline (i.e., target object discovery) and online (i.e., IOG with a physical robot arm) settings. Code and data are available at <https://github.com/gicheonkang/prograsp>.

I. INTRODUCTION

Recent advances in robotics and artificial intelligence (AI) have made intelligent robots ubiquitous in our daily lives. To get closer to non-expert users, robots should communicate with humans using natural language and make decisions based on the interaction. Notably, in the field of human-robot interaction, there have been extensive studies [1]–[5] on developing such robots under the umbrella of Interactive Object Grasping (IOG). A typical scenario of IOG starts mentioning the target object, such as “Give me the plastic bottle”, but there is more than one object in the scene that meets the instruction. The robot should disambiguate the target object by asking questions to the conversational partner and then perform object grasping.

While the progress of IOG is exciting, the current scenario limits the ability for robots to understand beyond the literal meaning of natural language instructions. Specifically, instructions in the existing scenario clearly specify the category of the target object (e.g., bottle). In other words, current IOG systems may work properly when the target object’s category is given. However, we humans often convey our *intended meanings* by relying on context to achieve communicative goals [6], [7]. For example, when we need some water and




Fig. 1. Overview of interactive object grasping with intention-oriented utterance. The initial utterance does not contain the target object’s category.

want our conversational partner to bring it, we briefly say “I am thirsty.” The partner then enriches the literal meaning of the utterance based on various types of shared context (e.g., visual information and dialogue context) and reason about context-appropriate behavior. This ability to interpret and use language in context to achieve goals is known as pragmatics [8]–[10].

We argue that the next-generation robotic system should have pragmatic reasoning ability – capture the user’s intention with contextual information and achieve the desired goal. Therefore, we introduce a new task, Pragmatic-IOG, to study pragmatic reasoning behavior in IOG. As shown in Figure 1, we consider a scenario where a human user begins a conversation with an *intention-oriented* utterance like “My device runs out of battery.” The robot should then find all valid object candidates (e.g., the red-colored object regions in Figure 1) via visual grounding [11] and ask a question for disambiguation. After receiving the user’s response, the robot pinpoints the target object and grasps the desired object. To study the problem, we propose a new dataset, called Intention-oriented Multi-modal Dialogue (IM-Dial). The IM-Dial dataset contains 800 images and 500 human-to-human dialogue data regarding 86 categories of everyday objects. The dialogue consists of the intention-oriented utterance and a series of question-and-answer pairs.

We propose a new robotic system that can reason about intention-oriented natural language utterances of the user and grasp desired objects through human-robot interaction, called **PR**agmatic **O**bject **G**rasping (PROGrasp). PROGrasp consists of four modules: (1) a visual grounding module (VG) that predicts the region coordinates of valid objects, (2) a question generation module (Q-gen) that learns to generate questions to identify the user’s intention, (3) an answer interpretation module (A-int) that interprets the human user’s response given the multi-modal context, and (4) an object grasping module (OG) to pick up the inferred object.

¹Interdisciplinary Program in AI, Seoul National University

²Interdisciplinary Program in Neuroscience, Seoul National University

³AI Institute, Seoul National University

This research was supported by the Institute of Information and Communications Technology Planning and Evaluation (IITP) (2021-0-01343-GSAI/40%, 2022-0-00953-PICA/30%, 2022-0-00951-LBA/10%, 2021-0-02068-AIHub/10%) and National Research Foundation of Korea (NRF) (RS-2023-00274280/10%) grant funded by the Korean government.

PROGrasp trains VG, Q-gen, and A-int on the IM-Dial data. After training, PROGrasp performs our proposed task by interacting with the human user. Specifically, VG first predicts a set of object region candidates given the intention-oriented utterance. Q-gen then generates a question, and the user responds to the question. Next, PROGrasp determines the target region among the region candidates based on how well each candidate region explains the visual and dialogue context, which we call *pragmatic inference*. We implement pragmatic inference as a multi-agent reasoning of VG and A-int where VG evaluates the likelihood of each region candidate given the visual and dialogue context, and A-int rescores alternative region candidates by interpreting the user’s response. Finally, OG computes the 3D coordinates of the inferred region and performs object grasping.

We conduct offline and online experiments on the IM-Dial dataset. In offline experiments, we study how well PROGrasp identifies the target object. PROGrasp significantly improves the accuracy of offline experiments by 35% compared with the baselines. Moreover, PROGrasp outperforms the powerful multimodal foundation model [12] on validation data. In online experiments, we use a physical robot arm to evaluate the success rate of object grasping. PROGrasp boosts the success rate by 17%. Furthermore, our system efficiently identifies the target object through fewer interactions than baselines. Finally, we perform qualitative analysis, visualizing diverse samples inferred by PROGrasp.

Our contributions are three-fold. First, we propose an interactive object-grasping system (*i.e.*, PROGrasp) that capably understands the human user’s intention and grasps the desired object through dialogue. Second, we introduce a new task Pragmatic-IOG with a novel dataset, Intention-oriented Multi-modal Dialogue (IM-Dial). Third, through extensive experiments, our robotic system validates its (1) efficacy in both offline and online experiments and (2) efficiency when identifying the target object via pragmatic inference.

II. RELATED WORK

Language-Guided Object Grasping. There has been extensive research on developing object-grasping systems that can understand natural language. Some studies [13]–[17] make robots manipulate objects only with initial language instruction, assuming that the instruction is enough to identify the desired object. However, natural language is inherently ambiguous [18]. Therefore, a line of research [1]–[5], which we call Interactive Object Grasping (IOG), considers the scenario where robots need more information to disambiguate the target object. They typically generate questions and perform object grasping based on the response from a human user. Our approach belongs to IOG, but it differs from previous studies in two aspects. First, prior work in IOG considers a scenario where the category of the target object is clearly specified. However, we design a task scenario that requires robots to understand the semantic meaning and, by extension, *the intended meaning* of the user’s utterance. Accordingly, intention-oriented utterances focus on the user’s

intention without specifying the category of the target object. Second, previous studies define the format of either questions or responses. Specifically, the questions are fixed (*e.g.*, “Which one?”) [2], [19] or based on templates [1], [3]–[5]. The answer formats are also binary [2], a single word [19], or based on a pre-defined pool [1], [3]–[5]. However, PROGrasp does not impose any constraints on the format of the questions and responses. Our Q-gen generates unconstrained questions without relying on any templates, and A-int understands various types of responses, enabling non-expert users to interact with the robot more naturally.

Pragmatics. There is a long history of research studying how linguistic meaning is affected by context [6], [7], [20] under the name of pragmatics. According to the work [8], there are four kinds of well-studied tasks in the field of pragmatics: reference games [7], [21], image captioning [22], [23], instruction following [24], [25], and grounded dialogue [26]–[29]. Our work belongs to the last category, but it is the first work that integrates grounded goal-oriented dialogue into a real-world robot arm for object grasping. Regarding computational modeling, PROGrasp shares the spirit with the Rational Speech Acts (RSA) [7], [9]. We propose a multi-agent reasoning method (*i.e.*, pragmatic inference) to identify target objects accurately and efficiently by interpreting the human user’s response.

III. METHOD

A. Background

As the Web-scale data sources [30], [31] are publicly available, finetuning the model pre-trained on such datasets to the specific task has become a de facto standard strategy in AI. Accordingly, there has been a lot of multi-modal pre-training methods [32]–[34] trained on the large-scale image-text pairs. We employ a simple yet powerful multi-modal sequence-to-sequence model, OFA [32], since it can cover various multi-modal tasks with a unified architecture. OFA is pre-trained on a wide range of multi-modal and uni-modal datasets with sequence-to-sequence learning [35], [36]. Specifically, the learning objective of OFA is to optimize $\max \sum_{i=1}^{|y|} \log P_{\theta}(y_i | y_{<i}, v, x)$, where x and y denote the input and target sequences, respectively. $y_{<i}$ denotes all tokens before the i -th token in the target sequence. θ is the model parameters, and v is visual information. The encoder encodes x and v and conveys the hidden states to the decoder. The decoder predicts the next token y_i given a set of preceding tokens $y_{<i}$ and the hidden states of the encoder. We train our proposed modules in PROGrasp (*i.e.*, VG, Q-gen, and A-int) by finetuning OFA on the IM-Dial dataset. More details can be found in the following Section.

B. Problem Statement

The goal of Pragmatic-IOG is to discover the target object’s region coordinates $r^* = \langle x_1, y_1, x_2, y_2 \rangle$ through human-robot interaction and pick up the object. We assume that a human user initially provides intention-oriented natural language utterance ℓ , and there is one target object in the given 2D image I . Our system asks natural language question q

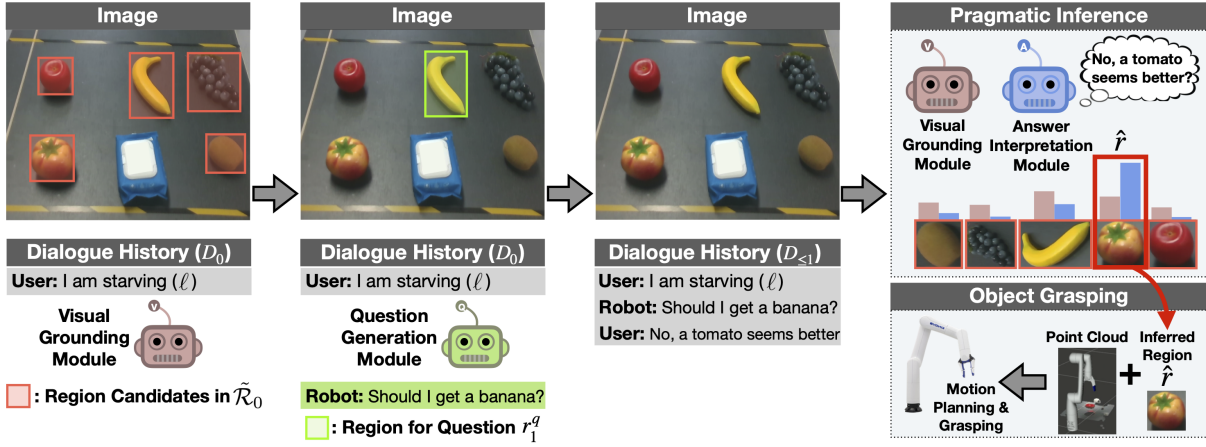


Fig. 2. **Illustration of the inference step in PROGrasp when $T = 1$.** VG first performs object grounding using the dialogue history. Q-gen then selects the object candidate to ask and generates a question. After obtaining the response from the user, VG cooperates with A-int to determine the target object region. The object grasping module finally grasps the object by computing the 3D coordinates of the target object.

to the user to identify the target object, and the user gives an answer a . The dialogue history \mathcal{D} is initialized to ℓ , and question and answer pairs at each round are appended to the dialogue history. Our robotic system predicts the region coordinates \hat{r} after T rounds of dialogue. Finally, it performs object grasping, computing the 3D coordinates of the target object based on \hat{r} and the point cloud.

C. Visual Grounding Module

The visual grounding module (VG) aims to localize objects based on the multi-modal context. Specifically, the Intention-oriented Multi-modal Dialogue (IM-Dial) dataset contains an image \mathcal{I} , a visually-grounded dialogue $\mathcal{D} = \{ \underbrace{\ell}_{\mathcal{D}_0}, \underbrace{(q_1, a_1)}_{\mathcal{D}_1}, \dots, \underbrace{(q_N, a_N)}_{\mathcal{D}_N} \}$, and region labels $\mathcal{R} = \{ \mathcal{R}_0, \dots, \mathcal{R}_N \}$ where each element $\mathcal{R}_n = \{ r_i^{vg} \}_{i=1}^{|\mathcal{R}_n|}$ is a set of object regions that VG should predict given the dialogue history $\mathcal{D}_{\leq n}$. $\mathcal{D}_{\leq n}$ denotes all dialogue data before or equal to the n -th round. Consequently, VG is trained to maximize the log-likelihood of the ground-truth regions. Formally, $\max \sum_{n=0}^N \log P_{\mathcal{V}}(\mathcal{R}_n | \mathcal{I}, \mathcal{D}_{\leq n})$. PROGrasp regards the concatenation of the region coordinates in \mathcal{R}_n as the target sequence in sequence-to-sequence learning and trains VG on top of the pre-trained OFA model [32]. After training, as shown in Figure 2, VG can predict $\tilde{\mathcal{R}}_0$ (i.e., the red-colored object regions) given \mathcal{D}_0 (i.e., “I am starving.”) and an image \mathcal{I} .

D. Question Generation & Answer Interpretation Modules

PROGrasp trains the question generation (Q-gen) and answer interpretation (A-int) modules to produce the utterances from the human questioner and the human answerer, respectively. First, Q-gen learns to generate human-annotated questions in the IM-Dial dataset, given an input image, the dialogue history, and an object region for questioning as follows: $\max \sum_{n=1}^N \log P_Q(q_n | \mathcal{I}, \mathcal{D}_{\leq n-1}, r_n^q)$. The object region r_n^q is annotated by the human questioner when collecting the IM-Dial dataset. Regarding sequence-to-sequence learning, the target sequence is the question q_n , and the triplet $(\mathcal{I}, \mathcal{D}_{\leq n-1}, r_n^q)$ is fed into the encoder. Next, the answer interpretation module (A-int), which is a proxy for the human

user, learns to generate the response of the human answerer. It is optimized by maximizing the log-likelihood of the ground-truth answer: $\max \sum_{n=1}^N \log P_{\mathcal{A}}(a_n | \mathcal{I}, r^*, q_n)$. Note that A-int takes the ground-truth region coordinates r^* during training and implicitly learns the semantic alignment between r^* and the question-and-answer pair (q_n, a_n) . Moreover, we assume that the answer distribution is independent of the dialogue history (i.e., $P_{\mathcal{A}}(a_n | \mathcal{I}, \mathcal{D}_{\leq n-1}, r^*, q_n) = P_{\mathcal{A}}(a_n | \mathcal{I}, r^*, q_n)$) by following the work [37]. We also implement both Q-gen and A-int by finetuning the pre-trained OFA model [32].

E. Inference Step

The inference step of PROGrasp is described in Algorithm 1 and Figure 2. PROGrasp obtains an image \mathcal{I} from a camera. A human user provides intention-oriented utterance ℓ . PROGrasp proceeds T rounds of dialogue in the inference step, and T is a hyperparameter. VG (i.e., $P_{\mathcal{V}}$) first predicts a set of object regions $\tilde{\mathcal{R}}$ and saves it to the superset \mathcal{R} . PROGrasp accumulates the object region candidates predicted in each round of dialogue to maximize the probability of having the target object’s region coordinates in the set \mathcal{R} . PROGrasp then samples an object region r_i^q from \mathcal{R} . The image, the dialogue history, and the sampled object region are fed into Q-gen (i.e., P_Q) to produce a question (e.g., “Should I get a banana?” in Figure 2). Next, the human user answers the question by checking whether the object mentioned in the question corresponds to the target object. After receiving the user’s response, PROGrasp saves the question-and-answer pairs to the dialogue history and evaluates each object region candidate r in the set \mathcal{R} . In the evaluation, VG computes the likelihood of each region candidate given the multi-modal context, and A-int rescores each candidate r whether it describes the question and the user’s response $(\tilde{q}_r, \tilde{a}_r)$. In other words, VG cooperates with A-int to determine the best region \hat{r} for the target based on the extent to which each candidate explains the visual and dialogue context. We call it *pragmatic inference* (see Figure 2). In line 7 at Algorithm 1, λ is a rationality parameter [21], [38], [39] in the range $[0, 1]$ that indicates the relative importance of the evaluation from A-int in pragmatic inference.

Algorithm 1 Pragmatic Object Grasping

Require: Modules for VG (P_V), Q-gen (P_Q), A-int (P_A)

Require: Module for object grasping (O)

Require: 2D RGB image I and dialogue history $\mathcal{D} \leftarrow \{\ell\}$

Require: A human user to interact with the robotic system

Require: The empty set of object regions $\mathcal{R} = \emptyset$

- 1: **for** $t \leftarrow 1$ to T **do**
 - 2: $\tilde{\mathcal{R}}_{t-1} \leftarrow P_V(\cdot|I, \mathcal{D}_{\leq t-1})$ where $\tilde{\mathcal{R}}_{t-1} = \{r_1, \dots, r_{|\tilde{\mathcal{R}}|}\}$
 - 3: $\mathcal{R} \leftarrow \mathcal{R} \cup \tilde{\mathcal{R}}_{t-1}$
 - 4: $\tilde{q}_t \leftarrow P_Q(\cdot|I, \mathcal{D}_{\leq t-1}, r_t^q)$ where $r_t^q \sim \mathcal{R}$
 - 5: The user provides an answer \tilde{a}_t to the question \tilde{q}_t
 - 6: $\mathcal{D} \leftarrow \mathcal{D} \cup \{\tilde{q}_t, \tilde{a}_t\}$
 - 7: $\hat{r} \leftarrow \operatorname{argmax}_{r \in \mathcal{R}} P_A(\tilde{a}_t|I, r, \tilde{q}_t)^\lambda \cdot P_V(r|I, \mathcal{D}_{\leq t})^{1-\lambda}$
 - 8: **end for**
 - 9: Grasp the object $O(\hat{r}, \mathcal{P})$ with \hat{r} and the point cloud \mathcal{P}
-

F. Object Grasping Module

The object grasping module (OG) first computes the 3D coordinates of the predicted object using the 2D region coordinates \hat{r} and the point cloud \mathcal{P} . Specifically, OG matches the 2D object region with the point cloud on the identical resolution and then segments points inside the region. We employ the RANSAC [40] to remove the table plane from the segmented points. The 3D target coordinates are computed by averaging the segmented points. Finally, OG computes the motion planning [41] and performs object grasping.

IV. EXPERIMENTS

A. Experimental Setup

Intention-oriented Multi-modal Dialogue Dataset. We evaluate our proposed method on the IM-Dial dataset, collected by the chatting between two players about images. The IM-Dial dataset consists of 800 images and 500 human-to-human dialogue data that cover 86 categories of everyday objects as shown in Figure 3. We divide the IM-Dial dataset into five splits: train, validation, test-seen, test-unseen, and test-cluttered. The train split contains 400 images and corresponding dialogue data for training. The validation split has 100 image and dialogue pairs. The test-seen, test-unseen, and test-cluttered data contain 100 pairs of images and intention-oriented utterances each. Note that these test splits do not have question-and-answer data, so the robotic system should identify the target object, interacting with a human user. The test-unseen split includes *never-seen-before* objects not observed in the training procedure. The goal of the test-unseen split is to evaluate the generalization ability of the robotic system. Furthermore, we define the test-cluttered as a cluttered version of the test-seen split where objects are arbitrarily placed (e.g., (a) in Figure 6).

Robotic Platform. We conduct online experiments using a physical robot arm, the 6-DoF Kinova Gen3 lite with a two-fingered gripper. Our system utilizes Intel Realsense Depth Camera D435 to get an RGB-D image. The remote server processes our proposed algorithm and communicates with the robotic platform. The robotic platform locally computes motion trajectory planning.



Fig. 3. The 86 categories of everyday objects used in the experiments.

Baselines. We compare PROGrasp with four methods:

- **Zero-Shot:** The zero-shot approach is a visual grounding model not trained on the IM-Dial dataset. We implement it by finetuning the pre-trained OFA model on the visual grounding dataset, RefCOCO [11]. This approach predicts the object region given an input image and intention-oriented utterance (*i.e.*, $\hat{r} = \operatorname{argmax}_{r \in \mathcal{R}} P_V(r|I, \mathcal{D} = \ell)$).
- **SilentGrasp:** SilentGrasp is an ablative model of PROGrasp that does not have the question generation (Q-gen) and answer interpretation (A-int) modules. It predicts the object region in the same way as Zero-Shot, but the visual grounding module (VG) is trained on the IM-Dial dataset.
- **LiteralGrasp:** LiteralGrasp is another ablative method of PROGrasp that does not have A-int. It is equivalent to PROGrasp without pragmatic inference (*i.e.*, $\lambda = 0$).
- **A-int-only:** This method is equivalent to PROGrasp that does not utilize VG in pragmatic inference (*i.e.*, $\lambda = 1$).

B. Offline Experiments

Evaluation Protocol. The offline experiment aims to verify how well the robotic system discovers the target object through human-robot natural language interaction. We measure the Intersection over Union (IoU) between the target object region r^* and the predicted region after the interaction \hat{r} . The IoU is defined as the overlapping region between the two divided by their union region. The percentage of examples with an IoU value greater than 0.5 is typically reported as Acc@0.5. However, the threshold value of 0.5 may not be a reliable indicator to estimate the success of object grasping since object grasping requires accurate prediction of the target coordinates. We thus additionally report Acc@0.9, which requires more tight alignment between r^* and \hat{r} .

Results on the Validation Split. We first evaluate PROGrasp and the compared methods on the validation split. As shown in Table 1, PROGrasp outperforms all compared methods on all evaluation metrics. It indicates our proposed components (*i.e.*, Q-gen, A-int, and pragmatic inference) play a crucial role in boosting performance. Moreover, comparing Zero-Shot and SilentGrasp, even the strong pre-trained model performs poorly without adapting to Pragmatic-IOG.

We further study a new setting for Zero-Shot and SilentGrasp, called Grounding from Dialog History (GDH). We naturally assume Zero-Shot and SilentGrasp can only access the intention-oriented utterance (*i.e.*, ℓ) in the inference phase since neither approach has a module for question generation. However, the initial utterance is insufficient to pinpoint the target object. In GDH, we assume that Zero-Shot and Silent-

TABLE I

Results on the offline experiments. Underlined scores indicate the performance of the runner-up method.

Method	GDH	Validation		Test-Seen		Test-Unseen		Test-Cluttered	
		Acc@0.5	Acc@0.9	Acc@0.5	Acc@0.9	Acc@0.5	Acc@0.9	Acc@0.5	Acc@0.9
Zero-Shot [32]		14%	4%	14%	7%	3%	2%	6%	4%
SilentGrasp		50%	44%	54%	45%	45%	31%	41%	22%
A-int-only		82%	<u>75%</u>	81%	66%	<u>78%</u>	<u>57%</u>	83%	40%
LiteralGrasp		84%	<u>75%</u>	<u>85%</u>	<u>74%</u>	<u>73%</u>	<u>54%</u>	<u>84%</u>	<u>41%</u>
Zero-Shot [32]	✓	51%	16%	-	-	-	-	-	-
SilentGrasp	✓	83%	72%	-	-	-	-	-	-
PROGrasp (ours)		87%	79%	90%	75%	83%	61%	88%	42%

Grasp can access the ground-truth human-to-human dialogue history, so we feed the entire dialogue history (*i.e.*, intention-oriented utterance and a set of question and answer pairs) into the models. As shown in Table 1, GDH significantly boosts Acc@0.9 of Zero-Shot (4%→16%) and SilentGrasp (44%→72%). The results indicate that additional question and answer pairs contain detailed information to identify the target object. Remarkably, PROGrasp outperforms SilentGrasp with GDH, although it does not require the ground-truth dialogue history. The results illustrate that PROGrasp works effectively, even in a more realistic scenario.

Results on the Test Splits. We compare PROGrasp with the compared methods on the test-seen, test-unseen, and test-cluttered splits. In Table 1, PROGrasp consistently yields improvements on all test splits compared with Zero-Shot, SilentGrasp, A-int-only, and LiteralGrasp. Specifically, compared with LiteralGrasp, PROGrasp improves Acc@0.5 on ten points (73%→83%) and Acc@0.9 on seven points (54%→61%) in the test-unseen split. We could not investigate the results of GDH on the test splits since the test splits do not have ground-truth dialogue data. Surprisingly, PROGrasp shows decent Acc@0.5 scores even in the test-cluttered and test-unseen splits, but relatively lower scores are observed on Acc@0.9. It illustrates that (1) accurately identifying partially occluded target objects and (2) generalizing a robotic system to previously unseen objects are challenging aspects of this task. We will discuss more details in the qualitative analysis.

Comparison with the Multimodal Foundation Model. We further identify the performance of the powerful multimodal foundation model, GPT-4V(ision) [12] on the validation split. GPT-4V is provided with images and detailed text prompts (see supplement). Table 2 shows the results. We observe that GPT-4V infers the target object well, but it poorly specifies the ground-truth region coordinates. We thus study a hybrid approach: (1) GPT-4V interacts with users and generates a distinctive caption of the best-fit object and (2) our VG model then predicts the coordinates \hat{r} based on the caption. As shown in Table 2, PROGrasp outperforms the approach on all evaluation metrics. It also demonstrates the significance of pragmatic inference.

Hyperparameter Study. We study how the hyperparameters in PROGrasp (*i.e.*, λ and T) affect performance. Note that

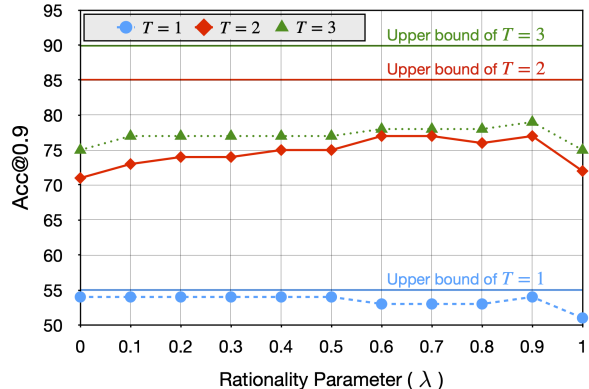


Fig. 4. Validation scores adjusting the hyperparameters, λ and T .

λ indicates the importance of the evaluation from A-int in pragmatic inference, and T denotes the number of interactions between PROGrasp and the human in Algorithm 1. As shown in Figure 4, we visualize the Acc@0.9 performance on the validation split. The blue, red, and green lines denote the results when $T = 1$, $T = 2$, and $T = 3$, respectively. We observe a huge performance gap between $T = 1$ and $T = 2$. It implies that many incorrect guesses in the first round of dialogue are corrected in the second round. Moreover, comparing $T = 2$ with $T = 3$, improvements seem saturated.

We also identify that performance varies depending on the value of the rationality parameter λ . LiteralGrasp is equivalent to $\lambda = 0$, and A-int-only corresponds to $\lambda = 1$. The best results are observed in $\lambda = 0.9$ across all T values. However, we do not see any merit in PROGrasp compared with LiteralGrasp when $T = 1$. To delve into this phenomenon, we visualize the upper bound performance for each T value. The upper bound is defined as the performance when the system perfectly selects the object region in a set of object region candidates \mathcal{R} . The upper bound performance of $T = 1$ is 55%, and PROGrasp and LiteralGrasp both show 54%. It illustrates that there is little room for improvement in $T = 1$. In contrast, the upper bound of $T = 2$ is 85%, and LiteralGrasp shows 71% in Figure 4. PROGrasp boosts 6% compared with LiteralGrasp (71%→77%). Likewise, we observe 4% gains (75%→79%) when $T = 3$. Unless stated otherwise, λ is 0.9, and T is 3.

Communicative Efficiency. Beyond task success, communicative efficiency is also an important criterion for pragmatic

TABLE II
Comparison with the Multimodal Foundation Model.

Method	Acc@0.1	Acc@0.5	Acc@0.9
GPT-4V [12]	29%	9%	1%
GPT-4V [12] + VG	82%	82%	68%
PROGrasp (ours)	87%	87%	79%

TABLE III
Analysis of Communicative Efficiency. ↓ indicates lower is better.

Method	Avg. # of Interactions ↓		
	Test-Seen	Test-Unseen	Test-Cluttered
Random	1.76	2.00	1.98
A-int	1.60	1.78	1.76
LiteralGrasp	1.55	1.78	1.71
PROGrasp (ours)	1.53	1.72	1.69

TABLE IV
Results on the online experiments.

Method	Ambiguous	Non-Ambiguous	Total
	Object Discovery / Success Rate		
SilentGrasp	42 / 30	78 / 38	56 / 33
LiteralGrasp	77 / 47	90 / 45	82 / 46
PROGrasp (ours)	80 / 53	90 / 45	84 / 50

systems [8]. Accordingly, we measure the average number of interactions (*i.e.*, question answering) required for the system to identify the target object. In this study, we assume that the dialogue immediately ends when the Intersection over Union (IoU) between the predicted and target regions is greater than 0.5. The efficiency can range from 1.0 to $T = 3$. We compare PROGrasp with three baselines: Random, A-int-only, and LiteralGrasp. The Random randomly selects the predicted object in the set of candidates (*i.e.*, \mathcal{R} in Algorithm 1). In Table 3, PROGrasp consistently improves the efficiency across all test splits. It indicates that our system efficiently identifies the target object through fewer interactions.

C. Online Experiments

Evaluation Protocol. We reproduce 100 images in the test-split and conduct online experiments to study how well the system picks the desired object up through human-robot dialogue. We compare PROGrasp with SilentGrasp and LiteralGrasp. The human rater evaluates object discovery and success rate. Object discovery measures whether the system correctly localizes the target.

Results. We divide 100 samples into two categories: Ambiguous (60) and Non-Ambiguous (40). The Ambiguous is a set of samples that require disambiguating over two objects given an initial utterance. The Non-Ambiguous corresponds to the other samples. In Table 4, PROGrasp achieves a total execution success rate of 50%, outperforming all baselines. It demonstrates that the superiority of PROGrasp’s target object discovery is successfully transferred to the success rate. Not surprisingly, PROGrasp is effective in the Ambiguous,

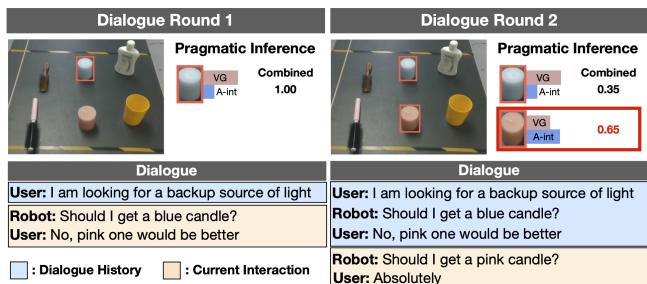


Fig. 5. Visualization of PROGrasp’s target object recovery.

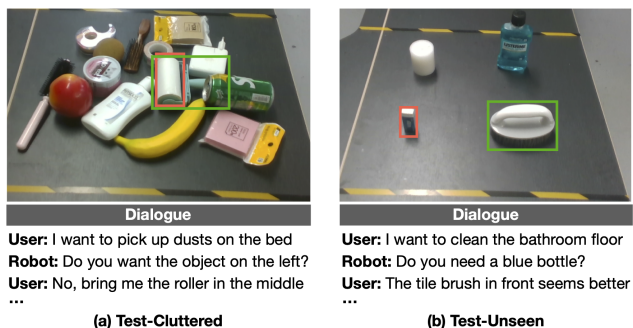


Fig. 6. Visualization of the failure cases.

boosting success rate of 23% compared with SilentGrasp. However, we observe many failure cases, although the target object is correctly localized. The system often drops the objects during lifting or fails to grasp them since objects are highly unstructured (see Figure 3). More precise object grasping will mitigate this issue. We leave it as a future work.

D. Qualitative Analysis

Results. In Figure 5, we visualize the inferred results from PROGrasp when $T = 2$. PROGrasp fails to find the target object (*i.e.*, a pink candle) in the first round of the dialogue, but it corrects the target in the second round by utilizing the question-and-answer pair from the first round as additional context and includes the target object as a candidate. Pragmatic inference finally selects the desired object region. This example clearly explains the performance gap between $T = 1$ and $T = 2$ in Figure 4. We also visualize two failure examples from the test-cluttered and test-unseen splits in Figure 6. The red and green boxes indicate the predicted and ground-truth regions, respectively. As in (a) at Figure 6, our system fails to identify the target region accurately due to the occlusion, which explains low scores in Acc@0.9 of the test-cluttered. Furthermore, PROGrasp sometimes makes an incorrect guess (*i.e.*, (b) in Figure 6) when observing never-seen-before objects, highlighting the need for further generalization in future work.

V. CONCLUSION

We present (1) a new task for interactive object grasping (Pragmatic-IOG) and (2) PROGrasp that effectively infers the user’s intention for Pragmatic-IOG. Experiments demonstrate the effectiveness of our proposed methods, including pragmatic inference. We expect the Pragmatic-IOG to open the door to developing human-centric and pragmatic robots.

REFERENCES

- [1] M. Shridhar and D. Hsu, "Interactive visual grounding of referring expressions for human-robot interaction," *arXiv preprint arXiv:1806.03831*, 2018.
- [2] J. Hatori, Y. Kikuchi, S. Kobayashi, K. Takahashi, Y. Tsuboi, Y. Unno, W. Ko, and J. Tan, "Interactively picking real-world objects with unconstrained spoken language instructions," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 3774–3781.
- [3] H. Zhang, Y. Lu, C. Yu, D. Hsu, X. La, and N. Zheng, "Invigorate: Interactive visual grounding and grasping in clutter," *arXiv preprint arXiv:2108.11092*, 2021.
- [4] Y. Yang, X. Lou, and C. Choi, "Interactive robotic grasping with attribute-guided disambiguation," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 8914–8920.
- [5] Y. Mo, H. Zhang, and T. Kong, "Towards open-world interactive disambiguation for robotic grasping," in *CoRL Workshop on Learning, Perception, and Abstraction for Long-Horizon Planning*, 2022.
- [6] J. R. Searle and J. R. Searle, *Speech acts: An essay in the philosophy of language*. Cambridge university press, 1969, vol. 626.
- [7] M. C. Frank and N. D. Goodman, "Predicting pragmatic reasoning in language games," *Science*, vol. 336, no. 6084, pp. 998–998, 2012.
- [8] D. Fried, N. Tomlin, J. Hu, R. Patel, and A. Nematzadeh, "Pragmatics in grounded language learning: Phenomena, tasks, and modeling approaches," *arXiv preprint arXiv:2211.08371*, 2022.
- [9] N. D. Goodman and M. C. Frank, "Pragmatic language interpretation as probabilistic inference," *Trends in cognitive sciences*, vol. 20, no. 11, pp. 818–829, 2016.
- [10] N. J. Smith, N. Goodman, and M. Frank, "Learning and using language via recursive pragmatic reasoning about other agents," *Advances in neural information processing systems*, vol. 26, 2013.
- [11] L. Yu, P. Poirson, S. Yang, A. C. Berg, and T. L. Berg, "Modeling context in referring expressions," in *European Conference on Computer Vision*. Springer, 2016, pp. 69–85.
- [12] OpenAI, "Gpt-4v(ision) system card," 2023, https://cdn.openai.com/papers/GPTV_System_Card.pdf.
- [13] R. Paul, J. Arkin, N. Roy, and T. M. Howard, "Grounding abstract spatial concepts for language interaction with robots," in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, 2017, pp. 4929–4933.
- [14] M. Shridhar and D. Hsu, "Grounding spatio-semantic referring expressions for human-robot interaction," *arXiv preprint arXiv:1707.05720*, 2017.
- [15] S. G. Venkatesh, A. Biswas, R. Upadrashta, V. Srinivasan, P. Talukdar, and B. Amrutur, "Spatial reasoning from natural language instructions for robot manipulation," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 11 196–11 202.
- [16] T. Nguyen, N. Gopalan, R. Patel, M. Corsaro, E. Pavlick, and S. Tellex, "Robot Object Retrieval with Contextual Natural Language Queries," in *Proceedings of Robotics: Science and Systems*, 2020.
- [17] J. Kim, G.-C. Kang, J. Kim, S. Shin, and B.-T. Zhang, "Gvcci: Lifelong learning of visual grounding for language-guided robotic manipulation," *arXiv preprint arXiv:2307.05963*, 2023.
- [18] S. T. Piantadosi, H. Tily, and E. Gibson, "The communicative function of ambiguity in language," *Cognition*, vol. 122, no. 3, pp. 280–291, 2012.
- [19] D. Whitney, E. Rosen, J. MacGlashan, L. L. Wong, and S. Tellex, "Reducing errors in object-fetching interactions through social feedback," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1006–1013.
- [20] H. P. Grice, "Logic and conversation," in *Speech acts*. Brill, 1975, pp. 41–58.
- [21] W. Monroe, R. X. Hawkins, N. D. Goodman, and C. Potts, "Colors in context: A pragmatic neural model for grounded language understanding," *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 325–338, 2017.
- [22] J. Andreas and D. Klein, "Reasoning about pragmatics with neural listeners and speakers," in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2016, pp. 1173–1182.
- [23] R. Cohn-Gordon, N. Goodman, and C. Potts, "Pragmatically informative image captioning with character-level inference," in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*. Association for Computational Linguistics, 2018, pp. 439–443.
- [24] D. Chen and R. Mooney, "Learning to interpret natural language navigation instructions from observations," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2011, pp. 859–865.
- [25] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, I. Reid, S. Gould, and A. Van Den Hengel, "Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 3674–3683.
- [26] H. De Vries, F. Strub, S. Chandar, O. Pietquin, H. Larochelle, and A. Courville, "Guesswhat?! visual object discovery through multimodal dialogue," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5503–5512.
- [27] J.-H. Kim, N. Kitaev, X. Chen, M. Rohrbach, B.-T. Zhang, Y. Tian, D. Batra, and D. Parikh, "CoDraw: Collaborative drawing as a testbed for grounded goal-driven communication," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 2019, pp. 6495–6513.
- [28] J. Y. Chai, L. She, R. Fang, S. Ottarson, C. Little, C. Liu, and K. Hanson, "Collaborative effort towards common ground in situated human-robot dialogue," in *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*, 2014, pp. 33–40.
- [29] G.-C. Kang, S. Kim, J.-H. Kim, D. Kwak, and B.-T. Zhang, "The dialog must go on: Improving visual dialog via generative self-training," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 6746–6756.
- [30] S. Changpinyo, P. Sharma, N. Ding, and R. Soricut, "Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 3558–3568.
- [31] C. Schuhmann, R. Vencu, R. Beaumont, R. Kaczmarczyk, C. Mullis, A. Katta, T. Coombes, J. Jitsev, and A. Komatsuzaki, "Laion-400m: Open dataset of clip-filtered 400 million image-text pairs," *arXiv preprint arXiv:2111.02114*, 2021.
- [32] P. Wang, A. Yang, R. Men, J. Lin, S. Bai, Z. Li, J. Ma, C. Zhou, J. Zhou, and H. Yang, "Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework," in *International Conference on Machine Learning*. PMLR, 2022, pp. 23 318–23 340.
- [33] J. Lu, D. Batra, D. Parikh, and S. Lee, "Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks," *Advances in neural information processing systems*, vol. 32, 2019.
- [34] H. Tan and M. Bansal, "LXMERT: Learning cross-modality encoder representations from transformers," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, 2019, pp. 5100–5111.
- [35] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," *Advances in neural information processing systems*, vol. 27, 2014.
- [36] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [37] S.-W. Lee, Y.-J. Heo, and B.-T. Zhang, "Answerer in questioner's mind: Information theoretic approach to goal-oriented visual dialog," *Advances in neural information processing systems*, vol. 31, 2018.
- [38] D. Fried, R. Hu, V. Cirik, A. Rohrbach, J. Andreas, L.-P. Morency, T. Berg-Kirkpatrick, K. Saenko, D. Klein, and T. Darrell, "Speaker-follower models for vision-and-language navigation," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [39] S. Shen, D. Fried, J. Andreas, and D. Klein, "Pragmatically informative text generation," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, 2019, pp. 4060–4067.
- [40] R. Schnabel, R. Wahl, and R. Klein, "Efficient ransac for point-cloud shape detection," in *Computer graphics forum*, vol. 26. Wiley Online Library, 2007, pp. 214–226.
- [41] D. Coleman, I. Sucas, S. Chitta, and N. Correll, "Reducing the barrier to entry of complex robotic software: a moveit! case study," *arXiv preprint arXiv:1404.3785*, 2014.