

# Zero-Shot Robotic Manipulation via 3D Gaussian Splatting-Enhanced Multimodal Retrieval-Augmented Generation

Zilong Xie<sup>1\*</sup>, Jingyu Gong<sup>1, 2, 3\*</sup>, Xin Tan<sup>1, 2</sup>, Zhizhong Zhang<sup>1, 3†</sup>, Yuan Xie<sup>1, 2</sup>

<sup>1</sup>School of Computer Science and Technology, East China Normal University, Shanghai, China

<sup>2</sup>Chongqing Key Laboratory of Precision Optics, Chongqing Institute of East China Normal University, Chongqing, China

<sup>3</sup>Shanghai Key Laboratory of Computer Software Evaluating and Testing, Shanghai, China

zilong6037@gmail.com, {jygong, xtan, zzzhang, yxie}@cs.ecnu.edu.cn

## Abstract

Existing end-to-end approaches of robotic manipulation often lack generalization to unseen objects or tasks due to limited data and poor interpretability. While recent Multimodal Large Language Models (MLLMs) demonstrate strong commonsense reasoning, they struggle with geometric and spatial understanding required for pose prediction. In this paper, we propose RobMRAG, a 3D Gaussian Splatting-Enhanced Multimodal Retrieval-Augmented Generation (MRAG) framework for zero-shot robotic manipulation. Specifically, we construct a multi-source manipulation knowledge base containing object contact frames, task completion frames, and pose parameters. During inference, a Hierarchical Multimodal Retrieval module first employs a three-priority hybrid retrieval strategy to find task-relevant object prototypes, then selects the geometrically closest reference example based on pixel-level similarity and Instance Matching Distance (IMD). We further introduce a 3D-Aware Pose Refinement module based on 3D Gaussian Splatting into the MRAG framework, which aligns the pose of the reference object to the target object in 3D space. The aligned results are reprojected onto the image plane and used as input to the MLLM to enhance the generation of the final pose parameters. Extensive experiments show that on a test set containing 30 categories of household objects, our method improves the success rate by 7.76% compared to the best-performing zero-shot baseline under the same setting, and by 6.54% compared to the state-of-the-art supervised baseline. Our results validate that RobMRAG effectively bridges the gap between high-level semantic reasoning and low-level geometric execution, enabling robotic systems that generalize to unseen objects while remaining inherently interpretable.

**Code** — <https://github.com/XieZilongAI/RobMRAG>

## Introduction

Robotic manipulation fundamentally requires precise interaction with diverse objects, where accurate prediction of contact points and 3D poses is critical for task success. Without reliable estimation of object poses, robots are often forced into trial-and-error adjustments, resulting in in-

efficiency and failure (Zhang et al. 2024). Therefore, robots must possess strong reasoning and generalization capabilities to accurately predict low-level action parameters and adapt to dynamic environments.

Although existing works (Geng et al. 2023a,b,c) have explored end-to-end learning for grasp pose prediction, their generalization remains limited due to constrained data scale and the black-box nature of deep models, which lack commonsense reasoning capabilities. Recent advances in Multimodal Large Language Models (MLLMs) (Zhang et al. 2023; Li et al. 2022, 2024a) have demonstrated promising performance in vision-language understanding and cross-modal reasoning. Several efforts (Huang et al. 2023; Zitkovich et al. 2023; Li et al. 2024b) utilize MLLMs for high-level instruction generation. For example, ManipLLM (Li et al. 2024b) employs chain-of-thought prompting and multi-task finetuning to guide category-level tasks. However, current MLLMs still lack sufficient understanding of geometric structure and spatial layout, limiting their effectiveness in precise grasp pose prediction, as shown in Figure 1(a).

Moreover, zero-shot generalization is a key capability for building general-purpose agents, enabling robots to manipulate unseen objects and adapt to varying embodiments and environments without task-specific training. Existing methods often rely on expert demonstrations (Vuong et al. 2023; Khazatsky et al. 2024) or transfer from human interaction data such as HOI (Liu et al. 2022; Grauman et al. 2022; Luo et al. 2022) and internet videos (Chen et al. 2025). RAM (Kuang et al. 2024) performs zero-shot manipulation via cross-domain affordance retrieval, yet its 2D-to-3D lifting pipeline is brittle under viewpoint or geometric shifts, and the MLLM’s reasoning remains outside the geometric core, limiting pose precision in complex scenes.

To address these challenges, we propose RobMRAG, a 3D Gaussian Splatting-Enhanced Multimodal Retrieval-Augmented Generation (MRAG) framework for zero-shot robotic manipulation (see Figure 1(b)). First, a multi-source manipulation knowledge base is built from simulation, robotic dataset, and Internet data, containing rich multimodal information such as object contact frames, task-completion frames, and contact poses. During inference, a hierarchical multimodal retrieval module first uses a three-priority hybrid retrieval strategy to find task-relevant op-

\*These authors contributed equally.

†Corresponding authors

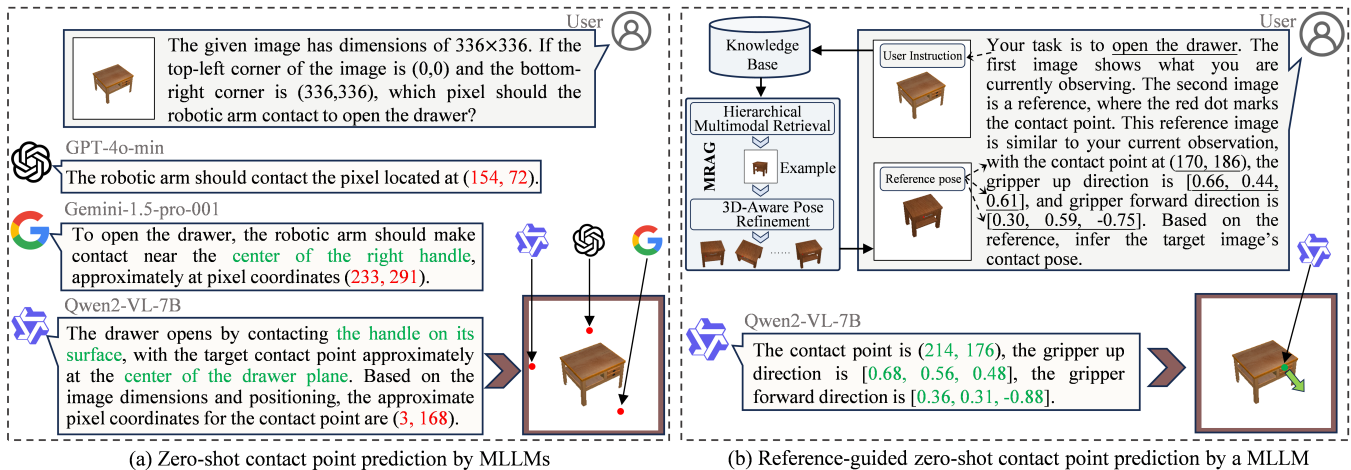


Figure 1: (a) Although multimodal large language models (MLLMs) possess the commonsense knowledge that opening a drawer requires grasping the handle and can roughly localize it, the predicted 2D contact points consistently fall outside the table itself, let alone on the handle. (b) We incorporate contact pose prediction using a multimodal retrieval-augmented generation (MRAG) framework. By retrieving relevant manipulation examples, our method enables more accurate inference of the contact pose on the target object.

eration prototypes. Then, pixel-level cosine similarity and Instance Matching Distance (IMD) identify the closest geometric reference. Within the MRAG framework, a 3D Gaussian Splatting-based pose refinement module is introduced to perform rigid transformations on retrieved reference poses for precise 3D alignment with target objects. The aligned poses are then reprojected onto the 2D image plane and fed into the MLLM to guide accurate manipulation pose generation. Extensive experiments show that RobMRAG outperforms state-of-the-art methods under both zero-shot and supervised settings, effectively bridging task understanding and concrete operation for open-world robotic manipulation.

The key contributions of this work are summarized as follows:

- We propose a Multimodal Retrieval-Augmented Generation (MRAG) framework for zero-shot robotic manipulation, which enables manipulation on unseen objects through a multi-source knowledge base.
- We integrate a 3D-Aware Pose Refinement module into the MRAG framework, enabling precise pose alignment between reference and target objects, thereby enhancing the geometric consistency of the retrieved results.
- Experimental results demonstrate that, on a test set comprising 30 categories of household objects, the proposed method achieves a 7.76% improvement in success rate compared with the SOTA zero-shot baseline, and a 6.54% improvement compared with the SOTA supervised baseline.

## Related Works

### Zero-shot Robotic Manipulation

Zero-shot robotic manipulation enables robots to perform novel tasks without task-specific training. Existing meth-

ods fall into two categories. The first attempts to directly learn transferable manipulation policies from human videos (Bharadhwaj et al. 2024; Chang, Prakash, and Gupta 2023; Bharadhwaj et al. 2023; Grauman et al. 2022; Xu et al. 2023; Chen et al. 2025), either through imitation learning by extracting intermediate representations such as hand trajectories and poses, or by constructing reward functions for reinforcement learning. However, these methods often rely on manually collected video datasets (Smith et al. 2020; Xiong et al. 2021) or require online fine-tuning (Bahl, Gupta, and Pathak 2022), which limits their applicability in real-world scenarios. The second category builds databases of diverse manipulation demonstrations (Nguyen et al. 2022; Ju et al. 2024; Kuang et al. 2024), retrieving demonstrations similar to the current task and adapting them to robot execution. Yet, the effectiveness of such methods heavily depends on the coverage of the demonstration database, which limits their generalization to novel tasks or unseen scenes.

In contrast, our method explicitly bridges the inherent gap between task-level intent and instance-level geometric parameters by constructing a multimodal manipulation knowledge base and integrating Hierarchical Multimodal Retrieval with a 3D-Aware Pose Refinement module.

### Multimodal Large Language Models for Robotics

Multimodal Large Language Models (MLLMs) unify visual perception, language understanding, and action planning into a single framework (Zhang et al. 2023; Li et al. 2022). Models like RT-2 (Zitkovich et al. 2023) and RoboMamba (Liu et al. 2024) leverage large-scale pretraining to translate natural language into executable control commands. Recent works (Zhen et al. 2024; Kim et al. 2024; Yue et al. 2024) adopt Transformer-based architectures for grounding semantics to actions. To improve adaptability in dynamic environments, while VoxPoser (Huang et al. 2023)

introduces language-driven affordance reasoning, dynamically constructing 3D semantic value maps to guide manipulation. Beyond task-specific designs, foundation models such as PaLM-E (Driess et al. 2023) and Flamingo (Alayrac et al. 2022) further extend cross-modal reasoning capabilities. ManipLLM (Li et al. 2024b) further enhances generalization and stability in grasp pose prediction by incorporating adapters and chain-of-thought reasoning mechanisms.

Building upon this foundation, we combine multimodal retrieval-augmented generation with MLLMs, significantly improving zero-shot generalization in robotic manipulation tasks.

## Multimodal Retrieval-Augmented Generation

Retrieval-Augmented Generation (RAG) mitigates limitations of Large Language Models (LLMs) like outdated knowledge and hallucination by leveraging external memory (Lewis et al. 2020; Guu et al. 2020). Extending to multimodal domains, Multimodal RAG (MRAG), which integrates both visual and textual inputs, has become a prominent research direction (Chen et al. 2022; Ma et al. 2024; Bonomo and Bianco 2025; Liu et al. 2025b; Yu et al. 2025; Liu et al. 2025c). Representative MRAG approaches, such as MuRAG (Chen et al. 2022), M2RAG (Ma et al. 2024), and MRAMG (Yu et al. 2025), typically enhance the quality of generation in tasks like visual question answering by retrieving image-text pairs from external memory, overcoming conventional RAG’s limitations in visual understanding. Meanwhile, the incorporation of graph neural networks and knowledge graphs (Dong et al. 2024; Edge et al. 2024; Guo et al. 2024; Liu et al. 2025a; Wu et al. 2025), exemplified by GraphRAG (Edge et al. 2024) and LightRAG (Guo et al. 2024), captures complex cross-modal relations for better semantic reasoning.

In this work, we integrate MRAG techniques with MLLMs in the domain of robotic manipulation, enabling high-precision prediction of operational poses.

## Method

Our goal is to develop a framework for zero-shot robotic manipulation that can accurately and efficiently respond to user instructions in unseen environments. To this end, as illustrated in Figure 2, we present RobMRAG, a 3D Gaussian Splatting-Enhanced Multimodal Retrieval-Augmented Generation framework for zero-shot robotic manipulation. In the following sections, we elaborate on the details of each component of RobMRAG.

### Multi-Source Knowledge Base Construction

We construct a multimodal manipulation knowledge base that integrates data from simulation, real-world or synthetic robotic dataset, and the Internet. Each manipulation case contains one or two key action frames: the object contact frame  $\mathcal{F}_{\text{ctc}}$ , capturing the moment when the robot (or human) correctly makes contact with the object and initiates the manipulation, and the task-completion frame  $\mathcal{F}_{\text{suc}}$ , showing the final successful state. It also includes a corresponding instruction text  $I_e$  describing the task. Specifically, for sim-

ulation data, we directly obtain precise 6D grasp parameters: the 2D contact point  $p_e = (x, y)$  and end-effector directions  $D_e = (d_u, d_f)$ , where the gripper up direction  $d_u = (x_u, y_u, z_u)$ . These data are collected from multiple successful executions across various object categories and manipulation types.

Real-world data are primarily sourced from human-object interaction video datasets (e.g. HOI4D (Liu et al. 2022)), where 2D contact points and motion directions are extracted via hand-keypoint detection and annotated as arrows on the contact-frame images. Internet data offer a more diverse range of manipulation scenarios, including examples extracted from instructional videos and animations. Although these cases often lack accurate 3D parameters, we apply a semi-automatic annotation pipeline to generate 2D grasp sketches. Ultimately, we construct a comprehensive multi-source manipulation knowledge base  $\mathcal{K} = \{I_e, \mathcal{F}_{\text{ctc}}, \mathcal{F}_{\text{suc}}, p_e, D_e\}$ , which furnishes rich cross-domain priors for subsequent retrieval augmentation.

### Hierarchical Multimodal Retrieval

Our framework employs a hierarchical retrieval strategy to identify the optimal reference example from the knowledge base. This process cascades from high-level textual semantics down to visual similarity and, finally, to fine-grained geometric matching, ensuring a precise instruction-to-example alignment.

**Textual Semantic Retrieval.** The first layer operates on the textual modality, utilizing a three-priority hybrid strategy to find semantically relevant candidates. Given a natural language instruction  $I = \{w_i\}_{i=1}^L$ , this strategy first executes its top priority: using sparse retrieval (BM25) to find manipulation examples with matching object names, primarily targeting our high-quality simulation data source. If no direct match is found (e.g., for a novel object category not in  $\mathcal{K}$ ), the system defaults to its second priority: dense retrieval based on semantic embeddings. At this stage, the instruction text is projected into a dense vector by a pre-trained language encoder, and the cosine similarity  $S_{\text{text}}$  between semantic embeddings is then computed:

$$S_{\text{text}}(I, I_e) = \frac{\hat{I} \cdot \hat{I}_e}{\|\hat{I}\| \|\hat{I}_e\|}. \quad (1)$$

Here,  $\hat{I} = \text{Enc}(I)$  and  $\hat{I}_e = \text{Enc}(I_e)$  are the semantic embeddings of the instruction and an example text, respectively. As a final fallback, should the top similarity score fall below a threshold  $\tau_{\text{den}}$ , the third priority is triggered: the retrieval scope is expanded to our broader robotic dataset and Internet data sources. This textual retrieval stage yields a candidate set of examples that are conceptually aligned with the given task.

**Visual Similarity Filtering.** Once the candidate set is obtained, the second layer refines it based on coarse visual similarity. The system computes the cosine similarity  $S_{\text{CLIP}}$  between the current observation image  $\mathcal{F}_{\text{obs}}$  and each candidate’s contact frame  $\mathcal{F}_{\text{ctc}}$  using a CLIP image encoder, selecting the top- $n$  visually similar samples:

$$S_{\text{CLIP}}(\mathcal{F}_{\text{obs}}, \mathcal{F}_{\text{ctc}}) = \frac{f_{\text{obs}} \cdot f_{\text{ctc}}}{\|f_{\text{obs}}\| \|f_{\text{ctc}}\|} \quad (2)$$

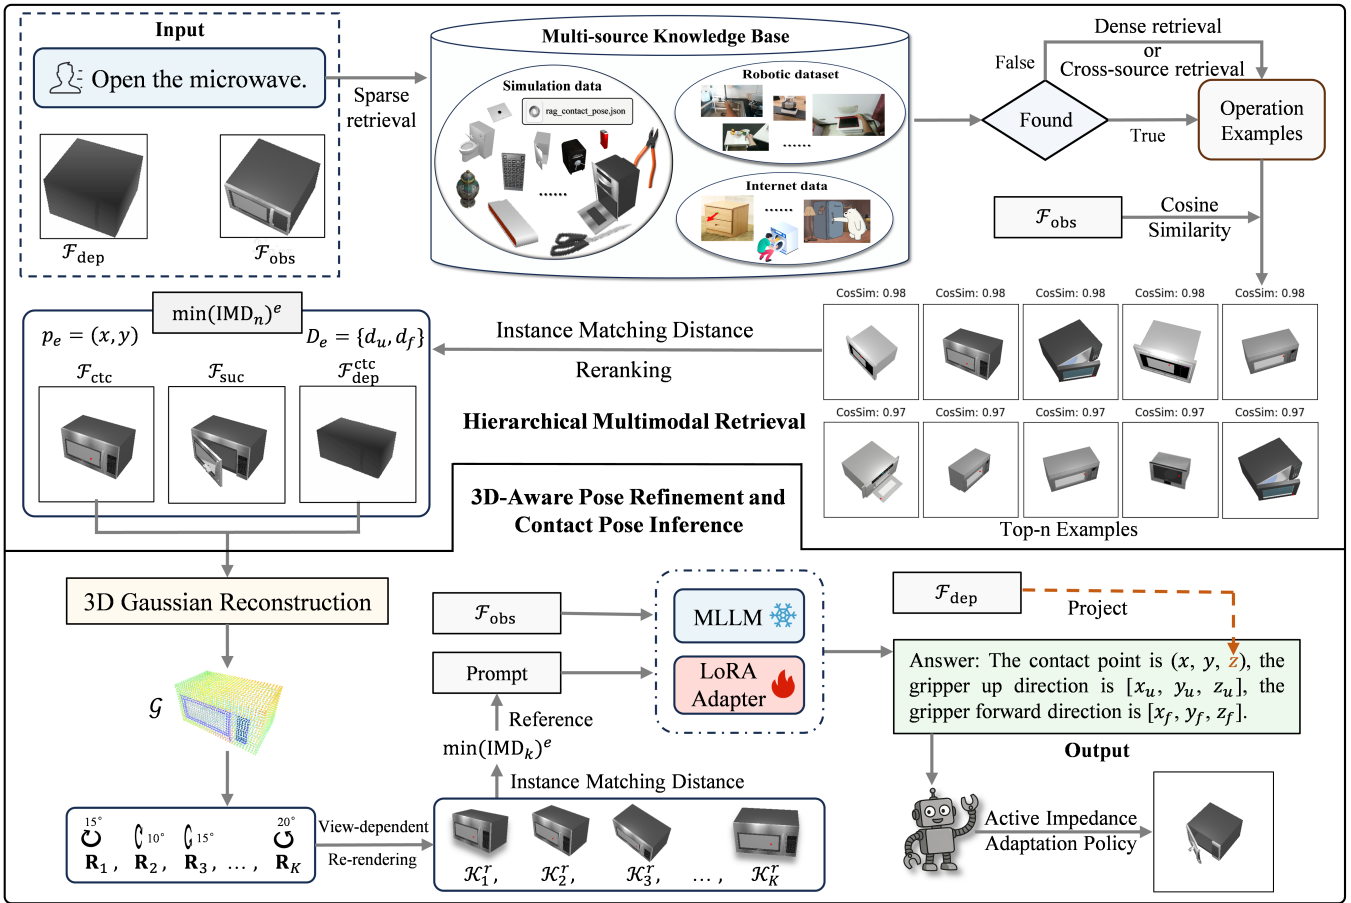


Figure 2: Overview of the RobMRAG: a 3D Gaussian splatting-enhanced multimodal retrieval-augmented generation framework.

where  $f_{\text{obs}}$  and  $f_{\text{ctc}}$  are the visual feature vectors extracted by CLIP from the  $\mathcal{F}_{\text{obs}}$  and  $\mathcal{F}_{\text{ctc}}$ , respectively. This step effectively filters out candidates that, while semantically related, are visually distinct from the current scene.

**Geometric Matching.** In the final layer, we perform fine-grained geometric matching on the remaining candidates to find the best reference prototype. We use Instance Matching Distance (IMD) to measure the geometric discrepancy between the observed object and each candidate by accounting for local feature consistency:

$$\text{IMD}(\mathcal{F}_{\text{obs}}, \mathcal{F}_{\text{ctc}}, \mathcal{M}_{\text{obs}}) = \sum_{p \in \mathcal{M}_{\text{obs}}} \|\mathbf{F}^{\text{obs}}(p) - \text{NN}(\mathbf{F}^{\text{obs}}(p), \mathbf{F}^{\text{ctc}})\|_2 \quad (3)$$

where  $\mathcal{M}_{\text{obs}}$  is the observed instance mask,  $\mathbf{F}(p)$  is the dense feature vector at pixel  $p$ , and  $\text{NN}(\cdot)$  finds the nearest neighbor match. The example with the minimum IMD is selected. If this score is below an empirically-determined threshold  $\tau_{\text{IMD}}$ , it is used directly as the final reference. Otherwise, it serves as a geometric prior for the subsequent pose refinement module.

### 3D-Aware Pose Refinement

For reference examples requiring further alignment, our framework employs a 3D-aware pose refinement module. We first generate a 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023; Gong et al. 2021a,b) representation from the reference’s RGB-D image using a pre-trained generative model TRELIS (Xiang et al. 2025). This allows us to re-render the object from the single input viewpoint. The core formulation expresses the object surface through a differentiable collection of Gaussian distributions  $\mathcal{G} = \{(\mu_i, \Sigma_i, c_i)\}_{i=1}^N$  where  $\mu_i \in \mathbb{R}^3$  denotes the Gaussian center position,  $\Sigma_i \in \mathbb{R}^{3 \times 3}$  controls the spatial distribution, and  $c_i$  represents the color attributes.

To eliminate the viewpoint discrepancy, our method applies a predefined set of small-angle rotational transformations  $\{\mathbf{R}_{\mathbf{k}}\}_{\mathbf{k}=1}^K$  to the reference grasp pose. This strategy is sufficient as the preceding IMD filtering has already ensured a high degree of initial geometric alignment. We first project the 2D reference contact point  $p_e$  to 3D space  $p'_e$  using depth information, then transform both the 3D reference contact point and orientation:

$$p'_k = \mathbf{R}_{\mathbf{k}} p'_e, \quad D'_k = D_e \otimes \text{Quat}(\mathbf{R}_{\mathbf{k}}) \quad (4)$$

where  $\otimes$  denotes quaternion multiplication. The differentiable rendering of 3DGS generates corresponding contact frames  $\{\mathcal{F}_{\text{ctc}}^r\}$  for each candidate pose after reprojecting the 3D contact points to the 2D image plane:

$$p_k^r = \pi(\mathbf{M}[\mathbf{R}_k | t] p_k') \quad (5)$$

where  $\pi(\cdot)$  is the camera projection function,  $\mathbf{M}$  is the intrinsic matrix, and  $t$  denotes the translational compensation. Finally, we obtain the re-rendered reference set  $\{\mathcal{K}_k^r\}_{k=1}^K$  and retain the pose whose frames yield the lowest  $\text{IMD}_k$  as the final reference inputs for the MLLM.

### Optional LoRA Fine-Tuning and Inference

The framework employs a joint loss function combining Masked Language Modeling (MLM) with pose regression tasks, achieving precise grasp pose prediction through Low-Rank Adaptation (LoRA) of the multimodal large language model.

MLM randomly masks grasp parameters in input text, forcing the model to infer masked numerical characters from context. The loss function is defined as character-level cross-entropy at masked positions:

$$\mathcal{L}_{\text{MLM}} = - \sum_{i \in \mathcal{M}} \sum_{c \in \mathcal{C}_i} \log P(c | w_{\setminus \mathcal{M}}, \mathcal{F}_{\text{obs}}, \mathcal{F}_{\text{ctc}}) \quad (6)$$

where  $\mathcal{M}$  denotes the set of masked positions,  $\mathcal{C}_i$  represents the character sequence at the  $i$ -th masked position,  $w_{\setminus \mathcal{M}}$  indicates the unmasked text context, and  $\mathcal{F}_{\text{obs}}/\mathcal{F}_{\text{ctc}}$  are the observation and reference images respectively.

The fine-tuning process supervises structured outputs with contact point  $(x, y)$  constrained by mean squared error and end-effector directions  $(d_u, d_f)$  regularized through cosine similarity:

$$\mathcal{L}_{\text{pose}} = \lambda_1 \|(x, y) - (\hat{x}, \hat{y})\|_2^2 + \lambda_2 (2 - d_u \cdot \hat{d}_u - d_f \cdot \hat{d}_f) \quad (7)$$

where  $\lambda_1, \lambda_2$  are balance weights,  $\hat{d}_u$  and  $\hat{d}_f$  denote predicted direction vectors. The total loss combines both components:  $\mathcal{L} = \mathcal{L}_{\text{MLM}} + \mathcal{L}_{\text{pose}}$ .

During inference, the framework first retrieves the most task-relevant reference through Hierarchical Multimodal Retrieval and 3D-Aware Pose Refinement module, which are then provided as additional inputs to the MLLM for generating the operational pose. The final output adopts a structured format containing normalized 2D contact point and two 3D end-effector directions (representing gripper up and forward directions). 2D contact point is projected to 3D manipulation space via depth map  $\mathcal{F}_{\text{dep}}$ , constructing complete end-effector poses. We apply active impedance adaptation policy (Li et al. 2024b) to adjust movement direction.

## Experiments

### Datasets and Evaluation Metrics

We construct interactive manipulation environments based on the SAPIEN simulator and the PartNetMobility dataset (Xiang et al. 2020), which provides part-level motion annotations for a wide range of object categories. Following the baseline setup (Li et al. 2024b), simulated manipulation is performed using a Franka Panda robotic arm

equipped with a suction gripper, rendered through the VulkanRenderer engine. To generate training data, we randomly sample contact points on movable parts and use the inverse surface normal as the initial operation direction, focusing on pull-type actions to ensure consistency between the motion and manipulation directions. In total, we collect approximately 20,000 successful manipulation samples across 20 object categories as our training set. Following the same procedure, we generate the test set, which is further divided into a Test Seen Split, containing objects of seen categories but with different instances, and a Test Unseen Split, containing entirely novel object categories.

The primary evaluation metric is manipulation average success rate (ASR) (Li et al. 2024b), defined as the proportion of successful trials exceeding a displacement threshold  $\delta$ . We use two levels:  $\delta = 0.01$  for verifying initial pose prediction, and  $\delta = 0.1$  for assessing sustained motion. Active impedance control is applied to adapt motion direction in response to interaction uncertainties.

### Competitor Methods

To validate our Multimodal Retrieval-Augmented Generation (MRAG) framework, we compare it against six robotic manipulation methods under identical train/test splits and end-effector configurations (suction grippers substituted for original parallel grippers in baselines). We compare our method with the following baselines: Where2Act (Mo et al. 2021) predicts pixel-level affordances for generalization; UMPNet (Xu, He, and Song 2022) generates 6DoF actions from monocular input via temporal causality; Flowbot3D (Eisner, Zhang, and Held 2022) infers 3D motion fields from point clouds for planning; Implicit3D (Zhong et al. 2023) extends Transporter to 3D using spatio-temporal keypoints; RAM (Kuang et al. 2024) is a zero-shot framework lifting retrieved 2D affordances to 3D actions; ManipLLM (Li et al. 2024b) uses MLLM chain-of-thought for planning and pose prediction; CrayonRobo (Li et al. 2025) is a prompt-driven VLA model using 2D visual prompts on images to predict SE(3) actions.

For evaluation, we consider the following settings: *Zero-shot*: Directly applies the method without task-specific training or fine-tuning. *All*: After fine-tuning, the model is evaluated on the complete test set to assess its generalization capability in real-world scenarios. *Local*: After fine-tuning, evaluation is performed only on samples where the contact point is correctly predicted, representing the theoretical upper bound given accurate contact point localization.

### Comparative Analysis

As shown in Table 1, our framework’s efficacy is first demonstrated in the *zero-shot* setting, where RobMRAG achieves an average success rate of 43.53% on the unseen split, surpassing the RAM baseline by 7.76 percentage points. This highlights the strength of our framework, which combines the advantages of MLLMs and MRAG. This design enables effective utilization of external knowledge bases, allowing the model to generalize to novel objects without task-specific training, as seen in high-performing tasks like Table (82.75%). In the fine-tuned *All* setting, our

Method	Test Seen Split															
RAM (Kuang et al. 2024) (Zero-shot) <sup>†</sup>	43.14	25.67	14.00	35.29	52.94	35.49	49.02	45.39	52.00	49.23	45.76	36.02	24.00	61.58	24.00	72.00
CrayonRobo (Li et al. 2025) (Zero-shot) <sup>†</sup>	17.65	21.59	28.00	47.06	62.18	33.33	57.45	58.54	48.00	49.19	37.24	57.39	26.00	73.33	12.00	68.00
Ours (Zero-shot)	62.75	47.45	20.00	66.67	50.98	17.65	59.18	72.05	64.00	43.14	49.02	38.00	22.00	67.85	18.00	78.00
Where2Act (Mo et al. 2021)	26.00	36.00	19.00	27.00	23.00	11.00	15.00	47.00	14.00	24.00	13.00	12.00	56.00	68.00	7.00	40.00
UMPNet (Xu, He, and Song 2022)	46.00	43.00	15.00	28.00	54.00	32.00	28.00	56.00	44.00	40.00	10.00	23.00	18.00	54.00	20.00	42.00
Flowbot3D (Eisner, Zhang, and Held 2022)	67.00	55.00	20.00	32.00	27.00	31.00	61.00	68.00	15.00	28.00	36.00	18.00	21.00	70.00	18.00	26.00
Implicit3D (Zhong et al. 2023)	53.00	58.00	35.00	55.00	28.00	<b>66.00</b>	58.00	51.00	52.00	57.00	45.00	34.00	41.00	54.00	39.00	43.00
ManipLLM (Li et al. 2024b)	<b>68.00</b>	<b>64.00</b>	<b>36.00</b>	77.00	43.00	62.00	65.00	61.00	65.00	52.00	53.00	40.00	<b>64.00</b>	71.00	<b>60.00</b>	64.00
ManipLLM (All) <sup>†</sup>	60.78	51.57	46.00	52.94	37.25	33.33	63.51	62.56	46.00	43.14	25.49	54.00	42.00	65.83	22.00	78.00
Ours (All)	66.67	51.56	22.00	<b>82.27</b>	<b>96.08</b>	18.73	<b>82.27</b>	<b>87.12</b>	<b>66.00</b>	<b>65.85</b>	<b>75.62</b>	<b>56.32</b>	32.00	<b>77.21</b>	22.00	<b>96.00</b>
ManipLLM (Local) <sup>†</sup>	72.73	67.89	47.83	63.41	44.96	64.00	75.37	82.43	63.16	54.29	26.00	84.62	63.64	72.07	29.04	80.80
Ours (Local)	74.07	66.01	24.29	85.05	98.00	21.86	100.0	100.0	87.24	72.12	88.24	100.0	71.43	82.12	31.25	97.96

Method	Test Seen Split					Test Unseen Split										
					AVG											
RAM (Kuang et al. 2024) (Zero-shot) <sup>†</sup>	32.45	52.64	21.61	38.93	40.56	42.38	25.60	42.83	56.07	30.47	38.00	32.00	43.06	18.37	28.96	35.77
CrayonRobo (Li et al. 2025) (Zero-shot) <sup>†</sup>	47.06	54.90	16.33	51.02	43.31	48.67	29.92	31.67	58.36	18.68	24.00	38.00	41.18	15.69	27.45	33.36
Ours (Zero-shot)	48.98	76.14	15.69	67.47	49.25	60.78	25.49	<b>82.75</b>	68.38	28.69	18.00	40.00	64.71	13.73	32.76	43.53
Where2Act (Mo et al. 2021)	13.00	18.00	13.00	40.00	26.00	18.00	<b>35.00</b>	38.00	28.00	5.00	21.00	17.00	20.00	15.00	15.00	21.00
UMPNet (Xu, He, and Song 2022)	22.00	33.00	26.00	64.00	35.00	42.00	20.00	35.00	42.00	29.00	20.00	26.00	28.00	25.00	15.00	28.00
Flowbot3D (Eisner, Zhang, and Held 2022)	17.00	53.00	29.00	42.00	37.00	23.00	10.00	60.00	39.00	27.00	42.00	28.00	51.00	13.00	23.00	32.00
Implicit3D (Zhong et al. 2023)	27.00	65.00	20.00	33.00	46.00	45.00	17.00	80.00	53.00	15.00	69.00	41.00	31.00	<b>30.00</b>	31.00	41.00
ManipLLM (Li et al. 2024b)	41.00	75.00	<b>44.00</b>	67.00	56.00	38.00	22.00	81.00	<b>86.00</b>	38.00	<b>85.00</b>	42.00	<b>83.00</b>	26.00	38.00	51.00
ManipLLM (All) <sup>†</sup>	24.49	64.78	31.37	58.92	48.20	29.41	17.65	80.27	78.51	29.03	72.00	28.00	66.67	15.69	36.20	45.34
Ours (All)	<b>72.92</b>	<b>83.51</b>	20.84	<b>68.06</b>	<b>62.15</b>	<b>65.21</b>	17.84	79.45	78.21	<b>56.09</b>	42.00	<b>86.00</b>	67.84	25.65	<b>57.13</b>	<b>57.54</b>
ManipLLM (Local) <sup>†</sup>	35.90	73.36	53.85	69.74	61.25	41.67	19.09	87.50	85.28	37.34	85.29	43.33	76.67	20.77	40.35	53.73
Ours (Local)	97.22	92.30	48.78	71.14	75.45	72.22	20.31	83.33	86.02	65.68	78.78	95.54	78.15	28.57	66.41	67.50

Table 1: Comparisons of our method against baseline methods, <sup>†</sup> indicates reproduced results (ASR, %).

model’s advantage becomes more pronounced: it achieves a 57.54% average SR on unseen categories, which is 6.54 percentage points higher than the strongest baseline, ManipLLM. On certain tasks, our model reaches exceptionally high success rates, such as Oven (86.00% vs. ManipLLM’s 42.00%) and Laptop (96.08% vs. 43.00%). In the *Local* evaluation setting, where a perfect contact point is provided, our pose generation process still demonstrates superior robustness. Specifically, we observe average success rate improvements of 14.20% on the seen split and 13.77% on the unseen split. These gains are attributed to our 3D-Aware Pose Refinement module, which aligns retrieved examples in 3D space using Gaussian Splatting and reprojects them to provide the MLLM with more accurate geometric context. This enriched spatial information leads to more precise 6-DoF pose estimation.

Despite the strong overall performance, our analysis also identifies challenges. The first stems from the inherent properties of the evaluation dataset itself. For instance, some tasks exhibit low success rates due to severely degraded visual input; certain Door images are reduced to a few feature lines with nearly 80% pixel loss, making robust feature extraction nearly impossible. Similarly, tasks like Faucet (25.65% SR) suffer from ambiguous instructions in the dataset, where the prompt does not distinguish between "lever adjustment" and "spout repositioning." These dataset-level issues present a fundamental challenge to any vision-language-based method. The second challenge, however, reflects a genuine limitation of our current framework: manipulation of objects with minuscule interaction regions. The low SR on tasks requiring sub-centimeter precision, such as Scissors (17.84%) and Pliers (32.00%), in-

Ablation Method	Test seen	Test unseen
<i>Main Component Ablation</i>		
w/o Retrieval (only vision input)	19.46	19.01
w/ Textual Retrieval + CosSim only	37.05	34.62
w/ Textual Retrieval + IMD only	45.55	41.27
w/ Textual Retrieval + CosSim + IMD	47.34	42.08
<b>Full MRAG (w/ 3D Pose Align)</b>	<b>49.25</b>	<b>43.53</b>
<i>Different MLLM Backbones</i>		
Qwen2-VL-7B-Instruct (base)	62.15	57.54
LLaMA3.2-11B-Vision-Instruct	62.93	58.18
Qwen2.5-VL-7B-Instruct	<b>65.52</b>	<b>60.63</b>

Table 2: Ablation studies on different components of our MRAG framework (ASR, %).

dicates that our model still struggles with the fine-grained physical reasoning and exacting grasping point prediction necessary for such millimeter-scale operations.

### Ablation Studies

We conduct a series of ablation studies to evaluate the contributions of different components in the MRAG framework. As shown in Table 2, the first block presents the zero-shot performance of the MLLM without utilizing our proposed framework. We observe that using only visual input without incorporating any additional signals leads to poor performance (19.46% seen / 19.01% unseen). Introducing multimodal retrieval based on either cosine similarity or Instance Matching Distance (IMD) significantly improves the results,




Figure 3: Ablation study on the number of retrieval candidates (Top-n) and rotation samples ( $K$ ).

and combining both further enhances generalization capability. Building on this, we further incorporate the 3D-Aware Pose Refinement module, which yields the best zero-shot results (49.25% seen / 43.53% unseen), highlighting the crucial role of spatial grounding in improving retrieval performance.

We then evaluate the framework’s performance with different MLLM backbones in the fine-tuned *All* setting. As shown in the table, leveraging more advanced models consistently enhances results, demonstrating our framework’s scalability. Specifically, upgrading the backbone from the base Qwen2-VL-7B-Instruct to Qwen2.5-VL-7B-Instruct achieves the overall peak performance of 65.52% on the seen split and 60.63% on the unseen split. Per-category success rates are detailed in the Appendix. These results clearly indicate that our RobMRAG architecture effectively capitalizes on the enhanced reasoning and visual understanding capabilities of more powerful MLLMs, establishing a robust and scalable foundation for robotic manipulation.

We conducted controlled ablation studies on the retrieval candidate number (top-n) and rotation sampling count ( $K$ ) separately on the test seen split, as shown in Figure 3. Experiments reveal: peak success rate of 62.15% occurs at top-n=5, with larger values degrading performance by 0.61–1.48% due to geometric noise; optimal  $K=6$  achieves the best performance, validating small-angle rotation compensation. Both parameters exhibit an initial increase followed by a decline, demonstrating the necessity to balance search and sampling breadth with geometric consistency. Excessive values reduce success rates and increase computational overhead at the same time.

### Qualitative Results

Figure 4 presents qualitative examples demonstrating the robustness of our three-priority hybrid retrieval strategy. The left example showcases the second-priority fallback: for the target object *Scissors*, which is absent from our simulation dataset, the system uses dense retrieval to successfully match *Pliers* as a functionally analogous counterpart for affordance reference. The right example illustrates the third-priority case: given a *Toilet* with no direct or semantically close match in simulation, the strategy expands the search to robotic and Internet datasets to retrieve a valid reference.

Figure 5 illustrates the geometric alignment process for




Figure 4: Visualization of the three-priority hybrid retrieval strategy: ① sparse retrieval (simulation), ② dense retrieval (simulation), and ③ cross-source retrieval (robotic/Internet).




Figure 5: Visualization of 3D-Aware Pose Refinement.

reference poses using 3D-Aware Pose Refinement module. Due to viewpoint discrepancies between the retrieved reference pose and the observed target object, we employ 3D Gaussian Splatting to generate candidate views by applying multiple rotational transformations. The IMD is then recalculated for each candidate, and the pose with the minimal IMD is selected to achieve precise alignment.

### Conclusion

We propose a novel robot manipulation framework based on Multimodal Retrieval-Augmented Generation (MRAG). By integrating a multi-source manipulation knowledge base, hierarchical multimodal retrieval, and a 3D-aware pose refinement module, our approach improves the average manipulation success rates by 7.76% over state-of-the-art baselines in a zero-shot setting across various objects. Furthermore, with LoRA fine-tuning, the performance improves by up to 6.54% compared to the baseline. While our method demonstrates strong capabilities in semantic understanding and pose adaptation, challenges remain in achieving millimeter-level precision and complex spatial reasoning. Future work will focus on enhancing fine-grained physical reasoning and spatial understanding.

## Acknowledgements

This work is supported by the National Natural Science Foundation of China (Grant No. 62176092, 62222602, 62302167, U23A20343, 62476090, 62502159), Natural Science Foundation of Shanghai (Grant No. 25ZR1402135), Shanghai Sailing Program (Grant No. 23YF1410500), Young Elite Scientists Sponsorship Program by CAST (Grant No. YESS20240780), the Chenguang Program of Shanghai Education Development Foundation and Shanghai Municipal Education Commission (Grant No. 23CGA34), Natural Science Foundation of Chongqing (Grant No. CSTB2023NSCQ-JQX0007, CSTB2023NSCQ-MSX0137, CSTB2025NSCQ-GPX0445), Open Project Program of the State Key Laboratory of CAD&CG (Grant No. A2501), Zhejiang University, Open Research Fund of Key Laboratory of Advanced Theory and Application in Statistics and Data Science-MOE, ECNU.

## References

- Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; et al. 2022. Flamingo: a visual language model for few-shot learning. *NeurIPS*, 35: 23716–23736.
- Bahl, S.; Gupta, A.; and Pathak, D. 2022. Human-to-robot imitation in the wild. *arXiv preprint arXiv:2207.09450*.
- Bharadhwaj, H.; Gupta, A.; Kumar, V.; and Tulsiani, S. 2024. Towards generalizable zero-shot manipulation via translating human interaction plans. In *ICRA*, 6904–6911. IEEE.
- Bharadhwaj, H.; Gupta, A.; Tulsiani, S.; and Kumar, V. 2023. Zero-shot robot manipulation from passive human videos. *arXiv preprint arXiv:2302.02011*.
- Bonomo, M.; and Bianco, S. 2025. Visual RAG: Expanding MLLM visual knowledge without fine-tuning. *arXiv preprint arXiv:2501.10834*.
- Chang, M.; Prakash, A.; and Gupta, S. 2023. Look ma, no hands! agent-environment factorization of egocentric videos. *NeurIPS*, 36: 21466–21486.
- Chen, H.; Sun, B.; Zhang, A.; Pollefeys, M.; and Leutenegger, S. 2025. VidBot: Learning Generalizable 3D Actions from In-the-Wild 2D Human Videos for Zero-Shot Robotic Manipulation. In *CVPR*, 27661–27672.
- Chen, W.; Hu, H.; Chen, X.; Verga, P.; and Cohen, W. W. 2022. MuRAG: Multimodal Retrieval-Augmented Generator for Open Question Answering over Images and Text. In *EMNLP*, 5558–5570.
- Dong, Y.; Wang, S.; Zheng, H.; Chen, J.; Zhang, Z.; and Wang, C. 2024. Advanced RAG Models with Graph Structures: Optimizing Complex Knowledge Reasoning and Text Generation. In *ISCEIC*, 626–630. IEEE.
- Driess, D.; Xia, F.; Sajjadi, M. S.; Lynch, C.; Chowdhery, A.; Wahid, A.; Tompson, J.; Vuong, Q.; Yu, T.; Huang, W.; et al. 2023. Palm-e: An embodied multimodal language model. *openreview*.
- Edge, D.; Trinh, H.; Cheng, N.; Bradley, J.; Chao, A.; Mody, A.; Truitt, S.; Metropolitansky, D.; Ness, R. O.; and Larson, J. 2024. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*.
- Eisner, B.; Zhang, H.; and Held, D. 2022. FlowBot3D: Learning 3D Articulation Flow to Manipulate Articulated Objects. In *RSS*.
- Geng, H.; Li, Z.; Geng, Y.; Chen, J.; Dong, H.; and Wang, H. 2023a. Partmanip: Learning cross-category generalizable part manipulation policy from point cloud observations. In *CVPR*, 2978–2988.
- Geng, H.; Xu, H.; Zhao, C.; Xu, C.; Yi, L.; Huang, S.; and Wang, H. 2023b. Gapartnet: Cross-category domain-generalizable object perception and manipulation via generalizable and actionable parts. In *CVPR*, 7081–7091.
- Geng, Y.; An, B.; Geng, H.; Chen, Y.; Yang, Y.; and Dong, H. 2023c. Rlafford: End-to-end affordance learning for robotic manipulation. In *ICRA*, 5880–5886. IEEE.
- Gong, J.; Xu, J.; Tan, X.; Song, H.; Qu, Y.; Xie, Y.; and Ma, L. 2021a. Omni-supervised point cloud segmentation via gradual receptive field component reasoning. In *CVPR*, 11673–11682.
- Gong, J.; Xu, J.; Tan, X.; Zhou, J.; Qu, Y.; Xie, Y.; and Ma, L. 2021b. Boundary-aware geometric encoding for semantic segmentation of point clouds. In *AAAI*, volume 35, 1424–1432.
- Grauman, K.; Westbury, A.; Byrne, E.; Chavis, Z.; Furnari, A.; Girdhar, R.; Hamburger, J.; Jiang, H.; Liu, M.; Liu, X.; et al. 2022. Ego4d: Around the world in 3,000 hours of egocentric video. In *CVPR*, 18995–19012.
- Guo, Z.; Xia, L.; Yu, Y.; Ao, T.; and Huang, C. 2024. Lightrag: Simple and fast retrieval-augmented generation. *openreview*.
- Guu, K.; Lee, K.; Tung, Z.; Pasupat, P.; and Chang, M. 2020. Retrieval augmented language model pre-training. In *ICML*, 3929–3938. PMLR.
- Huang, W.; Wang, C.; Zhang, R.; Li, Y.; Wu, J.; and Fei-Fei, L. 2023. Voxposer: Composable 3d value maps for robotic manipulation with language models. In *CoRL*.
- Ju, Y.; Hu, K.; Zhang, G.; Zhang, G.; Jiang, M.; and Xu, H. 2024. Robo-abc: Affordance generalization beyond categories via semantic correspondence for robot manipulation. In *ECCV*, 222–239. Springer.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3D Gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4): 139–1.
- Khazatsky, A.; Pertsch, K.; Nair, S.; Balakrishna, A.; Dasari, S.; Karamcheti, S.; Nasiriany, S.; Srirama, M. K.; Chen, L. Y.; Ellis, K.; et al. 2024. Droid: A large-scale in-the-wild robot manipulation dataset. In *RSS*.
- Kim, M. J.; Pertsch, K.; Karamcheti, S.; Xiao, T.; Balakrishna, A.; Nair, S.; Rafailov, R.; Foster, E.; Lam, G.; Sanketi, P.; et al. 2024. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*.

- Kuang, Y.; Ye, J.; Geng, H.; Mao, J.; Deng, C.; Guibas, L.; Wang, H.; and Wang, Y. 2024. Ram: Retrieval-based affordance transfer for generalizable zero-shot robotic manipulation. *arXiv preprint arXiv:2407.04689*.
- Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.-t.; Rocktäschel, T.; et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *NeurIPS*, 33: 9459–9474.
- Li, B.; Ge, Y.; Ge, Y.; Wang, G.; Wang, R.; Zhang, R.; and Shan, Y. 2024a. Seed-bench: Benchmarking multimodal large language models. In *CVPR*, 13299–13308.
- Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 12888–12900. PMLR.
- Li, X.; Xu, J.; Zhang, M.; Liu, J.; Shen, Y.; Ponomarenko, I.; Xu, J.; Heng, L.; Huang, S.; Zhang, S.; et al. 2025. Object-Centric Prompt-Driven Vision-Language-Action Model for Robotic Manipulation. In *CVPR*, 27638–27648.
- Li, X.; Zhang, M.; Geng, Y.; Geng, H.; Long, Y.; Shen, Y.; Zhang, R.; Liu, J.; and Dong, H. 2024b. Manipllm: Embodied multimodal large language model for object-centric robotic manipulation. In *CVPR*, 18061–18070.
- Liu, J.; Liu, M.; Wang, Z.; Lee, L.; Zhou, K.; An, P.; Yang, S.; Zhang, R.; Guo, Y.; and Zhang, S. 2024. Robomamba: Multimodal state space model for efficient robot reasoning and manipulation. *arXiv preprint arXiv:2406.04339*.
- Liu, J.; Meng, S.; Gao, Y.; Mao, S.; Cai, P.; Yan, G.; Chen, Y.; Bian, Z.; Shi, B.; and Wang, D. 2025a. Aligning Vision to Language: Text-Free Multimodal Knowledge Graph Construction for Enhanced LLMs Reasoning. *arXiv preprint arXiv:2503.12972*.
- Liu, J.; Tao, Y.; Wang, F.; Li, H.; and Qin, X. 2025b. SiQA: A Large Multi-Modal Question Answering Model for Structured Images Based on RAG. In *ICASSP*, 1–5. IEEE.
- Liu, P.; Liu, X.; Yao, R.; Liu, J.; Meng, S.; Wang, D.; and Ma, J. 2025c. Hm-rag: Hierarchical multi-agent multimodal retrieval augmented generation. In *ACM MM*, 2781–2790.
- Liu, Y.; Liu, Y.; Jiang, C.; Lyu, K.; Wan, W.; Shen, H.; Liang, B.; Fu, Z.; Wang, H.; and Yi, L. 2022. Hoi4d: A 4d egocentric dataset for category-level human-object interaction. In *CVPR*, 21013–21022.
- Luo, H.; Zhai, W.; Zhang, J.; Cao, Y.; and Tao, D. 2022. Learning affordance grounding from exocentric images. In *CVPR*, 2252–2261.
- Ma, Z.-A.; Lan, T.; Tu, R.-C.; Hu, Y.; Huang, H.; and Mao, X.-L. 2024. Multi-modal retrieval augmented multi-modal generation: A benchmark, evaluate metrics and strong baselines. *arXiv preprint arXiv:2411.16365*.
- Mo, K.; Guibas, L. J.; Mukadam, M.; Gupta, A.; and Tulsiani, S. 2021. Where2act: From pixels to actions for articulated 3d objects. In *ICCV*, 6813–6823.
- Nguyen, T.; Gopalan, N.; Patel, R.; Corsaro, M.; Pavlick, E.; and Tellex, S. 2022. Affordance-based robot object retrieval. *Autonomous Robots*, 46(1): 83–98.
- Smith, L.; Dhawan, N.; Zhang, M.; Abbeel, P.; and Levine, S. 2020. Avid: Learning multi-stage tasks via pixel-level translation of human videos. In *RSS*.
- Vuong, Q.; Levine, S.; Walke, H. R.; Pertsch, K.; Singh, A.; Doshi, R.; Xu, C.; Luo, J.; Tan, L.; Shah, D.; et al. 2023. Open x-embodiment: Robotic learning datasets and rt-x models. In *CoRL*.
- Wu, J.; Zhu, J.; Qi, Y.; Chen, J.; Xu, M.; Menolascina, F.; and Grau, V. 2025. Medical graph rag: Evidence-based Medical Large Language Model via Graph Retrieval-Augmented Generation. In *ACL*, 28443–28467.
- Xiang, F.; Qin, Y.; Mo, K.; Xia, Y.; Zhu, H.; Liu, F.; Liu, M.; Jiang, H.; Yuan, Y.; Wang, H.; et al. 2020. Sapien: A simulated part-based interactive environment. In *CVPR*, 11097–11107.
- Xiang, J.; Lv, Z.; Xu, S.; Deng, Y.; Wang, R.; Zhang, B.; Chen, D.; Tong, X.; and Yang, J. 2025. Structured 3d latents for scalable and versatile 3d generation. In *CVPR*, 21469–21480.
- Xiong, H.; Li, Q.; Chen, Y.-C.; Bharadhwaj, H.; Sinha, S.; and Garg, A. 2021. Learning by watching: Physical imitation of manipulation skills from human videos. In *IROS*, 7827–7834. IEEE.
- Xu, M.; Xu, Z.; Chi, C.; Veloso, M.; and Song, S. 2023. Xskill: Cross embodiment skill discovery. In *CoRL*, 3536–3555. PMLR.
- Xu, Z.; He, Z.; and Song, S. 2022. Universal manipulation policy network for articulated objects. *IEEE RA-L*, 7(2): 2447–2454.
- Yu, Q.; Xiao, Z.; Li, B.; Wang, Z.; Chen, C.; and Zhang, W. 2025. MRAMG-Bench: A BeyondText Benchmark for Multimodal Retrieval-Augmented Multimodal Generation. *arXiv preprint arXiv:2502.04176*.
- Yue, Y.; Wang, Y.; Kang, B.; Han, Y.; Wang, S.; Song, S.; Feng, J.; and Huang, G. 2024. Deer-vla: Dynamic inference of multimodal large language models for efficient robot execution. *NeurIPS*, 37: 56619–56643.
- Zhang, K.; Ren, P.; Lin, B.; Lin, J.; Ma, S.; Xu, H.; and Liang, X. 2024. Pivot-r: Primitive-driven waypoint-aware world model for robotic manipulation. *NeurIPS*, 37: 54105–54136.
- Zhang, R.; Han, J.; Liu, C.; Gao, P.; Zhou, A.; Hu, X.; Yan, S.; Lu, P.; Li, H.; and Qiao, Y. 2023. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*.
- Zhen, H.; Qiu, X.; Chen, P.; Yang, J.; Yan, X.; Du, Y.; Hong, Y.; and Gan, C. 2024. 3d-vla: A 3d vision-language-action generative world model. *arXiv preprint arXiv:2403.09631*.
- Zhong, C.; Zheng, Y.; Zheng, Y.; Zhao, H.; Yi, L.; Mu, X.; Wang, L.; Li, P.; Zhou, G.; Yang, C.; et al. 2023. 3d implicit transporter for temporally consistent keypoint discovery. In *ICCV*, 3869–3880.
- Zitkovich, B.; Yu, T.; Xu, S.; Xu, P.; Xiao, T.; Xia, F.; Wu, J.; Wohlhart, P.; Welker, S.; Wahid, A.; et al. 2023. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *CoRL*, 2165–2183. PMLR.