

AffordGrasp: In-Context Affordance Reasoning for Open-Vocabulary Task-Oriented Grasping in Clutter

Yingbo Tang^{1,3}, Shuaike Zhang^{2,3}, Xiaoshuai Hao^{3,✉,†}, Pengwei Wang³, Jianlong Wu²
Zhongyuan Wang³, Shanghang Zhang^{3,4,✉}

Abstract—Inferring the affordance of an object and grasping it in a task-oriented manner is crucial for robots to successfully complete manipulation tasks. Affordance indicates where and how to grasp an object by taking its functionality into account, serving as the foundation for effective task-oriented grasping. However, current task-oriented methods often depend on extensive training data that is confined to specific tasks and objects, making it difficult to generalize to novel objects and complex scenes. In this paper, we introduce *AffordGrasp*, a novel open-vocabulary grasping framework that leverages the reasoning capabilities of vision-language models (VLMs) for in-context affordance reasoning. Unlike existing methods that rely on explicit task and object specifications, our approach infers tasks directly from implicit user instructions, enabling more intuitive and seamless human-robot interaction in everyday scenarios. Building on the reasoning outcomes, our framework identifies task-relevant objects and grounds their part-level affordances using a visual grounding module. This allows us to generate task-oriented grasp poses precisely within the affordance regions of the object, ensuring both functional and context-aware robotic manipulation. Extensive experiments demonstrate that *AffordGrasp* achieves state-of-the-art performance in both simulation and real-world scenarios, highlighting the effectiveness of our method. We believe our approach advances robotic manipulation techniques and contributes to the broader field of embodied AI. Project website: <https://eqcy.github.io/affordgrasp/>.

I. INTRODUCTION

For robots assisting in daily tasks, understanding object affordances and selecting grasp areas and poses are vital for effective execution. For example, when pouring water, a robot must identify a cup’s handle as the graspable region to ensure success. This highlights the need for task-oriented grasping based on affordances, akin to how humans intuitively interact with objects—like holding a spoon or racket handle—and generalize to new objects. However, enabling robots to replicate this ability is challenging, particularly in accurately identifying affordances without extensive task-specific training. Overcoming this is key to developing robots that adapt to diverse real-world scenarios.

Existing methods [1], [2], [3], [4], [5], [6] transfer affordances [1], [2], [3] or task-oriented grasps [4], [5], [6] from training data to test objects by constructing affordance

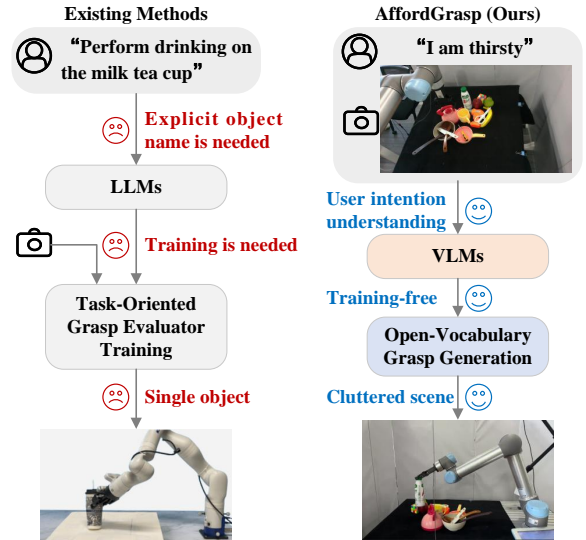


Fig. 1. Compared to existing task-oriented grasping methods, our approach offers three key advantages: (1) leveraging VLMs for affordance reasoning, enabling better understanding of user intentions from implicit language instructions and visual scene observations; (2) a training-free pipeline, eliminating the need for annotated data to train task-oriented grasp evaluators; and (3) the ability to handle cluttered scenes, rather than being limited to single-object grasping.

memories or learning task-specific skills. These approaches retrieve semantically and geometrically similar objects to transfer knowledge from seen to unseen instances. However, their reliance on time-consuming retrieval processes and limited affordance memory scales hinders their effectiveness in open-vocabulary scenarios with novel tasks and objects.

Recent research shows growing interest in open-vocabulary grasping [7], [8], [9], [10]. These methods generate grasp poses by jointly training vision-language grasping frameworks [7], [8], [9] or leveraging the visual grounding abilities of VLMs [10]. However, they focus solely on grasping entire objects, neglecting task-specific affordances. For finer-grained grasping, [11] uses open-vocabulary object detection [12] and part segmentation [13] to perform part-level grasps based on user’s part prompts. Other works [14], [15] employ Large Language Models (LLMs) to reason about object parts or affordances from language instructions, localizing regions via VLMs. However, existing methods often require explicit object names in prompts and struggle with implicit instructions (e.g., “I am thirsty”) without visual scene context, lacking the ability to interpret tasks and object functionalities in complex open-vocabulary environments.

¹Institute of Automation, Chinese Academy of Sciences.

²Harbin Institute of Technology (Shenzhen).

³Beijing Academy of Artificial Intelligence.

⁴State Key Laboratory of Multimedia Information Processing, School of Computer Science, Peking University.

Corresponding authors: ✉ Xiaoshuai Hao (xshao@baai.ac.cn), ✉ Shanghang Zhang (shanghang@pku.edu.cn).

†Project leader.

Although existing task-oriented or open-vocabulary grasping methods perform well in some scenarios, inferring appropriate open-vocabulary task-oriented grasps in cluttered environments remains challenging. We summarize the key difficulties as follows: (1) **Ambiguity in user instructions.** Existing methods often require explicit object names and tasks, lacking flexibility. In practice, user instructions frequently contain implicit goals, demanding deeper semantic understanding. (2) **Training dependency.** Some methods [5], [6] use LLMs to infer object properties and build task-object similarity, yet still rely on annotated data to train task-oriented grasp evaluators for generating desired poses. (3) **Cluttered scenes.** Current task-oriented grasping methods typically focus on single-object scenarios. Cluttered environments with multiple objects increase the complexity of identifying target objects, posing additional challenges.

To address the aforementioned challenges, we propose *AffordGrasp*, a novel affordance reasoning framework for open-vocabulary task-oriented grasping. As illustrated in Fig. 1, our method offers three key advantages over existing approaches. First, we tackle ambiguous user instructions by leveraging Vision-Language Models (VLMs) to reason about affordances from both language and visual inputs. By utilizing the inherent reasoning capabilities of VLMs, we establish relationships among tasks, objects, and affordances in complex scenes. Second, based on the affordance reasoning results, the target object and its visual affordance are grounded using VLMs for grasp generation, eliminating the need for additional training data. Finally, with a deep understanding of tasks and affordances, our method outperforms existing approaches in cluttered scenes, enabling robust and context-aware grasping. The main contributions of our work are as follows:

- We introduce *AffordGrasp*, a novel framework for open-vocabulary task-oriented grasping in cluttered scenes.
- Specifically, we propose an in-context affordance reasoning module that leverages the reasoning capabilities of Vision-Language Models (VLMs). This module extracts explicit tasks from implicit user instructions and infers object affordances from visual observations, significantly enhancing the understanding of tasks, objects, and their affordances through the integration of language and visual inputs.
- Importantly, our method is entirely open-vocabulary, requiring only arbitrary user instructions and RGB-D images. It functions without the need for additional training or fine-tuning, which makes it highly flexible and scalable.
- We validate the effectiveness of our approach through extensive experiments in both simulation and real-world robotic scenarios, demonstrating its robustness and practical applicability.

II. RELATED WORK

A. Visual Affordance Grounding

Visual Affordance Grounding involves identifying and localizing specific areas of an object in an image that enable

potential interactions. Researchers have explored various approaches to learn affordances, including annotated images [16], demonstration videos [17], [1], and cross-modal inputs [18]. Methods like Robo-ABC [2] and RAM [3] employ semantic correspondence and retrieval-based paradigms to generalize affordances across objects and domains, utilizing diverse data sources for zero-shot transfer. Additionally, GLOVER [19] fine-tunes a multi-modal LLM to predict visual affordances. However, these approaches often require varying degrees of data training, limiting their applicability in fully open-vocabulary scenarios.

B. Task-Oriented Grasping

Task-Oriented Grasping focuses on grasping specific object parts based on task constraints, enabling robots to perform subsequent manipulations. Traditional approaches rely on training with large-scale annotated datasets [20], [21]. To enhance generalization to novel objects, methods like GraspGPT [5] and FoundationGrasp [6] leverage the open-ended semantic knowledge of LLMs to bridge connections between dataset objects and novel ones, though they still depend on annotated datasets for training. LERF-TOGO [22] constructs a Language-Embedded Radiance Field (LERF) [23] and generates 3D relevancy heatmaps using DINO features [24] for zero-shot task-oriented grasping. However, it requires time-consuming multi-view image capture for 3D rendering and explicit object part specifications. ShapeGrasp [25] employs geometric decomposition, including attributes and spatial relationships, to grasp novel objects. Despite these advancements, existing methods primarily focus on single-object grasping, limiting their effectiveness in cluttered environments.

C. Robotic Grasping with LLMs and VLMs

The integration of LLMs and VLMs into robotic grasping has recently garnered significant interest. Approaches like LAN-Grasp [14] leverage LLMs to identify graspable object parts, while OVAL-Prompt [15] formulates this as affordance grounding. However, these methods rely exclusively on LLMs to infer grasping positions from language instructions, lacking integration with visual context. Consequently, they depend on explicit instructions and struggle in cluttered environments without advanced reasoning. Recent advancements [26], [27], [28] emphasize the strong contextual reasoning capabilities of VLMs. For example, ThinkGrasp [28] introduces a vision-language grasping system for cluttered scenes. While ThinkGrasp represents progress in clutter handling, it focuses on object-level grasping and overlooks task constraints and affordances, which are essential for fine-grained, task-oriented grasping.

III. METHOD

In this section, we introduce *AffordGrasp*, a novel framework for open-vocabulary task-oriented grasping. As illustrated in Fig. 2, our method begins by interpreting user instructions through a Vision-Language Model (VLM), which decomposes the task into three key steps: analyzing the

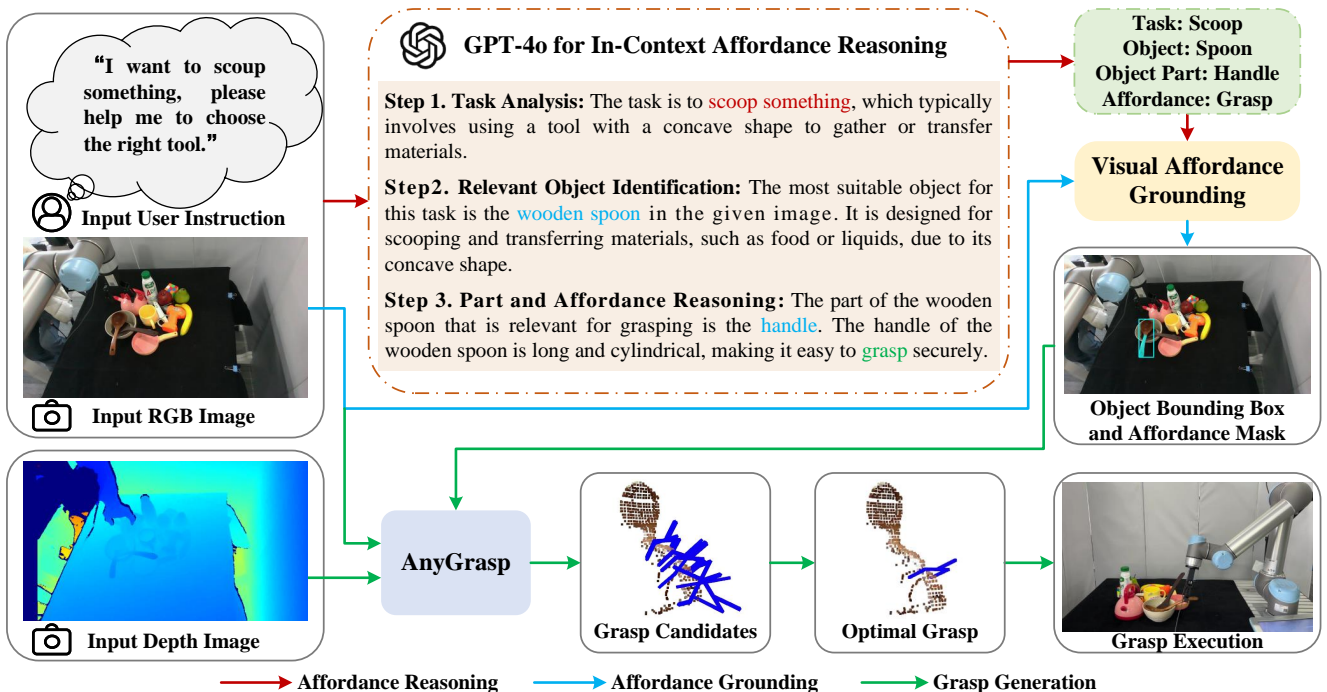


Fig. 2. **Overall Framework of AffordGrasp.** The framework processes user instructions and RGB-D scene observations to achieve open-vocabulary task-oriented grasping in clutter. We leverage GPT-4o [29] for in-context affordance reasoning, decomposing the process into three steps: (1) Extracting the task goal and functional requirements from implicit user instructions (e.g., “I want to scoop something”). (2) Identifying the most task-relevant object in the RGB image (e.g., a wooden spoon). (3) Decomposing the object into functional parts and selecting the optimal graspable part (e.g., the handle) based on its affordances. Based on the reasoning results, a visual affordance grounding module grounds the inferred object and part affordances into pixel-level masks. With the affordance mask and RGB-D images, we employ AnyGrasp [30] to generate task-oriented 6D grasp poses on the target part.

task goal, identifying the most relevant object, and reasoning about the optimal graspable part based on affordances. This reasoning process is grounded into pixel-level masks using a visual grounding module, which segments the object and its functional parts in the RGB image. Leveraging the RGB-D image and affordance mask, we employ AnyGrasp [30] to generate 6D grasp poses, aligning the affordance center to select the optimal grasp. By integrating language-guided reasoning, visual grounding, and geometry-aware grasp synthesis, *AffordGrasp* enables robust and interpretable task-oriented grasping in cluttered environments, effectively bridging high-level task semantics with low-level robotic execution.

A. In-Context Affordance Reasoning

Inferring affordances for task-oriented grasping in cluttered environments is challenging, as it requires interpreting user instructions, identifying task-relevant objects, and reasoning about functional parts and affordances. To address this, we propose an in-context affordance reasoning module powered by Vision-Language Models (VLMs), which decomposes the problem into three key steps: (1) task analysis, (2) relevant object identification, and (3) part and affordance reasoning. Specifically, we leverage the GPT-4o [29] model to analyze the underlying intention and implicit requirements of the task from user instructions, and then establish relationships among tasks, objects, and affordances in an in-context manner using visual observations.

Task Analysis. Given a language instruction \mathcal{L} and an RGB image of the scene \mathcal{I} , we first analyze the user instruction using GPT-4o to extract the explicit task \mathcal{T} .

Relevant Object Identification. Based on the requirements of the explicit task \mathcal{T} , the object \mathcal{O} most relevant to the task is identified in the image.

Part and Affordance Reasoning. The selected object \mathcal{O} is decomposed into functional parts $\mathcal{P} = \{p^1, p^2, \dots, p^n\}$, and the corresponding affordances $\mathcal{A} = \{a^1, a^2, \dots, a^n\}$ of each part are inferred. The optimal part p^* and its affordance a^* are selected based on their suitability for the task. The entire reasoning process can be formulated as follows:

$$\mathcal{T}, \mathcal{O}, p^*, a^* = \text{GPT-4o}(\mathcal{L}, \mathcal{I}). \quad (1)$$

Fig. 3 illustrates the prompt for the in-context affordance reasoning module, including an example of input and output. This prompt guides the Vision-Language Model (VLM) through a step-by-step reasoning process: (1) understanding user intent, (2) identifying the most relevant object, and (3) locating the optimal graspable part based on affordances. This method enhances interpretability and ensures task-oriented grasp generation. Our approach effectively handles complex scenarios with multiple objects, enabling the robot to prioritize relevant items while ignoring distractions. The reasoning output is structured as task, object, object part, and affordance, serving as input for subsequent visual affordance grounding and grasp pose generation.

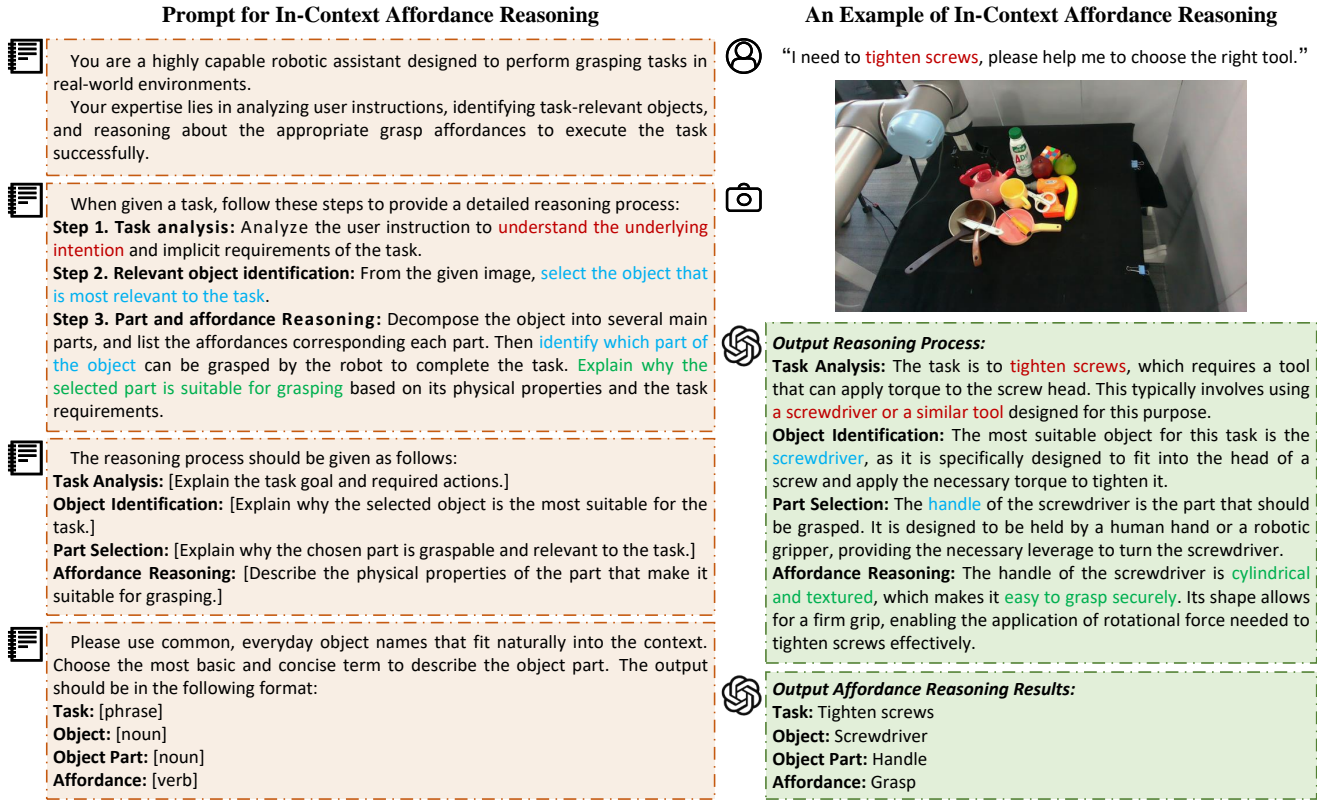


Fig. 3. The prompt for in-context affordance reasoning is designed, and an example is provided for illustration.

B. Visual Affordance Grounding

To ground the object part and corresponding affordance on the task-relevant object, we propose a visual affordance grounding module. Unlike traditional methods trained on closed-set semantic concepts [16], [17], our approach leverages the open-vocabulary part segmentation method VLPart [13], making it open-ended. The module takes as input the RGB image \mathcal{I} and the affordance reasoning results $\{\mathcal{O}, p^*, a^*\}$, and outputs an object bounding box with an affordance mask on the object part.

The process is divided into two steps:

- 1) **Object Localization:** VLPart locates the object bounding box $\mathcal{B}_{\mathcal{O}}$ and generates a masked image $\mathcal{M}_{\mathcal{B}_{\mathcal{O}}}$:

$$\mathcal{B}_{\mathcal{O}} = \text{VLPart}(\mathcal{I}, \mathcal{O}), \quad (2)$$

$$\mathcal{M}_{\mathcal{B}_{\mathcal{O}}}(i, j) = \begin{cases} \mathcal{I}(i, j) & \text{if } (i, j) \in \mathcal{B}_{\mathcal{O}}, \\ 0 & \text{if } (i, j) \notin \mathcal{B}_{\mathcal{O}}. \end{cases} \quad (3)$$

- 2) **Affordance Mask Prediction:** The affordance mask \mathcal{M}_{p^*} is predicted on the masked image using the optimal part p^* and affordance a^* :

$$\mathcal{M}_{p^*} = \text{VLPart}(\mathcal{M}_{\mathcal{B}_{\mathcal{O}}}, p^*, a^*). \quad (4)$$

This two-step approach effectively eliminates interference from other objects in cluttered scenes.

C. Grasp Pose Generation

We employ AnyGrasp [30] to generate 6D grasp poses based on visual affordance. First, the depth image is converted into a partial-view point cloud using camera intrinsics. The affordance mask from the visual affordance grounding module is then used to filter the point cloud, enabling task-oriented grasp generation. This strategy constrains the grasp generation process and effectively avoids interference from surrounding objects.

The point cloud of the object part with grasp affordance is processed by AnyGrasp to generate grasp candidates. AnyGrasp is a 6D grasp generation model trained on a large-scale dataset, which produces grasp poses from partial-view point cloud. Each grasp pose g is represented as:

$$g = [\mathbf{R}, \mathbf{t}, w], \quad (5)$$

where $\mathbf{R} \in \mathbb{R}^{3 \times 3}$ is the rotation matrix, $\mathbf{t} \in \mathbb{R}^{3 \times 1}$ is the translation vector, and $w \in \mathbb{R}$ is the minimum gripper width. The optimal grasp pose g^* is selected from candidates $\mathcal{G} = \{g_1, g_2, \dots, g_n\}$ using:

$$g^* = \arg \max_{g \in \mathcal{G}} \left(\frac{\text{score}(g)}{\|\mathbf{t}(g) - \mathbf{c}\|_2} \right), \quad (6)$$

where $\text{score}(g)$ and $\mathbf{t}(g)$ are the confidence score and translation vector of g , respectively. This prioritizes grasps closer to the affordance mask centroid $\mathbf{c} = [x, y, z]$ with higher confidence scores, ensuring alignment with the affordance's geometric center for improved stability.

TABLE I
SIMULATION RESULTS FOR VARIOUS METHODS IN SINGLE OBJECT GRASPING.

Methods	Cup	Spoon	Hammer	Bowl	Screwdriver	Scissors	Wine Glass	Average GSR
GraspNet [31]	0.68	0.42	0.60	0.84	0.84	0.14	0.44	0.57
ThinkGrasp [28]	0.70	0.92	0.86	0.90	0.94	0.48	0.46	0.75
AffordGrasp (Ours)	0.92	0.98	0.88	0.88	1.00	0.62	0.66	0.85

TABLE II
SIMULATION RESULTS OF VARIOUS METHODS FOR GRASPING IN CLUTTER.

Methods	Cup	Spoon	Hammer	Bowl	Screwdriver	Scissors	Wine Glass	Average GSR
ThinkGrasp [28]	0.70	0.88	0.76	0.96	0.16	0.30	0.00	0.54
AffordGrasp (Ours)	0.84	0.92	0.90	0.94	0.76	0.50	0.52	0.77



Fig. 4. **Simulation Cases of Grasping in Clutter:** The affordances of the target object are indicated with red stars.

IV. EXPERIMENT

A. Experiment Setup

Implementation Details. All experiments are conducted on a workstation with a 24GB GeForce RTX 4090 GPU. The simulation environment is built in PyBullet [32], featuring a UR5 arm with a ROBOTIQ-85 gripper and an Intel RealSense L515 camera. Raw images are resized to 224×224 pixels for segmentation. Our real-world setup includes a UR5 arm with an RS-485 gripper and an Intel RealSense L515 camera, calibrated in an eye-to-hand configuration. GraspNet [31] is used for grasp pose generation in simulations, while AnyGrasp [30] is employed for real-world experiments.

Baseline Methods. We compare our method with GraspNet [31] and AnyGrasp [30], both 6D grasping methods trained on large-scale real-world datasets; RAM [3], which generates 2D affordance maps and lifts them to 3D for

manipulation; and ThinkGrasp [28], a vision-language grasping system leveraging GPT-4o for reasoning in cluttered environments.

Evaluation Metrics. We use Grasp Success Rate (GSR), calculated as the percentage of successful grasps relative to total grasp executions, to evaluate performance in both simulation and real-world experiments.

B. Simulation Results

We compare our method with GraspNet [31] and ThinkGrasp [28] in simulation, using GraspNet for grasp generation. Experiments involve two tasks: **Grasping Single Object** and **Grasping in Clutter**. Each case includes 50 runs with identical scene settings for all methods.

Grasping Single Object involves placing objects in isolation on a table. Results in Tab. I show our method achieves an average grasp success rate (GSR) of 0.85, outperforming GraspNet and ThinkGrasp. For simple objects (*e.g.*, spoon,

TABLE III
COMPARISON RESULTS OF VARIOUS METHODS IN REAL-WORLD EXPERIMENTS.

Methods	Bottle	Mug	Spatula	Spoon	Knife	Pan	Kettle	Screwdriver	Average GSR
RAM [3]	6/10	1/10	0/10	0/10	0/10	0/10	0/10	0/10	0.14
AnyGrasp [30]	8/10	3/10	7/10	6/10	10/10	7/10	2/10	6/10	0.61
ThinkGrasp [28]	8/10	6/10	5/10	7/10	9/10	7/10	7/10	9/10	0.73
AffordGrasp (Ours)	8/10	8/10	9/10	8/10	10/10	9/10	6/10	8/10	0.83

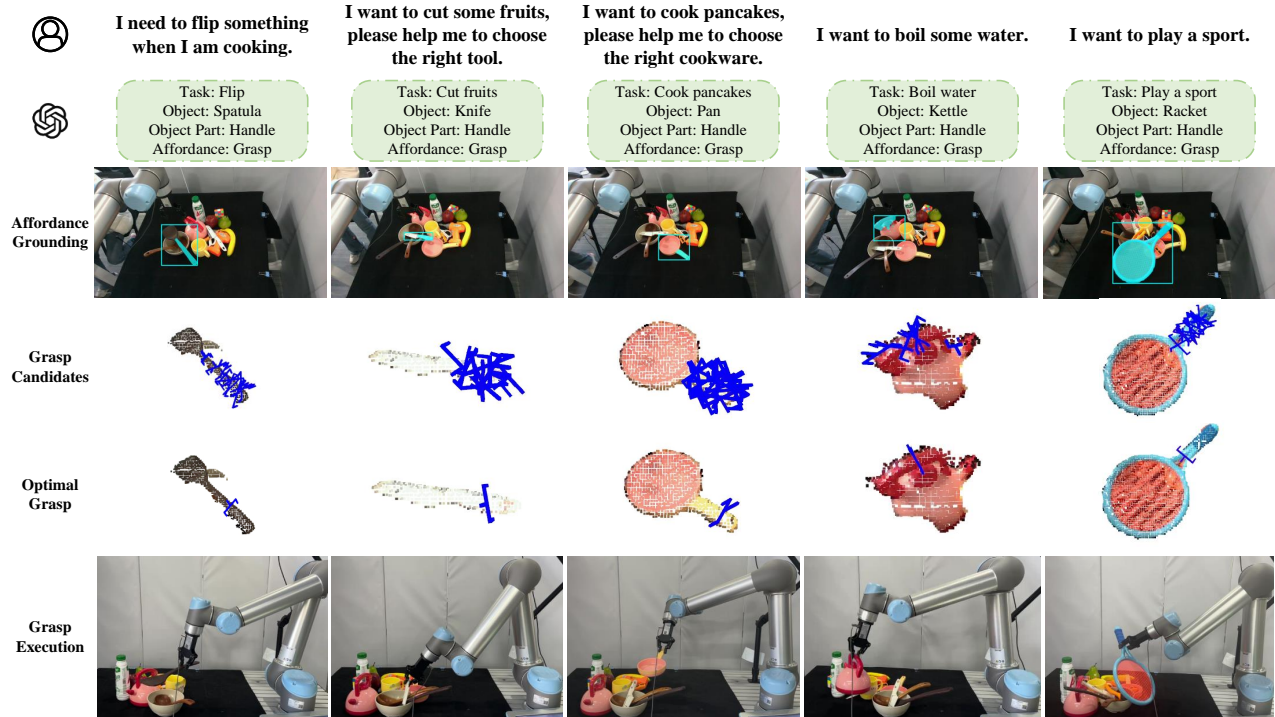


Fig. 5. **Real-world Examples of Grasping in Clutter:** The visualization of affordance grounding and task-oriented grasp generation are provided.

hammer, bowl), our method matches ThinkGrasp’s performance. For challenging objects like scissors and wine glass, our method achieves GSRs of 0.62 and 0.66, surpassing GraspNet by 0.48 and 0.22, and ThinkGrasp by 0.14 and 0.20, respectively. This demonstrates our method’s adaptability to complex geometries.

Grasping in Clutter involves randomly arranged objects, with target affordances marked by red stars (Fig. 4). Quantitative results in Tab. II show our method achieves an average GSR of 0.77, significantly outperforming ThinkGrasp’s 0.54. Our method excels with cups, spoons, hammers, and screwdrivers. For screwdrivers, ThinkGrasp obtains only 0.16 GSR, while our method achieves 0.76. ThinkGrasp also struggles with wine glasses, as it prioritizes the most salient object, leading to failures even when the target is visible.

C. Real-world Results

We evaluate our method on a diverse set of objects in real-world task-oriented grasping in clutter. Results in Tab. III show our method achieves the highest average grasp success rate (GSR) of 0.83, outperforming RAM, AnyGrasp, and

ThinkGrasp. For simple objects (*e.g.*, bottle), all methods achieve high GSRs. However, RAM struggles with complex geometries due to limited affordance memory, while ThinkGrasp shows competitive performance on specific objects like kettles and screwdrivers. Our method consistently excels across a broader range of objects, achieving GSRs of 0.80, 0.90, and 0.90 for mugs, spatulas, and pans, respectively, demonstrating robustness across varying shapes and functionalities.

Visualizations of affordance grounding and task-oriented grasp generation are shown in Fig. 5. Relevant objects and affordance masks are labeled with cyan bounding boxes and masks in RGB images. Grasp poses are generated on point clouds corresponding to affordance masks and visualized on the object point cloud. As shown in Fig. 5, grasp poses are densely distributed on object parts identified with grasp affordances by GPT-4o reasoning. The optimal grasp pose is selected using the filtering strategy detailed in Sec. III C for robotic grasp execution.

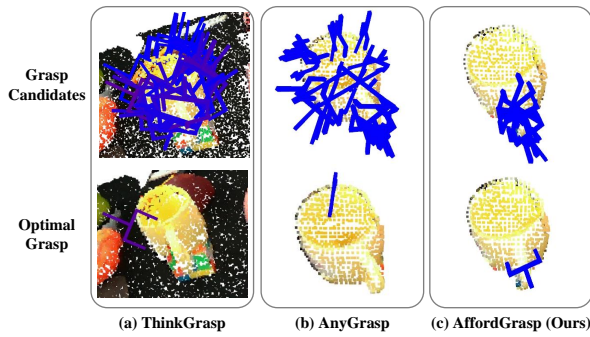


Fig. 6. The case study of mug in real-world experiments.

D. Case Study

To demonstrate the advantages of our method, we present a case study in Fig. 6. The first row shows grasp pose candidates, while the second row depicts the optimal grasp poses selected by different methods. ThinkGrasp selects grasps near the center of the cropped image due to its 3×3 grid strategy, which divides the image into numbered regions (1: top-left, 9: bottom-right). However, this often fails for objects with curved rims (e.g., mugs), as center-based grasps lack stability due to insufficient contact points or improper force distribution. Similarly, AnyGrasp generates arbitrary grasps without affordance reasoning, leading to unreliable task-oriented grasping (Fig. 6(b) and (c)). In contrast, our method integrates affordance reasoning and fine-grained grasp generation, producing precise and stable grasps. This highlights the importance of affordance reasoning and demonstrates the effectiveness of our method in real-world applications.

V. CONCLUSIONS

This paper introduces *AffordGrasp*, a novel open-vocabulary task-oriented grasping framework that leverages the reasoning capabilities of VLMs to infer object affordances and generate task-specific grasps in cluttered environments. By integrating in-context affordance reasoning and visual affordance grounding, our method effectively interprets user instructions, identifies task-relevant objects, and grounds part-level affordances for precise grasp generation. Extensive experiments in simulation and real-world scenarios demonstrate the effectiveness of our approach, outperforming existing methods in handling novel objects and complex scenes without requiring additional training data.

REFERENCES

- [1] S. Bahl, R. Mendonca *et al.*, “Affordances from human videos as a versatile representation for robotics,” in *CVPR*, 2023, pp. 13 778–13 790.
- [2] Y. Ju, K. Hu *et al.*, “Robo-abc: Affordance generalization beyond categories via semantic correspondence for robot manipulation,” in *ECCV*, 2024, pp. 222–239.
- [3] Y. Kuang, J. Ye *et al.*, “Ram: Retrieval-based affordance transfer for generalizable zero-shot robotic manipulation,” *arXiv preprint arXiv:2407.04689*, 2024.
- [4] W. Dong, D. Huang *et al.*, “Rtagrasp: Learning task-oriented grasping from human videos via retrieval, transfer, and alignment,” *arXiv preprint arXiv:2409.16033*, 2024.
- [5] C. Tang, D. Huang *et al.*, “Graspnet: Leveraging semantic knowledge from a large language model for task-oriented grasping,” *IEEE Robotics and Automation Letters*, 2023.

- [6] C. Tang, D. Huang, W. Dong, R. Xu, and H. Zhang, “Foundationgrasp: Generalizable task-oriented grasping with foundation models,” *arXiv preprint arXiv:2404.10399*, 2024.
- [7] M. Li, Q. Zhao, S. Lyu, C. Wang, Y. Ma, G. Cheng, and C. Yang, “Ovgnet: A unified visual-linguistic framework for open-vocabulary robotic grasping,” in *IROS*, 2024, pp. 7507–7513.
- [8] Y. Zheng, X. Chen, Y. Zheng *et al.*, “Gaussiangrasper: 3d language gaussian splatting for open-vocabulary robotic grasping,” *arXiv preprint arXiv:2403.09637*, 2024.
- [9] V. Bhat, P. Krishnamurthy *et al.*, “Hifi-cs: Towards open vocabulary visual grounding for robotic grasping using vision-language models,” *arXiv preprint arXiv:2409.10419*, 2024.
- [10] G. Tziafas and H. Kasaei, “Towards open-world grasping with large vision-language models,” *arXiv preprint arXiv:2406.18722*, 2024.
- [11] T. van Oort, D. Miller, W. N. Browne, N. Marticorena, J. Haviland, and N. Suenderhauf, “Open-vocabulary part-based grasping,” *arXiv preprint arXiv:2406.05951*, 2024.
- [12] S. Liu, Z. Zeng *et al.*, “Grounding dino: Marrying dino with grounded pre-training for open-set object detection,” in *ECCV*, 2024, pp. 38–55.
- [13] P. Sun, S. Chen *et al.*, “Going denser with open-vocabulary part segmentation,” in *ICCV*, 2023, pp. 15 453–15 465.
- [14] R. Mirjalili, M. Krawez, S. Silenzi, Y. Blei, and W. Burgard, “Langrasp: An effective approach to semantic object grasping using large language models,” in *First Workshop on Vision-Language Models for Navigation and Manipulation at ICRA 2024*.
- [15] E. Tong, A. Pipari, S. Lewis, Z. Zeng, and O. C. Jenkins, “Oval-prompt: Open-vocabulary affordance localization for robot manipulation through llm affordance-grounding,” *arXiv preprint arXiv:2404.11000*, 2024.
- [16] H. Luo, W. Zhai *et al.*, “Learning affordance grounding from exocentric images,” in *CVPR*, 2022, pp. 2252–2261.
- [17] J. Chen, D. Gao *et al.*, “Affordance grounding from demonstration video to target image,” in *CVPR*, 2023, pp. 6799–6808.
- [18] Y. Li, N. Zhao, J. Xiao *et al.*, “Laso: Language-guided affordance segmentation on 3d object,” in *CVPR*, 2024, pp. 14 251–14 260.
- [19] T. Ma, Z. Wang, J. Zhou, M. Wang, and J. Liang, “Glover: Generalizable open-vocabulary affordance reasoning for task-oriented grasping,” *arXiv preprint arXiv:2411.12286*, 2024.
- [20] A. Murali, W. Liu *et al.*, “Same object, different grasps: Data and semantic knowledge for task-oriented grasping,” in *CORL*, 2021, pp. 1540–1557.
- [21] Y. Song, P. Sun *et al.*, “Learning 6-dof fine-grained grasp detection based on part affordance grounding,” *arXiv preprint arXiv:2301.11564*, 2023.
- [22] A. Rashid, S. Sharma *et al.*, “Language embedded radiance fields for zero-shot task-oriented grasping,” in *7th Annual Conference on Robot Learning*, 2023.
- [23] J. Kerr, C. M. Kim *et al.*, “Lerf: Language embedded radiance fields,” in *ICCV*, 2023, pp. 19 729–19 739.
- [24] M. Caron, H. Touvron *et al.*, “Emerging properties in self-supervised vision transformers,” in *ICCV*, 2021, pp. 9650–9660.
- [25] S. Li, S. Bhagat, J. Campbell, Y. Xie, W. Kim, K. Sycara, and S. Stepputtis, “Shapegrasp: Zero-shot task-oriented grasping with large language models through geometric decomposition,” *arXiv preprint arXiv:2403.18062*, 2024.
- [26] S. Jin, J. Xu, Y. Lei, and L. Zhang, “Reasoning grasping via multi-modal large language model,” *arXiv preprint arXiv:2402.06798*, 2024.
- [27] J. Xu, S. Jin *et al.*, “Rt-grasp: Reasoning tuning robotic grasping via multi-modal large language model,” in *IROS*, 2024, pp. 7323–7330.
- [28] Y. Qian, X. Zhu *et al.*, “Thinkgrasp: A vision-language system for strategic part grasping in clutter,” *arXiv preprint arXiv:2407.11298*, 2024.
- [29] A. Hurst, A. Lerer *et al.*, “Gpt-4o system card,” *arXiv preprint arXiv:2410.21276*, 2024.
- [30] H.-S. Fang, C. Wang *et al.*, “Anygrasp: Robust and efficient grasp perception in spatial and temporal domains,” *IEEE Transactions on Robotics*, pp. 3929–3945, 2023.
- [31] H.-S. Fang, C. Wang, M. Gou, and C. Lu, “Graspnet-1billion: A large-scale benchmark for general object grasping,” in *CVPR*, 2020, pp. 11 444–11 453.
- [32] E. Coumans and Y. Bai, “Pybullet, a python module for physics simulation for games, robotics and machine learning,” 2016.