# IVLMap: Instance-Aware Visual Language Grounding for Consumer Robot Navigation

Jiacui Huang, Hongtao Zhang, Mingbo Zhao, Senior, IEEE, Wu Zhou, Senior, IEEE

Abstract-Vision-and-Language Navigation (VLN) is a challenging task that requires a robot to navigate in photo-realistic environments with human natural language promptings. Recent studies aim to handle this task by constructing the semantic spatial map representation of the environment, and then leveraging the strong ability of reasoning in large language models for generalizing code for guiding the robot navigation. However, these methods face limitations in instance-level and attribute-level navigation tasks as they cannot distinguish different instances of the same object. To address this challenge, we propose a new method, namely, Instance-aware Visual Language Map (IVLMap), to empower the robot with instance-level and attribute-level semantic mapping, where it is autonomously constructed by fusing the RGBD video data collected from the robot agent with special-designed natural language map indexing in the bird's-in-eye view. Such indexing is instancelevel and attribute-level. In particular, when integrated with a large language model, IVLMap demonstrates the capability to i) transform natural language into navigation targets with instance and attribute information, enabling precise localization, and ii) accomplish zero-shot end-to-end navigation tasks based on natural language commands. Extensive navigation experiments are conducted. Simulation results illustrate that our method can achieve an average improvement of 14.4% in navigation accuracy. Code and demo are released at https://ivlmap.github.io/.

Index Terms—Vision-and-Language Navigation, 3D Reconstruction, Zero-Shot Navigation, Visual Language Map

#### I. INTRODUCTION

**D**EVELOPING a robot that can cooperate with humans is of great importance in many real-world applications, such as Telsa Optimus, Mobile ALOHA [1], etc. A robot agent that can understand human language and navigate intelligently is therefore significant to benefit human society. Towards this end, Vision-and-Language Navigation (VLN) has been developed during the past few years, which is to empower a robot agent to navigate in photo-realistic environments according to natural language instructions, such as "navigate to the 3rd chair" or "navigate to the yellow sofa" [2], [3]. This requires the robot agent can interpret natural language from humans, perceive its visual environment, and utilize the information to navigate, where a key important issue is how to structure the visited environment and make global planning. To handle this case, a few recent approaches [4] utilize the topological map for structure. But these methods are difficult to represent the spatial relations among objects, causing detailed information may be lost; on the other hand, more recent works [5] model the navigation environment using the top-down semantic map,

This paper was produced by the IEEE Publication Technology Group. They are in Piscataway, NJ.

\*Corresponding Author

which represents spatial relations more precisely. However, the semantic concepts are extremely limited due to the pre-defined semantic labels.

In general, based on the environments for navigation, the VLN can be roughly divided into two categories, which include navigation in discrete and continuous environments. In discrete environments, such as those in R2R [2], REVERIE [6], and SOON [7], the VLN is conceptualized as a topology structure comprised of interconnected navigable nodes. The agent utilizes a connectivity graph to move between adjacent nodes by selecting a direction from the available navigable directions; in contrast, VLN in continuous environments, such as those in R2R-CE [8] and RxRCE [9], enable the agents to have the flexibility to navigate to any unobstructed point using a set of low-level actions (e.g., move forward 0.25m, turn left 15 degrees) instead of teleporting between fixed nodes. This approach is closer to real-world robot navigation, posing a more intricate challenge for the agent.

Take the instance illustrated in Fig.1, where the language instruction is to "navigate to the fourth black chair across from the table". To execute this command, we initially explore the entire room to identify all instances of tables, extracting their color attributes. Subsequently, we locate the position of the fourth table with the specified black color mentioned in the command. Completing such a navigation task becomes easily achievable if we have a global map of the scene, and the map should encompass information about each object, including its category, details, color, etc. Recently, VLMap [5] pioneers innovative zero-shot spatial navigation by constructing Semantic maps via indexing visual landmarks using natural language. However, VLMap is limited to navigating to the vicinity of the closest category to the robotic agent and cannot fulfill the envisioned, more common, and precise instance-level and attribute-level navigation needs in real-life scenarios.

In this paper, to address this challenge, we propose a new method, namely, Instance-aware Visual Language Map (IVLMap), to empower the robot with instance-level and attribute-level semantic mapping, where the IVLMap is autonomously constructed by fusing the RGBD video data with a specially-designed natural language map indexing in the bird's-in-eye view. Such indexing is instance-level and attribute-level. In this way, IVLMap can well separate different instances within the same category. When integrated with a large language model, it demonstrates the capability to i) transform natural language into navigation targets with instance and attribute information, enabling precise localization, and ii) accomplish zero-shot end-to-end navigation tasks based on natural language commands. The main contributions of the



Fig. 1. Under the guidance of our IVLMap, robotic agents can accomplish instance-level target navigation. Leveraging the map's information, the robot can navigate to specific objects based on their instance attributes(color, shape, etc). This capability enables the robotic agent to execute navigation tasks with a higher degree of precision, enhancing its ability to reach designated objects accurately.

proposed work are as follows:

- a novel instance-level and attribute-level map in bird'sin-eye review is developed for better comprehend the environment and robot navigation planning, where the instances in map are obtained by involving region matching and label scoring to achieve the category label for each SAM [10]-segmented mask, while the color attribute labels for each mask are also obtained by using the similar approach.
- 2) leveraging IVLMap, we propose a two-step method for locating landmarks. This involves an initial coarsegrained localization of the corresponding mask, followed by a fine-grained localization and navigation within the mask's designated region.
- 3) we established an interactive data collection platform, achieving real-time controllable data acquisition. This not only reduced data volume but also enhanced reconstruction efficiency. Furthermore, we conducted experiments in a real-world environment, providing valuable insights for the practical deployment of this method.

The rest of this work is organized as follows: In Section II, we will review some related work; in Section III, we will give the detailed description for presenting the proposed IVLMAP. Extensive simulations are conducted in Section IV and the final conclusions are drawn in Section V.

# II. RELATED WORK

**Semantic Mapping.** Advances in semantic mapping, driven by the integration of Convolutional Neural Networks (CNNs) [11] and dense Simultaneous Localization and Mapping (SLAM) [12], have progressed significantly. SLAM++ [13] introduced an object-oriented SLAM approach leveraging preexisting knowledge. Subsequent studies, e.g., [14], use Mask-RCNN for instance-level semantics in 3D volumetric maps. VLMaps [5] and NLMaps-Saycan [15] introduce scene representations queryable through natural language, using Visual Language Models (VLMs). While prior studies focus on semantic details, our work emphasizes pixel-level segmentation accuracy, categorizing each pixel precisely.

**Instance Segmentation.** Achieving instance-level segmentation is crucial in VLN, requiring precise identification and location of individual instances of similar objects. Innovative research [16] addresses human instance segmentation without distinct detection stages. Real-time solutions [17] introduce a Spatial Attention-Guided Mask (SAG-Mask) branch. Meta AI's Segment Anything Model (SAM) [10] excels in promptable segmentation, enabling zero-shot generalization to unfamiliar objects. SAM's object mask segmentation, while powerful, lacks individual mask labeling, limiting practical applications.

Vision-and-Language Navigation(VLN). Recent years have seen significant strides in VLN, driven by researchers such as Y. Long et al., who employed large-scale training to imbue extensive domain knowledge in models like VIM-Net [18]. Addressing fusion challenges, VIM-Net critically matches visual and linguistic information for accurate navigation [19]. Innovations like Talk2Nav, with dual attention mechanisms and spatial memory, tackle long-range VLN [20]. Challenges persist, including navigating unseen objects, interpreting language descriptions precisely, personalizing navigation for object attributes, and managing data-intensive requirements for end-to-end navigation [5], [21], [22].

LLM used in VLN. Large Language Models (LLMs) play a pivotal role in VLN, evident in NavGPT's reasoning capabilities [23]. LM-Nav uses GPT-3 for parsing instructions and navigating landmarks [4]. VELMA extends LLMs to streetlevel navigation, enriching AI in complex urban environments [24]. Studies like Vemprala et al.'s integration of ChatGPT into robotics applications and PaLM-E's multimodal reasoning with visual and textual inputs [25], [26] influence our project, where we achieved accurate navigation using LLMs guided by natural language.

#### III. METHOD

In our work, our objective is to construct a semantic map of the surrounding environment which also encompasses instance-level information and attribute information of objects. It is crucial to emphasize the inclusion of instance-level semantic information and object attributes in maps. This incorporation is vital for effectively processing linguistic commands commonly utilized in everyday language. In everyday life, robots commonly encounter commands like "navigate to the fourth black chair across from the table". Robots are required to identify "which of the chair" both in terms of its sequence



Fig. 2. The IVLMap pipeline consists of two main components. The first focuses on 3D reconstruction and constructing a visual language map. Building on this foundation, the second part integrates the Segment Anything Model (SAM), enhancing the map's representation with a segmentation-aware approach for more detailed information.

and attributes. Simultaneously, they are required to find "where is the chair" by paying attention to the spatial relationship between the chair and the table. Our method is proposed based on VLMap [5]. In the following subsections, we describe (i) how to build IVLMap (Sec.III-A), (ii) how to use these maps to localize open-vocabulary landmarks (Sec.III-B), (iii) How to employ VLMaps in tandem with large language models (LLMs) to enable instance level object goal navigation using natural language commands (sec.III-C). Our pipeline is visualized in Fig.2.

# A. IVLMap Creation

Construction of VLMap and 3D Reconstruction Map in Bird's-Eye View: The fundamental concept of VLMap involves integrating pre-trained visual language features into a three-dimensional (3D) reconstruction map. In our work, we utilize LSeg [27] as the visual-language model, a languagedriven semantic segmentation model that segments the RGB images based on a set of free-form language categories. As described in VLMap, VLMap is defiened as  $\mathcal{M} \in \mathbb{R}^{\tilde{H} \times \tilde{W} \times C}$ , where  $\tilde{H}$  and  $\tilde{W}$  represent the size of the top-down grid map, and C represents the length of the VLM embeddings vector for each grid cell. In the matrix  $\mathcal{M}$ , each cell encapsulates category information corresponding to pixels, which will furnish semantic support for the construction of our IVLMap.

In order to get the 3D Reconstruction Map in Bird's-Eye View, we, in line with the approach employed by VLMaps, obtain the corresponding coordinates of each pixel u in the RGB-D data frame within the grid map by formula

$$p_{\rm map}^x = \left\lfloor \frac{\bar{H}}{2} + \frac{P_W^x}{s} + 0.5 \right\rfloor, p_{\rm map}^y = \left\lfloor \frac{\bar{W}}{2} - \frac{P_W^z}{s} + 0.5 \right\rfloor$$
(1)

where  $\mathbf{P}_k \in \mathbb{R}^3$  is the pixel's 3D point position in the k-th frame,  $\mathbf{P}_w \in \mathbb{R}^3$  is the pixel's 3D point position in the world coordinate frame,  $p_{map}^x$  and  $p_{map}^y$  represent the coordinates of the projected point in the map  $\mathcal{M}$ . We define the 3D Reconstruction Map in Bird's-Eye View as  $\mathcal{B} \in \mathbb{R}^{\bar{H} \times \bar{W} \times 3}$ , where each cell  $\beta_{ij}$  at the coordinates of  $(p_{map}^x, p_{map}^y)$  denotes the color of the corresponding position in RGB image  $\mathcal{I}_k$ . In the context of navigation, robots only need to focus on obstacles that are approximately at their height or lower (assuming the height of the robot is denoted as h). Objects significantly taller than the robot, such as ceilings or ceiling fans, do not require much attention. Intuitively, there exist multiple 3D points projecting to the same grid location in the map. We define matrix  $\mathcal{H} \in \mathbb{R}^{\bar{H} \times \bar{W} \times 1}$  to represent the projection height of each pixel in every cell, it continuously updates with the processing of RGB-D data. Only pixels with a height lower than the robot's height h are considered. The descending update method is employed, causing the values in each cell of  $\mathcal{H}$  to gradually decrease during the reconstruction process. This ensures that all objects observed by the robot are taken into account in 3D Reconstruction Map  $\mathcal{B}$ .

Integrating Instance-level information to IVLMap: Formally, IVLMap is defined as  $\mathscr{M} = \{\mathscr{M}, \mathscr{U}, \mathscr{V}\}\)$ , where  $\mathscr{M}$ represents the VLMap (Sec.III-A Para-1),  $\mathscr{U} \in \mathbb{R}^{\bar{H} \times \bar{W}}$  holds instance information (represented by IDs), and  $\mathscr{V} \in \mathbb{R}^{\bar{H} \times \bar{W}}$ contains color information. Following Kirillov et al. [10], we utilize the pre-trained SAM model to segment the 3D reconstructed map  $\mathscr{B}$ , resulting in segmentation masks  $\bar{\mathcal{S}}$ , each comprising information like mask segmentation, area, and prediction accuracy. The key to achieving instance-level visual language maps is appending relevant attributes to each mask, effectively representing individual instance objects. The list of masks is denoted as  $\mathscr{S} = [\mathscr{S}_0, \mathscr{S}_1, \cdots, \mathscr{S}_N] \in \mathbb{R}^{\bar{H} \times \bar{W} \times N}$ , where each  $Si \in 0, 1^{\bar{H} \times \bar{W}}$  is the segmentation mask for the i-th instance. In the Pixel-Text Similarity Matrix P (Sec.III-B), each cell corresponds to a vector  $Pij = [p_0, p_1, \dots, p_N]$  of length N, with  $p_k$  representing the probability that the pixel at the current position belongs to the k-th category. The argmax<sup>1</sup> operation is applied to P, resulting in the Pixel-Label Map  $\bar{P} \in \mathbb{R}^{\bar{H} \times \bar{W} \times 1}$ , where each cell represents the label index of the corresponding pixel.

Due to data collection errors and model prediction inaccuracies, noisy points may exist in P, and instances belonging to the same category cannot be distinguished. Therefore, we employ SAM for segmenting the 3D Reconstruction Map in Bird's-Eye View and apply smoothing corrections for accurate instance segmentation. The regions in mask Si with a value of 1 represent segmented instances, and their category information is reflected in the corresponding regions of P. Through region matching, only the regions in  $\overline{P}$  that align with  $S_i$  are retained, while the rest are set to zero, yielding a modified matrix  $\overline{P'}$ . Following the approach by Chen et al, who used a top-k scoring model for obtaining category labels for SAM segmentation masks, we conduct a similar Top-k scoring operation on the category IDs in the corresponding regions of  $\overline{P'}$ . The combination of scores determines the category of the current instance. The unique<sup>2</sup> operation on  $\overline{P'}$  efficiently extracts unique labels  $\mathfrak{L} = [\mathfrak{l}_1, \mathfrak{l}_2, \cdots, \mathfrak{l}_n]$  and their corresponding counts  $\Phi = [\varphi_1, \varphi_2, \cdots, \varphi_n]$  within the matched regions, where n represents the total number of unique labels. Our scoring mechanism calculates, for each unique label within the masked region, the ratio of its count to the total number of pixels in that region, resulting in a score ranging from 0 to 1, as indicated by Equation 2.

$$\begin{cases} S^{c} = \frac{[\varphi_{1}, \varphi_{1}, \cdots, \varphi_{n}]}{\sum \mathfrak{C}} \\ \forall i, j \, (i < j \Longrightarrow \varphi_{i} > \varphi_{j}) \end{cases}$$
(2)

where  $S^c \in \mathbb{R}^n$  represents the score of each unique label. The second part of the formula signifies the descending sorting of  $S^c$ . Additionally, the order of elements in  $\mathfrak{L}$  undergoes a corresponding transformation along with the sorting of  $S^c$ . Finally, we perform the matching between category IDs and category strings. Since there is a one-to-one correspondence between  $\mathfrak{L}$  and  $\mathscr{C}$  (Sec.III-B), we only need to match each element in  $\mathfrak{L}$  to the corresponding string in  $\mathscr{C}$ . The entire implementation approach is outlined in Fig3.

Integrating additional attributes into the IVLMap: The essence of constructing the VLMap is to establish a correspondence between each pixel and a category label. Similarly, we can use the same approach to associate each pixel with a color label. Similar to Section III-B, we define the color label as  $\overline{\mathscr{C}} = [\overline{c}_1, \overline{c}_2, ..., \overline{c}_N]$ , continuing with the same operation, we can obtain Pixel-Color-Text Similarity. Naturally, the following steps can be referred to the process outlined in Fig.3 to ultimately obtain the color label attributes for each mask. As each mask corresponds to a unique object instance,



Fig. 3. A schematic diagram illustrating a Matching Algorithm. We matched the masks generated by the SAM model with the Pixel-Text Similarity obtained from VLMap, assigning labels to each mask. This facilitated the subsequent implementation of IVLMap.

the label\_id for each mask can be obtained through a unique counting id.The other attributes of the mask are relatively simple and won't be elaborated here. The final situation of each attribute is presented in Listing 1. Combining the steps mentioned above, we obtain the final IVLMap  $\mathcal{M}$ . It integrates semantic information from the VLMap, instance information from SAM segmentation masks, and color information of instances. This comprehensive map provides convenient assistance for our navigation tasks.

```
"segmentation":
array([[False, ..., False],
    [False, ..., False],
    · · · ,
    [False, ..., False],
    [False, ..., False]]),
"area": 1901,
"bbox": [172, 345, 70, 28],
"predicted_iou": 0.9936274290084839,
"point_coords": [[232.546875, 351.03125]]
"stability score": 0.9828214645385742,
"crop_box": [0, 0, 363, 478],
"label": "floor",
"label_id": 1,
"num_of_same_class": 5,
"color": "yellow"
```

Listing 1. An illustrative instance of IVLMaps' masks featuring diverse attributes in JSON format.

## B. Localizing Open-Vocabulary Landmarks

Despite VLMap's excellent open-vocabulary handling, we've further optimized it, leveraging LSeg's superior natural language-driven semantic segmentation. Precision is crucial in describing object categories during open-vocabulary implementation as LSeg's [27] Text encoder demands encoding potential labels. In daily life, non-experts may struggle to accurately list all object categories, emphasizing the necessity of transitioning from natural language to a precise category list. ChatGPT has established LLM's proven capability in comprehending and responding to human natural language standards [28]. To achieve our objective, we employ Llama2, a cutting-edge open-source language model developed by Meta [29]. In our work, we simply input natural language

<sup>&</sup>lt;sup>1</sup>The argmax operation in PyTorch returns the index of the maximum value along a specified axis in a tensor.

<sup>&</sup>lt;sup>2</sup>The numpy unique method efficiently identifies and returns the unique elements of an array along with the count of occurrences for each unique element, preserving their original order.

phrases(such as "Generate a map for a 1.5m high robot, especially considering items such as beds, sofas, chairs, tables, and other obstacles deemed necessary to consider") into the model. Llama, in conjunction with our contextual information (including object category labels and their sizes in the scene), combines this with the specified output specifications outlined by the prompt engineering to generate a list of object categories as output. Formally, we can define the object categories as  $\mathscr{C} = [c_1, c_2, ..., c_N]$ , where  $c_1$  represents the i-th category in text form, and N represents the total number of categories extracted by LLM. Building upon the foundation laid by Huang et al. [5], we leverage CLIP to encode the list of category characters, resulting in an embedding matrix  $\mathbf{E} \in \mathbb{R}^{M \times C}$ . The embedding matrix  $\mathcal{M}$  of VLMap is flattened into matrix  $\mathbf{Q} \in \mathbb{R}^{\overline{M}\overline{H} \times C}$ , and the pixel-to-category similarity matrix P is obtained through the multiplication  $P = Q * E^{T}$ . Each element  $P_{ij}$  in the matrix represents the similarity score between a pixel and a text category, reflecting the likelihood of the pixel belonging to that specific class. The obtained object categories & and the category similarity matrix P from the aforementioned steps are employed in the construction of IVLMap in section III-A. When locating a specific instance object of landmarks, we first identify the specific mask  $S_i$  from the mask set S. Then, we locate the region information of the mask in  $\mathcal{U}$  and  $\mathcal{V}$  of the IVLMap  $\mathcal{M}$ . Finally, we pinpoint the exact instance position in the VLMap  $\mathcal{M}$ .

# C. Zero-Shot Instance Level Object Goal Navigation from Natural Language

In this chapter, we discuss the implementation of instancelevel object goal navigation based on natural language, given a set of landmark descriptions specified by natural language instructions such as

First navigate to the nearest yellow chair, then navigate to the third table from left to right, and finally navigate to the black sofa.

There are numerous outstanding works in the field of zeroshot navigation, such as CoW [30], LM-Nav [4], VLMap [5]. Unlike their approaches, IVLMap enables navigation to instances of objects with specific attributes, such as "the third chair on the left" or "the nearest black sofa". Particularly, we employ a large language model (LLM) to interpret the given natural language commands and decompose them into subgoals with attributes including object name, object instance information, object color [4], [5], [31], [32]. We have established a comprehensive set of advanced function libraries tailored for IVLMap. Building upon the understanding of natural language commands by the LLM, these libraries are invoked to generate executable Python robot code [25], [33]. Additional mathematical computations may be performed if required. Below are snippets showcasing some commands and the Python code generated by the model(input command in blue, output Python code in black), the full prompt can be referred to in Appendix B.

Additionally, numerous large language models have emerged recently. Apart from utilizing the ChatGPT API for navigation tasks, we also achieved similar results using opensource models. We employ Llama2 [29], an advanced opensource artificial intelligence model, to accomplish this task. We utilized the Llama-2-13b-chat-hf<sup>3</sup> model from the Hugging Face repository. This model is fine-tuned for dialogue use cases and optimized for the Hugging Face Transformers format. To accelerate the model's inference speed, we employed GPTQ [34], a one-shot weight quantization method based on approximate second-order information, to quantize the model from 8 bits to 4 bits. The experiments demonstrate an average improvement of around 30% in the inference speed after quantization.

# navigate to the third yellow table
obj_attr = agent.get_obj_attributes('table',3,'yellow')
agent.face(obj_attr)
# Move to the red sofa then stop at the right side of it
obj_attr = agent.get_obj_attributes('sofa',0,'red')
agent.move_to_right(obj_attr)
# Go straight for 1 meter and then walk clockwise
around the red sofa for 2 turns.
agent.move_forward(1)
obj_attr = agent.get_obj_attributes( 'sofa', 0,'red' )
obj_contour = agent.get_specifed_obj_pos(obj_attr)
agent.face(obj_attr), agent.turn(-90)
for i in range(8):
agent.move_forward(obj_contour([i%4])
agent.turn(90)

#### IV. EXPERIMENT

#### A. Experimental Setup

We employ the Habitat simulator [35] alongside the Matterport3D dataset [2] to assess performance in multi-object and spatial goal navigation tasks. Matterport3D is a comprehensive RGB-D dataset, featuring 10,800 panoramic views from 194,400 RGB-D images across 90 building-scale scenes, designed for advancing research in scene understanding indoor environments. For map creation in Habitat, we capture 13,506 RGB-D frames spanning six distinct scenes and document the camera pose for each frame, utilizing the cmu-exploration environment we established(Sec.IV-B). Due to computational constraints, our navigation experiments and the execution of the Large Language Model(Llama2) are conducted on separate servers. Llama2 operates on a server equipped with two NVIDIA RTX 3090 GPUs. As for IVLMap experiment, our experimental setup comprises an NVIDIA GeForce RTX 2080 Ti GPU with 12GB VRAM. Communication between these two servers is established using the Socket.IO<sup>4</sup> protocol.

Baseline: We assess IVLMap in comparison to three baseline methods, all employing visual-language models and

<sup>&</sup>lt;sup>3</sup>https://huggingface.co/meta-llama/Llama-2-13b-chat-hf

<sup>&</sup>lt;sup>4</sup>Socket.IO is a real-time communication protocol built on WebSocket, providing event-driven bidirectional communication for seamless integration of interactive features in web applications, official website https://socket.io/.

demonstrating proficiency in zero-shot language-based navigation:

- VLMap [5] seamlessly integrates language and visuals, autonomously constructing maps, and excels in indexing landmarks from human instructions, enhancing language-driven robots for intuitive communication in diverse navigation scenarios with open-vocabulary mapping and natural language indexing.
- 2) Clip on Wheels (CoW) [30] achieves language-driven object navigation by creating a target-specific saliency map using CLIP and GradCAM [36]. It involves applying a threshold to saliency values, extracting a segmentation mask, and planning the navigation path based on this information.
- The CLIP-features-based map (CLIP Map) serves as an ablative baseline, projecting CLIP visual features onto the environment's feature map and generating object category masks through thresholding the feature similarity.

**Evaluation Metrics:** Similar to previous methods [5], [37] in VLN literature, we use the standard Success Ratemetric(SR) to measure the success ratio for the navigation task. We assessed our IVLMap's effectiveness by (i) presenting multiple navigation targets and (ii) using natural language commands. Success is defined as the agent stopping within a predefined distance threshold from the ground truth object.

#### B. Dataset acquisition and 3D Reconstruction

To construct a map in visual and language navigation tasks, it is crucial to acquire RGB images, depth information, and pose data from the robot or its agent while it is in motion. Common datasets on the internet, such as Matterport3D [2], Scannet [38], KITTI [39], may not be directly applicable to our scenario. Therefore, we undertook the task of collecting a dataset tailored to our specific requirements. Our data collection efforts were conducted in both virtual and real(Appendix.C) environments to ensure comprehensive coverage.



Fig. 4. We created an Interactive Dataset Collection Scheme by combining the cmu-exploration development environment with the Habitat simulator. This involves integrating cmu-exploration's autonomous exploration with Habitat robot agents for a unified dataset collection approach.

Interactive Data Collection in Virtual Habitats and CMU-Exploration Environment. CMU-Exploration [40], designed for autonomous navigation system development, offers various simulation environments and modules. Combined with the Habitat Simulator [35], it forms a platform where users develop and deploy navigation systems for real robots. Our interactive data collection system in this virtual environment utilizes CMU-Exploration's Joystick and Visualization Tools for waypoint setting. Waypoints undergo local planning, terrain and radar analysis, and state estimation, generating control commands for ROS-based robot motion. Simultaneously, robot pose data is sent to the Habitat simulator, providing RGB, depth, and pose information for direct sensor control. See Fig. 8 for an overview.

Compared to other black-box data collection methods in the Habitat simulator, where predefined routes or exploration algorithms are used, this approach offers strong controllability. It allows tailored responses to the environment, enabling the collection of fewer data points while achieving superior reconstruction results. For the same scene in Matterport3D, our approach achieves comparable results to the VLMap's original authors while reducing the data volume by approximately 8%. In certain areas, the reconstruction performance even surpasses that of the original authors. To compare the results, refer to Fig.5(a) and Fig.5(b), for more detailed results of our 3D reconstruction bird's-eye view, please refer to Appendix.D.



(a) IVLMap (Ours) (b) VLMap (baseline) Fig. 5. 3D Reconstruction Map in Bird's-Eye View

#### C. Multi-Object Navigation with given subgoals

To assess the localization and navigation performance of IVLMap, we initially conducted navigation experiments with given subgoals. In the navigation experiments conducted in the four scenes of the Matterport dataset, we curated multiple navigation tasks for each scene. Each navigation task comprises four subgoals. The robot is instructed to sequentially navigate to each subgoal of each task. We use the invocation of the "stop" function by the robot as a criterion. If the robot calls the "stop" function and its distance to the target is less than a threshold (set to 1 meter in our case), it is considered successful navigation to that subgoal. Successful completion of a navigation task is achieved when the robot successfully navigates to all four subgoals in sequence. It is noteworthy that we provided instance information of objects in the given subgoals to examine the effectiveness of our constructed IVLMap.

Our observations(Table.I) indicate that our approach outperforms all other baselines. It exhibits a slight improvement

Method	5LpN3gDmAk7 (subgoals: 28, tasks: 7				gTV8FGcVJC9 (subgoals: 36, tasks: 9					UwV83HsGsw3 (subgoals: 40, tasks: 10				Vt2qJdWjCF2 (subgoals: 40, tasks: 10										
	SN	SR	T_1	T_2	T_3	T_4	SN	SR	T_1	T_2	T_3	T_4	SN	SR	T_1	T_2	T_3	T_4	SN	SR	T_1	T_2	T_3	T_4
CoW	10	0.36	0.29	0.14	0.14	0.00	13	0.33	0.56	0.44	0.17	0.11	13	0.33	0.30	0.10	0.10	0.00	13	0.33	0.30	0.20	0.10	0.00
CLIP Map	8	0.29	0.29	0.14	0.00	0.00	11	0.33	0.56	0.33	0.11	0.11	11	0.28	0.20	0.10	0.00	0.00	12	0.30	0.20	0.10	0.00	0.00
VLMAP	15	0.54	0.43	0.29	0.14	0.00	20	0.56	0.67	0.56	0.22	0.22	17	0.43	0.60	0.30	0.30	0.10	18	0.45	0.50	0.40	0.10	0.00
IVLMap(Ours)	18	0.64	0.57	0.29	0.14	0.00	23	0.67	0.78	0.56	0.25	0.22	19	0.48	0.60	0.40	0.30	0.10	20	0.50	0.60	0.40	0.20	0.10
TABLE I																								

OUTCOME OF MULTI-OBJECT NAVIGATION WITH SPECIFIED SUBGOALS, DENOTED BY SN FOR SUCCESS NUMBER, SR FOR SUCCESS RATE, T\_K FOR ACHIEVING THE KTH SUBGOAL OUT OF THE TOTAL 4 SUBGOALS IN EACH TASK, AND TSR FOR TASK SUCCESS RATE, NAMELY T\_4.

in navigation performance compared to VLMap, while significantly surpassing the performance of CoW and CLIP Map. VLMap achieves zero-shot navigation by smoothly performing precise localization of landmarks. However, this baseline has limitations as it can only navigate to the nearest category to the robot agent, lacking the capability for precise instantiation navigation. As illustrated in the comparisons between Fig.6(a) and Fig.6(b), our proposed IVLMap incorporates instantiation information for each landmark, enabling precise instantiation navigation tasks. Consequently, the final navigation performance is significantly enhanced. During specific localization, our approach involves initially identifying the approximate region of the landmark from the  $\mathcal{U}$  and  $\mathcal{V}$  matrices of IVLMap  $\mathcal{M}($ Sec.III-A). Subsequently, further refinement is conducted using VLMap, optimizing the performance of VLMap and resulting in a significant improvement in navigation accuracy.



Fig. 6. Semantic segmentation results

Method	T_1	T_2	T_3	T_4				
CoW	0.39	0.22	0.08	0.03				
CLIP Map	0.33	0.19	0.06	0.0				
VLMap	0.56	0.39	0.19	0.08				
IVLMap (Ours) 0.64 0.44 0.25 0.14								
TABLE II								

OUTCOME OF ZERO-SHOT INSTANCE LEVEL OBJECT GOAL NAVIGATION FROM NATURAL LANGUAGE. T\_K FOR ACHIEVING THE KTH SUBGOAL OUT OF THE TOTAL 4 SUBGOALS IN EACH TASK.

# D. Zero-Shot Instance Level Object Goal Navigation from Natural Language

In these experiments, we assess IVLMaps' performance in comparison to alternative baselines concerning zero-shot instance-level object goal navigation initiated by natural language instructions. Our benchmark comprises 36 trajectories across four scenes, each accompanied by manually provided language instructions for evaluation purposes. In each language instruction, we provide instantiation and color information for navigation subgoals using natural language, such as "the first yellow sofa", "in between the chair and the sofa" or "east of the red table." Leveraging LLM, the robot agent extracts this information for localization and navigation. Each trajectory comprises four subgoals, and successful navigation to the proximity of a subgoal within a threshold range (set at 1m) is considered a success.



Fig. 7. Zero shot navigation diagram, where green dot represents the starting point and red dot represents the endpoint.

Analysis of Table.II reveals that in zero-shot level object goal navigation, our navigation accuracy is hardly affected. This is attributed to the initial parsing of natural language instructions using LLM, enabling precise extraction of physical attributes, ensuring robust performance in navigation. Moreover, as depicted in partial trajectory schematics of navigation tasks in Fig.7, our IVLMap achieves precise instance-level object navigation globally, a capability unmatched by other baselines.

# V. CONCLUSION

In this study, we introduce the Instance Level Visual Language Map (IVLMap), elevating navigation precision through instance-level and attribute-level semantic language instructions. Our approach is designed to enhance applicability in real-life scenarios, showing promising results in initial realworld robot applications. However, the mapping performance in dynamic environments requires improvement, prompting the exploration of real-time navigation using laser scanners. Our future goals include advancing towards 3D semantic maps to enable dynamic perception of object height, contributing to more accurate spatial navigation. Ongoing research efforts will focus on addressing these challenges.

# ACKNOWLEDGMENT

The work is supported by the National Natural Science Foundation of China (no. 61601112). It is also supported by the Fundamental Research Funds for the Central Universities and DHU Distinguished Young Professor Program.

#### REFERENCES

- Z. Fu, T. Z. Zhao, and C. Finn, "Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation," in *arXiv*, 2024.
- [2] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, I. Reid, S. Gould, and A. Van Den Hengel, "Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments," in *Proceedings of the IEEE conference on computer* vision and pattern recognition, 2018, pp. 3674–3683.
- [3] J. Gu, E. Stefani, Q. Wu, J. Thomason, and X. E. Wang, "Vision-andlanguage navigation: A survey of tasks, methods, and future directions," *arXiv preprint arXiv:2203.12667*, 2022.
- [4] D. Shah, B. Osiński, S. Levine *et al.*, "Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action," in *Conference* on Robot Learning. PMLR, 2023, pp. 492–504.
- [5] C. Huang, O. Mees, A. Zeng, and W. Burgard, "Visual language maps for robot navigation," in 2023 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2023, pp. 10608–10615.
- [6] Y. Qi, Q. Wu, P. Anderson, X. Wang, W. Y. Wang, C. Shen, and A. v. d. Hengel, "Reverie: Remote embodied visual referring expression in real indoor environments," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9982–9991.
- [7] F. Zhu, X. Liang, Y. Zhu, Q. Yu, X. Chang, and X. Liang, "Soon: Scenario oriented object navigation with graph-based exploration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 12 689–12 699.
- [8] Z. Wang, J. Li, Y. Hong, Y. Wang, Q. Wu, M. Bansal, S. Gould, H. Tan, and Y. Qiao, "Scaling data generation in vision-and-language navigation," in *Proceedings of the IEEE/CVF International Conference* on Computer Vision, 2023, pp. 12009–12020.
- [9] K. He, Y. Huang, Q. Wu, J. Yang, D. An, S. Sima, and L. Wang, "Landmark-rxr: Solving vision-and-language navigation with finegrained alignment supervision," *Advances in Neural Information Processing Systems*, vol. 34, pp. 652–663, 2021.
- [10] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo *et al.*, "Segment anything," *arXiv preprint arXiv:2304.02643*, 2023.
- [11] K. Fukushima, "Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position," *Biological cybernetics*, vol. 36, no. 4, pp. 193–202, 1980.
- [12] P. Estefo, J. Simmonds, R. Robbes, and J. Fabry, "The robot operating system: Package reuse and community dynamics," *Journal of Systems* and Software, vol. 151, pp. 226–242, 2019.
- [13] R. F. Salas-Moreno, R. A. Newcombe, H. Strasdat, P. H. Kelly, and A. J. Davison, "Slam++: Simultaneous localisation and mapping at the level of objects," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2013, pp. 1352–1359.
- [14] J. McCormac, R. Clark, M. Bloesch, A. Davison, and S. Leutenegger, "Fusion++: Volumetric object-level slam," in 2018 international conference on 3D vision (3DV). IEEE, 2018, pp. 32–41.
- [15] B. Chen, F. Xia, B. Ichter, K. Rao, K. Gopalakrishnan, M. S. Ryoo, A. Stone, and D. Kappler, "Open-vocabulary queryable scene representations for real world planning," in 2023 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2023, pp. 11509–11522.
- [16] S.-H. Zhang, R. Li, X. Dong, P. Rosin, Z. Cai, X. Han, D. Yang, H. Huang, and S.-M. Hu, "Pose2seg: Detection free human instance segmentation," in *Proceedings of the IEEE/CVF conference on computer* vision and pattern recognition, 2019, pp. 889–898.
- [17] Y. Lee and J. Park, "Centermask: Real-time anchor-free instance segmentation," in *Proceedings of the IEEE/CVF conference on computer* vision and pattern recognition, 2020, pp. 13 906–13 915.
- [18] Y. Long, X. Li, W. Cai, and H. Dong, "Discuss before moving: Visual language navigation via multi-expert discussions," *arXiv preprint* arXiv:2309.11382, 2023.
- [19] Z. Jia, K. Yu, J. Ru, S. Yang, and S. Coleman, "Vital information matching in vision-and-language navigation," *Frontiers in Neurorobotics*, vol. 16, p. 1035921, 2022.
- [20] A. B. Vasudevan, D. Dai, and L. Van Gool, "Talk2nav: Long-range vision-and-language navigation with dual attention and spatial memory," *International Journal of Computer Vision*, vol. 129, pp. 246–266, 2021.

- [21] P. Chen, D. Ji, K. Lin, R. Zeng, T. Li, M. Tan, and C. Gan, "Weaklysupervised multi-granularity map learning for vision-and-language navigation," *Advances in Neural Information Processing Systems*, vol. 35, pp. 38149–38161, 2022.
- [22] J. Krantz, S. Banerjee, W. Zhu, J. Corso, P. Anderson, S. Lee, and J. Thomason, "Iterative vision-and-language navigation," in *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023, pp. 14921–14930.
- [23] G. Zhou, Y. Hong, and Q. Wu, "Navgpt: Explicit reasoning in visionand-language navigation with large language models," arXiv preprint arXiv:2305.16986, 2023.
- [24] R. Schumann, W. Zhu, W. Feng, T.-J. Fu, S. Riezler, and W. Y. Wang, "Velma: Verbalization embodiment of llm agents for vision and language navigation in street view," *arXiv preprint arXiv:2307.06082*, 2023.
- [25] S. Vemprala, R. Bonatti, A. Bucker, and A. Kapoor, "Chatgpt for robotics: Design principles and model abilities. 2023," *Published by Microsoft*, 2023.
- [26] D. Driess, F. Xia, M. S. Sajjadi, C. Lynch, A. Chowdhery, B. Ichter, A. Wahid, J. Tompson, Q. Vuong, T. Yu *et al.*, "Palm-e: An embodied multimodal language model," *arXiv preprint arXiv:2303.03378*, 2023.
- [27] B. Li, K. Weinberger, S. Belongie, V. Koltun, and R. Ranftl, "Languagedriven semantic segmentation," 2023.
- [28] P. P. Ray, "Chatgpt: A comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope," *Internet* of Things and Cyber-Physical Systems, 2023.
- [29] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale *et al.*, "Llama 2: Open foundation and fine-tuned chat models," *arXiv preprint arXiv:2307.09288*, 2023.
- [30] S. Y. Gadre, M. Wortsman, G. Ilharco, L. Schmidt, and S. Song, "Clip on wheels: Zero-shot object navigation as object localization and exploration," arXiv preprint arXiv:2203.10421, vol. 3, no. 4, p. 7, 2022.
- [31] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman *et al.*, "Do as i can, not as i say: Grounding language in robotic affordances," *arXiv preprint arXiv:2204.01691*, 2022.
- [32] A. Zeng, M. Attarian, B. Ichter, K. Choromanski, A. Wong, S. Welker, F. Tombari, A. Purohit, M. Ryoo, V. Sindhwani *et al.*, "Socratic models: Composing zero-shot multimodal reasoning with language," *arXiv preprint arXiv:2204.00598*, 2022.
- [33] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. Florence, and A. Zeng, "Code as policies: Language model programs for embodied control," in 2023 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2023, pp. 9493–9500.
- [34] E. Frantar, S. Ashkboos, T. Hoefler, and D. Alistarh, "Gptq: Accurate post-training quantization for generative pre-trained transformers," arXiv preprint arXiv:2210.17323, 2022.
- [35] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik *et al.*, "Habitat: A platform for embodied ai research," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9339–9347.
- [36] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 618–626.
- [37] R. Schumann and S. Riezler, "Analyzing generalization of vision and language navigation to unseen outdoor areas," *arXiv preprint* arXiv:2203.13838, 2022.
- [38] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, "Scannet: Richly-annotated 3d reconstructions of indoor scenes," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 5828–5839.
- [39] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [40] C. Cao, H. Zhu, F. Yang, Y. Xia, H. Choset, J. Oh, and J. Zhang, "Autonomous exploration development environment and the planning algorithms," in 2022 International Conference on Robotics and Automation (ICRA). IEEE, 2022, pp. 8921–8928.

high-level functions	function description
get_nearest_obj_pos(object_name)	get the map position of the nearest front object.
get_obj_attributes(object_name, instance_idx, objext_color)	get the object attributes in Tuple( object_name,label_id,color)from IVLMap
get_specifed_obj_pos(object_with_attributes)	get the map position of the object with specifed attributes.
get_nearest_obj_contour(object_name)	get the contour turning points of the nearest front object on the map.
move_to(pos)	move to a position on the map.
move_to_left(objedt_with_attributes)	move to the left side of the object with specified attributes.
move_to_right(object_with_attributes)	move to the right side of the object with specified attributes.
with object on left(object_with_attributes)	turn until the specific object is on the robot's left side.
with object on right(object_with_attributes)	turn until the specific object is on the robot's right side.
move_in_between(object_a, object_b)	move in between two objects.
turn(angle)	turn right a certain angle. If the angle value is negative, turn left.
face(object_with_attributes)	turn until the object is with specific attributes in front of the robot.
turn_absolute(angle)	turn to absolute angle. 0 isnorth, 90 is east, -90 is west,180 is south.
move_north(object_with_attributes)	move to the north side of the object with specific attributes.
move_south(object_with_attributes)	move to the south side of the object with specific attributes.
move_east(object_with_attributes)	move to the east side of the object with specific attributes.
move_west(object_with_attributes)	move to the west side of the object with specific attributes.
move_to_object(object_with_attributes)	move to the object with specific attributes.
move_forward(dist)	move forward dist' meters.

TABLE III

LIST OF HIGH-LEVEL FUNCTIONS USED

### VI. APPENDIX

#### Appendix

#### A. Full List of Navigation High-Level Functions

On the basis of VLMap, we have developed a new high-level function library based on the distinctive features of IVLMap. Our full list of navigation high-level functions are listed in Table. III

# B. Full Prompts for LLM Generating Python Code system\_prompt

You are an assistant helping me with the mobile robot agent. When I ask you to do something, you are supposed to give me Python code that is needed to achieve that task with explanations of what that code does. You are only allowed to use the functions I have defined for you. You are not to use any other hypothetical functions that you think might exist. You can use simple Python functions from libraries such as math and numpy.

# attributes\_prompt

You should extract the object attributes from uer's command, here are some examples:

- I : navigate to the third yellow table.
- A: [(table,third,yellow)]

- A: [(chair,None,yellow),(sofa,None,black)]
- I: go to the kitchen and then go to the toilet.A: [(kitchen,None,None), (toilet,None,None)]
- Move to the west of the black chair, with the first I : red sofa on your right, move to the 2nd table, then
- turn right 90 degree, then find a table. (chair, None, black), (sofa, first, red), (table, sen
  - cond, None), (table, None, None)]

# function\_prompt

Here are some functions you can use to command the mobile robot agent.

# get the map position of the nearest front object.

 $agent.get\_nearest\_obj\_pos(object\_name)$ 

# get the object attributes in Tuple(ob-ject name,label id,color)from IVLMap

agent.get\_obj\_attributes(object\_name,instance\_idx, objext\_color)

# move forward dist' meters. If the angle value is negative, move backward.

 $agent.move\_forward(dist)$ 

The omitted function prompt can be referred to in Table III.

# example\_prompt

Here are some examples.

- Me: navigate to the third yellow table
- You: obj\_attr = agent.get\_obj\_attributes('table',3,'yellow') agent.face(obj\_attr)
- Me: Move to the red sofa then stop at the right side of it
- You: obj\_attr = agent.get\_obj\_attributes('sofa',0,'red') agent.move\_to\_right(obj\_attr)
- Me: Go straight for 1 meter and then walk clockwise around the red sofa for 2 turns.
- You: agent.move\_forward(1)

obj\_attr = agent.get\_obj\_attributes( 'sofa', 0, 'red'
)

obj\_contour = agent.get\_specifed\_obj\_pos(obj\_attr)
agent.face(obj\_attr), agent.turn(-90)

for i in range(8):

agent.move\_forward(obj\_contour([i%4])
agent.turn(90)

here are some examples for you to generate code in

python according to Human's command. You should

return your answer in Markdown format, especially

base coordinate system and the Camera coordinate system is necessary(Fig.9). The rotational relationship between the robot coordinate system and the camera coordinate system can be expressed using the rotation matrix shown in Equation.3.  $rot = \left[ \begin{array}{rrr} 0 & 0 & 1 \\ -1 & 0 & 0 \\ 0 & -1 & 0 \end{array} \right]$ 



# C. Experiment on Real Robot

low')

**ROS-based Smart Car Real-World Data Collection Scheme:** In current research on visual language navigation, the majority of work is implemented using simulators such as Habitat and AI2THOR, achieving notable results in virtual environments. To validate the effectiveness of the algorithm in real-world scenarios, we conducted corresponding experiments in actual environments. Our real-world data collection platform is illustrated in Fig.8. Before initiating the data collection process, we performed camera calibration to establish the transformation relationship between the robot base coordinate system and the camera coordinate system. During the data collection process, it is crucial to ensure that the robot and the host are on the same local network. The robot's movement is controlled using the laptop keyboard through the ROS communication mechanism. RGB and depth information is captured through the Astro Pro Plus camera, while the pose information is obtained from IMU and the robot's velocity encoder. These sensor data are published as ROS topics. Subsequently, ROS message\_filter<sup>5</sup> is utilized to synchronize the three different types of sensor data, ensuring they are roughly at the same timestamp.



Fig. 8. We built an intelligent robotic car with ROS, leveraging Jetson Nano for deep learning. It includes a high-precision Orbbec Astra Pro Plus monocular depth camera and Slamtec E300 LiDAR for RGB-D and LiDAR data capture.

In the real-world environment, where only odometer information reflecting the pose changes of the robot car is available, a coordinate transformation between the ROSRobot



Fig. 9. Habitat coordinate system, camera coordinate system, and ROSRobot coordinate system



Fig. 10. 3D Reconstruction Results in Real Environment, examples of RGB images captured by the camera on the left, the corresponding 3D reconstruction results are displayed on the right.

#### D. 3D Reconstruction in Bird's-Eye View

3D Reconstruction Bird's-Eye View of Different Scenes in the Dataset Collected with cmu-exploration can be refered at Fig. 11.

## E. IVLMap segmentation results

To compare the semantic maps between IVLMap (ours, depicted in orange) and VLMap (ground truth, represented in green), refer to the visual representation of the maps for a comprehensive understanding of the differences and similarities Fig.12.

(3)

<sup>&</sup>lt;sup>5</sup>ROS message\_filter is a library in the Robot Operating System (ROS) that provides tools for synchronizing and filtering messages from multiple topics in a flexible and efficient manner.





5ZKStnWn8Zo



759xd9YjKW5



0



5q7pvUzZiYa

Fig. 11. 3D Reconstruction Bird's-Eye View of Different Scenes



appliances

clothes

curtain plant

sink stairs

5ZKStnWn8Zo

Fig. 12. Semantic Map Comparison between IVLMap (Ours, Orange) and VLMap (GT, Green)