

Lookahead Exploration with Neural Radiance Representation for Continuous Vision-Language Navigation

Zihan Wang^{1,2}, Xiangyang Li^{1,2}, Jiahao Yang^{1,2}, Yeqi Liu^{1,2}, Junjie Hu^{3,4}, Ming Jiang⁵, Shuqiang Jiang^{1,2}

¹Institute of Computing Technology, Chinese Academy of Sciences, Beijing, 100190, China

²University of Chinese Academy of Sciences, Beijing, 100049, China

³Department of Computer Science, University of Wisconsin, Madison, WI, USA

⁴Department of Biostatistics and Medical Informatics, University of Wisconsin, Madison, WI, USA

⁵Department of Human-centered Computing, Indiana University, Indianapolis, IN, USA

zihan.wang@vipl.ict.ac.cn, lixiangyang@ict.ac.cn, {jiahao.yang, yeqi.liu}@vipl.ict.ac.cn,
 junjie.hu@wisc.edu, mj200@iu.edu, sqjiang@ict.ac.cn

Abstract

Vision-and-language navigation (VLN) enables the agent to navigate to a remote location following the natural language instruction in 3D environments. At each navigation step, the agent selects from possible candidate locations and then makes the move. For better navigation planning, the lookahead exploration strategy aims to effectively evaluate the agent's next action by accurately anticipating the future environment of candidate locations. To this end, some existing works predict RGB images for future environments, while this strategy suffers from image distortion and high computational cost. To address these issues, we propose the pre-trained hierarchical neural radiance representation model (HNR) to produce multi-level semantic features for future environments, which are more robust and efficient than pixel-wise RGB reconstruction. Furthermore, with the predicted future environmental representations, our lookahead VLN model is able to construct the navigable future path tree and select the optimal path via efficient parallel evaluation. Extensive experiments on the VLN-CE datasets confirm the effectiveness of our method. The code is available at <https://github.com/MrZihan/HNR-VLN>

1. Introduction

Vision-and-language navigation (VLN) tasks [5, 22, 23, 31] require an agent to understand natural language instructions and move to the destination. In the continuous environment setting (VLN-CE) [22], the navigation agent is free to traverse any unobstructed location with low-level actions (*i.e.*, turn left 15 degrees, turn right 15 degrees, or move forward 0.25 meters), similar to some visual navigation tasks [41, 42, 45–47]. As a result, the agent is more prone

Instruction: Walk past the bed and exit the bedroom through the door.
 Enter the toilet, and then stop before the window.

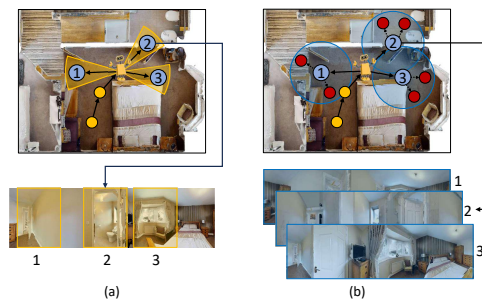


Figure 1. Illustration of different methods to represent the navigable candidate locations. (a) uses the single-view observation (yellow sector area). (b) uses the panorama of the candidate location (blue circular area) to anticipate the future environment.

to entering the visual blind area compared to the discrete environment setting with perfectly predefined navigable locations. This phenomenon raises a challenge to accurately represent future environments with visual occlusions, leading to incorrect action decisions.

As illustrated in Figure 1(a), previous approaches [8, 9, 25, 26] mainly rely on single-view visual observation of the current location to perceive candidate locations, thus leading to a *small restricted view* (yellow sector area) due to the visual occlusions. In contrast, as shown in Figure 1(b), a *comprehensive panorama* of the candidate location (blue circular area) is much more helpful to the agent for understanding the future environment and capturing critical visual cues in action decisions. *Lookahead exploration* [1, 11] aims to enable an agent to explore steps forward before making a navigation decision. Unlike other *lookahead* approaches, we adopt a strategy of exploring steps forward with environment anticipation, which helps current action decisions by predicting the future outcome of actions.

As generating future environment representations benefits much for action prediction, DREAMWALKER [39] proposes to imagine panoramic images of the navigable candidates. With simulated images through the visual generation model [20], it gains promising performance. However, one major drawback of this method is the large distortions between the generated panoramic images and the actual images, introducing *noisy* visual features to the agent. Due to the pixel-wise training objective, image generation model [20] tends to fit the RGB values of local pixels, rather than focusing on key environmental semantics. Indeed, for unseen 3D environments, accurate RGB reconstruction is insurmountably difficult due to the high information redundancy of RGB images. In addition, generating high-resolution panoramic images is computing-intensive and time-consuming, increasing the delay of agent’s responses.

To anticipate future environments with higher quality and faster speed, we propose a pre-trained Hierarchical Neural Radiance (HNR) Representation Model that produces multi-level semantic representations of future candidate locations instead of generating panoramic images. Our semantic representations are learned through a vision-language embedding model (*i.e.*, CLIP [32]) that compresses the redundant information of RGB images and extracts the critical visual semantics associated with the language. Specifically at each step of navigation, the fine-grained grid features [44] extracted by the CLIP model are saved into the feature cloud with 3D positions and orientations. To predict the semantic features of future environments, we sample points along camera rays, and aggregate features around these points to produce latent feature vectors and volume density, then use volume rendering techniques to composite these values into small-scale region features. Furthermore, multi-level encoders are adopted to produce multi-level semantic features for future environments, supervised by CLIP embeddings.

The advantages of our method over previous methods for future environment prediction are three-fold. First, our model directly predicts robust multi-level semantic features for future candidate locations, avoiding the difficulty of pixel-level image reconstruction in unseen environments as used in existing methods like RNR-Map [24, 35] and DREAMWALKER [39]. Second, as many empty regions in the future views can be caused by visual occlusions, we employ a hierarchical encoding method to predict features of these empty regions by integrating surrounding contexts at both region and view levels. Lastly, our method using volume rendering handles better spatial relationships in the 3D environment, which is challenging for 2D image generation methods [25, 39] that lack depth perception.

With the predicted high-quality future views of candidate locations, we propose a lookahead VLN model to evaluate the possible next actions. As shown in Figure 1(b), we

predict more navigable locations (red nodes) in future environments and integrate them into a future path tree. The cross-modal graph encoder parallelly evaluates the matching score of different path branches in this path tree and selects the optimal candidate locations (blue nodes) to move.

In this work, our main contributions include:

- We propose a hierarchical neural radiance representation model to produce multi-level semantic representations for future environments with better quality and efficiency.
- Utilizing predicted representations of future environments, we propose a lookahead VLN model to parallelly evaluate possible future paths in the path tree, thus improving navigation planning.
- Extensive experiments demonstrate the effectiveness of our methods over existing methods in continuous vision-and-language navigation tasks.

2. Related Work

Vision-and-Language Navigation (VLN). VLN [5, 22, 23, 31] has received significant attention in recent years. The VLN tasks include step-by-step instructions such as R2R [5] and RxR [23], navigation with dialogs such as CVDN [37], and navigation for remote object grounding such as REVERIE [31] and SOON [48]. VLN-CE tasks convert the topologically-defined VLN tasks into continuous environments, such as R2R-CE [22] and RxR-CE [23]. Lots of previous methods focus on the representations of the visited environment during navigation. Among them, the recurrent unit [5, 12, 15, 36, 40], explicitly encoded history sequence [8, 29], topological map [4, 9], top-down semantic map [7, 13, 18] and grid-based map [3, 27, 44] are usually adopted to represent the visited environment.

Although the representation of the visited environment in VLN has received continuous attention, future environmental representation and lookahead strategy have not been fully explored. VLN-SIG [25] generates the semantics of future navigation views using visual codebook [34] for better decision-making. DREAMWALKER [39] utilizes an image generation model [20] to generate panoramic images of future environments and predict future actions. ULN [11] and Active Exploration [38] explore steps forward for robust navigation decision-making, and Tactical Rewind [19] performs tree search for planning. In contrast to them, we propose a hierarchical neural radiance representation model to anticipate future environments and construct the navigable future path tree for long-term planning.

Navigation with Neural Radiance Fields. The neural radiance field (NeRF) [28] has gained significant popularity in various AI tasks. NeRF predicts the RGB color and density of a sampled point in a scene so that an image from an arbitrary viewpoint can be rendered. However, the traditional NeRF methods with implicit MLP network can only synthesize view images in seen scenes, which makes it dif-

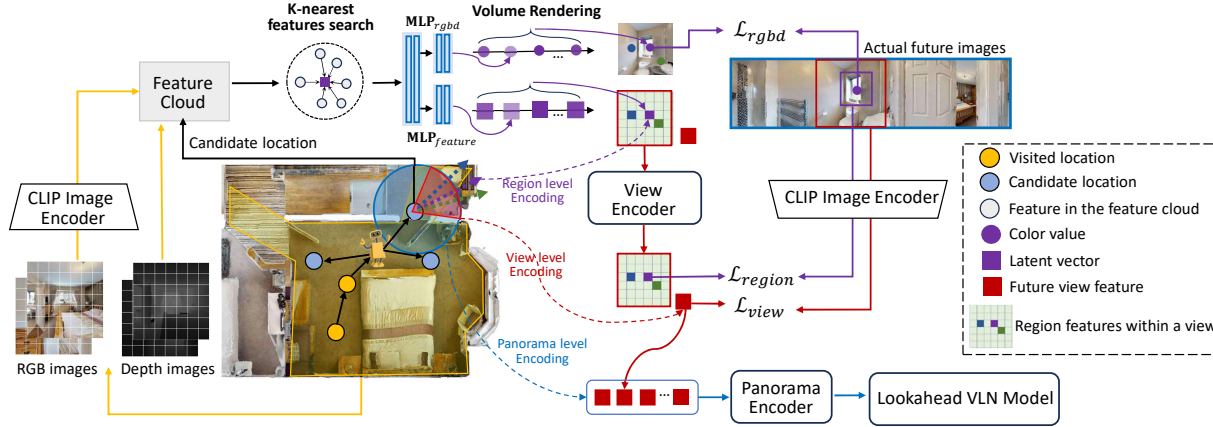


Figure 2. The framework of the hierarchical neural radiance representation model (HNR). The HNR model encodes the observed environments (yellow area) into the feature cloud. Through aggregating k-nearest features, the MLP network predicts the latent vector and volume density of sampled points along the rendered ray. A region-level representation is encoded by compositing these latent vectors via volume rendering, then a view encoder is used to encode all region-level representations within a future view (red area) and obtain an entire future view representation. All future views of the candidate location can be combined as a panorama (blue area) to support navigation.

difficult to generalize to the unseen scenes and adapt to many embodied AI tasks. To address this issue, GSN [10] proposes locally conditioned radiance fields, which encode observed images into a latent grid map, supporting rendering images in unseen environments. Based on GSN [10], RNR-Map [24] designs a localization framework using renderable neural radiance map for visual navigation, and Le-RNR-Map [35] embeds CLIP features into RNR-Map for query search with natural language. RNR-Map and Le-RNR-Map all encode features into a 2D latent grid map, which loses a lot of 3D geometric detail. The 3D feature cloud and hierarchical encoding in HNR provide better capabilities of spatial representation and multi-level context integration.

3. Method

3.1. Navigation Setups

The HNR model focuses on the VLN-CE [22, 23] tasks, where the agent navigates with low-level actions. Initialized at a starting location and given natural language instructions \mathcal{W} , the agent needs to explore the environment and reach the target location. At time step t , the agent observes panoramic RGB images $\mathcal{R}_t = \{r_{t,i}\}_{i=1}^{12}$ and the depth images $\mathcal{D}_t = \{d_{t,i}\}_{i=1}^{12}$ surrounding its current location (*i.e.*, 12 view images with 30 degrees separation).

During navigation, the agent’s visual observations are encoded and stored into the feature cloud. Meanwhile, a pre-trained waypoint predictor [16] is used to predict navigable candidates. For each navigable candidate, the HNR model predicts 12 future view representations using hierarchical encoding method. Firstly, HNR uses the volume rendering method to aggregate the features from the feature cloud and produce region-level embeddings. Secondly, region-level embeddings within the same future view are fed into the view encoder and obtain the entire future

view representation. These future representations of candidates help predict more and farther navigable locations (*i.e.*, lookahead nodes). By integrating lookahead node representations into a future path tree, the lookahead VLN model evaluates the match scores of different path branches and selects the optimal candidate location. Finally, a control module [4] is used to produce low-level actions to reach the selected candidate location.

3.2. Hierarchical Neural Radiance Representation

3.2.1 Feature Cloud Encoding

To encode the observed visual information, as shown in Figure 2, our HNR model stores the fine-grained visual features and their corresponding spatial information into the feature cloud \mathcal{M} . Specifically, at each navigation step t , for 12 observed RGB images $\mathcal{R}_t = \{r_{t,i}\}_{i=1}^{12}$, a pre-trained CLIP-ViT-B/32 [32] model is used to extract grid features $\{\mathbf{g}_{t,i} \in \mathbb{R}^{H \times W \times D}\}_{i=1}^{12}$. For convenience, all the subscripts (i, h, w) are denoted as j , where j ranges from 1 to J , and $J = 12 \cdot H \cdot W$. Through the downsized depth images $\{\mathbf{d}_{t,i} \in \mathbb{R}^{H \times W}\}_{i=1}^{12}$, each grid feature $\mathbf{g}_{t,j} \in \mathbb{R}^D$ is mapped to its 3D world position $P_{t,j} = [p_x, p_y, p_z]$ using camera pose $[\mathbf{R}, \mathbf{T}]$ and camera intrinsics \mathbf{K} as follows:

$$P_{t,j} = \left[\mathbf{d}_{t,j} \mathbf{R}^{-1} \mathbf{K}^{-1} \begin{bmatrix} h \\ w \\ 1 \end{bmatrix} - \mathbf{T} \right]^T \quad (1)$$

To better represent the spatial information of features in the feature cloud, we introduce the horizontal orientation $\theta_{t,j}$ and the size $s_{t,j}$ of each observed grid region:

$$s_{t,j} = 1/W \cdot \lceil \tan(\Theta_{HFOV}/2) \cdot \mathbf{d}_{t,j} \rceil \quad (2)$$

where W is the width of the feature map extracted by the CLIP model for each view, Θ_{HFOV} is the camera’s horizontal field-of-view. All these grid features and their spatial

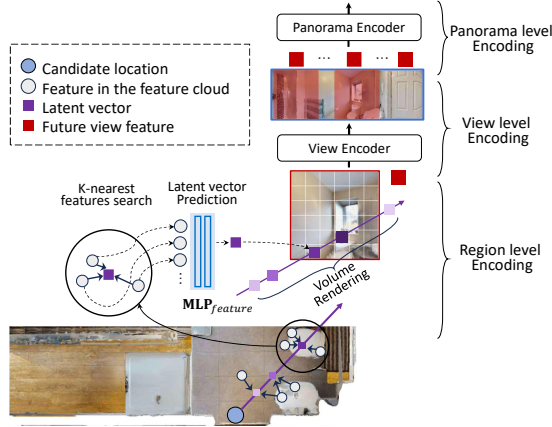


Figure 3. Illustration of the volume rendering method and hierarchical encoding.

information are stored into the feature cloud \mathcal{M} :

$$\mathcal{M}_t = \mathcal{M}_{t-1} \cup \{[\mathbf{g}_{t,j}, P_{t,j}, \theta_{t,j}, s_{t,j}]\}_{j=1}^J \quad (3)$$

3.2.2 Region Level Encoding via Volume Rendering

As shown in Figure 3, to produce a feature map $\mathbf{R} \in \mathbb{R}^{H \times W \times D}$ for each future view, given the encoded feature cloud \mathcal{M} and the camera pose of the future view, the HNR model predicts each region feature $\mathbf{R}_{h,w}$ via volume rendering method [28]. Specifically, for each region-level representation, the HNR model uniformly samples N points $\{\mathcal{P}_n | n = 1, \dots, N\}$ along the ray from the camera position \mathcal{P}_1 of the candidate location to the predicted region's center until 10 meters away. The KD-Tree algorithm [14] is used to search the k-nearest features $\{\mathbf{g}_k\}_{k=1}^K$ in feature cloud for each sampled point $\mathcal{P}_n = [\mathcal{X}_n, \mathcal{Y}_n, \mathcal{Z}_n]$. Then a $\text{MLP}_{feature}$ network is used to aggregate the k-nearest features of \mathcal{P}_n within radius R to produce a latent vector $\mathbf{r}_n \in \mathbb{R}^D$ and the volume density $\sigma_n \in \mathbb{R}^1$ as Equation 7. These latent vectors along the ray are composited into the region feature $\mathbf{R}_{h,w}$ via volume rendering as Equation 8.

To give the $\text{MLP}_{feature}$ network translational invariance for better generalization, the relative position P_k^{rel} and relative orientation θ_k^{rel} of k-nearest feature \mathbf{g}_k to sampled point \mathcal{P}_n is calculated as:

$$\begin{aligned} P_k^{rel} &= [x_k^{rel}, y_k^{rel}, z_k^{rel}] \\ &= [(x_k - \mathcal{X}_n) \cdot \cos\Theta_n + (y_k - \mathcal{Y}_n) \cdot \sin\Theta_n, \\ &\quad (y_k - \mathcal{Y}_n) \cdot \cos\Theta_n - (x_k - \mathcal{X}_n) \cdot \sin\Theta_n, \\ &\quad z_k - \mathcal{Z}_n] \end{aligned} \quad (4)$$

$$\theta_k^{rel} = [\sin(\theta_k - \Theta_n), \cos(\theta_k - \Theta_n)] \quad (5)$$

where $P_k = [x_k, y_k, z_k]$ is the world position of k-nearest feature \mathbf{g}_k , Θ_n denotes the horizontal orientation of the rendered ray. θ_k and s_k are the orientation and region size of \mathbf{g}_k . The positional embedding is encoded as:

$$\mathbf{q}_k = \text{LN}(\mathbf{W}_1[P_k^{rel}, \theta_k^{rel}, s_k]) \quad (6)$$

where LN denotes layer normalization and \mathbf{W}_1 is learnable parameters. With k-nearest features $\{\mathbf{g}_k\}_{k=1}^K$ and positional embeddings $\{\mathbf{q}_k\}_{k=1}^K$, the $\text{MLP}_{feature}$ network aggregate all k-nearest features of sampled point \mathcal{P}_n as follows:

$$[\mathbf{r}_n, \sigma_n] = \text{MLP}_{feature}(\{\mathbf{g}_k \oplus \mathbf{q}_k\}_{k=1}^K) \quad (7)$$

The volume density σ_n [28] can be interpreted as the differential probability of a ray terminating at point \mathcal{P}_n . As shown in Figure 3, the latent vector with a higher volume density has a higher contribution to the region feature. To reduce the computational cost, we adopt a *sparse sampling* strategy. Specifically, if the sampled point \mathcal{P}_n does not have any neighboring features within the radius \hat{R} , the latent vector \mathbf{r}_n and volume density σ_n will be directly set as zero, without using $\text{MLP}_{feature}$.

With all latent vectors $\{\mathbf{r}\}_{n=1}^N$ of N sampled points, we use the volume rendering method [28] to produce a region feature $\mathbf{R}_{h,w}$ for future view as follows:

$$\begin{aligned} \mathbf{R}_{h,w} &= \sum_{n=1}^N \tau_n (1 - \exp(-\sigma_n \Delta_n)) \mathbf{r}_n, \\ \text{where } \tau_n &= \exp(-\sum_{i=1}^{n-1} \sigma_i \Delta_i) \end{aligned} \quad (8)$$

τ_n represents volume transmittance, and Δ_n is the distance between adjacent sampled points. To enhance interaction among different region features, we use both region-level semantic alignment \mathcal{L}_{region} and view-level semantic alignment \mathcal{L}_{view} after view level encoding in Section 3.2.3.

Although the goal of region-level encoding is to generate regional semantic features, for image reconstruction and depth estimation, we also trained an MLP_{rgb} network to predict the color value \mathbf{c}_n and volume density $\hat{\sigma}_n$, then obtain the volume transmittance $\hat{\tau}$. As shown in Equation 9 and 10, the RGB pixel $\mathbf{C}_{h,w}$ and depth value $\mathbf{D}_{h,w}$ can be predicted via volume rendering. The rendering loss \mathcal{L}_{rgb} is the squared error between rendered pixels and ground truth.

$$\mathbf{C}_{h,w} = \sum_{n=1}^N \hat{\tau}_n (1 - \exp(-\hat{\sigma}_n \Delta_n)) \mathbf{c}_n \quad (9)$$

$$\mathbf{D}_{h,w} = \sum_{n=1}^N \hat{\tau}_n (1 - \exp(-\hat{\sigma}_n \Delta_n)) \mathbf{d}_n \quad (10)$$

where \mathbf{c}_n denotes color values, and \mathbf{d}_n denotes the distance between the sampled point \mathcal{P}_n and the camera position \mathcal{P}_1 .

3.2.3 View Level Encoding

The regional feature $\mathbf{R}_{h,w}$ obtained in Section 3.2.2 can only represent a small-scale region. To represent the entire future view and predict features of empty regions by integrating surrounding contexts, the regional feature map \mathbf{R}

together with a learnable view token \mathbf{V} is inputted into the view encoder and output the encoded $\hat{\mathbf{R}}$ and $\hat{\mathbf{V}}$. The view encoder consists of four-layer transformers.

As shown in Figure 2, to supervise the encoded region features $\hat{\mathbf{R}}$, we align each region feature $\hat{\mathbf{R}}_{h,w}$ with the CLIP embedding $\mathbf{R}_{h,w}^{gt}$ extracted from a small-scale cropped image within the actual future view. Similarly, the view feature $\hat{\mathbf{V}}$ is aligned with the CLIP embedding \mathbf{V}^{gt} extracted from an entire future view. During training, we randomly sample some region features and then minimize the loss between predicted features and actual CLIP embeddings, by maximizing cosine similarity as follows:

$$\mathcal{L}_{region} = \sum_{h,w} \left(1 - \frac{\hat{\mathbf{R}}_{h,w} \cdot \mathbf{R}_{h,w}^{gt}}{\|\hat{\mathbf{R}}_{h,w}\| \|\mathbf{R}_{h,w}^{gt}\|} \right) \quad (11)$$

$$\mathcal{L}_{view} = 1 - \frac{\hat{\mathbf{V}} \cdot \mathbf{V}^{gt}}{\|\hat{\mathbf{V}}\| \|\mathbf{V}^{gt}\|} \quad (12)$$

For better generalization ability, the HNR model is pre-trained in large-scale HM3D [33] dataset with 800 training scenes. Specifically, we randomly select a starting location in the scene and randomly move to a navigable candidate location at each step. At each step, parts of the unvisited candidate locations are randomly picked to predict future views. The model is optimized with the training loss $\alpha \mathcal{L}_{rgb} + \beta \mathcal{L}_{region} + \gamma \mathcal{L}_{view}$, where α, β, γ are the factors of proportionality. The maximum number of action steps per episode is set to 15.

3.2.4 Panorama Level Encoding

The pre-trained HNR model in Section 3.2.3 can be used to predict future view representations and combine them as panoramic representations of candidate locations through the panorama encoder, which is described in Figure 3. Given a navigable candidate location, the HNR model predicts 12 single-view semantic features and depth maps at 30 degrees separation. Furthermore, the waypoint predictor model [16] is used to predict navigable locations around the candidate location via these depth maps. The predicted future view features and navigable locations are encoded by the panorama encoder for lookahead exploration.

3.3. Architecture of the Lookahead VLN model

3.3.1 Node Embedding

Three types of nodes are utilized to structure the topological map, as shown in Figure 4. These nodes include visited nodes (yellow), candidate nodes (blue), and lookahead nodes (red). $\mathcal{V}_t^{visited}$ is represented as the current node embedding via average pooling 12 observed view representations. $\{\mathcal{V}_i^{visited}\}_{i=1}^t$ are all the visited node embeddings.

For each navigable candidate location, the HNR model predicts 12 future view features $\hat{\mathbf{V}}$ at 30 degrees separa-

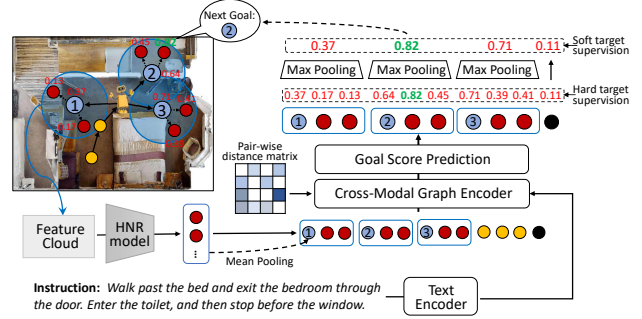


Figure 4. The framework of the lookahead VLN model. In addition to the stop embedding (black), three types of nodes are used to structure the topological map: visited nodes (yellow), candidate nodes (blue) and lookahead nodes (red).

tion. The panorama encoder consisting of a two-layer transformer is used to encode $\hat{\mathbf{V}}$ with corresponding position embedding and output features $\hat{\mathbf{V}}$. $\mathcal{V}^{candidate}$ is represented as the candidate node embedding via average pooling of $\hat{\mathbf{V}}$. The lookahead node embeddings are denoted as $\mathcal{V}^{lookahead}$, which are navigable parts of the future views $\hat{\mathbf{V}}$ with corresponding position embedding. The lookahead nodes connected to $\mathcal{V}_m^{candidate}$ are denoted as $\mathcal{V}_{\mathcal{N}(m)}^{lookahead}$. The ‘stop’ embedding \mathcal{V}^{stop} is added to the graph to denote a stop action and connect it with all other node embeddings.

3.3.2 Cross-Modal Graph Encoding

To encode the environmental topological map and evaluate future path branches in it, all node representations in Section 3.3.1 are fed into a 4-layer cross-modal transformer to conduct interaction. Each transformer layer consists of a cross-attention layer and a graph-aware self-attention layer (GASA). For cross-attention calculation, the encoded visited nodes $\mathcal{V}^{visited}$, candidate nodes $\mathcal{V}^{candidate}$, lookahead nodes $\mathcal{V}^{lookahead}$ and stop embedding \mathcal{V}^{stop} are used as query tokens, the word embeddings \mathcal{W} are used as key and value tokens. The standard self-attention layer only considers visual similarity among nodes, which may overlook nearby nodes that are more relevant than distant nodes. To this end, following [4, 9], we adopt the GASA layer that further takes into account the graph topology when computing inter-node attention for node encoding:

$$\text{GASA}(\mathcal{V}) = \text{Softmax}\left(\frac{\mathcal{V}\mathbf{W}_q(\mathcal{V}\mathbf{W}_k)^T}{\sqrt{d}} + E\mathbf{W}_e\right)\mathcal{V}\mathbf{W}_v \quad (13)$$

where \mathcal{V} denotes node representations, E is the pair-wise distance matrix obtained from the topological map, and $\mathbf{W}_q, \mathbf{W}_k, \mathbf{W}_e, \mathbf{W}_v$ are learnable parameters. After cross-modal graph encoding, the encoded node representations are denoted as $\hat{\mathcal{V}}^{visited}$, $\hat{\mathcal{V}}^{candidate}$, $\hat{\mathcal{V}}^{lookahead}$, and $\hat{\mathcal{V}}^{stop}$.

3.3.3 Lookahead Exploration and Action Prediction

The model predicts a navigation goal score for each node in the topological map as follows:

$$S = \text{FFN}(\hat{V}) \quad (14)$$

where **FFN** denotes a feed-forward network. To avoid unnecessary repeated visits to visited nodes, the scores for visited nodes are masked. With the candidate node $\hat{V}_m^{\text{candidate}}$ and those lookahead nodes $\hat{V}_{\mathcal{N}(m)}^{\text{lookahead}}$ connected to it, the VLN model predicts their goal scores $[S_m^{\text{candidate}}, S_{\mathcal{N}(m)}^{\text{lookahead}}]$ and take the maximum value among them as the goal score of the m -th path branch:

$$S^{\text{path}} = \{\text{Max}([S_m^{\text{candidate}}, S_{\mathcal{N}(m)}^{\text{lookahead}}])\}_m^M \quad (15)$$

Finally, the agent selects a candidate node to move according to the predicted goal scores (*i.e.*, select the path branch with the maximum score). The agent performs the stop action if the ‘stop’ score has the maximum value. If the selected candidate node is not adjacent to the current node, the agent computes the shortest path to the goal by performing Dijkstra’s algorithm on the topological map. The control module [4] is responsible for converting the topological plan into a series of low-level actions that guide the agent to the selected candidate node.

As shown in Figure 4, there are two navigation supervision strategies. The first one is the *Hard target* supervision $\mathcal{A}_{\text{hard}}$, which selects an unvisited node (candidate node or lookahead node) with the shortest distance to the destination. Then goal scores of all unvisited nodes are used to calculate the cross-entropy loss. The other one is the *Soft target* supervision $\mathcal{A}_{\text{soft}}$, which selects a candidate node with the shortest distance to the destination. The goal scores S^{path} (in Equation 15) of candidate nodes are calculated by max pooling all scores of the corresponding path branch shown in Figure 4. Then goal scores after max pooling are used to calculate the cross-entropy loss:

$$\mathcal{L}_{\text{nav}} = \text{CrossEntropy}(S^{\text{path}}, \mathcal{A}_{\text{soft}}) \quad (16)$$

4. Experiment

4.1. Datasets and Evaluation Metrics

We evaluate our model on the R2R-CE [22] and RxR-CE [23] datasets in continuous environments.

R2R-CE [22] is collected based on the discrete Matterport3D environments [6] with the Habitat simulator [33], enabling the agent to navigate in the continuous environments. It provides step-by-step instructions, and the average length is 32 words. The agent uses a 15° turning angle and the horizontal field-of-view is 90°.

RxR-CE [23] is a larger multilingual VLN dataset containing 126K instructions in English, Hindi, and Telugu. It

includes trajectories that are diverse in terms of length (average is 15 meters), which is more challenging in the continuous environments. The agent uses a 30° turning angle and the horizontal field-of-view is 79°.

There are several standard metrics [5] in VLN for evaluating the agent’s performance, including Trajectory Length (TL), Navigation Error (NE), Success Rate (SR), SR given the Oracle stop policy (OSR), Normalized inverse of the Path Length (SPL), Normalized Dynamic Time Warping (nDTW), and Success weighted by normalized Dynamic Time Warping (SDTW).

4.2. Comparison to State-of-the-Art Methods

Table 1 and 2 represent the performance of our proposed HNR model compared with existing VLN models on the R2R-CE and RxR-CE datasets respectively. Overall, HNR achieves state-of-the-art results in the majority of metrics, demonstrating the effectiveness of the proposed approach from diverse perspectives. As illustrated in Table 1, for the val unseen split of the R2R-CE dataset, our model outperforms our baseline method ETPNav [9] by 4% on SR and 2% on SPL. For the test unseen split, the proposed method outperforms ETPNav by 3% on SR and 2% on SPL. Meanwhile, as illustrated in Table 2, the proposed method also achieves the improvement of 2% in the majority of metrics on the RxR-CE dataset.

Compared with DREAMWALKER [39] in Table 1, which adopts a similar idea of lookahead exploration, our HNR model achieves performance improvement of about 10% on SR for all splits. On the one hand, our HNR constructs a unified future path tree to represent the future environment, which is better than the step-by-step path search strategy used in DREAMWALKER. The step-by-step search (*i.e.*, only evaluate one path branch at a time) makes it difficult to compare different future paths. On the other hand, the future environmental representations predicted by our HNR model are superior to those from the image generation model adopted by DREAMWALKER.

4.3. Ablation Study

As described in Section 1, the representations of future environments play a crucial role in lookahead exploration. In this section, we compare several popular representation strategies for future environments. (1) The most usual single-view representations [4], which rely on single-view visual observation of the current location to perceive candidates. (2) NeRF-based rendering methods [24] that focus on rendering RGB pixels for future view images and extract features via the CLIP model. (3) Image Generation method [20, 39] that imagine the panoramic RGBD images for candidate locations. (4) Our proposed hierarchical neural radiance representation model (HNR).

Methods	Val Seen					Val Unseen					Test Unseen				
	TL↓	NE↓	OSR↑	SR↑	SPL↑	TL↓	NE↓	OSR↑	SR↑	SPL↑	TL↓	NE↓	OSR↑	SR↑	SPL↑
CM ² [13]	12.05	6.10	50.7	42.9	34.8	11.54	7.02	41.5	34.3	27.6	13.9	7.7	39	31	24
WS-MGMap [7]	10.12	5.65	51.7	46.9	43.4	10.00	6.28	47.6	38.9	34.3	12.30	7.11	45	35	28
Sim-2-Sim [21]	11.18	4.67	61	52	44	10.69	6.07	52	43	36	11.43	6.17	52	44	37
CWP-RecBERT [16]	12.50	5.02	59	50	44	12.23	5.74	53	44	39	13.31	5.89	51	42	36
GridMM [44]	12.69	4.21	69	59	51	13.36	5.11	61	49	41	13.31	5.64	56	46	39
Reborn [2]	10.29	4.34	67	59	56	10.06	5.40	57	50	46	11.47	5.55	57	49	45
Ego ² -Map [17]	-	-	-	-	-	-	4.94	-	52	46	13.05	5.54	56	47	41
DREAMWALKER [39]	11.6	4.09	66	59	48	11.3	5.53	59	49	44	11.8	5.48	57	49	44
ScaleVLN [43]	-	-	-	-	-	-	4.80	-	55	51	-	5.11	-	55	50
BEVBert [3]	-	-	-	-	-	-	4.57	67	59	50	-	4.70	67	59	50
ETPNav [4]	11.78	3.95	72	66	59	11.99	4.71	65	57	49	12.87	5.12	63	55	48
HNR (Ours)	11.79	3.67	76	69	61	12.64	4.42	67	61	51	13.03	4.81	67	58	50

Table 1. Evaluation on the R2R-CE dataset.

Methods	Val Seen					Val Unseen					Test Unseen				
	NE↓	SR↑	SPL↑	NDTW↑	SDTW↑	NE↓	SR↑	SPL↑	NDTW↑	SDTW↑	NE↓	SR↑	SPL↑	NDTW↑	SDTW↑
CWP-CMA [16]	-	-	-	-	-	8.76	26.59	22.16	47.05	-	10.40	24.08	19.07	37.39	18.65
CWP-RecBERT [16]	-	-	-	-	-	8.98	27.08	22.65	46.71	-	10.40	24.85	19.61	37.30	19.05
Reborn [2]	5.69	52.43	45.46	66.27	44.47	5.98	48.60	42.05	63.35	41.82	7.10	45.82	38.82	55.43	38.42
ETPNav [4]	5.03	61.46	50.83	66.41	51.28	5.64	54.79	44.89	61.90	45.33	6.99	51.21	39.86	54.11	41.30
HNR (Ours)	4.85	63.72	53.17	68.81	52.78	5.51	56.39	46.73	63.56	47.24	6.81	53.22	41.14	55.61	42.89

Table 2. Evaluation on the RxR-CE dataset.

#	Representation Methods	NE↓	OSR↑	SR↑	SPL↑
1	Single View	4.71	64.71	57.21	49.15
2	NeRF Rendering	4.79	65.14	56.55	48.61
3	Image Generation	4.68	66.01	58.35	50.96
4	HNR	4.42	67.48	60.74	51.27
5	Ground truth	4.13	71.29	63.13	54.59
6	HNR w/o \mathcal{L}_{region}	4.55	66.78	60.20	50.74
7	HNR w/o Positional embeddings	4.78	64.87	56.17	48.91

Table 3. Comparison among different candidate location representation methods on the val unseen split of the R2R-CE dataset.

Comparisons among different representation methods.

For fair comparisons, rows 2-4 in Table 3 adopt the same lookahead VLN model shown in Section 3.3 for navigation planning but use different methods for future environments prediction. Traditional single-view representation (row 1) is restricted to a limited observable view and is difficult to comprehensively represent future environments due to the visual occlusions. The lookahead exploration strategy (rows 2-4) helps evaluate the agent’s future paths by anticipating the future environment for better navigation planning. NeRF-based RGB rendering method [24] has a low image reconstruction accuracy in unseen environments due to the visual occlusions and high information redundancy of RGB images, resulting in undesirable performance gains. The image generation method [39] has better performance gains than NeRF Rendering in unseen environments but is still inferior to our HNR model. To avoid the difficulty of pixel-level prediction, the proposed HNR model takes the CLIP embeddings as the prediction objective and predicts

multi-level semantics through hierarchical encoding.

Upper bound of our lookahead exploration strategy.

The ground truth (GT) representations (row 5) are extracted from the actual future view images via the CLIP model. It demonstrates the upper bound of our lookahead VLN model with GT future view prediction and also confirms the effectiveness of the lookahead exploration strategy.

Region level semantic alignment.

As shown in Table 3, without the training objective \mathcal{L}_{region} of region-level semantic alignment (row 6) in Section 3.2.3, the performance of the HNR model has degraded. Hierarchical encoding and multi-level semantic alignment help HNR integrate surrounding contexts and predict features of empty regions caused by visual occlusions.

Quality analysis of future environment prediction.

As illustrated in Figure 5, we evaluate the quality of predicted representations by comparing the embeddings’ cosine similarity between the predicted future views and ground truth. We observe a downward trend in cosine similarity as the distance between candidate location and agent increases. Combining Table 3 and Figure 5, it’s obvious that the better quality of future environmental representations makes the higher navigation performance gains. Using region-level alignment, our HNR method has the best representation quality and its cosine similarity is above 0.8 overall, supporting the lookahead exploration well within 5 meters.

Positional embeddings in region level encoding.

The positional embeddings calculated by Equation 6 represent

the spatial relationships for region feature encoding. Without the position and orientation of the k-nearest features relative to the sampled point (row 7) in Table 3, $MLP_{feature}$ network cannot accurately estimate the volume density and fails to perceive spatial relations in the 3D environment, resulting in a significant performance decrease.

Number of k-nearest features. Table 4 shows the effect of the number of k-nearest features on performance. We observe an upward trend in performance as the number of aggregated k-nearest features increases. Although the many more aggregated features make the region semantic prediction more accurate, the performance gains of the number over 4 are marginal, so we finally take the K value as 4 for lower computational cost.

Sparse sampling and k-nearest features search. Table 5 illustrates two important strategies used in our HNR to improve computational efficiency. The first one is the sparse sampling strategy described in Section 3.2.2, the $MLP_{feature}$ is not used to aggregate features in empty areas of the feature cloud. Due to the high proportion of empty space in 3D indoor environments, this strategy can greatly reduce the redundant calculation. To parallelly search the k-nearest features of all sampled points, the KD-Tree algorithm is adopted in HNR, as described in Section 3.2.2. Since the k-nearest features search has a heavy computational cost in our HNR model, we use the CUDA implementation of KD-Tree algorithm [14] to accelerate it, much faster than other nearest neighbor search methods [30] (i.e., w/o KD-Tree) in Table 5.

Performance analysis of the lookahead exploration. In this part, we analyze different future path evaluation strategies and navigation supervision methods. As shown in Table 6, $S^{lookahead}$ indicates that the VLN model uses the scores of the lookahead nodes to evaluate the future paths, and $S^{candidate}$ indicates the usage of the candidate node scores, as shown in Equation 15. *Soft target* and *Hard target* are two different types of navigation supervision described in Section 3.3.3. In Table 6, row 1 has the lowest performance without the lookahead exploration strategy. Row 2 doesn't use the lookahead node scores to evaluate the future paths and gain marginal performance improvement, confirming the necessity of the lookahead node representation. The performance of row 5 is superior to rows 3-4, showing the *Soft target* supervision is the better navigation supervision. The lookahead node closest to the destination (i.e., *Hard target*) is not sure of the highest semantic match score with the instruction due to visual occlusions and prediction noise, compulsively overfitting the *Hard target* compromises the navigation performance.




Figure 5. Average cosine similarity between predicted future views and ground truth at different distances between candidate locations and agent, on the val unseen split of the R2R-CE dataset.

Number of nearest features	NE↓	OSR↑	SR↑	SPL↑
1	4.68	65.63	57.86	49.21
2	4.57	66.99	59.49	50.81
4	4.42	67.48	60.74	51.27
6	4.37	67.70	60.58	51.30

Table 4. The effect of different numbers of nearest features in the HNR model on the val unseen split of the R2R-CE dataset.

HNR	w/o Sparse sampling	w/o KD-Tree
87.3 Hz (11.5 ms)	5.9 Hz (169.5ms)	23.1 Hz (42.3 ms)

Table 5. Runtime analysis measured on one RTX 3090 GPU.

# $S^{lookahead}$	$S^{candidate}$	Soft	Hard	NE↓	OSR↑	SR↑	SPL↑
1				4.71	64.71	57.21	49.15
2	✓	✓		4.71	66.43	57.75	49.54
3	✓	✓	✓	4.81	65.85	56.72	48.41
4	✓	✓	✓	4.62	67.92	59.87	50.03
5	✓	✓	✓	4.42	67.48	60.74	51.27

Table 6. Ablation study of the lookahead VLN model.

5. Conclusion

In this paper, we propose a lookahead exploration method for continuous vision-language navigation. Our proposed HNR model predicts multi-level future environmental representations through volume rendering and hierarchical encoding, supervised by CLIP embeddings. With predicted future representations, the lookahead VLN model constructs the navigable future path tree and selects the optimal path branch via efficient parallel evaluation. Extensive experiments demonstrate the accuracy of the HNR model's representations and the excellent performance of the lookahead VLN model. The lookahead exploration method helps navigation planning by predicting the future outcome of actions, which has great research potential for VLN and Embodied AI tasks.

Acknowledgment. This work was supported in part by the National Natural Science Foundation of China under Grants 62125207, 62102400, 62272436, and U23B2012, in part by the National Postdoctoral Program for Innovative Talents under Grant BX20200338.

References

- [1] Arpit Agarwal, Katharina Muelling, and Katerina Fragkiadaki. Model learning for look-ahead exploration in continuous control. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3151–3158, 2019. **1**
- [2] Dong An, Zun Wang, Yangguang Li, Yi Wang, Yicong Hong, Yan Huang, Liang Wang, and Jing Shao. 1st place solutions for rxr-habitat vision-and-language navigation competition (cvpr 2022). *arXiv preprint arXiv:2206.11610*, 2022. **7**
- [3] Dong An, Yuankai Qi, Yangguang Li, Yan Huang, Liang Wang, Tieniu Tan, and Jing Shao. Bevbort: Multimodal map pre-training for language-guided navigation. In *ICCV*, pages 2737–2748, 2023. **2, 7**
- [4] Dong An, Hanqing Wang, Wenguan Wang, Zun Wang, Yan Huang, Keji He, and Liang Wang. Etpnav: Evolving topological planning for vision-language navigation in continuous environments. *arXiv preprint arXiv:2304.03047*, 2023. **2, 3, 5, 6, 7**
- [5] 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 *CVPR*, pages 3674–3683, 2018. **1, 2, 6**
- [6] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. In *3DV*, pages 667–676, 2017. **6**
- [7] Peihao Chen, Dongyu Ji, Kunyang Lin, Runhao Zeng, Thomas H Li, Mingkui Tan, and Chuang Gan. Weakly-supervised multi-granularity map learning for vision-and-language navigation. In *NeurIPS*, 2022. **2, 7**
- [8] Shizhe Chen, Pierre-Louis Guhur, Cordelia Schmid, and Ivan Laptev. History aware multimodal transformer for vision-and-language navigation. In *NeurIPS*, pages 5834–5847, 2021. **1, 2**
- [9] 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 *CVPR*, pages 16537–16547, 2022. **1, 2, 5, 6**
- [10] Terrance DeVries, Miguel Angel Bautista, Nitish Srivastava, Graham W Taylor, and Joshua M Susskind. Unconstrained scene generation with locally conditioned radiance fields. In *ICCV*, pages 14304–14313, 2021. **3**
- [11] Weixi Feng, Tsu-Jui Fu, Yujie Lu, and William Yang Wang. Uln: Towards underspecified vision-and-language navigation. *arXiv preprint arXiv:2210.10020*, 2022. **1, 2**
- [12] 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 *NeurIPS*, 2018. **2**
- [13] Georgios Georgakis, Karl Schmeckpeper, Karan Wanchoo, Soham Dan, Eleni Miltsakaki, Dan Roth, and Kostas Daniilidis. Cross-modal map learning for vision and language navigation. In *CVPR*, 2022. **2, 7**
- [14] Thomas Grandits, Alexander Effland, Thomas Pock, Rolf Krause, Gernot Plank, and Simone Pezzuto. GEASI: Geodesic-based earliest activation sites identification in cardiac models. *International Journal for Numerical Methods in Biomedical Engineering*, 37(8):e3505, 2021. **4, 8**
- [15] Yicong Hong, Qi Wu, Yuankai Qi, Cristian Rodriguez-Opazo, and Stephen Gould. Vln bert: A recurrent vision-and-language bert for navigation. In *CVPR*, pages 1643–1653, 2021. **2**
- [16] 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 *CVPR*, 2022. **3, 5, 7**
- [17] Yicong Hong, Yang Zhou, Ruiyi Zhang, Franck Dornoncourt, Trung Bui, Stephen Gould, and Hao Tan. Learning navigational visual representations with semantic map supervision. In *ICCV*, pages 3055–3067, 2023. **7**
- [18] Chenguang Huang, Oier Mees, Andy Zeng, and Wolfram Burgard. Visual language maps for robot navigation. In *ICRA*, London, UK, 2023. **2**
- [19] 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. **2**
- [20] Jing Yu Koh, Honglak Lee, Yinfei Yang, Jason Baldrige, and Peter Anderson. Pathdreamer: A world model for indoor navigation. In *ICCV*, pages 14738–14748, 2021. **2, 6**
- [21] Jacob Krantz and Stefan Lee. Sim-2-sim transfer for vision-and-language navigation in continuous environments. In *ECCV*, 2022. **7**
- [22] Jacob Krantz, Erik Wijmans, Arjun Majumdar, Dhruv Batra, and Stefan Lee. Beyond the nav-graph: Vision-and-language navigation in continuous environments. In *ECCV*, 2020. **1, 2, 3, 6**
- [23] Alexander Ku, Peter Anderson, Roma Patel, Eugene Ie, and Jason Baldrige. Room-across-room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In *EMNLP*, pages 4392–4412, 2020. **1, 2, 3, 6**
- [24] Obin Kwon, Jeongho Park, and Songhwai Oh. Renderable neural radiance map for visual navigation. In *CVPR*, pages 9099–9108, 2023. **2, 3, 6, 7**
- [25] Jialu Li and Mohit Bansal. Improving vision-and-language navigation by generating future-view image semantics. In *CVPR*, pages 10803–10812, 2023. **1, 2**
- [26] Xiangyang Li, Zihan Wang, Jiahao Yang, Yaowei Wang, and Shuqiang Jiang. Kerm: Knowledge enhanced reasoning for vision-and-language navigation. In *CVPR*, pages 2583–2592, 2023. **1**
- [27] Rui Liu, Xiaohan Wang, Wenguan Wang, and Yi Yang. Bird’s-eye-view scene graph for vision-language navigation. In *ICCV*, pages 10968–10980, 2023. **2**
- [28] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. **2, 4**

- [29] Alexander Pashevich, Cordelia Schmid, and Chen Sun. Episodic transformer for vision-and-language navigation. In *ICCV*, 2021. [2](#)
- [30] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of machine Learning research*, 12:2825–2830, 2011. [8](#)
- [31] 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 *CVPR*, pages 9982–9991, 2020. [1](#), [2](#)
- [32] 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 *ICML*, pages 8748–8763, 2021. [2](#), [3](#)
- [33] Santhosh K Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alex Clegg, John Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X Chang, et al. Habitat-matterport 3d dataset (hm3d): 1000 large-scale 3d environments for embodied ai. *arXiv preprint arXiv:2109.08238*, 2021. [5](#), [6](#)
- [34] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, pages 8821–8831. PMLR, 2021. [2](#)
- [35] Francesco Taioli, Federico Cunico, Federico Girella, Riccardo Bologna, Alessandro Farinelli, and Marco Cristani. Language-enhanced rnr-map: Querying renderable neural radiance field maps with natural language. In *ICCV*, pages 4669–4674, 2023. [2](#), [3](#)
- [36] Hao Tan, Licheng Yu, and Mohit Bansal. Learning to navigate unseen environments: Back translation with environmental dropout. In *NAACL*, pages 2610–2621, 2019. [2](#)
- [37] Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer. Vision-and-dialog navigation. In *PMLR*, 2020. [2](#)
- [38] Hanqing Wang, Wenguan Wang, Tianmin Shu, Wei Liang, and Jianbing Shen. Active visual information gathering for vision-language navigation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16*, pages 307–322. Springer, 2020. [2](#)
- [39] Hanqing Wang, Wei Liang, Luc Van Gool, and Wenguan Wang. Dreamwalker: Mental planning for continuous vision-language navigation. In *ICCV*, pages 10873–10883, 2023. [2](#), [6](#), [7](#)
- [40] 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 *CVPR*, pages 6629–6638, 2019. [2](#)
- [41] Xiaohan Wang, Yuehu Liu, Xinhang Song, Beibei Wang, and Shuqiang Jiang. Generating explanations for embodied action decision from visual observation. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 2838–2846, 2023. [1](#)
- [42] Xiaohan Wang, Yuehu Liu, Xinhang Song, Beibei Wang, and Shuqiang Jiang. Camp: Causal multi-policy planning for interactive navigation in multi-room scenes. *Advances in Neural Information Processing Systems*, 36, 2024. [1](#)
- [43] Zun Wang, Jialu Li, Yicong Hong, Yi Wang, Qi Wu, Mohit Bansal, Stephen Gould, Hao Tan, and Yu Qiao. Scaling data generation in vision-and-language navigation. In *ICCV*, pages 12009–12020, 2023. [7](#)
- [44] Zihan Wang, Xiangyang Li, Jiahao Yang, Yeqi Liu, and Shuqiang Jiang. Gridmm: Grid memory map for vision-and-language navigation. In *ICCV*, pages 15625–15636, 2023. [2](#), [7](#)
- [45] Sixian Zhang, Xinhang Song, Yubing Bai, Weijie Li, Yakui Chu, and Shuqiang Jiang. Hierarchical object-to-zone graph for object navigation. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 15110–15120. IEEE, 2021. [1](#)
- [46] Sixian Zhang, Weijie Li, Xinhang Song, Yubing Bai, and Shuqiang Jiang. Generative meta-adversarial network for unseen object navigation. In *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXIX*, pages 301–320, 2022.
- [47] Sixian Zhang, Xinhang Song, Weijie Li, Yubing Bai, Xinyao Yu, and Shuqiang Jiang. Layout-based causal inference for object navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10792–10802, 2023. [1](#)
- [48] Fengda Zhu, Xiwen Liang, Yi Zhu, Qizhi Yu, Xiaojun Chang, and Xiaodan Liang. Soon: Scenario oriented object navigation with graph-based exploration. In *CVPR*, pages 12689–12699, 2021. [2](#)