

# GeoLanG: Geometry-Aware Language-Guided Grasping with Unified RGB-D Multimodal Learning

Rui Tang<sup>1\*</sup>, Guankun Wang<sup>1\*</sup>, Long Bai<sup>1</sup>, Huxin Gao<sup>1</sup>, Jiewen Lai<sup>1</sup>, Chi Kit Ng<sup>1</sup>,  
Jiazheng Wang<sup>2</sup>, Fan Zhang<sup>2</sup>, Hongliang Ren<sup>†1,3</sup>

**Abstract**—Language-guided grasping has emerged as a promising paradigm for enabling robots to identify and manipulate target objects through natural language instructions, yet it remains highly challenging in cluttered or occluded scenes. Existing methods often rely on multi-stage pipelines that separate object perception and grasping, which leads to limited cross-modal fusion, redundant computation, and poor generalization in cluttered, occluded, or low-texture scenes. To address these limitations, we propose GeoLanG, an end-to-end multi-task framework built upon the CLIP architecture that unifies visual and linguistic inputs into a shared representation space for robust semantic alignment and improved generalization. To enhance target discrimination under occlusion and low-texture conditions, we explore a more effective use of depth information through the Depth-guided Geometric Module (DGGM), which converts depth into explicit geometric priors and injects them into the attention mechanism without additional computational overhead. In addition, we propose Adaptive Dense Channel Integration, which adaptively balances the contributions of multi-layer features to produce more discriminative and generalizable visual representations. Extensive experiments on the OCID-VLG dataset, as well as in both simulation and real-world hardware, demonstrate that GeoLanG enables precise and robust language-guided grasping in complex, cluttered environments, paving the way toward more reliable multimodal robotic manipulation in real-world human-centric settings. Our code and dataset are publicly available at <https://github.com/Tomry1114/GeoLanG>.

## I. INTRODUCTON

In recent years, robots have been increasingly deployed in applications such as household assistance, elderly care, and warehouse logistics [1]–[5]. As shown in Fig. 1, a central challenge in this field is enabling robots to perform adaptive, explainable, and reliable grasping in open-world environments, where multiple objects may occlude or interfere with each other. Traditional RGB-based grasping pipelines typically separate object detection, segmentation, and grasp planning [6], [7]. While effective in controlled settings, these approaches often exhibit limited semantic understanding,

\* Equal contribution

This work was supported in part by Ministry of Science and Technology (MOST) of China Key Project 2025YFE0122500, Guangdong Basic and Applied Basic Research Foundation (Grant No. 2025A1515011594); National Natural Science Foundation of China (Grant No. 62403402); NSFC Distinguished Young Scientists Fund – Category A (Grant No. T252500134); Hong Kong Research Grants Council (Grant Nos. C4026-21GF, R4020-22, 14200425, 14206125, 14204524, 14203323). (Corresponding to: H. Ren, hren@ee.cuhk.edu.hk.)

<sup>1</sup> The Chinese University of Hong Kong, Hong Kong SAR, China.

<sup>2</sup> The Theory Lab, Central Research Institute, 2012 Labs, Huawei Technologies Co. Ltd., Hong Kong SAR, China.

<sup>3</sup> Shenzhen Loop Area Institute, Shenzhen, China.




Fig. 1. Language-Guided Multimodal Perception for 6-DoF Robotic Grasping in Cluttered Environments.

poor generalization to cluttered or occluded scenes, and cumulative errors across stages. Language-guided grasping has emerged as a promising paradigm that uses natural language instructions to guide perception and planning. Contrastive language-image pretraining models, such as CLIP [8], provide strong capabilities for aligning visual and linguistic modalities. Recent approaches integrate CLIP with preprocessed object regions and Transformer encoders to construct multimodal fusion frameworks [9]. Although these methods improve semantic understanding, several limitations remain. First, many approaches rely on external object and grasp detectors, which are susceptible to cascading errors. Second, commonly adopted visual encoders such as CLIP-ResNet and CLIP-ViT each face intrinsic drawbacks: CLIP-ResNet, constrained by its static hierarchical structure, cannot dynamically model object scale variations and deformations, thereby limiting its representation of multi-scale and non-rigid objects [10], [11]; CLIP-ViT, while offering global context modeling, incurs high computational cost due to quadratic self-attention and struggles to capture fine-grained structures under fixed patch partitioning, restricting its ability to represent local details and deformations. These limitations highlight the need for a visual framework that combines efficient global modeling with dynamic features for robust language-guided grasping.

Accurate spatial modeling is crucial in cluttered or occluded scenes, where objects overlap or exhibit substantial visual similarity. Recent studies have introduced depth information to enhance grasping by providing structural cues such as object shape, spatial relationships, and occlusion reason-

ing [12], [13]. However, depth maps from real-world sensors are often noisy or incomplete due to reflections, translucent materials, or missing regions. Existing RGB-D methods typically adopt dual-encoder architectures that process RGB and depth separately. While this extracts complementary features, it also incurs computational redundancy. Furthermore, low-quality or inconsistent depth data can hinder RGB feature learning [14]–[16], weakening spatial reasoning and target discrimination. Thus, it is essential to design strategies that use depth as geometric guidance while preserving RGB integrity and avoiding redundant computation [17].

Beyond deep guidance, effective language-guided grasping also hinges on precise cross-modal alignment that unifies semantic understanding with spatial reasoning. While language offers semantic cues for target identification and depth provides structural guidance, associating textual descriptions with corresponding visual regions remains highly challenging in cluttered or occluded environments [18]. In prevailing approaches, the vision encoder is typically frozen and only high-level features are passed to the connector to reduce computational cost, leaving informative low- and mid-level representations underutilized. Moreover, common cross-modal fusion strategies, such as concatenation or cross-attention, fail to fully exploit multi-layer visual features [19], thereby diluting task-relevant signals and overlooking critical semantic or spatial cues [20]. Approaches that either select a subset of visual layers or concatenate all layers further exacerbate this issue, as they produce high-dimensional representations that increase computational overhead and complicate optimization [10]. Collectively, these limitations hinder precise text-to-vision alignment in complex grasping scenarios, ultimately constraining the discriminative power and generalization ability of multimodal representations.

To address these challenges, we introduce GeoLanG (Geometry-Aware Language-Guided Grasping), a unified multi-task visual-language framework to perform reliable grasping in open-world environment. Our contributions are summarized as follows:

- We present GeoLanG, an end-to-end framework that integrates RGB-D perception with natural language understanding into a shared representation space, enabling robust semantic alignment and improved generalization in cluttered and complex environments.
- To enhance spatial reasoning, we introduce the Depth-guided Geometric Module (DGGM), which encodes depth as explicit geometric priors directly into the attention mechanism, improving robustness against occlusions and target discrimination without the need for a dedicated depth encoder.
- To facilitate effective cross-modal fusion, we propose Adaptive Dense Channel Integration (ADCI), which adaptively aggregates multi-layer visual features and emphasizes task-relevant signals, yielding more discriminative and generalizable visual representations.
- Extensive experiments on the OCID-VLG dataset demonstrate that GeoLanG surpasses existing methods in both language segmentation and grasping metrics,

while maintaining robustness under occlusions and ambiguous instructions, validating the effectiveness of our unified multi-task approach.

## II. METHOD

### A. Overall Architecture of GeoLanG

As shown in Fig. 2, GeoLanG is an end-to-end vision-language grasping framework that is specifically designed to enable precise language-guided grasping in complex open scenes. In this framework, given an RGB-D image and a language query, the text encoder first extracts high-level semantic features, while simultaneously, the RGB encoder processes multi-scale visual information.

We follow the CLIP paradigm to align visual and linguistic features into a shared semantic space. For the image encoder, existing CLIP visual backbones face notable challenges in complex grasping scenarios. CLIP-ResNet relies on a rigid hierarchical design that restricts multi-scale adaptability, while CLIP-ViT conducts self-attention over fixed-size patches, leading to high computational cost and limited sensitivity to fine-grained spatial cues. Therefore, we employ CLIP-VMamba [21], which integrates the complementary advantages of ViT and CNN architectures: the ViT backbone captures long-range dependencies and global context, whereas the CNN-style modules preserve high-resolution local details. This hybrid design effectively represents non-rigid and multi-scale objects in cluttered open scenes. In the CLIP-VMamba visual encoder, we retain the last three scales of features, excluding the lowest-level feature to balance computational cost and semantic richness. Formally, the retained multi-scale features can be denoted as  $C = \{C_i \in \mathbb{R}^{H/s_i \times W/s_i \times C_i} \mid i = 1, 2, 3\}$ , where  $C_1$  preserves spatial details for precise localization, and  $C_2$  and  $C_3$  encode higher-level semantic concepts. We set the downsampling factors to  $s_1 = 8$ ,  $s_2 = 16$ , and  $s_3 = 32$ , respectively. Then, the Depth-Guided Geometric Module (DGGM) incorporates depth-derived geometric priors to enhance structural cues. These spatial-enhanced features are optimized through Adaptive Dense Channel Integration (ADCI) to produce multi-scale visual representations  $C_v \in \mathbb{R}^{N \times C}$ , where  $N = H/16 \cdot W/16$ . Meanwhile, for language inputs, the textual instructions  $P = \{w_1, \dots, w_T\}$  are encoded with CLIP-BERT [22] to obtain token embeddings  $C_t \in \mathbb{R}^{K \times C}$  and a sentence embedding  $C_s \in \mathbb{R}^{C'}$ , capturing both local and global textual semantics.

In multi-task decoding,  $C_v$  is fused with the sentence embedding  $C_s$  through the Multimodal Fusion Neck to produce pixel-wise multimodal representations  $C_m \in \mathbb{R}^{N \times C}$ , which are subsequently fed into the downstream multi-task decoder with  $C_t$  to produce  $C_c \in \mathbb{R}^{N \times C}$ . In the segmentation branch,  $C_c$  and  $C_s$  are projected into a shared space and supervised via a binary cross-entropy loss applied to the dot product of the projected embeddings. The final segmentation mask is obtained by reshaping and upsampling the projected features, while the grasp head predicts grasp configurations from the same shared features, enabling joint optimization of perception and action reasoning.

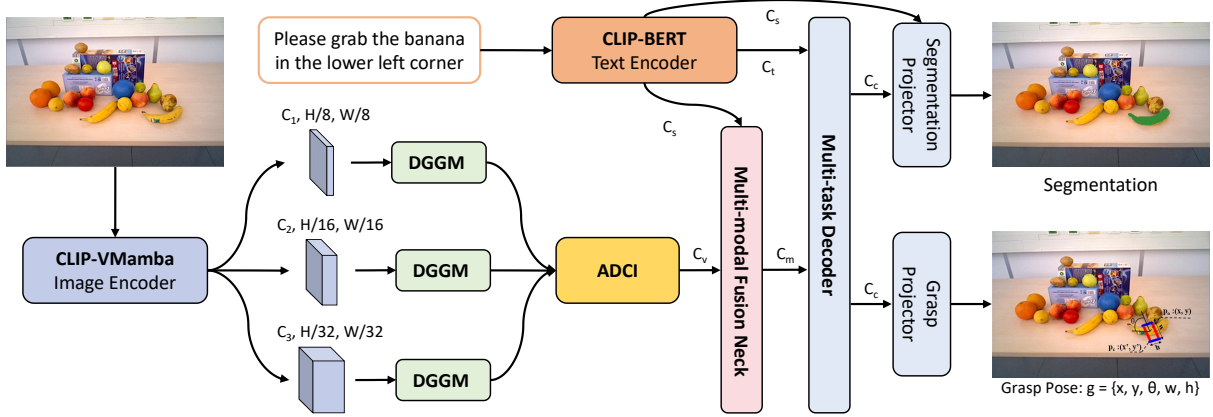


Fig. 2. An overview of the GeoLanG framework. Given an RGB-D image and a language query, the text encoder extracts high-level semantic features while the RGB encoder processes multi-scale visual information. Depth-derived geometric priors are incorporated via the Depth-Guided Geometric Module (DGGM) to enhance structural cues. Multi-layer visual features are then optimized through ADCI before being fed into dual-path projectors that generate pixel-level segmentation masks and refine the grasp pose for the target object.

### B. Depth-guided Geometric Module

To effectively incorporate depth cues while avoiding excessive computational overhead, we introduce depth-derived geometric priors [23] and fuse them with RGB features from the V-Mamba encoder (Fig. 3). Specifically, a 2D input RGB image of size  $h \times w$  is evenly partitioned into  $H \times W$  patches, where  $H$  and  $W$  denote the number of patches along rows and columns, respectively. Each patch  $Q_{m,n}$  is indexed by its spatial coordinates  $(m, n)$ . When the corresponding depth map is available, the patch at the same position encodes its relative distance from the camera plane. By integrating these two modalities, we explicitly model the geometric relationships among all patches and embed them into the self-attention mechanism, capable of adapting to multi-scale image features.

For the depth prior, we apply average pooling to all pixels in the depth patch at  $(m, n)$  to obtain its representative depth value  $D_{m,n}$ . The depth difference between any two patches is defined as:

$$\Delta D_{m,n,m',n'} = |D_{m,n} - D_{m',n'}| \quad (1)$$

where  $\Delta D \in \mathbb{R}^{HW \times HW}$  forms the depth relationship matrix. The spatial distance is calculated through the Manhattan distance between patch coordinates:

$$\Delta S_{m,n,m',n'} = |m - m'| + |n - n'| \quad (2)$$

where  $\Delta S \in \mathbb{R}^{HW \times HW}$  is the spatial relationship matrix. To unify depth and spatial priors, the fused geometry prior  $\mathcal{G}$  is constructed by:

$$\mathcal{G} = \lambda_1 \cdot \Delta D + \lambda_2 \cdot \Delta S \quad (3)$$

Where  $\lambda_1, \lambda_2$  are learnable parameters balancing the contributions of depth and spatial priors, the resulting  $\mathcal{G} \in \mathbb{R}^{HW \times HW}$  encodes comprehensive 3D geometric relationships among all patches.

For the integration of the geometry prior  $\mathcal{G}$  into the attention mechanism, given an RGB feature map  $X \in \mathbb{R}^{HW \times C}$ ,

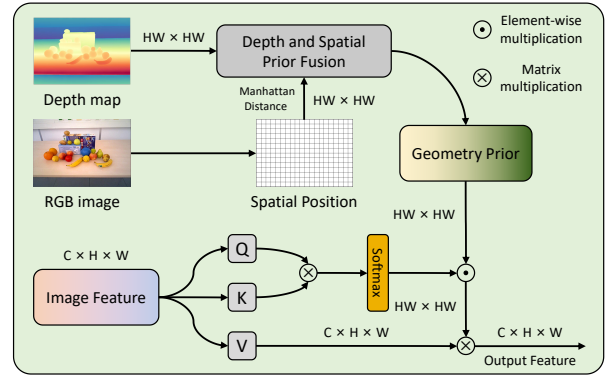


Fig. 3. The overview of the Depth-guided Geometric Module (DGGM). The pink rectangle represents image features extracted from the RGB encoder, the green rectangle denotes the learned geometry prior, and the yellow rectangle indicates the spatial prior computed from depth and RGB.  $\otimes$  represents matrix multiplication and  $\odot$  represents element-wise multiplication. The output feature integrates visual information with geometric and spatial cues to enhance multi-scale representations.

we can obtain the query, key, and value matrices by linear projections:  $Q, K, V \in \mathbb{R}^{HW \times d}$ . We integrate the geometry prior  $\mathcal{G}$  via a decay factor  $\eta \in (0, 1)$  to emphasize nearby patches:

$$\hat{X} = (\text{Softmax}(QK^\top) \odot \eta \mathcal{G}) V^\top \quad (4)$$

Smaller values in  $\eta \mathcal{G}$  correspond to larger geometric distances, suppressing irrelevant key-value pairs and emphasizing geometrically relevant ones.

This formulation can be extended to multi-head attention by assigning different decay rates  $\eta$  for each head. Through DGGM, our model effectively captures spatial dependencies, object shape variations, and low-texture regions, while maintaining computational efficiency.

### C. Adaptive Dense Channel Integration

Prevailing paradigms in visual representation learning often freeze backbones and only propagate high-level semantic features, which discards valuable low- and mid-level cues. In

this work, we propose the Adaptive Dense Channel Integration (ADCI) module to holistically harness this multi-layer information. ADCI first preserves dense cross-layer connections while adaptively balancing contributions from different layers. Given  $L$  image features  $\{C_1, \dots, C_L\}$  extracted from V-Mamba, they are divided into  $G$  groups, each containing  $M = L/G$  consecutive layers. Learnable weights  $\alpha_i$  within each group adaptively adjust layer contributions based on input features, enabling dynamic information aggregation. For the  $g$ -th group of  $M$  consecutive layers, the group-level feature  $GC_g$  is computed as a weighted combination of its constituent features:

$$GC_g = \sum_{i=(g-1)M+1}^{gM} \alpha_i C_i, \quad 1 \leq g \leq G \quad (5)$$

where the adaptive weights  $\alpha_i$  are predicted by a lightweight gating network and normalized to satisfy

$$\sum_{i=(g-1)M+1}^{gM} \alpha_i = 1 \quad (6)$$

The gating network first summarizes each feature map  $C_i$  into a compact descriptor  $z_i$  using global average pooling (GAP):

$$z_i = \text{GAP}(C_i) \quad (7)$$

Which is then passed through a lightweight two-layer MLP to produce unnormalized scores  $s_i$ :

$$s_i = W_2 \sigma(W_1 z_i + b_1) + b_2 \quad (8)$$

Where  $W_1$  and  $W_2$  are learnable parameters and  $\sigma$  denotes a non-linear activation function (e.g., ReLU). Softmax normalization within each group yields the adaptive weights  $\alpha_i$ :

$$\alpha_i = \frac{\exp(s_i)}{\sum_{j=(g-1)M+1}^{gM} \exp(s_j)} \quad (9)$$

Finally, the group-level features  $\{GC_1, \dots, GC_G\}$  are concatenated with the last-layer feature  $C_L$  and projected through an MLP to generate the final visual embedding:

$$e_v = \text{MLP}\left(\text{Concatenate}([GC_1, \dots, GC_G, C_L], \text{dim=channel})\right) \quad (10)$$

This formulation allows ADCI to balance contributions across layers while preserving dense connections, ensuring informative low- and mid-level features are effectively aggregated. Consequently, it yields richer visual representations that enhance cross-modal expressiveness for downstream tasks while maintaining efficiency. Algorithm 1 shows the flowchart for ADCI.

### III. EXPERIMENT

#### A. Dataset and Experiment setup

To evaluate our approach on language-guided segmentation and grasping, we use the OCID-VLG benchmark [24]. As shown in Fig. 4, OCID-VLG extends the OCID cluttered indoor scenes with multimodal annotations, bridging

---

#### Algorithm 1: Adaptive Dense Channel Integration (ADCI) for Mamba Features

---

**Input:** Multi-layer visual features  $\{C_1, \dots, C_L\}$  from VMamba image encoder

**Output:** Final visual embedding  $e_v$

**Step 1: Grouping.** Partition  $\{C_1, \dots, C_L\}$  into  $G$  groups, each with  $M = L/G$  layers.

**for**  $g \leftarrow 1$  **to**  $G$  **do**

**Step 2: Feature descriptors.** For each feature  $C_i$  in group  $g$ :  
 $z_i = \text{GAP}(C_i)$ ; // Global average pooling

**Step 3: Gating network.** Compute score via 2-layer MLP:

$$s_i = W_2 \sigma(W_1 z_i + b_1) + b_2$$

**Step 4: Adaptive weights.** Apply softmax within the group:

$$\alpha_i = \frac{\exp(s_i)}{\sum_{j=(g-1)M+1}^{gM} \exp(s_j)}$$

**Step 5: Group feature aggregation.**

$$GC_g = \sum_{i=(g-1)M+1}^{gM} \alpha_i C_i$$

**Step 6: Final embedding.** Concatenate group features with last layer  $C_L$ :

$$e_v = \text{MLP}\left(\text{Concatenate}([GC_1, \dots, GC_G, C_L], \text{dim=channel})\right)$$


---

language supervision and robotic grasping. It contains 1,763 RGB-D tabletop scenes across 31 categories and 58 fine-grained instances, with labels covering semantic attributes, spatial relations, 2D/3D locations, and 4-DoF grasps. By incorporating CLEVR-style data generation, the dataset provides 89,639 triplets of language expressions, segmentation masks/bounding boxes, and grasp annotations, enabling unified multimodal supervision.



Fig. 4. Examples from the OCID-VLG dataset [24]. (a) Single item (b) Multiple items cluttered scene (c) Multiple objects cluttered and overlapping scenes.

1) *Experimental Setup:* We initialize the visual and textual encoders with pre-trained CLIP-VMamba and BERT weights, respectively. The resolution of input images is  $416 \times 416$ , while text sequences are tokenized using byte-pair encoding (BPE) [25], [26] with a maximum length of 20 tokens. Training is performed on four NVIDIA RTX 4090 GPUs using the Adam optimizer for 50 epochs, with an initial learning rate of  $1 \times 10^{-4.5}$  and a batch size of 64.

TABLE I  
COMPARISON EXPERIMENTS AGAINST SOTA SOLUTIONS ON THE OCID-VLG DATASET.

Method	Segmentation Task						Grasping Task	
	IoU	Pr@50	Pr@60	Pr@70	Pr@80	Pr@90	J@1	J@N
VLG	76.35	89.84	85.38	77.63	59.67	20.36	78.51	85.19
GraspCLIP	77.20	91.34	87.85	80.88	62.52	13.50	78.32	84.73
CLIPort	78.11	94.03	90.84	82.73	60.08	15.68	83.90	88.75
CTNet	80.94	<b>97.45</b>	95.43	89.63	67.05	16.66	84.32	90.91
CROG	80.77	95.87	93.83	86.69	65.14	<b>16.85</b>	81.64	88.05
GeoLanG (Ours)	<b>85.77</b>	96.73	<b>95.45</b>	<b>89.82</b>	<b>68.23</b>	16.70	<b>87.32</b>	<b>92.13</b>

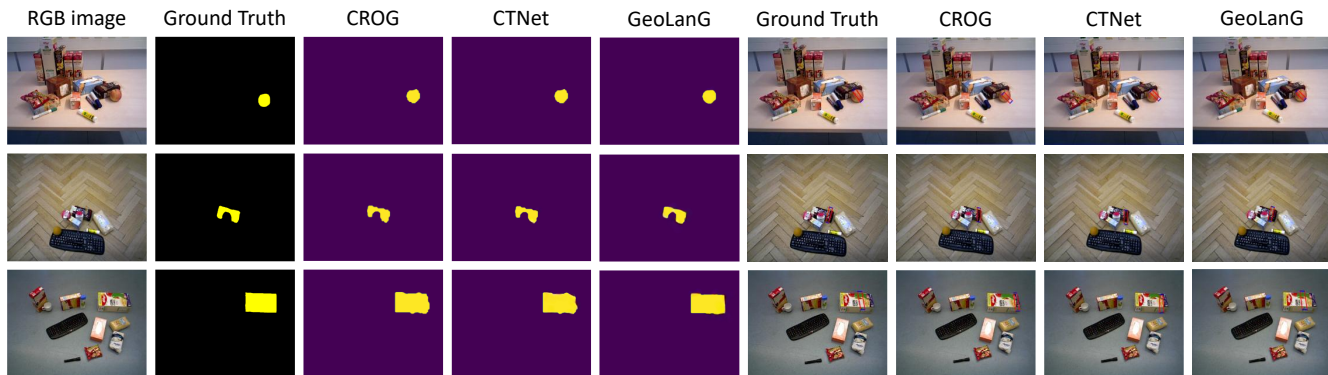


Fig. 5. Qualitative comparison against SOTA solutions on the OCID-VLG dataset.

2) *Evaluation Metrics*: For referring image segmentation (RIS) [27], [28], we adopt Intersection-over-Union (IoU) and Precision@ $X$ . IoU measures the average overlap between predicted masks and ground-truth annotations, while Precision@ $X$  computes the percentage of predictions whose IoU exceeds thresholds  $X \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ . For referring grasp synthesis (RGS), we report the Jacquard index  $J@N$  [29]–[31], which evaluates the proportion of the top- $N$  grasp predictions that achieve an intersection-over-union greater than 0.25 with the ground-truth grasp rectangle and an orientation difference within  $30^\circ$ .

### B. Comparison with baselines on OCID-VLG

We first evaluate our proposed GeoLanG framework on the OCID-VLG dataset and compare it with several CLIP-based baselines, including CROG [24], GraspCLIP [32], CLIPort [33], and CTNet [34]. As shown in Table I, GeoLanG achieves superior performance on both segmentation and grasping metrics. For segmentation, the model attains an IoU of 85.77%, substantially outperforming CTNet (80.94%) and CROG (80.77%), demonstrating its ability to produce more accurate pixel-level object masks. Under stricter evaluation, GeoLanG maintains consistent robustness, achieving 89.82% at Pr@70 compared with 89.63% for CTNet and 86.69% for CROG, indicating effective multi-scale feature aggregation and precise localization even under challenging conditions. Although its performance at Pr@90 (16.70%) is slightly lower than CROG (16.85%), the overall stability across varying thresholds highlights the model’s balanced precision-recall trade-off. In the grasping task, GeoLanG reaches the highest performance with J@1 of 87.32% and J@N of 92.13%, surpassing CTNet by 3.00 and 1.22 points, respec-

tively. CLIPort achieves competitive grasp metrics (J@1 = 83.90%, J@N = 88.75%) but underperforms, primarily due to coarser pixel-level segmentation and weaker depth-aware cues, which constrain reliable grasp prediction in cluttered or occluded scenes. These results demonstrate that integrating RGB-D priors with dense multi-scale feature fusion and depth-guided geometric reasoning enables GeoLanG to achieve both accurate segmentation and grasp prediction, yielding strong generalization in complex environments.

Since CROG and CTNet address the same task as proposed GeoLanG and incorporate attribute-aware segmentation, we further conduct a qualitative comparison. As illustrated in Fig. 5, GeoLanG delivers more precise fine-grained recognition, generating segmentation masks that closely follow object boundaries. This advantage is especially evident in challenging cases such as small objects, blurry edges, and occlusions (e.g., first and second rows). In contrast, CROG and CTNet often yield oversimplified masks (e.g., retaining only central regions) or misaligned boundaries, indicating weaker semantic alignment under complex visual conditions. For grasp pose prediction, CROG and CTNet often display grasp drift or incorrectly target neighboring objects, while GeoLanG consistently localizes grasp points on small objects within cluttered scenes. These advantages stem from GeoLanG’s enhanced spatial reasoning and depth-aware geometric perception, which strengthen its robustness in cluttered and occluded environments.

### C. Comparison with baselines on unseen items (OCID-VLG)

To evaluate generalization, we comprehensively test GeoLanG on the *novel instance* split of OCID-VLG (Table II). GeoLanG achieves the highest overall performance, with

TABLE II  
COMPARISON EXPERIMENTS AGAINST SOTA SOLUTIONS ON THE *novel instance* SPLIT OF OCID-VLG.

Method	Segmentation Task						Grasping Task	
	IoU	Pr@50	Pr@60	Pr@70	Pr@80	Pr@90	J@1	J@N
VLG	69.70	77.03	70.16	63.31	49.58	14.78	66.42	74.91
GraspCLIP	69.13	83.21	76.44	63.80	42.20	8.43	63.98	58.16
CLIPort	71.17	80.51	74.23	67.75	53.77	18.55	66.24	74.03
CTNet	74.38	87.34	81.63	73.89	57.23	17.83	72.56	79.06
CROG	73.28	82.98	72.22	81.41	56.61	<b>20.20</b>	71.36	81.41
GeoLanG (Ours)	<b>81.25</b>	<b>96.32</b>	<b>93.48</b>	<b>88.19</b>	<b>63.78</b>	15.42	<b>83.96</b>	<b>90.72</b>

TABLE III  
ABLATION STUDY ON IMAGE ENCODER, DEPTH-GUIDED GEOMETRIC MODULE AND ADAPTIVE DENSE CHANNEL INTEGRATION.

CLIP-VMamba	DGGM	ADCI	Segmentation Task						Grasping Task	
			IoU	Pr@50	Pr@60	Pr@70	Pr@80	Pr@90	J@1	J@N
✗	✗	✗	80.77	95.87	93.83	86.69	65.14	16.85	81.64	88.05
✓	✗	✗	83.25	96.23	94.39	86.93	67.04	16.91	83.44	90.76
✓	✓	✗	84.88	96.46	95.27	88.38	67.52	<b>17.02</b>	85.35	91.28
✓	✓	✓	<b>85.77</b>	<b>96.73</b>	<b>94.45</b>	<b>89.82</b>	<b>68.23</b>	16.70	<b>87.32</b>	<b>92.13</b>

IoU 81.25%, J@1 83.96%, and J@N 90.72%, clearly surpassing CTNet by 6.87, 11.66, and 11.66 points on IoU, J@1, and J@N, respectively. In addition, GeoLanG achieves strong and stable segmentation at mid thresholds (Pr@50 = 96.32%, Pr@60 = 93.48%, Pr@70 = 88.19%) while still outperforming baselines in grasping, despite slightly lower Pr@90 (15.42%) than CROG (20.20%). Overall, GeoLanG demonstrates robust and consistent segmentation-grasping predictions on novel objects. These performances primarily stem from the Depth-guided Geometric Module, which effectively uses RGB-D priors to maintain segmentation and grasping accuracy for occluded or unseen objects, and the Adaptive Dense Channel Integration, which adaptively aggregates multi-layer features for precise localization.

#### D. Ablation Studies

To validate the effectiveness of our proposed components in GeoLanG, we conduct an ablation study on the CLIP-VMamba backbone, DGGM, and ADCI, with results summarized in Table III. When using CLIP-ResNet instead of CLIP-VMamba, the model struggles to capture sufficient spatial reasoning and cross-modal semantics, leading to significantly lower grasping and segmentation performance (J@1 = 81.64%, J@N = 88.05%, IoU = 80.77%). Incorporating CLIP-VMamba alleviates this issue, as its stronger representation learning boosts performance across all metrics, confirming its importance for multi-scale visual-textual alignment. Building on this, the incorporation of DGGM enables explicit modeling of RGB-D geometric priors, leading to notable gains in structural sensitivity and robustness, as reflected by higher Pr@90 (17.02%). Further, the integration of ADCI strengthens multi-layer feature aggregation, allowing the model to capture fine-grained localization cues and ultimately achieve the best overall trade off across segmentation and grasping tasks, particularly at higher precision thresholds (e.g., Pr@80) (J@1 = 87.32%, J@N = 92.13%, IoU = 85.77%, Pr@80 = 68.23%). These results demonstrate that each module provides complementary benefits CLIP VMamba enhances representation capacity, DGGM enriches

geometric reasoning, and ADCI facilitates dense feature fusion—together driving robust and high-precision referring grasping in challenging scenarios.




Fig. 6. Simulation experiments on the RoboDK platform.

#### E. Robot Experiments

We further evaluate GeoLanG on a table-cleaning task through both simulation and real-world experiments. To mitigate the observation domain gap between CLIP pretraining and GeoLanG training on the OCID-VLG dataset, we fine-tune the vision-language encoders end-to-end on the target data and incorporate depth-guided geometric features, which reduce sensitivity to appearance variations and improve cross-domain robustness. The simulation study is conducted on the RoboDK platform using an eye-to-hand configuration for efficient setup (Fig. 6), whereas the real-world experiment employs a DOBOT Nova 2 robot equipped with parallel-jaw grippers and an Intel RealSense D405 camera in an eye-in-hand configuration (Fig. 7(c)), enabling accurate grasp

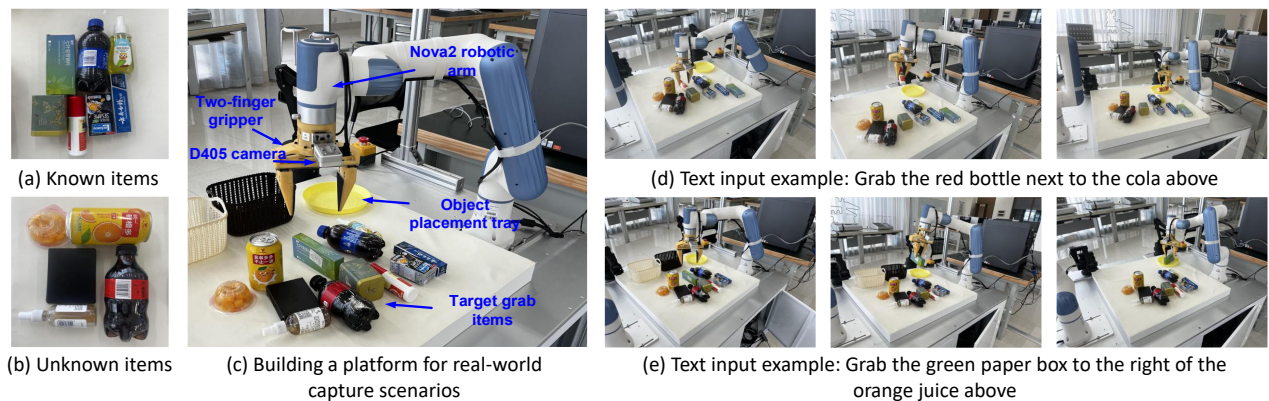


Fig. 7. Real-world experiments on the DOBOT Nova 2 robo.

TABLE IV  
SIMULATION AND REAL ROBOT EXPERIMENTAL RESULTS

Setting	Toothbrush	Mug	Fruit	Bottle	Spray	Medicine	Snack Box	Total	Sim
Isolated Scene (Count)	3	5	7	3	2	10	10	40	40
Segmentation IOU	3(100%)	4 (80%)	6 (85.71%)	3 (100%)	2 (100%)	6 (60%)	10(100%)	34 (85%)	98.21%
Grasping Accuracy	3(100%)	4 (80%)	6 (85.71%)	2 (66.67%)	2 (100%)	5 (50%)	10(100%)	33 (82.5%)	100%
Cluttered Scene (Count)	3	5	7	3	2	10	10	40	40
Segmentation IOU	3(100%)	5 (100%)	6 (85.71%)	2 (66.67%)	2 (100%)	6 (60%)	9 (90%)	32 (80%)	94.21%
Grasping Accuracy	2 (66.67%)	4 (80%)	4 (57.14%)	1 (33.33%)	1 (50%)	4 (40%)	8 (80%)	24 (60%)	92%

execution. The eye-in-hand configuration is adopted for practical deployment and accurate viewpoint alignment during grasp execution, while the perception and grasp prediction are performed in a single-shot manner without closed-loop visual feedback. Example trials in both isolated and cluttered scenes are shown in Fig. 7(d) and (e). Notably, although the robotic arm offers 6-DoF motion, 4-DoF grasp annotations are sufficient, as the remaining degrees of freedom primarily influence approach orientation rather than grasp location or width.

We design two experimental settings to evaluate GeoLanG: *isolated* scenes, where objects are spatially separated, and *cluttered* scenes, where objects are densely packed. To further increase the challenge, each scene includes distractors from the same category as the query object (Fig. 7(a)(b)). The evaluation covers eight object categories, consisting of four seen classes (toothbrush, mug, fruit, beverage bottle) and three unseen classes (spray, medicine, snack box). In total, 40 scenes are constructed (20 isolated, 20 cluttered) with each scene paired with a corresponding language instruction.

As shown in Table IV, GeoLanG achieves high performance in simulation, with grasping accuracy reaching 95% across all categories. In real-world isolated scenes, segmentation and grasping accuracies are 85% and 82.5%, respectively, while in cluttered scenes they decrease to 80% and 60%, reflecting the impact of sensor noise, calibration errors, and occlusions. Detailed per-category results show that objects with minimal occlusion, such as toothbrushes and spray bottles, achieve near-perfect segmentation and grasping, whereas heavily occluded or small objects, such as medicine boxes, exhibit lower performance. Despite unseen objects and complex occlusions, GeoLanG consistently

localizes attribute concepts and uses spatial relationships to distinguish target objects. Overall, the results highlight GeoLanG’s strong generalization to novel objects, robustness under both virtual and real-world imperfections, and effectiveness in executing table-cleaning tasks across both structured and cluttered environments.

#### IV. DISCUSSION

GeoLanG demonstrates competitive performance on language-guided grasping under standard benchmark settings. The current framework, however, adopts a simplified formulation by focusing on 4-DoF planar grasping, which reduces the complexity of gripper orientation and interaction dynamics. While this design is consistent with common practice in existing benchmarks, it constrains the range of manipulation scenarios. The evaluation protocol reflects this formulation. Experiments are primarily conducted on the OCID-VLG benchmark and set of real-world table-cleaning tasks, offering a controlled but relatively narrow coverage of real-world variability. Consequently, the behavior of the model in more diverse or dynamic environments has yet to be fully characterized. These considerations point to several directions for future work, including extending the framework to full 6-DoF grasping and validating its performance in broader and more dynamic real-world settings. Such extensions would enable a more comprehensive assessment of the generality of GeoLanG for complex language-conditioned manipulation tasks.

#### V. CONCLUSIONS

This work introduces GeoLanG, an end-to-end vision-language grasping framework that integrates RGB-D perception with natural language understanding. GeoLanG adopts

the CLIP paradigm, employing a CLIP-VMamba visual encoder and a CLIP-BERT text encoder to establish a shared representation space, thereby strengthening cross-modal alignment and improving semantic discrimination. To make full use of depth cues, the Depth-Guided Geometric Module injects geometric priors into RGB features, strengthening spatial awareness and improving robustness under occlusion and low-texture conditions. Besides, the Adaptive Dense Channel Integration module further integrates multi-layer features, boosting the discriminative power of visual representations. Extensive experiments demonstrate the effectiveness of our proposed methods in challenging scenarios, showing improved target localization, enhanced object understanding, and efficient representation learning for downstream tasks including grasp detection, segmentation, and instruction parsing.

## REFERENCES

- [1] J. Hatori, Y. Kikuchi, S. Kobayashi, K. Takahashi, Y. Tsuboi, Y. Unno, W. Ko, and J. Tan, "Interactively picking real-world objects with unconstrained spoken language instructions," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 3774–3781.
- [2] M. Shridhar, D. Mittal, and D. Hsu, "Ingress: Interactive visual grounding of referring expressions," *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 217–232, 2020.
- [3] P. Asgharian, A. M. Panchea, and F. Ferland, "A review on the use of mobile service robots in elderly care," *Robotics*, vol. 11, no. 6, 2022. [Online]. Available: <https://www.mdpi.com/2218-6581/11/6/127>
- [4] C. K. Ng, L. Bai, G. Wang, Y. Wang, H. Gao, C. Jin, T. Zeng, H. Ren *et al.*, "Endovla: Dual-phase vision-language-action for precise autonomous tracking in endoscopy," in *9th Annual Conference on Robot Learning*.
- [5] M. Hersh, "Overcoming barriers and increasing independence—service robots for elderly and disabled people," *International Journal of Advanced Robotic Systems*, vol. 12, no. 8, p. 114, 2015.
- [6] Z. Zhang, C. Zhou, Y. Koike, and J. Li, "Single rgb image 6d object grasping system using pixel-wise voting network," *Micromachines*, vol. 13, no. 2, 2022. [Online]. Available: <https://www.mdpi.com/2072-666X/13/2/293>
- [7] R. Tang, S. Guo, Y. Qiu, H. Chen, L. Huang, M. Yong, L. Zhou, and L. Guo, "Optimizing drug delivery in smart pharmacies: A novel framework of multi-stage grasping network combined with adaptive robotics mechanism," *arXiv preprint arXiv:2410.00753*, 2024.
- [8] 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.
- [9] H. Wang, G. Yang, S. Zhang, J. Qin, Y. Guo, B. Xu, Y. Jin, and L. Zhu, "Video-instrument synergistic network for referring video instrument segmentation in robotic surgery," *IEEE Transactions on Medical Imaging*, 2024.
- [10] S. Zhang, B. Li, and C. Yin, "Cross-modal sentiment sensing with visual-augmented representation and diverse decision fusion," *Sensors*, vol. 22, no. 1, p. 74, 2021.
- [11] S. Shen, L. H. Li, H. Tan, M. Bansal, A. Rohrbach, K.-W. Chang, Z. Yao, and K. Keutzer, "How much can clip benefit vision-and-language tasks?" *arXiv preprint arXiv:2107.06383*, 2021.
- [12] H. Lin, S. Peng, J. Chen, S. Peng, J. Sun, M. Liu, H. Bao, J. Feng, X. Zhou, and B. Kang, "Prompting depth anything for 4k resolution accurate metric depth estimation," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 17 070–17 080.
- [13] B. Yu, Q. Yuan, K. Li, H. Kasaei, and M. Cao, "Co-navgpt: Multi-robot cooperative visual semantic navigation using vision language models," *Preprint*, 2025.
- [14] L. Bai, B. Cui, L. Wang, Y. Li, S. Yao, S. Yuan, Y. Wu, Y. Zhang, M. Q.-H. Meng, Z. Li *et al.*, "V<sup>2</sup>-sfmlerner: Learning monocular depth and ego-motion for multimodal wireless capsule endoscopy," *IEEE Transactions on Automation Science and Engineering*, 2025.
- [15] S. Bodenstedt, M. Wagner, L. Mündermann, H. Kenngott, B. Müller-Stich, M. Breucha, S. T. Mees, J. Weitz, and S. Speidel, "Prediction of laparoscopic procedure duration using unlabeled, multimodal sensor data," *International Journal of Computer Assisted Radiology and Surgery*, vol. 14, pp. 1089–1095, 2019.
- [16] X. Li, S. Wang, C. Chen, H. Wei, Y. Shi, and H. Mo, "Roboflamingo-plus: Fusion of depth and rgb perception with vision-language models for enhanced robotic manipulation," in *2025 IEEE International Conference on Real-time Computing and Robotics (RCAR)*. IEEE, 2025, pp. 311–316.
- [17] D. Marin, J.-H. R. Chang, A. Ranjan, A. Prabhu, M. Rastegari, and O. Tuzel, "Token pooling in vision transformers," *arXiv preprint arXiv:2110.03860*, 2021.
- [18] X. Yi, H. Xu, H. Zhang, L. Tang, and J. Ma, "Text-if: Leveraging semantic text guidance for degradation-aware and interactive image fusion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 27 026–27 035.
- [19] R. Tang, H. Yin, G. Wang, L. Bai, A. Wang, H. Gao, J. Wang, and H. Ren, "Geo-repnet: Geometry-aware representation learning for surgical phase recognition in endoscopic submucosal dissection," in *2025 International Conference on Information and Automation (ICIA)*. IEEE, 2025, pp. 359–364.
- [20] J. Chen, G. Wang, W. Liu, X. Zhong, Y. Tian, and Z. Wu, "Perception reinforcement using auxiliary learning feature fusion: A modified yolov8 for head detection," in *2023 China Automation Congress (CAC)*. IEEE, 2023, pp. 4709–4714.
- [21] W. Huang, Y. Shen, and Y. Yang, "Clip-mamba: Clip pretrained mamba models with ood and hessian evaluation," *arXiv preprint arXiv:2404.19394*, 2024.
- [22] J. Lei, L. Li, L. Zhou, Z. Gan, T. L. Berg, M. Bansal, and J. Liu, "Less is more: Clipbert for video-and-language learning via sparse sampling," in *CVPR*, 2021.
- [23] B.-W. Yin, J.-L. Cao, M.-M. Cheng, and Q. Hou, "Dformerv2: Geometry self-attention for rgb-d semantic segmentation," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 19 345–19 355.
- [24] G. Tziafas, X. Yucheng, A. Goel, M. Kasaei, Z. Li, and H. Kasaei, "Language-guided robot grasping: Clip-based referring grasp synthesis in clutter," in *7th Annual Conference on Robot Learning*, 2023.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [26] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [27] Z. Wang, Y. Lu, Q. Li, X. Tao, Y. Guo, M. Gong, and T. Liu, "Cris: Clip-driven referring image segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 11 686–11 695.
- [28] H. Ding, C. Liu, S. Wang, and X. Jiang, "Vision-language transformer and query generation for referring segmentation," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 16 321–16 330.
- [29] A. Depierre, E. Dellandréa, and L. Chen, "Jacquard: A large scale dataset for robotic grasp detection," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 3511–3516.
- [30] Y. Xu, M. Kasaei, H. Kasaei, and Z. Li, "Instance-wise grasp synthesis for robotic grasping," *arXiv preprint arXiv:2302.07824*, 2023.
- [31] S. Ainetter and F. Fraundorfer, "End-to-end trainable deep neural network for robotic grasp detection and semantic segmentation from rgb," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 452–13 458.
- [32] C. Tang, D. Huang, L. Meng, W. Liu, and H. Zhang, "Task-oriented grasp prediction with visual-language inputs," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 4881–4888.
- [33] 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.
- [34] J. Xie, J. Liu, S. Huang, C. Wang, and F. Zhou, "Listen, perceive, grasp: Clip-driven attribute-aware network for language-conditioned visual segmentation and grasping," *IEEE Transactions on Automation Science and Engineering*, 2024.