

Interactive Navigation in Environments with Traversable Obstacles Using Large Language and Vision-Language Models

Zhen Zhang^{1,2}, Anran Lin³, Chun Wai Wong¹, Xiangyu Chu^{1,2,4}, Qi Dou³, and K. W. Samuel Au^{1,2}

Abstract—This paper proposes an interactive navigation framework by using large language and vision-language models, allowing robots to navigate in environments with traversable obstacles. We utilize the large language model (GPT-3.5) and the open-set Vision-language Model (Grounding DINO) to create an action-aware costmap to perform effective path planning without fine-tuning. With the large models, we can achieve an end-to-end system from textual instructions like “Can you pass through the curtains to deliver medicines to me?”, to bounding boxes (e.g., curtains) with action-aware attributes. They can be used to segment LiDAR point clouds into two parts: traversable and untraversable parts, and then an action-aware costmap is constructed for generating a feasible path. The pre-trained large models have great generalization ability and do not require additional annotated data for training, allowing fast deployment in the interactive navigation tasks. We choose to use multiple traversable objects such as curtains and grasses for verification by instructing the robot to traverse them. Besides, traversing curtains in a medical scenario was tested. All experimental results demonstrated the proposed framework’s effectiveness and adaptability to diverse environments.

I. INTRODUCTION

Robot navigation is critical in deploying robots in an obstacle-involved workspace. Besides the long-sought need for safety and efficiency [1]–[3], interaction with humans to satisfy real-time needs is becoming more appealing. For example, a patient can interact with a medicine-delivery dog and ask it to enter the ward area by passing through the surrounding curtains, as shown in Fig. 1. Without real-time interaction and extra treatments, it is challenging for the robotic dog to pass through the curtain because the curtain is normally considered as an obstacle in its motion planner based on its sensor data (e.g., LiDAR) and thus no feasible path can be found. Recently, large language models have provided an alternative solution to real-time human-robot interaction, but using them to interact with robots for navigation in various environments remains open. In this paper, we focus on large-model-based interactive navigation for robots, especially in environments with traversable obstacles like curtains and grasses for a robotic dog.

Interacting with robots for navigation has been studied via many interfaces such as gestures, teleoperation devices,

This work was supported in part by Chow Yuk Ho Technology Centre of Innovative Medicine, The Chinese University of Hong Kong, in part by the Multiscale Medical Robotics Centre, AIR@InnoHK, and in part by the Research Grants Council (RGC) of Hong Kong under Grant 14211320. ¹: Department of Mechanical and Automation Engineering, The Chinese University of Hong Kong, Hong Kong SAR. ²: Multiscale Medical Robotics Centre, Hong Kong SAR. ³: Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong SAR. ⁴: Computer Aided Medical Procedures, Technical University of Munich, Germany. (Corresponding authors: Xiangyu Chu and Qi Dou)




Fig. 1. An example of interactive navigation. When a medicine-delivery dog arrives at a room, a patient can further interact with the robotic dog and ask it to walk to the bed by passing through the traversable curtains. Such an interaction can help navigate the robot to the place where it cannot normally reach and therefore meet humans’ real-time needs.

and natural languages. [4]–[6] recognized simple human gestures by using onboard cameras to control robots’ behavior like their forward-backward direction. Besides, different teleoperation devices such as joysticks [7] and tablets [8] can be used for navigating robots to desired places. We find these two kinds of interfaces may not be available for patients in bed. A more natural interface is to use voices, i.e., talking to the robot. Early works on robot navigation using natural languages relied on natural language processing techniques to obtain semantic information such as landmarks [9]. Recently, the birth of Large Language Models (LLMs) greatly improved the capability of extracting semantic information [10]–[13]. Although LLMs have shown much better generalization, their output cannot be understood by navigation algorithms. For example, a motion planner cannot understand “curtains” extracted by LLMs. To bridge the gap, people started to combine language and vision-language large models. [14] used Vision-and-Language Models (VLMs) to associate images with texts from LLMs and then a visual navigation model is used to generate a feasible path based on the images from VLMs. Since the combination of a pre-trained language and a vision-language model can parse complex high-level instructions for robots, in this paper, we choose to use this paradigm to deal with human’s natural language in real-time interaction. In our case, we cannot use visual navigation methods to generate feasible paths because the landmarks in visual navigation methods (e.g., [14]–[17]) are normally assumed to be untraversable,

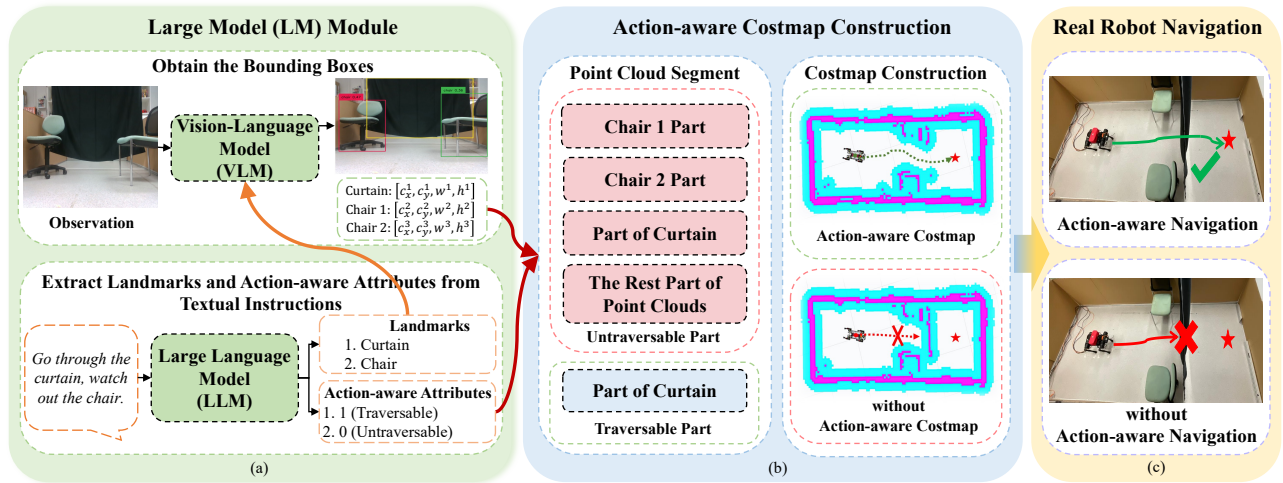


Fig. 2. Proposed large-model-based interactive navigation framework for the robots in environments with traversable obstacles. (a) The large model module uses LLM and VLM to extract the landmarks’ bounding boxes and action-aware attributes from texts/speeches. (b) The action-aware costmap is constructed with the segmented point clouds based on the output of the large model module. (c) A feasible path can be planned with an action-aware costmap, while no feasible path would be found if not considering the landmarks’ action-aware attribute.

thus no feasible path can be found when the goal (e.g., a bed) cannot be detected (see Fig. 1).

Constructing costmaps is an alternative solution to generate a feasible path. There are some works on connecting images with costmaps. For example, [9] made instance segmentation to extract semantic objects of interest in the current image. Then the obtained masks are fed to the depth image for objects’ localization, and finally, the costmap would be updated to include the objects like a moving human. [18] also built different obstacle maps for different embodiments given images. For example, a table is an obstacle for a wheeled robot but not for a drone. Compared to visual navigation that relies on goal images, constructing costmaps is more fundamental, and shifts the feasible path generation to the navigation planner. However, constructing costmaps based on the landmark images is not sufficient in our task because the landmarks are not traversable in default and thus the robot would get stuck in front of the landmarks. Unlike [14] where only landmarks were extracted and verbs/commands were disregarded, we propose to extract action-aware attributes corresponding to the landmarks and then use them for costmap construction. For example, “pass through the curtains” means the curtains are traversable due to “pass through”. Such an action-aware property integrated into building the costmaps allows the robot to pass through the curtains.

In this paper, we propose an action-aware interactive navigation system based on pre-trained large models (i.e., LLM and VLM). LLM is used to extract landmarks and corresponding action-aware attributes in texts, and VLM is used to detect the landmarks in the picture and obtain the corresponding bounding boxes. These large models help to make the transformation from speeches and texts to bounding boxes and their action-aware attributes without any fine-tuning and additional training data. [19] used large models to classify different navigation subtasks and required multiple

action-specific navigators. Different from [19], we use the action-aware attributes of landmarks obtained from large models to fundamentally build an action-aware costmap for generating a feasible path and thus achieve more intelligent and flexible interactive navigation that meets humans’ real-time needs. We summarize the contributions of this paper:

- 1) We propose a large-model-based interactive navigation framework for robots to plan a feasible path in environments with traversable objects.
- 2) Besides landmarks, action-aware attributes in textual instructions are extracted to assist sensor data segmentation. Such an action-aware attribute allows us to construct an action-aware costmap for interactive navigation.
- 3) Extensive experiments demonstrate the effectiveness of the proposed framework and its generalization on different traversable objects and scenarios.

II. PROPOSED FRAMEWORK

A. Problem Formulation

Given a set of landmarks and their corresponding action-aware attributes $\{(\ell_1, P_{a_1|\ell_1}), (\ell_2, P_{a_2|\ell_2}), \dots, (\ell_n, P_{a_n|\ell_n})\}$ extracted by LLM from language-based interaction. $P_{a_i|\ell_i} = 1$ denotes that the landmark ℓ_i associated with the action a_i is traversable and $P_{a_i|\ell_i} = 0$ denotes that the landmark ℓ_i associated with the action a_i is untraversable. With a goal in the static map, our objective is to find a feasible path including a series of waypoints $\{w_1, w_2, \dots, w_m\}$. In this path, all untraversable landmarks’ action-aware attributes are respected (i.e., maintaining a safe distance from these landmarks) and traversable landmarks are fully utilized to achieve the goal.

Note that finding a feasible path is not a trivial problem when we consider landmarks’ action-aware attributes. For example, in Fig. 1, the bed cannot be detected if using visual sensors before traversing the curtains and entering the ward area, thus no feasible path can be found. Without human

interaction, the curtains are still regarded as obstacles and it takes more effort to get close to the bed.

B. Framework Overview

To solve the problem, we propose a large-model-based interactive navigation framework that can map language commands to a feasible path, as shown in Fig. 2. This framework consists of three parts: a large model module, action-aware costmap construction, and a navigation planner.

The large model module includes a LLM and a VLM. LLMs can learn the grammar and semantics of natural language by learning massive corpus data. In this work, the pre-trained LLM does not require additional data for training and fine-tuning. It can be used to extract landmarks ℓ_i and their corresponding action-aware attributes $P_{a_i|\ell_i}$ in the text. With VLMs, we can detect objects in images using language [20]–[23]. The key is to associate the semantic information of the language with the image feature information of the object. The output of the VLM is the information on the bounding boxes of the landmarks that occur in the language instructions.

The costmap construction part utilizes the output of the large model module, including the landmarks’ bounding boxes and action-aware attributes, to build up an action-aware costmap. The costmap is a container that can store information about the environment and can be further used for path planning [24], [25]. Moreover, the layered costmap [26] can include multiple costmap layers according to different sensor data and semantic information, which can solve the path planning problem in complex scenarios [27]–[29]. In the proposed framework, we construct such a flexible layered costmap with semantic information related to the landmarks in the form of bounding boxes and their action-aware attributes obtained from the large model module.

With an action-aware costmap, many off-the-shelf path planning methods can be employed for the part of robot navigation. In this paper, we choose to use the A star (A^*) search algorithm [30]. A feasible path including a series of waypoints $\{w_1, w_2, \dots, w_m\}$ can be obtained, which is not attainable without action-aware costmaps (see Fig. 2 (b)).

III. LARGE MODEL DEPLOYMENT

In this work, we use GPT-3.5 as the LLM to extract the landmarks and corresponding action-aware attributes in textual instructions which can be transferred from speeches, and then use the VLM (Grounding-DINO [23]) to locate the bounding boxes corresponding to the landmarks, including a center point coordinate (c_x, c_y) , width w , and height h in pixel for each bounding box, denoted as $B = \{c_x, c_y, w, h\}$.

A. Large Language Model

We deploy GPT-3.5 by accessing OpenAI’s API¹. To perform interactive navigation, we introduce an **action-aware** attribute corresponding to a landmark. We define that objects in the real-world environment can have two kinds of action-aware attributes. One is “traversable”, which means a robot

can traverse the landmark. Another is “untraversable”, and a robot cannot traverse the landmark as much previous work assumed.

To describe the action-aware attribute of a landmark ℓ , we define:

$$P_{a|\ell} = \begin{cases} 1 & \rightarrow \ell \text{ is traversable under } a \\ 0 & \rightarrow \ell \text{ is untraversable under } a \end{cases} \quad (1)$$

where a is an action (e.g., “go through” and “traverse”) related to the landmark.

With GPT-3.5, we can extract landmarks and action-aware attributes from free-form instructions. The large model is robust to different descriptions in the input prompts and is also greatly effective in parsing instructions. In the following, we make prompt engineering and then show a test example.

Example 1: I ask: “Please go through the curtain and watch out for the medicine trolley”. You should respond “curtain” and “1”; “trolley” and “0”.

Example 2: I ask: “Please pass through the curtain but be careful of the table in the middle of the room”. You should respond “curtain” and “1”; “table” and “0”.

Here is an example. The test textual instruction is: “Go through the curtain, and watch out the chair.”. The test output is: *curtain and 1; chair and 0*.

B. Vision-Language Model

The landmarks obtained from the LLM output are fed to the VLM in the form of a simple prompt “*curtain, chair*”. We utilize Grounding DINO [23] to obtain the bounding box B for each landmark. Specifically, for each (*Image, Text*) pair, image and text features are extracted by an image backbone and a text backbone respectively, and then fed into a feature enhancer module to fuse cross-modality features. After that, cross-modality queries are selected using a language-guided query selection module and fed into a cross-modality decoder. The output queries of the last decoder layer are used to predict object boxes $B = \{c_x, c_y, w, h\}$ and extract corresponding phrases (e.g., *curtain*). Therefore, the final



Fig. 3. Examples of output of the large model module. The initial instruction for these examples is: “Please go through the curtain and be careful of the chair”. We obtain the landmarks and the corresponding traversable attributes, *curtain and 1* (in orange); *chair and 0* (in red), using GPT-3.5. Then, we can ground the related bounding boxes with [23].

outputs of the large model module for each landmark form the input, $\{c_x, c_y, w, h, P_{a|\ell}\}$, for the following action-aware costmap construction. Figure 3 shows some examples. It should be noted that in the example in the lower right corner, due to the limitation of the VLM, a certain part of the object

¹<https://openai.com/api/>

may be output as a bounding box (see the two bounding boxes of the curtain). With the help of the method in Sec. IV-B, we can also handle this situation well.

IV. ACTION-AWARE COSTMAP CONSTRUCTION

With LLM and VLM, we obtain the 2D bounding boxes $\{c_x, c_y, w, h\}$ of the landmarks in the picture and the action-aware attribute $P_{a|\ell}$ corresponding to each landmark ℓ . To find a feasible path, we propose an action-aware costmap construction method. In our system, we use a LiDAR mapping sensor. First, We transform the LiDAR point clouds to the image plane. Then, according to the bounding boxes and their corresponding action-aware attributes obtained by the large model module, the LiDAR point clouds are segmented into two parts: the traversable part and the untraversable part. Lastly, we use the two parts of LiDAR point clouds as input sensor sources to build an action-aware costmap and update the map in real-time. The costmap construction pipeline accompanied by the processing based on large models can be found in Algorithm 1.

Algorithm 1: Action-aware Costmap Construction with Large Language and Vision-Language Models

Input : Textual instructions T_t or Voice V_t , camera image I_t , current LiDAR point clouds PC_t . t is the current step.

Output: An action-aware costmap M_t

```

1  elif  $V_t$  is not None  $T_t = \text{Speech\_recognition}(V_t)$ ;
2   $T_t = T_t$ ;
3   $\{(\ell_1, P_{a_1|\ell_1}), (\ell_2, P_{a_2|\ell_2}), \dots, (\ell_n, P_{a_n|\ell_n})\} = \text{LLM}(T_t)$ ;
4  while  $\| \text{current\_position} - \text{target} \|_2 > \text{Threshold}$  do
5       $\{B_1, B_2, \dots, B_m\} = \text{VLM}(I_t, (\ell_1, \ell_2, \dots, \ell_n))$ , where
6           $B_i = \{c_x^i, c_y^i, w^i, h^i\}$ ;
7           $P_{A_i} = P_{a_i|\ell_i}$ , where
8           $A_i = \{c_x^i - \frac{w}{2} \leq x \leq c_x^i + \frac{w}{2}, c_y^i - \frac{h}{2} \leq y \leq c_y^i + \frac{h}{2}\}$ ;
9      if  $\{A_1, A_2, \dots, A_m\}$  is not empty then
10          $[PC_t^{Tra}, PC_t^{Untra}] = \text{Segment}(PC_t, \{(A_1, P_{A_1}), \dots, (A_m, P_{A_m})\})$ ;
11          $M_t = \text{UpdateCostmap}(M_{t-1}, (PC_t^{Tra}, PC_t^{Untra}))$ ;
12     else
13          $M_t = \text{UpdateCostmap}(M_{t-1}, PC_t)$ ;
14     end
15 end

```

A. Transformation from LiDAR to Camera

To segment the LiDAR point clouds with the obtained bounding boxes, the LiDAR point clouds need to be transformed into the image plane. We can use the extrinsic matrix and intrinsic matrix to project a 3D point cloud onto the 2D image plane of the camera. Given a LiDAR point $p = (x_i, y_i, z_i)$ in the Cartesian coordinate system with z -axis pointing upward, it can be transformed into the image plane and obtain a coordinate $s = (u, v)$ [31]:

$$\begin{bmatrix} u & v & 1 \end{bmatrix}^T = \begin{bmatrix} \frac{u'}{w'} & \frac{v'}{w'} & 1 \end{bmatrix}^T, \quad (2)$$

$$\begin{bmatrix} u' \\ v' \\ w' \end{bmatrix} = \underbrace{\begin{bmatrix} f s_x & k s_y & u_0 \\ 0 & f s_y & v_0 \\ 0 & 0 & 1 \end{bmatrix}}_{\text{intrinsic}} \begin{bmatrix} \mathbf{I}_{3 \times 3} \\ \mathbf{0}_{1 \times 3} \end{bmatrix}^T \underbrace{\begin{bmatrix} R_L^C & t_L^C \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix}}_{\text{extrinsic}} \begin{bmatrix} x_i \\ y_i \\ z_i \\ 1 \end{bmatrix},$$

where f is the focal length, s_x and s_y are the scale in x and y respectively, k is the rotation of shear in the y direction to than in x , R_L^C and t_L^C are the rotation matrix and translation vector from LiDAR frame to camera frame, (u_0, v_0) is the distance from the image center to the image plane coordinate system origin, and w is the distance from camera centre to the LiDAR point along z -axis. For simplification, we can define such a transformation as $s = \mathbb{T}(p)$.

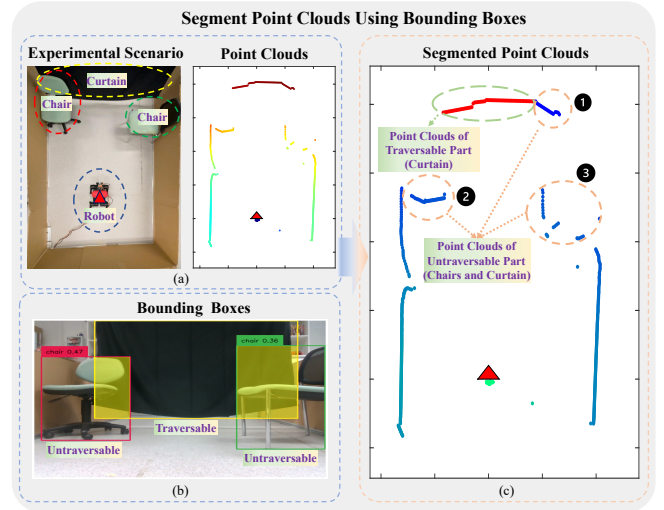


Fig. 4. (a) The experimental scenario in bird's eye view and the point clouds obtained from LiDAR Mapping sensor. (b) The yellow areas are the intersection between the bounding boxes of untraversable and traversable objects. (c) The point clouds colored red are the traversable part belonging to the curtain. The rest is the untraversable part belonging to part of the curtain (1), chairs (2, 3), and the environment.

B. LiDAR Point Clouds Segmentation

After projecting the LiDAR point clouds PC_t into the image plane I_t , we can determine the action-aware attribute for each LiDAR point by checking if $s = \mathbb{T}(p) \in \mathcal{A}$, where $\mathcal{A} = \text{Area}\{c_x - \frac{w}{2} \leq u \leq c_x + \frac{w}{2}, c_y - \frac{h}{2} \leq v \leq c_y + \frac{h}{2}\}$. Specifically, we create the masks with the same size as the bounding boxes and then use them to segment the LiDAR point clouds belonging to different landmarks on the image plane. In practice, there are multiple overlapping bounding boxes, so one LiDAR point may belong to multiple bounding boxes. Based on the action-aware attributes of the landmarks $P_{a|\ell}$, the action-aware attribute for each LiDAR point in the bounding box is known, i.e., $P_A = P_{a|\ell}$. Thus, a function Segment is defined as

$$p \in \begin{cases} PC_t^{Tra}, & P_{A_i} = 1 \text{ if } s \in \mathcal{A}_i, \forall i, \forall p \in PC_t, \\ PC_t^{Untra}, & \text{else} \end{cases} \quad (3)$$

where $i = 1, \dots, m$ and m is the number of bounding boxes. A LiDAR point belonging to the traversable part PC_t^{Tra} requires that all action-aware attributes of its occupied bounding boxes should be 1. When the bounding box with the traversable attribute overlaps with the bounding box with the untraversable attribute (see yellow intersection areas in Fig. 4 (b)), the overlapping part needs to be considered as

the untraversable part PC_t^{Untra} for safety. From Fig. 4 (c), it is observed that LiDAR point clouds are segmented into two parts (blue for untraversable and red for traversable) using bounding boxes with action-aware attributes. Especially, the circle 1 (❶) in Fig. 4 (c) is untraversable, although it is part of the traversable curtain.

C. Costmap Building

To build the costmap, we use the two parts of the LiDAR point clouds: PC_t^{Tra} and PC_t^{Untra} . For the LiDAR points in PC_t^{Tra} , the cost on the costmap is 0, which is considered to be *FREESPACE*, while for the LiDAR points in PC_t^{Untra} , its cost on the costmap is 254, which is considered to be *LETHALSPACE*. With an action-aware costmap, a feasible path can be found since the cost of the LiDAR points related to the curtain has been set to 0. Otherwise, no feasible path can be obtained, see the costmap construction in Fig. 2 (b).

V. EVALUATION

We tested the performance of our interactive navigation framework on the Unitree A1 quadruped robot platform equipped with a 2D LiDAR mapping sensor (*Slamtec Mapper M2M2*) and a depth camera (*Intel RealSense D435i*). The legged robot control used an open-source package [32]. We ran the interactive navigation framework on a PC equipped with an Intel Core i7-13700K CPU. We demonstrated the results of the entire interactive navigation system in both simulated and real-world scenarios. Especially, in real-world scenarios, we conducted extensive experiments on curtains of different materials, types, and shapes to test the reliability and adaptability of the proposed framework. We also evaluated the performance of other objects with traversable properties such as grasses.

A. Framework Evaluation in a Simulated Environment




Fig. 5. (a) We test our interactive navigation system in a simulated in-door hospital ward environment with static obstacles. (b-c) The simulated legged robot can obtain the RGB image and point clouds, respectively.

As shown in Fig. 5, we used an indoor hospital scene in the Gazebo simulator [33] to verify the feasibility of the framework. It contains static obstacles in the room such as “elder lady patient”, “medical instrument”, “walker”, “orange wooden wall”, “freezer condenser”, “cabinet”, and “white table”. To be noticed, the “orange wooden wall” separates the hospital ward into two spaces.

In the simulation, we assumed that the action-aware attribute of the “orange wooden wall” is traversable and our textual instruction can be “Go through the orange wooden wall”. Given a goal behind the wall, no feasible path can be planned without the action-aware costmap, as shown in Fig. 6

(a). However, we found part of the “orange wooden wall” can be cleared in the costmap if constructing an action-aware costmap and a straightforward path that crosses “orange wooden wall” can be obtained, as shown in Fig. 6 (b), demonstrating the effectiveness of the proposed framework.




Fig. 6. (a) Without using the action-aware costmap, the orange wooden wall would be regarded as an untraversable obstacle and marked in the costmap. Thus, no feasible path is planned. (b) A feasible path can be planned by using the action-aware costmap because part of the orange wooden wall is regarded as a traversable obstacle and cleared in the costmap.

B. Framework Evaluation in Real-world Environments

We conducted experiments in two real-world scenarios: a laboratory scenario and a medical scenario at the Multi-Scale Medical Robotics Center (MRC: <https://www.mrc-cuhk.com>). We set up a simulation-like scenario in both and ran the SLAM algorithm [34] to construct a global map and annotate the goal (e.g., a bed in the scenario of MRC) in the map. For the traversable objects, We prepared 4 curtains with different colors and textures, which are common in our daily life, and we also prepared two different types of artificial grass.

1) *Go through the curtain*: First, we conducted experiments with 4 different types, materials, and colors of curtains in the laboratory scenario. We used “Go through the curtain and watch out the chair.” as the textual instruction. Given an arbitrary target behind the curtain, no feasible path could be planned without using our framework because the curtain separated the experimental scenario into two parts. With our framework, an action-aware costmap based on the textual instruction could be built and there would be a feasible path. The robot could pass through the curtain and reach the given target location following the path, as shown in Fig. 7 (a-d).

Then we conducted the experiment of going through the curtain in a medical scenario at MRC. It contains a “hospital bed”, “medical curtain”, a “warning sign”, and a “medical trolley”. We used “Deliver the medicine to me. By the way, you can go through the curtain and watch out the warning sign.” as the textual instruction. As shown in Fig. 8, the robot avoided the warning sign, then passed through the curtain successfully, and finally reached the desired location.

2) *Go through grass*: We prepared two types of artificial grass: *setaria viridis* and *phragmites australis*. We mixed two kinds of grass and used “Go through the grass and watch out the chair.” as the textual instruction. Similar to the curtain experiments, a feasible path could be planned and the robot went through the grass using our framework (Fig. 7 (e)).



Fig. 7. (a-d) Experiments of going through different curtains. (e) Experiment of going through the grass.



Fig. 8. Experiment of the medical scenario at MRC. The upper right corner of each picture shows the first-person point of view of the robot.

VI. CONCLUSION AND FUTURE WORK

In the paper, we have presented a large-model-based interactive navigation framework, mapping textual instructions to feasible paths for navigating robots in the environment with traversable objects. To obtain a feasible path, we propose to construct an action-aware costmap by segmenting the LiDAR point clouds into traversable and untraversable parts, according to the bounding boxes of the landmarks and action-aware attributes from the large model module (LLM+VLM). Extensive experiments have shown the effectiveness and generalization of the proposed framework in different environments including a laboratory scenario and a medical

scenario at MRC.

In the future, we plan to use a 3D LiDAR to replace the current 2D one to restore the 3D information of the environment for better sensing. To achieve more precise and robust action-aware segmentation of sensor data, we will deploy a state-of-the-art segmentation model (e.g., Segment Anything Model [35]) to our LM module to produce high-quality object masks for improving the capability of segmenting sensor data (e.g., LiDAR point clouds) to the background and objects. Besides, we will upgrade the legged robot to a legged manipulator (e.g., [36]), allowing us to manipulate traversable objects and thus enhancing interactive navigation.

REFERENCES

- [1] Lorenz Wellhausen, René Ranftl, and Marco Hutter. Safe robot navigation via multi-modal anomaly detection. *IEEE Robotics and Automation Letters*, 5(2):1326–1333, 2020.
- [2] Veronica Vulcano, Spyridon G Tarantos, Paolo Ferrari, and Giuseppe Oriolo. Safe robot navigation in a crowd combining nmpc and control barrier functions. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 3321–3328. IEEE, 2022.
- [3] Tianrui Guan, Divya Kothandaraman, Rohan Chandra, Adarsh Jagan Sathyamoorthy, Kasun Weerakoon, and Dinesh Manocha. Ga-nav: Efficient terrain segmentation for robot navigation in unstructured outdoor environments. *IEEE Robotics and Automation Letters*, 7(3):8138–8145, 2022.
- [4] Xuan Liu, Kashif Nazar Khan, Qamar Farooq, Yunhong Hao, and Muhammad Shoaib Arshad. Obstacle avoidance through gesture recognition: Business advancement potential in robot navigation socio-technology. *Robotica*, 37(10):1663–1676, 2019.
- [5] Tayyab Naseer, Jürgen Sturm, and Daniel Cremers. Followme: Person following and gesture recognition with a quadcopter. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 624–630. IEEE, 2013.
- [6] Muaammar Hadi Kuzman Ali, M Asyraf Azman, Zool Hilmi Ismail, et al. Real-time hand gestures system for mobile robots control. *Procedia Engineering*, 41:798–804, 2012.
- [7] Patrick Stokto, Stefan Krumpen, Max Schwarz, Christian Lenz, Sven Behnke, Reinhard Klein, and Michael Weinmann. A vr system for immersive teleoperation and live exploration with a mobile robot. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3630–3637. IEEE, 2019.
- [8] Samuel A Olatunji, Andre Potenza, Andrey Kiselev, Tal Oron-Gilad, Amy Loutfi, and Yael Edan. Levels of automation for a mobile robot teleoperated by a caregiver. *ACM Transactions on Human-Robot Interaction (THRI)*, 11(2):1–21, 2022.
- [9] Zhe Hu, Jia Pan, Tingxiang Fan, Ruigang Yang, and Dinesh Manocha. Safe navigation with human instructions in complex scenes. *IEEE Robotics and Automation Letters*, 4(2):753–760, 2019.
- [10] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [11] Ross Taylor, Marcin Kardas, Guillem Cucurull, Thomas Scialom, Anthony Hartshorn, Elvis Saravia, Andrew Poulton, Viktor Kerkez, and Robert Stojnic. Galactica: A large language model for science. *arXiv preprint arXiv:2211.09085*, 2022.
- [12] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- [13] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022.
- [14] Dhruv Shah, Błażej Osiniński, Sergey Levine, et al. Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action. In *Conference on Robot Learning*, pages 492–504. PMLR, 2023.
- [15] Dhruv Shah, Ajay Sridhar, Nitish Dashora, Kyle Stachowicz, Kevin Black, Noriaki Hirose, and Sergey Levine. Vint: A foundation model for visual navigation. *arXiv preprint arXiv:2306.14846*, 2023.
- [16] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton 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*, pages 3674–3683, 2018.
- [17] Mateus V Gasparino, Arun N Sivakumar, Yixiao Liu, Andres EB Velasquez, Vitor AH Higuti, John Rogers, Huy Tran, and Girish Chowdhary. Wayfast: Navigation with predictive traversability in the field. *IEEE Robotics and Automation Letters*, 7(4):10651–10658, 2022.
- [18] Chenguang Huang, Oier Mees, Andy Zeng, and Wolfram Burgard. Visual language maps for robot navigation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10608–10615. IEEE, 2023.
- [19] Peihao Chen, Xinyu Sun, Hongyan Zhi, Runhao Zeng, Thomas H. Li, Gaowen Liu, Mingkui Tan, and Chuang Gan. a^2 nav: Action-aware zero-shot robot navigation by exploiting vision-and-language ability of foundation models, 2023.
- [20] Ting Chen, Saurabh Saxena, Lala Li, David J Fleet, and Geoffrey Hinton. Pix2seq: A language modeling framework for object detection. *arXiv preprint arXiv:2109.10852*, 2021.
- [21] Yu Du, Fangyun Wei, Zihe Zhang, Miaojing Shi, Yue Gao, and Guoqi Li. Learning to prompt for open-vocabulary object detection with vision-language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14084–14093, 2022.
- [22] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10965–10975, 2022.
- [23] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023.
- [24] Léonard Jaillet, Juan Cortés, and Thierry Siméon. Sampling-based path planning on configuration-space costmaps. *IEEE Transactions on Robotics*, 26(4):635–646, 2010.
- [25] Jim Mainprice, E. Akin Sisbot, Léonard Jaillet, Juan Cortés, Rachid Alami, and Thierry Siméon. Planning human-aware motions using a sampling-based costmap planner. In *2011 IEEE International Conference on Robotics and Automation*, pages 5012–5017, 2011.
- [26] David V Lu, Dave Hershberger, and William D Smart. Layered costmaps for context-sensitive navigation. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 709–715. IEEE, 2014.
- [27] Marina Kollmitz, Kaijen Hsiao, Johannes Gaa, and Wolfram Burgard. Time dependent planning on a layered social cost map for human-aware robot navigation. In *2015 European Conference on Mobile Robots (ECMR)*, pages 1–6, 2015.
- [28] Yoichi Morales, Atsushi Watanabe, Florent Ferreri, Jani Even, Tetsushi Ikeda, Kazuhiro Shinozawa, Takahiro Miyashita, and Norihiro Hagita. Including human factors for planning comfortable paths. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6153–6159, 2015.
- [29] Omar A. Islas Ramírez, Harmish Khambhaita, Raja Chatila, Mohamed Chetouani, and Rachid Alami. Robots learning how and where to approach people. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 347–353, 2016.
- [30] Peter E Hart, Nils J Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.
- [31] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2 edition, 2004.
- [32] Qiayuan Liao and Shang Yangxing. legged_control. https://github.com/qiayuanliao/legged_control, 2022.
- [33] Nathan Koenig and Andrew Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *2004 IEEE/RSJ international conference on intelligent robots and systems (IROS)(IEEE Cat. No. 04CH37566)*, volume 3, pages 2149–2154. IEEE, 2004.
- [34] Wolfgang Hess, Damon Kohler, Holger Rapp, and Daniel Andor. Real-time loop closure in 2d lidar slam. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 1271–1278. IEEE, 2016.
- [35] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023.
- [36] Xiangyu Chu, Shengzhi Wang, Minjian Feng, Jiayi Zheng, Yuxuan Zhao, Jing Huang, and K. W. Samuel Au. Model-free large-scale cloth spreading with mobile manipulation: Initial feasibility study. In *2023 IEEE 19th International Conference on Automation Science and Engineering (CASE)*, pages 1–6, 2023.