

OCID-Ref: A 3D Robotic Dataset with Embodied Language for Clutter Scene Grounding

Ke-Jyun Wang* Yun-Hsuan Liu* Hung-Ting Su Jen-Wei Wang
Yu-Siang Wang Winston H. Hsu Wen-Chin Chen
National Taiwan University

Abstract

To effectively apply robots in working environments and assist humans, it is essential to develop and evaluate how visual grounding (VG) can affect machine performance on occluded objects. However, current VG works are limited in working environments, such as offices and warehouses, where objects are usually occluded due to space utilization issues. In our work, we propose a novel OCID-Ref dataset featuring a referring expression segmentation task with referring expressions of occluded objects. OCID-Ref consists of 305,694 referring expressions from 2,300 scenes with providing RGB image and point cloud inputs. To resolve challenging occlusion issues, we argue that it's crucial to take advantage of both 2D and 3D signals to resolve challenging occlusion issues. Our experimental results demonstrate the effectiveness of aggregating 2D and 3D signals but referring to occluded objects still remains challenging for the modern visual grounding systems. OCID-Ref is publicly available at <https://github.com/lluma/OCID-Ref>

1 Introduction

Visual grounding (VG), which aims to locate the object according to a structured language query, is a crucial task in natural language processing (NLP), computer vision (CV), and robotics. Recent VG studies most focus on *web-crawled images* such as (Kazemzadeh et al., 2014; Krishna et al., 2017; Mao et al., 2016; Yu et al., 2016). However, VG for human-robot interaction (HRI) is less explored. Most of the images in existing VG datasets are people and daily necessities, e.g., RefCOCO contains mainly persons, cars, and cats, which are separated and therefore easier to detect. Nevertheless, working spaces such as offices or warehouses, where robots are usually applied to assist works, are usually crowded, and objects are overlapped with each

* Equal contribution.

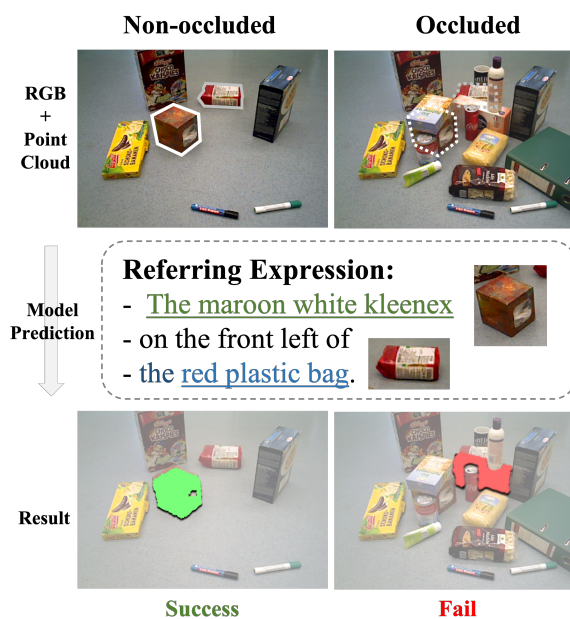


Figure 1: A hard case where visual grounding (VG) network fails to predict the occluded object in clutter scene. Our dataset provides more such cases than other datasets, which are commonly seen in the working spaces, like offices and warehouses.

other to utilize space better. Therefore, objects in working environments are often occluded and hard to detect.

Previous work (Ralph and Moussa, 2005) suggested that a system that uses language for human-computer interaction can help non-professionals instruct robots to complete technical work and collaborate. Recent research pointed out that VG plays an important role in HRI. (Shridhar and Hsu, 2018) utilized VG to resolve ambiguity in grasping tasks. (Matuszek) studied how the robot learns about objects and tasks in an environment via nature language queries. Therefore, explicit language instructions and good referring (grounding) expressions are pivotal in human-robot interaction and improve communication between non-expert humans and

robots.

Some efforts have been made to collect VG datasets. RefCOCO (Yu et al., 2016) and Cops-Ref (Chen et al., 2020b) utilize web-crawled images and manually label language expressions. A limitation is that images alone do not provide precise position cues, which are essential for various downstream robotic tasks such as grasping. A recent work, Sun-Spot (Mauceri et al., 2019), utilizes a depth channel for object detection and referring expression segmentation tasks. Another existing dataset, ScanRefer (Chen et al., 2020a), uses more accurate multi-view point clouds for 3D signals. However, both Sun-Spot and ScanRefer do not address *occlusion* issues, which is ordinary in working spaces and more challenging due to more compositions of shapes of each object. As shown in figure 1, when an object (*the red plastic bag*) is blocked in an occluded environment, the shape of the object could be deformed and increase VG difficulty.

Observing this, we propose a novel OCID-Ref dataset with two key features: (1) For each scene, we utilize both RGB image and point cloud to provide multi-modal signals for learning system development. (2) OCID-Ref scenes have higher clutter level compared to existing datasets, as shown in figure 2. Hence, the model capability for resolving challenging occlusion issues could be evaluated. To the best of our knowledge, OCID-Ref is the only existing dataset supporting the above features, and therefore allows VG task in grasping scenario.

Experimental results demonstrate that occluded scenes are more challenging to modern VG baselines. We observe 27% to 34% performance drops on referring expression segmentation tasks. Also, utilizing 3D information continually improves performance across all clutter levels. Furthermore, fusing 2D and 3D features reach the best performance on all clutter levels. We suggest that OCID-Ref dataset could pave a new path for VG research in HRI and benefit the research community and application developments.

2 Dataset and Task

To open up a new way for VG research in HRI, we collect a novel OCID-Ref dataset by the following steps: (1) We leverage a robotic object cluttered indoor dataset, OCID (Suchi et al., 2019), which consists of complex clutter-level scenes with rich 3D point cloud data and the point-wise instance




Figure 2: **Clutter level ratio.** Our proposed dataset, OCID-Ref, has more challenging and practical examples (touching and stacked) than ScanRefer (Chen et al., 2020a) and SunSpot (Mauceri et al., 2019) for visual grounding task.

labels for each occluded objects. (2) We manually annotate fine-grained attributes and relations such as color, shape, size relation or spatial relation. (3) We generate referring expressions based on annotated attributes and relations with a similar scene-graph generation system from (Yang et al., 2020) and (Chen et al., 2020c). In this section, we will describe more details on our data collection and the scene-graph generation method we adopt to generate the referring expressions.

2.1 Data Collection

A proper dataset to evaluate and develop VG models in a working environment requires two properties: (1) cluttered scenes and (2) 3D signals. To point out the important of these two properties, we conduct a pilot experiment of grasp detection¹. We observe that using 3D cues significantly boosts performance, the geometric features extracted from point cloud data benefit the robots on visual perception (e.g., object grasping or object tracking). Also, we see a severe performance drop in occluded scenes.

Therefore, to provide scenes with occluded objects to develop and evaluate learning systems, we leverage an existing robotic 3D dataset, OCID (Suchi et al., 2019), which has higher clutter level scenes and sequential object-level scenes that help robots better understand the instance difference between two subsequent scenes.

Hence, we choose OCID as our original dataset, and extend it with extra semantic annotations such as attributes (e.g., color, texture, shape) and rela-

¹See Appendix A for details

	# Scenes/# Images	# Obj. Cat.	Dis. Score	Data Format	# Expressions	AvgLen
RefCOCO	19,994	80	4.9	RGB	142,210	3.5
Cops-Ref	75,299	508	20.3	RGB	148,712	14.4
Sun-Spot	1,948	38	2.46	RGB-D	7,990	14.04
ScanRefer	703	18	4.64	3D	46,173	17.91
OCID-Ref	2,300	58	3.36	3D	267,339	8.56

Table 1: Statistic comparison of previous 2D, RGB-D, 3D referring datasets and the OCID-Ref in terms of the number of scenes or images (#Scenes/#Images), number of object categories(#Obj. Cat.), Distractor score(Dis. Score), Data format, number of expressions (#Expressions), and average lengths of the expressions(AvgLen). Our OCID-Ref is the first dataset featuring both 3D signals and object occlusion, which are both crucial for visual grounding for HRI.

	Exemplar Templates
Common Sentence	The <Attr> <Obj>. The <Attr> object / item.
Relational Sentence	The <Obj> <Rel>. The <Attr> <Obj> <Rel>. The <Attr> <Obj1> <Rel> <Obj2>.

Table 2: Show the exemplar templates we use to generate the free-form referring expressions.

tions (e.g., color relation, spatial relation, etc.) for all the objects in dataset. We design an online web-based annotation tool to collect these extra labels, and dispatch the labeling tasks over the annotation specialists from a professional data service company. Additionally, we ensure each task is randomly assigned to three trained workers and verified by one checker. The overall tasks take around two months to finish.

2.2 Referring Expression Generation

Gathering the labels we annotated and following the method from the scene-graph based referring expression generation system. In detail, first, we build up the scene graph for each scene in OCID-Ref, and the nodes and edges in the graph represent the attributes and relations, respectively. Second, we design several textual templates (Table 2) to have various sentence structures. Third, we leverage the conventional incremental algorithm(Dale and Reiter, 1995) and functional programs to generate reasonable REs. That is, we add attributes and relations into our conditional set until it conforms with the specific unambiguous condition. Finally, we generate the total of 305,694 referring expressions with an average length of 8.56, and for details, there are an average of 14.71 expressions per object instance and 113.07 expressions per scene.

2.3 Dataset Statistics

OCID-Ref uses the same scenes as OCID, containing 2D object segmentation and 3D object bounding boxes for 2300 fully built-up indoor cluttered scenes. Each object is associated with more than 20 relationships with other objects in the same scene, including 3D spatial relations, 2D spatial relations, comparative color relations, and comparative size relations. Table 1 shows the basic statistic comparison of the previous 2D, RGB-D,3D referring datasets and the OCID-Ref. To evaluate the difficulty of REC, we follow Cops-Ref to calculate the number of candidate objects of the same categories as the target object(Distractor score) for all scenes. Though there are only 3.36 same candidates in an average of OCID-Ref, lower than 4.64 of ScanRefer, we attribute this difference to the dataset characteristic that our scenes are components of one by one sequence with few objects in the first few scenes. To evaluate the referring performance from no clutter to dense clutter scenes, we follow OCID to separate the scenes into three cluttered levels, free, touching, and stacked, from clearly separated to physically touching to being on top of each other. We also split the val split of ScanRefer into three clutter level.

3 Experiments

We conduct referring expression segmentation experiments on our collected OCID-Ref dataset and ScanRefer (Chen et al., 2020a) dataset. We compare different modalities, clutter levels, and regular expression lengths and provide a comprehensive analysis to pave a new path for future research. We also conduct the grasp experiment using different modality data as input, and the details are described in Appendix A.

	All	Free(F)	Touching(T)	Stacked(S)	Decrease rate(F->S) ↓
2D	0.512 / 0.501	0.673 / 0.660	0.450 / 0.496	0.450 / 0.433	33.23% / 34.29%
3D	0.588 / 0.580	0.745 / 0.751	0.589 / 0.583	0.507 / 0.489	31.97% / 34.92%
Fusion	0.634 / 0.637	0.763 / 0.769	0.64 / 0.651	0.551 / 0.540	27.71% / 29.75%

Table 3: Referring expression segmentation performance (pr@0.25) on OCID-Ref. The performance is negatively associated with the clutter level, indicating that the occlusion could make the VG task more challenging. Also, leveraging 3D signals enhances the overall performance on all clutter levels.




Figure 3: Compute the angles related to 3D relation on spherical coordinates.

$$R(\theta, \phi) = \begin{cases} 4 & \text{if } \theta < 15 \\ 5 & \text{if } \theta > -15 \\ 16 + r(\phi) & \text{if } -65 < \theta < -15 \\ 8 + r(\phi) & \text{if } 15 < \theta < 65 \\ r(\phi) & \text{otherwise} \end{cases} \quad (1)$$

$$r(\phi) = 6 + \left(\frac{\phi + 22.5}{45} \right) \quad (2)$$

3.1 Setup

Baseline We run our experiments with a modern graph-based DGA (Yang et al., 2019) model. We compare 2D (RGB), 3D (point cloud) and 2D+3D input signals.

Feature Extraction For 2D inputs, we use ResNet-101 based Faster-RCNN as our 2D feature extractor and pre-train the extractor on OCID to extract the ROI features from the pool5 layer as the 2D visual features, and use the original DGA’s settings for node feature and edge feature on the graph. For 3D inputs, we utilize point-wise features extracted from PointNet (Charles et al., 2017) as the 3D version of the visual feature for each node in the graph. Also, we change the box information from 2D to 3D with box center, box bounds, and

box volume. The relations for the edges are modified with 3D relationships between objects instead of 2D relationships. Figure 3 and equation 1, 2 shows how we compute the angles related to 3D relation on spherical coordinates.

2D and 3D Fusion To utilize advantages from both 2D and 3D signals, we implement a handy fusion module. We take max-pooling on the point features to aggregate them into a global scene feature and concatenate it to the 2D visual feature as a new visual feature for each object instance. Afterward, we fuse the box information into (2D box center, 2D box bounds, 2D box area, 3D box center, 3D box bounds, 3D box volume) to preserve the location information from two distinct coordinates. The edge representation is defined as the same as the 3D version.

Evaluation Metric We use Acc@0.25IoU as our metric to measure the thresholded accuracy where the positive predictions have a higher intersection over union (IoU) with the ground truths than the thresholds.

3.2 Quantitative Analysis

Clutter Levels Table 3 compares 2D (RGB), 3D (point cloud) models and Fusion model performance on OCID-Ref dataset. Obviously, all models struggle against the highly occluded *stacked* subset (Fourth column). The 27 to 34 % of performances drop from free to stacked subset indicates that occlusion, which occurs in working environments, is a challenge for modern VG models. Table 4 shows model performance on ScanRefer dataset, and the result is consistent with OCID-Ref dataset, where stacked performance is dropped from 0.465 to 0.320 for the unique scenario and from 0.198 to 0.131 for the multiple scenario. The results suggest that tackling occlusion is crucial for future research and applications in working environments.

Input Modality As shown in table 3, for single modality models, the 3D model (Second row) con-

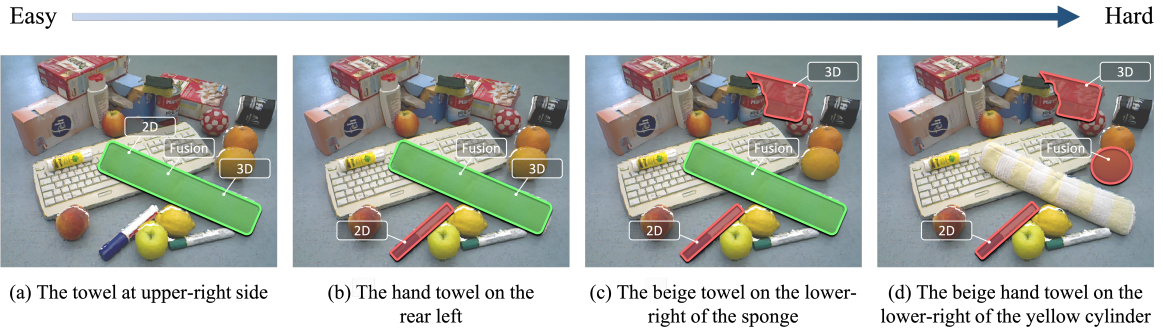


Figure 4: Qualitative results from 2D, 3D, and the fusion methods. Predicted masks with an IOU score higher than 0.25 are marked in green, otherwise in red. Examples are tested in the same cluttered scene with referring expressions in different difficulty levels. Fusion method produces better results than 2D and 3D method.

stantly outperforms the 2D model (First row) in all clutter levels and indicates that accurate spatial information is crucial. Furthermore, aggregating 2D and 3D signals (Third row) reaches the best performance and suffers less performance drop from free to stacked. Therefore, we suggest future work to explore an effective way to utilize and fuse 2D and 3D signals to tackle our challenging dataset.

Referring Expression Length Table 5 compares the performance of short (not more than 12 wordpieces) and long (equal or more than 12 wordpieces). We observe that all models perform worse when the expressions are long.

3.3 Qualitative Analysis

Figure 4 shows results produced by 2D, 3D baseline, and the fusion model. First, in figure 4-d we discover that all three methods fail when the RE is long and complicated. The fusion method successfully localizes the towel in the scene with 2D and 3D spatial descriptions (refer to figure 4-c), while the 3D method has difficulty identifying what is "lower-right." Unsurprisingly, we observe that the 2D method fails on the query with the 3D relation "rear" (refer to figure 4-b). Figure 4-d also shows the failure cases of the fusion method, indicating that our model cannot handle all spatial relations to distinguish between ambiguous objects. 2D and 3D get better performance when the query RE consisted mainly of the common sentences and relationships regarding the whole scene. The failure case suggests that our fusion and localization module can still be improved to utilize the 2D information better and decrease the 3D features' misuse.

	Free	Touching	Stacked
unique	0.465	0.407	0.32
multiple	0.198	0.179	0.131

Table 4: The performance on ScanRefer in different clutter level.

	2D	3D	Fusion
short	0.508	0.592	0.645
long	0.484	0.562	0.580

Table 5: Referring expression segmentation performance on different length of the referring expression.

4 Conclusion

In this work, we propose a novel OCID-Ref dataset for VG with both 2D (RGB) and 3D (point cloud) and occluded objects. OCID-Ref consists of 305,694 referring expressions from 2,300 scenes with providing RGB image and point cloud inputs. Experimental results demonstrate the difficulty of occlusion and suggest the advantages of leveraging both 2D and 3D signals. We are excited to pave a new path for VG researches and applications.

5 Acknowledgement

This work was supported in part by the Ministry of Science and Technology, Taiwan, under Grant MOST 110-2634-F-002-026 and Qualcomm Technologies, Inc. We benefit from NVIDIA DGX-1 AI Supercomputer and are grateful to the National Center for High-performance Computing. We also thank Yu-Kai Huang, Hsin-Ying Lee for his insightful suggestion on the figures.

References

- R. Charles, Hao Su, Kaichun Mo, and Leonidas Guibas. 2017. [Pointnet: Deep learning on point sets for 3d classification and segmentation](#). pages 77–85.
- Dave Zhenyu Chen, Angel X Chang, and Matthias Nießner. 2020a. [Scanrefer: 3d object localization in rgb-d scans using natural language](#). *16th European Conference on Computer Vision (ECCV)*.
- Zhenfang Chen, Peng Wang, Lin Ma, Kwan-Yee K Wong, and Qi Wu. 2020b. [Cops-ref: A new dataset and task on compositional referring expression comprehension](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10086–10095.
- Zhenfang Chen, Peng Wang, Lin Ma, Kwan-Yee K Wong, and Qi Wu. 2020c. [Cops-ref: A new dataset and task on compositional referring expression comprehension](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Robert Dale and Ehud Reiter. 1995. [Computational interpretations of the gricean maxims in the generation of referring expressions](#). *Cognitive Science*, 19(2):233–263.
- Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara L. Berg. 2014. [Referit game: Referring to objects in photographs of natural scenes](#). In *EMNLP*.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. [Visual genome: Connecting language and vision using crowdsourced dense image annotations](#). *International journal of computer vision*, 123(1):32–73.
- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan Yuille, and Kevin Murphy. 2016. [Generation and comprehension of unambiguous object descriptions](#). In *CVPR*.
- Cynthia Matuszek. [Grounded language learning: Where robotics and nlp meet](#).
- Cecilia Mauceri, Martha Palmer, and Christoffer Heckman. 2019. [Sun-spot: An rgb-d dataset with spatial referring expressions](#). In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0.
- Maria Ralph and Medhat Moussa. 2005. [Human-robot interaction for robotic grasping: A pilot study](#). In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 454–459. IEEE.
- Mohit Shridhar and David Hsu. 2018. [Interactive visual grounding of referring expressions for human-robot interaction](#). In *Proceedings of Robotics: Science and Systems*.
- Markus Suchi, Timothy Patten, David Fischinger, and Markus Vincze. 2019. [Easylab: A semi-automatic pixel-wise object annotation tool for creating robotic RGB-D datasets](#). In (Suchi et al., 2019), pages 6678–6684.
- Sibei Yang, Guanbin Li, and Yizhou Yu. 2019. [Dynamic graph attention for referring expression comprehension](#).
- Sibei Yang, Guanbin Li, and Yizhou Yu. 2020. [Graph-structured referring expressions reasoning in the wild](#).
- Licheng Yu, Patrick Poirson, Shan Yang, Alexander C. Berg, and Tamara L. Berg. 2016. [Modeling context in referring expressions](#). In *Computer Vision – ECCV 2016*, pages 69–85, Cham. Springer International Publishing.