

Bridging Zero-shot Object Navigation and Foundation Models through Pixel-Guided Navigation Skill

Wenzhe Cai, Siyuan Huang, Guanran Cheng, Yuxing Long, Peng Gao, Changyin Sun, Hao Dong

Abstract—Zero-shot object navigation is a challenging task for home-assistance robots. This task emphasizes visual grounding, commonsense inference and locomotion abilities, where the first two are inherent in foundation models. But for the locomotion part, most works still depend on map-based planning approaches. The gap between RGB space and map space makes it difficult to directly transfer the knowledge from foundation models to navigation tasks. In this work, we propose a Pixel-guided Navigation skill (PixNav), which bridges the gap between the foundation models and the embodied navigation task. It is straightforward for recent foundation models to indicate an object by pixels, and with pixels as the goal specification, our method becomes a versatile navigation policy towards all different kinds of objects. Besides, our PixNav is a pure RGB-based policy that can reduce the cost of home-assistance robots. Experiments demonstrate the robustness of the PixNav which achieves 80+% success rate in the local path-planning task. To perform long-horizon object navigation, we design an LLM-based planner to utilize the commonsense knowledge between objects and rooms to select the best waypoint. Evaluations across both photorealistic indoor simulators and real-world environments validate the effectiveness of our proposed navigation strategy. More details are accessible via our project website <https://sites.google.com/view/pixnav/>.

I. INTRODUCTION

Zero-shot object navigation is considered fundamental for home-assistance robots. Before a robot can interact with objects and complete many complex tasks, the robot should be able to navigate and find the target objects with a high success rate. Besides, with all kinds of different human demands, the robot navigation policy should be able to generalize to a variety of objects. This task requires three essential abilities: visual grounding for novel objects recognition, commonsense inference for semantic exploration and path-planning for robot locomotion. For the first two necessities, recent prosperity in foundation models provide powerful visual and language perception capabilities [1], [2], [3], [4], [5], [6]. But for the locomotion part, most of the zero-shot object navigation works still tethered to the map-based planning approaches [7], [8], [9], [10], [11]. There are two main limitations of such navigation systems. Firstly, as most foundation models are trained with internet RGB images, it is difficult to directly integrate such prior knowledge into a map-based locomotion method. Secondly, map-based approaches require accurate depth sensing and localization, which often comes at a

Wenzhe Cai, Guanran Cheng and Changyin Sun are with the School of Automation, Southeast University. Yuxing Long and Hao Dong are with Hyperlane Lab, CFCs, School of CS, Peking University and National Key Laboratory for Multimedia Information Processing. Siyuan Huang and Peng Gao are with the Shanghai AI Laboratory. This work was supported by National Natural Science Foundation of China (No. 62136001) and National Youth Talent Support Program (8200800081).

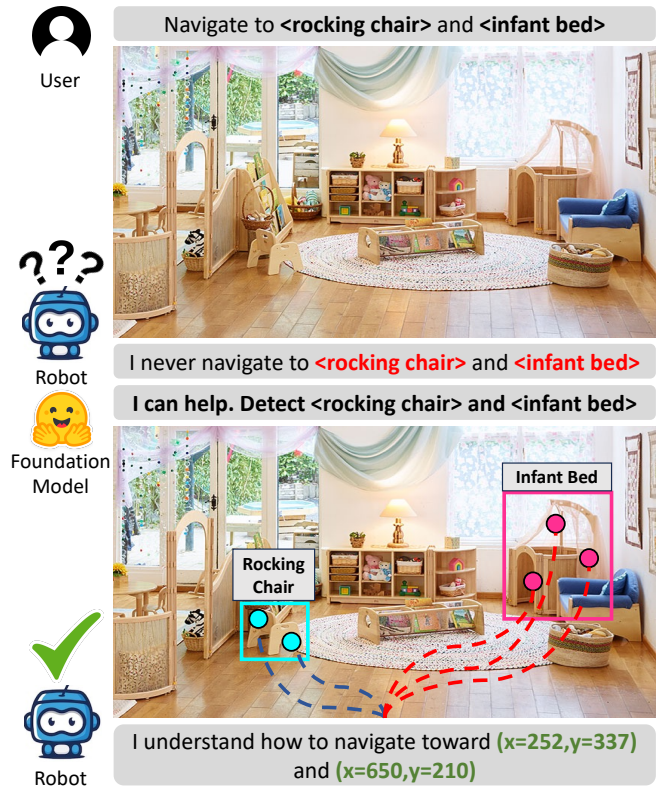


Fig. 1. For the home-assistance robot, humans may ask for searching and interacting with uncommon objects like ‘a rocking chair’ and ‘a infant bed’. It is challenging for the robot to develop a multi-modal navigation policy that can accommodate various types of objects using both text and image inputs. However, since foundation models are capable of zero-shot image understanding, any objects can be indicated by a single pixel when utilizing such foundation models. Therefore, we train a navigation policy with pixels as goal specifications, enabling the robot to navigate towards arbitrary objects. Since each object typically comprises hundreds of pixels, this approach offers a wide range of possible navigation trajectories, which greatly enlarges the scale of the navigation dataset.

higher cost and can be easily disturbed compared with RGB cameras. This raises a pertinent question: Can we utilize the superpower of foundation models and develop a pure RGB-based intelligent navigation system?

In this work, we posit that pixels serve as an ideal bridge, seamlessly connecting foundational models to the navigation task. Since most foundation models are trained with RGB images, aligning such foundation models with an RGB-based navigation skill is much easier than a map-based skill. Meanwhile, the task of navigating towards arbitrary objects can be converted into navigating towards a designated pixel. Therefore, we propose a **Pixel-guided Navigation skill (PixNav)** to complete the last-mile navigation process. Different from the previous object-goal [12], [13] skill, which

uses a word or one-hot vector as the goal specification, PixNav avoids learning the multi-modal alignment for the specific navigation task but still versatile for navigating to any kind of objects. Moreover, to ensure the generalization performance of a learning-based robotics policy, the dataset scale is a crucial factor [14], [15]. For our PixNav, it is much easier to collect large-scale demonstrations than vanilla object navigation. For example, in Fig 1, we can only get one trajectory for the task "navigating towards infant bed" with an explicit focus on "infant bed" in the object navigation setting. However, we can gather multiple navigation trajectories for PixNav by assigning different pixels belonging to the "infant bed" as the goal. Given that each image comprises thousands of pixels, the potential for generating diverse trajectories is vast. Compared with object navigation, gathering PixNav trajectories doesn't need explicit semantic annotations, which further reduces the cost of data collection. Generating the optimal actions towards the pixel is also simple: Each pixel is an underlying 3-D coordinates and we can use off-the-shelf planning methods to achieve this point-goal navigation for dataset collection.

The remaining question of PixNav is how to navigate towards the out-of-sight target. For indoor navigation scenarios, the house's layout and object placements typically adhere to human preferences, and such commonsense priors have already been embedded into large-language models (LLMs). Therefore, we designed a hierarchical policy that employs LLM as the planner to guide the exploration sequences across rooms. To make the planning consistent and rational, we designed a step-by-step prompt template to evoke the reasoning ability of LLM. With our designed prompt, the LLM demonstrates its potential in self-localization and structural memory organization. Our hierarchical policy achieves competitive performance compared with the mapping-based methods in zero-shot object navigation without depth modality as input. In summary, our contributions are listed below:

- **Pixel Navigation** - We propose a universal RGB-based navigation policy to replace the function of the path-planning method in map-based navigation.
- **Chain of Thought of LLM for Navigation** - We design a step-by-step prompt template and discover the potential of LLM in self-localization and structural memory organization.
- **Foundation Model Powered Navigation** - Our approach bridges the gap between the embodied navigation tasks and the foundation models. Our hierarchical policy achieves competitive performance compared with the mapping-based methods.

II. RELATED WORK

A. Open-Vocabulary Object Navigation

Searching for an object is a common task in indoor environments. While there is a rich body of work dedicated to the object navigation task [13], [16], [17], [18], [12], [19], most are limited in a close-set of objects. In response, many researchers study the more realistic open-vocabulary

object navigation task. ZSON [20] transfers an image-goal navigation policy for object-goal navigation tasks with the text-image alignment from the CLIP [21] encoder. And SSNet [22] explicitly introduces an encoding network to extract the feature for the target object tags. Without a large-scale object navigation dataset, the end-to-end policy network suffers in learning the multi-modal alignment for such embodied task and achieved limit performance. Therefore, many map-based approaches [7], [23], [24], [9] are proposed where a high-level policy selects a promising coordinate for object searching and the low-level policy complete the point-goal navigation with path-planning. Yet, the map-based approaches are highly dependent on the accuracy of depth-sensor and self-localization. On the contrary, we propose a fundamental pixel navigation policy that can replace the traditional path-planning method and make the approach become a pure RGB-based solution.

B. Foundation Models on Vision and Language

Recent advancements have seen significant strides in foundational models across both visual and language domains. For example, GroundingDINO [1] and Owl-ViT [2] can accomplish zero-shot detection, RAM [25] can tag almost all the appeared objects in the image. SAM [3] pioneers in delivering fine-granular segmentation for a diverse set of objects. As for language, LLM shows surprising performance on almost all the tasks, including text understanding, generation [26], [27] and even planning for robotics field [28], [29], [30], [31]. However, there are many tasks that require both vision and language inputs. To this end, many works introduce the large-language model as the backbone and develop vision-language models for general-purpose multi-modal tasks [32], [6], [5]. Within the robotics arena, PALM-E [33] is a specialized foundation model that grounds the images and language into several embodied tasks. In our work, we utilize the LLama-Adapter [5] to convert images into text and employ GPT-4 to facilitate reasonable textual planning. Furthermore, leverage the robust zero-shot recognition capabilities of foundational models to enhance performance.

C. Large-Language Model for Visual Navigation

For many visual navigation tasks, the agent is required to achieve efficient exploration by linking the goals and the observed semantic clues. Leveraging LLM's powerful commonsense reasoning ability [27], [33], [34], many works start to explore the ways to introduce the LLM's planning capability into visual navigation tasks. For instance, L3MVN [9] transcribes surrounding information of map boundaries into text, enabling the LLM to evaluate the scores of the frontiers for better exploration. ESC [8] further feeds both room-level and object-level information to LLM to select an exploration target. Similarly, LGX [10] also captions the image observation into text but uses LLM to make a search plan for an arbitrarily described object. StructNav [11] expands the scene memory with point clouds, scene graphs, and semantic occupancy maps, making more reasonable planning with LLM for object navigation. NavGPT [35]

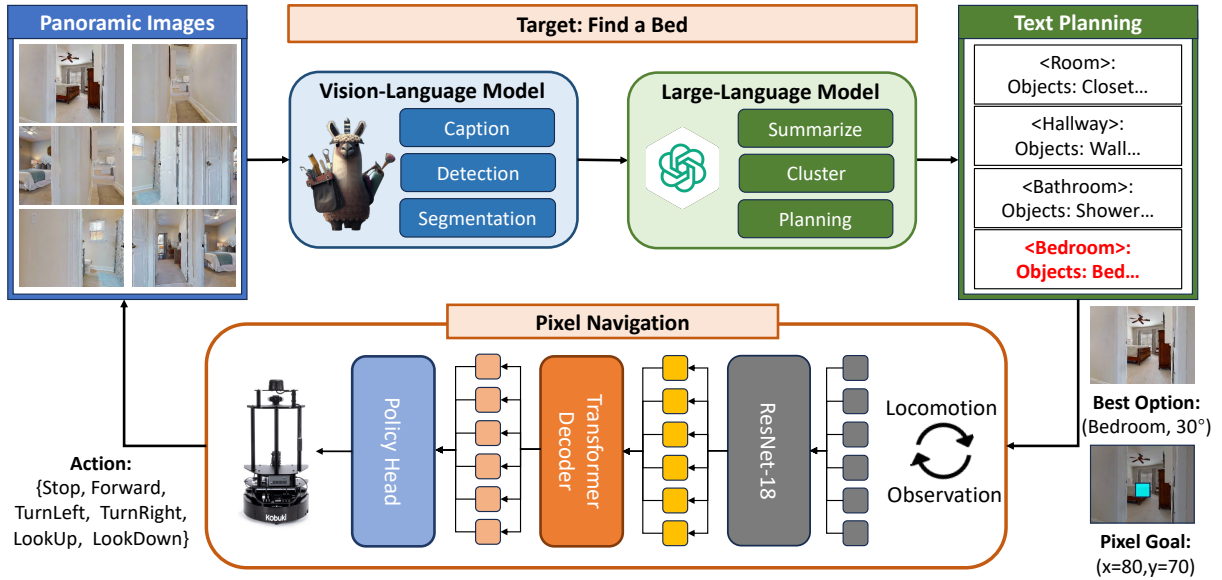


Fig. 2. The pipeline of our RGB-centric strategy for zero-shot object navigation. In each cycle, the agent first gets the panoramic images of the surroundings. This can be finished by controlling the robot by turning around or installing multiple RGB cameras. And a vision-language model translates this visual data into a textual description. Utilizing a systematically crafted step-by-step prompt template, the LLM then strategizes the most optimal next step for the target location. Then, the target location will be indicated as a pixel and conveyed to the navigation policy. The PixNav policy will continuously receive observation images and perform actions until arriving at the goal area.

discovers a way to directly use the planning and reasoning capability of LLMs on the instruction-following navigation agent with prompt engineering. We follow the idea to introduce the LLM to promote the efficiency of an object navigation agent, but differ from those studies, we design a step-by-step prompt template and discover the potential of LLM in self-localization and structural memory organization, which can be further extended in other navigation tasks. However, our unique contribution lies in a crafted step-by-step prompt template, which could guide LLMs to pinpoint the agent’s location and adeptly organize memory. And we believe these findings have potential implications for a broader range of navigation tasks.

III. METHOD

A. Path-Planning with Pixel Navigation

In pursuit of a universally applicable navigation strategy, we introduce the pixel-guided navigation skill (PixNav) tailored for zero-shot object navigation tasks. The PixNav uses a pixel in the initial frame as the goal and executes primitive actions navigating towards the corresponding area indicated by the pixel. Our approach leverages an action space composed of six primitive actions: {Stop, MoveAhead, TurnLeft, TurnRight, LookUp, LookDown}. Notably, while our model currently operates in this discrete action space, PixNav has the potential to be extended to continuous settings. To help the model deal with the partial observability and tracking the pixel in the first frame, we use a transformer-based network with all history images $I_{0 \sim T} = \{I_0, I_1 \dots I_T\}$ and the goal pixel $p_g = (x_i, y_i)$ as inputs. All the parameters are optimized from scratch and the architecture is summarized at the bottom part in Fig 2. **Training Data.** The large-scale dataset is essential for the robot learning [14]. Regarding the demonstration data collection, our PixNav skill offers marked advantages over

the object-goal navigation approaches. The latter demands a unique object instance for each navigation trajectory, complicating the process through the necessity for object instance annotation prior to generating each trajectory. Such object instances are generally sparse within a single image which limits the potential collected training trajectories. Conversely, PixNav directly utilizes the vast potential of the pixel space, yielding much diverse navigation trajectories. The integration of the depth sensor further eases this approach, by converting each pixel to a 3D coordinate. Then, off-the-shelf path planners can be utilized, which is especially straightforward in simulation environments. In our work, we use the Habitat-Lab [36], [37] with HM3D [38] dataset to generate the training trajectories for PixNav. At the start of each episode, we randomly sample a pixel from the RGB observation and use the underlying path planner to interpolate the actions. In total, we generate a dataset with more than 100,000 trajectories to train the policy and we constrain the trajectory length within 64-time steps. The size of RGB observation in the trajectories is 160x120.

Architecture Details. To encode all the history images and the target pixel, we first convert the pixel goal $p_g = (x_i, y_i)$ into a mask image M_g where the values within the bounding box $(x_i - \delta, y_i - \delta, x_i + \delta, y_i + \delta)$ are set to 1 and the rest are set to 0. Here, δ is a parameter used to control the masked area and in our work, we set it to 2. We then concatenate the mask image with the initial image in a channel-wise manner and employ a 4-channel ResNet18 to process the combined input $v_g = (I_0, M_g)$, yielding a 768-dimensional goal token $\Phi_\theta(v_g)$. For the RGB observation images, we use another 3-channel ResNet18 to obtain a sequence of 512-dimensional observation tokens $[\varphi_\theta(I_0), \varphi_\theta(I_1) \dots \varphi_\theta(I_T)]$. Both ResNet18 models are initialized randomly. To ensure the information of the goal pixel can be utilized by the transformer, we introduce

a goal-fusion layer to generate a 256-dimensional feature from $\Phi_\theta(v_g)$ and concatenate it with each observation tokens. Subsequently, a 4-layer transformer decoder is employed to integrate the global information across all tokens. Denote the output tokens as $[s_0, s_1 \dots s_T]$. These output tokens are fed into a policy layer to predict the optimal action $\pi_\theta(s_i)$. As the robot performs the navigation task, the changes in viewpoint will cause the movement of the pixel goal in the subsequent frames. To help the model understand such transformation, we introduce a tracking layer to predict the corresponding pixel in every RGB observation $f_\theta(s_i)$. Moreover, to aid the model in understanding the 3-D structure, we introduce a temporal distance layer to predict the remaining time steps $d_\theta(s_i)$ to each the pixel goal. All the policy layer, tracking layer, and temporal distance layer are 1-layer MLP built upon output tokens from transformer decoder $[s_0, s_1 \dots s_T]$. The parameters of the entire network are 54.5M.

Training Objective. To train the entire network, we incorporate three objectives. The first is the imitation learning objective for $\pi_\theta(s_i)$, the second is MSE loss for the corresponding pixel prediction $f_\theta(s_i)$, and the third is MSE loss for temporal distance prediction $d_\theta(s_i)$. Because the temporal distance can vary from 1 to 64, we scale the loss function for temporal distance by dividing it by 10. In summary, the objectives can be formulated as follows:

$$L_{il} = \mathbb{E}_\tau[\hat{a}_t \log \pi_\theta(s_t)] \quad (1)$$

$$L_{distance} = \mathbb{E}_\tau[(\hat{d}_t - d_\theta(s_i))^2] \quad (2)$$

$$L_{track} = \mathbb{E}_\tau[(\hat{p}_t - f_\theta(s_i))^2] \quad (3)$$

The $\hat{a}_t, \hat{d}_t, \hat{p}_t$ are the ground truth for best actions, temporal distance, and the corresponding pixel. Therefore, all the training processes are supervised learning. The entire loss is written as follows:

$$L = L_{il} + L_{distance} + L_{track} \quad (4)$$

B. Perception with Foundation Model

Despite the commonsense inference and reasoning capabilities contained in LLMs, navigation tasks generally operate with image-based observations. To bridge this modality gap, we employ the vision-language model, LLama-Adapter [5], as a conduit for image summarization. Specifically, we prompt the VLM with two structured queries: "Describe the room type in the image" and "Describe the objects with details in the image". Then, each image can be converted into a structured textual representation, as demonstrated in Fig 3. To make a comprehensive planning, a panoramic understanding is essential. To this end, the robot undergoes a series of rotations and captures six distinct images which are converted into six paragraphs as the input for LLM later. In order to allocate one pixel as the goal, we introduce two additional foundation models, namely GroundingDINO [1] and SAM [3] to capture the find-grained segmentation masks. Once an object target is proposed by LLM, the center point of the object mask will be assigned as the pixel goal for the low-level navigation policy.




Fig. 3. An example of the VLM translation. The VLM is prompted with room-level and object-level queries. The robot performs a series of rotations, capturing six distinct images to ensure a comprehensive panoramic view.

C. Planning with Large-Language Model

For the long-horizon object navigation task, we design an LLM-based planner to decide which direction is most promising for searching the target object. For example, a bed is usually placed in a bedroom, and the agent needs to traverse a corridor to reach the bedroom. To achieve this, we design step-to-step prompts to help the LLM understand an indoor navigation scenario. The entire planning procedure is divided into three stages: Summarize, Clustering, and Planning, as shown in Fig 4. Firstly, as prior works show structured text is beneficial for LLM’s planning and reasoning [30], [29], we ask the LLM to summarize the VLM captions into JSON format which contains {‘ID’, ‘Angle’, ‘Type’, ‘Objects’}. Each JSON data encapsulates room-level information. However, the six surrounding images do not generally represent six unique room instances and several images can be grouped under a single room. To cluster the data, we use LLM to estimate the robot’s current room and its position, and then build the inter-room connections. The same type of rooms with an edge connection will be merged into one. We also cluster the same type of rooms with adjacent angles. Then, a graph of the local layout in the house is generated by the LLM as shown at the top of Fig 4. With this graph, the LLM can perform consistent exploration behaviors among rooms instead of roaming in one room. We design a planning prompt for the LLM to balance between exploration and exploitation for searching the target object. Our planner proposes the subgoal after periodic time steps or the low-level policy completes navigating towards the assigned goal.

IV. EXPERIMENTAL RESULTS

A. Benchmarks and Metrics

We employed the large-scale indoor dataset HM3D [38] to evaluate our proposed method. The dataset contains 80 scenes for training and 20 scenes for validation. For the collection of PixNav demonstration trajectories, we utilized the 80 training scenes with the camera height 0.88m and horizontal field of view (HFOV) 79°. The evaluation for object navigation follows the settings in objectnav-challenge-2022 [39]. For the object navigation task, we report the performance with respect to Success and SPL, which is a metric evaluating the navigation path efficiency. To evaluate the performance of PixNav, we randomly initialized the agent in 20 validation scenes and sampled a pixel goal in the first RGB frame. The PixNav evaluation is divided into two levels, one is 1-3m, and

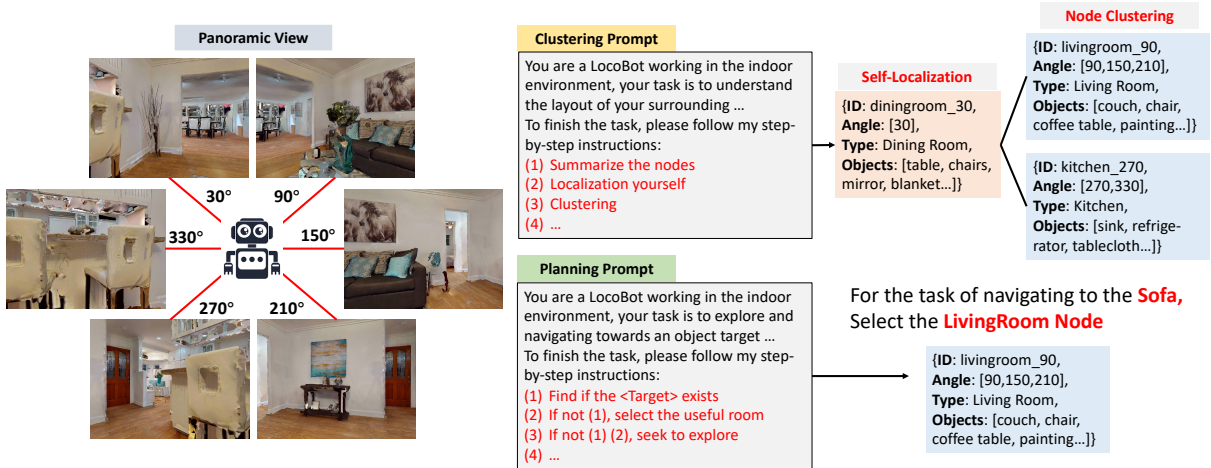


Fig. 4. The step-by-step prompt template for LLM to make a reasonable plan for navigation. After translating the panoramic images into text, we first use the LLM to summarize the captions into highly structured data and ask LLM to estimate the robot’s location according to the commonsense of indoor environments. With the robot location, we can cluster the data into room-level instances and thus the layout can be described as a graph. Based on this graph, we then ask the LLM to provide a plan considering both exploitation and exploration.

the other is 3-5m, which represents the underlying geodesic distance from the assigned pixel goals to the initial robot location. We further report the distance to goal (DTG), the geodesic distance from the location at the end of the trajectory to the pixel goal. For each level, we sample 500 episodes for evaluation.

B. Baselines

We consider both non-zero-shot object navigation methods and zero-shot object navigation methods to verify our idea.

- **Habitat-Web** [16]. A closed-set object navigation baseline with human demonstrations for imitation learning.
- **OVRL** [18]. A closed-set object navigation baseline with the large-scale indoor dataset for representation learning.
- **ESC** [8] A map-based zero-shot object navigation baseline with LLM for frontier planning.
- **ZSON** [20] An RGB-based zero-shot object navigation baseline with CLIP [21] as the image and object goal encoders.

C. Performance on Zero-Shot ObjectNav

We compare our method with baseline results in Tab I and example visualization of trajectories are shown in Fig 5. Our approach achieves competitive performance even without depth and localization as inputs. Our PixNav can consistently track the initial pixel target and adaptively look down to avoid collision. The long-horizon trajectories show heuristic exploration behavior. For example, in the third figure, the agent can first go out of one room and enter another, then walk back to the corridor to search for a bathroom. It proves the effectiveness of our LLM-based planner. How to properly select the pixel as the goal for the PixNav is essential. We found that assigning the center pixel of the selected room ‘floor’ is a good strategy.

D. Ablation Study on PixNav

For real-world applications, robots may have different embodiment and camera settings. To become a general RGB-based path-planning module, the PixNav policy should be

TABLE I
PERFORMANCE ON OBJECT NAVIGATION TASKS WITH HM3D DATASET.

| Setting | Mapless | Method | Success | SPL |
|----------------------------|---------|-------------|-------------|-------------|
| <i>Close-Set ObjectNav</i> | ✓ | Habitat-Web | 41.5 | 16.0 |
| | ✓ | OVRL | 62.0 | 26.8 |
| <i>Open-Set ObjectNav</i> | ✓ | ZSON | 25.5 | 12.6 |
| | ✗ | ESC | 39.2 | 22.3 |
| | ✓ | Ours | 37.9 | 20.5 |

able to transfer to various camera configurations. Therefore, we adjusted the camera settings in the simulator with different heights and HFOV and re-evaluated the PixNav performance shown in Table II. We found that even if we introduce a severe change in camera height (50% relative change) and HFOV (25% relative change), our policy can still achieve a high success rate in 1m-3m navigation tasks and degenerate little in most of the cases. This indicates that our PixNav can achieve zero-shot generalization to new camera settings and learn a universal representation for navigation. But what makes our policy robust? We conduct ablation studies on three components of our method, which are tracking prediction, temporal distance prediction, and goal fusion network. The performance is shown in Tab III. Most of the components have a minor influence on the 1-3m scenarios, which shows the superiority of the transformer-based policy network architecture. However, tracking prediction and goal fusion significantly influence the 3-5m scenarios, emphasizing the need to explicitly remind the network not to overlook the inconspicuous pixel goal in order to achieve long-term navigation with PixNav.

E. Analysis of LLM Planner

To achieve long-horizon navigation tasks, the LLM planner is also an essential part. Traditional map-based approaches prove that mapping and localization are valuable skills for navigation tasks. If the LLM also possesses these skills, the potential of our LLM-based RGB policy is promising



Fig. 5. Trajectory visualization of both PixNav policy and long-horizon object navigation. The top part represents the first-person view of PixNav trajectory. Even though the robot’s view changes constantly, our model can still predict the original goal and walk towards it. The bottom part represents the top-down view of the entire trajectory. Our LLM-based planner helps the robot achieve consistent and semantic exploration towards the target.

TABLE II

PERFORMANCE ON DIFFERENT VARIATIONS OF CAMERA SETTINGS

| Camera | 1m-3m | | | 3m-5m | | |
|--------------|---------------|----------------|------------------|---------------|----------------|------------------|
| | SR \uparrow | SPL \uparrow | DTS \downarrow | SR \uparrow | SPL \uparrow | DTS \downarrow |
| (0.88, 79°) | 0.88 | 0.88 | 0.54 | 0.42 | 0.41 | 1.61 |
| (0.48, 79°) | 0.79 | 0.76 | 0.61 | 0.26 | 0.25 | 1.95 |
| (1.28, 79°) | 0.78 | 0.75 | 0.86 | 0.33 | 0.32 | 1.83 |
| (0.88, 60°) | 0.71 | 0.70 | 0.82 | 0.14 | 0.13 | 2.32 |
| (0.88, 100°) | 0.80 | 0.78 | 0.69 | 0.33 | 0.32 | 1.88 |

TABLE III

THE ABLATION STUDY OF THE PIXNAV IN HM3D DATASET

| Ablation | 1m-3m | | | 3m-5m | | |
|--------------|---------------|----------------|------------------|---------------|----------------|------------------|
| | SR \uparrow | SPL \uparrow | DTS \downarrow | SR \uparrow | SPL \uparrow | DTS \downarrow |
| Ours | 0.88 | 0.88 | 0.54 | 0.42 | 0.41 | 1.61 |
| w/o tracking | 0.82 | 0.81 | 0.63 | 0.33 | 0.31 | 1.75 |
| w/o distance | 0.86 | 0.85 | 0.60 | 0.38 | 0.35 | 1.69 |
| w/o fusion | 0.84 | 0.84 | 0.63 | 0.34 | 0.33 | 1.91 |

for navigation field. To probe on the abilities of LLM, we implement a quantitative analytic to verify whether the LLM owns such spatial-awareness of the indoor environments. The experiments contain two skills evaluations: Room Localization and Room Clustering. We ask 3 humans to see the panoramic images of the robot’s location and answer two questions: "Which room is the robot in?" and "Describe the number and type of the rooms in the surroundings." We compare the answers from LLM and humans for 100 episodes. The accuracy of the two skills is shown in Tab IV.

TABLE IV

THE ANALYSIS OF LLM ABILITIES IN LOCALIZATION AND MAPPING

| LLM Ability | Human Eval (%) |
|-------------------|----------------|
| Room Localization | 68.4 |
| Room Clustering | 47.1 |

F. Sim2Real Transfer

The real-world experiments are conducted with an iRobot-Create3 Mobile base equipped with an OAK-D camera. However, we only use the RGB observations from the camera for the experiments. The action configuration of the robot follows the discrete action space setting, but contains 4 actions {Stop, MoveAhead 0.25m, TurnLeft 30°, TurnRight 30°} which is different from the simulation. This is because the camera owns a fixed pose relative to the mobile base. We conduct experiments to verify the function of the PixNav policy. Please refer to our website for the details of real-world experiments.

V. CONCLUSION

In this work, we take a significant stride towards developing a pure RGB-based solution for visual navigation tasks through the utilization of a pixel-guided navigation skill. As we demonstrate PixNav skill can be used for navigating any kinds of objects and exhibits generalization capabilities across different camera settings, using **Pixel** as the goal specification is possible for building a versatile low-level path-planning policy. And we found that even the most advanced large language models (GPT-4) still perform poorly in two navigation related tasks. Therefore, in the future, we will investigate how to further improve the robustness of the PixNav skill and try to grounding the LLM capabilities into navigation planning tasks.

REFERENCES

- [1] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, C. Yue Li, J. Yang, H. Su, J.-J. Zhu, and L. Zhang, "Grounding dino: Marrying dino with grounded pre-training for open-set object detection," *arXiv preprint arXiv:2303.05499*, 2023.
- [2] M. Minderer, A. A. Gritsenko, A. Stone, M. Neumann, D. Weissenborn, A. Dosovitskiy, A. Mahendran, A. Arnab, M. Dehghani, Z. Shen, X. Wang, X. Zhai, T. Kipf, and N. Houlsby, "Simple open-vocabulary object detection with vision transformers," *European Conference on Computer Vision (ECCV)*, 2022.
- [3] 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," *International Conference on Computer Vision (ICCV)*, 2023.
- [4] B. Li, K. Q. Weinberger, S. J. Belongie, V. Koltun, and R. Ranftl, "Language-driven semantic segmentation," *International Conference on Learning Representations (ICLR)*, 2022.
- [5] R. Zhang, J. Han, A. Zhou, X. Hu, S. Yan, P. Lu, H. Li, P. Gao, and Y. Qiao, "Llama-adapter: Efficient fine-tuning of large language models with zero-initialized attention," *International Conference on Learning Representations (ICLR)*, 2024.
- [6] W. Dai, J. Li, D. Li, A. M. H. Tiong, J. Zhao, W. Wang, B. Li, P. Fung, and S. C. H. Hoi, "Instructblip: Towards general-purpose vision-language models with instruction tuning," *Advances In Neural Information Processing Systems (NeurIPS)*, 2023.
- [7] 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," *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [8] K.-Q. Zhou, K. Zheng, C. Pryor, Y. Shen, H. Jin, L. Getoor, and X. E. Wang, "Esc: Exploration with soft commonsense constraints for zero-shot object navigation," in *International Conference on Machine Learning (ICML)*, 2023.
- [9] B. Yu, H. Kasaei, and M. Cao, "L3m3v: Leveraging large language models for visual target navigation," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.
- [10] V. S. Dorbala, J. F. Mullen Jr, and D. Manocha, "Can an embodied agent find your 'cat-shaped mug'?" *llm-based zero-shot object navigation*, *IEEE Robotics and Automation Letters (RA-L)*, 2023.
- [11] J. Chen, G. Li, S. Kumar, B. Ghanem, and F. Yu, "How to not train your dragon: Training-free embodied object goal navigation with semantic frontiers," *Robotics, Science and Systems (RSS)*, 2023.
- [12] A. Khandelwal, L. Weihs, R. Mottaghi, and A. Kembhavi, "Simple but effective: Clip embeddings for embodied ai," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [13] D. S. Chaplot, D. Gandhi, A. Gupta, and R. Salakhutdinov, "Object goal navigation using goal-oriented semantic exploration," in *Advances In Neural Information Processing Systems (NeurIPS)*, 2020.
- [14] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, and T. J. et al., "Rt-1: Robotics transformer for real-world control at scale," *Robotics, Science and Systems (RSS)*, 2023.
- [15] H. Walke, K. Black, A. Lee, M. J. Kim, M. Du, C. Zheng, T. Zhao, P. Hansen-Estruch, Q. H. Vuong, A. W. He, V. Myers, K. Fang, C. Finn, and S. Levine, "Bridgedata v2: A dataset for robot learning at scale," *7th Annual Conference on Robot Learning (CoRL)*, 2023.
- [16] R. Ramrakhya, E. Undersander, D. Batra, and A. Das, "Habitat-web: Learning embodied object-search strategies from human demonstrations at scale," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [17] R. Ramrakhya, D. Batra, E. Wijmans, and A. Das, "Pirlnav: Pre-training with imitation and rl finetuning for objectnav," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [18] K. Yadav, R. Ramrakhya, A. Majumdar, V.-P. Berges, S. Kuhar, D. Batra, A. Baevski, and O. Maksymets, "Offline visual representation learning for embodied navigation," in *Workshop on Reinforcement Learning at ICLR 2023*, 2023.
- [19] M. Zhu, B. Zhao, and T. Kong, "Navigating to objects in unseen environments by distance prediction," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [20] A. Majumdar, G. Aggarwal, B. Devnani, J. Hoffman, and D. Batra, "Zson: Zero-shot object-goal navigation using multimodal goal embeddings," in *Advances In Neural Information Processing Systems (NeurIPS)*, 2022.
- [21] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning transferable visual models from natural language supervision," in *International Conference on Machine Learning (ICML)*, 2021.
- [22] Q. Zhao, L. Zhang, B. He, H. Qiao, and Z. yong Liu, "Zero-shot object goal visual navigation," *IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [23] C. Huang, O. Mees, A. Zeng, and W. Burgard, "Visual language maps for robot navigation," *IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [24] 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," *IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [25] Y. Zhang, X. Huang, J. Ma, Z. Li, Z. Luo, Y. Xie, Y. Qin, T. Luo, Y. Li, S. Liu, Y. Guo, and L. Zhang, "Recognize anything: A strong image tagging model," *arXiv preprint arXiv:2306.03514*, 2023.
- [26] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, and G. Lample, "Llama: Open and efficient foundation language models," *ArXiv*, 2023.
- [27] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. J. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, "Language models are few-shot learners," *Advances In Neural Information Processing Systems (NeurIPS)*, 2020.
- [28] W. Huang, P. Abbeel, D. Pathak, and I. Mordatch, "Language models as zero-shot planners: Extracting actionable knowledge for embodied agents," *International Conference on Machine Learning (ICML)*, 2022.
- [29] I. Singh, V. Blukis, A. Mousavian, A. Goyal, D. Xu, J. Tremblay, D. Fox, J. Thomason, and A. Garg, "Progprompt: Generating situated robot task plans using large language models," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- [30] K. Rana, J. Haviland, S. Garg, J. Abou-Chakra, I. Reid, and N. Sunderhauf, "Sayplan: Grounding large language models using 3d scene graphs for scalable task planning," *7th Annual Conference on Robot Learning (CoRL)*, 2023.
- [31] S. Huang, Z. Jiang, H.-W. Dong, Y. J. Qiao, P. Gao, and H. Li, "Instruct2act: Mapping multi-modality instructions to robotic actions with large language model," *ArXiv*, 2023.
- [32] D. Zhu, J. Chen, X. Shen, X. Li, and M. Elhoseiny, "Minigt-4: Enhancing vision-language understanding with advanced large language models," *International Conference on Representation Learning (ICLR)*, 2024.
- [33] D. Driess, F. Xia, M. S. M. Sajjadi, C. Lynch, A. Chowdhery, B. Ichter, A. Wahid, J. Tompson, Q. H. Vuong, T. Yu, W. Huang, Y. Chebotar, P. Sermanet, D. Duckworth, S. Levine, V. Vanhoucke, K. Hausman, M. Toussaint, K. Greff, A. Zeng, I. Mordatch, and P. R. Florence, "Palme: An embodied multimodal language model," *International Conference on Machine Learning (ICML)*, 2023.
- [34] OpenAI, "Gpt-4 technical report," *ArXiv*, vol. abs/2303.08774, 2023.
- [35] G. Zhou, Y. Hong, and Q. Wu, "Navgpt: Explicit reasoning in vision-and-language navigation with large language models," *ArXiv*, 2023.
- [36] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik, D. Parikh, and D. Batra, "Habitat: A Platform for Embodied AI Research," in *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [37] A. Szot, A. Clegg, E. Undersander, E. Wijmans, Y. Zhao, J. Turner, N. Maestre, M. Mukadam, D. Chaplot, O. Maksymets, K. Gokaslan, V. Vondrus, S. Dharur, F. Meier, W. Galuba, A. Chang, Z. Kira, V. Koltun, J. Malik, M. Savva, and D. Batra, "Habitat 2.0: Training home assistants to rearrange their habitat," in *Advances In Neural Information Processing Systems (NeurIPS)*, 2021.
- [38] S. K. Ramakrishnan, A. Gokaslan, E. Wijmans, O. Maksymets, A. Clegg, J. Turner, E. Undersander, W. Galuba, A. Westbury, A. X. Chang, M. Savva, Y. Zhao, and D. Batra, "Habitat-matterport 3d dataset (hm3d): 1000 large-scale 3d environments for embodied ai," *NeurIPS Datasets and Benchmarks*, 2021.
- [39] K. Yadav, J. Krantz, R. Ramrakhya, S. K. Ramakrishnan, J. Yang, A. Wang, J. Turner, A. Gokaslan, V.-P. Berges, R. Mootaghi, O. Maksymets, A. X. Chang, M. Savva, A. Clegg, D. S. Chaplot, and D. Batra, <https://aihabitat.org/challenge/2022/>, 2022.