

Prioritized Semantic Learning for Zero-shot Instance Navigation

Xinyu Sun^{1*}, Lizhao Liu^{3*}, Hongyan Zhi, Ronghe Qiu¹, and Junwei Liang^{1,2†} 

¹ AI Thrust, The Hong Kong University of Science and Technology (Guangzhou)

² Department of Computer Science and Engineering, The Hong Kong University of Science and Technology

³ Tencent AI Lab, Shenzhen, China

{csxinyusun, selizhaoliu}@gmail.com, junweiliang@hkust-gz.edu.cn

Abstract. We study zero-shot instance navigation, in which the agent navigates to a specific object without using object annotations for training. Previous object navigation approaches apply the image-goal navigation (**ImageNav**) task (go to the location of an image) for pretraining, and transfer the agent to achieve object goals using a vision-language model. However, these approaches lead to issues of semantic neglect, where the model fails to learn meaningful semantic alignments. In this paper, we propose a Prioritized Semantic Learning (PSL) method to improve the semantic understanding ability of navigation agents. Specifically, a semantic-enhanced PSL agent is proposed and a prioritized semantic training strategy is introduced to select goal images that exhibit clear semantic supervision and relax the reward function from strict exact view matching. At inference time, a semantic expansion inference scheme is designed to preserve the same granularity level of the goal-semantic as training. Furthermore, *for the popular HM3D environment*, we present an Instance Navigation (**InstanceNav**) task that requires going to a specific object instance with detailed descriptions, as opposed to the Object Navigation (**ObjectNav**) task where the goal is defined merely by the object category. Our PSL agent outperforms the previous state-of-the-art by 66% on zero-shot **ObjectNav** in terms of success rate and is also superior on the new **InstanceNav** task. Code will be released at <https://github.com/XinyuSun/PSL-InstanceNav>.

Keywords: Zero-Shot Instance Navigation · Zero-Shot Object Navigation · Image Goal Navigation · Prioritized Semantic Learning

1 Introduction

Visual navigation is a fundamental skill for embodied agents to travel efficiently in complex 3D environments [66]. Among different kinds of visual navigation tasks, Zero-shot Object Navigation (ZSON) is a promising task and paves the

* Equal contribution.

† Corresponding author.

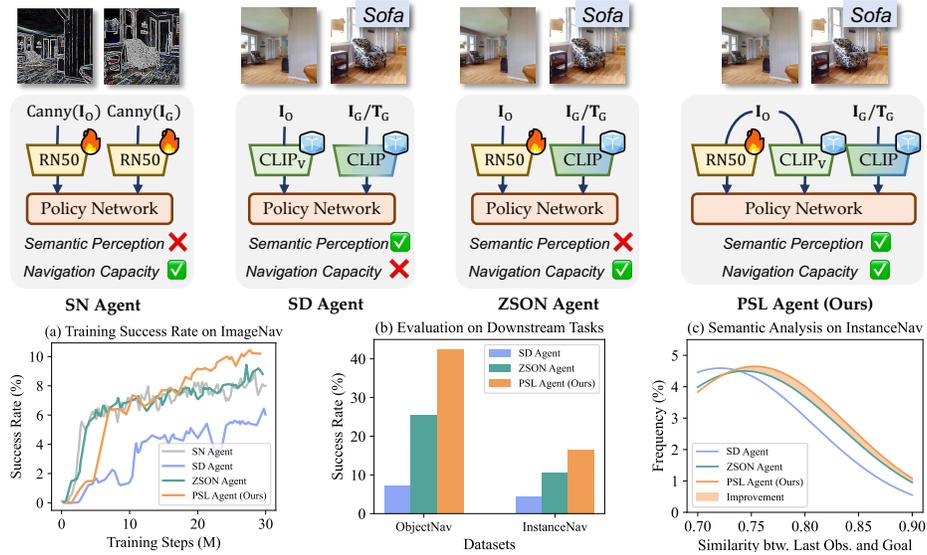


Fig. 1: Pilot studies on four navigation agents: Semantic-Non-dominant (SN) agent, Semantic-Dominant (SD) agent, ZSON [38] agent, and our PSL agent. (a) The agent trained on **ImageNav** task does not necessarily need to learn the semantic information to obtain a high success rate. (b)-(c) Our PSL agent achieves both strong semantic perception and navigation capacity.⁴

way for developing a general embodied agent [2, 12, 13, 25, 28, 48, 50, 55, 61], since it requires zero scene object annotations for training. However, the criteria of the current ZSON task is to evaluate whether the agent can go to the object that matches solely the given category, which remains a considerable gap from *real-life applications* that require the identification of a specific object. Therefore, we propose to extend ZSON to Zero-shot Instance Navigation (ZSIN) task in the HM3D environment [43, 59], which requires the agent to go to a unique object instance given detailed text descriptions. Both the ZSON and ZSIN tasks are very challenging due to the zero-shot constraint, in which the category information of goal objects is not feasible during training.

To achieve zero-shot visual navigation, pioneering works [1, 38] turn to pre-trained vision-language models such as CLIP [42] for its zero-shot capability and propose to train the agent on Image-goal Navigation (**ImageNav**) pretext task that requires the agent to go to a randomly sampled goal image [66]. To be specific, using the CLIP vision encoder to obtain a semantic goal from the goal image for training, the ZSON agent [38] can be transferred to **ObjectNav** by switching the semantic goal into the CLIP text embedding. Although the ZSON agent demonstrates solid navigation capabilities, we argue that it lacks strong semantic perception abilities.

⁴ The SN agent does not support **InstanceNav** task as it lacks a semantic goal encoder.

To investigate how the agent architecture affects its semantic perception ability and navigation capacity, we conduct pilot studies on four different agent architectures and the results are presented in Fig. 1. We intriguingly find a Semantic-Non-dominant (SN) agent, with Canny operator⁵ and learnable ResNet50 [16] encoders to obtain observation and goal embeddings, is able to achieve a competitive success rate on the `ImageNav` task as the ZSON agent (see Fig. 1 (a)). In contrast, the Semantic-Dominant (SD) agent with two fixed CLIP vision encoders gets the worst results. Therefore, we conclude that *the ImageNav pre-training task does not necessarily require the agent to perceive semantic information*. In this sense, the ZSON agent with only a learnable observation encoder may possess inferior semantic perception ability, limiting the performance on navigation tasks that heavily rely on finding semantic clues. Therefore, how to effectively improve the semantic perception ability of the navigation agent in the zero-shot setting remains an important but unresolved issue.

In this paper, we propose a Prioritized Semantic Learning (PSL) method to improve the semantic perception and understanding ability of the navigation agent. Our PSL approach is comprised of three parts: a PSL agent architecture, a prioritized semantic training strategy, and a semantic expansion inference scheme. **First**, we equip the agent with an additional CLIP vision encoder to encode the semantics in the observation. Subsequently, a Semantic Perception Module (SPM) is proposed to comprehend the semantic differences between the observation and goal images. **Second**, we take advantage of the entropy-minimization technique to select goal images with clear semantic supervision. Moreover, we relax the reward function to focus more on semantic correspondence rather than strict geometric matching. **Third**, considering that the proposed SPM learns to comprehend the semantic differences of image embeddings during training, we develop a semantic expansion inference scheme that retrieves image embeddings using text queries. The retrieved image embeddings expand the text embeddings with rich visual priors, facilitating more precise indications. Without bells and whistles, the proposed PSL method achieves state-of-the-art performance on both widely-evaluated `ObjectNav` and our extended `InstanceNav` benchmarks. For example, in the `ObjectNav` task, PSL outperforms the ZSON baseline [38] and the ESC method [64] that uses large-language models by **16.9%** and **3.2%** success rate, respectively. Furthermore, we verify the proposed PSL agent has a stronger semantic perception ability by visualizing the semantic similarities between the agent’s last observation and the goal images in Fig. 1 (c). Our main contributions are summarized as follows:

- We investigate the semantic perception ability of different navigation agent architectures and reveal that the commonly-used `ImageNav` pre-training task does not necessarily require the agent to learn semantic information.
- We propose a prioritized semantic learning method that includes a strong agent architecture, a novel training strategy and an effective inference scheme to improve the agents’ semantic perception ability in zero-shot object/instance navigation tasks.

⁵ We consider that the Canny operator has undermined the semantics in the image.

- Extensive experiments and sufficient ablation studies on both **ObjectNav** and **InstanceNav** benchmarks demonstrate the superior performance of the proposed PSL method over the state-of-the-art.

2 Related Work

2.1 Visual Navigation

Visual navigation is a fundamental and crucial task for robots to perform various embodied tasks [4, 26, 27, 47, 53]. In visual navigation tasks, the agent receives a geometry [45, 54] or semantic [5, 66] goal and an RGB observation from the camera attached to the agent, and then makes efforts to navigate to this goal according to the observation in each time step. These tasks can be divided into different categories, including point-goal navigation [45, 54], object-goal navigation [6, 38, 44, 57, 58, 64], image-goal navigation [1, 10, 15, 18, 24, 40, 56–58, 66], and vision-language navigation [3, 8, 20, 22, 23], *etc.* Taking image-goal navigation [56] as an example, the agent needs to explore the environment, locate the target object specified by the image, and navigate to its proximity. A common trail for these visual navigation tasks is to train the agent in the simulator [45, 49] with photo-realistic environments [7, 43, 59] and physical engine, and provide the agent with the goal indication as well as its corresponding trajectory annotation. Since labeling objects in the 3D scenes is labor expensive, recent works [1, 6, 38, 64] have explored how to design an object-goal navigation agent without human annotations. These methods are easy to scale up [38] and show potential to solve the open set challenge [37] in locating objects. In this work, we focus on a more challenging task of visual navigation, zero-shot instance navigation, in which the agent is required to navigate to an object specified by detailed descriptions, for example, an instance image [21] or text.

2.2 Vision-Language Model

Vision-language models (VLM) provide mutual understanding for both vision and language modalities [51] and thereby naturally serve as a fundamental bedrock for vision navigation. To build a visual navigation agent, two lines of work are widely employed, namely, representation-based and module-based methods. Representation-based methods [17, 30, 42] aim to co-embed the image and text into a common semantic space. CLIP [42] and ALIGN [17] apply contrastive loss to train the image and text encoders on large-scale pair-wise image text datasets. ALBEF [30] further proposes to leverage a momentum distillation mechanism that alleviates the weak correlation between noisy image text pairs to improve multimodal semantic learning. As for the module-based methods [14, 29, 31, 34, 35, 63], they aim to learn a strong multimodal reasoner to simultaneously comprehend vision-language modalities and produce text responses, which are commonly used as the navigation controller [6, 64]. Early works such as VisualBERT [31] apply mask language loss on both image patches

and word embeddings to pre-train a multimodal reasoning model. To incorporate the strong reasoning ability of large language models, recent works like BLIP-2 [29], LLaMA-Adapter [14, 63] and LLaVA [34, 35] perform visual instruction tuning on pretrained large language models. However, using these module-based vision language models is compute-intensive, making them unsuitable for real-life robotic scenarios that require nearly real-time response. In this paper, we construct the navigation agent using the representation-based VLM and aim to improve the semantic perception ability of the agent.

3 Motivation of Prioritized Semantic Learning

`ImageNav` is an important pre-training task to achieve Zero-Shot Instance Navigation (ZSIN) as it provides supervision signals to navigate in 3D environments with visual observation and a semantic indication [38]. However, we argue that the `ImageNav` task can be solved without semantic clues, resulting in a sub-optimal pre-training objective. To verify this, we conduct a preliminary experiment on the HM3D `ImageNav` dataset that is used to train a zero-shot object navigation agent in previous literature [38]. We set up four different agents, namely the Layout-Only (LO) agent, Semantic-Only (SO) agent, ZSON agent, and our PSL agent. Among them, the Semantic-Only agent is equipped with two frozen CLIP ResNet-50 encoders to obtain semantic-level observations and goals. The Layout-Only agent is adopted with two trainable ResNet-50 encoders and we remove the semantic information in both observation and goal images by using a Sobel operator to calculate gradient images, which leave only layout information like contours and edges (see Fig. 1 for more details). Results in Fig. 1 (a) show that the agent does not necessarily need to learn semantic information to arrive at an image goal, merely relying on the layout information to perform view matching can obtain a high success rate.

4 Prioritized Semantic Learning

4.1 Preliminaries

We follow ZSON [38] to adopt image-goal navigation pre-training and perform the zero-shot evaluation on downstream tasks. During pre-training, the image-goal navigation (`ImageNav`) objective provides supervision signals to update the agent. Specifically, for each episode with a randomly sampled start point and goal point, the agent is asked to navigate to the goal destination according to a goal image. Empowered by the cross-modal alignment ability of Vision-and-Language models (*i.e.*, CLIP [42]), we can evaluate the trained agent on the object-goal navigation (`ObjectNav`) dataset by replacing the image embeddings with text embeddings.

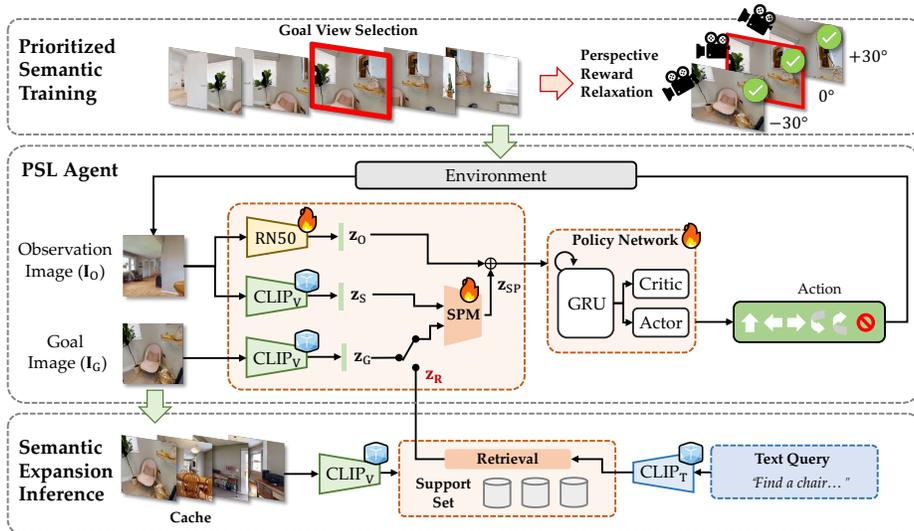


Fig. 2: Overview of our PSL. 1) During ImageNav pre-training, we provide clear semantic supervision signals with the Prioritized Semantic Training strategy. 2) Our PSL agent exploits the Semantic Perception Module (SPM) to achieve both strong semantic understanding and navigation capacity. 3) During inference, a Semantic Expansion Inference scheme is incorporated to ensure the same semantic granularity of the goal-embedding between training and testing.

4.2 Prioritized Semantic Learning Agent Architecture

We propose a Prioritized Semantic Learning (PSL) agent that emphasizes the reasoning of the semantics differences in the goal and observation. The key components of PSL agent include observation and goal encoders, a semantic perception module and a recurrent policy network.

Observation and Goal Encoders. Following the ZSON baseline, given the observation image I_O and the goal image I_G , we apply a learnable ResNet50 encoder and a fixed CLIP encoder to obtain the observation embedding z_O and the goal embedding z_G , respectively. Since with only a trainable ResNet50 as the observation encoder may not effectively learn the semantic information, we additionally leverage a frozen CLIP encoder to extract semantic-level observation z_S . Note that the ResNet50 are initialized from the one pre-trained by OVRL [58].

Semantic Perception Module. We then introduce the semantic perception module, an MLP layer that encodes the semantic correspondence between z_G and z_S into a low-dimension feature. Specifically, it takes both $z_G \in \mathbb{R}^{C_1}$ and $z_S \in \mathbb{R}^{C_1}$ as input and produce semantic perception embedding $z_{SP} \in \mathbb{R}^{C_2}$, where $C_2 < 2 \times C_1$. The semantic perception module reduces the feature dimension, serving as a bottleneck to condense the critical semantic correspondence in both the goal image and observation image.



Fig. 3: Illustration of the perspective relaxation and goal view selection approach in our training strategy. The goal view image set is expanded with additional yaw and tilt views using perspective relaxation. Top-10 views with minimum entropy are circled in red with their ranking.

Policy Network. Based on the semantic perception embedding \mathbf{z}_{SP} and the observation embedding \mathbf{z}_O , we train a navigation policy π_θ using reinforcement learning (RL):

$$\mathbf{s}_t, \mathbf{h}_t = \pi_\theta(\mathbf{z}_{SP} \oplus \mathbf{z}_O \oplus \mathbf{a}_{t-1} | \mathbf{h}_{t-1}) \quad (1)$$

where \mathbf{s}_t is the embedding of the agent’s current state, which is then used to predict an action distribution among six action categories, including MOVE_FORWARD, TURN_LEFT, TURN_RIGHT, STOP, LOOK_UP and LOOK_DOWN. \mathbf{h}_{t-1} denotes hidden state of the recurrent layers in policy π_θ from the previous step. Following previous methods [1, 38], we adopt an actor-critic network to predict state value c_t and action a_t using s_t and train it end-to-end using PPO [46]. More details can be found in the Appendix.

4.3 Prioritized Semantic Training Strategy

Previous literature [19, 21] has pointed out that a large portion of image goals in the ImageNav task lead to ambiguity (e.g. looking at walls). We provide statistical analysis on the distribution of ImageNav goal images to support this speculation (see Fig. II in Appendix). In particular, We select 6 object categories that widely exist in HM3D scenes, as well as two additional categories “wall” and “room” to represent images with ambiguous semantic content. Then we use CLIP to perform zero-shot classification on all goal images by selecting the category with maximum image-text similarity. As presented in previous literature, in the original HM3D ImageNav dataset, most of the goal images cannot be categorized as common object categories; instead, they are classified as meaningless indications (i.e., “wall” and “room”). We believe that the unreasonable distribution of goal images aggravates the semantic neglect issue.

Entropy-minimized Goal View Selection. To alleviate the semantic neglect issue, we propose an entropy-based goal view selection method to ensure a dominant object exists in the goal images. Specifically, we randomly select a base orientation at the goal point and rotate the agent for Ω times uniformly to render

images in different view angles. Then we obtain image embeddings $\{\mathbf{v}_\omega \mid \omega \in \Omega\}$ using CLIP [42] vision encoder and calculate the similarity between image embeddings with all text embeddings of the candidate object categories. We follow the standard setting of Object-Goal navigation approaches [57, 58, 64] to use 6 common objects \mathcal{C} in the 3D simulator and compute query vectors $\{\mathbf{q}_c \mid c \in \mathcal{C}\}$ using CLIP text embedding. The objective is to select a goal view ω^* that minimizes the entropy of probability distribution among all object categories:

$$\omega^* = \arg \min_{\omega \in \Omega} - \frac{1}{\log(|\mathcal{C}|)} \sum_{c \in \mathcal{C}} \mathbf{p}_c \log \mathbf{p}_c, \quad (2)$$

where $\mathbf{p}_c = \text{softmax}(g(\mathbf{v}_\omega, \mathbf{q}_c))$ is the class probability w.r.t. c^{th} object category, $g(\mathbf{a}, \mathbf{b}) = \tau \cdot \frac{\mathbf{a}^T \mathbf{b}}{\|\mathbf{a}\|_2 \|\mathbf{b}\|_2}$ is the scaled cosine similarity function and τ is the temperature. As shown in Fig. II in the Appendix, the goal images with ambiguous semantics are significantly reduced. Additionally, image distribution among all object categories becomes more balanced.

Perspective Reward Relaxation. Choosing goal views with rich semantic content enables the agent to focus on meaningful regions in the goal images. However, the ImageNav reward [1, 38] still motivates the agent to strive for an exact match between the goal image and the observation image. We tackle this challenge by introducing a perspective reward relaxation technique that cancels exact matching. Specifically, we augment the goal view selection process by rendering additional circles of images around the agent at different pitch angles. As illustrated in Fig. 3, we then perform goal view selection among all rendered views that vary in pitch and yaw angles. Furthermore, we rewrite the **reward function** of PPO training by relaxing the perspective matching requirement:

$$\begin{aligned} R_t^{\text{PSL}} = & \underbrace{\gamma^{suc} * \mathbb{1}\{d_t < \epsilon^d\}}_{\text{reach the goal location or not}} + \underbrace{\gamma^{suc} * \mathbb{1}\{d_t < \epsilon^d\} * \mathbb{1}\{\text{extract}_{\mathbf{Y}}(\mathbf{a}_t) < \epsilon^a\}}_{\text{match the goal view or not}} \\ & + \underbrace{r_d(d_t, d_{t-1}) + \mathbb{1}\{d_t < \epsilon^d\} * \text{extract}_{\mathbf{Y}}(r_a(\mathbf{a}_t, \mathbf{a}_{t-1}))}_{\text{closer to the goal or not}} - \gamma^{delay}, \end{aligned} \quad (3)$$

which is comprised of four parts. **First**, a success reward $\gamma^{suc} = 5$ is given when distance to goal metric d_t is smaller than threshold ϵ^d . **Second**, an angle success reward is given when the agent reaches the goal location and faces the goal orientation, where $\text{extract}_{\mathbf{Y}}(\mathbf{a}_t)$ means the component of quaternion \mathbf{a}_t rotating around the \mathbf{Y} axis. We only encourage the agent to head to the target in the x-z plane while ignoring its pitch heading. **Third**, a dense reward compose of reduced distance r_d and angle r_a to the target at time step t . **Last**, a delay penalty γ^{delay} is to encourage efficiency.

4.4 Semantic Expansion Inference Scheme

In the inference stage, the agent receives a series of text descriptions of the target object. Existing methods directly replace the image-goal embedding with the text-goal embedding of the input text description to perform zero-shot evaluation. However, the granularity gap [33] between image and text embeddings could considerably impair the evaluation performance [32]. Since it is non-trivial to obtain text annotations for all goal images (28.8M) during training, we opt to enrich the text goal with fine-grained visual features to reduce the gap between image-goal pre-training and text-goal evaluation. We name this method semantic expansion inference.

During training, we maintain a support set $\mathcal{V} = \{\mathbf{v} \in \mathbb{R}^d \mid \mathbf{v}_i \neq \mathbf{v}_j, g(\mathbf{v}_i, \mathbf{v}_j) < \lambda\}$ without using any category information. As we expect the selected embeddings to be as dissimilar as possible from the ones already in the set, we set a threshold $\lambda = 0.8$ to filter representative image embeddings, resulting in around 0.1M vectors. These embeddings approximate the distribution of goal images, and they encode rich information from a variety of object instances, regarding their detailed attributes (*e.g.*, color, shape, texture, and material). *During inference*, given a language description of the target object, we leverage the CLIP text encoder to produce a feature \mathbf{z}_T . Then we query a goal embedding using the retrieval operation:

$$\mathbf{z}_R = \sum_{\mathbf{v}_i \in \mathcal{V}} \frac{\exp(g(\mathbf{z}_T, \mathbf{v}_i))}{\sum_{\mathbf{v}_j \in \mathcal{V}} \exp(g(\mathbf{z}_T, \mathbf{v}_j))} * \mathbf{v}_i. \tag{4}$$

We perform a weighted sum based on the similarity score to avoid inferior performance produced by the nearest neighbor querying [32].

5 Experiments

In this section, we first compare the proposed PSL agent with previous zero-shot methods in the `ObjectNav` task that target finding arbitrary object instances of the same category, and then investigate how it performs when instructed to navigate to a specific object in the `InstanceNav` task. We also conduct extensive ablation experiments to study the effectiveness of different modules in our method empirically.

5.1 Experimental Setup

Datasets. During pre-training, we utilize the 7,200,000 episodes generated by ZSON [38] from HM3D [59]. We randomly select 4 goal images for each episode from 10 candidates with minimal entropy in Eqn. 2. For `ObjectNav` evaluation, we adopt the `habitat2022` challenge dataset with 2000 episodes and 6 object categories. For `InstanceNav`, we first evaluate the agent on IIN dataset [21], an image-goal track with 795 unique instances in 1000 test episodes, and then extend it to the text-goal setting. The agent is tested on 20 scenes excluded from the 800 training scenes.

Agent Configuration. We follow the configuration in [38] with an agent height of 0.88m, radius of 0.18m, and a single 640×480 RGB sensor with a 79° horizontal field-of-view (HFOV) placed 0.88m from the ground. For **InstanceNav**, in the image-goal setting, the agent receives a goal image captured in a random camera configuration that does not match the agent’s own. In the text-goal setting, the agent receives text description of both intrinsic and extrinsic attributes (e.g., “Find a double bed and its material is a beige bamboo frame. There are only two objects around the bed: a table and a window.”). For **ObjectNav**, the agent receives the category text as instruction like “Find a bed”.

Evaluation Metrics. We evaluate the agent’s performance by two metrics: Success Rate (SR) and Success Rate Weighted by Path Length (SPL). An episode is marked as a success if the agent navigates to the goal object and executes the **STOP** action within 0.1m to any viewpoints of the goal object. The SR metric is the average number of success indicators among all episodes. The SPL metric is defined as $\frac{1}{N} \sum_{i=1}^N \text{SR}_i \frac{l_i^*}{\max(l_i, l_i^*)}$, where l_i denotes to the agent’s path length and l_i^* denotes to the ground-truth path length. We report the average number over three runs with different random seeds.

5.2 Experiments on the ObjectNav Task

In Table 1, we report the zero-shot evaluation results on the HM3D **ObjectNav** task. Our PSL agent outperforms the ZSON baseline [38] by 16.9% in SR (42.4% vs. 25.5%), thanks to our training strategy and inference scheme design. More importantly, our method exceeds the LLM-based methods in the first time. Compared to the PixelNav method [6] that learns to navigate to a point deduced by a powerful multi-modal LLM [41], our PSL agent achieves +4.5% improvement in SR. Furthermore, taking into account mapping-based methods that perform LLM-guided frontier-based exploration [60], we see +6.0% improvements against L3MVN [62] and +3.2% against ESC [64]. Our method shows great potential to be applied in the real environment as it does not require accurate GPS coordinates to build a map. Besides, our LLM-free architecture is more efficient.

Table 1: Comparison with state-of-the-art methods on the ObjectNav task. Our PSL surpasses both LLM-based and Mapping-based methods in terms of Success Rate (SR).

| Method | with Mapping | with LLM | LLM | Extra Sensors | SR | SPL |
|--------------|--------------|----------|---------|---------------|-------------|-------------|
| L3MVN [62] | ✓ | ✓ | GPT-2 | Depth, GPS | 35.2 | 16.5 |
| PixelNav [6] | ✗ | ✓ | GPT-4 | - | 37.9 | 20.5 |
| ESC [64] | ✓ | ✓ | GPT-3.5 | Depth, GPS | 39.2 | 22.3 |
| CoW [11] | ✓ | ✗ | - | Depth, GPS | 6.1 | 3.9 |
| ProcTHOR [9] | ✓ | ✗ | - | Depth, GPS | 13.2 | 7.7 |
| ZSON [38] | ✗ | ✗ | - | - | 25.5 | 12.6 |
| PSL (Ours) | ✗ | ✗ | - | - | 42.4 | 19.2 |



Fig. 4: Comparison of the **InstanceNav** task with the **ObjectNav** task. The **ObjectNav** task allows the agent to navigate to any chair in the scene, while the **InstanceNav** task has only one target object specified by detailed attribute descriptions.

5.3 Experiments on the InstanceNav Task

The **ObjectNav** task assigns multiple available destinations for each episode. It is not sure whether the agent can distinguish object instances. In comparison, we extend the zero-shot object navigation task to the zero-shot instance navigation in the *widely-used HM3D environment*, as illustrated in Fig. 4. In this task, the agent receives a textual description of the targeted instance or an instance image as the goal. It needs to understand the detailed attributes of an object instance and navigate to an exact location.

Dataset Preparation. Previous works on the instance-level object navigation task focus on discrete graph-based environments [65] and are limited by complicated human annotation process [11]. In this work, instead, we utilize the recent success of the Generative Vision-and-Language Models (GVLMs) [36, 41, 52] to automatically build the open-vocabulary text-goal setting for the **InstanceNav** task on the popular HM3D environment. Specifically, we randomly select a goal view for each validation episode to render an image of the goal object in the Instance Image Navigation (IIN) dataset [21]. Each episode in the IIN dataset corresponds to a unique goal object instance. To specify each object instance, we follow a previous attempt [39] to separate text descriptions into two aspects: Intrinsic Attributes and Extrinsic Attributes. Intrinsic attributes cover inherent characteristics of the object, such as shape, color, and material. Extrinsic attributes describe the environment surrounding the object, which is used to determine instances with similar intrinsic attributes. We instruct a GVLM model (*i.e.*, CogVLM [52]) to generate both types of attributes according to the instance image with a hand-crafted prompt (See Sec. A in the Appendix for details). The original ground truth trajectories are preserved to construct the additional **InstanceNav** test set. In total, the test set of the text-goal setting comprises 1,000 episodes featuring 795 unique objects across 36 scenes.

Results under the Text-goal Setting. To demonstrate our PSL agent’s ability to understand complicated text descriptions and navigate to specified instances, we compare it with state-of-the-art zero-shot object-goal navigation methods under the text-goal setting of the `InstanceNav` task. In Table 2, we find that most of the existing zero-shot object navigation methods struggle in this task, including both map-based methods CoW [11] (1.8% in SR) and LLM-based method ESC [64] (6.5%). All these methods perform inferior to the methods that receive semantic goals (*i.e.*, ours and ZSON), indicating the effectiveness of end-to-end approaches that are pre-trained on large-scale unlabeled data in this task. Similar to our approach, we also transfer a well-trained ImageNav method (OVRL [58]) to perform the text-goal evaluation using our semantic expansion inference scheme. We find that though good at the ImageNav task, the OVRL model performs poorly when given retrieved image goals, indicating that the learnable goal encoder and observation encoder failed to encode semantic-level information for the `InstanceNav` task. In comparison, our PSL further brings a significant improvement over the ZSON baseline (+5.9% in SR).

Table 2: Comparison with state-of-the-art methods in the text-goal track of the `InstanceNav` task. We report the baseline results based on the released code and models. [†]We perform evaluation with our proposed semantic expansion inference scheme.

| Method | Backbone | with LLM | with Mapping | Extra Sensors | SR | SPL |
|------------------------|-----------|----------|--------------|---------------|-------------|------------|
| CoW [11] | ViT-Base | ✗ | ✓ | Depth, GPS | 1.8 | 1.1 |
| GoW [64] | ViT-Base | ✗ | ✓ | Depth, GPS | 7.2 | 4.2 |
| ESC [64] | ViT-Base | ✓ | ✓ | Depth, GPS | 6.5 | 3.7 |
| OVRL [†] [58] | ResNet-50 | ✗ | ✗ | - | 3.7 | 1.8 |
| ZSON [38] | ResNet-50 | ✗ | ✗ | - | 10.6 | 4.9 |
| PSL (Ours) | ResNet-50 | ✗ | ✗ | - | 16.5 | 7.5 |

Results under the Image-goal Setting. We then compare our PSL agent with other methods under the original Image-goal setting [21]. The agent is tasked to navigate to an object instance specified by an image. This task is different from the `ImageNav` task since the goal images are shot from perspectives with different camera configurations, including different pitch angles and HFOV. Existing state-of-the-art methods for this task heavily rely on supervision signals from scene object annotations and labeled viewpoints. For a fair comparison, we mainly focus on the performance of zero-shot approaches. Results in Table 3 show that our method brings significant improvements on state-of-the-art image-goal navigation methods (+13.1% compared to FGPrompt [56] and +8.4% compared to ZSON [38]). Meanwhile, we significantly reduce the gap with the supervised methods like OVRL-V2 in terms of SPL (11.4% *vs.* 11.8%), even though we use a much smaller vision backbone (ResNet-50 *vs.* ViT-Base).

Table 3: Comparison with state-of-the-art methods in the image-goal track of the InstanceNav task. In this track, the ‘‘Supervised’’ mark means human labels on the objects are used. †We re-implement OVRL based on released pre-trained weight.

| Method | Backbone | Supervised | Pre-training Data | SR | SPL |
|---------------|-----------|------------|-------------------|-------------|-------------|
| RL Agent [21] | ResNet-18 | ✓ | - | 8.3 | 3.5 |
| OVRL-V2 [57] | ViT-Base | ✓ | Gibson | 24.8 | 11.8 |
| OVRL-V2 [57] | ViT-Base | ✗ | Gibson | 0.6 | 0.2 |
| OVRL† [58] | ResNet-50 | ✗ | HM3D | 8.0 | 4.2 |
| FGPrompt [56] | ResNet-9 | ✗ | HM3D | 9.9 | 2.8 |
| ZSON [38] | ResNet-50 | ✗ | HM3D | 14.6 | 7.3 |
| PSL (Ours) | ResNet-50 | ✗ | HM3D | 23.0 | 11.4 |

5.4 Ablation Study

Effectiveness of Prioritized Semantic Learning. We first conduct ablation studies under the image-goal setting of the InstanceNav task. We start from the ZSON [38] baseline, with a fixed CLIP goal encoder and a learnable ResNet-50 observation encoder. Then we apply the entropy-minimized goal view selection algorithm to select meaningful goal images during pre-training data pre-processing. In Table 4, this modification brings +2.1% improvement in success rate. We then add the perspective reward relaxation technique, which enhances goal orientation with an alternative pitch angle. This approach aims to prevent the agent from being immersed in exactly matching the goal view and observations, neglecting the semantic correlation between them. However, we discovered that with the vanilla model design, the augmented image goal confuses the agent and decreases the navigation success rate (from 12.7% to 10.8%). We argue that this phenomenon is due to the inferior agent design that neglects to reason semantic differences in the goal images and observation images. Incorporating with the PSL agent alleviates this issue. We found that it boosts the navigation success rate by +3.8% with the goal view selection approach, but more impressively, it cooperates well with the perspective relaxation mechanism and contributes a lot to the improvement (+9.3% in SR and +4.2% in SPL).

Ablation on Semantic Expansion Inference. To alleviate the granularity inconsistency issue of text and image embeddings during testing, we propose a semantic expansion inference scheme for our PSL agent. We perform elaborate ablation studies under the text-goal setting to investigate its effectiveness. We compare two different implementations of support set in the semantic expansion inference scheme, including collecting goal images from the image-goal navigation dataset, which refers to the ImageNav support set; and from the instance image navigation dataset, namely the IIN support set. In Table 5, we start with ablating the detailed configurations of the retrieval module and have two main findings: 1) using the support set images from the limited categories of objects as the dataset defined yields inferior results (11.1% in SR). Instead, the diversity of

Table 4: Ablation studies of different components in our Prioritized Semantic Learning (PSL) agent and Prioritized Semantic Training (PST) strategy under the Image-Goal setting of the `InstanceNav` task. The default entry is marked in gray. “SPM”: Semantic Perception Module; “GVS”: Goal View Selection; “PRR”: Perspective Reward Selection.

| | PSL | | PST | | ZSIN-image | | ZSIN-text | | ZSON | |
|---------------|-----|-----|-----|-------------|-------------|-------------|------------|-------------|-------------|--|
| | SPM | GVS | PRR | SR | SPL | SR | SPL | SR | SPL | |
| ZSON | ✗ | ✗ | ✗ | 12.7 | 6.5 | 10.6 | 6.5 | 25.5 | 12.6 | |
| | ✓ | ✗ | ✗ | 19.5 | 7.9 | 13.0 | 5.6 | 33.7 | 15.8 | |
| PSL | ✗ | ✓ | ✗ | 14.8 | 7.7 | 11.8 | 6.1 | 30.4 | 14.7 | |
| (Ours) | ✓ | ✓ | ✗ | 16.5 | 6.5 | 12.3 | 5.7 | 35.0 | 18.1 | |
| | ✓ | ✓ | ✓ | 22.0 | 10.7 | 16.5 | 7.5 | 42.4 | 19.2 | |

the ImageNav support set brings +1.3% improvements. 2) directly replacing the image features with text queries performs poorly (6.6% in SR) due to the modality gap [33]. Retrieving image features to augment the text queries yields better goal embeddings that share the same level of semantic granularity as training, thus significantly improving the navigation success rate by +9.9%.

Table 5: Ablation studies of the Semantic Expansion Inference (SEI) scheme under the Text-Goal setting of the `InstanceNav` task.

| PSL Agent | SEI Scheme | Support Set | #Supp. Vec. | SR | SPL |
|-----------|------------|-------------|-------------|-------------|------------|
| ✗ | ✓ | IIN | 3.5K | 11.1 | 5.4 |
| ✗ | ✓ | ImageNav | 0.1M | 12.4 | 6.6 |
| ✓ | ✗ | ImageNav | 0.1M | 6.6 | 2.7 |
| ✓ | ✓ | ImageNav | 0.1M | 16.5 | 7.5 |

6 Conclusion

We present PSL, a zero-shot approach for instance navigation. PSL is built upon the open-vocabulary and cross-modal ability of CLIP, and targets addressing the limited semantic perceiving ability of existing methods. Our method is composed of three different parts, including a PSL agent with a semantic perception module to reason the semantic differences between the observation and goal, a prioritized semantic training strategy with clear semantic supervision to train the PSL agent, and a semantic expansion inference scheme to maintain the granularity of the goal-semantic. We investigate how this agent acts when instructed to navigate to a specific object by extending the zero-shot object navigation task to a zero-shot instance navigation task. We find that most existing methods struggle in this task, while our PSL agent outperforms other approaches.

Acknowledgements

This work was supported by the Meituan Academy of Robotics Shenzhen, the Guangzhou Municipal Science and Technology Project (No. 2024A03J0619) and the National Natural Science Foundation of China (No. 62306257).

References

1. Al-Halah, Z., Ramakrishnan, S.K., Grauman, K.: Zero experience required: Plug & play modular transfer learning for semantic visual navigation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17010–17020 (2022)
2. Anderson, P., Wu, Q., Teney, D., Bruce, J., Johnson, M., Sünderhauf, N., Reid, I., Gould, S., Van Den Hengel, A.: Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3674–3683 (2018)
3. Anderson, P., Wu, Q., Teney, D., Bruce, J., Johnson, M., Sünderhauf, N., Reid, I.D., Gould, S., van den Hengel, A.: Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2018)
4. Batra, D., Chang, A.X., Chernova, S., Davison, A.J., Deng, J., Koltun, V., Levine, S., Malik, J., Mordatch, I., Mottaghi, R., Savva, M., Su, H.: Rearrangement: A challenge for embodied AI. CoRR **abs/2011.01975** (2020)
5. Batra, D., Gokaslan, A., Kembhavi, A., Maksymets, O., Mottaghi, R., Savva, M., Toshev, A., Wijmans, E.: Objectnav revisited: On evaluation of embodied agents navigating to objects. arXiv preprint arXiv:2006.13171 (2020)
6. Cai, W., Huang, S., Cheng, G., Long, Y., Gao, P., Sun, C., Dong, H.: Bridging zero-shot object navigation and foundation models through pixel-guided navigation skill. CoRR **abs/2309.10309** (2023)
7. Chang, A., Dai, A., Funkhouser, T., Halber, M., Niessner, M., Savva, M., Song, S., Zeng, A., Zhang, Y.: Matterport3d: Learning from rgb-d data in indoor environments. arXiv preprint arXiv:1709.06158 (2017)
8. Chen, P., Sun, X., Zhi, H., Zeng, R., Li, T.H., Liu, G., Tan, M., Gan, C.: a^2 nav: Action-aware zero-shot robot navigation by exploiting vision-and-language ability of foundation models. arXiv preprint arXiv:2308.07997 (2023)
9. Deitke, M., VanderBilt, E., Herrasti, A., Weihs, L., Ehsani, K., Salvador, J., Han, W., Kolve, E., Kembhavi, A., Mottaghi, R.: Proctor: Large-scale embodied AI using procedural generation. In: Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., Oh, A. (eds.) Proceedings of the International Conference on Neural Information Processing Systems (2022)
10. Du, Y., Gan, C., Isola, P.: Curious representation learning for embodied intelligence. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2021)
11. Gadre, S.Y., Wortsman, M., Ilharco, G., Schmidt, L., Song, S.: CLIP on wheels: Zero-shot object navigation as object localization and exploration. CoRR **abs/2203.10421** (2022)

12. Gan, C., Gu, Y., Zhou, S., Schwartz, J., Alter, S., Traer, J., Gutfreund, D., Tenenbaum, J.B., McDermott, J.H., Torralba, A.: Finding fallen objects via asynchronous audio-visual integration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10523–10533 (2022)
13. Gan, C., Schwartz, J., Alter, S., Mrowca, D., Schrimpf, M., Traer, J., De Freitas, J., Kubilius, J., Bhandwaldar, A., Haber, N., et al.: Threedworld: A platform for interactive multi-modal physical simulation. arXiv preprint arXiv:2007.04954 (2020)
14. Gao, P., Han, J., Zhang, R., Lin, Z., Geng, S., Zhou, A., Zhang, W., Lu, P., He, C., Yue, X., et al.: Llama-adapter v2: Parameter-efficient visual instruction model. arXiv preprint arXiv:2304.15010 (2023)
15. Hahn, M., Chaplot, D.S., Tulsiani, S., Mukadam, M., Rehg, J.M., Gupta, A.: No rl, no simulation: Learning to navigate without navigating. In: Proceedings of the International Conference on Neural Information Processing Systems (2021)
16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)
17. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: Proceedings of the International Conference on Machine Learning. pp. 4904–4916. PMLR (2021)
18. Kim, N., Kwon, O., Yoo, H., Choi, Y., Park, J., Oh, S.: Topological semantic graph memory for image-goal navigation. In: Conference on Robot Learning. pp. 393–402. PMLR (2023)
19. Krantz, J., Gervet, T., Yadav, K., Wang, A., Paxton, C., Mottaghi, R., Batra, D., Malik, J., Lee, S., Chaplot, D.S.: Navigating to objects specified by images. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10916–10925 (2023)
20. Krantz, J., Gokaslan, A., Batra, D., Lee, S., Maksymets, O.: Waypoint models for instruction-guided navigation in continuous environments. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15162–15171 (2021)
21. Krantz, J., Lee, S., Malik, J., Batra, D., Chaplot, D.S.: Instance-specific image goal navigation: Training embodied agents to find object instances. arXiv preprint arXiv:2211.15876 (2022)
22. Krantz, J., Wijmans, E., Majumdar, A., Batra, D., Lee, S.: Beyond the nav-graph: Vision-and-language navigation in continuous environments. In: European Conference on Computer Vision. pp. 104–120. Springer (2020)
23. Ku, A., Anderson, P., Patel, R., Ie, E., Baldrige, J.: Room-Across-Room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In: Conference on Empirical Methods for Natural Language Processing (2020)
24. Kwon, O., Kim, N., Choi, Y., Yoo, H., Park, J., Oh, S.: Visual graph memory with unsupervised representation for visual navigation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15890–15899 (2021)
25. Li, C., Xia, F., Martín-Martín, R., Lingelbach, M., Srivastava, S., Shen, B., Vainio, K., Gokmen, C., Dharan, G., Jain, T., et al.: igibson 2.0: Object-centric simulation for robot learning of everyday household tasks. arXiv preprint arXiv:2108.03272 (2021)
26. Li, C., Xia, F., Martín-Martín, R., Lingelbach, M., Srivastava, S., Shen, B., Vainio, K.E., Gokmen, C., Dharan, G., Jain, T., Kurenkov, A., Liu, C.K., Gweon, H., Wu, J., Fei-Fei, L., Savarese, S.: igibson 2.0: Object-centric simulation for robot learning

- of everyday household tasks. In: Conference on Robot Learning. vol. 164, pp. 455–465 (2021)
27. Li, C., Zhang, R., Wong, J., Gokmen, C., Srivastava, S., Martín-Martín, R., Wang, C., Levine, G., Lingelbach, M., Sun, J., Anvari, M., Hwang, M., Sharma, M., Aydin, A., Bansal, D., Hunter, S., Kim, K., Lou, A., Matthews, C.R., Villa-Renteria, I., Tang, J.H., Tang, C., Xia, F., Savarese, S., Gweon, H., Liu, K., Wu, J., Fei-Fei, L.: BEHAVIOR-1K: A benchmark for embodied AI with 1, 000 everyday activities and realistic simulation. In: Conference on Robot Learning. vol. 205, pp. 80–93 (2022)
 28. Li, C., Zhang, R., Wong, J., Gokmen, C., Srivastava, S., Martín-Martín, R., Wang, C., Levine, G., Lingelbach, M., Sun, J., et al.: Behavior-1k: A benchmark for embodied ai with 1,000 everyday activities and realistic simulation. In: Conference on Robot Learning. pp. 80–93. PMLR (2023)
 29. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. arXiv preprint arXiv:2301.12597 (2023)
 30. Li, J., Selvaraju, R., Gotmare, A., Joty, S., Xiong, C., Hoi, S.C.H.: Align before fuse: Vision and language representation learning with momentum distillation. Proceedings of the International Conference on Neural Information Processing Systems **34**, 9694–9705 (2021)
 31. Li, L.H., Yatskar, M., Yin, D., Hsieh, C.J., Chang, K.W.: Visualbert: A simple and performant baseline for vision and language. arXiv preprint arXiv:1908.03557 (2019)
 32. Li, W., Zhu, L., Wen, L., Yang, Y.: Decap: Decoding clip latents for zero-shot captioning via text-only training. arXiv preprint arXiv:2303.03032 (2023)
 33. Liang, V.W., Zhang, Y., Kwon, Y., Yeung, S., Zou, J.Y.: Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. In: Proceedings of the International Conference on Neural Information Processing Systems. vol. 35, pp. 17612–17625 (2022)
 34. Liu, H., Li, C., Li, Y., Lee, Y.J.: Improved baselines with visual instruction tuning. arXiv preprint arXiv:2310.03744 (2023)
 35. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. arXiv preprint arXiv:2304.08485 (2023)
 36. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. arXiv preprint arXiv:2304.08485 (2023)
 37. Ma, S., Wang, Y., Wei, Y., Fan, J., Sun, X., Chen, P., Zhang, E.: A simple knowledge distillation framework for open-world object detection. arXiv preprint arXiv:2312.08653 (2023)
 38. Majumdar, A., Aggarwal, G., Devnani, B., Hoffman, J., Batra, D.: ZSON: zero-shot object-goal navigation using multimodal goal embeddings. In: Proceedings of the International Conference on Neural Information Processing Systems (2022)
 39. Majumdar, A., Xia, F., Batra, D., Guibas, L., et al.: Findthis: Language-driven object disambiguation in indoor environments. In: Conference on Robot Learning (2023)
 40. Mezghani, L., Sukhbaatar, S., Lavril, T., Maksymets, O., Batra, D., Bojanowski, P., Alahari, K.: Memory-augmented reinforcement learning for image-goal navigation. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (2022)
 41. OpenAI: GPT-4 technical report. CoRR **abs/2303.08774** (2023)
 42. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable

- visual models from natural language supervision. In: Meila, M., Zhang, T. (eds.) *Proceedings of the International Conference on Machine Learning*. vol. 139, pp. 8748–8763 (2021)
43. Ramakrishnan, S.K., Gokaslan, A., Wijmans, E., Maksymets, O., Clegg, A., Turner, J., Undersander, E., Galuba, W., Westbury, A., Chang, A.X., et al.: Habitat-matterport 3d dataset (hm3d): 1000 large-scale 3d environments for embodied ai. *arXiv preprint arXiv:2109.08238* (2021)
 44. Ramrakhya, R., Batra, D., Wijmans, E., Das, A.: Pirlnav: Pretraining with imitation and rl finetuning for objectnav. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 17896–17906 (2023)
 45. Savva, M., Malik, J., Parikh, D., Batra, D., Kadian, A., Maksymets, O., Zhao, Y., Wijmans, E., Jain, B., Straub, J., Liu, J., Koltun, V.: Habitat: A platform for embodied AI research. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2019)
 46. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017)
 47. Srivastava, S., Li, C., Lingelbach, M., Martín-Martín, R., Xia, F., Vainio, K.E., Lian, Z., Gokmen, C., Buch, S., Liu, C.K., Savarese, S., Gweon, H., Wu, J., Fei-Fei, L.: BEHAVIOR: benchmark for everyday household activities in virtual, interactive, and ecological environments. In: *Conference on Robot Learning*. vol. 164, pp. 477–490 (2021)
 48. Srivastava, S., Li, C., Lingelbach, M., Martín-Martín, R., Xia, F., Vainio, K.E., Lian, Z., Gokmen, C., Buch, S., Liu, K., et al.: Behavior: Benchmark for everyday household activities in virtual, interactive, and ecological environments. In: *Conference on Robot Learning*. pp. 477–490. PMLR (2022)
 49. Szot, A., Clegg, A., Undersander, E., Wijmans, E., Zhao, Y., Turner, J., Maestre, N., Mukadam, M., Chaplot, D.S., Maksymets, O., Gokaslan, A., Vondrus, V., Dharur, S., Meier, F., Galuba, W., Chang, A.X., Kira, Z., Koltun, V., Malik, J., Savva, M., Batra, D.: Habitat 2.0: Training home assistants to rearrange their habitat. In: *Proceedings of the International Conference on Neural Information Processing Systems* (2021)
 50. Thomason, J., Murray, M., Cakmak, M., Zettlemoyer, L.: Vision-and-dialog navigation. In: *Conference on Robot Learning*. pp. 394–406. PMLR (2020)
 51. Udandarao, V., Burg, M.F., Albanie, S., Bethge, M.: Visual data-type understanding does not emerge from scaling vision-language models. In: *Proceedings of the International Conference on Learning Representations* (2023)
 52. Wang, W., Lv, Q., Yu, W., Hong, W., Qi, J., Wang, Y., Ji, J., Yang, Z., Zhao, L., Song, X., et al.: Cogvlm: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079* (2023)
 53. Weihs, L., Deitke, M., Kembhavi, A., Mottaghi, R.: Visual room rearrangement. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5922–5931 (2021)
 54. Wijmans, E., Kadian, A., Morcos, A., Lee, S., Essa, I., Parikh, D., Savva, M., Batra, D.: DD-PPO: learning near-perfect pointgoal navigators from 2.5 billion frames. In: *Proceedings of the International Conference on Learning Representations* (2020)
 55. Wu, Y., Wu, Y., Gkioxari, G., Tian, Y.: Building generalizable agents with a realistic and rich 3d environment. *arXiv preprint arXiv:1801.02209* (2018)
 56. Xinyu, S., Peihao, C., Jugang, F., Thomas, H.L., Jian, C., Mingkui, T.: Fgprompt: Fine-grained goal prompting for image-goal navigation. In: *Proceedings of the International Conference on Neural Information Processing Systems* (2023)

57. Yadav, K., Majumdar, A., Ramrakhya, R., Yokoyama, N., Baevski, A., Kira, Z., Maksymets, O., Batra, D.: Ovrl-v2: A simple state-of-art baseline for imagenav and objectnav. arXiv preprint arXiv:2303.07798 (2023)
58. Yadav, K., Ramrakhya, R., Majumdar, A., Berges, V., Kuhar, S., Batra, D., Baevski, A., Maksymets, O.: Offline visual representation learning for embodied navigation. CoRR **abs/2204.13226** (2022)
59. Yadav, K., Ramrakhya, R., Ramakrishnan, S.K., Gervet, T., Turner, J., Gokaslan, A., Maestre, N., Chang, A.X., Batra, D., Savva, M., et al.: Habitat-matterport 3d semantics dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4927–4936 (2023)
60. Yamauchi, B.: A frontier-based approach for autonomous exploration. In: Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97. pp. 146–151. IEEE (1997)
61. Yenamandra, S., Ramachandran, A., Yadav, K., Wang, A., Khanna, M., Gervet, T., Yang, T.Y., Jain, V., Clegg, A.W., Turner, J., et al.: Homerobot: Open-vocabulary mobile manipulation. arXiv preprint arXiv:2306.11565 (2023)
62. Yu, B., Kasaei, H., Cao, M.: L3mvm: Leveraging large language models for visual target navigation. arXiv preprint arXiv:2304.05501 (2023)
63. Zhang, R., Han, J., Zhou, A., Hu, X., Yan, S., Lu, P., Li, H., Gao, P., Qiao, Y.: Llama-adapter: Efficient fine-tuning of language models with zero-init attention. arXiv preprint arXiv:2303.16199 (2023)
64. Zhou, K., Zheng, K., Pryor, C., Shen, Y., Jin, H., Getoor, L., Wang, X.E.: ESC: exploration with soft commonsense constraints for zero-shot object navigation. In: Proceedings of the International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 202, pp. 42829–42842. PMLR (2023)
65. Zhu, F., Liang, X., Zhu, Y., Yu, Q., Chang, X., Liang, X.: Soon: Scenario oriented object navigation with graph-based exploration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12689–12699 (2021)
66. Zhu, Y., Mottaghi, R., Kolve, E., Lim, J.J., Gupta, A., Fei-Fei, L., Farhadi, A.: Target-driven visual navigation in indoor scenes using deep reinforcement learning. In: IEEE International Conference on Robotics and Automation (2017)