DRAGON: A Dialogue-Based Robot for Assistive Navigation with Visual Language Grounding

Shuijing Liu, Aamir Hasan, Kaiwen Hong, Runxuan Wang, Peixin Chang, Zachary Mizrachi, Justin Lin, D. Livingston McPherson, Wendy A. Rogers, and Katherine Driggs-Campbell

***Abstract*— Persons with visual impairments (PwVI) have difficulties understanding and navigating spaces around them. Current wayfinding technologies either focus solely on navigation or provide limited communication about the environment. Motivated by recent advances in visual-language grounding and semantic navigation, we propose DRAGON, a guiding robot powered by a dialogue system and the ability to associate the environment with natural language. By understanding the commands from the user, DRAGON is able to guide the user to the desired landmarks on the map, describe the environment, and answer questions from visual observations. Through effective utilization of dialogue, the robot can ground the user's free-form descriptions to landmarks in the environment, and give the user semantic information through spoken language. We conduct a user study with blindfolded participants in an everyday indoor environment. Our results demonstrate that DRAGON is able to communicate with the user smoothly, provide a good guiding experience, and connect users with their surrounding environment in an intuitive manner. Video and code are at https://sites.google.com/view/dragon-wayfinding/home.**

I. Introduction

Wayfinding, defined as helping people orient themselves in an environment and guiding them from place to place, is a longstanding challenge for persons with visual impairments (PwVI) [1], [2]. A guiding robot that can verbally interact with PwVI and connect language to the world, such as finding a destination or helping the user understand the environment, has the potential to improve the quality of their lives and to also reduce the load on their caregivers [2]–[4].

A large body of previous wayfinding guides tackles joint dynamics and path planning for the human-robot team [5]–[7]. Another line of work pairs wayfinding with communication signals such as navigation instructions [8], [9] and basic environment information [10], [11]. As a step further, other wayfinding technologies recognize and convey the semantic meaning of the surrounding environment such as naming the landmarks [12]–[14]. However, these methods require special environmental setups, such as multiple RFID tags and bluetooth beacons. We aim to remove dependence on these types of special infrastructure by integrating advances in visual-language grounding into conversational wayfinding.

On the other hand, technologies in vision-language navigation and voice-controlled robots have made significant progress [15]–[17]. These navigation agents are able to perform various tasks according to natural language commands such as "Bring me a cup." with simple onboard sensors. This is usually achieved by encoding visual landmarks in a semantic map and associating language with these landmarks during navigation, which is referred to as visual-language grounding [16], [18]. However, these general-purpose frameworks assume that humans can provide step-by-step navigation instructions. Thus, they are not built for PwVI that often need help in perceiving the environment and planning paths. Although attempts have been made to integrate vision or language models in wayfinding [11], [19], [20], how to build a guiding robot that can intuitively exchange semantic information with users remains an open challenge.

Fig. 1: DRAGON identifies the intents of the user through dialogue, grounds language with the environment, and guides the user to their desired goal.

In this paper, we propose DRAGON, a **D**ialogue-based **R**obot for **A**ssistive navigation with visual-language **Groun**ding. In Fig. 1, DRAGON uses speech to communicate with the user and a physical handle for navigation guidance. The dialogue and navigation can be executed simultaneously. When the user gives a speech command, Speech Recognition (SR) and Natural Language Understanding (NLU) modules first extract the user's intents and desired destinations. The main grounding functionalities include: (1) finding users' desired destinations with a visual-language model [21] and guiding them to the destinations; (2) describing nearby objects with an object detector [22]; and (3) answering questions from users using a Visual Question Answering (VQA) model [23]. With dialogue and grounding, DRAGON can effectively navigate users to their desired destinations and help them gain awareness of their surroundings in an intuitive manner.

To find users' intended goals on a map, we propose a novel landmark recognition module based on CLIP [21]. After a straightforward mapping process, the landmark recognizer is able to select the landmark whose image best matches

S. Liu, A. Hasan, K. Hong, R. Wang, P. Chang, J. Lin, Z. Mizrachi, D. McPherson and K. Driggs-Campbell are with the Department of Electrical and Computer Engineering at the University of Illinois Urbana-Champaign. W. A. Rogers is with the Department of Applied Health Sciences at the University of Illinois Urbana-Champaign. emails: {sliu105, aamirh2, kaiwen2, runxuan6,pchang17, jklin6, zdm3, dlivm, wendyr, krdc}@illinois.edu

the user descriptions. Our landmark recognizer is able to associate flexible and open-vocabulary commands with few constraints on user expressions. If the description is ambiguous, our system will disambiguate user intents through additional dialogue. Then, the corresponding goal location is passed to the path planners for navigation guidance, which allows for easy integration of our model into the navigation stack for mobile robots. Combined with the navigation stack, the powerful and reliable landmark recognizer is essential to ensure the success and user experience of DRAGON.

Our main contributions are as follows: (1) As an interactive navigation guide for PwVI, DRAGON enables voice-based dialogue, which carries rich information and has grounding capabilities; (2) We propose a novel landmark mapping and recognition method that can associate free-form language commands with image landmarks. Our method can be easily plugged into the navigation stack of mobile robots; and (3) A user study (N=5) with blindfolded participants demonstrates that DRAGON is able to understand user intents through dialogue and guide them to desired destinations in an intuitive manner. To the best of our knowledge, our work is the first to show that visual-language grounding via dialogue benefits robotic assistive navigation.

II. Related Works

A. Wayfinding robots and technologies

Navigation guidance: To guide PwVI from point A to point B following a planned path, unactuated devices, such as smartphones and wearables, rely on haptic or audio feedback to give instructions such as going straight and turning right [9], [12], [13], [20]. However, delays and misunderstandings might lead to inevitable deviations, which take time and effort to recover from [10], [12]. On the other hand, robots provide a physical holding point, which offers kinesthetic feedback to minimize deviations and reduce the mental load of users [6], [24], [25]. Such physical guidance can be combined with aforementioned verbal or haptic navigation instructions to further improve performance at the cost of a more expensive system [8], [11]. To ensure both efficiency and low cost, we mount a handle on our robot to give intuitive real-time steering feedback in navigation.

Semantic communication: A large part of blind navigation technologies ignores exchanging environmental information with users [6]–[8]. To deal with this issue, CaBot applies object recognition to describe the user's neighborhood, yet the user cannot hold conversations with the robot or choose their destinations [11]. To enable users to choose a semantic goal (*e.g.* a restroom), some works mark points of interest using bluetooth beacons [12], [13] or RFID tags [14], [24], which requires heavy instrumentation. As an alternative, extracting semantic information from ego-centric camera images is much cheaper and easier. For example, SeeWay uses skybox images to represent landmarks [20]. Similarly, Landmark AI offers semantic-related functionalities including describing the environment, reading road signs, and recognizing landmarks using a phone camera [19]. However, these phone applications are not robots and thus cannot physically guide users or provide a stable mounting point for cameras. In contrast, Table I shows that DRAGON brings conversational wayfinding to the next level: A robot can simultaneously offer physical guidance and enable users to trigger a variety of functionalities through dialogue.

B. Command following navigation

Tremendous efforts have been made in understanding and grounding human language instructions for various robotic tasks [15]–[17]. In command following navigation, a modular pipeline usually consists of three modules: (1) an NLU system to map instructions to speaker intent; (2) a grounding module to associate the intent with physical entities; and (3) a SLAM and a planner to generate feasible trajectories [18], [26], [27]. Other works attempt to learn end-to-end policies from simulated environments or datasets [28]–[31]. However, due to sim-to-real gaps in perception, language, and planning, deploying these policies to the real world remains an open challenge for applications in the low data regime such as wayfinding [32]. Therefore, we adopt the modular pipeline to ensure performance in the real world.

C. Semantic landmark recognition

Understanding the semantic meanings of a scene is a vital step towards interactive navigation [16], [18]. Some works reconstruct volumetric maps for the environment, where each grid is associated with a semantic label [16], [27], [33]. Other works build more abstract scene graphs [34], [35]. However, implementing these methods on a real robot is expensive, as they require accurately calibrated depth cameras and high-performing instance segmentation models.

TABLE I: Benchmark for conversational wayfinding technologies. A ✓ means that the functionality is implemented. A ❍ means partial implementation. A blank cell means the functionality is absent. (In [24], the users have to enter a number sequence into a keypad to specify their destinations. [19] can only describe a fixed set of pre-mapped landmarks and can only answer two fixed questions.)

Method	User-chosen semantic goals	Speech dialogue Input	Speech dialogue Output	Enviornment description	VQA	Form	Environmental Instrumentation
GuideBeacon [12]	✓	✓	✓			Phone application	Bluetooth beacons
NavCog3 [13]	✓	✓	✓	❍	❍	Phone application	Bluetooth beacons
LandmarkAI [19]	✓	✓	✓	✓	❍	Phone application	GPS
SeeWay [20]	✓	✓	✓			Phone application	WiFi
Robotic Shopping cart [24]	❍		✓			Robot	RFID tags
CaBot [11]			✓	✓		Robot	Remote joystick
Ballbot [8]	✓	✓	✓			Robot	WiFi + Remote computer
Ours	✓	✓	✓	✓	✓	Robot	WiFi + Remote computer

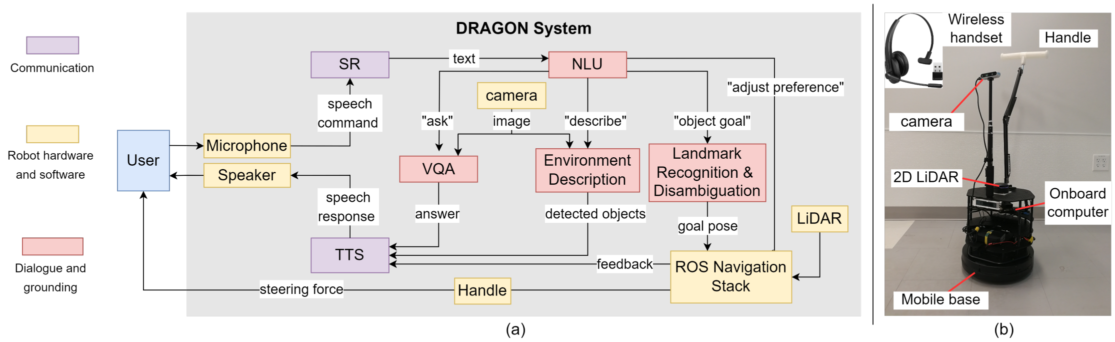

Fig. 2: **An overview of the system and platform of DRAGON**. (a) Submodules, message passing, and user interface. (b) The robot platform.

Another line of work collects images as landmarks to create topological graphs [36]–[38]. In navigation, the goal location is retrieved by computing the similarity between a goal image and all stored landmarks. However, the above works only consider image goals, which are less natural than language in human-centered applications. Inspired by Shah *et al*. [18] and Huang *et al*. [16], we use CLIP [21] to associate image landmarks with users' language commands. Compared with previous works that use closed vocabulary object detectors, which are limited to a predefined set of semantic classes [33]–[35], our method is able to handle more flexible and open-vocabulary commands. We use CLIP to select landmarks and keep traditional cost maps for planning, enabling easy integration of our method into the navigation stack of mobile robots.

III. System Overview

In this section, we describe the setup and configuration of our robot guide. Fig. 2(a) shows an overview of our proposed system consisting of three main components: (1) The TurtleBot platform (yellow); (2) Audio communication interface (purple); (3) Dialogue and grounding modules (red). The modules communicate with each other through ROS.

A. Robot platform

Overview: We use the Turtlebot2i as our robot platform. As shown in Fig. 2(b), the robot is fitted with the following sensors and equipment: (1) An RP-Lidar A3 laser range finder is mounted on the top of the robot structure for SLAM; (2) An Intel RealSense D435i camera is mounted on the top of a monopod facing forward for scene description and question answering; (3) A wireless headset is used to communicate with the user. The headset is lightweight and maximally protects the users' privacy, while the absence of wires avoids tripping hazards; (4) A T-shaped handle is attached to the top rear side of the robot as a holding point for the user's arm. The handle allows users to choose their preferred holding configurations such as one hand or two hands. The robot is connected to a desktop computer which provides more computation resources through WiFi.

Planning and Navigation: The robot operations are managed by the standard ROS `move_base` navigation stack. Before navigation, we create a 2D occupancy map of the environment using laser-based SLAM and mark the semantic landmarks at the same time (see Fig. 3 and Sec. IV-B for details). At the beginning of each trial, the goal pose is obtained from the dialogue with the user (further specified in Sec. IV). During navigation, adaptive Monte Carlo localization is used to localize the robot on the map. We use the dynamic window approach (DWA) [39] and A* as local and global planners, respectively. The minimum translational velocity is restricted to be non-negative to prevent the robot from moving backward and colliding with the user. The maximum velocity of the robot can be adjusted by the user (see Sec. IV-D for details).

B. Audio Communication Interface

Speech is a natural choice for human-robot communication, particularly in cases where the human has limited vision [2]–[4]. To this end, as shown in purple in Fig. 2(a), we develop an audio communication interface that consists of: (1) **Input:** When audio is captured by the `audio_capture` package, the OpenAI Whisper speech recognition model [40] transcribes speech commands to text, which are passed to the NLU module. The SR module continuously transcribes the audio from the microphone and publishes the text transcriptions to a ROS topic in real time; (2) **Output:** We use the Google text-to-speech (TTS) service to convert the text output from the visual-language modules and navigation module to speech, which is then narrated to the user via the headset. The TTS is another ROS topic that converts and plays the synthesized sound constantly.

C. Dialogue and grounding

In the dialogue system, when SR publishes a transcribed sentence, an NLU first extracts intents and landmarks (if any) from the sentence. Then, based on the intents, the information from the sentence is passed to different visual-language grounding modules which output language responses to TTS and/or the navigation module which changes navigation behaviors. By doing so, the robot fulfills the user's request. The dialogue and grounding modules are detailed in Sec. IV.

IV. Dialogue And Grounding

The goal of DRAGON is to connect the user with the environment through conversation. In this section, we describe how our dialogue system understands user language (Sec. IV-A), maps and localizes semantic landmarks (Sec. IV-B), provides information about the environment (Sec. IV-C), and adjusts the navigation preference of the user (Sec. IV-D). Our grounding system is visualized in the red parts of Fig. 2(a).

A. Natural language understanding (NLU)

The NLU takes a transcribed sentence as input and outputs user intents and entities of interest. Table II shows all possible user intents. The entities are locations, objects, and object attributes which include the material and functionalities of an object. We use Dual Intent and Entity Transformer for intent classification and entity recognition [41]. We train the model using a custom dataset with 1092 sentences. For each intent, we collect various expressions including misspelled and phonetically similar phrases, which makes our NLU robust to the nuances of human language and the errors caused by the SR. For example, "a think" and "a sink" both refer to the kitchen sink. We also collected expressions for multi-intents and unknown intents so that the NLU can fulfill a request containing multiple intents and ignore noise input. For instance, "Hello robot, can you take me to a sofa?" will both activate the robot and set an object goal. Once the intent and entities are extracted, the corresponding downstream module is activated and the robot informs the user via language feedback. The NLU may pass additional input arguments to modules such as extracted entities or the whole sentence.

During navigation, the landmark recognition is triggered if the user intent is *Object goal* and the NLU extracts a goal object from the input sentence. If the user mentions additional information about the landmark, we use simple prompt engineering to make the description more specific. For example, locations and attributes of objects can be added to the description by "a chair in the office" or "a gaming chair." In addition, the robot uses clarification dialogue to disambiguate the desired landmark if the input description does not contain any object. If the user only provided the location or attributes without mentioning the object name (*e.g.* "Take me to the kitchen"), our system provides hints to encourage the user to provide more specific descriptions (*e.g.* "What object are you looking for in the kitchen?"). If there are multiple similar objects in different landmarks, our system disambiguates the user's preferred landmark (*e.g.* "What kind of chair are you looking for? A dining chair, an office chair, or a sofa?"). After choosing a unique landmark, our system confirms the goal with the user (*e.g.* "Do you wish to go to a dining chair?"). No further action is taken until the user affirms the goal. With the disambiguation and confirmation dialogue, the NLU is able to precisely capture the user's desired destination, which is crucial for the whole navigation experience.

B. Landmark mapping and recognition

To guide the user to their object goals, we first record the images and locations of landmarks during SLAM. Then, we use a fine-tuned CLIP model to match the user's description with goal images, whose corresponding location is sent to the navigation stack for navigation guidance.

The landmark mapping process is performed simultaneously with SLAM. During SLAM, when the robot is at a landmark that might be a point of interest, we simply save the current robot pose in the map frame and an RGB image of the landmark to the disk with a single key press. No labels or text descriptions are needed at this stage. The resulting landmark map is shown in Fig. 3.

TABLE II: All user intents and their descriptions.

Intents	Descriptions
Greet	Wake up the robot and begin an interaction.
Object goal	Go to a specific object landmark. May contain entities including objects and attributes.
Location goal	Go to a rough goal location (kitchen, lounge, etc) without mentioning a specific object. May contain location entities.
Affirm	Confirm the goal.
Deny	Deny the goal.
Describe	Ask for a description of the surrounding environment.
Ask	Ask a question about the surrounding environment.
Pause	Pause the current navigation.
Resume	Resume the current navigation.
Accelerate	Increment velocities, up to a limit.
Decelerate	Decrement velocities, down to a limit.
Unknown	The text does not belong to any intents above (i.e. be noise, chitchat, etc) and is ignored by the robot.

During navigation, this module is activated when the extracted intent is *Object goal*. After the goal is confirmed by the NLU, the CLIP model selects the landmark whose image has the highest similarity score with the descriptions of landmarks. To obtain the image-text similarity score, a text encoder and an image encoder first convert the input text and all images to vector embeddings. Then, the text and image similarity score is computed by the cosine similarity between the pairwise text and image embeddings. The image with the highest similarity score is selected as the goal. Finally, the corresponding location of the chosen landmark on the map is sent to an action client, which sets the goal for the robot.

The zero-shot performance of pre-trained CLIP models is not satisfactory in our environment due to distribution shifts. As shown in Fig. 3, the objects in the images are frequently cropped due to the low mounting point of our camera and the close distance between the camera and the objects. In addition, the descriptions of landmarks from a PwVI might be vaguer than those in public datasets (*e.g.* "a chair" v.s. "a blue chair in front of a white wall"). To this end, we fine-tune the CLIP model with a custom dataset containing 544 image and text description pairs from our environment with a 8×10^{-6} learning rate for 35 epochs. By using an open-vocabulary model to recognize landmarks, DRAGON is able to handle free-form language and is not limited to a fixed set of object classes. Thus, the user input expressions are less restricted, making the grounding module easier for non-experts to use.

C. Environment understanding modules

To help the user gain awareness of their surroundings, we use an object detector [22] to describe the objects (activated if the intent is *Describe*) and a VQA model [42] to answer the user's questions (activated if the intent is *Ask*). Both models take the current camera image as input.

The output of the object detector consists of a list of detected instances, their object classes, confidence scores, and bounding boxes. To avoid narrating a long list and to keep the description concise, we post-process the output as follows. We first apply non-maximum suppression and filter out the detected instances with low confidence scores. Then, for the remaining instances, we keep the top three classes with the largest average bounding boxes, and list the object class names together with the numbers of objects (*e.g.* "2 chairs, 1 person, and 1 table").

The VQA model takes the current camera image and the user's question from the SR and outputs a short answer to the question. Similar to CLIP, we collect a dataset of 10252 (image, question, answer) triplets to fine-tune the VQA model for 20 epochs. To handle free-form user expressions, the dataset contains cases where multiple questions have the same meaning but different phrasing (*e.g.* "Is any person in front of me?" and "Anyone here?").

The processed output of the object detector and the answer from the VQA model is sent to the TTS topic, which is narrated to the user in real time. Since both models can only take a single RGB image, our system cannot provide depth-based information or detect anything out of the camera view.

D. Navigation preference customization

To accommodate the different walking paces of users and to avoid tiring the user during navigation, the robot can change its speed (activated if the intent is *Accelerate* or *Decelerate*), take a pause (*Pause*), and resume (*Resume*). To pause the robot, our system stores and cancels the current goal from the action client in the navigation stack. To resume, the stored goal is sent to the action client again. To update the speed, we change the maximum translational and rotational velocities of the DWA local planner by updating its configuration through the `dynamic_reconfigure` package.

V. Experiments

This section describes a user study to evaluate our system. We also describe our baseline and evaluation metrics.

A. Baseline

We compare the CLIP-based landmark recognizer with a closed-vocabulary object detector as the baseline [22] [1]. The vocabulary size, or the number of classes, of the detector is more than 1200 and it is fine-tuned with the same number of data as CLIP. In the baseline, the landmark images are passed into the object detector, which outputs the class names of detected objects. During navigation, the baseline chooses the landmark with the highest number of objects mentioned by the user. Since the vocabulary of object detectors is fixed, the baseline is unable to incorporate an object's attributes or locations obtained from disambiguation. All other modules are the same for our system and the baseline.

B. User Study

Environment: All experiments were conducted in an everyday indoor environment in a university building. Three routes were created with furniture acting as obstacles. Fig. 3 provides a layout of the environment and the three routes. The routes are designed to have varying levels of difficulties for the system to correctly interpret the destination. Specifically, the goal landmark A of Route 1 contains simple objects, whereas landmark B of Route 2 contains more complicated objects, and landmark C of Route 3 contains a transparent door that is hard to recognize.

Fig. 3: **The map of our environment with semantic landmarks.** The images are landmarks with locations marked by red dots. The orange lines are the three routes in the user study. The red squares are the starting locations of routes.

TABLE III: Example expressions and their corresponding landmarks from CLIP v.s. the detector. The landmark labels are from Fig. 3. Underlined expressions are collected from the user study.

Landmark	CLIP	Detector
A	door, exit, entrance, gate glass door, automatic door	poster
B	sofa, couch, coach, fabric chair,relaxing chair thermostat, climate control	sofa, thermostat
C	sink, think,sync, faucet soap, hand wash, water pipe paper towel dispenser, bowls kitchen countertop, drying rack	faucet bottle dispenser bowl

Participants: The user study was conducted with N=5 participants (mean age=26; 3 males; 2 females; all participants were university students). All participants provided informed consent under IRB #23565, which was approved by University of Illinois Urbana Champaign on February 27th, 2023. All participants have full (corrected) vision and are asked to wear a blindfold to simulate a visual impairment. While our true target population are PwVI, the purpose of this pilot study is to validate the capabilities of DRAGON. Running experiments with PwVI is left for future work.

Procedure: Participants were first familiarized with the goals of the study and requested to fill a demographic and robot technology survey. Then, participants were provided with a test run to get familiar with the system and its intricate navigation feedback mechanism. To begin the trial, the users were asked to command the robot to take them to a predetermined goal destination. Participants were not constrained in either the vocabulary or the sentence structure

[1]Open-vocabulary object detectors exist. We choose a closed-vocabulary detector to represent a common closed-vocabulary grounding model.

Semantic goal recognition | Speed adjustment | Environment description | VQA

Fig. 4: **An example navigation trial with human-robot dialogue in the user study.** In the dialogue boxes, "H" denotes the human and "R" denotes the robot. The camera view is shown in the lower right corner.

of their speech commands. The users were also informed that they could interact with the robot (*e.g.* ask for a description of their surroundings) at any point of the navigation. After each route, we used a short questionnaire to measure the participant's perception of the system. A strictly structured post-survey interview was conducted after participants finished all three routes to collect their feedback with the system. The same procedure was performed for CLIP and the detector. The order of which method was tested first was randomized for each participant to minimize the bias introduced due to the order of testing. The full survey can be found at this link.

C. Metrics

Objective Metrics: We measure the accuracy of all interactions during the user study, including 312 NLU, 30 landmark recognition (LR) and navigation trials, 15 environment description (EnvDes), 74 VQA, and 15 navigation preference adjustment (NavAdj). The NLU is considered correct when the extracted intent and entities (if any) are both correct. We also measure the accuracy of the NLU by taking the correctness of SR into account to analyze the effect of wrong SR. The effect of wrong NLU outputs is ignored when evaluating its downstream modules. The correctness of answers from VQA is based on the camera images, not on the information out of the camera view. To compare the two LR methods, we measure the success rates of LR and the resulting navigation.

Subjective metrics: For both methods, we compare the scores for categories from the short questionnaire in Table VI. The difference in scores for each participant was aggregated and analyzed to discount individual biases. We evaluate user preferences for the other modules through a simple Likert scale analysis on the responses from the post-survey interview. Additionally, participants' feedback is summarized for qualitative analysis.

VI. Results

In this section, we discuss the results of our user study. A demonstration can be found in the supplementary video with results from each module and navigation trials that test the whole system. Fig. 4 also provides an example trial along with the dialogue during the user study.

A. Quantitative Evaluation

LR and navigation: As seen in Table IV, the success rate of navigation is 100% if LR succeeds. This dependency indicates that the performance of LR is the key factor for navigation in the DRAGON system.

For LR, as shown in Table IV, our CLIP model with disambiguation outperforms the detector baseline by achieving 100% success rate in LR and navigation with fewer rounds of dialogue on average. We attribute this result to the fact that CLIP is an open vocabulary model that can take free-form query text, which is essential for our task because the user may use different expressions to refer to the same landmark. On the contrary, a closed vocabulary object detector can only handle a fixed set of object classes with limited expressions. For example, in Table III, although both models can handle different objects that belong to the same landmark, CLIP can associate synonyms, such as "sofa" and "couch", and wrong transcriptions, such as "coach", to the correct landmark. In contrast, the closed-vocabulary detector can only handle strictly fixed expressions. The detector misidentifies some objects such as the transparent door in Landmark C after fine-tuning. Since our target users are usually non-experts, the baseline sometimes needs the user to rephrase multiple times to recognize the goal, which causes the user to run out of patience, and results in failure or more rounds of dialogue.

Besides CLIP, the disambiguation dialogue also contributes to the performance. With disambiguation, additional information such as the material and functionality of objects can be merged into the query text, such as "fabric chair" and "relaxing chair" as shown in Table III. These rich descriptions are helpful in distinguishing landmarks that have the same objects with different attributes, such as the different types of chairs in Landmark A, D, and E in Fig. 3 with fewer rounds of user rephrasing.

NLU: In Table V, the overall accuracy of NLU is over 15% higher than SR, as the NLU is trained with incorrectly transcribed text and thus can work even when SR is incorrect.

TABLE IV: Success rates (%) of LR and navigation (including overall success rate, and success rate if LR is correct), and the average number of dialogue rounds for a successful LR.

Method	LR		Navigation	
	Overall	# rounds	Overall	Correct LR
Ours	100	2.4	100	100
Baseline	46.67	3.71	46.67	100

TABLE V: Accuracies (%) of the SR, NLU (including overall accuracy, accuracy if SR is correct and if SR is wrong), EnvDes with fully correct and partially correct number of objects, VQA, and navigation adjustment modules.

SR	NLU			EnvDes		VQA	NavAdj
	Overall	Correct SR	Wrong SR	Full	Partial		
70.19	85.26	93.61	65.59	45.45	75.76	82.43	100

However, we do notice that NLU performs better with correct SR. The common failure cases of NLU occur when (1) The SR mistakenly breaks a sentence into two halves (*e.g.* "Is there anything?" and "To my right." are treated as two sentences); and (2) The NLU does not correctly extract intents from noisy transcriptions and chitchat, which can happen during the user study. Thus, we believe that a better SR engine would vastly benefit the performance of the whole system. However, since DRAGON will not begin navigation until the user confirms the goal in the dialogue, the wrong SR and NLU have little effect on navigation.

Other Modules: The system's environment descriptions are sometimes inaccurate due to errors in the object detector such as: (1) detecting incorrect number of objects (*e.g.* 3 wall sockets, when there was only 1 present); and (2) incorrect object classifications of rare or uncommon objects (*e.g.* a building information tablet was classified as a poster). Although we use non-maximum suppression and confidence score threshold to reduce the errors, they are hard to entirely eliminate due to the data distribution shift and the blurry images caused by the robot motion. Nevertheless, in Table V, the model is able to output a list of objects with correct class names in 75.76% of the cases, which might be more important to the user than a correct number of objects.

The VQA module accurately answers the user's questions in 82.43% of the cases. The model fails in cases where the user asks questions that the robot cannot answer based on a single RGB image. For example, without precise depth information the VQA model only answers "far" or "close" if the question is "How far is the person from me?". Without a wider field of view, the model outputs objects on the front side if the question is "What is on my right?".

B. Qualitative Evaluation

In Table VI, participants showed an increasing preference for DRAGON with CLIP over the detector in all user experience categories across all routes. Specifically, participants reported a 32% improvement with a mean score difference of 1.60 ± 0.89 in the overall experience and a mean score difference of 1.40 ± 0.89 in the communication experience. The difference increases as the goal landmark contains more complicated objects in Route 2, and objects that are difficult to detect in Route 3, where the failures in LR significantly lower the user score for the detector based system. Particularly, participants noted that DRAGON with CLIP understood their intent, asked good follow-up questions, and correctly guided them to their destination. In contrast, the closed-vocabulary detector failed at these aspects and occasionally was unable to recognize destinations even though they existed. Participants also noted that the failures in intent understanding led to a frustrating communication experience with the detector.

One user in particular mentioned that the CLIP based model "... was able to actually understand me, so it accurately took me to the location and correctly answer [sic] my questions." while the detector based model "... would confirm the location I wanted to go to but could not find [sic; participant meant understand] the right location". However, users also mentioned potential improvements for DRAGON including more detailed descriptions of the environment, a quicker response time, and warnings of potential dangers such as "We're going through a door."

TABLE VI: Mean user experience scores on a scale of 1 to 10.

Use experience category	Route 1		Route 2		Route 3	
	CLIP	Detector	CLIP	Detector	CLIP	Detector
Ease of following	**8.8**	8.6	**8.8**	5.6	**9.2**	1.0
Navigational Experience	**8.4**	7.4	**7.6**	4.8	**8.8**	1.0
Intent Understanding	7.6	**8**	**7.6**	4.6	**8.4**	3.4

For the user experience categories that are the same for both LR methods, such as the 'intuitiveness of communication interface' and the ability of the system to aid in 'gaining awareness of the environment', participants reported average scores of 7.07 ± 2.17 and 6.07 ± 3.21, respectively. As evidenced by these scores, the users' opinions regarding these two categories were positive, due to the inclusion of the dialogue and grounding modules. However, participants highlighted minor inaccuracies in the environment descriptions and the slow pace of communication due to processing times and network delays as potential issues.

VII. Conclusion and Future Work

In conclusion, we present DRAGON, a first-of-its-kind guide robot that fulfills user intents and familiarizes the user with their surroundings through interactive dialogue. We use CLIP to retrieve landmark destinations from commands and provide visual information through language. The user study shows promising communication, grounding, and navigation performance of DRAGON. Our work suggests that visual-language grounding and dialogue can greatly improve human-robot interaction.

To extend DRAGON and address its limitations, we point out the following directions for future work. First, the current dialogue system is rule-based with fixed behaviors for each intent. Replacing hard-coded rules with adaptive learning-based policies should generalize to more complex user behaviors and more subtasks. Second, the environment understanding modules provide limited information. Future informative descriptions should include object relationships in images, incorporate information from the map and other sensors, and inform users about the planned path and potential dangers. Finally, the physical interface of the platform should be redesigned to improve the user experience. DRAGON demonstrates the feasibility of dialogue and visual-language grounding in assistive navigation that future research in dialogue management, computer vision, and robotics can explore further.

Acknowledgements

This work was supported by a Research Support Award from the University of Illinois Urbana-Champaign Campus Research Board. We thank Dr. Samuel Olatunji and Dr. Peter Du for their help on the user study, Ye-Ji Mun and Zhe Huang and Neeloy Chakraborty for their feedback on the paper draft, and all participants in the user study.

References

[1] V. Kulyukin, C. Gharpure, J. Nicholson, and S. Pavithran, "Rfid in robot-assisted indoor navigation for the visually impaired," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2004.

[2] M. A. Bayles, T. Kadylak, S. Liu, A. Hasan, W. Liang, K. Hong, K. Driggs-Campbell, and W. A. Rogers, "An interdisciplinary approach: Potential for robotic support to address wayfinding barriers among persons with visual impairments," in *Human Factors and Ergonomics Society Annual Meeting (HFES)*, 2022.

[3] S. Azenkot, C. Feng, and M. Cakmak, "Enabling building service robots to guide blind people a participatory design approach," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2016.

[4] P. Bhat and Y. Zhao, "'i was confused by it; it was confused by me:' exploring the experiences of people with visual impairments around mobile service robots," *ACM on Human-Computer Interaction*, 2022.

[5] E. OhnBar, K. Kitani, and C. Asakawa, "Personalized dynamics models for adaptive assistive navigation systems," in *Conference on Robot Learning (CoRL)*, 2018.

[6] A. Nanavati, X. Z. Tan, J. Connolly, and A. Steinfeld, "Follow the robot: Modeling coupled human-robot dyads during navigation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.

[7] A. Xiao, W. Tong, L. Yang, J. Zeng, Z. Li, and K. Sreenath, "Robotic guide dog: Leading a human with leash-guided hybrid physical interaction," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.

[8] Z. Li and R. Hollis, "Toward a ballbot for physically leading people: A human-centered approach," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.

[9] L. Jin, H. Zhang, and C. Ye, "A wearable robotic device for assistive navigation and object manipulation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.

[10] J. Wilson, B. N. Walker, J. Lindsay, C. Cambias, and F. Dellaert, "Swan: System for wearable audio navigation," in *IEEE International Symposium on Wearable Computers*, 2007.

[11] J. Guerreiro, D. Sato, S. Asakawa, H. Dong, K. M. Kitani, and C. Asakawa, "Cabot: Designing and evaluating an autonomous navigation robot for blind people," in *International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS)*, 2019.

[12] S. A. Cheraghi, V. Namboodiri, and L. Walker, "Guidebeacon: Beacon-based indoor wayfinding for the blind, visually impaired, and disoriented," in *IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 2017.

[13] D. Sato, U. Oh, K. Naito, H. Takagi, K. Kitani, and C. Asakawa, "Navcog3: An evaluation of a smartphone-based blind indoor navigation assistant with semantic features in a large-scale environment," in *International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS)*, 2017.

[14] V. Kulyukin, C. Gharpure, J. Nicholson, and G. Osborne, "Robot-assisted wayfinding for the visually impaired in structured indoor environments," *Autonomous robots*, 2006.

[15] P. Chang, S. Liu, and K. Driggs-Campbell, "Learning visual-audio representations for voice-controlled robots," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.

[16] C. Huang, O. Mees, A. Zeng, and W. Burgard, "Visual language maps for robot navigation," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.

[17] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, *et al.*, "Do as i can, not as i say: Grounding language in robotic affordances," in *Conference on Robot Learning (CoRL)*, 2022.

[18] D. Shah, B. Osiński, brian ichter, and S. Levine, "LM-nav: Robotic navigation with large pre-trained models of language, vision, and action," in *Conference on Robot Learning (CoRL)*, 2022.

[19] M. Saha, A. J. Fiannaca, M. Kneisel, E. Cutrell, and M. R. Morris, "Closing the gap: Designing for the last-few-meters wayfinding problem for people with visual impairments," in *International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS)*, 2019.

[20] Z. Yang, L. Yang, L. Kong, A. Wei, J. Leaman, J. Brooks, and B. Li, "Seeway: Vision-language assistive navigation for the visually impaired," in *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2022.

[21] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, "Learning transferable visual models from natural language supervision," in *International Conference on Machine Learning (ICML)*, 2021.

[22] X. Zhou, R. Girdhar, A. Joulin, P. Krähenbühl, and I. Misra, "Detecting twenty-thousand classes using image-level supervision," in *European Conference on Computer Vision (ECCV)*, 2022.

[23] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh, "Vqa: Visual question answering," in *IEEE International Conference on Computer Vision (ICCV)*, 2015.

[24] V. A. Kulyukin and C. Gharpure, "Ergonomics-for-one in a robotic shopping cart for the blind," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2006.

[25] Y. Zhang, Z. Li, H. Guo, L. Wang, Q. Chen, W. Jiang, M. Fan, G. Zhou, and J. Gong, "'i am the follower, also the boss': Exploring different levels of autonomy and machine forms of guiding robots for the visually impaired," in *ACM Conference on Human Factors in Computing Systems (CHI)*, 2023.

[26] R. Liu and X. Zhang, "A review of methodologies for natural-language-facilitated human-robot cooperation," *International Journal of Advanced Robotic Systems*, vol. 16, no. 3, 2019.

[27] S. Y. Min, D. S. Chaplot, P. Ravikumar, Y. Bisk, and R. Salakhutdinov, "Film: Following instructions in language with modular methods," in *International Conference on Learning Representations (ICLR)*, 2022.

[28] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[29] C. Chen, U. Jain, C. Schissler, S. V. A. Gari, Z. Al-Halah, V. K. Ithapu, P. Robinson, and K. Grauman, "Soundspaces: Audio-visual navigation in 3d environments," in *European Conference on Computer Vision (ECCV)*, 2020.

[30] P. Chang, S. Liu, H. Chen, and K. Driggs-Campbell, "Robot sound interpretation: Combining sight and sound in learning-based control," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.

[31] S. Y. Gadre, M. Wortsman, G. Ilharco, L. Schmidt, and S. Song, "Cows on pasture: Baselines and benchmarks for language-driven zero-shot object navigation," *arXiv*, 2022.

[32] H. Zhu, J. Yu, A. Gupta, D. Shah, K. Hartikainen, A. Singh, V. Kumar, and S. Levine, "The ingredients of real world robotic reinforcement learning," in *International Conference on Learning Representations (ICLR)*, 2020.

[33] J. McCormac, A. Handa, A. Davison, and S. Leutenegger, "Semanticfusion: Dense 3d semantic mapping with convolutional neural networks," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.

[34] N. Hughes, Y. Chang, and L. Carlone, "Hydra: A real-time spatial perception system for 3D scene graph construction and optimization," in *Robotics: Science and Systems*, 2022.

[35] S.-C. Wu, J. Wald, K. Tateno, N. Navab, and F. Tombari, "Scenegraphfusion: Incremental 3d scene graph prediction from rgb-d sequences," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.

[36] N. Savinov, A. Dosovitskiy, and V. Koltun, "Semi-parametric topological memory for navigation," in *International Conference on Learning Representations (ICLR)*, 2018.

[37] X. Meng, N. Ratliff, Y. Xiang, and D. Fox, "Scaling local control to large-scale topological navigation," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.

[38] D. S. Chaplot, R. Salakhutdinov, A. Gupta, and S. Gupta, "Neural topological slam for visual navigation," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.

[39] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics and Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.

[40] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," 2022.

[41] T. Bunk, D. Varshneya, V. Vlasov, and A. Nichol, "Diet: Lightweight language understanding for dialogue systems," *arXiv*, 2020.

[42] W. Kim, B. Son, and I. Kim, "Vilt: Vision-and-language transformer without convolution or region supervision," in *International Conference on Machine Learning (ICML)*, 2021.