



저작자표시-비영리-변경금지 2.0 대한민국

이용자는 아래의 조건을 따르는 경우에 한하여 자유롭게

- 이 저작물을 복제, 배포, 전송, 전시, 공연 및 방송할 수 있습니다.

다음과 같은 조건을 따라야 합니다:



저작자표시. 귀하는 원저작자를 표시하여야 합니다.



비영리. 귀하는 이 저작물을 영리 목적으로 이용할 수 없습니다.



변경금지. 귀하는 이 저작물을 개작, 변형 또는 가공할 수 없습니다.

- 귀하는, 이 저작물의 재이용이나 배포의 경우, 이 저작물에 적용된 이용허락조건을 명확하게 나타내어야 합니다.
- 저작권자로부터 별도의 허가를 받으면 이러한 조건들은 적용되지 않습니다.

저작권법에 따른 이용자의 권리는 위의 내용에 의하여 영향을 받지 않습니다.

이것은 [이용허락규약\(Legal Code\)](#)을 이해하기 쉽게 요약한 것입니다.

[Disclaimer](#)

M.S. THESIS

Meta-Explore: Exploratory Hierarchical
Vision-and-Language Navigation
Using Scene Object Spectrum Grounding

장면 객체 스펙트럼 기반의 계층적 탐색을 이용한
시각-언어 네비게이션

August 2023

DEPARTMENT OF ELECTRICAL AND COMPUTER
ENGINEERING
COLLEGE OF ENGINEERING
SEOUL NATIONAL UNIVERSITY

Minyoung Hwang

M.S. THESIS

Meta-Explore: Exploratory Hierarchical
Vision-and-Language Navigation
Using Scene Object Spectrum Grounding

장면 객체 스펙트럼 기반의 계층적 탐색을 이용한
시각-언어 네비게이션

August 2023

DEPARTMENT OF ELECTRICAL AND COMPUTER
ENGINEERING
COLLEGE OF ENGINEERING
SEOUL NATIONAL UNIVERSITY

Minyoung Hwang

Meta-Explore: Exploratory Hierarchical
Vision-and-Language Navigation
Using Scene Object Spectrum Grounding

장면 객체 스펙트럼 기반의 계층적 탐색을 이용한
시각-언어 네비게이션

지도교수 오 성 회

이 논문을 공학석사학위논문으로 제출함

2023 년 05 월

서울대학교 대학원

전기정보공학부

황 민 영

황 민 영의 공학석사 학위논문을 인준함

2023 년 05 월

위 원 장	<hr/>	최 진 영	(인)
부위원장	<hr/>	오 성 회	(인)
위 원	<hr/>	양 인 순	(인)

Abstract

Vision-and-Language Navigation (VLN) is an emerging field that focuses on developing intelligent agents capable of following natural-language instructions to navigate based on visual observations of the environment. The main challenge in VLN is how to understand natural-language instructions in an unseen environment. The limitation of conventional VLN algorithms is that if an action is mistaken, the agent fails to follow the instructions or explores unnecessary regions, leading the agent to an irrecoverable path. This thesis presents Meta-Explore, a hierarchical navigation method deploying an exploitation policy to correct misled recent actions. We show that an exploitation policy, which moves the agent toward a well-chosen local goal among unvisited but observable states, outperforms a method which moves the agent to a previously visited state. We also highlight the demand for imagining regretful explorations with semantically meaningful clues. The key to our approach is understanding the object placements around the agent in spectral-domain. Specifically, we present a novel visual representation, called scene object spectrum (SOS), which performs category-wise 2D Fourier transform of detected objects. Combining exploitation policy and SOS features, the agent can correct its path by choosing a promising local goal. We evaluate our method in three VLN benchmarks: R2R, SOON, and REVERIE. Meta-Explore outperforms other baselines and shows significant generalization performance. In addition, local goal search using the proposed spectral-domain SOS features significantly improves the success rate by 17.1% and SPL by 20.6% against the state-of-the-art method of the SOON benchmark.

Keywords: Visual Navigation, Vision-and-Language Navigation, Computer Vision, Object Detection, Hierarchical Exploration, Fourier Transform, Cross-Modal Transformer

Student Number: 2021-24215

To my parents.

Contents

Abstract	i
Chapter 1 Introduction	1
1.1 Study Background	1
1.2 Purpose of Research	2
Chapter 2 Meta-Explore: Exploratory Hierarchical Vision-and-Language Navigation Using Scene Object Spectrum Grounding	3
2.1 Related Work	7
2.1.1 Vision-and-Language Navigation	7
2.1.2 Exploration-Exploitation	7
2.1.3 Visual Representations	8
2.2 Method	9
2.2.1 Problem Formulation	9
2.2.2 Meta-Explore	10
2.2.3 Training Details	18
2.3 Navigation Experiments	19
2.3.1 Experiment Settings	19

2.3.2	Evaluation Metrics	20
2.3.3	Baselines and Implementation Details	20
2.3.4	Comparison with Navigation Baselines	21
2.3.5	Local Goal Search using SOS Features	23
2.3.6	Ablation Study	25
2.4	Discussion	25
Appendix A Appendices		26
A.1	Model Details	26
A.1.1	Algorithm Details	26
A.1.2	Exploitation Module	26
A.1.3	Implementation Details	30
A.2	Experiment Setup	31
A.2.1	Dataset Statistics	31
A.2.2	Data Preprocessing	31
A.2.3	Baselines	32
A.3	Navigation Experiments	33
A.3.1	Detailed Analyses in R2R	33
A.3.2	Detailed Analyses in SOON	34
A.3.3	Evaluation Results in REVERIE benchmark	35
A.3.4	Local Goal Search	38
A.4	Ablation Study	38
A.4.1	Language-triggered Hierarchical Exploration	38
A.4.2	Image-Goal Navigation in Continuous Domain	41
A.4.3	VLN in Continuous Domain	46
요약		60

List of Figures

Figure 2.1	Hierarchical Exploration. At each episode, a natural-language instruction is given to the agent to navigate to a goal location. The agent explores the environment and constructs a topological map by recording visited nodes ● and next step reachable nodes ✱. Each node consists of the position of the agent and visual features. o_t denotes the observation at time t . The agent chooses an unvisited local goal to solve the regretful exploration problem.	4
Figure 2.2	Local Goal Search for Exploitation. The local goal is likely to be chosen as the closest node to the global goal. Existing methods only backtrack to a visited node (left). We expand the searchable area by including unvisited but reachable nodes (right).	5

Figure 2.3	Network Architecture.	Three types of visual features: panoramic (yellow), object image (aquamarine), and object spectrum (red) are encoded. The color in each parenthesis denotes the color describing the corresponding feature. The cross-modal transformer encodes language and spatial visual features as hidden state H_t . A mode selector gives <i>explore</i> or <i>exploit</i> command to the agent by predicting the explore probability $P_{explore}$. The selected navigation module outputs an action a_t from the possible n_{cand} candidate nodes.	10
Figure 2.4	Navigation Modules.	Mode selector estimates $P_{explore}$, i.e., the probability to explore, and chooses between exploration and exploitation modules. The selected navigation module outputs the next action a_t	12
Figure 2.5	Scene Object Spectrum (SOS).	The agent calculates <i>scene object spectrum</i> (SOS) features for efficient exploitation. SOS features incorporate semantic information observed in a single panoramic image by performing category-wise 2D FFT. . . .	15
Figure 2.6	Toy Example.	Monotonic alignment between language instruction and visual observation is desirable. Yellow dots \bullet in the nodes describe the ground truth trajectory. Based on the node at $t = 3$, the similarity matrix can show either monotonic or non-monotonic alignment between object tokens and SOS features. The green circles \oplus describe the possible candidates A, B for next action.	18

Figure A.1	Bounding box coordinate transformation. Front-view visual observations from different angles at the same location. Each bounding box shows the ‘chair’ detection. We use coordinate transformation to convert coordinates into panoramic view.	28
Figure A.2	Bounding box statistics. We collect width, height, and size of detected bounding boxes. The histograms show statistics for ‘chair’ objects. Yellow line and red line show the median and average values of each distribution, respectively.	28
Figure A.3	Relationship between navigation score (S_{nav}) and normalized distance sum (nDS) in R2R. We measure navigation scores for augmented trajectories which include both successful and failed trajectories. R and Q illustrate an example case of a ground truth trajectory and a query trajectory. Maximum hop of a query trajectory is 15. Trajectories with high nDS scores also have high navigation scores.	29
Figure A.4	Local goal search scenarios in R2R. Ground truth trajectory (orange) and current trajectory at time $t = 6$ (blue) are shown in the left. Traj. denotes trajectory. The number next to each node denotes the navigation score S_{nav} of the shortest path trajectory from the start node to the corresponding node. If the local goal is chosen from the previously visited nodes, the local goal becomes the node with $S_{nav} = 0.11$. If the local goal is chosen from the unvisited but observed nodes, the local goal becomes the node with $S_{nav} = 0.22$	37

Figure A.5	Sample image captions. (a), (b), (c), and (d) show captions that successfully describe the scenes. (e) and (f) shows failure cases of caption generation. For successful language-triggered hierarchical exploration, image captions should correctly describe the scenes. However, current image captioning methods often generates misdescribed captions, thereby leading to a low navigation performance.	39
Figure A.6	Exploitation by searching a local goal. In the exploit mode, the agent aims to escape from the stranded local area. It first searches for the most similar node to the goal. Then, it finds an optimal local goal which is unexplored and also similar to the goal image. We use SOS features to compare with the target image.	41
Figure A.7	Experiment visualization for image-goal navigation task in continuous environment. The mode selector detects stuck event at $t = 76$ and switches the explore mode to exploit mode. Then, the agent returns toward the local goal, which is chosen as a position nearby one of the nodes in the previously constructed graph.	45

List of Tables

Table 2.1	Comparison and evaluation results of the baselines and our model in the R2R Navigation Task. Gray shaded rows describe hierarchical navigation baselines. Three memory types: Rec (recurrent), Seq (sequential), and Top. Map (topological map)	21
Table 2.2	Comparison and evaluation results of the baselines and our model in the SOON Navigation Task.	22
Table 2.3	Comparison of Exploitation Policies. (R2R)	24
Table 2.4	Comparison of Exploitation Policies. (SOON)	24
Table A.1	Object Parsing Examples. For each dataset, object tokens are extracted from the instructions. Target objects are inferred from the instructions using VQA. Words that have similar meanings are unified into a single object word for categorization.	31

Table A.2	Comparison and evaluation results of the baselines and our model in REVERIE Navigation Task. Gray shaded rows describe hierarchical navigation baselines. Three memory types: Rec (recurrent), Seq (sequential), and Top. Map (topological map)	36
Table A.3	Comparison of Exploitation Policies. (REVERIE)	37
Table A.4	Comparison and evaluation results of the local goal search methods using different target and candidate domains. (R2R)	40
Table A.5	Evaluation results for Image-goal Navigation Task. (SR: success rate, SPL: success weighted by path length)	44
Table A.6	Evaluation results in the VLN-CE val unseen split.	46

Acknowledgements

I would like to express my sincere gratitude to my advisor, Songhwai Oh, for his invaluable guidance and support throughout my master's program. His expertise and unwavering encouragement helped me to complete this research and write this thesis. Additionally, I extend my gratitude to Yoonseon Oh and Kyungjae Lee, who served as co-advisors during my master's program. Their consistent interest in my research and specific feedback helped me to develop essential skills as a researcher. I would also like to thank all the participants in my study for their time and willingness to share their experiences. This work would not have been possible without their contribution. Lastly, I am deeply grateful to my cherished friends and family, whose unwavering support and diverse perspectives have continuously inspired and encouraged me throughout this journey.

Chapter 1

Introduction

1.1 Study Background

With the recent progress in artificial intelligence, the need for developing an agent that can understand natural human language and act intelligently has risen [1]. Such an agent can help humans with various tasks: from daily tasks such as household chores to tasks that involve high risk, such as working in hazardous regions. This motivates vision-and-language navigation (VLN) [2–4], which is an emerging field of research that focuses on developing an embodied agent that follows natural language instructions to navigate in unseen environments to arrive at the target location. Conventional VLN algorithms, however, have a critical drawback that it is hard to handle mistaken actions effectively. When an agent makes a wrong action and deviates from the intended path, this may lead to failure in following the instructions and potentially getting stuck in an irrecoverable path.

1.2 Purpose of Research

This research aims to develop a navigation method that enables the agent to recover mistaken actions in VLN. In Chapter 2, we propose a hierarchical navigation method that solves the challenges in conventional VLN algorithms and enhances the generalization performance so that the agent can navigate well in unseen environments. We focus on understanding the object placements around the agent and finding a promising exploitation policy that makes the agent correct its path based on semantically meaningful clues.

Chapter 2

Meta-Explore: Exploratory Hierarchical Vision-and-Language Navigation Using Scene Object Spectrum Grounding

Visual navigation in indoor environments has been studied widely and shown that an agent can navigate in unexplored environments [5]. By recognizing the visual context and constructing a map, an agent can explore the environment and solve tasks such as moving towards a goal or following a desired trajectory. With the increasing development in human language understanding, vision-and-language navigation (VLN) [2] has enabled robots to communicate with humans using natural languages. The high degree of freedom in natural language instructions allows VLN to expand to various tasks, including (1) following fine-grained step-by-step instructions [2, 6–16] and (2) reaching a target location described by goal-oriented language instructions [17–23].

A challenging issue in VLN is the case when an action is mistaken with respect to the given language instruction [24–29]. For instance, if the agent is asked to turn right at the end of the hallway but turns left, the agent may end

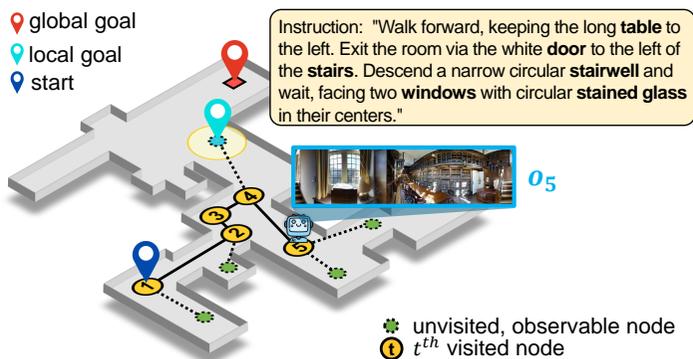


Figure 2.1: **Hierarchical Exploration.** At each episode, a natural-language instruction is given to the agent to navigate to a goal location. The agent explores the environment and constructs a topological map by recording visited nodes \circ and next step reachable nodes $*$. Each node consists of the position of the agent and visual features. o_t denotes the observation at time t . The agent chooses an unvisited local goal to solve the regretful exploration problem.

up in irrecoverable paths. Several existing studies solve this issue via hierarchical exploration, where the high-level planner decides when to explore and the low-level planner chooses what actions to take. If the high-level planner chooses to explore, the agent searches unexplored regions, and if it chooses to exploit, the agent executes the best action based on the previous exploration. Prior work [24–26] returns the agent to the last successful state and resumes exploration. However, such methods take a heuristic approach because the agent only backtracks to a recently visited location. The agent does not take advantage of the constructed map and instead naively uses its recent trajectory for backtracking. Another recent work [29] suggests graph-based exploitation, which uses a topological map to expand the action space in global planning. Still, this method assumes that the agent can directly jump to a previously visited node. Since this method can perform a jump action at every timestep, there is no trigger that explicitly decides when to explore and when to exploit. Therefore, we address the importance of time scheduling for exploration-exploitation and

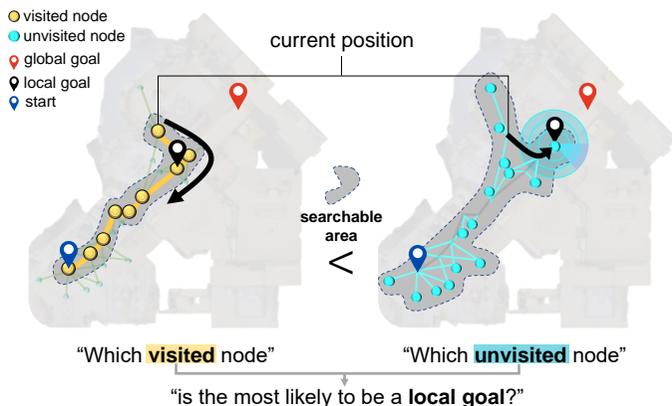


Figure 2.2: **Local Goal Search for Exploitation.** The local goal is likely to be chosen as the closest node to the global goal. Existing methods only backtrack to a visited node (left). We expand the searchable area by including unvisited but reachable nodes (right).

efficient global planning using a topological map to avoid reexploring visited regions.

We expand the notion of hierarchical exploration by proposing Meta-Explore, which not only allows the high-level planner to choose when to correct misled local movements but also finds an unvisited state inferred to be close to the global goal. We illustrate the overview of hierarchical exploration in Figure 2.1. Instead of backtracking, we present an exploitation method called local goal search. We show that it is more efficient to plan a path to a local goal, which is the most promising node from the unvisited but reachable nodes. We illustrate the difference between conventional backtracking and local goal search in Figure 2.2. Based on our method, we show that exploration and exploitation are not independent and can complement each other: (1) to overtake regretful explorations, the agent can perform exploitation and (2) the agent can utilize the constructed topological map for local goal search. We also highlight the demand for imagining regretful explorations with semantically meaningful clues. Most VLN tasks

require a level of understanding objects nearby the agent, but previous studies simply encode observed panoramic or object images [2, 4, 6, 19–21, 24–37]. In this paper, we present a novel semantic representation of the scene called *scene object spectrum* (SOS), which is a matrix containing the arrangements and frequencies of objects from the visual observation at each location. Using SOS features, we can sufficiently estimate the context of the environment. We show that the proposed spectral-domain SOS features manifest better linguistic interpretability than conventional spatial-domain visual features. Combining exploitation policy and SOS features, we design a navigation score that measures the alignment between a given language instruction and a corrected trajectory toward a local goal. The agent compares local goal candidates and selects a near-optimal candidate with the highest navigation score from corrected trajectories. This involves high-level reasoning related to the landmarks (e.g., bedroom and kitchen) and objects (e.g., table and window) that appear in the instructions.

The main contributions of this paper are as follows:

- We propose a hierarchical navigation method called Meta-Explore, deploying an exploitation policy to correct misled recent actions. The agent searches for an appropriate local goal instead of reversing the recent action sequence.
- In the exploitation mode, the agent uses a novel scene representation called *scene object spectrum* (SOS), which contains the spectral information of the object placements in the scene. SOS features provide semantically meaningful clues to choose a near-optimal local goal and help the agent to solve the regretful exploration problem.
- We evaluate our method on three VLN benchmarks: R2R [2], SOON [19],

and REVERIE [20]. The experimental results show that the proposed method, Meta-Explore, improves the success rate and SPL in test splits of R2R, SOON and val split of REVERIE. The proposed method shows better generalization results compared to all baselines.

2.1 Related Work

2.1.1 Vision-and-Language Navigation

In VLN, an agent encodes the natural language instructions and follows the instructions, which can be either (1) a fine-grained step-by-step instruction the agent can follow [2, 6, 7], (2) a description of the target object and location [19, 20], or (3) additional guidance given to the agent [4, 21]. These tasks require the agent to recognize its current location using some words in the natural-language instructions. Prior work [2, 30–33] show that an agent can align visual features to language instructions via neural networks and use the multimodal output embeddings to generate a suitable action at each timestep. Most VLN methods utilize cross-modal attention, either with recurrent neural networks [2, 30] or with transformer-based architectures [31–33]. For sequential action prediction, Hong *et al.* [34] further use recurrent units inside transformer architectures, while Pashevich *et al.* [35] and Chen *et al.* [36] use additional transformers to embed past observations and actions.

2.1.2 Exploration-Exploitation

In an unseen environment, the agent must maximize the return without knowing the true value functions. One of the solutions to this problem is to switch back and forth between exploration and exploitation [38]. In the exploration mode, the agent gathers more information about the environment. On the other hand, the agent uses information collected during exploration and chooses the

best action for exploitation. Ecoffet *et al.* [39] reduced the exploration step by archiving the states and exploring again from the successful states. Pislár *et al.* [40] addressed the various scheduling policies and demonstrated their method on Atari games. Recent work [41,42] successfully demonstrates the effectiveness of hierarchical exploration in image-goal navigation.

Like commonly used greedy navigation policies, VLN tasks also deal with the problem of maximizing the chance to reach the goal without knowing the ground truth map. Several VLN methods employ the concept of exploitation to tackle this problem. Ke *et al.* [37] look forward to several possible future trajectories and decide whether to backtrack or not and where to backtrack. Others [24–26] estimate the progress to tell whether the agent becomes lost and make the agent backtrack to a previously visited location to restart exploration. However, previous studies do not take into account what should be done in the exploitation mode. In order to handle this problem, we propose a hierarchical navigation method which determines the scheduling between exploration and exploitation.

2.1.3 Visual Representations

Popular visual encoding methods via ResNet [43] and ViT [44] can be trained to learn rotation-invariant visual features. Both methods learn to extract visual features with high information gain for global and local spatial information. The high complexity of the features leads to low interpretability of the scene and therefore requires the agent to use additional neural networks or complex processing to utilize them. On the other hand, traditional visual representation methods such as Fourier transform use spectral analysis, which is highly interpretable and computationally efficient. One drawback of the traditional methods is that they fail to maximize the information gain. Nonetheless, an

appropriate use of essential information can be helpful for high-level decision making and enables more straightforward interpretation and prediction of the visual features. One traditional navigation method, Sturz *et al.* [45] used Fourier transform to generate rotation-invariant visual features. However, no research has transformed the spectral information of the detected objects to represent high-level semantics from visual observations. Focusing on the fact that 2D Fourier transform can extract morphological properties of images [46], we can find out the shape or structure of detected objects through 2D Fourier transform. In this paper, we decompose the object mask into binary masks by object categories and perform a 2D Fourier transform on each binary mask.

2.2 Method

2.2.1 Problem Formulation

We deal with VLN in discrete environments, where the environment is given as an undirected graph $G_e = \{V, E\}$. V denotes a set of N navigable nodes, $\{V_i\}_{i=1}^N$, and E is the adjacency matrix describing connectivity among the nodes in V . We denote the observation at node V_i as O_i . The agent uses a panoramic RGB image observation o_t and current node v_t , which are collected at time t . The agent either moves to a neighboring node or executes a **stop** action. a_t denotes the action at time t . The objectives of VLN are categorized as follows: (1) to follow language instructions [2] and (2) to find a target object described by language instructions in a fixed time T [19, 20]. We present a general hierarchical exploration method that can be applied to both tasks. We also enhance the navigation policy by extracting cross-domain visual representations from the environments, i.e., spatial-domain and spectral-domain representations. To balance the information loss and interpretability of the visual feature, we adopt multi-channel fast Fourier transform (FFT) to encode semantic masks of the

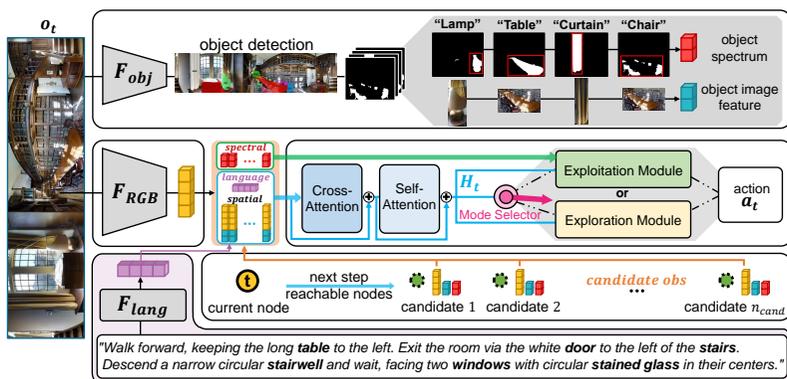


Figure 2.3: **Network Architecture.** Three types of visual features: panoramic (yellow), object image (aquamarine), and object spectrum (red) are encoded. The color in each parenthesis denotes the color describing the corresponding feature. The cross-modal transformer encodes language and spatial visual features as hidden state H_t . A mode selector gives *explore* or *exploit* command to the agent by predicting the explore probability $P_{explore}$. The selected navigation module outputs an action a_t from the possible n_{cand} candidate nodes.

detected objects into category-wise spectral-domain features.

2.2.2 Meta-Explore

We design a learnable hierarchical exploration method for VLN called Meta-Explore, which decides (1) when to *explore* or *exploit* and (2) a new imagined local goal to seek during exploitation. The overall network architecture of the proposed Meta-Explore is shown in Figure 2.3. Given a language instruction L , the agent navigates in the environment until it finds the target described in L . Meta-Explore consists of a mode selector and two navigation modules corresponding to two modes: exploration and exploitation. At each timestep, the mode selector chooses to explore or exploit. At $t = 0$, the mode is initialized to exploration. In the *exploration mode*, the agent outputs an action toward a neighboring node to move the agent toward the goal. When the mode selector recognizes that the agent is not following the instruction successfully, the mode

is switched to exploitation. In the *exploitation mode*, the agent seeks a new *local goal* with the highest correspondence against the language instructions from the previously unvisited candidate nodes using spectral-domain visual features. The agent moves toward the local goal by planning a path. After the agent arrives at the local goal, the mode is reset to exploration. The explore-exploit switching decision occurs through the mode selector by estimating the probability to explore. The agent repeats this explore-exploit behavior until it determines that the target is found and decides to stop.

Mode Selector

At time t , the agent observes visual features about the current node v_t and several reachable nodes. We call the nodes reachable at the current timestep as candidate nodes. n_{cand} denotes the number of candidate nodes. We use a cross-modal transformer with n_L layers to relate visual observations to language instructions. The cross-modal transformer takes the visual features of nodes in the constructed topological map at time t , G_t , and outputs cross-modal embedding H_t to encode visual observations with L . We concatenate location encoding and history encoding [27] to the visual features as node features to consider the relative pose from v_t and the last visited timestep of each node, respectively. Each word is encoded via a pretrained language encoder [47], which is used for general vision-language tasks.

The cross-modal transformer consists of cross-attention layer

$$\text{L2V_Attn}(\hat{W}, \hat{V}) = \text{Softmax}(\hat{W}\Theta_q^\dagger(\hat{V}\Theta_k^\dagger)^T/\sqrt{d})\hat{V}\Theta_v^\dagger, \quad (2.1)$$

and self-attention layer

$$\text{SelfAttn}(X) = \text{Softmax} = ((X\Theta_q(X\Theta_k)^T + A\Theta_e + b_e)/\sqrt{d})X\Theta_v, \quad (2.2)$$

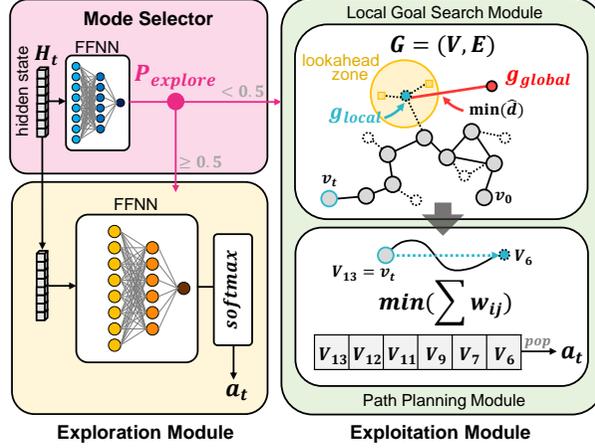


Figure 2.4: **Navigation Modules.** Mode selector estimates $P_{explore}$, i.e., the probability to explore, and chooses between exploration and exploitation modules. The selected navigation module outputs the next action a_t .

where \hat{W} , \hat{V} , X , and A denote word, visual, node representations and adjacency matrix of G_t , respectively. The (query, key, value) weight matrices of self-attention and cross-attention layers are denoted as $(\Theta_q, \Theta_k, \Theta_v)$ and $(\Theta_q^\dagger, \Theta_k^\dagger, \Theta_v^\dagger)$, respectively. The final cross-modal embedding at time t after passing through n_L transformer layers is denoted as H_t . To encourage the monotonic increasing relationship between language and visual attentions at each timestep, we define a correlation loss $L_{corr} = \sum_{t=1}^T \|\text{L2V_Attn} - I_{n_x}\|_1$ for training the cross-modal transformer, where n_x denotes the dimension of the H_t and I_{n_x} denotes an identity matrix of size $n_x \times n_x$.

As illustrated in Figure 2.4, the mode selector estimates the probability to explore $P_{explore}$ given the cross-modal hidden state H_t . We denote the mode selector as S_{mode} and use a two-layer feed-forward neural network. Given H_t , S_{mode} outputs the exploration probability as $P_{explore} = 1 - S_{mode}(H_t)$. If $P_{explore}$ is greater than 0.5, the exploration policy outputs a probability distribution for reachable nodes at the next step. At time $t+1$, the agent moves to the node with

the highest probability. If $P_{explore} < 0.5$, the agent determines that the current trajectory is regretful, so the agent should traverse to find a local goal, which is the most likely to be the closest node to the global goal. The exploitation policy mainly utilizes object-level features to search for the local goal with high-level reasoning. After the local goal is chosen, the path planning module outputs an action following the shortest path to the local goal.

To train the mode selector, we require additional demonstration data other than the ground truth trajectory, such that it switches between exploration and exploitation. We generate the demonstration data from the ground truth trajectories, with additional detours. For the detours, we stochastically select candidate nodes other than the ground truth paths and add the trajectory that returns to the current viewpoint. The imitation learning loss for training the mode selector is defined as $L_{mode} = \sum_{t=1}^T \mathbb{1}(m_t \neq gt_t)$, where m_t is the mode of the agent, 0 for exploitation and 1 for exploration. gt_t is 1 if the current node is in the shortest ground truth trajectory and $gt_t = 0$, otherwise.

Exploration Module

In the exploration mode, the agent follows the following sequential operations: topological map construction, self-monitoring, and an exploration policy. To improve the exploration, we adopt self-monitoring [24] to predict the current progress of exploration to enhance the exploration policy itself. Prior work [24, 25] has shown that auxiliary loss using self-monitoring can regularize the exploration policy.

Topological Map Construction. The agent constructs graph G_t by classifying nodes into two types: (1) visited nodes and (2) unvisited but observable nodes. At current time t , the agent at node $v_t \in \{V_i\}_{i=1}^N$ observes $N(v_t)$ neigh-

bor nodes as next step candidates at time $t + 1$. The visited nodes consist of visual features of their own and the neighboring nodes from panoramic RGB observations. The unvisited nodes can be observed only if they are connected to at least one visited node. The topological map records the positions and visual features of observed nodes at each timestep. By knowing the positions of nodes in G_t , the agent can plan the shortest path trajectory between two nodes.

Self-Monitoring. We use a progress monitor to estimate the current navigation progress at each episode. Self-monitoring via estimating current progress helps the agent choose the next action that can increase the progress. The estimated progress $\hat{p}_t = F_{progress}(H_t)$ is the output of a feed-forward neural network, given H_t as input. We measure the ground truth progress p_t as the ratio between the current distance to the goal and the shortest path length of the episode subtracted from 1, described as $1 - \frac{d_{geo}(v_t, v_{goal})}{d_{geo}(v_0, v_{goal})}$, where $d_{geo}(a, b)$ is the geodesic distance between a and b . v_0 , v_t , and v_{goal} denote initial, current, and goal positions, respectively. We add progress loss $L_{progress} = \sum_{t=1}^T (\hat{p}_t - p_t)^2$ to train the progress monitor while training the exploration policy.

Exploration Policy. The exploration policy $F_{explore}$ estimates the probability of moving to the candidate nodes at the next step. The agent chooses the action a_t at time t based on the estimated probability distribution among candidate nodes, described as $a_t = \arg \max_{V_i} (F_{explore}([H_t]_i))$. $F_{explore}$ is implemented via a two-layer feed-forward network with the cross-modal hidden state H_t given as input. The output of $F_{explore}$ becomes a probability distribution over possible actions. To only consider unvisited nodes, we mask out the output for visited nodes. For training, we sample the next action from the probability distribution instead of choosing a node with the highest probability. We describe the training details in Section 2.2.3.

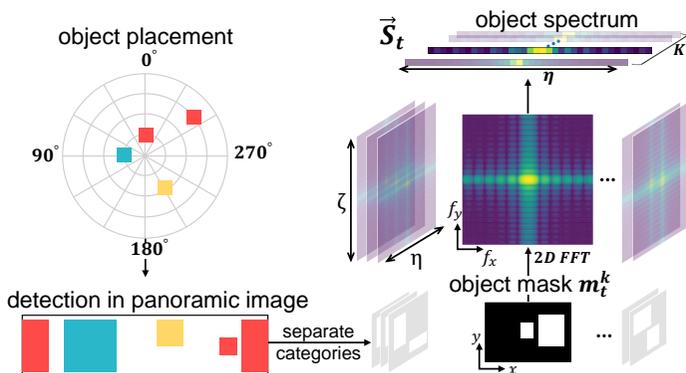


Figure 2.5: **Scene Object Spectrum (SOS)**. The agent calculates *scene object spectrum* (SOS) features for efficient exploitation. SOS features incorporate semantic information observed in a single panoramic image by performing category-wise 2D FFT.

Exploitation Module

In the exploitation mode, the agent requires high-level reasoning with identifiable environmental clues to imagine regretful exploration cases. To find clues in an object-level manner, we present a novel visual representation by capturing object information in the spectral-domain. The novel representation is more easily predictable than spatial features such as RGB image embeddings. The agent can take advantage of the predictability by expanding the searchable area to find a local goal. We choose the local goal as the closest node to the global goal in the feature space.

Spectral-Domain Visual Representations. Common navigation policies can lead the agent toward the node with the highest similarity to the target. However, even with a good learned policy, the agent can act in a novice manner in unseen environments. In this paper, we seek extra information from the environment for generalizable high-level reasoning to resolve the issue. As illustrated in Figure 2.5, *scene object spectrum* (SOS) incorporates semantic information observed in a single panoramic image by generating a semantic mask for each

object category and applying Fourier transform to each semantic mask. The semantic mask for object class k at time t is calculated as a binary mask $[m_t^k]_{ij}$ that detects the object at pixel (i, j) . Suppose there are a total of K object categories. When multiple objects are detected for one object category, the binary mask appears as a union of the bounding boxes of the detected objects. We define **FFT** as a channel-wise 2D fast Fourier transform that receives K binary semantic masks and outputs K spectral-domain features, where K is the number of object classes. Then, SOS feature $\vec{S}_t = [s_t^1, \dots, s_t^K]^T$ can be defined as $s_t^k = \log |\mathbf{FFT}(m_t^k)|$. For simplicity, we perform mean pooling on the vertical spectral axis and normalize the output. The final SOS feature has shape (K, η) , where η is the maximum horizontal frequency.

Local Goal Search Using Semantic Clues. We argue that returning to a previously visited node does not guarantee the agent escapes from the local optima. Instead of backtracking to a previously visited node, the agent searches for a local goal to move towards. If the agent plans a path and moves towards the local goal, the agent does not need to repeat unnecessary actions in visited regions after the exploitation ends. Additionally, searching for a local goal takes full advantage of the topological map by utilizing the connections among the observed nodes. To expand the searchable area further, we let the agent choose the local goal from previously unvisited and unchosen candidate nodes.

To choose a local goal, we first score the corrected trajectories to measure the alignment with the language instruction L . We use SOS features as semantic environmental clues to estimate the navigation score S_{nav} of the corrected trajectory, which is the shortest path trajectory from the initial node to the local goal in the constructed topological map. To simplify, we convert the language instruction into a list of objects $W^o = [w_1^o, \dots, w_B^o]$ consisting of $B[\leq K]$ object categories (e.g., desk, cabinet, and microwave). We approximate the cor-

responding reference SOS features as $[\hat{\delta}(w_1^o), \dots, \hat{\delta}(w_B^o)]$ where the i^{th} row of $\hat{\delta}(w_k^o)$ is defined as follows:

$$[\hat{\delta}(w_k^o)]_i = \mathbb{1}(k = i)\lambda(\delta(w_k^o))\text{sinc}\left(\frac{j}{2} - \frac{\eta}{4}\right). \quad (2.3)$$

$\lambda(\delta(w_k^o))$ denotes the average width of detected bounding boxes of object w_k^o in the environment. A detailed approximation process is explained in the supplementary material. To simulate a corrected trajectory $\mathcal{T}' = (v'_1, \dots, v'_{t'})$, we calculate the SOS features $[\vec{S}'_1, \dots, \vec{S}'_{t'}]$ corresponding to the nodes in \mathcal{T}' . We measure the similarity between two object spectrum features via the cosine similarity of the flattened vectors. Finally, the navigation score S_{nav} of \mathcal{T}' is computed as:

$$S_{nav}(\mathcal{T}') = \frac{\sum_{i=1}^B \sum_{j=1}^{t'} \left(\frac{\hat{\delta}(w_i^o)}{|\hat{\delta}(w_i^o)|} \cdot \frac{\vec{S}'_j}{|\vec{S}'_j|} \right) \left((\hat{\delta}(w_i^o) - \hat{\delta}(\bar{w}^o)) \cdot (\vec{S}'_j - \bar{\vec{S}}') \right)}{\sqrt{\frac{t'}{B} \cdot \sum_{i=1}^B (\hat{\delta}(w_i^o) - \hat{\delta}(\bar{w}^o))^2 \sum_{j=1}^{t'} (\vec{S}'_j - \bar{\vec{S}}')^2}}, \quad (2.4)$$

where $\hat{\delta}(\bar{w}^o)$ and $\bar{\vec{S}}'$ denote the average values of SOS features $\hat{\delta}(w_i^o)$ and \vec{S}'_j , respectively. This equation can also be interpreted as a pseudo correlation-coefficient function between object list W^o and trajectory \mathcal{T}' . The exploitation policy selects the node with the highest navigation score as the local goal from the previously unvisited candidates.

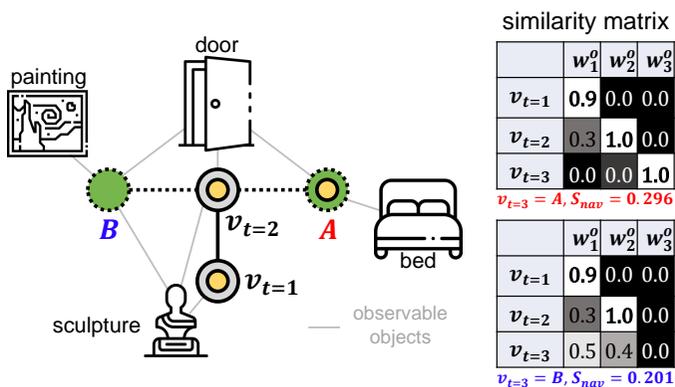


Figure 2.6: **Toy Example.** Monotonic alignment between language instruction and visual observation is desirable. Yellow dots \bullet in the nodes describe the ground truth trajectory. Based on the node at $t = 3$, the similarity matrix can show either monotonic or non-monotonic alignment between object tokens and SOS features. The green circles \clubsuit describe the possible candidates A, B for next action.

Figure 2.6 illustrates a simple scenario of entering a room. Suppose W^o is given as [sculpture, door, bed] and the agent has to compare two trajectories $\mathcal{T}_1 = (v_1, v_2, A)$ and $\mathcal{T}_1 = (v_1, v_2, B)$. Each similarity matrix in Figure 2.6 has the (t, j) element as the similarity between the SOS feature of V_t and $\hat{\delta}(w_j^o)$, which is calculated as $\hat{\delta}(w_j^o) \cdot \vec{S}_t^o$. Notably, the similarity matrix shows monotonic alignment and the navigation score is higher when the next action is chosen correctly.

2.2.3 Training Details

We use [27] for pretraining the visual encoder with panoramic RGB observations. We use the DAgger algorithm [48] to pretrain the navigation policy and the mode selector. To prevent overfitting, we iteratively perform teacher forcing and student forcing to choose the action from the exploration policy. Imitation learning loss is calculated as $L_{IL} = \sum_{t=1}^T -\log p(a_t^* | a_t)$ and object

grounding loss is calculated as $L_{OG} = -\log p(\text{obj}^* | \text{obj}_{pred})$, where obj^* denotes the ground truth and obj_{pred} denotes the predicted object location. The total loss function is defined as $L_{total} = L_{mode} + L_{progress} + L_{corr} + L_{IL} + L_{OG}$. We further finetune the agent via A2C [49]. The exploration policy selects the action a_t with probability p_t^a . Reinforcement learning loss is defined as $L_{RL} = -\sum_t a_t^s \log(p_t^a) A_t - \lambda \sum_t a_t^* \log(p_t^a)$. To train the mode selector, progress monitor, and exploration policy in an end-to-end manner, we use the total loss function as $L_{fine} = L_{mode} + L_{progress} + L_{RL}$. The exploitation policy searches the path toward the local goal from the constructed navigation graph. Thus, the exploitation policy is not learned.

2.3 Navigation Experiments

2.3.1 Experiment Settings

We evaluate our method on three VLN benchmarks, Room-to-Room(R2R) [2], SOON [19], and REVERIE [20].

R2R evaluates the visually-grounded natural navigation performance of the agent. The agent must navigate to the predefined goal point given image observations and language instructions in an unseen environment.

SOON is also a goal-oriented VLN benchmark. Natural language instructions in SOON have an average length of 47 words. The agent should locate the target location and detect the location of an object to find the target object.

REVERIE is a goal-oriented VLN benchmark that provides natural language instruction about target locations and objects. In REVERIE, the agent is given an instruction referring to a remote object with an average length of 21 words. With this instruction and a panoramic observation from the environment, the agent should navigate to the location the instruction describes and find the correct object bounding box among the predefined object bounding boxes.

2.3.2 Evaluation Metrics

Navigation performance

We evaluate algorithms using the trajectory length (TL), success rate (SR), and success weighted by inverse path length (SPL) [50], and oracle success rate (OSR) for the navigation performance comparison. An episode is recorded as a success if the agent takes a `stop` action within 3 m of the target location. TL is the average path length in meters. SR is denoted as the number of successes divided by the total number of episodes, M . SPL is calculated as $\frac{1}{M} \sum_{i=1}^M S_i \frac{l_i}{\max(p_i, l_i)}$, where S_i denotes the success as a binary value. p_i and l_i denote the shortest path and actual path lengths for the i^{th} episode. OSR uses the oracle stop policy instead of the stop policy of the agent.

Object grounding performance

We also evaluate the object grounding performance of the agent by the success rate of finding the target object (FSR) and the target finding success weighted by inverse path length (FSPL)¹ [19, 20]. FSPL is calculated as follows:

$$\text{FSPL} = \frac{1}{N} \sum_{i=1}^N S_i^{nav} S_i^{loc} \cdot l_i^{nav} / \max(l_i^{nav}, l_i^{gt}), \quad (2.5)$$

where S_i^{nav} is whether the agent navigates to the target, S_i^{loc} is whether the agent finds a target object bounding box, and l_i^{nav} and l_i^{gt} are the navigation trajectory length and ground truth trajectory length, respectively.

2.3.3 Baselines and Implementation Details

We compare our method with several other baselines as follows. For each task, we compare our method with a number of baselines that use various types of memory (recurrent, sequential, and topological map). For methods implemented

¹Identical with its original term, Remote Grounding Success (RGS).

Methods	Memory	Exploit	Val Seen				Val Unseen				Test Unseen			
			SR↑	SPL↑	TL↓	NE↓	SR↑	SPL↑	TL↓	NE↓	SR↑	SPL↑	TL↓	NE↓
Random	-	-	16	-	9.58	9.45	16	-	9.77	9.23	13	12	9.89	9.79
Human	-	-	-	-	-	-	-	-	-	-	11.85	1.61	86	76
Seq2Seq [2]	Rec	✗	6.0	39	11.33	-	22	-	8.39	7.84	20	18	8.13	7.85
VLN \odot BERT [34]	Rec	✗	72	68	11.13	2.90	63	57	12.01	3.93	63	57	12.35	4.09
SMNA \uparrow [24]	Rec	homing	69	63	11.69	3.31	47	41	12.61	5.48	61	56	-	4.48
Regretful-Agent [25]	Rec	homing	69	63	-	3.23	50	41	-	5.32	48	40	-	5.69
FAST (short) [37]	Rec	homing	-	-	-	-	56	43	21.17	4.97	54	41	22.08	5.14
FAST (long) [37]	Rec	homing	70	04	188.06	3.13	63	02	224.42	4.03	61	03	196.53	4.29
HAMT-e2e [36]	Seq	✗	76	72	11.15	2.51	66	61	11.46	2.29	65	60	12.27	3.93
DUET [27]	Top. Map	✗	79	73	12.32	2.28	72	60	13.94	3.31	69	59	14.73	3.65
SSM [29]	Top. Map	jump	71	62	14.7	3.10	62	45	20.7	4.32	61	46	20.4	4.57
Meta-Explore (Ours)	Top. Map	local goal	81	75	11.95	2.11	72	62	13.09	3.22	71	61	14.25	3.57

Table 2.1: Comparison and evaluation results of the baselines and our model in the R2R Navigation Task. Gray shaded rows describe hierarchical navigation baselines. Three memory types: Rec (recurrent), Seq (sequential), and Top. Map (topological map)

with a hierarchical navigation framework, we compare the specific exploitation methods: homing, jump, and local goal search. Homing makes the agent backtrack, and jump makes the agent jump to a previously visited node. The hyperparameters and detailed model architecture of Meta-Explore are described in the supplementary material.

2.3.4 Comparison with Navigation Baselines

We compare our method with navigation baselines². We focus on the success rate and SPL. Rendered results and detailed analyses with other evaluation metrics are provided in the supplementary material.

R2R. Table 2.1 compares the proposed Meta-Explore with baselines for the R2R navigation task. We categorize the baseline methods based on the type of constructed memory and the type of exploitation. Our method outperforms other exploration-only baselines over all types of validation and test splits in suc-

²† indicates reproduced results.

Methods	Memory	Exploit	Val Seen Instruction				Val Seen House				Test Unseen House			
			SR↑	SPL↑	OSR↑	FSPL↑	SR↑	SPL↑	OSR↑	FSPL↑	SR↑	SPL↑	OSR↑	FSPL↑
Human	-	-	-	-	-	-	-	-	-	-	90.4	59.2	91.4	51.1
Random	Rec	✗	0.0	1.5	0.1	1.4	0.1	0.0	0.4	0.9	2.1	0.4	2.7	0.0
Speaker-Follower [30]	Rec	✗	97.9	97.7	97.8	24.5	61.2	60.4	69.4	9.1	7.0	6.1	9.8	0.6
RCM [51]	Rec	✗	84.0	82.6	89.1	10.9	62.4	60.9	72.7	7.8	7.4	6.2	12.4	0.7
AuxRN [26]	Rec	✗	98.4	97.4	98.7	13.7	68.8	67.3	78.5	8.3	8.1	6.7	11.0	0.5
GBE w/o GE	Top. Map	✗	89.5	88.3	91.8	24.2	62.5	60.8	73.0	6.7	11.4	8.7	18.8	0.8
GBE [19]	Top. Map	✗	98.4	97.9	98.6	44.2	76.3	62.5	64.1	7.3	11.9	10.2	19.5	1.4
GBE [†]	Top. Map	✗	-	-	-	-	19.5	13.3	28.5	1.2	12.9	9.2	21.5	0.5
DUET [27]	Top. Map	✗	94.0	91.6	90.0	31.1	36.3	22.6	50.9	3.8	33.4	21.4	43.0	4.2
Meta-Explore (Ours)	Top. Map	local goal	100.0	99.1	96.0	33.9	44.7	34.8	52.7	8.9	39.1	25.8	48.7	4.0

Table 2.2: Comparison and evaluation results of the baselines and our model in the SOON Navigation Task.

cess rate and SPL. Compared with hierarchical baselines SMNA [24], Regretful-Agent [25], FAST [37], and SSM [29], Meta-Explore improves success rate and SPL by at least 16.4% and 8.9%, respectively. The main difference is that Meta-Explore constructs a topological map during exploration and uses the map for local goal search in exploitation. On the contrary, homing exploitation policies in SMNA, Regretful-Agent, and FAST only rely on the current trajectory, instead of taking advantage of the constructed memory. Jump exploitation in SSM uses a topological map to search a successful previous node, but it makes an unrealistic assumption that the agent can directly jump to a previously visited distant node and unfairly saves time. In our approach, we plan a path to the local goal based on the topological map. The experiment results reveal that even if we design a hierarchical navigation framework, exploration and exploitation are not entirely separate but they can complement each other.

SOON, REVERIE. Table 2.2 compares Meta-Explore with baselines in the SOON navigation task. While the proposed method does not improve performance in val seen split, Meta-Explore outperforms other baselines in the test unseen split of SOON for success rate by 17.1% and SPL by 20.6%. The result implies that for the goal-oriented VLN task, high performance in train or val

seen splits can be the overfitted result. Because the agent can be easily overfitted to the training data, making a generalizable model or providing a deterministic error-correction module for inference is essential. Meta-Explore chooses the latter approach by correcting the trajectory via exploitation in regretful cases. The evaluation results in the REVERIE navigation task are described in the supplementary material. Meta-Explore shows improvement in the val split of REVERIE for success rate and SPL, but the improvement in the test split is lower than the results in R2R and SOON. We found 252 meaningless object categories (e.g., verbs, adjectives, and prepositions) and 418 replaceable object categories (e.g., typographical errors and synonyms) in the REVERIE³ dataset. Because our exploitation method utilizes object-based parsing of the given instruction to match with the detected object categories, the effectiveness of the proposed method is lessened due to inaccuracies and inconsistencies in the dataset. We expect to have higher performance if the mistakes in the dataset are fixed.

2.3.5 Local Goal Search using SOS Features

To discuss the significance of modeling exploitation policy, we conduct specific experiments about choosing the local goal for R2R and SOON. We evaluate our method using different types of local goal search, as shown in Table 2.3 and 2.4. Oracle denotes a method which selects a local goal using the ground truth trajectory. The performance of the oracle provides the achievable performance for each dataset. The results imply that local goal search using either spatial or spectral visual representations is more effective than random local goal search. The results show that local goal search using spectral visual representations, i.e., SOS features, lead the agent to desirable nodes the most. We also compare

³10.7% and 41.2% of a total of 46,476 words in the bounding box dataset correspond to meaningless and replaceable object categories, respectively.

Local Goal	Val Seen					Val Unseen				
	SR↑	SPL↑	OSR↑	TL↓	NE↓	SR↑	SPL↑	OSR↑	TL↓	NE↓
Oracle	81.88	74.12	87.46	13.06	1.93	75.95	62.53	84.16	14.00	2.71
Random	79.33	72.67	85.31	13.19	2.22	70.97	59.45	80.16	14.92	3.34
Homing	80.22	73.63	85.60	12.51	2.14	71.65	60.60	80.33	13.91	3.26
Spatial	79.63	73.14	85.60	12.99	2.22	71.56	60.01	80.33	14.90	3.27
Spectral	80.61	75.15	85.80	11.95	2.11	71.78	61.68	80.76	13.09	3.22

Table 2.3: Comparison of Exploitation Policies. (R2R)

Local Goal	Val Seen House				Test Unseen House			
	SR↑	SPL↑	OSR↑	FSPL↑	SR↑	SPL↑	OSR↑	FSPL↑
Oracle	54.42	37.96	63.72	11.01	48.38	28.45	62.98	4.74
Random	24.78	11.97	34.96	3.08	24.19	7.41	35.84	1.29
Homing	42.04	27.72	48.23	10.18	35.40	19.18	51.62	3.14
Spatial	32.30	11.60	39.38	1.90	26.11	10.58	39.23	1.43
Spectral	44.69	34.84	52.65	8.89	39.09	25.80	48.67	4.01

Table 2.4: Comparison of Exploitation Policies. (SOON)

local goal search with homing and the difference between the performance of the two methods is most noticeable in the test split of the SOON navigation task. As shown in Table 2.4, choosing the local goal with only spatial-domain features, the navigation performance does not improve compared to homing. On the contrary, spectral-domain local goal search shows significant improvement against homing by 10.4% in success rate, 34.5% on SPL, and 27.4% on FSPL. The results imply that using spectral-domain SOS features helps high-level decision making, thereby enhancing the navigation performance. To further show the effectiveness of SOS features, we provide sample local goal search scenarios in the supplementary material.

2.3.6 Ablation Study

We conduct an ablation study to compare the proposed method against language-triggered hierarchical exploration. Results in the supplementary material show that among the three representation domains, spatial, spectral, and language, the spectral-domain features enhance navigation performance the most. Additionally, to implicate further applications of Meta-Explore in continuous environments, we evaluate our method on the photo-realistic Habitat [52] simulator to solve image-goal navigation and vision-and-language navigation tasks. Implementation details and results are included in the supplementary material. Results show that our method outperforms baselines in both tasks.

2.4 Discussion

In this Chapter 2, we have proposed Meta-Explore, a hierarchical navigation method for VLN, by correcting mistaken short-term actions via efficient exploitation. In the exploitation mode, the agent is directed to a local goal which is inferred to be the closest to the target. A topological map constructed during exploration helps the agent to search and plan the shortest path toward the local goal. To further search beyond the frontier of the map, we present a novel visual representation called *scene object spectrum* (SOS), which compactly encodes the arrangements and frequencies of nearby objects. Meta-Explore achieves the highest generalization performance for test splits of R2R, SOON, and val split of REVERIE navigation tasks by showing less overfitting and high success rates. We plan to apply Meta-Explore for VLN tasks in continuous environments in our future work.

Appendix A

Appendices

We provide additional details and analyses of the proposed method in this supplementary material. Section A.1 provides model details. Section A.2 provides detailed settings and data preprocessing for experiments. Section A.3 provides evaluation results with detailed analyses. Section A.4 provides implementation details and detailed results for the ablation study.

A.1 Model Details

A.1.1 Algorithm Details

Algorithm 1 summarizes the overall hierarchical exploration process. The mode selector supervises the process and chooses whether the agent should explore or exploit at each time step.

A.1.2 Exploitation Module

Reference SOS Features.

In the proposed method, we approximate the reference SOS feature of an object token by using prior information about objects in the training data. For in-

Algorithm 1: Meta-Explore

```

 $P_{explore} \leftarrow 1$ 
 $Success \leftarrow False$ 
Initialize  $G_t$  and node features
while  $t < T$  do
    Update  $G_t$ 
    Update node features
     $H_t \leftarrow$  cross-modal embedding at time  $t$ 
    if  $P_{explore} \geq 0.5$  then
         $a_t \leftarrow \arg \max_{V_i} (F_{explore}([H_t]_i))$ 
         $\hat{p}_t \leftarrow F_{progress}(H_t)$ 
         $t \leftarrow t + 1$ 
    else
         $V_{local} \leftarrow$  unvisited but observed nodes in  $G_t$ 
         $v_{local} \leftarrow \arg \max_{v' \in V_{local}} (S_{nav}(\tau'(v_0, v')))$ 
         $\tau \leftarrow PathPlanning(v_t, v_{local})$ 
        while not arrived at  $v_{local}$  do
             $a_t \leftarrow pop(\tau)$ 
             $t \leftarrow t + 1$ 
        end
    end
     $P_{explore} \leftarrow 1 - S_{mode}(H_t)$ 
    if  $a_t$  is stop and  $d(v_t, v_{goal}) < d_{success}$  then  $Success \leftarrow True$ 
end

```

stance, for the ‘chair’ object, we collected the widths and heights of the detected bounding boxes as shown in Figure A.1. Figure A.2 shows two representative values: median and mean for each distribution. We choose the median values, which minimizes the L1 error, to represent the reference bounding box of each object.

To generate rotation-invariant SOS features, we convert the four vertices of the bounding box detected from the front view image of size 640×480 to the vertices of a bounding box detected from the panoramic view image of size 2048×512 using coordinate transformations. To simplify the implementation,

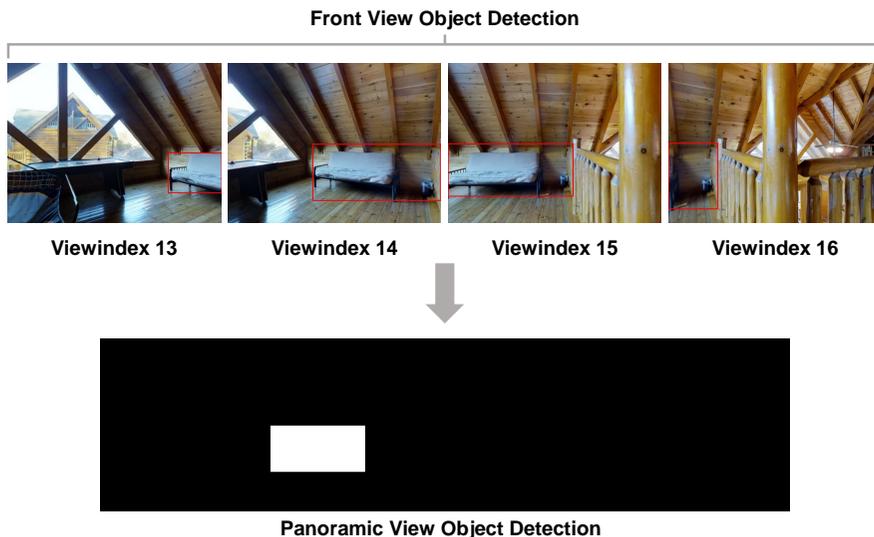


Figure A.1: **Bounding box coordinate transformation.** Front-view visual observations from different angles at the same location. Each bounding box shows the ‘chair’ detection. We use coordinate transformation to convert coordinates into panoramic view.

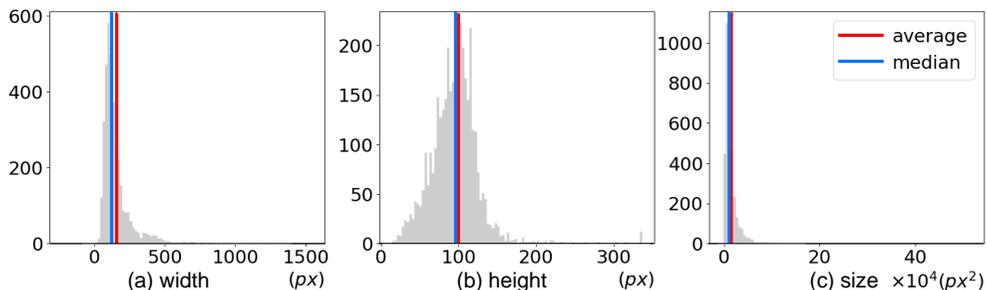


Figure A.2: **Bounding box statistics.** We collect width, height, and size of detected bounding boxes. The histograms show statistics for ‘chair’ objects. Yellow line and red line show the median and average values of each distribution, respectively.

we assume that the converted bounding box has a rectangle shape with the vertices transformed into coordinates in a panoramic view. The reference SOS feature is calculated as the logarithmic magnitude of the Fourier transform of the panoramic mask with mean pooling on the vertical spectral axis. Consid-

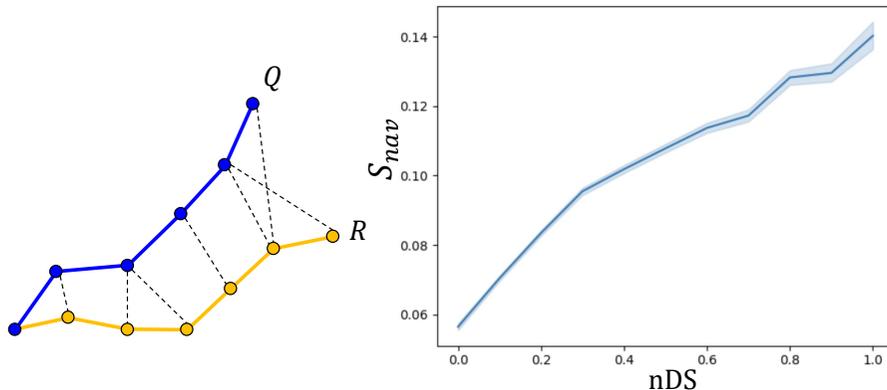


Figure A.3: **Relationship between navigation score (S_{nav}) and normalized distance sum (nDS) in R2R.** We measure navigation scores for augmented trajectories which include both successful and failed trajectories. R and Q illustrate an example case of a ground truth trajectory and a query trajectory. Maximum hop of a query trajectory is 15. Trajectories with high nDS scores also have high navigation scores.

ering that the shift in the spatial-domain only affects the phase of the Fourier transform, the location of a reference bounding box does not matter.

Navigation Score

To compare local goal candidates, we design a navigation score of a corrected trajectory τ' as equation A.1. This metric can also be interpreted as a weighted correlation coefficient among SOS features and object tokens weighted by the similarities between them.

$$S_{nav}(\mathcal{T}') = \frac{\sum_{i=1}^B \sum_{j=1}^{t'} \left(\frac{\hat{\delta}(w_i^o)}{|\hat{\delta}(w_i^o)|} \cdot \frac{\vec{S}'_j}{|\vec{S}'_j|} \right) ((\hat{\delta}(w_i^o) - \hat{\delta}(\bar{w}^o)) \cdot (\vec{S}'_j - \bar{\vec{S}}'))}{\sqrt{\frac{t'}{B} \cdot \sum_{i=1}^B (\hat{\delta}(w_i^o) - \hat{\delta}(\bar{w}^o))^2 \sum_{j=1}^{t'} (\vec{S}'_j - \bar{\vec{S}}')^2}} \quad (\text{A.1})$$

Figure A.3 shows the relationship between the navigation score and an evaluation metric in the R2R navigation task. Both metrics measure how similar the

current trajectory is to the ground truth trajectory. We generate 49,986 augmented trajectories with an average length of $8.23m$ based on 5596 ground truth trajectories. To generate various samples, we separate each augmented trajectory (v_1, v_2, \dots, v_t) into t augmented trajectories (v_1) , (v_1, v_2) , \dots , and (v_1, v_2, \dots, v_t) . The final 421,383 augmented trajectories include trajectories with 1 to 15 nodes and include both successful and unsuccessful trajectories. We classify the trajectories with the normalized distance sum (nDS) between ground truth trajectory R and a query trajectory Q as follows:

$$\text{nDS}(R, Q) = \exp \left(- \frac{\sum_{v_i \in R} \min_{u_j \in Q} d(v_i, u_j) + \sum_{u_j \in Q} \min_{v_i \in R} d(u_j, v_i)}{\frac{|R|+|Q|}{2} d_{\text{success}}} \right), \quad (\text{A.2})$$

which requires the ground truth information of R . $d(u, v)$ denotes the geodesic distance between two nodes, u and v , and d_{success} denotes the success distance. The plot in Figure A.3 shows a linear relationship between the nDS and the navigation score. The results imply that the proposed navigation score effectively scores the augmented trajectories even though it only relies on the given target instruction and observation from the augmented paths, without any location information about the nodes on the ground truth trajectory.

A.1.3 Implementation Details

We use ViT-B/16 [44] pretrained on ImageNet to extract features from the viewpoint panoramic images. We use pretrained LXMERT [47] for the language encoder and cross-modal transformer. We implement the mode selector as a two-layer feed-forward network.

Dataset	Instruction	Object Tokens	Target Object
R2R	“Walk through the kitchen. Go past the sink and stove stand in front of the dining table on the bench side.”	[“kitchen”, “sink”, “stove”, “stand”, “dish”, “table”, “chair”]	“chair”
SOON	“This is a brand new white, rectangular wooden table, which is above a few chairs, under a pot of flowers. It is in a very neat study with many books.”	[“book”, “chair”, “pitcher”, “flower”, “table”, “table”]	“table”
REVERIE	“Go to the bedroom with the fireplace and bring me the lowest hanging small picture on the right wall across from the bedside table with the lamp on it”	[“bedroom”, “fireplace”, “bed”, “table”, “stand”, “art”]	“art”

Table A.1: **Object Parsing Examples.** For each dataset, object tokens are extracted from the instructions. Target objects are inferred from the instructions using VQA. Words that have similar meanings are unified into a single object word for categorization.

A.2 Experiment Setup

A.2.1 Dataset Statistics

R2R. The average length of instructions is 32 words. The average path length of the ground truth trajectory of each instruction is six steps. The number of train, val seen, val unseen, and test episodes are 14,025, 1020, 2349, and 4173.

SOON. The average path length of the ground truth trajectory of each instruction is four to seven steps. The number of train, validation seen instruction, validation seen house, validation unseen house episodes are 3085, 245, 195, and 205.

REVERIE. The average path length of the ground truth trajectory of each instruction is 9.5 steps. The number of train, val seen, val unseen, and test episodes are 10,466, 1423, 3521, and 6292.

A.2.2 Data Preprocessing

To calculate reference SOS features, we preprocess object tokens from language instructions. Using a pretrained visual question answering (VQA) model [53] with the question “*What is the target object? Answer in one word.*”, we extract tar-

get objects from the instructions in R2R and REVERIE datasets. For SOON dataset, the target object names are already given. After extracting target objects, we perform object parsing for the instructions as shown in Table A.1. The final object tokens are sorted by order of appearance in the instructions for R2R and REVERIE. For SOON, considering that the full instruction is divided into 5 parts: object name, object attribute, object relationship, target area, and neighbor areas, we sort the object tokens by reversed order of sentences.

A.2.3 Baselines

Seq2Seq [2] uses sequence-to-sequence action prediction to generate actions from the agent trajectory. Speaker-Follower [30] uses the speaker model to augment natural language instructions and evaluate the candidate action sequence. FAST [37] uses both local and global signals to look forward the unobserved environment during exploration and backtrack to the originally visited nodes when needed. SMNA [24] uses visual-textual co-grounding module that encodes the past instructions and the instructions and actions to be done. SMNA also uses a progress monitor to estimate the current progress of the agent relative to the total instructions. Regretful-Agent [25] improves SMNA via two modules. The regret module decides whether to continue to explore or rollback to previous state by a learned policy, and the progress marker decides the direction the agent should head to by selecting visited nodes with progress estimates. RCM [51] applies reinforcement learning to enforce the global matching between the agent trajectory and the given natural language instruction. Via cycle-reconstruction reward, RCM allows the agent to comprehend the natural language instruction and penalize paths that do not match with the given instructions. FAST-MATTN [20] introduces a Navigator-Pointer model to both navigate to the target point and to localize the object from the navigation point according to the language guidance.

AuxRN [26] introduces four auxiliary tasks that help learning the navigation policy: a trajectory retelling task, a progress estimation task, an angle prediction task, and cross-modal matching task, and improves navigation success by aligning representations in these unseen domains with seen domain. HAMT [36] uses transformer instead of a recurrent unit to predict actions from a long-range trajectory of observations and actions. Airbert [33] uses ViLBert [32] to measure the correlation between the language instructions and the viewpoint trajectories. VLN \circ BERT [34] adds a recurrent unit in the transformer to predict the action from the trajectory. SIA [54] first pretrains the agent to learn the cross-modality between object grounding task and scene grounding task, and then generates real action sequences with memory-based attention. SSM [29] integrates information during exploration and constructs a scene memory and chooses the most probable node among visited nodes during backtracking. GBE [19] models the navigation state as a graph and explores the environment based on the navigation graph. DUET [27] uses two models, a local encoder and a global map planner, to fuse the local observations and coarse scale encoding for planning actions.

A.3 Navigation Experiments

In this section, we analyze the evaluation results of navigation experiments with different evaluation metrics. The results are provided in the paper.

A.3.1 Detailed Analyses in R2R

Navigation Error (NE). Navigation error (NE) is measured as the average distance between the final location of the agent and the target location of episode in meters. Because each episode is recorded as success if NE is less than $3m$, NE is strongly related with the success rate. Meta-Explore shows the lowest NE in the val seen and test unseen splits of the R2R navigation task. The results imply that

hierarchical exploration with local goal search helps the agent arrive to the target location closer than other baselines.

Trajectory Length (TL). Among all the R2R navigation baselines, Seq2Seq shows the lowest TL. However, Seq2Seq shows low success rate and low SPL in all data splits. Compared to navigation baselines with SPL higher than 50% in the test split, VLNBERT, SMNA, and HAMT-e2e show lower TL than Meta-Explore. However, all three of these methods show a lower success rate, SPL, and NE than Meta-Explore. According to R2R [2], train episodes show a wide range of average trajectory length from $5m$ to $25m$, while the test episodes have an average trajectory length of $9.93m$. This implies that the agent is trained with longer trajectories than the test split trajectories, thereby the navigation policy might have learned to navigate longer paths better than shorter paths.

A.3.2 Detailed Analyses in SOON

Oracle Success Rate (OSR). In the SOON navigation task, Meta-Explore achieves the highest OSR in the test split while it does not improve the OSR in the val seen instruction and val seen house splits. The proposed method shows a significant generalization result compared to the baselines. AuxRN shows the highest OSR in both the val seen instruction split and the val seen house split as 78.5% and 97.8%, respectively, but shows the OSR in the test unseen split as 11.0%. On the other hand, Meta-Explore shows OSR as 96.0%, 52.7%, and 48.7% in the val seen instruction, val seen house, and test unseen splits, respectively. Meta-Explore outperforms AuxRN on OSR by 442.7% in the test split.

Object Grounding Performance (FSPL). Following [19], we measure the object grounding performance with the target finding success weighted by path length (FSPL). Although Meta-Explore show the highest success rate and SPL in the val seen instruction and test splits, it does not improve FSPL over baseline

methods. We expect to achieve better performance on FSPL if the agent uses the SOS features as deterministic clues to find the target object at the end of each episode.

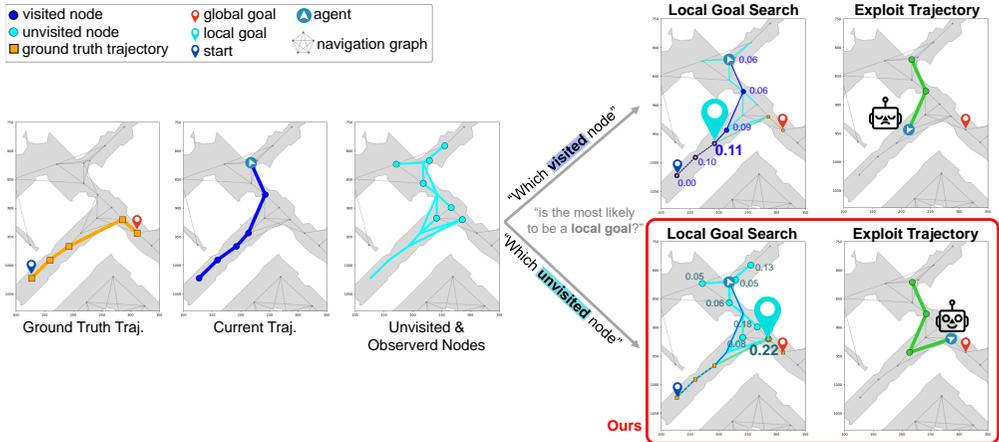
A.3.3 Evaluation Results in REVERIE benchmark

Table A.2 compares Meta-Explore with the baselines in the REVERIE navigation task. While the proposed method does not improve performance in the val seen split, Meta-Explore outperforms other baselines in the val unseen on success rate, SPL, FSR, FSPL, and OSR. However, the improvement of performance is lower than the improvements shown in R2R and SOON benchmarks. We found 252 meaningless object categories (e.g., verbs, adjectives, and prepositions) and 418 replaceable object categories (e.g., typographical errors and synonyms) in the REVERIE dataset. 10.7% and 41.2% of a total of 46,476 words in the bounding box dataset correspond to meaningless and replaceable object categories, respectively. Because our exploitation method utilizes object-based parsing of the given instruction to match with the detected object categories, the effectiveness of the proposed method is lessened due to inaccuracies and inconsistencies in the dataset. We expect to have higher performance if the mistakes in the dataset are fully fixed. To provide evidence for this hypothesis, we evaluate Meta-Explore with a modified dataset, which is partially fixed. Typographical errors are fixed and words that have similar meanings are unified into a single object category. For instance, ‘blackboard’, ‘whiteboard’, and ‘bulletin’ are all unified into ‘board’. The results are shown as the performance of **Meta-Explore*** in Table A.2. The results imply that the proposed method can effectively enhance the SPL by classifying the detected objects correctly, using the modified dataset.

Comparison between exploitation policies in the REVERIE navigation task is shown in Table A.3. Among the four exploitation methods: random, spatial,

Local Goal	Val Seen					Val Unseen				
	SR↑	SPL↑	FSR↑	OSR↑	TL↓	SR↑	SPL↑	FSR↑	OSR↑	TL↓
Oracle	79.20	64.17	62.83	84.05	18.53	59.07	38.23	40.36	66.86	26.71
Random	0.21	0.04	0.00	20.31	46.02	1.11	0.18	0.34	26.70	0.05
Homing	68.45	50.54	55.24	73.23	17.95	43.60	28.25	29.59	49.28	25.64
Spatial	67.53	40.21	54.25	70.91	26.92	40.90	23.25	27.61	45.84	26.92
Spectral	71.68	63.90	57.34	73.79	13.84	47.49	34.03	32.32	51.21	22.12

Table A.3: Comparison of Exploitation Policies. (REVERIE)



spectral local goal search and homing, spectral-domain local goal search shows the highest performance. The results in Table A.3 are consistent with the results in R2R and SOON.

A.3.4 Local Goal Search

In this section, we provide sample local goal search scenarios. Figure A.4 shows two scenarios of choosing the local goal when the agent moves to a wrong direction. The agent is given an instruction “*Turn right and turn right again after the desk on the right. Wait next to the cabinets and microwave.*”. In both scenarios, we assume that the agent chooses the local goal as the node with the highest navigation score among the possible candidates. If the local goal is chosen from the previously visited nodes, the agent has to move back toward the explored regions. In contrast, if the local goal is chosen from unvisited but observed nodes, the agent can choose a local goal which is close to the global goal. The two scenarios imply that the local goal search in Meta-Explore is more effective than exploitation methods that return the agent to a previously visited node.

A.4 Ablation Study

A.4.1 Language-triggered Hierarchical Exploration

In the proposed method, the target instruction and local goal candidates are compared in spectral-domain using SOS features. Since semantic information can also be expressed in language-domain, we further experiment with the local goal search method using synthesized language captions from visual observations in the R2R navigation task. We compare three types of representation domains: spatial, spectral, and language, which are implemented as panoramic RGB image embeddings, SOS features, and sentence embeddings, respectively. To compare features in different domains, we transfer the source domain to another using augmentation or cross-domain similarity.

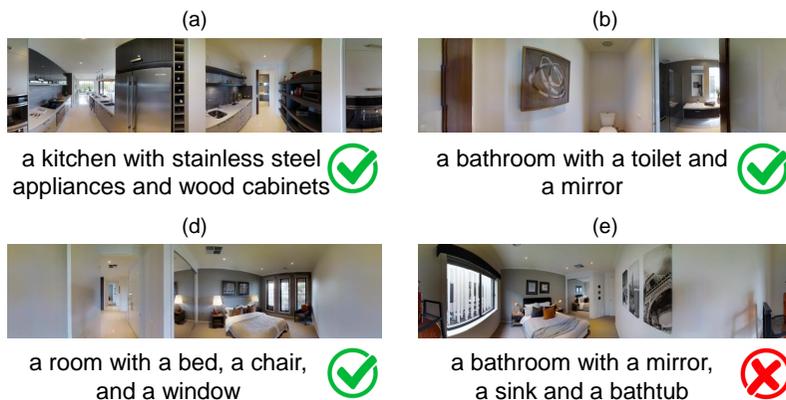


Figure A.5: **Sample image captions.** (a), (b), (c), and (d) show captions that successfully describe the scenes. (e) and (f) shows failure cases of caption generation. For successful language-triggered hierarchical exploration, image captions should correctly describe the scenes. However, current image captioning methods often generates misdescribed captions, thereby leading to a low navigation performance.

Implementation Details

We address that the agent can use image captioning to extract contextual information from visual observations such as room type, color, and object placements. To compare local goal candidates and target instruction in language domain, we use pretrained ViT [44] and GPT-2 [55] to generate the caption for each viewpoint as Figure A.5. The Figure shows four successful cases and two failure cases of image captions. To find a local goal using the generated captions, we calculated the similarities between the captions corresponding to local goal candidates and the target instruction using a fine-tuned sentence transformer ‘all-MiniLM-L6-v2’ [56]. The local goal is chosen as the candidate with the highest similarity. Additionally, we use pretrained CLIP [57] to evaluate local goal search based on cross-modal similarities between the visual observations of local goal candidates and the target instruction.

Domains		Val Seen		Val Unseen	
Nav. Target	Local Goal	SR	SPL	SR	SPL
Lang.	X	79.92	72.79	70.63	59.81
Lang.	Spatial	78.84	71.96	71.05	58.86
Lang.	Lang. Aug.	77.96	70.77	69.52	57.26
Spectral Aug.	Spectral	80.61	75.15	71.78	61.68

Table A.4: **Comparison and evaluation results of the local goal search methods using different target and candidate domains. (R2R)**

Experiment Results

Table A.4 shows the evaluation results of the local goal search methods using different target and candidate domains in R2R navigation task. *Nav. Target* denotes the target of VLN, initially given as language. *Lang. Aug.* denotes language captions generated from images. *Spectral Aug.* denotes reference SOS features generated from language instructions. Among the three representation domains, the spectral-domain features enhance navigation performance the most. This implies that hierarchical exploration is most effective when used with spectral visual features. Table 2.1 and Table 2.2 in the paper also show the improvement of navigation performance by using both hierarchical exploration and spectral visual features over DUET [58], which uses the same ViT-B/16 to extract spatial visual features, resulting in 17.1% increase in SR and 20.6% increase in SPL in the SOON test unseen split.

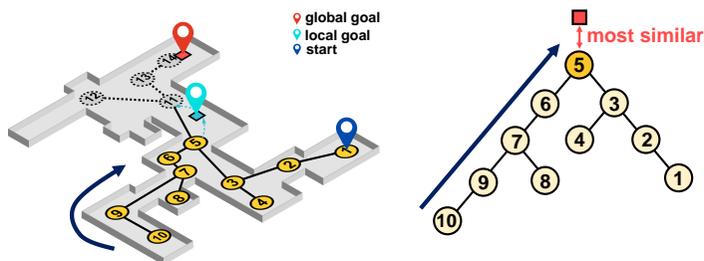


Figure A.6: **Exploitation by searching a local goal.** In the exploit mode, the agent aims to escape from the stranded local area. It first searches for the most similar node to the goal. Then, it finds an optimal local goal which is unexplored and also similar to the goal image. We use SOS features to compare with the target image.

A.4.2 Image-Goal Navigation in Continuous Domain

To implicate further applications of Meta-Explore in a continuous domain, we evaluate our method on the photo-realistic Habitat [52] simulator with continuous action space with realistic noises to solve an image-goal navigation task. The objective is to arrive at the target location of the given goal image in an unseen environment. We mainly focus on the effectiveness of hierarchical exploration using local goal search in this experiment. The results are shown in Table A.5.

Exploration-Exploitation Selection

We extend Meta-Explore to continuous environments to address the impact of hierarchical exploration in realistic environments. The mode selector decides when to explore and exploit. In the exploration mode, the agent explores around a local area until the meta-controller decides to stop the exploration. The exploration module consists of graph construction module and navigation module. We use recurrent action policy that takes the current and target image features and outputs low-level actions for exploration. We illustrate that the explore-exploit switching decision occurs in stuck scenarios, such as entering a small place or getting stranded in a corner. Figure A.6 shows the overview of exploitation in image-goal navigation

by searching a local goal. When the control mode is changed to exploitation mode, the agent returns to the closest previously visited node. Then, the agent finds a local goal among the nodes in the constructed topological map and moves toward the local goal using dijkstra’s algorithm [59]. The local goal is chosen as the node which has the most similar SOS feature with the SOS feature of the target image based on cosine similarity. The agent repeats this explore-exploit behavior until it finds the goal. This explore-exploit switching decision increases the navigation success rate.

Experiment Details

We evaluate Meta-Explore in the Gibson dataset [60] with Habitat [52] simulator to solve an image-goal navigation task. Habitat simulator allows the agent to navigate in photo-realistic indoor environments. The exploration policy of the agent is trained using 72 scenes. We evaluate Meta-Explore using 14 unseen scenes. We use panoramic RGBD observations and construct image-based graph memory. To construct a context frequency vector, we detect objects via Mask2Former [61] pretrained in ADE-20K dataset [62], to effectively detect the objects that are generally located in indoor scenes. We use a discrete action space, `{stop, move forward, turn left, turn right}` for navigation. With `move forward` action, an agent moves forward by 0.25 m, while `turn left` and `turn right` denotes a 10° rotation, counter-clockwise and clockwise, respectively. The difficulty of each episode is determined by the geodesic distance between the initial and the goal location; *easy*: 1.5 m~3 m, *medium*: 3 m~5 m, and *hard*: 5 m~10 m. The actuation noise model [41] is also applied to the agent in order to evaluate in realistic situations. We also demonstrated navigation experiments in the real world using a Jackal robot. The episodes are sampled from simulation point goal episodes with all difficulties; *easy*, *medium* and *hard*. We demonstrate both

straight and curved trajectories to evaluate that our model is not task-specific. We used the model only trained in Habitat simulator with Gibson dataset. To collect panoramic RGBD observations, we use one panoramic RGB camera and four front-view RGBD cameras. In order to implement collision avoidance similar to the construction of navigable mesh in Habitat simulator, we implemented a collision avoidance module by clipping the action value based on the depth image observation.

Baselines

We compare our image-goal navigation policy with various baselines. Active Neural SLAM (ANS) constructs a top-down metric map and uses a hierarchical structure consisting of global and local policies. The global policy outputs long-term goals, which are used to generate short-term goals. The local policy uses a geometric path planner to navigate to a short-term goal. NTS [63] constructs a topological graph during exploration and plans subgoals with graph localization and planning, while navigating to the node with local point goal navigation policy. Neural Planner [64] constructs a graph using an estimated connectivity probability calculated from the neural network. VGM [65] uses unsupervised image-based graph memory representation to compare the similarity between goal image and the current observation image. We adapt VGM for graph construction and local navigation policy. PCL [66] encoder with ResNet18 [43] backbone network is used as the visual encoder for VGM [65].

Evaluation Metrics

We evaluate both success rate (SR) and success weighted by inverse path length (SPL) [50]. An episode is recorded as success if the agent takes a `stop` action within 1 m of the target location. SR is denoted as the number of successes divided

Methods	Exploit	Need Pose Info.	Domain		Easy		Medium		Hard		Overall	
			spatial frequency		SR	SPL	SR	SPL	SR	SPL	SR	SPL
VGM [65]	✗	no	RGBD	✗	0.86	0.80	0.81	0.68	0.61	0.46	0.76	0.64
Neural Planner [64]	✓	global	RGBD	✗	0.72	0.41	0.65	0.39	0.42	0.27	0.60	0.36
NTS [63]	✓	global	RGBD	✗	0.87	0.65	0.58	0.38	0.43	0.26	0.63	0.43
ANS [41]	✓	global	RGBD	✗	0.74	0.21	0.68	0.23	0.30	0.11	0.58	0.18
Meta-Explore (homing)	✓	local	RGBD	SOS	0.82	0.61	0.83	0.61	0.70	0.48	0.78	0.57
Meta-Explore (localgoal)	✓	no	RGBD	SOS	0.94	0.84	0.88	0.63	0.71	0.18	0.84	0.55

Table A.5: **Evaluation results for Image-goal Navigation Task.**
(SR: success rate, SPL: success weighted by path length)

by the total number of episodes, E . SPL is calculated as $\frac{1}{E} \sum_{i=1}^E S_i \frac{l_i}{\max(p_i, l_i)}$. S_i denotes the success as a binary value. p_i and l_i denote the shortest path and actual path length for the i^{th} episode. For each task difficulty, SR and SPL are measured separately.

Experiment Results

Detailed comparisons with the baseline methods are shown in Table A.5. The results show that the continuous version Meta-Explore and SOS features help navigation and the exploitation mode provides corrections for misled exploration or undesirable actions. Compared with the exploration policy baseline VGM [65], Meta-Explore shows an enhancement in the overall success rate by 10.5%. The results imply that local goal search helps the agent escape from the current location when the agent recurrently explores a local area but cannot find the target location. Exploitation can reduce unnecessary exploration and help the agent reach the target goal before the maximum time horizon. Among two methods of exploitation, local goal search outperforms homing, presumably because of the noisy actuation model used in the simulator. Due to the noisy actions, the agent can hardly return to a previously visited location by directly reversing the action sequence.

Comparing our method with other graph-based hierarchical navigation methods, Meta-Explore outperforms ANS, Neural Planner, and NTS in the success rate.

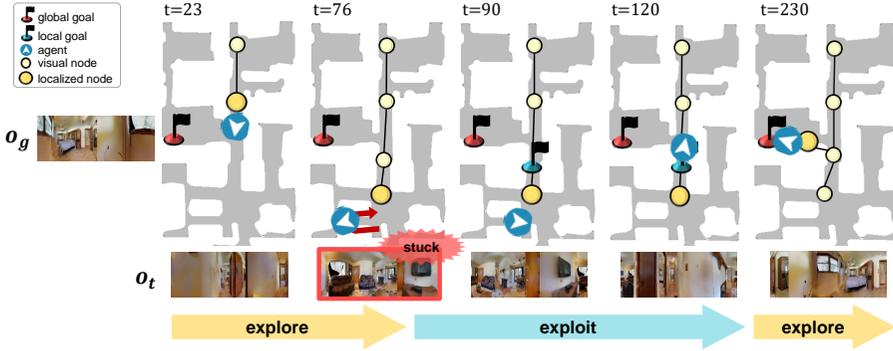


Figure A.7: **Experiment visualization for image-goal navigation task in continuous environment.** The mode selector detects stuck event at $t = 76$ and switches the explore mode to exploit mode. Then, the agent returns toward the local goal, which is chosen as a position nearby one of the nodes in the previously constructed graph.

Our model shows lower performance in SPL for *hard* episodes while the success rate is higher than the baselines. This implies that the exploitation mode of the proposed method allows the agent to explore more uncovered areas. Meanwhile, the proposed method appears to yield a positive impact for *easy* episodes, with the increase on both success rate by 9.3% and SPL by 5.0%. Specifically, our method outperforms ANS in terms of both success rate and SPL across all episodes. When compared to Neural Planner and NTS, our approach shows better performance in both success rate and SPL for *easy* and *medium* episodes, while outperforming Neural Planner and NTS in success rate for *hard* episodes. On the other hand, the proposed method shows lower SPL for *hard* episodes than NTS and Neural Planner. This implies that Meta-Explore tends to explore uncovered areas in both successful and unsuccessful episodes, which could be the result of using the SOS features to understand scenes. Comparing the proposed method using different exploitation methods (homing and local goal search) shows that searching for a local goal leads the agent to better escape from a local area. Figure A.7 shows a

Methods	Memory	Exploit	SR \uparrow	SPL \uparrow	OSR \uparrow	TL \downarrow	NE \downarrow
VLN-CE [11]	Rec	\times	32	30	40	8.64	7.37
HCM † [68]	Rec	\times	-	-	43	15.61	8.93
SASRA [69]	Semantic Map	\times	24	22	-	7.89	8.32
Conti-CMA † [67]	Top. Map	\times	41	35	51	10.90	6.20
Meta-Explore (Ours)	Top. Map	local goal	49	38	54	14.88	4.25

Table A.6: Evaluation results in the VLN-CE val unseen split.

simple scenario of image-goal navigation using Meta-Explore. The mode selector detects a regretful situation when the agent is recurrently exploring a local area but cannot find the target location. Hierarchical exploration via local goal search helps the agent overcome the situation and move toward the global goal in fixed time.

A.4.3 VLN in Continuous Domain

Image-goal navigation results in complex settings (continuous environments with noisy actions, max \sim 300 steps) imply that our model can be transferred to long-horizon VLN with noisy actions. We further extend the proposed method in continuous environments to solve the VLN-CE [11] task. In the VLN-CE [11] task, our agent constructs a topological map by using Conti-CMA [67] as a baseline to find reachable nodes (i.e., waypoints) and reuses the map in the exploitation mode. We compare our continuous version Meta-Explore with various navigation baselines¹: VLN-CE [11], HCM [68], SASRA [69], and Conti-CMA [67]. We evaluate algorithms using the success rate (SR), success weighted by inverse path length (SPL), oracle success rate (OSR), trajectory length (TL), and navigation error (NE), following the definitions of the evaluation metrics in the paper.

Experiment Results

Results in Table A.6 show that our method outperforms other baselines by at least 19.5% in the success rate, 8.6% in SPL, and 5.9% in OSR. We excluded the results

¹ \dagger indicates reproduced results.

of HCM for SR and SPL because HCM measures SR, SPL using oracle stop in the official code, which is not allowed in other baselines. We address that our model can be transferred to long-horizon (max. step 300) VLN with noisy actions in complex settings, as demonstrated by image-goal navigation results in Sec. A.4.2.

Bibliography

- [1] Jing Gu, Eliana Stefani, Qi Wu, Jesse Thomason, and Xin Eric Wang. Vision-and-language navigation: A survey of tasks, methods, and future directions. *arXiv preprint arXiv:2203.12667*, 2022.
- [2] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton Van Den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3674–3683, 2018.
- [3] Tao Chen, Saurabh Gupta, and Abhinav Gupta. Learning exploration policies for navigation. In *Proceedings of the International Conference on Learning Representations*, 2019.
- [4] Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer. Vision-and-dialog navigation. In *Proceedings of the Conference on Robot Learning*, pages 394–406, 2020.
- [5] Fengda Zhu, Yi Zhu, Vincent Lee, Xiaodan Liang, and Xiaojun Chang. Deep learning for embodied vision navigation: A survey. *arXiv preprint arXiv:2108.04097*, 2021.

- [6] Howard Chen, Alane Suhr, Dipendra Misra, Noah Snaveley, and Yoav Artzi. Touchdown: Natural language navigation and spatial reasoning in visual street environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12538–12547, 2019.
- [7] Alexander Ku, Peter Anderson, Roma Patel, Eugene Ie, and Jason Baldridge. Room-Across-Room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In *Conference on Empirical Methods for Natural Language Processing*, 2020.
- [8] Vihan Jain, Gabriel Magalhaes, Alexander Ku, Ashish Vaswani, Eugene Ie, and Jason Baldridge. Stay on the path: Instruction fidelity in vision-and-language navigation. In *Transactions of the Association for Computational Linguistics*, pages 1862–1872, Florence, Italy, July 2019. Association for Computational Linguistics.
- [9] An Yan, Xin Eric Wang, Jiangtao Feng, Lei Li, and William Yang Wang. Cross-lingual vision-language navigation. *arXiv preprint arXiv:1910.11301*, 2019.
- [10] Keji He, Yan Huang, Qi Wu, Jianhua Yang, Dong An, Shuanglin Sima, and Liang Wang. Landmark-rxr: Solving vision-and-language navigation with fine-grained alignment supervision. In *Proceedings of the International Conference on Neural Information Processing Systems*, volume 34, pages 652–663. Curran Associates, Inc., 2021.
- [11] Jacob Krantz, Erik Wijmans, Arjun Majundar, Dhruv Batra, and Stefan Lee. Beyond the nav-graph: Vision and language navigation in continuous environments. In *Proceedings of the European Conference on Computer Vision*, 2020.

- [12] Piotr Mirowski, Andras Banki-Horvath, Keith Anderson, Denis Teplyashin, Karl Moritz Hermann, Mateusz Malinowski, Matthew Koichi Grimes, Karen Simonyan, Koray Kavukcuoglu, Andrew Zisserman, et al. The streetlearn environment and dataset. *arXiv preprint arXiv:1903.01292*, 2019.
- [13] Karl Moritz Hermann, Mateusz Malinowski, Piotr Mirowski, Andras Banki-Horvath, Keith Anderson, and Raia Hadsell. Learning to follow directions in street view. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11773–11781, Apr. 2020.
- [14] Arun Balajee Vasudevan, Dengxin Dai, and Luc Van Gool. Talk2nav: Long-range vision-and-language navigation with dual attention and spatial memory. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, volume 129, pages 246–266. Springer, 2021.
- [15] Dipendra Misra, Andrew Bennett, Valts Blukis, Eyvind Niklasson, Max Shatkhin, and Yoav Artzi. Mapping instructions to actions in 3D environments with visual goal prediction. In *Proceedings of the Conference on Empirical Methods for Natural Language Processing*, pages 2667–2678, Brussels, Belgium, October-November 2018. Association for Computational Linguistics.
- [16] Ta-Chung Chi, Minmin Shen, Mihail Eric, Seokhwan Kim, and Dilek Hakkani-tur. Just ask: An interactive learning framework for vision and language navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2459–2466, 2020.
- [17] Yi Wu, Yuxin Wu, Georgia Gkioxari, and Yuandong Tian. Building generalizable agents with a realistic and rich 3d environment. *arXiv preprint arXiv:1801.02209*, 2018.

- [18] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1–10, 2018.
- [19] Fengda Zhu, Xiwen Liang, Yi Zhu, Qizhi Yu, Xiaojun Chang, and Xiaodan Liang. Soon: Scenario oriented object navigation with graph-based exploration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12689–12699, 2021.
- [20] Yuankai Qi, Qi Wu, Peter Anderson, Xin Wang, William Yang Wang, Chunhua Shen, and Anton van den Hengel. Reverie: Remote embodied visual referring expression in real indoor environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9982–9991, 2020.
- [21] Khanh Nguyen and Hal Daumé III. Help, anna! visual navigation with natural multimodal assistance via retrospective curiosity-encouraging imitation learning. In *Proceedings of the Conference on Empirical Methods for Natural Language Processing*, November 2019.
- [22] Khanh Nguyen, Debadepta Dey, Chris Brockett, and Bill Dolan. Vision-based navigation with language-based assistance via imitation learning with indirect intervention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2019.
- [23] Alane Suhr, Claudia Yan, Jack Schluger, Stanley Yu, Hadi Khader, Marwa Mouallem, Iris Zhang, and Yoav Artzi. Executing instructions in situated collaborative interactions. In *Proceedings of the Conference on Empirical Methods for Natural Language Processing*, pages 2119–2130, 2019.

- [24] Chih-Yao Ma, Jiasen Lu, Zuxuan Wu, Ghassan AlRegib, Zsolt Kira, Richard Socher, and Caiming Xiong. Self-monitoring navigation agent via auxiliary progress estimation. In *Proceedings of the International Conference on Learning Representations*, 2019.
- [25] Chih-Yao Ma, Zuxuan Wu, Ghassan AlRegib, Caiming Xiong, and Zsolt Kira. The regretful agent: Heuristic-aided navigation through progress estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6732–6740, 2019.
- [26] Fengda Zhu, Yi Zhu, Xiaojun Chang, and Xiaodan Liang. Vision-language navigation with self-supervised auxiliary reasoning tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10012–10022, 2020.
- [27] Shizhe Chen, Pierre-Louis Guhur, Makarand Tapaswi, Cordelia Schmid, and Ivan Laptev. Think global, act local: Dual-scale graph transformer for vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16537–16547, 2022.
- [28] Zhiwei Deng, Karthik Narasimhan, and Olga Russakovsky. Evolving graphical planner: Contextual global planning for vision-and-language navigation. In *Proceedings of the International Conference on Neural Information Processing Systems*, volume 33, pages 20660–20672, 2020.
- [29] Hanqing Wang, Wenguan Wang, Wei Liang, Caiming Xiong, and Jianbing Shen. Structured scene memory for vision-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8455–8464, June 2021.

- [30] Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, and Trevor Darrell. Speaker-follower models for vision-and-language navigation. In *Proceedings of the International Conference on Neural Information Processing Systems*, 2018.
- [31] Xiujun Li, Chunyuan Li, Qiaolin Xia, Yonatan Bisk, Asli Celikyilmaz, Jianfeng Gao, Noah A Smith, and Yejin Choi. Robust navigation with language pretraining and stochastic sampling. In *Proceedings of the Conference on Empirical Methods for Natural Language Processing*, pages 1494–1499, 2019.
- [32] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Proceedings of the International Conference on Neural Information Processing Systems*, 2019.
- [33] Pierre-Louis Guhur, Makarand Tapaswi, Shizhe Chen, Ivan Laptev, and Cordelia Schmid. Airbert: In-domain pretraining for vision-and-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1634–1643, 2021.
- [34] Yicong Hong, Qi Wu, Yuankai Qi, Cristian Rodriguez-Opazo, and Stephen Gould. Vln \odot bert: A recurrent vision-and-language bert for navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1643–1653, 2021.
- [35] Alexander Pashevich, Cordelia Schmid, and Chen Sun. Episodic transformer for vision-and-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15942–15952, 2021.

- [36] Shizhe Chen, Pierre-Louis Guhur, Cordelia Schmid, and Ivan Laptev. History aware multimodal transformer for vision-and-language navigation. In *Proceedings of the International Conference on Neural Information Processing Systems*, pages 5834–5847, 2021.
- [37] Liyiming Ke, Xiujun Li, Yonatan Bisk, Ari Holtzman, Zhe Gan, Jingjing Liu, Jianfeng Gao, Yejin Choi, and Siddhartha Srinivasa. Tactical rewind: Self-correction via backtracking in vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6741–6749, 2019.
- [38] James G March. Exploration and exploitation in organizational learning. *Organization science*, 2(1):71–87, 1991.
- [39] Adrien Ecoffet, Joost Huizinga, Joel Lehman, Kenneth O Stanley, and Jeff Clune. First return, then explore. *Nature*, 590(7847):580–586, 2021.
- [40] Miruna Pislari, David Szepesvari, Georg Ostrovski, Diana L Borsa, and Tom Schaul. When should agents explore? In *Proceedings of the International Conference on Learning Representations*, 2022.
- [41] Devendra Singh Chaplot, Dhiraj Gandhi, Saurabh Gupta, Abhinav Gupta, and Ruslan Salakhutdinov. Learning to explore using active neural slam. In *Proceedings of the International Conference on Learning Representations*, 2020.
- [42] Meera Hahn, Devendra Singh Chaplot, Shubham Tulsiani, Mustafa Mukadam, James M Rehg, and Abhinav Gupta. No rl, no simulation: Learning to navigate without navigating. In *Proceedings of the International Conference on Neural Information Processing Systems*, volume 34, pages 26661–26673. Curran Associates, Inc., 2021.

- [43] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [44] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proceedings of the International Conference on Learning Representations*, 2021.
- [45] Wolfgang Stürzl and Hanspeter A Mallot. Efficient visual homing based on fourier transformed panoramic images. *Robotics and Autonomous Systems*, 54(4):300–313, 2006.
- [46] Jean Serra. *Mathematical Morphology*, pages 1–16. Springer International Publishing, Cham, 2020.
- [47] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *Proceedings of the Conference on Empirical Methods for Natural Language Processing*, pages 5100–5111, 2019.
- [48] Stéphane Ross, Geoffrey J Gordon, and J Andrew Bagnell. No-regret reductions for imitation learning and structured prediction. *arXiv preprint arXiv:1011.0686*, 2010.
- [49] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *Proceedings of the International Conference on Machine Learning*, pages 1928–1937, 2016.

- [50] Peter Anderson, Angel Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, et al. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*, 2018.
- [51] Xin Wang, Qiuyuan Huang, Asli Celikyilmaz, Jianfeng Gao, Dinghan Shen, Yuan-Fang Wang, William Yang Wang, and Lei Zhang. Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6629–6638, 2019.
- [52] Andrew Szot, Alex Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah Maestre, Mustafa Mukadam, Devendra Chaplot, Oleksandr Maksymets, Aaron Gokaslan, Vladimir Vondrus, Sameer Dharur, Franziska Meier, Wojciech Galuba, Angel Chang, Zsolt Kira, Vladlen Koltun, Jitendra Malik, Manolis Savva, and Dhruv Batra. Habitat 2.0: Training home assistants to rearrange their habitat. In *Proceedings of the International Conference on Neural Information Processing Systems*, 2021.
- [53] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2425–2433, 2015.
- [54] Xiangru Lin, Guanbin Li, and Yizhou Yu. Scene-intuitive agent for remote embodied visual grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7036–7045, June 2021.

- [55] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [56] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2019.
- [57] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning*, pages 8748–8763, 2021.
- [58] S. Chen et al. Think global, act local: Dual-scale graph transformer for vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [59] Edsger W Dijkstra. A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1):269–271, 1959.
- [60] Fei Xia, Amir R Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: Real-world perception for embodied agents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9068–9079, 2018.
- [61] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.

- [62] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 633–641, 2017.
- [63] Devendra Singh Chaplot, Ruslan Salakhutdinov, Abhinav Gupta, and Saurabh Gupta. Neural topological slam for visual navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12875–12884, 2020.
- [64] Edward Beeching, Jilles Dibangoye, Olivier Simonin, and Christian Wolf. Learning to plan with uncertain topological maps. In *Proceedings of the European Conference on Computer Vision*, pages 473–490. Springer, 2020.
- [65] Obin Kwon, Nuri Kim, Yunho Choi, Hwiyeon Yoo, Jeongho Park, and Songhwai Oh. Visual graph memory with unsupervised representation for visual navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15890–15899, 2021.
- [66] Junnan Li, Pan Zhou, Caiming Xiong, and Steven Hoi. Prototypical contrastive learning of unsupervised representations. In *Proceedings of the International Conference on Learning Representations*, 2021.
- [67] Yicong Hong, Zun Wang, Qi Wu, and Stephen Gould. Bridging the gap between learning in discrete and continuous environments for vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [68] Muhammad Zubair Irshad, Chih-Yao Ma, and Zsolt Kira. Hierarchical cross-modal agent for robotics vision-and-language navigation. In *2021 IEEE In-*

ternational Conference on Robotics and Automation (ICRA), pages 13238–13246. IEEE, 2021.

- [69] Muhammad Zubair Irshad, Niluthpol Chowdhury Mithun, Zachary Seymour, Han-Pang Chiu, Supun Samarasekera, and Rakesh Kumar. Semantically-aware spatio-temporal reasoning agent for vision-and-language navigation in continuous environments. *Proceedings of the International Conference on Pattern Recognition (ICPR)*, 2021.

요약

시각-언어 네비게이션(vision-and-language navigation, VLN)은 환경에 대한 시각 정보를 기반으로 사람으로부터 주어진 자연어 명령을 따라 길을 찾는 지능형 에이전트를 개발하는 분야이다. 시각-언어 네비게이션 분야의 주요 과제는 학습에 사용되지 않은 새로운 환경에서도 주어진 명령을 효과적으로 이해하는 것이다. 기존 시각-언어 네비게이션 방법론들의 한계는 에이전트가 잘못된 행동을 수행하여 올바른 길로부터 벗어났을 때, 불필요한 영역을 탐색하게 되거나 돌이킬 수 없는 경로로 진입하는 것이다. 본 논문에서는 에이전트가 잘못된 행동을 수행하였을 때, 이를 감지하고 원래의 경로를 다시 찾아갈 수 있도록 하는 익스플로이테이션 정책(exploitation policy)을 이용하는 메타-익스플로어라는 계층적 탐색 방법을 제안한다. 본 논문에서는 에이전트가 방문한 적은 없지만 관찰은 가능했던 장소들 중 가장 적절한 단기 목표 지점(local goal)을 설정하는 방식을 제안하며, 해당 방식이 기존 논문들의 이전에 방문했던 장소로 되돌아가는 회귀적 방식보다 성능적으로 우수함을 보인다. 또한, 잘못된 탐색을 돌이킬 때에 물체와 같은 의미론적 단서를 가지고 단기 목표 지점을 찾아야 한다는 것을 강조한다. 본 논문에서 제안하는 방식의 핵심은 에이전트 주변에 배치된 물체들을 주파수 영역의 정보로 변환하여 이해하는 것이다. 구체적으로는, 객체별 이차원 푸리에 변환을 수행하여 구하는 장면 객체 스펙트럼(scene object spectrum, SOS)이라는 새로운 시각적 표현 방식을 제안한다. 본 논문에서는 장면 객체 스펙트럼을 계층적 탐색에 활용하였을 때 에이전트가 단기 목표 지점을 보다 성공적으로 찾아내어 잘못된 경로를 수정할 수 있음을 보인다. 본 제안 방식은 R2R, REVERIE, SOON 총 세 가지 시각-언어 네비게이션 벤치마크에서 성능을 평가하였다. 본 논문에서 제안하는 메타-익스플로어는 기존 방식들보다 높은 성능뿐만 아니라 매우 우수한 일반화 성능을 통해 에이전트가 경험해보지 못한 환경에서도 성공적으로 동작함을 보인다. 특히, SOON 벤치마크에서는 기존 최고 수준의 알고리즘 대비 네비게이션

성공률이 17.1%, 경로 길이 대비 성공률(success weighted by path length, SPL)이 20.6% 향상되었다.

주요어: 시각 기반 네비게이션, 시각-언어 네비게이션, 컴퓨터 비전, 물체 탐지, 계층적 탐색, 푸리에 변환, 크로스모달 트랜스포머

학번: 2021-24215