

# REGNav: Room Expert Guided Image-Goal Navigation

Pengna Li\*, Kangyi Wu\*, Jingwen Fu, Sanping Zhou<sup>†</sup>

Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University  
 {sauerfisch, wukangyi747600, fu1371252069}@stu.xjtu.edu.cn, spzhou@xjtu.edu.cn

## Abstract

Image-goal navigation aims to steer an agent towards the goal location specified by an image. Most prior methods tackle this task by learning a navigation policy, which extracts visual features of goal and observation images, compares their similarity and predicts actions. However, if the agent is in a different room from the goal image, it's extremely challenging to identify their similarity and infer the likely goal location, which may result in the agent wandering around. Intuitively, when humans carry out this task, they may roughly compare the current observation with the goal image, having an approximate concept of whether they are in the same room before executing the actions. Inspired by this intuition, we try to imitate human behaviour and propose a Room Expert Guided Image-Goal Navigation model (REGNav) to equip the agent with the ability to analyze whether goal and observation images are taken in the same room. Specifically, we first pre-train a room expert with an unsupervised learning technique on the self-collected unlabelled room images. The expert can extract the hidden room style information of goal and observation images and predict their relationship about whether they belong to the same room. In addition, two different fusion approaches are explored to efficiently guide the agent navigation with the room relation knowledge. Extensive experiments show that our REGNav surpasses prior state-of-the-art works on three popular benchmarks.

**Code** — <https://github.com/leeBooMla/REGNav>

## Introduction

Image-goal navigation (ImageNav) (Zhu et al. 2017) is an emerging embodied intelligence task, where the agent is placed in an unseen environment and needs to navigate to an image-specified goal location using visual observations. Due to its widespread applications in last mile delivery, household robots, and personal robots (Wasserman et al. 2023; Majumdar et al. 2022; Krantz et al. 2023b), it has raised increasing research attention in recent years.

Despite its broad applications, this task remains highly challenging. Since the environment map is unknown, the

\*These authors contributed equally.

<sup>†</sup>Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

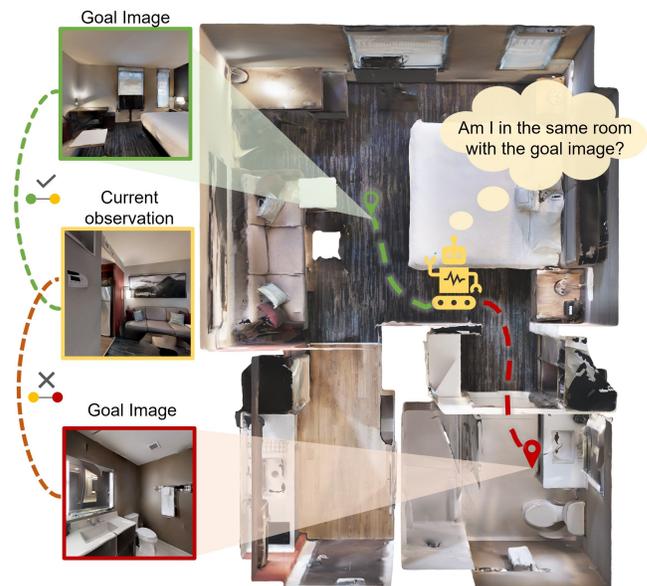


Figure 1: We solve the task of image-goal navigation, where an agent (the yellow robot) is required to navigate to a location depicted by a goal image. To accomplish this, our agent tries to compare the current observation with the goal image and tease out whether the current location is in the same room with the goal image before executing actions.

agent must reason the likely location of the goal image to navigate to. This requires the agent to perceive the environment efficiently, compare the current observation with the goal image and find the associations before taking the action. However, the complex spatial structure of unseen environments often leads to significant discrepancies between the agent's actual location and the goal location (e.g. in different rooms). In such cases, the goal image and the current observation may have little overlap and it becomes challenging for the agent to identify their similarities and associations. This results in the agent failing to reason the goal location, thus taking meaningless actions, such as back-tracking or aimlessly wandering. The key to solving this issue is to extract the spatial information from the observations to help reason the spatial relationship with the goal image.

To incorporate spatial information into robot operations, the modular methods (Hahn et al. 2021; Chaplot et al. 2020b; Krantz et al. 2023b; Lei et al. 2024) employ GPS or depth sensors to construct a geometric or occupancy map and localize the agent (SLAM (Durrant-Whyte and Bailey 2006)) in an unfamiliar environment. However, these methods heavily rely on depth sensors or GPS to provide spatial information, limiting their scalability in real-world deployment. Alternatively, the learning-based methods (Yadav et al. 2023b; Sun et al. 2024) attempt to learn an end-to-end navigation policy solely relying on the RGB sensors. These methods directly extract representations of the goal and visual observation to predict the corresponding action. Although the learning-based methods have shown great potential in this task, they have difficulty in exploiting the spatial relationship between the goal and current observation with only RGB sensors available, limiting their performance.

Imagine that when humans are given the task of finding a place depicted by an image in an unseen environment as illustrated in Figure 1, to find the shortest path to the goal location, humans always try to coarsely estimate spatial relationships about whether the current location is in the same room as the goal image first before further comparing their fine-grained semantics. If not, humans tend to find the door to get out of the current room, which can reduce meaningless actions and move to the target more quickly. Inspired by this, we want to imitate human behaviour and enable the agent to have the ability first to evaluate the coarse spatial relationships between the goal and visual observation, thereby mitigating the issue of the agent’s invalid actions.

To devise such a solution, we have to figure out *what we can rely on to infer the spatial relationships with only RGB images available*. As shown in Figure 1, we observe that different rooms within a house, such as a bedroom, bathroom, and kitchen, often have their specific styles, *e.g.*, decoration style, furniture, floor, and wall. These variations are primarily due to the different functions and requirements of each room. For instance, bedrooms tend to prioritize comfort and they might contain soft lighting, warm colours, and carpeting. Bathrooms often have tiles on the wall and floor for waterproofing and cleaning. This observation suggests that it’s possible for the agent to identify the room style information from the visual signals. The style information can be used to determine whether the current observation is located in the same room as the goal image.

Learning a model to extract the style information from observation images requires a large amount of annotated image data. However, acquiring the supervision signals is costly and may raise fairness concerns. To address this challenge, we attempt to train a Room Expert with an unsupervised learning method to identify the hidden room style information. Specifically, we utilize the unsupervised clustering with *must-link* and *cannot-link* constraints to pre-train a room-style encoder and a room relation network based on the intuition that if two points are far apart, they are likely located in different rooms. From the pre-trained model, the agent can obtain the room style representation of the goal and visual observation and obtain their relation about whether they belong to the same room.

In this paper, we propose a Room Expert Guided Image-Goal Navigation model (REGNav) to explicitly empower the agent with the ability to analyze the spatial relationships between the goal and observation images through a pre-trained room expert model. Specifically, we adopt a two-stage learning scheme: 1) pre-train a room-style expert offline, and 2) incorporate the ability of the pre-trained room expert to learn an efficient navigation policy. The room expert pre-training stage involves adopting an unsupervised learning technique to train a style encoder and a relation network on a large-scale self-collected dataset of images from the indoor environment dataset Gibson (Xia et al. 2018). The collected training images share the same parameters and settings as the observations captured by the agent camera. In the latter stage, we explore two different fusion approaches to efficiently guide the agent navigation with the room relation knowledge. We freeze the parameters of the room expert and train the visual encoder and the navigation policy in the Habitat simulator (Savva et al. 2019). Extensive experiments demonstrate that our proposed method can achieve more successful navigation.

We conclude the main contributions of our paper below:

- We discuss the issue of the agent’s wandering around and explore the feasibility of reasoning the spatial relationships from the pure RGB images.
- We observe that room-style information can be the link between the visual signals and spatial relationships. A novel unsupervised method with *must-link* and *cannot-link* constraints is devised to pre-train a room expert to extract room style and predict the spatial relationships.
- Finally, We present REGNav, an efficient image-goal navigation framework, equipping the agent with the ability to reason spatial relationships.

## Related Works

**Visual navigation.** Visual navigation (Krantz et al. 2023b; Kwon, Park, and Oh 2023; Pelluri 2024; Li and Bansal 2023; Liu et al. 2024; Sun et al. 2024; Wang et al. 2024; Zhao et al. 2024) requires an agent to navigate based on visual sensors. It can be categorized into several types, including Visual-and-Language navigation (VLN), Object Navigation, Image Goal Navigation, etc. Some works (Anderson et al. 2018b; Chen et al. 2021; Li, Tan, and Bansal 2022; Krantz et al. 2023a) focus on VLN, which uses additional natural language instructions to depict the navigation targets. These works either depend on detailed language instructions (Wang et al. 2023; Li et al. 2023) or require conversations with humans (Zhang et al. 2024a; Thomason et al. 2020) during the navigation process, leading to low usability. Object Navigation is proposed with a given object category as the target (Chaplot et al. 2020a; Mayo, Hazan, and Tal 2021; Du et al. 2023; Zhang et al. 2024b). However, this kind of method can only reach the surrounding area of an object and cannot accurately arrive at a specific location. Given the reasons above, we address the Image Goal Navigation task where an arbitrary image is provided as the target and only an RGB sensor is utilized during the navigation process. The agent must reach the location depicted in the goal

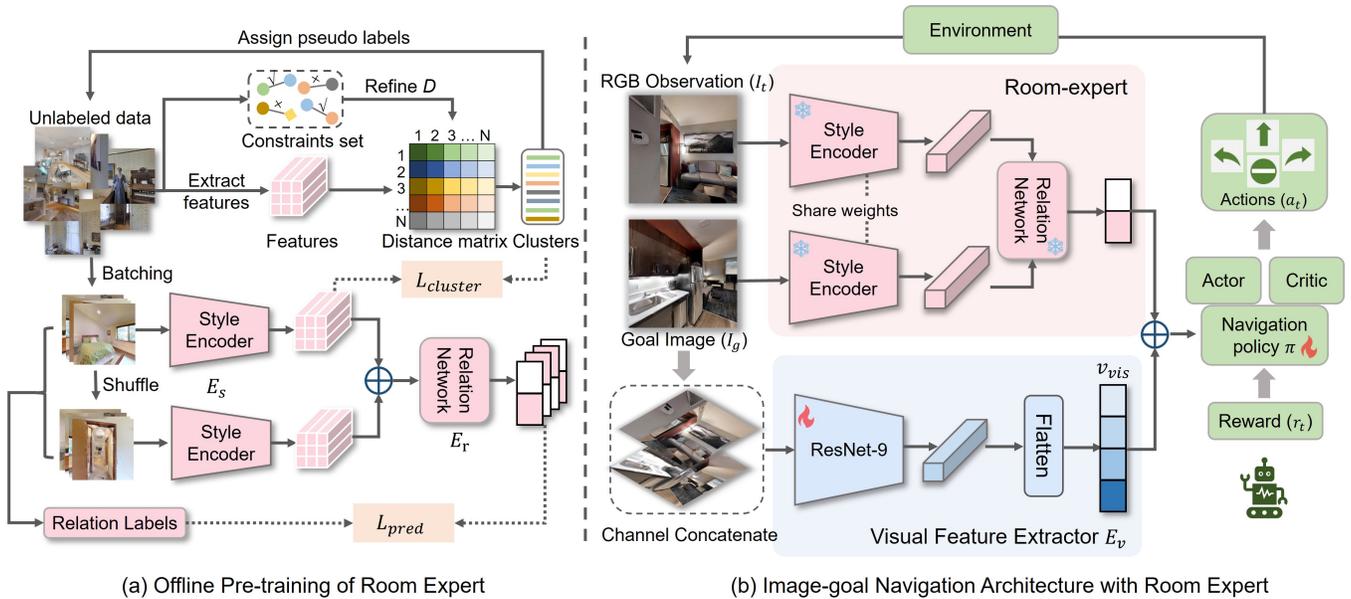


Figure 2: The overview of our REGNav. (a) Pre-training the Room Expert offline. We employ an unsupervised clustering method to train a style encoder and a relation network to extract style representation and predict the relationships. We use the constraints set deduced from the unlabeled data to refine the feature distance matrix to obtain more reliable pseudo labels. (b) The image-goal navigation architecture with Room Expert. We lock the Room Expert and proceed to train the visual encoder and navigation policy. The visual feature extractor regards the channel concatenation of the observation and goal image as input. The navigation policy takes the concatenation of the relation flag (2-dimension) and the fused feature as input.

image. We study how to make full use of the knowledge in observation to improve navigation performance.

**Reinforcement learning in visual navigation.** Since the image navigation method with reinforcement learning (RL) can learn directly from interacting with the environment in an end-to-end manner, it has gained a great population in recent years (Du, Gan, and Isola 2021; Majumdar et al. 2022). Some methods aim to enhance the representation capability of the feature extractors before the RL policy. (Sun et al. 2024) explore fusion methods to guide the observation encoder to focus on goal-relevant regions. (Sun et al. 2025) propose a prioritized semantic learning method to improve the agents’ semantic ability. Some works (Li and Bansal 2023; Wang et al. 2024) utilized the pre-training strategy to enforce the agent to have an expectation of the future environments. However, if the agent has a large distance from the goal, these methods may fail to extract useful knowledge from the observations. Some methods try to incorporate additional memory mechanisms to enable long-term reasoning and exploit supplementary knowledge from previous states. (Mezghan et al. 2022) trained a state-embedding network to take advantage of the history with external memory. (Qiao et al. 2022) devised a history-and-order pre-training paradigm to exploit past observations and support future prediction. (Kim et al. 2023) inserted semantic information into topological graph memory to obtain a thorough description of history states. (Li et al. 2024) classified history states into three types to ensure both diversity and long-term memory. However, these methods have no spatial awareness if the

agent has never been to the area near the target. On the contrary, we aim to equip the agent with spatial awareness and enable it to analyze whether the observation is in the same space as the goal.

**Auxiliary knowledge in visual navigation.** Image-Goal Navigation requires the agent to navigate to an image-specified goal location in an unseen environment using visual observations. Only depending on a single RGB sensor has raised the challenge and makes the task difficult even for humans (Paul, Roy-Chowdhury, and Cherian 2022). To release the difficulty, auxiliary knowledge is introduced. (Liu et al. 2024) enables the agent to interact with a human for help when it’s unable to solve the task. (Li et al. 2023) utilizes an external pre-trained image description model to provide additional knowledge. (Kim et al. 2023) introduces a pre-trained semantic segmentation model to extract objects in both observations and targets. All of these methods require external datasets to acquire auxiliary knowledge, which may have fairness concerns. In contrast, we devised a Room Expert trained with an unsupervised clustering method without using any additional dataset. Our Room Expert effectively empowers the agent with spatial awareness to analyze the spatial relationship between observations and the goal location, improving the navigation performance.

## Proposed Method

**Task setup.** ImageNav tasks involve directing an agent to a destination depicted in a target image  $I_g$ , taken at the goal location. Initially positioned at a random starting point  $p_0$ , the

agent is equipped solely with this goal image  $I_g$  from the environment. At each time step  $t$ , it perceives the environment through an egocentric RGB image  $I_t$ , captured by an onboard RGB sensor. Then the agent takes an action conditioned on  $v_t$  and  $v_g$ . These actions denoted as  $a_t$ , are guided by a trained policy in reinforcement learning frameworks.

**Reward** After acting  $a_t$ , a reward  $r_t$  is given to the agent, encouraging it to reach the goal location through the shortest path (Al-Halah, Ramakrishnan, and Grauman 2022). Both the reduced distance and the reduced angle in radians are utilized to provide the reward to the agent. The overall reward function for time step  $t$  can be formulated as:

$$r_t = r_d(d_t, d_{t-1}) + \mathcal{I}(d_t \leq d_s)r_\alpha(\alpha_t, \alpha_{t-1}) - \gamma, \quad (1)$$

where  $r_d$  and  $r_\alpha$  are the reduced distance to the goal from the current position and the reduced angle in radians to the goal view from the current view respectively relative to the previous one,  $\gamma$  represents a slack reward to encourage efficiency and  $\mathcal{I}$  denotes an indicator. What’s more, the agent receives a maximum success reward  $R_s$  if it reaches the goal and stops within a distance  $d_s$  from the goal location and an angle  $\alpha_s$  from the goal angle. The success reward can be formulated as:

$$R_s = 5 \times [\mathcal{I}(d_t \leq d_s) + \mathcal{I}(d_t \leq d_s \text{ and } \alpha_t \leq \alpha_s)], \quad (2)$$

Here, we set the  $d_s = 1\text{m}$  and  $\alpha_s = 25^\circ$ .

## REGNav

In this section, we detail the REGNav methodology. It adopts a two-stage learning strategy: 1) pre-training room expert in an unsupervised manner. We first collect a new room-relation image dataset from the indoor dataset Gibson (Xia et al. 2018). Based on this, we learn a room expert composed of a style encoder and a relation network in a novel clustering method with constraints set. 2) learning a navigation policy conditioned room expert. We design and explore two different fusion manners to train the visual encoder and navigation policy with the room expert frozen.

### Room Expert Pre-Training

**Dataset collection.** To obtain the room style representation from observation and goal images, our Room Expert needs to be trained with images taken in different rooms. To avoid focusing on the varied objects between rooms instead of the room style to analyze room relation, images should also be taken from different angles in the same room (since these images will have completely different objects while they still represent the same room). To ensure the generalizability of the room style representation, images taken in different scenes or houses should also be available. Currently, no publicly available dataset meets the aforementioned requirements. MP3D dataset (Chang et al. 2017) has provided room annotations. However, previous image-goal navigation models (Sun et al. 2024; Majumdar et al. 2022) only use Gibson training episode (Mezghan et al. 2022) to train the agent and use MP3D testing episode (Al-Halah, Ramakrishnan, and Grauman 2022) to evaluate. Directly using room annotations

from the MP3D dataset will cause fairness concerns. Therefore, images are collected from the training episodes of Gibson (Mezghan et al. 2022). Specifically, for a given training episode  $E_m$ , we first extract the starting location  $p_{ms}$  and the target location  $p_{mt}$ . Then, agents equipped with a single egocentric RGB sensor are put in these two locations and take images from varied angles. Lastly, we annotate the collected images  $\{I_i\}_N$  with scene identity  $\{S_i\}_N$  which indicates the 3D scene or house, episode identity  $\{E_i\}_N$ , and episode difficulty  $\{Ed_i\}_N$ . We observed that some images collected in this way may contain little room-style information (e.g. when the RGB sensor is too close to the wall, the images taken will be completely black or white.). These blank images, if used in the training process, will provide confusing guidance. To discard these blank images, we input the collected images to SAM (Kirillov et al. 2023) to get object masks for the whole image. A threshold is set as the minimum object number. Those images whose object mask number is smaller than the threshold are regarded as blank images and are discarded from the dataset. In this way, we build a self-collected dataset to support the training process of the Room Expert to get room style representation. More details can be found in the Appendix.

**Unsupervised learning with constraints.** Due to the lack of room annotation in the Gibson dataset, we devise a Room Expert composed of a room style encoder and a room relation network trained using an unsupervised clustering algorithm with *must-link* and *cannot-link* to exploit the collected dataset and obtain room-style representation.

We observe that the Gibson training episodes from (Mezghan et al. 2022) have provided the level of difficulty depending on the distance between the start and target locations: *easy* (1.5-3m), *medium* (3-5m) and *hard* (5-10m). Intuitively, if the two locations are far apart (*hard*), they are most likely in different rooms. Based on this intuition, four rules of room relationship between two arbitrary images  $I_i$  and  $I_j$  are summarized and a distance refine matrix  $M$  with size equal to  $N \times N$  is pre-built where  $N$  is the number of all collected images:

- (1) If  $S_i \neq S_j$ , then  $I_i$  and  $I_j$  are not in the same room (*cannot-link*), set  $M_{i,j} = -1$ ;
- (2) If the two images are taken at the same location, then  $I_i$  and  $I_j$  are definitely in the same room (*must-link*), set  $M_{i,j} = 1$ ;
- (3) If  $E_i = E_j$  and  $Ed_i = Ed_j = \textit{easy}$ , then  $I_i$  and  $I_j$  are probably in the same room, set  $M_{i,j} = 0.5$ ;
- (4) If  $E_i = E_j$  and  $Ed_i = Ed_j = \textit{hard}$ ,  $I_i$  and  $I_j$  are probably in different rooms, set  $M_{i,j} = -0.5$ .

We build the Unsupervised Room Style Representation Learning based on the four rules above. The framework of Room Expert is illustrated in Figure 2 (a). Generally, a memory dictionary that contains the cluster feature representations is built and the contrastive loss and cross-entropy loss are utilized to train the Room Expert. Specifically, a standard ResNet-50 (He et al. 2016) pre-trained on ImageNet (Deng et al. 2009) is used as the backbone for the room-style encoder to extract feature vectors of all the room images. Based on these, we calculate the pair-wise distance

matrix  $D$  between feature vectors. Then we refine the distance matrix through the pre-built matrix  $M$  which serves as the constraints set for feature vectors. The refinement process can be defined as follows:

$$RefinedDistance = D - \gamma M, \quad (3)$$

where  $\gamma$  the refinement hyper-parameter. Based on the refined distance matrix, we adopt InfoMap (Rosvall and Bergstrom 2008) clustering algorithm to cluster similar features and assign pseudo labels. With the annotations, we could employ contrastive loss for feature encoder optimization. In this paper, we use the cluster-level contrastive loss (Dai et al. 2022), which is formulated by

$$L_{cluster} = -\log \frac{\exp(E_s(I_i) \cdot \phi_+ / \tau)}{\sum_{k=1}^K \exp(E_s(I_i) \cdot \phi_k / \tau)}, \quad (4)$$

where  $E_s$  represents the style encoder.  $K$  is the number of cluster representations and  $\phi_k$  denotes the cluster centroid defined by the mean feature vectors of each cluster.  $\phi_+$  is a cluster centre which shares the same label with  $I_i$ . Two different image features  $I_i$  and  $I_j$  are concatenated as input to the room relation network  $E_r$  to predict their relation about whether they are taken in the same room. We employ the cross-entropy loss as the relation prediction loss for relation network and style encoder training. The relation predict loss is defined by:

$$L_{pred} = -\sum_{n=1}^N y_i \cdot \log(E_r(E_s(I_i), E_s(I_j))) + (1 - y_i) \cdot \log(1 - E_r(E_s(I_i), E_s(I_j))), \quad (5)$$

where  $E_r$  denotes the relation network and  $y_i$  is the relation labels. We jointly adopt the contrastive loss and the relation prediction loss for the room expert training. In summary, the overall objective can be formulated as follows

$$L_{total} = L_{cluster} + \omega L_{pred}, \quad (6)$$

where  $\omega$  represents the hyper-parameter used to balance the two losses.

## Navigation Policy Learning

We follow FGPrompt-EF (Sun et al. 2024) to set up only one visual feature encoder. It concatenates the 3-channel RGB observation  $I_t$  with the goal image  $I_g$  on the channel dimension and takes the concatenated 6-channel image as the input of the visual feature encoder. We formulate the encoder output as follows:

$$v_{vis} = E_v(I_t \oplus I_g), \quad (7)$$

where  $E_v$  is the visual feature encoder and  $\oplus$  denotes the channel-wise concatenating.

In this section, we train the visual encoder and navigation policy conditioned on the pre-trained room expert. Two different fusion methods are designed and explored to enable the agent with spatial relation awareness. A naive solution to fuse the knowledge from the room expert is to directly fuse the room-style embedding from the room-style encoder with the visual feature  $v_{vis}$ . We call this *Implicit Fusion*. These

fused features are then fed into the navigation policy  $\pi$  to determine the action  $a_t$ . In this case, the fusion mechanism can be written as:

$$a_t = \pi(I_{fusion}(v_{vis}, E_s(I_t), E_s(I_g))) \quad (8)$$

Implicit fusion manner requires the agent to distinguish the room relation from room-style embeddings. It's more straightforward to directly give the agent the room relation between the observation and target images and this leads to *Explicit Fusion*. Specifically, the room-style embeddings of the observation and target images  $E_s(I_t)$ ,  $E_s(I_g)$  are firstly fed into the room relation network  $E_r$  to obtain the spatial relation, as illustrated in Figure 2 (b). The agent is trained to take actions considering this spatial relation. This process can be formulated as:

$$relation(I_g, I_t) = E_r(E_s(I_t), E_s(I_g)), \quad (9)$$

$$a_t = \pi(E_{fusion}(v_{vis}, relation)), \quad (10)$$

The explicit fusion manner is more direct for the navigation policy. More details can be found in the Appendix.

## Experiments

**Dataset and evaluation metric.** We conduct all of the experiments on the Habitat simulator (Savva et al. 2019; Szot et al. 2021). We train our agent on the Gibson dataset (Xia et al. 2018) with the dataset split provided by (Mezghan et al. 2022). The dataset provides diverse indoor scenes, consisting of 72 training scenes and 14 testing scenes. We test our agent on the Matterport 3D (Chang et al. 2017) and Habitat-Matterport 3D dataset (Ramakrishnan et al. 2021) to validate the cross-domain generalization ability of our agent.

For evaluation, we utilize the Success Rate (SR) and Success Weighted by Path Length (SPL) (Anderson et al. 2018a). SPL balances the efficiency and success rate by calculating the weighted sum of the ratio of the shortest navigation path length to the predicted path length. In an episode, the success distance is within 1m and the maximum steps are set to 500.

**Implementation details.** We follow the agent setting of ZER (Al-Halah, Ramakrishnan, and Grauman 2022). The height of agent is set to 1.5m and the radius is 0.1m. The agent has a single RGB sensor with a 90° FOV and 128×128 resolution. The action space consists of MOVE\_FORWARD by 0.25m, TURN\_LEFT, TURN\_RIGHT by 30° and STOP. For the pre-training stage, we use the Adam optimizer with weight decay 5e-4 and batch size 64 to train the style encoder and relation network for 20 epochs. We set the refinement hyper-parameter  $\gamma$  as an adaptive parameter. See the Appendix for the detailed calculation. We set the balance parameter  $\omega$  equal to 1. For navigation learning, we train our REGNav for 500M steps on 8×3090 GPUs. Other hyperparameters follow the ZER.

**Baseline.** We build our method on FGPrompt-EF (Sun et al. 2024), which involves an agent containing a ResNet-9 encoder for extracting visual features and a policy network composed of a 2-layer GRU (Chung et al. 2014).

**Comparison with SOTAs on Gibson.** We report the results averaged over 3 random seeds. (The variances are less than

Method	Reference	Backbone	Sensor	Memory	SPL $\uparrow$	SR $\uparrow$
ZER	CVPR22	ResNet-9	1RGB	$\times$	21.6%	29.2%
ZSON	NIPS22	ResNet-50	1RGB	$\times$	28.0%	36.9%
OVRL	ICLRW23	ResNet-50	1RGB	$\times$	27.0%	54.2%
OVRL-V2	arXiv23	ViT-Base	1RGB	$\times$	58.7%	82.0%
FGPrompt-MF	NeurIPS23	ResNet-9	1RGB	$\times$	62.1%	90.7%
FGPrompt-EF	NeurIPS23	ResNet-9	1RGB	$\times$	66.5%	90.4%
REGNav	This paper	ResNet-9	1RGB	$\times$	<b>67.1%</b>	<b>92.9%</b>

Table 1: Comparison with state-of-the-art methods without external memory on Gibson. 1RGB denotes that only the front RGB sensor is available for the agent and the observation type is one RGB image. All results of these methods are obtained from the overall test set on Gibson.

Method	Reference	Backbone	Sensor(s)	Memory	Easy		Medium		Hard	
					SPL $\uparrow$	SR $\uparrow$	SPL $\uparrow$	SR $\uparrow$	SPL $\uparrow$	SR $\uparrow$
VGM	ICCV21	ResNet-18	4RGB-D	$\checkmark$	79.6%	86.1%	68.2%	81.2%	45.6%	60.9%
TSGM	CoRL22	ResNet-18	4RGB-D	$\checkmark$	83.5%	91.1%	68.1%	82.0%	50.0%	70.3%
Mem-Aug	IROS22	ResNet-18	4RGB	$\checkmark$	63.0%	78.0%	57.0%	70.0%	48.0%	60.0%
MemoNav	CVPR24	ResNet-18	4RGB-D	$\checkmark$	-	-	-	-	57.9%	74.7%
REGNav	This paper	ResNet-9	1RGB	$\times$	<b>71.4%</b>	<b>97.5%</b>	<b>69.4%</b>	<b>95.4%</b>	<b>59.4%</b>	<b>87.1%</b>

Table 2: Comparison with state-of-the-art methods using memory on Gibson. 4RGB denotes that the agent takes a panoramic image from 4 RGB sensors as the observation type. 4RGB-D means that depth image can be used as additional input. The results are evaluated on the easy, medium and hard set of Gibson.

Method	MP3D		HM3D	
	SPL $\uparrow$	SR $\uparrow$	SPL $\uparrow$	SR $\uparrow$
Mem-Aug	3.9%	6.9%	1.9%	3.5%
ZER	10.8%	14.6%	6.3%	9.6%
FGPrompt-MF	44.3%	75.3%	38.8%	73.8%
FGPrompt-EF	48.8%	75.7%	42.1%	<b>75.2%</b>
REGNav	<b>50.2%</b>	<b>78.0%</b>	<b>44.0%</b>	<b>75.2%</b>

Table 3: Comparison of cross-domain evaluation on Matterport 3D (MP3D) and HabitatMatterport 3D (HM3D). All methods are trained on the Gibson and directly tested on these two unseen datasets without finetuning.

1e-3). As demonstrated in Table 1, we first compare our proposed methods with recent state-of-the-art image-goal navigation methods without additional external memory, which don’t use the agent’s depth or pose sensor. These methods includes ZER (Al-Halah, Ramakrishnan, and Grauman 2022), ZSON (Majumdar et al. 2022), OVRL (Yadav et al. 2023b), OVRL-V2 (Yadav et al. 2023a) and FGPrompt (Sun et al. 2024). Our REGNav shows a promising result with SPL = 67.1% and SR = 92.9% on the overall Gibson dataset.

We also provide several recent memory-based methods for comparison, including VGM (Kwon et al. 2021), TSGM (Kim et al. 2023), Mem-Aug (Mezghan et al. 2022) and MemoNav (Li et al. 2024). Mem-Aug categorized the test episodes of Gibson into three levels of difficulty We evaluate our REGNav on the corresponding set and report the results in Table 2. our proposed method illustrates superior performance, outperforming the memory-based meth-

Models	Components		MP3D-Accuracy $\uparrow$
	Clean Data	Refine Dist	
RE-1			57.0%+0.0%
RE-2	$\checkmark$		57.4%+0.4%
RE-3		$\checkmark$	57.9%+0.5%
RE-4	$\checkmark$	$\checkmark$	<b>58.4%+0.5%</b>

Table 4: Comparison of the room expert (RE) and its variants. *Clean Data* refers to the dataset cleaning before the model training. *Refine Dist* is using the constraints set to refine the feature distance matrix.

Datasets	Implicit fusion		Explicit Fusion	
	SPL $\uparrow$	SR $\uparrow$	SPL $\uparrow$	SR $\uparrow$
Gibson	47.4%	77.0%	<b>67.1%</b>	<b>92.9%</b>
MP3D	33.2%	59.9%	<b>50.2%</b>	<b>78.0%</b>
HM3D	27.7%	55.8%	<b>44.0%</b>	<b>75.2%</b>

Table 5: Ablation study of Fusion Manners.

ods by a large margin, which indicates the capacity of REGNav to effectively leverage the style information.

**Cross-domain evaluation.** To prove the domain generalization capability of our REGNav, we evaluate the Gibson-trained models on the Matterport 3D (MP3D) and HabitatMatterport 3D (HM3D) without extra finetuning. Cross-domain evaluation is an extremely challenging setting due to the visual domain gap between these datasets. Table 3 reports the comparison results. The results of Mem-Aug (Mezghan et al. 2022) and ZER (Al-Halah, Ramakr-

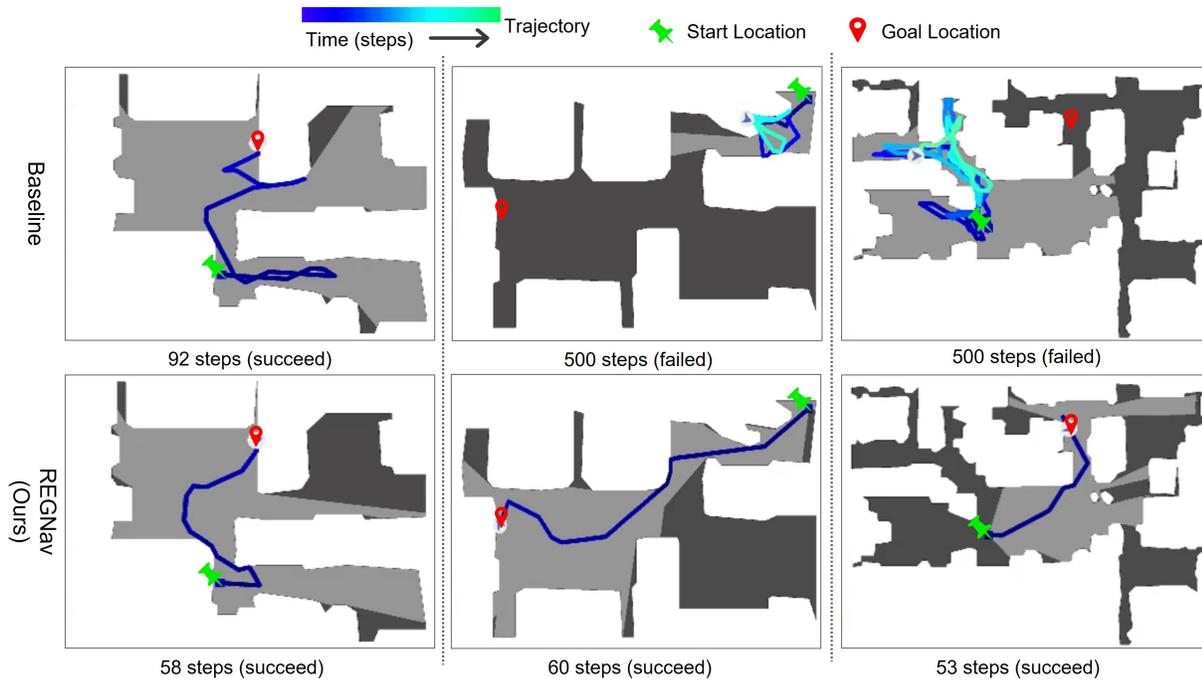


Figure 3: **The visualization results of example episodes from a top-down view.** The lines originating from the green locations refer to the agent’s trajectories, where the colour changes as the steps. The grey regions on the top-down map represent the areas explored by the agent’s camera. Compared with the baseline, our REGNav plans more efficient navigation paths.

ishnan, and Grauman 2022) are cited from their paper, while the results of FGPrompt (Sun et al. 2024) are evaluated from the trained models that FGPrompt released. Compared with previous methods, our REGNav achieves comparable performance on SPL and SR, which shows that focusing on spatial information can lead to better generalization.

**Ablation study on Room Expert training scheme.** We investigate the necessity of cleaning the dataset using SAM (Kirillov et al. 2023) and refining the feature distance matrix using *must-link* and *cannot-link* constraints set in the pre-training room expert stage. Due to the lack of pair annotations in Gibson, we follow the data collection technique in Gibson to collect a validation dataset in MP3D which has room labels. All the results are trained in the Gibson-collected dataset with unsupervised clustering and evaluated in the MP3D-collected validation set with real labels. We report the relation accuracy of input pair images as the evaluation metric. As shown in Table 4, using both data cleaning and distance refinement are superior to the counterparts, validating the effectiveness of these components.

**Comparison of different fusion manners.** We investigate two different fusion methods of incorporating the room-level information of visual observations into the semantic information. Implicit fusion refers to using the room-style embedding of the pre-trained model while explicit fusion is to directly use the room relation between the goal and observation. We report the comparison results in Table 5 and the more straightforward explicit fusion manner can obtain better performance. This validates that the direct room relation prior can empower the agent with more successful and effi-

cient navigation than the implicit representation.

**Visualization.** To qualitatively analyze the effect of our proposed method, we visualize the navigation results using top-down maps. We compare our REGNav with the baseline (FGPrompt-EF) for different scenes in the Gibson test set in Figure 3. When there exist certain discrepancies between the goal location and start location, due to the lack of spatial relationship priors, the agent of FGPrompt needs to take more steps and frequently wander around, especially in narrow pathways. In contrast, REGNav can analyze the spatial relationships and reason the relative goal location. Therefore, it can efficiently reduce redundant actions and achieve shorter navigation paths, which validates the superiority of REGNav in planning better paths. We also provide more visualization and analysis in the Appendix.

## Conclusion

In this paper, we introduced REGNav to address the issue of the agent’s meaningless actions for the ImageNav task. Our motivation draws on human navigation strategies, enabling agents to evaluate spatial relationships between goal and observation images through a pre-trained room expert model. This model uses unsupervised learning to extract room style representations, determining whether the current location belongs to the same room as the goal and guiding the navigation process. Our experimental results highlight REGNav’s superior performance in planning efficient navigation paths, particularly in complex environments where traditional models struggle with spatial discrepancies.

## Acknowledgments

This work was supported in part by National Natural Science Foundation of China under Grants 62088102 and 12326608, Natural Science Foundation of Shaanxi Province under Grant 2022JC-41, Key R&D Program of Shaanxi Province under Grant 2024PT-ZCK-80, and Fundamental Research Funds for the Central Universities under Grant XTR042021005.

## References

- Al-Halah, Z.; Ramakrishnan, S. K.; and Grauman, K. 2022. 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*, 17031–17041.
- Anderson, P.; Chang, A.; Chaplot, D. S.; Dosovitskiy, A.; Gupta, S.; Koltun, V.; Kosecka, J.; Malik, J.; Mottaghi, R.; Savva, M.; et al. 2018a. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*.
- Anderson, P.; Wu, Q.; Teney, D.; Bruce, J.; Johnson, M.; Sünderhauf, N.; Reid, I.; Gould, S.; and Van Den Hengel, A. 2018b. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3674–3683.
- Chang, A.; Dai, A.; Funkhouser, T.; Halber, M.; Niessner, M.; Savva, M.; Song, S.; Zeng, A.; and Zhang, Y. 2017. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*.
- Chaplot, D. S.; Gandhi, D. P.; Gupta, A.; and Salakhutdinov, R. R. 2020a. Object goal navigation using goal-oriented semantic exploration. *Advances in Neural Information Processing Systems*, 33: 4247–4258.
- Chaplot, D. S.; Salakhutdinov, R.; Gupta, A.; and Gupta, S. 2020b. Neural topological slam for visual navigation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12875–12884.
- Chen, S.; Guhur, P.-L.; Schmid, C.; and Laptev, I. 2021. History aware multimodal transformer for vision-and-language navigation. *Advances in neural information processing systems*, 34: 5834–5847.
- Chung, J.; Çağlar Gülçehre; Cho, K.; and Bengio, Y. 2014. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. *ArXiv*, abs/1412.3555.
- Dai, Z.; Wang, G.; Yuan, W.; Zhu, S.; and Tan, P. 2022. Cluster contrast for unsupervised person re-identification. In *Proceedings of the Asian conference on computer vision*, 1142–1160.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Du, H.; Li, L.; Huang, Z.; and Yu, X. 2023. Object-goal visual navigation via effective exploration of relations among historical navigation states. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2563–2573.
- Du, Y.; Gan, C.; and Isola, P. 2021. Curious representation learning for embodied intelligence. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10408–10417.
- Durrant-Whyte, H.; and Bailey, T. 2006. Simultaneous localization and mapping: part I. *IEEE robotics & automation magazine*, 13(2): 99–110.
- Hahn, M.; Chaplot, D. S.; Tulsiani, S.; Mukadam, M.; Rehg, J. M.; and Gupta, A. 2021. No rl, no simulation: Learning to navigate without navigating. *Advances in Neural Information Processing Systems*, 34: 26661–26673.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Kim, N.; Kwon, O.; Yoo, H.; Choi, Y.; Park, J.; and Oh, S. 2023. Topological semantic graph memory for image-goal navigation. In *Conference on Robot Learning*, 393–402. PMLR.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.
- Krantz, J.; Banerjee, S.; Zhu, W.; Corso, J.; Anderson, P.; Lee, S.; and Thomason, J. 2023a. Iterative vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14921–14930.
- Krantz, J.; Gervet, T.; Yadav, K.; Wang, A.; Paxton, C.; Mottaghi, R.; Batra, D.; Malik, J.; Lee, S.; and Chaplot, D. S. 2023b. Navigating to objects specified by images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10916–10925.
- Kwon, O.; Kim, N.; Choi, Y.; Yoo, H.; Park, J.; and Oh, S. 2021. Visual graph memory with unsupervised representation for visual navigation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 15890–15899.
- Kwon, O.; Park, J.; and Oh, S. 2023. Renderable neural radiance map for visual navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9099–9108.
- Lei, X.; Wang, M.; Zhou, W.; Li, L.; and Li, H. 2024. Instance-aware Exploration-Verification-Exploitation for Instance ImageGoal Navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16329–16339.
- Li, H.; Wang, Z.; Yang, X.; Yang, Y.; Mei, S.; and Zhang, Z. 2024. MemoNav: Working Memory Model for Visual Navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17913–17922.
- Li, J.; and Bansal, M. 2023. Improving vision-and-language navigation by generating future-view image semantics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10803–10812.

- Li, J.; Tan, H.; and Bansal, M. 2022. Envedit: Environment editing for vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15407–15417.
- Li, X.; Wang, Z.; Yang, J.; Wang, Y.; and Jiang, S. 2023. Kerm: Knowledge enhanced reasoning for vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2583–2592.
- Liu, X.; Paul, S.; Chatterjee, M.; and Cherian, A. 2024. CAVEN: An Embodied Conversational Agent for Efficient Audio-Visual Navigation in Noisy Environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 3765–3773.
- Majumdar, A.; Aggarwal, G.; Devnani, B.; Hoffman, J.; and Batra, D. 2022. Zson: Zero-shot object-goal navigation using multimodal goal embeddings. *Advances in Neural Information Processing Systems*, 35: 32340–32352.
- Mayo, B.; Hazan, T.; and Tal, A. 2021. Visual navigation with spatial attention. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16898–16907.
- Mezghan, L.; Sukhbaatar, S.; Lavril, T.; Maksymets, O.; Batra, D.; Bojanowski, P.; and Alahari, K. 2022. Memory-augmented reinforcement learning for image-goal navigation. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 3316–3323. IEEE.
- Paul, S.; Roy-Chowdhury, A.; and Cherian, A. 2022. Avlen: Audio-visual-language embodied navigation in 3d environments. *Advances in Neural Information Processing Systems*, 35: 6236–6249.
- Pelluri, N. 2024. Transformers for Image-Goal Navigation. *arXiv preprint arXiv:2405.14128*.
- Qiao, Y.; Qi, Y.; Hong, Y.; Yu, Z.; Wang, P.; and Wu, Q. 2022. Hop: History-and-order aware pre-training for vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15418–15427.
- Ramakrishnan, S. K.; Gokaslan, A.; Wijmans, E.; Maksymets, O.; Clegg, A.; Turner, J.; Undersander, E.; Galuba, W.; Westbury, A.; Chang, A. X.; et al. 2021. Habitat-matterport 3d dataset (hm3d): 1000 large-scale 3d environments for embodied ai. *arXiv preprint arXiv:2109.08238*.
- Rosvall, M.; and Bergstrom, C. T. 2008. Maps of random walks on complex networks reveal community structure. *Proceedings of the national academy of sciences*, 105(4): 1118–1123.
- Savva, M.; Kadian, A.; Maksymets, O.; Zhao, Y.; Wijmans, E.; Jain, B.; Straub, J.; Liu, J.; Koltun, V.; Malik, J.; et al. 2019. Habitat: A platform for embodied ai research. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9339–9347.
- Sun, X.; Chen, P.; Fan, J.; Chen, J.; Li, T.; and Tan, M. 2024. FGPrompt: fine-grained goal prompting for image-goal navigation. *Advances in Neural Information Processing Systems*, 36.
- Sun, X.; Liu, L.; Zhi, H.; Qiu, R.; and Liang, J. 2025. Prioritized semantic learning for zero-shot instance navigation. In *European Conference on Computer Vision*, 161–178. Springer.
- Szot, A.; Clegg, A.; Undersander, E.; Wijmans, E.; Zhao, Y.; Turner, J.; Maestre, N.; Mukadam, M.; Chaplot, D. S.; Maksymets, O.; et al. 2021. Habitat 2.0: Training home assistants to rearrange their habitat. *Advances in neural information processing systems*, 34: 251–266.
- Thomason, J.; Murray, M.; Cakmak, M.; and Zettlemoyer, L. 2020. Vision-and-dialog navigation. In *Conference on Robot Learning*, 394–406. PMLR.
- Wang, Z.; Li, J.; Hong, Y.; Wang, Y.; Wu, Q.; Bansal, M.; Gould, S.; Tan, H.; and Qiao, Y. 2023. Scaling data generation in vision-and-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12009–12020.
- Wang, Z.; Li, X.; Yang, J.; Liu, Y.; Hu, J.; Jiang, M.; and Jiang, S. 2024. Lookahead Exploration with Neural Radiance Representation for Continuous Vision-Language Navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13753–13762.
- Wasserman, J.; Yadav, K.; Chowdhary, G.; Gupta, A.; and Jain, U. 2023. Last-mile embodied visual navigation. In *Conference on Robot Learning*, 666–678. PMLR.
- Xia, F.; R. Zamir, A.; He, Z.-Y.; Sax, A.; Malik, J.; and Savarese, S. 2018. Gibson env: real-world perception for embodied agents. In *Computer Vision and Pattern Recognition (CVPR), 2018 IEEE Conference on*. IEEE.
- Yadav, K.; Majumdar, A.; Ramrakhya, R.; Yokoyama, N.; Baevski, A.; Kira, Z.; Maksymets, O.; and Batra, D. 2023a. Ovrl-v2: A simple state-of-art baseline for imagenav and objectnav. *arXiv preprint arXiv:2303.07798*.
- Yadav, K.; Ramrakhya, R.; Majumdar, A.; Berges, V.-P.; Kuhar, S.; Batra, D.; Baevski, A.; and Maksymets, O. 2023b. Offline visual representation learning for embodied navigation. In *Workshop on Reincarnating Reinforcement Learning at ICLR 2023*.
- Zhang, C.; Li, M.; Budvytis, I.; and Liwicki, S. 2024a. DiLoc: An Iterative Approach to Embodied Dialog Localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12585–12593.
- Zhang, S.; Yu, X.; Song, X.; Wang, X.; and Jiang, S. 2024b. Imagine Before Go: Self-Supervised Generative Map for Object Goal Navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16414–16425.
- Zhao, G.; Li, G.; Chen, W.; and Yu, Y. 2024. OVERNAV: Elevating Iterative Vision-and-Language Navigation with Open-Vocabulary Detection and StructurEd Representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16296–16306.
- Zhu, Y.; Mottaghi, R.; Kolve, E.; Lim, J. J.; Gupta, A.; Fei-Fei, L.; and Farhadi, A. 2017. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *2017 IEEE international conference on robotics and automation (ICRA)*, 3357–3364. IEEE.