

ManipVQA: Injecting Robotic Affordance and Physically Grounded Information into Multi-Modal Large Language Models

Siyuan Huang*, Iaroslav Ponomarenko*, Zhengkai Jiang, Xiaoqi Li, Xiaobin Hu
Peng Gao, Hongsheng Li, Hao Dong

Abstract—While the integration of Multi-modal Large Language Models (MLLMs) with robotic systems has significantly improved robots’ ability to understand and execute natural language instructions, their performance in manipulation tasks remains limited due to a lack of robotics-specific knowledge. Conventional MLLMs are typically trained on generic image-text pairs, leaving them deficient in understanding affordances and physical concepts crucial for manipulation. To address this gap, we propose ManipVQA, a novel framework that infuses MLLMs with manipulation-centric knowledge through a Visual Question-Answering (VQA) format. This approach encompasses tool detection, affordance recognition, and a broader understanding of physical concepts. We curated a diverse dataset of images depicting interactive objects, to challenge robotic understanding in tool detection, affordance prediction, and physical concept comprehension. To effectively integrate this robotics-specific knowledge with the inherent vision-reasoning capabilities of MLLMs, we leverage a unified VQA format and devise a fine-tuning strategy. This strategy preserves the original vision-reasoning abilities while incorporating the newly acquired robotic insights. Empirical evaluations conducted in robotic simulators and across various vision task benchmarks demonstrate the robust performance of ManipVQA. The code and dataset are publicly available at <https://github.com/SiyuanHuang95/ManipVQA>.

I. INTRODUCTION

Recently, Multi-modal Large Language Models (MLLMs) such as OpenAI’s GPT-4V [1], Google’s Gemini and SPHINX-X [2], have significantly advanced the capabilities for understanding and interpreting both text and images. These models achieve this improvement by aligning multi-modal encoders with Large Language Models (LLMs) through training on numerous text-image pairs to enhance a comprehensive understanding of both modalities. As a result, MLLMs exhibit promising potential in addressing common sense reasoning and demonstrate remarkable generalization in vision tasks. However, the application of these models in manipulation tasks [3], [4], such as robotic affordance understanding and the design of associated robotic systems, continues to present significant challenges.

Robotic manipulation [4], [5] is defined as a robot’s capacity to perceive its environment and discern the potential actions applicable to various objects. The direct application of existing MLLMs, pre-trained on common scene reasoning,

* indicates the equal contribution.

Siyuan Huang, Peng Gao, and Hongsheng Li are with the Shanghai AI Laboratory. Iaroslav Ponomarenko, Xiaoqi Li and Hao Dong are with Peking University. Siyuan Huang is also with Shanghai Jiao Tong University and Hongsheng Li is also with the MMLab, CUHK. Zhengkai Jiang is with UCAS and Xiaobin Hu is with TUM.




Fig. 1: The ManipVQA model generates predictions in a unified VQA format, combining tool-object and affordance detection, affordance grounding, and an understanding of physical object properties. The displayed masks are created using SAM-HQ, based on the bounding boxes predicted by ManipVQA. The contact point, represented by a red star.

to robotic manipulation, does not yield satisfactory performance due to the absence of low-level action samples in their pretraining data. Previous research has addressed robotic affordance by prompting MLLMs to process scene images and then generate a sequence of robotic actions, but the performance has been suboptimal. Affordance grounding [6] aims to localize the regions in objects where actions are possible. This task also faces the challenge of establishing an explicit link with object parts due to the diversity of interactive affordances. Prior research [7] employed LLMs to get the answer on affordance with prompt engineering. Despite the potential utility of MLLMs in robotic applications, their effectiveness is limited by several challenges. Classic MLLMs [2] are generally trained on generic image-text pairs, thereby lacking critical robotic knowledge for understanding object affordances and their physical properties. This deficiency in specialized knowledge impedes their performance in manipulation tasks, consequently restricting

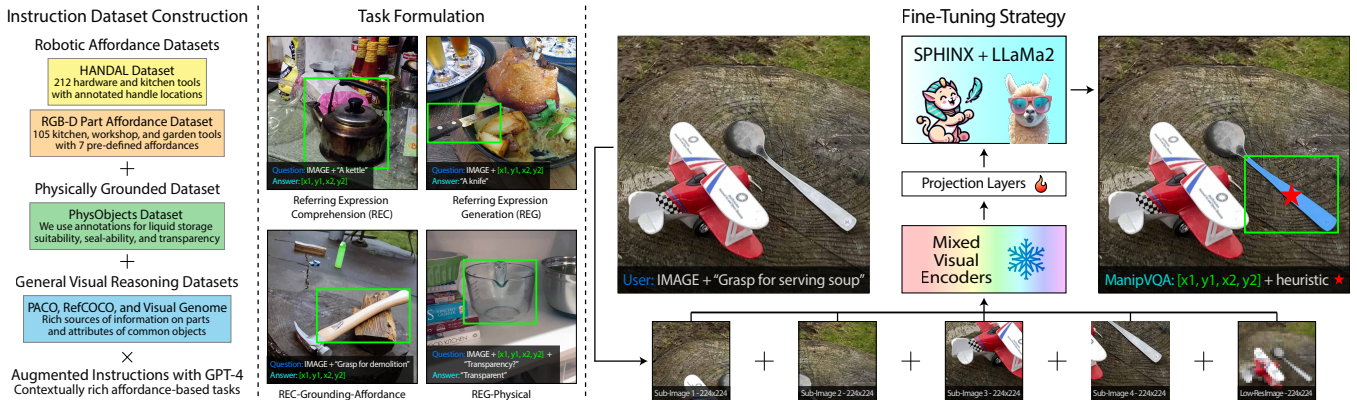


Fig. 2: Overview of ManipVQA: We created a comprehensive vision-language dataset by merging existing datasets and expanding affordance grounding tasks with ChatGPT. To ensure consistency with existing VQA datasets, we adopted a similar VQA format. Using this dataset, we then fine-tuned an MLLM. When combined with a heuristic policy, the enhanced MLLM can perform a wide range of manipulation tasks.

the types of tasks that robots can execute and the precision of their execution. The gap between the capabilities of MLLMs and the demands of robotic systems constitutes a significant challenge that needs to be addressed.

To address this gap, we present ManipVQA, an innovative framework devised to equip MLLMs with manipulation-centric knowledge via a Visual Question-Answering paradigm. This integration is achieved through a unified VQA approach and a fine-tuning strategy that preserves the original vision-reasoning capabilities of MLLMs while infusing them with critical insights aimed at robotic tasks, enhancing tool detection, affordance recognition, and physical concepts understanding. To achieve this, we collect a diverse set of images featuring interactive objects, thus encompassing a broad range of challenges related to object detection, affordance, and physical concept prediction.

Empirical assessments performed in robotic simulators and across various vision task benchmarks substantiate the robust performance of ManipVQA. Our research makes several significant contributions to the fields of robotics and machine learning. First, we propose a novel approach to robotic manipulation and affordance understanding tasks, which addresses the shortcomings of existing methodologies. Second, we are dedicated to fostering the research community and have made our datasets, codes, and models publicly available.

II. RELATED WORK

A. Large-Object Datasets

The field of computer vision has been significantly advanced by datasets such as ImageNet [8] and PACO [9], which have been crucial for progress in image classification, object detection, and segmentation. While these datasets enhance object instance detection and segmentation and provide a broad semantic understanding of objects, they often overlook the granularity of object parts and attributes, leading to a gap in robotic applications. To address this gap, datasets like HANDAL [10] and PhysObjects [11] have been introduced, focusing on robotic needs by providing annotations for part masks, object attributes, and affordances.

HANDAL emphasizes manipulable object pose estimation with 308k annotated frames, while PhysObjects offers detailed annotations of household object properties. To unify these advancements, ManipVQA merges these datasets into a cohesive VQA dataset format, granting models the ability to perform visual reasoning on general images as well as affordance reasoning for manipulation tasks, thereby enhancing the model’s visual understanding for robotic applications.

B. Multi-Modal Large Language Model

Building on the foundation of extensive LLMs like LLaMa [12], which exhibit remarkable language reasoning capabilities across various tasks, MLLMs have significantly expanded the scope of language processing by incorporating the ability to understand visual stimuli [2], [13]. A notable innovation in this field is SPHINX [13], which excels in multi-modal tasks by aligning ensembled visual features with language embeddings through projection layers. Despite these advances, the application of MLLMs to robotic manipulation remains in its early stages. Some recent initiatives such as [4], [6], [14] have attempted to integrate robotic-domain knowledge into MLLMs. However, these efforts often overlook critical physical information [4] or do not fully embody a robotic-centric approach [6]. In contrast, our proposed ManipVQA integrates essential physical knowledge and affordance reasoning capabilities from a robotic perspective into MLLMs. We present this information in a direct, human-readable, purely linguistic format, enhancing the utility of MLLMs for practical robotic applications.

C. Language-Driven Robotics Manipulation

Instruction-based policies in robotics serve as a bridge between high-level human communication and low-level robotic control, offering intuitive interaction and the potential for skill transfer and complex task planning. One approach [15], [16] involves using a pre-trained text encoder to extract text embeddings that encapsulate linguistic knowledge, which are then used to inform robotic policies. For example, CLIPort has obtained the ability to understand semantics and manipulate objects by encoding text input

Phys. Concepts	Descriptions	Type
Transparency	Object’s light transmission ability	Levels
Liquids Storage	Object’s liquid-holding capacity	Bool
Seal-ability	Object’s closure state	Bool

TABLE I: Physical concepts used in ManipVQA.

through CLIP [17]. Another research direction leverages the reasoning capabilities of LLMs for planning [18], [19] or even directly generating the policy codes [3], [20], [21] by prompt engineering. For example, CaP [22] generates policy codes with detailed comments and context-specific examples to guide LLM output. Furthermore, some studies have trained large-scale policy models [23]–[25] using extensive robotic datasets, aiming to learn robotic control policies from diverse data directly. Despite the progress in these areas, a gap remains in the explicit consideration of affordance grounding and physical reasoning, which are essential for proficient manipulation in robotics. Aiming to bridge this gap, our ManipVQA injects robotic affordance and physically grounded information into MLLM while preserving its original reasoning ability.

III. METHOD

A. Modeling of Affordances and Physical Concepts

In robotic manipulation, understanding and modeling object affordances is essential for enabling robots to interact effectively with their environment. We extend upon the graspable affordance model $\mathbf{A}_{\text{grasp}}$ for common tool objects \mathbf{O}_{tool} , forming the tuple $(\mathbf{A}_{\text{grasp}}, \mathbf{O}_{\text{tool}})$, as delineated by HANDAL [10]. In our representation, each (partial) object is paired with a bounding box, described by the coordinates $[x_{\min}, y_{\min}, x_{\max}, y_{\max}]$, which specify the top-left and bottom-right corners, respectively.

To ensure our model’s adaptability across different contexts, we normalize these coordinates relative to the image dimensions, enhancing the model’s generalizability and robustness. Furthermore, we recognize that affordances may vary depending on the specific task \mathbf{T} at hand. For instance, distinct regions of a tool may be utilized for different functions. To address this, we form a tuple $(\mathbf{A}_{\mathbf{T}}, \mathbf{O}_{\text{tool}})$, which associates task-specific affordances $\mathbf{A}_{\mathbf{T}}$ with the tool object. The task \mathbf{T} is succinctly described using a brief natural language sentence to encapsulate its functionality.

In addition, we incorporate the modeling of physical concepts, denoted as \mathbf{P}_i . These concepts are quantified using discrete levels or boolean values, drawing inspiration from the methodology presented in PhysObjects [11]. Each physical concept \mathbf{P}_i is linked to its corresponding object \mathbf{O} , resulting in the tuple $(\mathbf{P}_i, \mathbf{O})$. Table I enumerates the physical concepts under consideration, alongside their respective descriptions. These concepts include but are not limited to, transparency, liquid storage capacity, and seal-ability. Each concept is crucial for the robot’s ability to perceive and interact with various objects within its environment.

B. Instruction Dataset Construction

In the development of the ManipVQA model, we focused on creating a comprehensive training dataset by combining

various publicly available datasets. The goal was to equip the model with a strong understanding of robotic affordances, advanced visual reasoning, and a solid foundation in physically grounded knowledge.

1) **Robotic Affordance Datasets:** The HANDAL Dataset [10] serves as a foundation base for imparting the ManipVQA model with the ability to discern objects and identify graspable components. It comprises over 212 real-world objects, each with annotated handle locations. To further refine the model’s capability in understanding complex affordances, we integrated the RGB-D Part Affordance Dataset [26], which differentiates seven pre-defined affordances such as grasp, cut, scoop, contain, pound, support, and wrap-grasp. While the AGD20K Dataset [27] represents another extensive affordance resource, it is predominantly labeled with broad human-centric action terms, such as *drink* or *sit-on*. Therefore, we have excluded AGD20K from the training regime and instead selected it for zero-shot evaluation, assessing the model’s generalization capabilities.

2) **Physically Grounded Dataset:** The PhysObjects Dataset [11] was employed to infuse the model with a nuanced understanding of physical properties. This dataset features eight core physical concepts, but we selectively utilized annotations relevant to our focus—specifically, those related to liquid storage suitability, seal-ability, and transparency as listed in Table I.

3) **General Visual Reasoning Datasets:** To maintain and extend the model’s capabilities in general visual reasoning, we incorporated datasets like PACO [9], RefCOCO [28], and the Visual Genome [29], which provide a rich source of information on parts and attributes of common objects.

4) **Augmented Instructions with GPT-4:** Considering the limitations of existing annotations in the robotic affordance datasets—primarily confined to partial masks in [10] and basic action types with associated masks in [26], we employed GPT-4 to generate complex and contextually rich affordance-based tasks. This augmentation enables our model to not only learn explicit commands but also to interpret and execute complex, implicitly defined tasks. Listing 1 showcases examples of the prompts used for generating these advanced affordance grounding tasks.

C. Task Formulation

The ManipVQA training protocol integrates a pair of principal vision-language tasks: Referring Expression Comprehension (REC) and Referring Expression Generation (REG). Following [30], REC involves the model receiving an image accompanied by a natural language description and subsequently predicting the bounding box coordinates that delineate the specified target within the image. Conversely, REG prompts the model to produce a descriptive natural language statement about an area within an image, defined by provided bounding box coordinates. To advance ManipVQA’s proficiency in recognizing robotic affordances and discerning object physical properties, we have augmented the task framework with:

Capabilities	Tasks	Examples of Task Templates	Source	Num.
Gen. Visual Reasoning	REC/REG: Object	User: Please provide a short description of this region: BBox . ManipVQA: A ratchet.	H/P/C/V	36K
Robotic	REC/REG: Affordance	User: Please provide bounding box coordinates of this region: handle of a screwdriver. ManipVQA: BBox .	H/R	26K
Aff. Understanding	REC-Grounding: Affordance	User: Please provide bounding box coordinates of this region: grasp for tightening bolts. ManipVQA: BBox .	H/R	15K
Phys. Gr. Understanding	REG-Phys: Liq./Seal./Transp.	User: Please provide a short description of whether this object can contain liquid: BBox . ManipVQA: True.	Phys	7K

TABLE II: Overview of the ManipVQA Dataset: This table summarizes the tasks, their associated capabilities, example templates, and the number of samples in each dataset. The capabilities are represented by the abbreviations: “Gen.” for *General*, “Aff.” for *Affordance*, and “Phys. Gr.” for *Physically Grounded*. The source datasets are identified by the acronyms: H for HANDAL [10], P for PACO [9], C for RefCOCO [28], R for RGB-D Part Affordance Dataset [26], and Phys for PhysObjects [11]. The bounding box (BBox) format is represented by $[x_{min}, y_{min}, x_{max}, y_{max}]$.

<p>Role: You are an analytical assistant specializing in robotic affordance grounding. Your expertise is in creating tasks that facilitate the training of robotic policies, enabling robots to reason about task execution, such as determining the appropriate part of an object to grasp.</p> <p>Task Description: You will be provided with the name of a tool that can be attached to a robotic arm. The robot is expected to use this tool to perform a variety of everyday tasks. Along with the tool name, you will receive a list of tasks that have already been generated for this tool.</p> <p>Guidelines: Diversity: Aim for a wide range of tasks, ensuring that there is no overlap with previous ones. Daily Tasks: Tasks should be common and representative of the ones encountered in daily life. Leakage Avoidance: Ensure that the generated tasks do not explicitly mention the name of the tool object.</p> <p>Examples: ...</p> <p>Instruction: With the provided OBJECT_NAME, generate five new affordance grounding tasks. Use the HISTORY of generated tasks as a reference to ensure compliance with the diversity guideline. Output should be in the JSON format with the object name as the key.</p>
--

Listing 1: Example prompt for generating an affordance grounding tasks for a specific tool object using ChatGPT.

- **REC-Grounding-Affordance:** This task refines the model’s capacity to identify functional parts of objects based on their usage descriptions. It presents the challenge of localizing these parts without directly naming the object or its components, a step towards intuitive affordance recognition for robots.
- **REC-Physical:** This task broadens the model’s attribute recognition by requiring it to pinpoint objects based on their physical properties and to engage with related inquiries. This is essential for detailed robotic perception and manipulation.

These additional tasks enhance the core REC and REG tasks, together cultivating a robust skill set tailored for practical robotic deployment. Detailed instances of these task formats can be found in Table II.

D. MLLM Fine-Tuning Strategy

1) **Model Architecture:** We adopt the MLLM, SPHINX [13] as our primary architecture, with LLaMA2 [12] serving as the language backbone. Given

the necessity for both global and local visual grounding in robotic tasks, we integrate the visual encoder from CLIP [17] to extract local semantic features and the Q-Former [31] for summarizing visual features. Spatial alignment is facilitated using projection layers, and global features are merged with local ones through channel-wise concatenation. We acknowledge that the standard image resolution for pre-trained visual encoders, typically 224×224 , is insufficient for detailed visual perception. This limitation is particularly significant for robotic affordance reasoning, which often requires fine-grained visual grounding of object parts, such as tool handles or machine buttons. To augment ManipVQA’s region-level grounding capabilities, we employ a sub-images patching strategy following [13]. Specifically, we partition a 448×448 image into four 224×224 sub-images taken from each corner, thereby preserving intricate visual details. The resulting image tokens are then positioned before the language instructions to provide visual context for the ensuing prompts.

2) **Fine-Tuning Strategy:** As elucidated in Sec. III-A, we model both affordance and physical concepts within natural language representations and training samples are formatted in line with the general VQA framework, as delineated in Sec. III-C. As a result, the training objective employs a unified cross-entropy loss, diverging from the approaches in [4], [6]. To maintain the model’s broad visual reasoning proficiency, we amalgamate general visual reasoning exercises with tasks specific to robotics.

IV. EXPERIMENTS

A. Implementation Details

1) **Training Details:** We fine-tuned the ManipVQA model using the SPHINX framework [13] on eight NVIDIA A100 (80GB) GPUs. The fine-tuning was completed in a single epoch, which took approximately 6 hours. During this phase, we kept the visual encoders frozen to maintain the integrity of the pre-trained features. For training, we used the pre-trained SPHINX-1K model from the official repository, with a batch size of 4 and a learning rate of 2×10^{-5} .

2) **Connected with Robotic Policy:** The objective of ManipVQA is to augment the generalizability of robotic control policies. It could be used in a language-free format during high-level decision-making for physically grounded knowledge. For affordance localization, while the initial

Method	Task	Hardware Tools										Kitchen Tools							AVG
		Ha	Pf	Ps	Pl	Pd	Ra	Sd	Wa	Wc	La	Mc	Mug	Pan	Sp	St	Ut	Wh	
HANDAL	Obj-B	0.76	0.48	0.42	0.77	0.74	0.67	0.60	0.63	0.79	0.81	0.53	0.53	0.80	0.46	0.66	0.74	0.79	0.66
Ours		0.96	0.96	0.94	0.89	0.98	0.93	0.94	0.93	0.92	0.97	0.95	0.90	0.97	0.90	0.94	0.88	0.95	0.94
Ours	Aff-B	0.70	0.67	0.64	0.71	0.33	0.71	0.76	0.56	0.69	0.82	0.51	0.52	0.59	0.54	0.81	0.72	0.65	0.64
HANDAL	Obj-M	0.62	0.38	0.27	0.67	0.75	0.44	0.56	0.48	0.52	0.67	0.49	0.55	0.81	0.40	0.62	0.57	0.75	0.56
LISA		0.84	0.63	0.66	0.70	0.93	0.68	0.76	0.79	0.80	0.77	0.67	0.78	0.83	0.77	0.58	0.79	0.84	0.75
Ours		0.71	0.70	0.61	0.65	0.57	0.52	0.82	0.66	0.62	0.57	0.53	0.55	0.43	0.55	0.46	0.58	0.55	0.58
LISA	Aff-M	0.67	0.59	0.56	0.54	0.43	0.48	0.62	0.62	0.61	0.41	0.41	0.56	0.45	0.58	0.40	0.49	0.49	0.59
Ours		0.75	0.70	0.54	0.53	0.44	0.56	0.80	0.65	0.68	0.62	0.60	0.52	0.62	0.65	0.64	0.58	0.68	0.62
LISA	Gr-Aff-M	0.57	0.42	0.41	0.48	0.35	0.41	0.62	0.54	0.56	0.37	0.35	0.43	0.36	0.39	0.37	0.41	0.48	0.44
Ours		0.82	0.69	0.49	0.42	0.42	0.54	0.83	0.66	0.64	0.71	0.66	0.59	0.73	0.71	0.78	0.78	0.63	0.65

TABLE III: Evaluation results on HANDAL Dataset. Task abbreviations are as follows: “Obj.” for *Complete Object Detection*, “Aff.” for *Robotic Affordance Detection*, and “Gr” for *Grounded Detection*. The letter “B” denotes bounding box format, and “M” denotes mask representation format. Object abbreviations are listed in sequence: Hammer, Pliers-Fixed Joint, Pliers-Slip Joint, Pliers-Locking, Power Drill, Ratchet, Screwdriver, Wrench-Adjustable, Wrench-Combinational, Label, Measuring Cup, Mug, Pan, Spatula, Strainer, Utensil, and Whisk.

bounding box identifies areas of potential manipulation, it may also encompass extraneous elements like background features. To achieve a more precise affordance map, we leverage the SAM-HQ variant [32], which uses the initial bounding box as a “box prompt” for more accurate segmentation. Heuristic methods are then employed to determine the contact point. In our experiments, the ground truth (GT) surface normals, essential for rotation estimation, were presumed to be accessible either through RGB-D sensing or an alternative pre-trained model.

B. Experimental Setup

We aim to systematically evaluate the performance of our approach, encompassing both robotic affordance grounding and robotic manipulation tasks.

1) *Robotic Affordance Detection*: Our evaluation is primarily conducted on the HANDAL Dataset [10], assessing both bounding box average precision (AP) and pixel-wise segmentation AP. Unlike the baseline model in the HANDAL, which detects only the whole object, our ManipVQA is capable of identifying both the entire object and its manipulable parts, namely the affordance regions. Additionally, we compare our approach to LISA [33], which integrates LLM and SAM decoder.

2) *Physical Concept Grounding*: We utilize the PhysObjects Dataset [11] to evaluate physical concept grounding capabilities. We benchmark ManipVQA against PG-InstructBLIP [11], a fine-tuned version of InstructBLIP on PhysObjects, and the latest and most advanced MLLM, GPT-4V. Due to the limited localization capacity of GPT-4V, we pair it with Set-of-Mark [34], where bounding boxes and index numbers are explicitly annotated on the target objects within the input images.

3) *General Affordance Grounding*: Although ManipVQA is trained solely on a robotic affordance dataset, we are interested in its generalization capabilities on broader affordance grounding datasets, such as on AGD20K [27]. This exploration is motivated by the robust reasoning and generalization potential of LLMs. Our method is evaluated on AGD20K and follows its metrics, including KLD, SIM, and NSS. We compared with

AffordanceLLM [6], Cross-View-AG [27], LOCATE [35] and 3DOI [36].

4) *Robotic Manipulation Tasks*: We further integrate ManipVQA with a basic robotic control policy, as detailed in Section IV-A, and use the manipulation success rate as a metric to gauge its practicality in robotic manipulation tasks using the PartNet-Mobility Dataset within the SAPIEN [37] simulator. Our experimental setup is modeled after ManipLLM [4], utilizing identical metrics. However, we employ ground truth (GT) surface normals from the simulator for rotation estimation. We compare our method with Where2Act [38], FlowBot3D [39], and ManipLLM [4]. For additional details on the experimental settings, please refer to ManipLLM [4].

C. Results

1) *Robotic Affordance Detection Evaluation*: As illustrated in Table III, our ManipVQA achieves remarkable performance in the detection of both complete objects and their affordances in a unified framework. Enhanced by the SAM-HQ, our model attains superior results in tasks involving affordance detection and grounding with a mask. However, it falls short in full object segmentation, likely due to SAM-HQ’s tendency for over-segmentation.

2) *Physical Concept Grounding Evaluation*: Table IV presents the evaluation results on the PhysObjects [11]. It is noteworthy that even the advanced MLLM GPT-4V encounters challenges with tasks that require an understanding of physical concepts. This struggle probably stems from its limited capacity for precise localization and a deficiency in visual physical reasoning. With the integration of SoM [34], the GPT-4V has better localization ability while its performance remains suboptimal. Our ManipVQA outperforms PG-InstructCLIP [11] which is also fine-tuned on PhysObjects, and we hypothesize that this enhanced performance can be attributed to the more powerful LLM and the ensemble of vision encoders deployed within our MLLM.

3) *General Affordance Grounding*: Table V presents the evaluation results on the AGD20K Dataset [27] using the *Hard* split as defined by [6]. Remarkably, although our ManipVQA framework is trained solely on robotic




Fig. 3: Illustration of the affordance detection and grounding tasks using the HANDAL Dataset. The first four columns show the results for affordance detection, and the last two columns present the results for the affordance grounding task. In the affordance grounding task, a description of the tool’s intended use is given instead of explicit part names. Blue masks are produced by SAM-HQ based on ManipVQA’s predicted bounding boxes, and red masks are LISA’s output.



Fig. 4: Illustration for the general affordance grounding tasks on the AGD20K [27]. Notably, the action *Sit-On* is unseen in our affordance training split. Blue masks are produced by SAM-HQ based on ManipVQA’s predicted bounding boxes.

Methods	Trans.	Liquid Stor.	Seal.	AVG
GPT-4V	35.0	52.8	49.7	45.8
SoM [34] + GPT-4V	34.5	55.5	53.1	47.7
PG-InstructBLIP [11]	83.8	89.1	80.6	84.5
Ours	93.5	85.6	91.7	90.3

TABLE IV: Physical evaluation results on PhysObject [11].

affordances, such as grasping, it still demonstrates competitive performance on the general affordance grounding task. Notably, our method achieves the best NSS score, indicating strong model performance, but also the highest KLD, which is undesirable. We attribute this high KLD to the tendency of the SAM-series models to over-segment images. Additionally, the discrepancy between the ground truth represented as a heatmap, and our output, a segmentation mask, likely contributes to the lower KLD score. Our method also shows a promising ability to distinguish between fine-grained affordances associated with the same object class. For example, it can differentiate between *Cut-with* and *Hold* actions for a knife, which involve the blade and handle,

Method	SIM \uparrow	NSS \uparrow	KLD \downarrow
Cross-View-AG [27]	0.209	0.138	2.092
LOCATE [35]	0.282	0.276	1.829
3DOI [36]	0.200	0.549	4.017
AffordanceLLM [6]	0.361	0.947	1.661
Ours	0.246	1.735	12.67

TABLE V: General affordance evaluation results on AGD20K [27].

respectively, as illustrated in Fig. 5.

4) **Robotic Manipulation in Simulator:** Table VI illustrates the zero-shot performance of our model in SAPIEN [37] when combined with a heuristic-based control policy. The model’s success is mainly because it retains commonsense reasoning from the MLLM and incorporates affordance knowledge, which enables effective robotic manipulation without task-specific fine-tuning.

D. Further Analysis

1) **Ablation Studies:** To dissect the contributions of the designs in ManipVQA, we performed a series of ablation

Method	Training Categories														
Where2Act [38]	0.26	0.36	0.19	0.27	0.23	0.11	0.15	0.47	0.14	0.24	0.12	0.56	0.68	0.07	0.40
FlowBot3D [39]	0.67	0.55	0.20	0.32	0.27	0.31	0.61	0.68	0.15	0.28	0.18	0.21	0.70	0.18	0.26
ManipLLM [4]	0.68	0.64	0.36	0.77	0.43	0.62	0.65	0.61	0.65	0.52	0.40	0.64	0.71	0.60	0.64
Ours	0.67	0.87	0.46	0.91	0.56	0.42	0.69	0.79	0.41	0.53	0.69	1.00	0.53	0.17	0.58

Method	Training Categories					Testing Categories								AVG	
Where2Act [38]	0.13	0.18	0.13	0.40	0.18	0.35	0.38	0.28	0.05	0.21	0.17	0.20	0.15	0.15	0.25
FlowBot3D [39]	0.17	0.53	0.29	0.42	0.23	0.10	0.60	0.39	0.27	0.42	0.28	0.51	0.13	0.23	0.35
ManipLLM [4]	0.41	0.75	0.44	0.67	0.38	0.22	0.81	0.86	0.38	0.85	0.42	0.83	0.26	0.38	0.57
Ours	0.20	0.56	0.47	0.75	0.68	0.93	0.92	0.82	0.32	0.58	0.71	0.81	0.69	0.51	0.63

TABLE VI: Performance evaluation within the SAPIEN simulator using PartNet-Mobility Dataset. Our model achieves robust performance without fine-tuning on the SAPIEN samples.



Fig. 5: Strong reasoning capability of ManipVQA to distinguish subtle features within affordances of the same object class. Blue masks are produced by SAM-HQ based on ManipVQA’s predicted bounding boxes.

studies. The results are presented in Table VII. It was observed that the SOTA MLLM, SPHINX, cannot execute vision-based reasoning tasks in robotics without the ManipVQA Dataset. Furthermore, without the mixture of general vision data during the fine-tuning process, there is a noticeable decline in its physical understanding and affordance reasoning capabilities. This suggests that the current ManipVQA Dataset alone may not provide sufficiently large or diverse samples for effectively fine-tuning an MLLM. Additionally, the absence of visual ensembles leads to a significant drop in the model’s ability to reason about affordances, likely because robotic affordance reasoning often requires detailed part-level understanding.

Manip. Data	Vis. Ens	Mix. Train	Phys.	Aff. Box	Aff. Mask
	✓		39.7	0.16	0.31
✓	✓		84.2	0.48	0.30
✓		✓	86.7	0.40	0.61
✓	✓	✓	90.3	0.64	0.62

TABLE VII: Ablation Studies. “Manip. Data” indicates the use of the ManipVQA Dataset, “Vis. Ens” represents the employment of a visual encoder ensemble, and “Mix. Train” refers to the inclusion of a mixed general visual dataset during fine-tuning. “Phys.” assesses the model’s physical grounding; and “Aff.” denotes the model’s affordance reasoning capabilities.

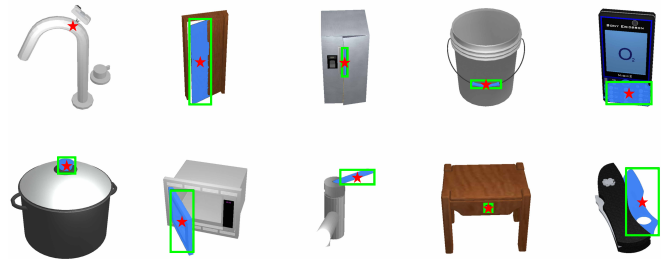


Fig. 6: Visualizations of ManipVQA’s predictions within the SAPIEN simulator. The first sub-image depicts the experimental setup, where the model predicts the gripper’s contact point, represented by a red star, to achieve the desired movement of a specific object part. The green bounding boxes denote ManipVQA’s predictions, while the blue masks are obtained using SAM-HQ. The contact point is heuristically determined as the geometric center of the blue mask.

2) *Impact on Pre-Existing Vision Reasoning Ability:* When fine-tuning a model on a specialized dataset like ManipVQA, it is essential to consider the potential impact on its pre-established general vision reasoning skills. To investigate this, we evaluated the model on the val split of RefCOCO+ [28]. The post-fine-tuning accuracy was recorded at 81.8%. This performance is to be compared against the accuracy of the pre-trained model from [2] at 86.6%. Despite a slight drop, the model retains a robust general vision reasoning ability.

V. CONCLUSION

This study seeks to reconcile the disparity between the capabilities of existing MLLMs and the demands of robotic systems. We present ManipVQA, a novel approach designed to equip MLLMs with manipulation-centric knowledge via a visual question-answering paradigm. Our approach involves the collection of a diverse set of images featuring interactive objects, thus encompassing a broad range of challenges related to object detection, affordance, and physical concept prediction. Empirical assessments performed in robotic simulators and across various vision task benchmarks substantiate the efficacy and resilience of ManipVQA.

REFERENCES

- [1] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, “Gpt-4 technical report,” *arXiv preprint arXiv:2303.08774*, 2023. **1**
- [2] P. Gao, R. Zhang, C. Liu, L. Qiu, S. Huang, W. Lin, S. Zhao, S. Geng, Z. Lin, P. Jin *et al.*, “Sphinx-x: Scaling data and parameters for a family of multi-modal large language models,” *arXiv preprint arXiv:2402.05935*, 2024. **1, 2, 7**
- [3] W. Huang, C. Wang, R. Zhang, Y. Li, J. Wu, and L. Fei-Fei, “Voxposer: Composable 3d value maps for robotic manipulation with language models,” *arXiv preprint arXiv:2307.05973*, 2023. **1, 3**
- [4] X. Li, M. Zhang, Y. Geng, H. Geng, Y. Long, Y. Shen, R. Zhang, J. Liu, and H. Dong, “Manipllm: Embodied multimodal large language model for object-centric robotic manipulation,” *arXiv preprint arXiv:2312.16217*, 2023. **1, 2, 4, 5, 7**
- [5] X. Jia, D. Blessing, X. Jiang, M. Reuss, A. Donat, R. Litoutikov, and G. Neumann, “Towards diverse behaviors: A benchmark for imitation learning with human demonstrations,” *arXiv preprint arXiv:2402.14606*, 2024. **1**
- [6] S. Qian, W. Chen, M. Bai, X. Zhou, Z. Tu, and L. E. Li, “Affordancellm: Grounding affordance from vision language models,” *arXiv preprint arXiv:2401.06341*, 2024. **1, 2, 4, 5, 6**
- [7] W. Xia, D. Wang, X. Pang, Z. Wang, B. Zhao, and D. Hu, “Kinematic-aware prompting for generalizable articulated object manipulation with llms,” *arXiv preprint arXiv:2311.02847*, 2023. **1**
- [8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” pp. 248–255, 2009. **2**
- [9] V. Ramanathan, A. Kalia, V. Petrovic, Y. Wen, B. Zheng, B. Guo, R. Wang, A. Marquez, R. Kovvuri, A. Kadian, A. Mousavi, Y. Song, A. Dubey, and D. Mahajan, “Paco: Parts and attributes of common objects,” 2023. **2, 3, 4**
- [10] A. Guo, B. Wen, J. Yuan, J. Tremblay, S. Tyree, J. Smith, and S. Birchfield, “Handal: A dataset of real-world manipulable object categories with pose annotations, affordances, and reconstructions,” 2023. **2, 3, 4, 5**
- [11] J. Gao, B. Sarkar, F. Xia, T. Xiao, J. Wu, B. Ichter, A. Majumdar, and D. Sadigh, “Physically grounded vision-language models for robotic manipulation,” *arXiv preprint arXiv:2309.02561*, 2023. **2, 3, 4, 5, 6**
- [12] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, “Llama: Open and efficient foundation language models,” *arXiv preprint arXiv:2302.13971*, 2023. **2, 4**
- [13] Z. Lin, C. Liu, R. Zhang, P. Gao, L. Qiu, H. Xiao, H. Qiu, C. Lin, W. Shao, K. Chen *et al.*, “Sphinx: The joint mixing of weights, tasks, and visual embeddings for multi-modal large language models,” *arXiv preprint arXiv:2311.07575*, 2023. **2, 4**
- [14] X. Li, M. Liu, H. Zhang, C. Yu, J. Xu, H. Wu, C. Cheang, Y. Jing, W. Zhang, H. Liu *et al.*, “Vision-language foundation models as effective robot imitators,” *arXiv preprint arXiv:2311.01378*, 2023. **2**
- [15] M. Shridhar, L. Manuelli, and D. Fox, “Cliport: What and where pathways for robotic manipulation,” in *Conference on Robot Learning*. PMLR, 2022, pp. 894–906. **2**
- [16] —, “Perceiver-actor: A multi-task transformer for robotic manipulation,” in *Conference on Robot Learning*. PMLR, 2023, pp. 785–799. **2**
- [17] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, “Learning transferable visual models from natural language supervision,” in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763. **3, 4**
- [18] J. Wu, R. Antonova, A. Kan, M. Lepert, A. Zeng, S. Song, J. Bohg, S. Rusinkiewicz, and T. Funkhouser, “Tidybot: Personalized robot assistance with large language models,” *Autonomous Robots*, vol. 47, no. 8, pp. 1087–1102, 2023. **3**
- [19] W. Cai, S. Huang, G. Cheng, Y. Long, P. Gao, C. Sun, and H. Dong, “Bridging zero-shot object navigation and foundation models through pixel-guided navigation skill,” *arXiv preprint arXiv:2309.10309*, 2023. **3**
- [20] S. Huang, Z. Jiang, H. Dong, Y. Qiao, P. Gao, and H. Li, “Instruct2act: Mapping multi-modality instructions to robotic actions with large language model,” *arXiv preprint arXiv:2305.11176*, 2023. **3**
- [21] Y. J. Ma, W. Liang, G. Wang, D.-A. Huang, O. Bastani, D. Jayaraman, Y. Zhu, L. Fan, and A. Anandkumar, “Eureka: Human-level reward design via coding large language models,” *arXiv preprint arXiv:2310.12931*, 2023. **3**
- [22] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. Florence, and A. Zeng, “Code as policies: Language model programs for embodied control,” *arXiv preprint arXiv:2209.07753*, 2022. **3**
- [23] A. Brohan, Y. Chebotar, C. Finn, K. Hausman, A. Herzog, D. Ho, J. Ibarz, A. Irpan, E. Jang, R. Julian *et al.*, “Do as i can, not as i say: Grounding language in robotic affordances,” in *Conference on Robot Learning*. PMLR, 2023, pp. 287–318. **3**
- [24] D. Driess, F. Xia, M. S. Sajjadi, C. Lynch, A. Chowdhery, B. Ichter, A. Wahid, J. Tompson, Q. Vuong, T. Yu *et al.*, “Palm-e: An embodied multimodal language model,” *arXiv preprint arXiv:2303.03378*, 2023. **3**
- [25] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, X. Chen, K. Choremanski, T. Ding, D. Driess, A. Dubey, C. Finn *et al.*, “Rt-2: Vision-language-action models transfer web knowledge to robotic control,” *arXiv preprint arXiv:2307.15818*, 2023. **3**
- [26] A. Myers, C. L. Teo, C. Fermüller, and Y. Aloimonos, “Affordance detection of tool parts from geometric features,” 2015. **3, 4**
- [27] H. Luo, W. Zhai, J. Zhang, Y. Cao, and D. Tao, “Learning affordance grounding from exocentric images,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 2252–2261. **3, 5, 6**
- [28] S. Kazemzadeh, V. Ordonez, M. Matten, and T. Berg, “Referitgame: Referring to objects in photographs of natural scenes,” in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 787–798. **3, 4, 7**
- [29] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, M. S. Bernstein, and F.-F. Li, “Visual genome: Connecting language and vision using crowdsourced dense image annotations,” 2016. **3**
- [30] K. Chen, Z. Zhang, W. Zeng, R. Zhang, F. Zhu, and R. Zhao, “Shikra: Unleashing multimodal llm’s referential dialogue magic,” *arXiv preprint arXiv:2306.15195*, 2023. **3**
- [31] 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 *International conference on machine learning*. PMLR, 2023, pp. 19730–19742. **4**
- [32] L. Ke, M. Ye, M. Danelljan, Y.-W. Tai, C.-K. Tang, F. Yu *et al.*, “Segment anything in high quality,” *Advances in Neural Information Processing Systems*, vol. 36, 2024. **5**
- [33] X. Lai, Z. Tian, Y. Chen, Y. Li, Y. Yuan, S. Liu, and J. Jia, “Lisa: Reasoning segmentation via large language model,” *arXiv preprint arXiv:2308.00692*, 2023. **5**
- [34] J. Yang, H. Zhang, F. Li, X. Zou, C. Li, and J. Gao, “Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v,” *arXiv preprint arXiv:2310.11441*, 2023. **5, 6**
- [35] G. Li, V. Jampani, D. Sun, and L. Sevilla-Lara, “Locate: Localize and transfer object parts for weakly supervised affordance grounding,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 10922–10931. **5, 6**
- [36] S. Qian and D. F. Fouhey, “Understanding 3d object interaction from a single image,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 21753–21763. **5, 6**
- [37] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang *et al.*, “Sapient: A simulated part-based interactive environment,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11097–11107. **5, 6**
- [38] K. Mo, L. J. Guibas, M. Mukadam, A. Gupta, and S. Tulsiani, “Where2act: From pixels to actions for articulated 3d objects,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 6813–6823. **5, 7**
- [39] B. Eisner, H. Zhang, and D. Held, “Flowbot3d: Learning 3d articulation flow to manipulate articulated objects,” *arXiv preprint arXiv:2205.04382*, 2022. **5, 7**