



ESceme: Vision-and-Language Navigation with Episodic Scene Memory

Qi Zheng^{1,2} · Daqing Liu³ · Chaoyue Wang³ · Jing Zhang² · Dadong Wang⁴ · Dacheng Tao²

Received: 17 July 2023 / Accepted: 24 June 2024 / Published online: 26 July 2024
© The Author(s) 2024

Abstract

Vision-and-language navigation (VLN) simulates a visual agent that follows natural-language navigation instructions in real-world scenes. Existing approaches have made enormous progress in navigation in new environments, such as beam search, pre-exploration, and dynamic or hierarchical history encoding. To balance generalization and efficiency, we resort to memorizing visited scenarios apart from the ongoing route while navigating. In this work, we introduce a mechanism of Episodic Scene memory (ESceme) for VLN that wakes an agent's memories of past visits when it enters the current scene. The episodic scene memory allows the agent to envision a bigger picture of the next prediction. This way, the agent learns to utilize dynamically updated information instead of merely adapting to the current observations. We provide a simple yet effective implementation of ESceme by enhancing the accessible views at each location and progressively completing the memory while navigating. We verify the superiority of ESceme on short-horizon (R2R), long-horizon (R4R), and vision-and-dialogue (CVDN) VLN tasks. Our ESceme also wins first place on the CVDN leaderboard. Code is available: <https://github.com/qizhust/esceme>.

Keywords Vision-and-language navigation · Embodied AI · Scene memory · Episodic memory · Multimodal fusion

1 Introduction

With breakthroughs in computer vision and natural language understanding, the embodiment hypothesis that an intelligent agent is born from its interaction with environments (Smith & Gasser, 2005) is now attracting more and more attention to embodied AI tasks such as vision-and-language navigation (VLN). VLN is firstly defined in Anderson et al. (2018b) towards the goal of a robot carrying out general verbal instructions, where an agent is required to follow natural-language instructions based on what it sees and adapt to previously unseen environments. VLN has developed various

settings, such as fine-grained and short-horizon navigation (e.g., R2R Anderson et al., 2018b and RxR Ku et al., Ku et al., 2020a), long-horizon navigation (e.g., R4R Jain et al., 2019), vision-and-dialogue navigation (e.g., CVDN Thomason et al. 2020), and navigation with high-level instructions (e.g., REVERIE Qi et al., 2020b). Compared with non-embodied VL tasks such as visual question answering (Antol et al., 2015) and visual captioning (Chen et al., 2015; Xu et al., 2016), VLN agents suffer from domain shifts and changing observations during multi-step decision-making in the scenarios.


A vanilla Seq2Seq pipeline (Anderson et al., 2018b) that implicitly encodes path history with LSTMs (Hochreiter & Schmidhuber, 1997) shows moderate navigating ability. Since then, VLN performance has been considerably improved by pre-training (Hao et al., 2020; Hong et al., 2021; Chen et al., 2021b; Qiao et al., 2022), data augmentations (Fried et al., 2018; Tan et al., 2019; Li et al., 2022a), and algorithms that explicitly track past decisions along the trajectory (Chen et al., 2021b; Wang et al., 2021; Chen et al., 2022b). These methods learn enhanced representations by training VLN agents in each episode but ignore the dynamics of navigating over the whole data. Different strategies, including modified beam search (Fried et al., 2018) and pre-

Communicated by Dima Damen.

✉ Jing Zhang
jingzhang.cv@gmail.com

Qi Zheng
qiz@szu.edu.cn

- ¹ College of Electronics and Information Engineering, Shenzhen University, 518000 Shenzhen, China
- ² School of Computer Science, The University of Sydney, Sydney, NSW 2008, Australia
- ³ JD Explore Academy, Beijing, China
- ⁴ DATA 61, CSIRO, Sydney, NSW 2122, Australia



Instruction 1: Walk toward and past the ping pong table. Continue in that direction until you reach the sofa with two green pillows. Stop in front of the table that is in front of that sofa.

Instruction 2: Walk through the office past the work tables. Go to the right and stop next to the white bookshelf outside the conference room.

Fig. 1 The longer blue trajectory shows an agent carrying out instruction 1. The next time, the agent enters this scene to conduct the second instruction along the shorter red path. EScheme allows it to recall the visited nodes (i.e., the blue ones B_1 and B_3) at where it is standing (A) and choose the neighboring node B_1 that will see “the white bookshelf” in one more step at C. Finally, it navigates towards the red dash route and reaches the target (Color figure online)

exploration (Wang et al., 2019; Tan et al., 2019; Majumdar et al., 2020; Zhu et al., 2020a), are devised to specifically increase adaptation to unseen environments at the cost of efficiency. Specifically, beam search significantly extends route length and involves much more interactions with the environment; pre-exploration takes extra steps to gather information and train the agent with auxiliary objectives before it can conduct given instructions. Such strategies incur burdensome time and computational expenses in practical usage.

In this work, we propose a navigation mechanism with Episodic Scene memory (EScheme) to balance generalization and efficiency by exploiting the dynamics of navigating all the episodes. EScheme requires no extra annotations or heavy computation and is agent-agnostic. We encode observation, instruction, and path history separately and update the scene memory during navigation via candidate enhancing. By preserving the memory among episodes, EScheme envisions the agent seeing a bigger picture in each decision. This way, the agent learns to utilize dynamically updated information instead of merely adapting to the current observations. Then during inference, it predicts actions with the progressively completed memory. A demonstration is shown in Fig. 1. When carrying out an instruction at Location A, the agent is to select one from the adjacent nodes B_1 – B_5 to navigate. It recalls the episodic scene memory, i.e., the blue route of a completed trajectory, and chooses Node B_1 that will see “the white bookshelf” in one more step at C.

We verify the superiority of EScheme in short-horizon navigation with fine-grained instruction (R2R), long-horizon navigation (R4R), and vision-and-dialog navigation (CVDN). We find that EScheme notably benefits navigation with longer routes (R4R and CVDN), promoting both successful reaching and path fidelity. Our method achieves the highest Goal Progress in the CVDN challenge. Besides a fair comparison with existing approaches under a single run, we test the per-

formance with an approximately complete memory, where the agent fully updates its scene memory in the first round of navigation over all the episodes. We denote it as EScheme*, which serves as the upper bound of EScheme. We observe a further improvement in EScheme*, which indicates better-completed memory magnifies the advantage of EScheme. We hope this work can inspire further explorations in modeling episodic scene memory for VLN.

Since EScheme does not introduce any extra time or steps before following the instruction in inference, it is fair to compare it with its counterparts in the single-run setting. Very different from pre-exploration optimizing the parameters of an agent before solving the task, EScheme only renews its episodic memory while conducting instructions and requires no back-propagation operations. Moreover, EScheme neither involves beam search nor changes the local action space in sequential decision-making. These properties make EScheme both efficient and effective in reality use. Our contributions are summarized as follows:

- We devise the first navigation mechanism with episodic scene memory (EScheme) for VLN to balance generalization and efficiency.
- We provide a simple yet effective implementation of EScheme via candidate enhancing, tested with two navigation architectures and two inferring strategies.
- We verify the superiority of EScheme in short-horizon (R2R), long-horizon (R4R), and vision-and-dialog (CVDN) navigation, and achieve a new state-of-the-art.

2 Related Work

2.1 Vision-and-Language Navigation

Since Anderson et al. (2018b) defined the VLN task and provided an LSTM-based sequence-to-sequence baseline (Seq2Seq), numerous approaches have been developed. A branch of methods improves navigation via data augmentation, such as SF (Fried et al., 2018), EnvDrop (Tan et al., 2019), and EnvEdit (Li et al., 2022a). As for agent training, Wang et al. (2018) model the environment to provide planned-ahead information during navigation. RCM (Wang et al., 2019) provides an intrinsic reward for reinforcement learning via an instruction-trajectory matching critic. Wang et al. (2020b) jointly train an agent on VLN and vision-dialog navigation (MT-RCM+EnvAg). To fully use available semantic information in the environment, AuxRN (Zhu et al., 2020a) devises four self-supervised auxiliary reasoning tasks. TDSTP (Zhao et al., 2022) introduces an extra target location estimation during finetuning to achieve reliable path planning. Many methods explore more effective feature representations and architectures, such as PTA (Cornia &

Cucchiara, 2019), OAAM (Qi et al., 2020a), NvEM (An et al., 2021), RelGraph (Hong et al., 2020), MTVM (Lin et al., 2022b), and SEvol (Chen et al., 2022a).

Some methods construct and reason about a graph of navigation while conducting an episode, such as NTS (Chaplot et al., 2020) and RECON (Shah et al., 2022) in the ImageGoal space and ETPNav (An et al., 2023) and CMTP (Chen et al., 2021a) in the VLN space. VLN-SIG (Li & Bansal, 2023) adds the tasks of generating semantics for future navigation views in pre-training and fine-tuning, and contributes to a more powerful agent backbone. KERM (Li et al., 2023) introduces knowledge described by text to aid action prediction, which is useful mainly in seen environments. GridMM (Wang et al., 2023) builds a grid memory with fine-grained features and adopts a global action space, which improves the success rate but suffers from a much longer trajectory length.

Inspired by the breakthrough of large-scale pre-trained BERT (Kenton & Toutanova, 2019) in natural language processing tasks, PRESS (Li et al., 2019) replaces RNNs with pre-trained BERT to encode instructions and achieves a non-trivial improvement in unseen environments. PREVELENT (Hao et al., 2020) pre-trains BERT from scratch using image-text-action triplets and further boosts the performance. RecBERT (Hong et al., 2021) integrates a recurrent unit into a BERT model to be time-aware. Chen et al. (2021b) propose the first VLN network that allows a sequence of historical memory and can be optimized end-to-end (HAMT). HOP (Qiao et al., 2022) designs trajectory order modeling and group order modeling tasks to model temporal order information in pre-training. CSAP (Wu et al., 2022) proposes trajectory-conditioned masked fragment modeling and contrastive semantic-alignment modeling tasks for pre-training. ADAPT (Lin et al., 2022a) explicitly learns action-level modality alignment with action prompts. There are also some works specially designed for vision-and-dialog navigation, such as VISITRON (Shrivastava et al., 2022), SCoA (Zhu et al., 2021), and CMN (Zhu et al., 2020b).

The differences between the proposed ESceme and previous graph- or map-construction approaches are twofold. First, they construct path-level memory along a route in a single conduction. ESceme maintains scene-level memory from multiple episodes in the same scenario. Second, they use the path-level memory in planning by extending the agent's action space from local to global. ESceme does not change the agent's action space. Instead, it improves navigation by increasing the information of each node, which is the core idea that makes ESceme perform better than path-level memory methods (e.g., EGP and SSM). Scene memory is also studied in other works (Datta et al., 2022; Georgakis et al., 2022; Li et al., 2022b; Krantz et al., 2023; Vasudevan et al., 2021; Gupta et al., 2017).

The method that is the most closely related to ours is IVLN (Krantz et al., 2023). Our setting is actually identical to

IVLN, which reorganizes episodes into tours. We store scene IDs during inference instead of explicitly organizing episodes according to their IDs, yet the two ways result in the same effect. Although both explore the impact of episodic memory and compare with the same baseline model, i.e., HAMT, our work provides a more effective design of the memory mechanism and obtains better performance due to candidate enhancing. Instead, the episodic memory in IVLN encodes the memory map as a whole and observes even worse results than the path-level memory baseline (cf. Table 2 in IVLN).

2.2 Exploration Strategies in VLN

As the navigation graph is pre-defined in discrete VLN, diverse strategies are adopted other than the regularly used single-run. For example, Fried et al. (2018) modifies the standard beam search to select the final navigation route, which notably increases navigation success at the cost of unbearable trajectory lengths. More efficient pre-exploration methods are studied. For instance, a progress monitor is trained to discard unfinished trajectories during inference (Ma et al., 2019a). Ma et al. (2019b) learn a regret module to decide when to backtrack. Ke et al. (2019) compare partial paths with global information considered and backtrack only when necessary. AcPercep (Wang et al., 2020a) learns an exploration policy to gather visual information for navigation. Although these methods improve searching efficiency, they heavily depend on manually designed or heuristic rules. Deng et al. (2020) define a global action space for the first time and build a graphical representation of the environment for elegant exploration/backtracking. Wang et al. (2021) extend EnvDrop (Tan et al., 2019) with an external structured scene memory (SSM) to promote exploration in the global action space.

Pre-exploration, which allows an agent to pre-explore unseen environments before navigating, is first introduced in Wang et al. (2019) as a setting different from single-run and beam search. The obtained information functions in diverse ways. RCM (Wang et al., 2019) uses the exploration experience in self-supervised imitation learning. EnvDrop (Tan et al., 2019) exploits the environment information for data augmentation via back-translation. VLN-BERT (Majumdar et al., 2020) provides the agent with a global view for optimal route selection. AuxRN (Zhu et al., 2020a) finetunes the agent in unseen environments with auxiliary tasks.

3 Method

3.1 Problem Formulation

Given an instruction X_i , e.g., “*Turn around and walk to the right of the room...*”, an agent starts from the initial location

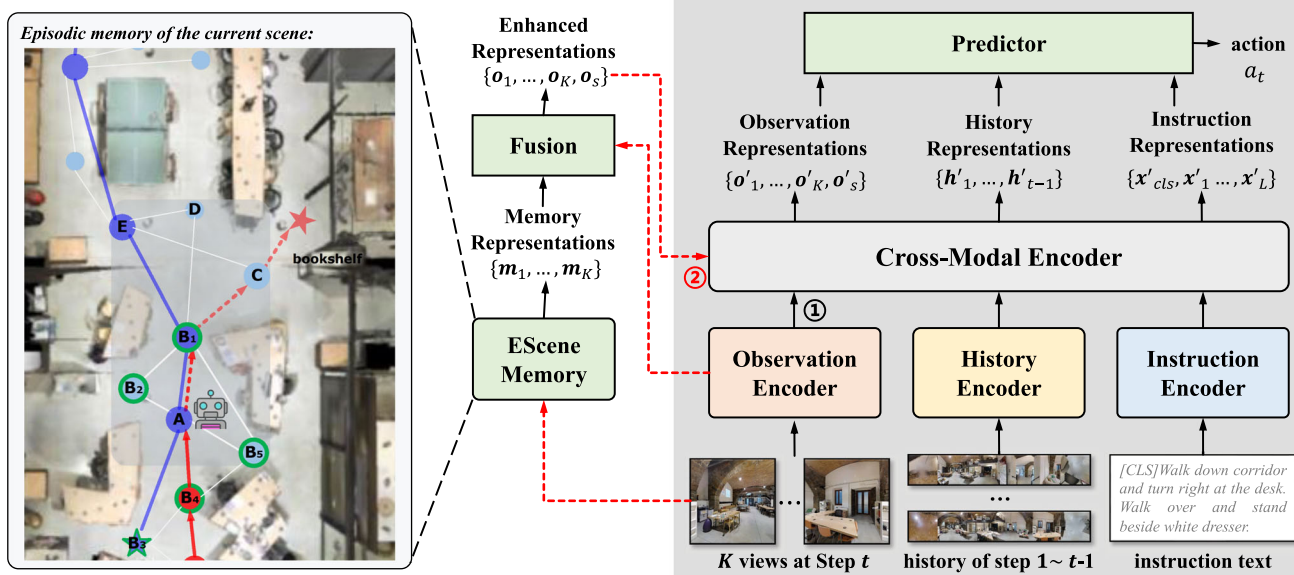


Fig. 2 An overview of the Episodic Scene memory mechanism for VLN. On the left is partial episodic memory for the current scene, which gets updated in navigation 1) following the previous instruction, i.e., the blue route $B_3 \rightarrow A \rightarrow B_1 \rightarrow E$, and 2) following the current instruction from Step 1 to $t - 1$, i.e., the solid trajectory $B_4 \rightarrow A$. The cyan nodes B_2, B_5, C , and D are those viewed but not visited. The shadow box shows the memory of node B_1 , which has six adjacent neighbors, i.e., A, B_2, B_5, C, D , and E . The integration of these nodes consists of the memory of B_1 . At Step t , the agent stands at Node A and is expected to choose one node from B_1 to B_5 . Given obser-

vation from K views, each view retrieves its memory in EScene and produces $\{m_1, \dots, m_K\}$. The memory representation then fuses with original encoded observations, which yields $\{o_1, \dots, o_K, o_s\}$. o_s is the representation for STOP. The enhanced observations, instruction text, and history from Step 1 to $t - 1$ compose the input to a navigation network to predict the action $a_t = i \in \{1, \dots, K, s\}$. Generally, a navigation network uses the encoded features of the original K views as the input to the cross-modal encoder, i.e., the output ①. Our EScene exploits the enhanced observations from ② (Color figure online)

of route R_i . It observes a panoramic view of the environment Y_i . The panoramic view consists of $K=36$ single viewpoints, each of which is accompanied by an orientation (θ, ϕ) indicating heading and elevation and a binary navigable signal. The agent selects a viewpoint from the navigable ones and moves to the next location with new observations. This process repeats until the agent takes the STOP action.

In a regular VLN task, there is a set of training samples $\mathcal{D} = \{(Y_1, X_1, R_1), \dots, (Y_{N_1}, X_{N_1}, R_{N_1})\}$, where (X_i, R_i) is the instruction-route pair in an environment Y_i . The set $\{Y_1, \dots, Y_{N_1}\}$ denotes the seen environments during training. An agent is expected to learn navigation with \mathcal{D} and carry out instructions in unseen scenarios given by $\mathcal{D}^u = \{(Y_1^u, X_1), \dots, (Y_{N_2}^u, X_{N_2})\}$. The set $\{Y_1^u, \dots, Y_{N_2}^u\}$ denotes the unseen environments for test.

For a sequence prediction problem, history is an important source of information apart from observations and instructions. The shadow part in Fig. 2 shows a decision step by a general navigation approach that follows the pretraining-finetuning branch and encodes path history, represented by HAMT(Chen et al., 2021b). We denote the vanilla features of K single views extracted by the observation encoder as $\{f_1, \dots, f_K, f_s\}$, which can be obtained by concatenating the separate features of encoded RGB images and orientations. f_s is appended to allow a STOP action. Together

with history representations $\{h_1, \dots, h_{t-1}\}$ from the history encoder and text representations $\{x_{cls}, x_1, \dots, x_L\}$ from the instruction encoder, the features of the observations ① are input into a cross-modal encoder for multi-modal fusion. A predictor block takes in the cross-modal representations $\{o'_1, \dots, o'_K, o'_s\}$, $\{h'_1, \dots, h'_{t-1}\}$, and $\{x'_{cls}, x'_1, \dots, x'_L\}$ to predict action a_t .

Due to potential differences between seen and unseen environments, such as the appearance and layout of the scenario and the display of objects, an agent trained in the above way suffers from decreased decision ability. The mistake accumulates along the path, which incurs a heavy drop in successful navigation in new environments. Since strategies such as pre-exploration and beam search that exploit extra clues in a new scene are too expensive for a deployed robot, we propose a mechanism of episodic scene memory to balance accuracy and efficiency. Figure 2 provides an overview of the proposed EScene mechanism. By retrieving episodic memory for the K views at Step t , EScene replaces the vanilla encoded observations with enhanced representations for cross-modal encoding and action prediction, i.e., ① → ②. In the following sections, we detail how to build the episodic scene memory and promote observations with the memory in navigation.

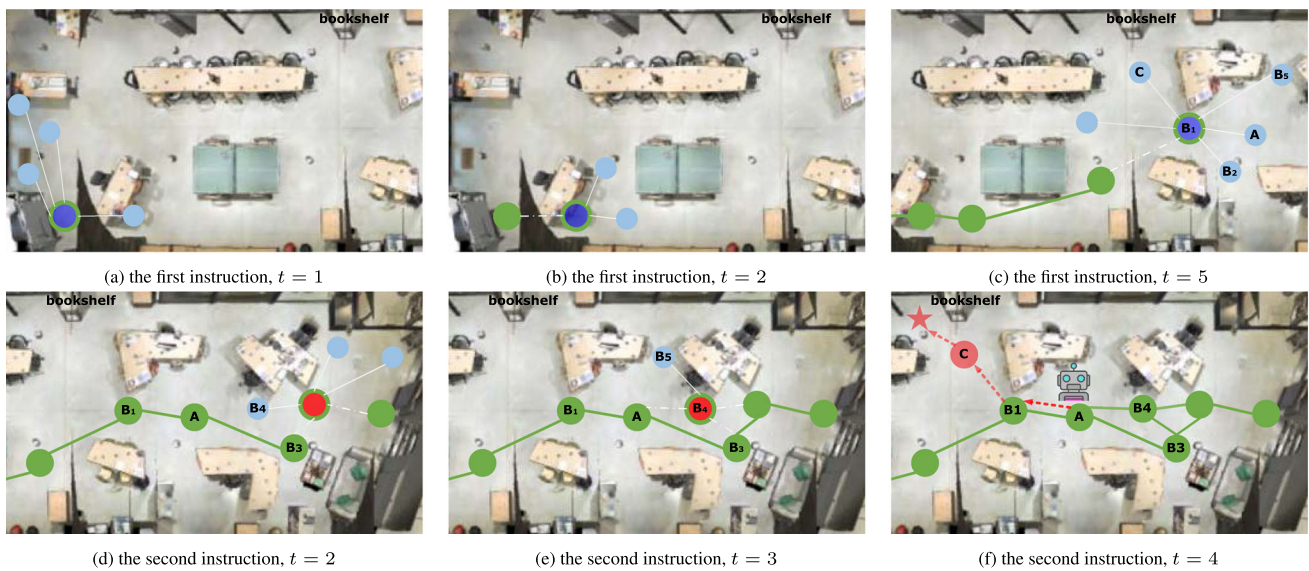


Fig. 3 Episodic memory construction of a scene during navigation. EScene at the beginning of each time step is presented in the figures, which comprises green nodes and edges and is empty at the beginning of $t = 1$. The blue nodes indicate the current location of following the first instruction at each time step, and the red ones correspond to the second instruction. The small cyan nodes mark the remaining navigable

viewpoints of the current location. Nodes with green boundary are the chosen viewpoints in each time step. EScene at the end of that time step is updated by the node with green boundary and the dashed lines connected to its existing nodes. Please refer to Fig. 1 for a complete global graph of the scene, which is unavailable to the agent either in navigation or EScene construction (Color figure online)

3.2 Episodic Scene Memory Construction

We initialize the episodic memory of Scene Y with an empty graph $\mathcal{G}_Y^{(0)} = (\mathcal{V}_Y^{(0)} = \emptyset, \mathcal{E}_Y^{(0)} = \emptyset)$ if an agent has never seen the scene. Namely, for the first instruction in Scene Y , an agent starts navigation with an empty episodic memory. As shown in Fig. 3a, the start location has four neighbors and is added to \mathcal{G}_Y at the end of $t=1$ by $\mathcal{V}_Y^{(1)} = \{V_1\}$. Node feature \mathbf{m}_{V_1} is an integration of its neighbors,

$$\mathbf{m}_{V_1} = \text{pooling}(\mathbf{f}_{V_{1,i}}), \quad (1)$$

where $i \in \{1, 2, 3, 4\}$ in Fig. 3a, $\mathbf{f}_{V_{1,i}} \in \mathbb{R}^d$ is d -dim plain representations of the i -th neighbor view from the observation encoder, and $\mathbf{m}_{V_1} \in \mathbb{R}^d$. The pooling function can be either *max* or *mean* pooling along the number of features. It is worth noting that obtaining $\mathbf{f}_{V_{1,i}}$ does not involve extra computation since these features have been calculated in offline feature extraction. The agent selects its right neighbor to navigate, and at the end of $t=2$, the visited node is added to \mathcal{G}_Y by $\mathcal{V}_Y^{(2)} = \{V_1, V_2\}$, $\mathcal{E}_Y^{(2)} = \{e_{12}\}$, with node feature \mathbf{m}_{V_2} calculated similarly as Eq. (1). We set all edges $e_{jk} = 1$.

While following the first instruction, the agent updates its episodic scene memory \mathcal{G}_Y accordingly, i.e., the green nodes and edges in Fig. 3b, c. At the end of $t = 5$, $\mathcal{V}_Y^{(5)} = \{V_1, V_2, \dots, V_5\}$, $\mathcal{E}_Y^{(5)} = \{e_{12}, e_{23}, e_{34}, e_{45}\}$. When the agent is directed to the second instruction in Scene Y , its memory in previous visits is preserved in \mathcal{G}_Y and is updated

at the end of each time step accordingly as Fig. 3d, e demonstrate. In Fig. 3f, since the agent's location A has been added to EScene in conducting the first instruction, there is no update to \mathcal{G}_Y . The agent stores episodic memory for each scene separately in similar ways. Therefore, we omit the subscript Y for simplicity.

3.3 EScene Navigation by Candidate Enhancing

In addition to information from instruction, current observation, and route history, an agent can refer to its episodic scene memory in decision-making at each step.

Since the node representation in EScene integrates information within the neighborhood, it is expected to envision the agent with a bigger picture of the current location. Therefore, we devise a candidate-enhancing (CE) mechanism to improve navigation. A flowchart of CE is shown in Fig. 2. Faced with K candidate views at Step t , the agent retrieves their representations \mathbf{m}_k , $k \in \{1, \dots, K\}$ from episodic memory $\mathcal{G}^{(t-1)}$,

$$\mathbf{m}_k = \begin{cases} \mathbf{m}_{V_j} & \text{if the } k\text{-th view is } V_j \in \mathcal{V}^{(t-1)} \\ \mathbf{0} & \text{otherwise.} \end{cases} \quad (2)$$

Then the Fusion block integrates the EScene representations with the plain features $\{\mathbf{f}_1, \dots, \mathbf{f}_K\}$ to produce enhanced can-

didate viewpoints,

$$\mathbf{o}_k = \text{MLP}([\mathbf{m}_k; \mathbf{f}_k]), \quad (3)$$

where $[\cdot; \cdot]$ denotes concatenation along feature dimension. The MLP function is a two-layer non-linear projection from \mathbb{R}^{2d} to \mathbb{R}^d . Following Chen et al. (2021b); Zhao et al. (2022), type embedding that distinguishes visual and linguistic signals, navigable embedding that indicates the navigability of each candidate view, and orientation encoding are added to \mathbf{o}_k . A zero vector $\mathbf{o}_s \in \mathbb{R}^d$ is appended as the feature for STOP action.

Finally, together with encoded history features, the enhanced candidate representations $\{\mathbf{o}_1, \dots, \mathbf{o}_K, \mathbf{o}_s\}$ are input to the cross-modal encoder to merge linguistic information from encoded text features. The agent predicts the distribution of action a_t via a two-layer non-linear Predictor block,

$$P(a_t = k \in \{1, \dots, K, s\}) = \frac{e^{\text{MLP}(\mathbf{o}'_k \odot \mathbf{x}'_{cls})}}{\sum_{j \in \{1, \dots, K, s\}} e^{\text{MLP}(\mathbf{o}'_j \odot \mathbf{x}'_{cls})}}, \quad (4)$$

where \odot is element-wise multiplication of two vectors \mathbf{o}'_k and $\mathbf{x}'_{cls} \in \mathbb{R}^d$, and the two-layer non-linear MLP block maps the result to a scalar $\in \mathbb{R}$. Following Tan et al. (2019); Chen et al. (2021b), we train the framework end-to-end by a mixture of Imitation Learning and Reinforcement Learning (A2C Mnih et al. 2016) loss,

$$\mathcal{L} = -\alpha \sum_{t=1}^{T^*} \log P(a_t = a_t^*) - \sum_{t=1}^T \log P(\tilde{a}_t)(r_t - v_t), \quad (5)$$

where T^* and T are the length of the annotated route and predicted path, respectively. \tilde{a}_t is sampled action. r_t is the discount reward, and v_t is the state value given by a two-layer (MLP) critic network.

4 Experiments

4.1 Experimental Setup

4.1.1 Datasets and Metrics

We conduct experiments on the following three VLN tasks for evaluation.

- (1) Short-horizon with fine-grained instructions. R2R¹ (Anderson et al., 2018b) constructs on Matterport3D (Chang et al., 2017) and has 7,189 direct-to-goal trajectories with an average of 10m. Each path is associated with three

instructions of 29 words on average. The train, val seen, val unseen, and test unseen splits include 61, 56, 11, and 18 houses, respectively.

- (2) Long-horizon with fine-grained instructions. R4R² (Jain et al., 2019) is generated by joining existing trajectories in R2R with others that start close by where they end. Compared to R2R, it has longer paths and instructions and reduced shortest-path bias. The train, val seen, and val unseen have 233,613, 1,035, and 45,162 samples, respectively.
- (3) Vision-dialog navigation. CVDN³ (Jain et al., 2019) requires an agent to navigate given a target object and a dialog history. It has 7k trajectories and 2,050 navigation dialogs, where the paths and language contexts are also longer than those in R2R. The train, val seen, val unseen, and test splits contain 4,742, 382, 907, and 1,384 instances, respectively.

Following standard criteria (Chen et al., 2021b; Anderson et al., 2018b, a), we evaluate the R2R dataset with Trajectory Length (TL), Navigation Error (NE), Success Rate (SR), and Success weighted by Path Length (SPL). TL is the average length of an agent's navigation route in meters, NE is the mean shortest path distance between the agent's stop location and the target, and SR measures the ratio of navigation that stops less than three meters from the goal. SPL normalizes SR by the ratio between the path length of ground truth and the navigated, which balances accuracy and efficiency and becomes the key metric for the R2R dataset. We adopt three additional metrics, Coverage weighted by Length Score (CLS), normalized Dynamic Time Warping (nDTW), and Success weighted by nDTW (SDTW), to assess path fidelity on the R4R dataset. As for vision-dialog navigation on CVDN, the primary evaluation metric is Goal Progress (GP) in meters.

4.1.2 Implementation Details

We adopt the encoders from Chen et al. (2021b) in comparison by default, where the text, history, and cross-modal encoders have nine, two, and four transformer layers, respectively. Features of single views are extracted offline using finetuned ViT-B/16 released by Chen et al. (2021b). For a fair comparison, we set the feature dimension $d=768$, the ratio of imitation learning loss $\alpha=0.2$, and train the EScheme framework for 100K iterations on each dataset with a batch size of 8 and a learning rate of $1e-5$. All the experiments run on a single NVIDIA V100 GPU. We adopt max pooling and single-run by default in comparison with other methods, and provide the results of mean pooling and inferring twice in

¹ <https://github.com/peteanderson80/Matterport3DSimulator>

² <https://github.com/google-research/google-research/tree/master/r4r>

³ <https://github.com/mmurray/cvdm/>

ablation studies and supplementary material, with qualitative examples and failure cases included.

For Reinforcement Learning, the action space is restricted to navigable locations (loosely equal to viewpoints) from each node, which is implemented by first predicting the log probability distribution over all the K viewpoints plus a STOP token and then setting the non-navigable ones as $-\infty$. The policy is given by $\pi(a_t | \{o'_i\}_1^s, \{h'_i\}_1^{t-1}, \{x'_i\}_{cls}^L)$, and sampling is conducted according to the restricted log probability. For Imitation Learning, the shortest path planner provides expert demonstrations, which are directly available from the simulator.

4.1.3 Fair Comparison

We consider deploying an agent in new environments to execute *a series of* language instructions. Admittedly, this definition is slightly different from existing methods, whereas it 1) preserves the original setup of *unseen* environment, i.e., the agent never sees the environment before deployment, and 2) is more practical in real scenarios, e.g., housework robots. Meanwhile, the proposed episodic memory leads to initialization change: the agent conducts the first episode with empty memory and the following episodes with its own estimates. The comparisons we made in the paper aim to verify the superiority of the proposed episodic memory instead of just showing an instantiation of EScheme surpassing its counterparts. This way, we inevitably compare it with existing path memory since this is a novel memory mechanism. Our *inter-episode* memory requires no extra time or computation while maintaining partial episodic memory via initialization, which is worth further exploration.

The directly available location ID, which we use to retrieve enhanced features for the current node, is universally adopted by 1) implicit path-level memory methods (e.g., HAMT) to retrieve accessible candidates, and 2) explicit path-level memory methods (e.g., EGP, SSM) to extend action space. Our comparisons introduce the essential signal, i.e., scan id as environment index, to be fair in showing the superiority of episodic memory over path-level memory, where *unseen* scenes refer to those never appearing in training/validation. In discrete environments, the usage of location ID inevitably leads to “the agent knowing the current location is exactly something it sees before”. The proposed EScheme can easily extend to continuous settings by combining with waypoint prediction methods (e.g., CWP Hong et al., 2022) that surpass most semantic-map-based approaches. Moreover, the proposed episodic memory mechanism can transfer to continuous scenes by maintaining a global map via the widely studied visual SLAM.

4.2 Comparison to State-of-the-Art

4.2.1 Results on R2R Dataset

Table 1 compares the proposed EScheme with existing methods on the R2R dataset. We can see that the pretraining-finetuning paradigm (e.g., RecBERT (Hong et al., 2021), HAMT (Chen et al., 2021b), ADAPT (Lin et al., 2022a), CSAP (Wu et al., 2022), TDSTP (Zhao et al., 2022)) largely improves the performance of VLN in unseen environments. EScheme achieves the highest SPL on the unseen splits. It surpasses the baseline model HAMT (Chen et al., 2021b) by about 5% SPL on the validation and test unseen environments and even outperforms TDSTP (Zhao et al., 2022) that involves auxiliary training tasks. Besides, EScheme brings a relative decrease of 6.4% and 4.1% in NE on validation and test unseen split, respectively. The results demonstrate the efficacy of episodic scene memory in generalization to unseen scenarios with short instructions.

We also compare with the most recent works, including VLN-SIG (Li & Bansal, 2023), KERM (Li et al., 2023), and GridMM (Wang et al., 2023). The proxy pre-training task involved in VLN-SIG shows no advantage in unseen environments. KERM surpasses all the methods on the validation-seen split but drops much more heavily on unseen splits. GridMM achieves the highest SR and slightly lower SPL than ours in unseen scenarios yet takes a much longer trajectory length.

4.2.2 Results on R4R Dataset

We evaluate the proposed EScheme on the R4R dataset to examine if the generalization promotion is maintained in long-horizon navigation tasks. The results are listed in Table 2. Our EScheme outperforms existing state-of-the-art by a large margin, i.e., a relative improvement of 6.4% in SPL, 7.0% in CLS, 7.3% in nDTW, and 9.1% in SDTW. It indicates that EScheme improves not only navigation success but also path fidelity. Although good at carrying out short instructions, TDSTP (Zhao et al., 2022) suffers a heavy drop in long-horizon navigation regarding path fidelity compared with its baseline model HAMT (Chen et al., 2021b). It reveals that goal-related auxiliary tasks such as target prediction benefit reaching the target location but undermine the ability to follow instructions. Equipped with EScheme, an agent has a promoted ability to travel the expected route in long-horizon navigation. Besides, a consistent advantage of pretraining-based methods can be observed on this dataset.

Table 1 Comparison with state-of-the-art methods on R2R dataset

Method	Validation seen				Validation unseen				Test unseen			
	TL	NE↓	SR↑	SPL↑	TL	NE↓	SR↑	SPL↑	TL	NE↓	SR↑	SPL↑
Seq2Seq (Anderson et al., 2018b)	11.33	6.01	39	–	8.39	7.81	22	–	8.13	7.85	20	18
SF (Fried et al., 2018)	–	3.36	66	–	–	6.62	35	–	14.82	6.62	35	28
AcPercep (Wang et al., 2020a)	19.7	3.20	70	52	20.6	4.36	58	40	21.6	4.33	60	41
PRESS (Li et al., 2019)	10.57	4.39	58	55	10.36	5.28	49	45	10.77	5.49	49	45
SSM (Wang et al., 2021)	14.7	3.10	71	62	20.7	4.32	62	45	20.4	4.57	61	46
EnvDrop (Tan et al., 2019)	11.00	3.99	62	59	10.70	5.22	52	48	11.66	5.23	51	47
OAAM (Qi et al., 2020a)	10.20	–	65	62	9.95	–	54	50	10.40	–	53	50
AuxRN (Zhu et al., 2020a)	–	3.33	70	67	–	5.28	55	50	–	5.15	55	51
PREVALENT (Hao et al., 2020)	10.32	3.67	69	65	10.19	4.71	58	53	10.51	5.30	54	51
RelGraph (Hong et al., 2020)	10.13	3.47	67	65	9.99	4.73	57	53	10.29	4.75	55	52
NvEM (An et al., 2021)	11.09	3.44	69	65	11.83	4.27	60	55	12.98	4.37	58	54
NvEM+SEvol (Chen et al., 2022a)	11.97	3.56	67	63	12.26	3.99	62	57	13.40	4.13	62	57
CSAP (Wu et al., 2022)	11.29	2.80	74	70	12.59	3.72	65	59	13.30	4.06	62	57
RecBERT (Hong et al., 2021)	11.13	2.90	72	68	12.01	3.93	63	57	12.35	4.09	63	57
ADAPT (Lin et al., 2022a)	11.39	2.70	74	69	12.33	3.66	66	59	13.16	4.11	63	57
HOP (Qiao et al., 2022)	11.26	2.72	75	70	12.27	3.80	64	57	12.68	3.83	64	59
TDSTP (Zhao et al., 2022)	–	2.34	77	73	–	3.22	70	63	–	3.73	67	61
HAMT (Chen et al., 2021b)	11.15	2.51	76	72	11.46	3.62	66	61	12.27	3.93	65	60
VLN-SIG (Li & Bansal, 2023)	–	–	–	–	–	–	68	62	–	–	65	60
KERM (Li et al., 2023)	12.16	2.19	80	74	13.54	3.22	72	61	14.60	3.61	70	59
GridMM (Wang et al., 2023)	–	–	–	–	13.27	2.83	75	64	14.43	3.35	73	62
ESceme (Ours)	10.65	2.57	76	73	10.80	3.39	68	64	11.89	3.77	66	63

Best results are marked as bold

ESceme (Ours) is built with HAMT (Chen et al., 2021b) architecture by default

Table 2 Comparison on the val unseen split of R4R dataset

Method	NE↓	SR↑	SPL↑	CLS↑	nDTW↑	SDTW↑
SF (Fried et al., 2018)	8.47	24	12	30	–	–
EnvDrop (Tan et al., 2019)	–	29	–	34	–	9
PTA (Cornia & Cucchiara, 2019)	8.25	24	10	37	32	10
RCM (Wang et al., 2019)	8.08	26	21	35	30	13
SSM (Wang et al., 2021)	8.27	32	–	53	39	19
NvEM (An et al., 2021)	6.85	38	28	41	36	20
NvEM+SEvol (Chen et al., 2022a)	6.90	39	29	41	36	20
RelGraph (Hong et al., 2020)	7.43	36	26	41	47	34
EGP (Deng et al., 2020)	8.00	30.2	–	44.4	37.4	17.5
TDSTP (Zhao et al., 2022)	6.32	43.3	40.6	46.4	42.1	25.5
RecBERT (Hong et al., 2021)	6.67	43.6	–	51.4	45.1	29.9
CSAP (Wu et al., 2022)	6.21	43.0	–	58.6	51.9	31.5
HAMT (Chen et al., 2021b)	6.09	44.6	40.6	57.7	50.3	31.8
ESceme (Ours)	5.84	45.6	43.2	62.7	55.7	34.7

Best results are marked as bold

4.2.3 Results on CVND Dataset

Table 3 compares ESceme with state-of-the-art methods on the vision-and-dialog navigation task. CVDN provides longer instructions and trajectories than R2R and more

complicated instructions than R4R. The proposed ESceme achieves the best goal process in both seen and unseen scenarios and wins first place on the leaderboard. HAMT (Chen et al., 2021b) shows an obvious advantage over other pretraining-based methods such as PREVALENT (Hao et

Table 3 Results of goal process (GP) in meters on CVDN dataset

Method	Val seen	Val unseen	Test unseen
Seq2Seq (Anderson et al., 2018b)	5.92	2.10	2.35
PREVALENT (Hao et al., 2020)	–	3.15	2.44
CMN (Zhu et al., 2020b)	7.05	2.97	2.95
VISITRON (Shrivastava et al., 2022)	5.11	3.25	3.11
HOP (Qiao et al., 2022)	–	4.41	3.24
SCoA (Zhu et al., 2021)	7.11	2.85	3.31
MT-RCM+EnvAg (Wang et al., 2020b)	5.07	4.65	3.91
HAMT (Chen et al., 2021b)	6.91	5.13	5.58
ESceme (Ours)	8.34	5.42	5.99

Best results are marked as bold

Table 4 Ablation studies of ESceme construction on R2R dataset to compare the effect of different pooling functions

	pooling	Validation seen				Validation unseen			
		TL	NE↓	SR↑	SPL↑	TL	NE↓	SR↑	SPL↑
HAMT (Chen et al., 2021b)	–	11.02±0.10	2.52±0.10	75.0±0.9	71.7±0.7	11.72±0.34	3.63±0.05	65.7±0.7	60.9±0.7
ESceme	mean	11.13±0.16	2.59±0.09	75.1±0.7	71.9±0.6	11.49±0.27	3.50±0.03	67.1±0.5	62.3±0.5
ESceme	max	10.81±0.12	2.60±0.12	75.6±0.4	72.6±0.4	11.18±0.23	3.44±0.03	67.4±0.5	63.2±0.5

al., 2020), and even surpasses those counterparts specially designed for vision-and-dialog navigation, e.g., CMN (Zhu et al., 2020b), VISITRON (Shrivastava et al., 2022), and SCoA (Zhu et al., 2021). Our ESceme brings a relative improvement of 20.7%, 5.7%, and 7.3% over the baseline HAMT (Chen et al., 2021b) in val seen, val unseen, and test unseen environments, respectively.

4.3 Ablation Studies & Analysis

4.3.1 Different ESceme Constructions

We evaluate the effect of different pooling functions in Table 4. Candidate Enhancing with mean pooling brings a relative improvement of 2.3% in SPL for unseen navigation and behaves similarly in seen environments. Integrated with max pooling, Candidate Enhancing further boosts the performance in unseen environments, which produces a 3.8% relative increase compared to the HAMT (Chen et al., 2021b) baseline. The results demonstrate the efficacy of the proposed Candidate Enhancing, which improves observation representations via direct injection and fusion, and max pooling, which preserves more distinguishable features of each view. Appendix A discusses a different implementation of the proposed episodic scene memory by Graph Encoding.

4.3.2 Different Navigation Architectures & Inferring Strategies

The proposed ESceme is devised to be model-agnostic and should be compatible with any navigation network that has an

observation input. To validate this property, we build ESceme upon TDSTP (Zhao et al., 2022) that achieves the highest SR on the R2R dataset and list the results in Table 5. ESceme improves navigation in both seen and unseen environments by 4.9% and 1.4% in SPL, respectively.

As introduced in Sect. 3.3, the agent starts with an empty episodic scene memory during inference, and the memory keeps updating. If we let the agent renew its memory thoroughly by going through all the episodes and then evaluate its navigation performance, it will have a much more complete episodic memory. We present the results of ESceme* in Table 5. We can see that the nearly completed memory further boosts the performance in unseen environments by 1.3% and 2.1% regarding SPL for ESceme upon HAMT (Chen et al., 2021b) and TDSTP (Zhao et al., 2022), respectively. More results of ESceme* are in supplementary material, with slighter improvements observed for longer-horizon navigation. The results demonstrate that an agent learns to assist navigation with partial and persistently updated episodic memory.

The observation that the performance of ESceme* is only slightly better than that of ESceme has two sides. On the one hand, it indicates that the agent has learned to use the dynamically accumulative episodic memory instead of working until collecting the complete memory. On the other hand, the slight gain of ESceme* indicates possible bottlenecks in the encoder/cross-encoder architecture, the frozen vision encoder, and the scale of datasets.

More effects of the proposed episodic scene memory are present in Appendix B and F. Comparison with pre-exploration methods shows that ESceme* is more robust to

Table 5 Ablation studies of navigation architectures and inferring strategies on R2R dataset

Method	Validation seen			Validation unseen			Params (MB)	GPU (GB)	Time (ms)
	TL	NE↓	SPL↑	TL	NE↓	SPL↑			
HAMT (Chen et al., 2021b)	11.02±0.10	2.52±0.10	71.7±0.7	11.72±0.34	3.63±0.05	60.9±0.7	651.5	8.5	29.4
+ ESceme	10.81±0.12	2.60±0.12	72.6±0.4	11.18±0.23	3.44±0.03	63.2±0.5	+6.8	+0.1	+1.4
+ ESceme*	10.77±0.13	2.58±0.12	72.8±0.4	10.89±0.14	3.35±0.05	64.0±0.4	+6.8	+0.1	+32.2
TDSTP (Zhao et al., 2022)	13.09±0.37	2.42±0.08	70.9±0.7	14.28±0.44	3.22±0.09	62.1±0.6	657.2	10.5	46.9
+ ESceme	11.80±0.26	2.34±0.10	74.4±0.6	13.86±0.21	3.31±0.07	63.0±0.8	+6.8	+0.1	+1.8
+ ESceme*	11.83±0.23	2.33±0.08	74.8±0.7	13.38±0.28	3.21±0.05	64.3±0.7	+6.8	+0.1	+50.5

ESceme* denotes navigating with a nearly completed scene memory by first going through all the episodes. For a scene in R2R covering 92 visited nodes on average, the maximum episodic memory cost in CPU is about 1.5MB

unseen scenarios. Ablation on graph re-initialization verifies that episodic scene memory contributes to decision-making in both seen and unseen environments. The observation in the IVLN benchmark is consistent with our discussion in Sect. 2 and our experimental results in Sect. 4.2, and validates the superiority of the proposed ESceme.

4.3.3 Computational Efficiency

We present model size, GPU usage, and time cost during inference on the R2R dataset in Table 5. Either upon HAMT (Chen et al., 2021b) or TDSTP (Zhao et al., 2022), the proposed ESceme brings about 1.0% extra parameters and memory occupation in GPU. In a single-run setting, ESceme slightly increases the computational time by 4.8% when built on top of HAMT. Compared with HAMT, the TDSTP baseline costs more time by 59.5% and GPU by 23.5%. Accordingly, our ESceme only raises the time cost by 3.8% and almost no extra GPU consumption. With better-completed memory, ESceme* further boosts navigation performance in new environments at the expense of double the time. We can see that ESceme achieves a good trade-off between efficiency and efficacy in a single run. The proposed episodic memory mechanism consumes marginal ($\leq 0.1\%$) computation and parameters. For D -dim features, K nodes/scene, and N scenes, increased cost of space and parameters are about $3.81e^{-6} \times DK N$ and $1.14e^{-5} \times D^2$, respectively.

4.3.4 Order of Executing Instructions

Since ESceme learns with dynamically updated episodic memory while conducting instructions, the order of execution has little impact on overall performance. Table 6 lists navigating performance with shuffled episodes on the val unseen split in all the datasets, which indicates the stability of ESceme.

4.3.5 Success Variation During Inference

Figure 4 compare SPL and CLS curves of different methods to visualize the variation of navigation quality in infer-

Table 6 $\bar{x} \pm \sigma$ scores of shuffled episodes with five random seeds on the val unseen split of the datasets

(R2R) NE↓	SPL↑	(R4R) SPL↑	CLS↑	(CVDN) GP↑
3.39±0.03	63.7±0.3	43.2±0.07	62.7±0.1	5.57±0.11

ring progress. On the short-horizon navigation dataset R2R, HAMT (Chen et al., 2021b) oscillates around 62 and drops in the last 1/5 progress. The decrease could result from more tough samples at the end. TDSTP (Zhao et al., 2022) presents a more stable oscillation around 62, owing to a global action space and an auxiliary goal-related task. Starting from a moderate navigation ability, an agent with ESceme benefits greatly from memory updates and maintains a high success rate with completed memory.

On the long-horizon VLN dataset R4R, TDSTP (Zhao et al., 2022) shares a similar oscillation around 41 with HAMT (Chen et al., 2021b) in SPL. TDSTP preserves a relatively more stable success rate at the cost of much lower CLS, which reveals that goal-related auxiliary task undermines the ability of instruction following. Our ESceme shows a sharp increase within the first 4/5 navigation and has remained stable since then. We attribute the excellent promotion on R4R to two reasons, 1) long-horizon navigation involves more action steps, so a slight increase in navigation ability results in a big difference in final performance; 2) the sample density of a scene from R4R is much higher than that from the R2R dataset.

4.4 Qualitative Analysis

To intuitively demonstrate the benefit of the proposed episodic scene memory, we provide a visualization example in Fig. 5. It shows the panoramic views and top-down overviews of navigation. The last step of HAMT and TDSTP navigates to a visible corner of the bedroom. Instead, ESceme understands the instruction better. It takes a step to walk down to the end of the hall and then turns left to the bedroom.




Fig. 4 Navigation quality w.r.t. inferring progress. The x-axis indicates the ratio of samples tested, and the y-axis is the smoothed average of SPL or CLS. We use the default order for all the methods. Navigation with EScheme improves over time




Fig. 5 Panoramic views and top-down overviews of navigation. Mistakes during navigation are marked with red boxes for panorama and red arrows for top-down trajectories. The star indicates the target location.

Our EScheme strictly follows the instruction “walk down to the end of hall” and waits at the door of the bedroom (Color figure online)



(a) ground-truth trajectory

(b) trajectory predicted by ESceme

Fig. 6 Failure case in R2R val unseen split. The instruction is “Leave sitting room and head towards the kitchen, turn right at living room and enter. Walk through living room to dining room and enter. Turn left and

head to front door. Exit the house and stop on porch.” After correctly predicting the first three actions, ESceme failed to enter the dining room and got lost

A failure case of ESceme is shown in Fig. 6, where the instruction is “Leave sitting room and head towards the kitchen, turn right at living room and enter. Walk through living room to dining room and enter. Turn left and...” After correctly predicting the first three actions, ESceme failed to enter the dining room and got lost. It indicates that the representations for the viewpoints are not distinguishable enough to capture some fine-grained difference between the dining room and the living room.

5 Conclusion

In this paper, we devise the first VLN mechanism with episodic scene memory (ESceme) and propose a simple yet effective implementation via candidate enhancing. We show that an agent with ESceme improves navigation ability in short-horizon, long-horizon, and vision-and-dialog navigation. Our method outperforms the existing state-of-the-art and wins first place in the CVDN leaderboard, bringing a marginal increase in memory, parameters, and inference time. We hope this work can inspire further explorations on episodic memory in VLN and related fields, e.g., build-

ing the memory in continuous environments and with more advanced techniques such as neural SLAM.

5.1 Limitations

Although we show the effectiveness of the proposed episodic scene memory, there are still several limitations. First, the agent requires knowledge of environmental identity to build episodic memory for each scene. It is inevitable but supported by practical demands where an agent conducts multiple instructions in one scenario. Second, the “location ID” information is directly available from the simulator and the dataset, which is accurate and free of noise. For the case where location ID is unknown in advance, the episodic scene memory can be built by adding a discrete mapping process analogous to SLAM. No specific location ID is required, and the rough global position of each node can be dynamically estimated using the angle of each navigable viewpoint. Third, the architecture of a navigation agent and the training data limit the efficacy of a complete scene memory. We hope the proposed episodic scene memory can be explored in more advanced and diverse architectures.

Table 7 Ablation studies of EScheme construction on R2R dataset

	GE	CE	pooling	Validation seen				Validation unseen			
				TL	NE↓	SR↑	SPL↑	TL	NE↓	SR↑	SPL↑
HAMT (Chen et al., 2021b)	–	–	–	11.02±0.10	2.52±0.10	75.0±0.9	71.7±0.7	11.72±0.34	3.63±0.05	65.7±0.7	60.9±0.7
EScheme	✓	✗	mean	11.20±0.18	2.56±0.11	75.7±0.9	72.3±0.6	11.64±0.05	3.60±0.06	65.9±0.5	60.9±0.6
EScheme	✗	✓	mean	11.13±0.16	2.59±0.09	75.1±0.7	71.9±0.6	11.49±0.27	3.50±0.03	67.1±0.5	62.3±0.5
EScheme	✗	✓	max	10.81±0.12	2.60±0.12	75.6±0.4	72.6±0.4	11.18±0.23	3.44±0.03	67.4±0.5	63.2±0.5

We compare the effect of graph encoding (GE) and candidate enhancing (CE), and different pooling functions

Appendix A: EScheme Navigation by Graph Encoding

Intuitively, the memory can be injected into the cross-modal encoder via a separate branch. We denote the solution as Graph Encoding (GE) and list experimental results. Figure 7 demonstrates EScheme-assisted navigation by adding a graph encoding (GE) branch to the cross-modal encoder. At the current location where the agent stands, a local window is masked to avoid repetition with the path history from time 1 to $t-1$. Thus, the searched episodic memory graph includes six nodes and three edges, i.e., $\mathcal{G}^{(t-1)} = \{\mathcal{V}^{(t-1)}, \mathcal{E}^{(t-1)}\}$. We adopt 3-WL GNNs (Maron et al., 2019; Dwivedi et al., 2020) that can distinguish two non-isomorphic graphs to encode the memory graph, where the input $G \in \mathbb{R}^{n \times n \times (1+d)}$ is given by

$$G_{ijk} = \begin{cases} e_{ij} & \text{if } k = 1 \\ m_i & \text{if } j = i \text{ for } k > 1 \\ 0 & \text{otherwise,} \end{cases} \quad (6)$$

where n is the number of nodes in $\mathcal{V}^{(t-1)}$. $m_i \in \mathbb{R}^d$ is the representation of the node V_i , with detailed calculations presented in Section 3.2. $e_{ij} = 1$ if V_i and V_j are connected, else $e_{ij} = 0$. The graph is encoded by

$$G' = [(W_1 G) \odot (W_2 G); (W_3 G)], \quad (7)$$

where $W_{1\sim 3} \in \mathbb{R}^{(1+d) \times (d/2)}$ are two-layer MLPs. \odot denotes element-wise multiplication and $[\cdot; \cdot]$ is the concatenation along feature dimension. The final encoded feature to the cross-modal encoder is $\sum_{i=1}^n \sum_{j=1}^n G'_{ij} \in \mathbb{R}^d$.

We evaluate the superiority of Candidate Enhancing over Graph Encoding and the effect of different pooling functions in Table 7. First, Graph Encoding with mean pooling slightly increases navigation success in seen environments with almost no promotion in unseen scenarios. We infer that Graph Encoding adjusts the representation of observations in cross-modal encoding and does not align well with the remaining branches to provide complementary information, resulting in a limited effect.

Appendix B: Effects of the Episodic Scene Memory

We thoroughly compare navigating with progressively completed and nearly complete episodic memory on three datasets in Tables 8 and 9. EScheme conducts instructions in a single-run setting, where the agent dynamically updates memory in inference. EScheme* first goes through all the episodes to build a nearly complete memory at the beginning of the evaluation. EScheme* improves navigating in new environments by 1.6% (SPL) on test unseen split of

Table 8 Results of different inferring strategies on R2R dataset

Method	Validation seen				Validation unseen				Test unseen			
	TL	NE↓	SR↑	SPL↑	TL	NE↓	SR↑	SPL↑	TL	NE↓	SR↑	SPL↑
ESceme	10.65	2.57	76	73	10.80	3.39	68	64	11.89	3.77	66	63
ESceme*	10.62	2.57	76	73	10.65	3.36	68	64	11.77	3.69	68	64

Table 9 Results of different inferring strategies on R4R and CVDN datasets

Method	R4R Val unseen						CVDN		
	NE↓	SR↑	SPL↑	CLS↑	nDTW↑	SDTW↑	Val seen	Val unseen	Test unseen
ESceme	5.84	45.6	43.2	62.7	55.7	34.7	8.34	5.42	5.99
ESceme*	5.83	45.8	43.4	62.9	55.8	34.8	8.39	5.72	6.17

Table 10 Comparison between ESceme* and pre-exploration methods on R2R dataset

Method	Validation seen				Validation unseen			
	TL	NE↓	SR↑	SPL↑	TL	NE↓	SR↑	SPL↑
Follower (Fried et al., 2018)	10.40	3.68	65	62	9.57	5.20	52	49
Speaker (Fried et al., 2018)	11.19	3.80	61	56	10.71	4.25	55	49
VLN-BERT (Majumdar et al., 2020)	10.28	3.73	70	66	9.60	4.10	59	55
Airbert (Guhur et al., 2021)	11.09	2.68	75	70	11.78	4.01	62	56
ESceme*	10.62	2.57	76	73	10.65	3.36	68	64

Best results are marked as bold

Table 11 Ablation of memory re-initialization on R2R dataset

Method	Validation seen				Validation unseen			
	TL	NE↓	SR↑	SPL↑	TL	NE↓	SR↑	SPL↑
ESceme	10.65	2.57	76	73	10.80	3.39	68	64
re-initialize	13.88	3.53	64	60	14.88	4.33	55	51

Best results are marked as bold

Table 12 Comparison on RxR dataset

Method	Validation seen				Validation unseen			
	SR↑	SPL↑	nDTW↑	SDTW↑	SR↑	SPL↑	nDTW↑	SDTW↑
HAMT	59.4	58.9	65.3	50.9	56.5	56.0	63.1	48.3
ESceme	64.8	61.4	68.8	56.2	62.0	58.4	67.6	53.5

Best results are marked as bold

ESceme (Ours) is built with HAMT architecture by default

the R2R dataset. As for vision-dialog navigation CVDN, the improvement in val unseen and test unseen is 5.5% and 3.0%, respectively. On the long-horizon navigation dataset R4R, the relative increase is about 0.5%.

Overall, ESceme* further promotes generalization to novel scenarios, indicating that ESceme benefits from the nearly complete scene memory. On the other hand, the small gap between ESceme and ESceme* shows that the agent has learned to utilize progressively completed memory in navigation.

Besides, Table 10 lists the comparison with pre-exploration methods. The pre-exploration methods achieve very competitive results on val seen split while suffering from a heavier drop in unseen environments. In Table 11, we test ESceme with the memory graph re-initialized at every episode. The

results on the R2R dataset verify that the ESceme agent indeed benefits from episodic memory for decision-making in both seen and unseen environments.

Appendix C: Experiments on RxR Dataset

Results on RxR⁴ (Ku et al., 2020b) in Table 12 indicate that longer instructions and trajectories add sufficient knowledge to the proposed episodic memory to overcome coreference challenge and promote navigation.

⁴ <https://github.com/google-research-datasets/RxR>

Algorithm 1: EScene construction and candidate enhancing

```

1 def update( $G_Y$ , ob):
2   #  $G_Y$ : episodic memory of Scene Y
3   # ob: observation structure that includes information of the current viewpoint and  $K$  views of a panorama
4   if ob.viewpoint not in  $G_Y$ :
5      $G_Y$ .update(ob.viewpoint), feat=pooling(ob.navigable_feats) # Update nodes by pooling ViT features
6     for node in ob.navigable_views:
7       if node in  $G_Y$ :
8          $G_Y$ .add_edge(ob.viewpoint, node) # Update edges
9   return  $G_Y$ 
10
11 def candidate_enhancing( $G_Y$ , ob):
12    $m$  =  $G_Y$ .search(ob.navigable_views).feat # Search memory representations for currently navigable views
13    $o$  = MLP(torch.cat([ob.navigable_feats,  $m$ ], dim=1)) # Enhance candidate representations
14    $o$  =  $o$  + embeddings # Optionally add embeddings from orientation and navigable type
15   return  $o$ 

```

Appendix D: Pseudo-Code Implementation

We provide the pseudo-code of EScene construction and candidate enhancing in Algorithm 1. EScene requires easy implementation and can be integrated with any navigation networks that encode the observation.

Appendix E: Qualitative Examples and Failure Cases

We present the navigating process to provide a more intuitive comparison with HAMT (Chen et al., 2021b) and TDSTP (Zhao et al., 2022). Figures 8 and 9 are two navigation examples on R2R dataset, and Figs. 10 and 11 illustrate two examples on R4R dataset. All the examples are tested in unseen environments. For short-horizon navigation, our EScene outperforms its counterparts regarding stopping precision. For long-horizon navigation, our EScene shows an improved ability to follow instructions that require a forward and back trip and arrives at the target location. We attribute these advantages to the episodic memory of the scenes.

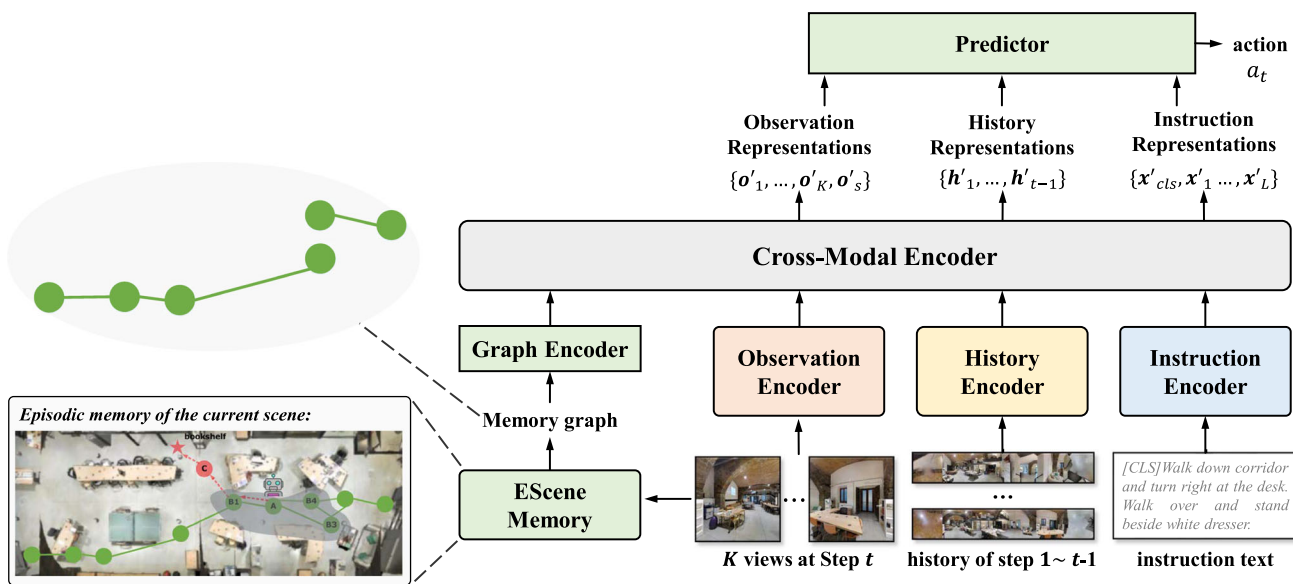


Fig. 7 An overview of EScene-assisted navigation by graph encoding. First, Episodic memory is built in the same way as that for candidate enhancing (c.f. Section 3.2). Then, the agent searches the episodic mem-

ory for the current viewpoint and obtains thememory graph by masking a local window. The encoded memory composes a separate branch to the cross-modal encoder




Fig. 8 Panoramic views and top-down overviews of navigation. Mistakes during navigation are marked with red boxes for panorama and red arrows for top-down trajectories. The star indicates the target loca-

tion. Our ESceme strictly follows the instruction “go up two steps” and waits on the third step (Color figure online)



Fig. 9 Panoramic views and top-down overviews of navigation. Mistakes during navigation are marked with red boxes for panorama and red arrows for top-down trajectories. The star indicates the target location. Our EScene stops at the right place (Color figure online)



Fig. 10 The top-down trajectory of navigation. Mistakes during navigation are marked with red. The star indicates the target location. Our EScene moves forward to the right place and then back and arrives at the second bathroom (Color figure online)

Figures 12 showcase one more situation where EScene failed to follow the instructions. The instruction is “Go down the stairs. Go into the room straight ahead on the slight left. Wait there.” EScene succeeded in going downstairs but failed to determine the *slight left* direction and entered the wrong room. The result indicates difficulties in understanding finer-grained instructions and distinguishing finer-grained visual observations in physical scenarios, as discussed in Sect. 5 Limitations.

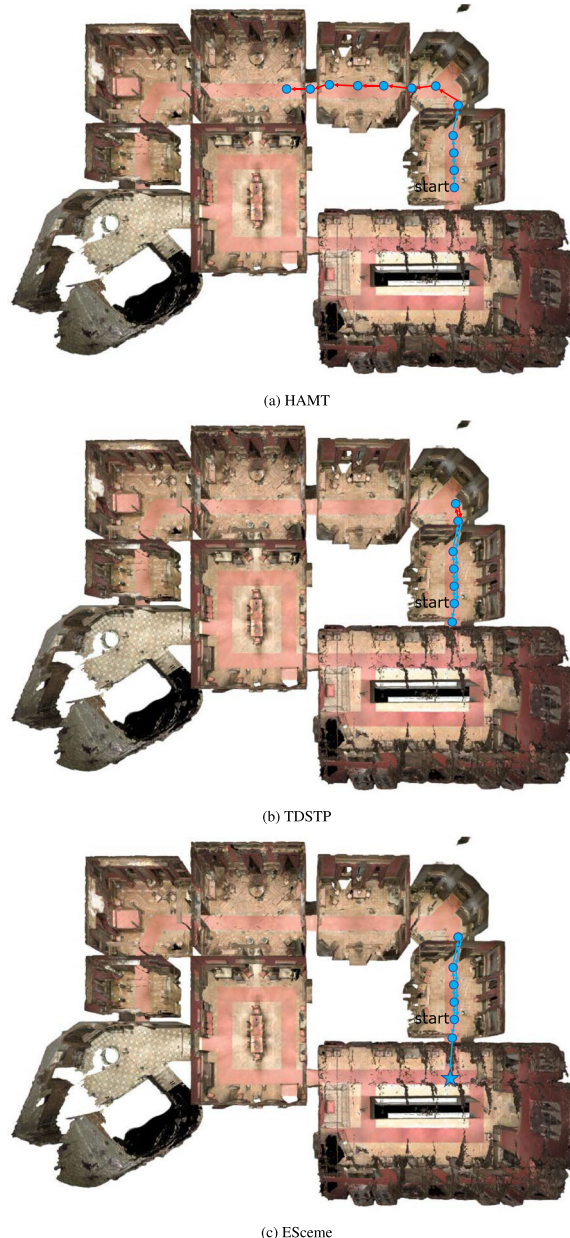


Fig. 11 The top-down trajectory of navigation. Mistakes during navigation are marked with red. The star indicates the target location. Our EScememoves forward to the right place and then back and stops inside the double doors (Color figure online)

Appendix F: Comparison with IVLN

Our setting is identical to IVLN, which reorganizes episodes into tours. Table 13 is the direct comparison using the IVLN benchmark. IVLN decreases the performance in both seen and unseen environments, yet our EScene promotes navigation in unseen scenarios while maintaining the performance



Fig. 12 Failure case in R2R val unseen split. The instruction is “Go down the stairs. Go into the room straight ahead on the slight left. Wait there.” EScene succeeded in going downstairs but failed to determine the *slight left* direction and entered the wrong room

Table 13 Comparison with IVLN (Krantz et al., 2023)

Method	Validation seen				Validation unseen			
	TL	NE↓	SR↑	SPL↑	TL	NE↓	SR↑	SPL↑
HAMT	10.1	4.2	63	61	9.4	4.7	56	54
IVLN	9.4	5.8	45	43	10.0	6.2	39	36
EScene	10.0	4.4	63	61	9.2	4.6	58	56

Best results are marked as bold

in seen ones. The conclusion is consistent with our observations in Sect. 4.

Acknowledgements This work was partially supported by the Start-up Funding of Shenzhen University and a CSIRO top-up scholarship.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- An, D., Qi, Y., Huang, Y., Wu, Q., Wang, L., & Tan, T. (2021). Neighbor-view enhanced model for vision and language navigation. In *ACMMM*, pp. 5101–5109.
- An, D., Wang, H., Wang, W., Wang, Z., Huang, Y., He, K., & Wang, L. (2023). Etpnav: Evolving topological planning for vision-language navigation in continuous environments. *arXiv preprint arXiv:2304.03047*.
- Anderson, P., Chang, A., Chaplot, D. S., Dosovitskiy, A., Gupta, S., Koltun, V., Kosecka, J., Malik, J., Mottaghi, R., & Savva, M., et al. (2018). On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*.
- Anderson, P., Wu, Q., Teney, D., Bruce, J., Johnson, M., Sünderhauf, N., Reid, I., Gould, S., & Van Den Hengel, A. (2018). Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *CVPR*, pp. 3674–3683.
- Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C. L., & Parikh, D. (2015). Vqa: Visual question answering. In *ICCV*, pp. 2425–2433.
- Chang, A., Dai, A., Funkhouser, T., Halber, M., Niessner, M., Savva, M., Song, S., Zeng, A., & Zhang, A. (2017). Matterport3d: Learning from rgb-d data in indoor environments. In *International Conference on 3D Vision*, pp. 667–676.
- Chaplot, D. S., Salakhutdinov, R., Gupta, A., & Gupta, S. (2020). Neural topological slam for visual navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12875–12884.
- Chen, J., Gao, C., Meng, E., Zhang, Q., & Liu, S. (2022). Reinforced structured state-evolution for vision-language navigation. In *CVPR*, pp. 15450–15459.
- Chen, K., Chen, J. K., Chuang, J., Vázquez, M., & Savarese, S. (2021). Topological planning with transformers for vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11276–11286.

- Chen, S., Guhur, P.-L., Schmid, C., & Laptev, I. (2021). History aware multimodal transformer for vision-and-language navigation. In *NeurIPS*, 34, 5834–5847.
- Chen, S., Guhur, P.-L., Tapaswi, M., Schmid, C., & Laptev, I. (2022). Think global, act local: Dual-scale graph transformer for vision-and-language navigation. In *CVPR*, pp. 16537–16547.
- Chen, X., Fang, H., Lin, T.-Y., Vedantam, R., Gupta, S., Dollár, P., & Zitnick, C. L. (2015). Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Cornia, F. L. B. M., & Cucchiara, M. C. R. (2019). Perceive, transform, and act: Multi-modal attention networks for vision-and-language navigation. *arXiv preprint arXiv:1911.12377*.
- Datta, S., Dharur, S., Cartillier, V., Desai, R., Khanna, M., Batra, D., & Parikh, D. (2022). Episodic memory question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19119–19128.
- Deng, Z., Narasimhan, K., & Russakovsky, O. (2020). Evolving graphical planner: Contextual global planning for vision-and-language navigation. In *NeurIPS*, 33, 20660–20672.
- Dwivedi, V. P., Joshi, C. K., Laurent, T., Bengio, Y., & Bresson, X. (2020). Benchmarking graph neural networks. *arXiv preprint arXiv:2003.00982*.
- Fried, D., Hu, R., Cirik, V., Rohrbach, A., Andreas, J., Morency, L.-P., Berg-Kirkpatrick, T., Saenko, K., Klein, D., & Darrell, T. (2018). Speaker-follower models for vision-and-language navigation. In *NeurIPS*, volume 31.
- Georgakis, G., Schmeckpeper, K., Wanchoo, K., Dan, S., Mitsakaki, E., Roth, D., & Daniilidis, K. (2022). Cross-modal map learning for vision and language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15460–15470.
- Guhur, P.-L., Tapaswi, M., Chen, S., Laptev, I., & Schmid, C. (2021). Airbert: In-domain pretraining for vision-and-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1634–1643.
- Gupta, S., Davidson, J., Levine, S., Sukthankar, R., & Malik, J. (2017). Cognitive mapping and planning for visual navigation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2616–2625.
- Hao, W., Li, C., Li, X., Carin, L., & Gao, J. (2020). Towards learning a generic agent for vision-and-language navigation via pre-training. In *CVPR*, pp. 13137–13146.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735–1780.
- Hong, Y., Rodriguez, C., Qi, Y., Wu, Q., & Gould, S. (2020). Language and visual entity relationship graph for agent navigation. In *NeurIPS*, 33, 7685–7696.
- Hong, Y., Wu, Q., Qi, Y., Rodriguez-Opazo, C., & Gould, S. (2021). Vln bert: A recurrent vision-and-language bert for navigation. In *CVPR*, pp. 1643–1653.
- Hong, Y., Wang, Z., Wu, Q., & Gould, S. (2022). 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*, pp. 15439–15449.
- Jain, V., Magalhaes, G., Ku, A., Vaswani, A., Ie, E., & Baldrige, J. (2019). Stay on the path: Instruction fidelity in vision-and-language navigation. In *ACL*, pp. 1862–1872.
- Ke, L., Li, X., Bisk, Y., Holtzman, A., Gan, Z., Liu, J., Gao, J., Choi, Y., & Srinivasa, S. (2019). Tactical rewind: Self-correction via backtracking in vision-and-language navigation. In *CVPR*, pp. 6741–6749.
- Kenton, J. D. M.-W. C., & Toutanova, L. K. (2016). Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pp. 4171–4186.
- Krantz, J., Banerjee, S., Zhu, W., Corso, J., Anderson, P., Lee, S., & Thomason, J. (2023). Iterative vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14921–14930.
- Ku, A., Anderson, P., Patel, R., Ie, E., & Baldrige, J. (2020). Room-across-room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In *EMNLP*.
- Ku, A., Anderson, P., Patel, R., Ie, E., & Baldrige, J. (2020). Room-Across-Room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In *Conference on Empirical Methods for Natural Language Processing (EMNLP)*.
- Li, J., & Bansal, M. (2023). Improving vision-and-language navigation by generating future-view image semantics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10803–10812.
- Li, J., Tan, H., & Bansal, M. (2022). Envedit: Environment editing for vision-and-language navigation. In *CVPR*, pp. 15407–15417.
- Li, M., Wang, Z., Tuytelaars, T., & Moens, M.-F. (2022). Layout-aware dreamer for embodied referring expression grounding. *arXiv preprint arXiv:2212.00171*.
- Li, X., Li, C., Xia, Q., Bisk, Y., Celikyilmaz, A., Gao, J., Smith, N., & Choi, Y. (2019). Robust navigation with language pretraining and stochastic sampling. In *EMNLP-IJCNLP*.
- Li, X., Wang, Z., Yang, J., Wang, Y., & Jiang, S. (2023). Kerm: Knowledge enhanced reasoning for vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2583–2592.
- Lin, B., Zhu, Y., Chen, Z., Liang, X., Liu, J., & Liang, X. (2022). Adapt: Vision-language navigation with modality-aligned action prompts. In *CVPR*, pp. 15396–15406.
- Lin, C., Jiang, Y., Cai, J., Qu, L., Haffari, G., & Yuan, Z. (2022). Multimodal transformer with variable-length memory for vision-and-language navigation. In *ECCV*.
- Ma, C.-Y., Lu, J., Wu, Z., AlRegib, G., Kira, Z., Socher, R., & Xiong, C. (2019). Self-monitoring navigation agent via auxiliary progress estimation. In *ICLR*.
- Ma, C.-Y., Wu, Z., AlRegib, G., Xiong, C., & Kira, C. (2019). The regretful agent: Heuristic-aided navigation through progress estimation. In *CVPR*, pp. 6732–6740.
- Majumdar, A., Shrivastava, A., Lee, S., Anderson, P., Parikh, D., & Batra, D. (2020). Improving vision-and-language navigation with image-text pairs from the web. In *ECCV*, pp. 259–274. Springer.
- Maron, H., Ben-Hamu, H., Serviansky, H., & Lipman, Y. (2019). Provably powerful graph networks. *NeurIPS*, 32.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., & Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. In *ICML*, pp. 1928–1937. PMLR.
- Qi, Y., Pan, Z., Zhang, S., Hengel, A. v. d., & Wu, Q. (2020). Object-and-action aware model for visual language navigation. In *ECCV*, pp. 303–317. Springer.
- Qi, Y., Wu, Q., Anderson, P., Wang, X., Wang, W. Y., Shen, C., & Hengel, A. v. d. (2020). Reverie: Remote embodied visual referring expression in real indoor environments. In *CVPR*, pp. 9982–9991.
- Qiao, Y., Qi, Y., Hong, Y., Yu, Z., Wang, P., & Wu, Q. (2022). Hop: History-and-order aware pre-training for vision-and-language navigation. In *CVPR*, pp. 15418–15427.
- Shah, D., Eysenbach, B., Rhinehart, N., & Levine, S. (2022). Rapid exploration for open-world navigation with latent goal models. In *Conference on Robot Learning*, pp. 674–684. PMLR.
- Shrivastava, A., Gopalakrishnan, K., Liu, Y., Piramuthu, R., Tür, G., Parikh, D., & Hakkani-Tur, D. (2022). Visitron: Visual semantics-aligned interactively trained object-navigator. In *Findings of the Association for Computational Linguistics: ACL, 2022*, 1984–1994.
- Smith, L., & Gasser, M. (2005). The development of embodied cognition: Six lessons from babies. *Artificial life*, 11(1–2), 13–29.

- Tan, H., Yu, L., & Bansal, M. (2019). Learning to navigate unseen environments: Back translation with environmental dropout. In *NAACL*.
- Thomason, J., Murray, M., Cakmak, M., & Zettlemoyer, L. (2020). Vision-and-dialog navigation. In *Conference on Robot Learning*, pp. 394–406. PMLR.
- Vasudevan, A. B., Dai, D., & Van Gool, L. (2021). Talk2nav: Long-range vision-and-language navigation with dual attention and spatial memory. *International Journal of Computer Vision*, 129, 246–266.
- Wang, H., Wang, W., Shu, T., Liang, W., & Shen, J. (2020). Active visual information gathering for vision-language navigation. In *ECCV*, pp. 307–322. Springer.
- Wang, H., Wang, W., Liang, W., Xiong, C., & Shen, J. (2021). Structured scene memory for vision-language navigation. In *CVPR*, pp. 8455–8464.
- Wang, X., Xiong, W., Wang, H., & Wang, W. Y. (2018). Look before you leap: Bridging model-free and model-based reinforcement learning for planned-ahead vision-and-language navigation. In *ECCV*, pp. 37–53.
- Wang, X., Huang, Q., Celikyilmaz, A., Gao, J., Shen, D., Wang, Y.-F., Wang, W. Y., & Zhang, L. (2019). Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation. In *CVPR*, pp. 6629–6638.
- Wang, X. E., Jain, V., Ie, E., Wang, W. Y., Kozareva, Z., & Ravi, S. (2020). Environment-agnostic multitask learning for natural language grounded navigation. In *ECCV*, pp. 413–430. Springer.
- Wang, Z., Li, X., Yang, J., Liu, Y., & Jiang, S. (2023). Gridmm: Grid memory map for vision-and-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15625–15636.
- Wu, S., Fu, X., Wu, F., & Zha, Z.-J. (2022). Cross-modal semantic alignment pre-training for vision-and-language navigation. In *ACMMM*, pp. 4233–4241.
- Xu, J., Mei, T., Yao, T., & Rui, Y. (2016). Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*, pp. 5288–5296.
- Zhao, Y., Chen, J., Gao, C., Wang, W., Yang, L., Ren, H., Xia, H., & Liu, S. (2022). Target-driven structured transformer planner for vision-language navigation. In *ACMMM*, pp. 4194–4203.
- Zhu, F., Zhu, Y., Chang, X., & Liang, X. (2020). Vision-language navigation with self-supervised auxiliary reasoning tasks. In *CVPR*, pp. 10012–10022.
- Zhu, Y., Zhu, F., Zhan, Z., Lin, B., Jiao, J., Chang, X., & Liang, X. (2020). Vision-dialog navigation by exploring cross-modal memory. In *CVPR*, pp. 10730–10739.
- Zhu, Y., Weng, Y., Zhu, F., Liang, X., Ye, Q., Lu, Y., & Jiao, J. (2021). Self-motivated communication agent for real-world vision-dialog navigation. In *ICCV*, pp. 1594–1603.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.