

DREAMWALKER: Mental Planning for Continuous Vision-Language Navigation

Hanqing Wang^{1,2} Wei Liang^{1,4*} Luc Van Gool² Wenguan Wang^{3*}

¹Beijing Institute of Technology ²Computer Vision Lab, ETH Zurich ³ReLER, CCAI, Zhejiang University

⁴Yangtze Delta Region Academy of Beijing Institute of Technology, Jiaxing

<https://github.com/hanqingwangai/Dreamwalker>

Abstract

VLN-CE is a recently released embodied task, where AI agents need to navigate a freely traversable environment to reach a distant target location, given language instructions. It poses great challenges due to the huge space of possible strategies. Driven by the belief that the ability to anticipate the consequences of future actions is crucial for the emergence of intelligent and interpretable planning behavior, we propose DREAMWALKER — a **world model** based VLN-CE agent. The world model is built to summarize the visual, topological, and dynamic properties of the complicated continuous environment into a discrete, structured, and compact representation. DREAMWALKER can simulate and evaluate possible plans entirely in such internal abstract world, before executing costly actions. As opposed to existing model-free VLN-CE agents simply making greedy decisions in the real world, which easily results in shortsighted behaviors, DREAMWALKER is able to make strategic planning through large amounts of “mental experiments.” Moreover, the imagined future scenarios reflect our agent’s intention, making its decision-making process more transparent. Extensive experiments and ablation studies on VLN-CE dataset confirm the effectiveness of the proposed approach and outline fruitful directions for future work.

1. Introduction

For decades, the AI community has strived to develop intelligent robots that can understand human instructions and carry them out. As a small step towards this long-held goal, vision-language navigation (VLN) [6] — the task of entailing autonomous agents to navigate in *never-before-seen* 3D environments with language instructions — gained growing attention. In the standard VLN setting, agent’s movement is constrained to a small set of pre-defined sparse locations. As pointed out by [45], such over-simplified, discrete task setup involves many unrealistic assumptions such as known

*Corresponding authors.

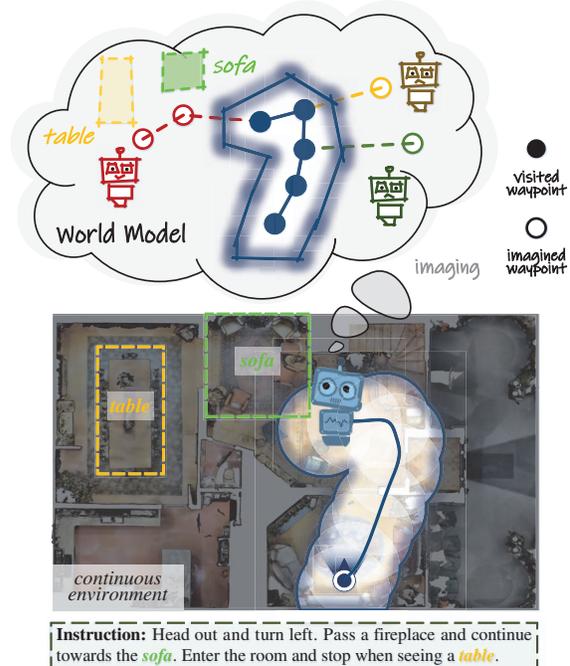


Figure 1: In partially observable, continuous VLN environments, DREAMWALKER maps its surrounding into a discrete and structured abstraction. In this internal world, it is able to conduct mental planning (👁️🧭) by imagining future scenarios, before taking real action.

topology, perfect localization, and deterministic transition.

To better reflect the challenges of real world navigation, Krantz *et al.* [45] update the discrete VLN to a continuous version – VLN-CE (VLN in continuous environments), where the agent is free to traverse any unobstructed location with low-level actions. VLN-CE proved much more challenging than its discrete counterpart: the performance gap between the state-of-the-arts in the two settings is more than 20%, in terms of episode success rate. The main challenge posed by VLN-CE lies in the demand of strategic planning in continuous environments with low-level actions.

As a direct response, we developed a *world model* based VLN-CE agent, called DREAMWALKER. Previous studies in cognitive science [17, 34, 35] suggest that humans build a mental model of the local surrounding, based on our limited

senses. This internal world model summarizes our knowledge about the environment and serves as the basis for many high-level meta-skills, *e.g.*, reasoning, planning, decision-making, and interpretation. The world model theory is one source of the idea of model-based Reinforcement Learning (RL) [73] and promotes many recent advances in robot control [69, 62, 58, 91]. Keeping this grand idea in head, we let DREAMWALKER explicitly abstract crucial characteristics of its continuous surrounding environment to a *discrete, structured* representation (Fig. 1). This allows DREAMWALKER to “imagine” a lot of future possible navigation plans and evaluate the corresponding consequences entirely in the mind, before taking actual low-level actions in the real world. In this way, DREAMWALKER takes the challenge of VLN-CE head-on: mental planning with discrete world model enables efficient navigation behavior in continuous environments.

Technically, the world model is built upon agent’s past experiences and can make predictions about the future. It contains two parts: **i)** An *environment graph* (EG) is constructed as a composition of selected or predicted waypoints and their typological relations. EG collects agent’s temporary knowledge about its surrounding. **ii)** A learnable *scene synthesizer* (SS) predicts future observations from a waypoint with multiple steps. SS embeds agent’s stable knowledge about environments, such as general room layout rules and transition dynamics, into its network parameters. Based on the world model, DREAMWALKER synthesizes various future navigation trajectories, and assesses their progress towards the final target location. Then, the best mental plan is found by Monte Carlo Tree Search [41] and executed in the continuous world with low-level actions. With the navigation proceeds, EG is further updated for making a new round of mental planning.

Notably, our DREAMWALKER significantly distinguishes itself from prior VLN-CE solutions [67, 29, 43, 44] in the following aspects: **i)** Recent advanced solutions are essentially model-free methods. While in principle a representation of the environment could be *implicitly* learned through model-free RL, the reinforcement signal may be too weak to quickly learn such a representation and how to make use of it. In contrast, our agent plans its actions within an explicit, and abstract model of the continuous environment. **ii)** Existing agents navigate by greedily and reactively choosing between a small set of nearby waypoints, based on their hidden state which compresses past observations. They tend to be shortsighted, due to the absence of reliable strategies for capturing information for achieving the future [23]. Yet DREAMWALKER can use the world model to anticipate the impacts of possible actions and plan strategic behavior. **iii)** The future scenarios created by the world model explain the intention of DREAMWALKER in a way that human can understand, making its behaviors more interpretable [9, 79].

Extensive experiments on VLN-CE dataset [45] confirm that our DREAMWALKER gains promising performance with

the appealing ability of real-time behavioral interpretation. This work is expected to foster future research in developing more strategic, robust, and interpretable VLN-CE agents.

2. Related Work

VLN in Discrete Environments. The release of R2R dataset [6] has stimulated the emergence of various VLN systems and datasets. In particular, early VLN agents are built upon recurrent neural network based sequence-to-sequence model [6, 18, 77, 88, 86, 31, 52, 93, 38, 53, 84, 82, 2], while the recent ones explore graph neural networks [15, 83] and non-local transformer models [30, 60, 10] for long-term planning, with the assist of environment map building [92, 11, 51], cross-modal matching [81, 94, 87], and multi-modal pre-training [54, 26, 64, 65, 3] techniques. Besides the advance in model design, several more challenging VLN datasets, such as R4R [32], RxR [46], and REVERIE [63], are developed to address long-term navigation [32, 46] and concise instruction guided navigation [63].

These remarkable efforts follow the seminal work of R2R to assume the agent can ‘perfectly’ teleport to and from a fixed small set of locations (pre-stored in a sparse navigation graph) in the Matterport3D [8] environment. Although facilitating the evolution and evaluation of VLN algorithms, such discrete task setup is too simplified to cover the practical challenges that a robot would encounter while navigating the real world, such as environment topology acquirement and localization error. In this work, we focus on performing VLN in continuous environments, setup by [45].

VLN in Continuous Environments (VLN-CE). Krantz *et al.* [45] lift the discrete R2R VLN task setup to the continuous setting — VLN-CE, where the embodied agent is initiated in freely traversable 3D environments and must execute low-level actions to follow natural language navigation instructions. After throwing away the unrealistic assumption of the navigation graph, the VLN task becomes more challenging and closer to the real-world. Later approaches attempt to reproduce the success of the abstract VLN in the continuous counterpart, by building a high-level action space based on online prediction of waypoints [67, 29, 43, 80, 4]. More recently, [44] explores transferring pre-trained VLN policies to continuous environments, demonstrating advantages over training VLN-CE policies from scratch.

Our agent is favored due to **i)** its model-based nature, and **ii)** the ability of strategic and interpretable planning. First, our agent learns and builds an explicit model of the world, freeing navigation policy learning from environment modeling. Second, by visualizing possible futures, our agent is able to develop advanced plan before moving and explain its behaviors. In contrast, current agents make only greedy decisions purely relying on their latent state. Their planning ability is rather weak and the decision mode is non-transparent.

World Model. Equipping robot machines with the ability to

build world models is viewed as a key step towards the next-generation of AI [34, 50]. Capturing high-level aspects of the environment, world models help enable AI agents for reasonable decision-making through simulation (a.k.a., imagination) of possible futures. Towards this direction, a bunch of approaches have been recently developed to predict forward dynamics/representations [59, 37, 12, 20, 39, 49], perform sampling-based planning [72, 61, 91], conduct self-simulation based policy learning [89, 85, 22], and reconcile the advantages of model-based and model-free RL regimes [66, 56, 75, 47]. Some others leverage world models to anticipate (short-term) targets for goal-conditioned policy learning [58, 57], boost the learning of world models by introducing belief of state uncertainty [13, 20, 19], informative state representation [48, 25, 76], and model regularization [50, 14, 40]. Although world models find successful applications in Atari [36, 24, 23, 25] and robot motion planning [16, 69], many of them are restricted to relatively simple tasks or low-dimensional environments, pointed out by [23, 25, 58, 42]. To date, world models have been rarely investigated in visually-rich, embodied tasks. The few exceptions [49, 42] simply use a viewpoint synthesis network as the world model. They either pre-access the entire environment typology [42], or only treat the synthesized observations as sub-goals without sampling-based planning [49]. In [1], NeRF-based world model is adopted for drone navigation planning, which requires pre-exploration of the environment. For ours, in addition to the challenging task setup, our world model is organized as a structured summarization of environment layout and a compact representation of environment evolution, instead of [49, 42] encoding all the knowledge about the environment into latent network parameters. More importantly, our explicit and discrete abstraction of the continuous environments offers a suitable testbed for mental experiments. Through sampling and assessing numerous possible solutions entirely in the mind, our agent can conduct strategic and global navigation planning before moving.

Monte Carlo Tree Search (MCTS). MCTS [41] is a heuristic search algorithm for approximating optimal choices in large search spaces. It has demonstrated great success in creating game-playing AI agents and solving sequential decision tasks, such as Atari games [21, 71] and Go [73, 74, 71].

In this work, we adopt MCTS to search for possible navigation plans based on the world model. MCTS has so far been rarely explored in the field of VLN-CE. In addition, in many application scenarios of MCTS, all the aspects of the world states are fully observable. However, the continuous VLN-CE environments are partially observable, and hence we run MCTS in a discrete, synthesized world space.

3. Methodology

Task Setup and Notations. In VLN-CE [45], AI agents are required to navigate in never-before-seen 3D environments

to target positions, according to language instructions. The environments are assumed to be continuous open space. At each step, the agent must choose a low-level action from $\{\text{turn-left } 15^\circ, \text{turn-right } 15^\circ, \text{move-forward } 0.25\text{m}, \text{stop}\}$, given the instruction X and 360° panoramic RGBD observation Y . The navigation is successful only if the agent selects **stop** within 3m of the target location within 500 steps.

Waypoint Action Space. Recent VLN-CE solutions [67, 29, 43, 44] adopt a high-level waypoint action space. During navigation, the agent uses a *Waypoint Predictor* to get a heatmap of 120 angles-by-12 distances, which highlights navigable waypoints in its surrounding. Each angle is of 3 degrees, and the distances range from 0.25 meters to 3.00 meters with 0.25 meters separation, respectively corresponding to the turning angle and forward step size of the low-level action space (Fig. 2(f)). In this way, the problem of inferring low-level controls is translated to the task of waypoint prediction and selection. Please refer to [29] for more details.

Algorithmic Overview. Our DREAMWALKER agent solves VLN-CE through **i)** building a world model (§3.1) that explicitly abstracts intrinsic properties of continuous environments into a structured, discrete representation; **ii)** conducting strategic planning in the discrete mental world before taking actual actions in the continuous environment (§3.2).

3.1. World Model: Structured, Discrete, and Compact Abstraction of Continuous Environments

To efficiently plan and act in the highly complex world, humans develop a mental model to represent our knowledge about the surrounding, based on our past daily experiences and current information perceived by limited sense [22]. In view of this, our world model is built as combination of two parts: **i)** An *environment graph* (EG) – a structured and discrete representation of agent’s temporary knowledge about the visual landmarks and typologies of the current partially observed environment; and **ii)** A *scene synthesizer* (SS) that encodes agent’s stable knowledge about some general rules of environments such as transition dynamics and room layout, which are learned from training experiences and utilized to predict the unobserved portions of the current environment.

Environment Graph (EG). EG is organized as an *episodic graph* $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where nodes $v \in \mathcal{V}$ denote previously visited waypoints and edges $e_{u,v} \in \mathcal{E}$ store geometric relations between waypoints (Fig. 2(a)(g)). At the beginning of current navigation episode, \mathcal{G} only contains one single node — the starting location. Then, the agent predicts a set of accessible waypoints, and selects one of them to navigate. After reaching the selected waypoint, \mathcal{G} will be updated by involving this waypoint. Hence EG evolves with the navigation proceeds. Concretely, for each node $v \in \mathcal{V}$, its embedding is the feature of the corresponding waypoint observation:

$$v = Y_{pv}, \quad (1)$$

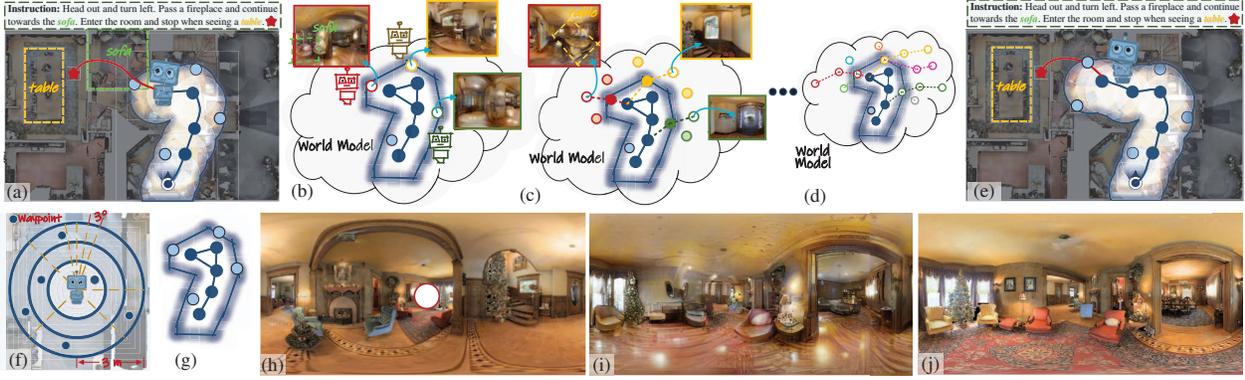


Figure 2: (a) Top-down view of current navigation, where \bullet indicates previously visited waypoints and \circ refers to detected but unvisited waypoints. (b) DREAMWALKER synthesizes future observations at unvisited waypoints $\circ\circ$ through SS. (c-d) With the synthesized observations, DREAMWALKER further extends the synthesized trajectories and looks deeper into the future. (e) DREAMWALKER selects the best mental plan for execution. After reaching the selected waypoint, it starts next-round planning. (f) Top-down view of the waypoint action space. (g) EG G of (a). (h-j) DREAMWALKER images its future observation at the unvisited waypoint \circ , based on its current observation. The synthesized and real observations at \circ are given in (i) and (j), respectively. For clarity, only RGB observation is provided. Notably, the imagined scenarios explain agent’s inner decision mode in a way that human can understand, leading to improved interpretability.

where p_v refers to the location of waypoint v , and the coordinate of the start point is set as $(0, 0)$; \mathbf{Y}_{p_v} indicates the embedding of the panoramic observation Y_{p_v} at location p_v .

For each edge $e_{u,v} \in \mathcal{E}$, its embedding $e_{u,v}$ encodes topological relations between waypoints u and v :

$$e_{u,v} = (\cos \theta_{u,v}, \sin \theta_{u,v}, \Delta p_{u,v}), \quad (2)$$

where $\theta_{u,v}$ and $\Delta p_{u,v} = p_u - p_v$ refer to the angle and distance between u and v , respectively. Note that the connectivity among waypoints is also captured by \mathcal{G} . An edge, $e_{u,v}$, exists only if u and v are connected, that is, they are detected as valid in the waypoint prediction heatmap of each other.

As such, EG is built upon the information gathered during current navigation episode. Hence it represents the agent’s temporary knowledge about its observed surroundings, and organizes them in a concise, structured, and discrete manner. **Scene Synthesizer (SS)**. SS is a generative network that predicts future scenes based on agent’s past observations only, without having to navigate them (Fig. 2(h-j)). Given panoramic RGBD observation Y_p perceived at position p , SS is to render a plausible, full-resolution panoramic RGBD observation $\hat{Y}_{p'}$ at an unvisited position p' . The position p' of interest is a waypoint of p , thus p and p' are spatially close. Like previous world structure-aware video synthesis methods [55, 42], we first project Y_p to a 3D point cloud using the depth information and re-project this point cloud into the observation space at position p' , so as to obtain a sparse, geometry-aligned RGBD panoramic image $Y_{p \rightarrow p'}$. The SS network takes as inputs both the observation Y_p perceived at p and the reconstructed observation $Y_{p \rightarrow p'}$, and synthesizes the observation $\hat{Y}_{p'}$ for the unvisited waypoint position p' :

$$\hat{Y}_{p'} = \text{Scene_Synthesizer}(Y_p, Y_{p \rightarrow p'}). \quad (3)$$

SS learns from experiences of ‘viewing’ numerous VLN-CE training environments. Its parameters encode agent’s statis-

tical knowledge about some fundamental constraints in the world, such as the transition dynamics and general rules behind room arrangement (e.g., ‘dishwasher is typically located in the kitchen’).

Working Mode of World Model. Integrating EG and SS together leads to a powerful world model. EG stores agent’s understanding of the observed portions of the environment. Based on working memory, SS makes use of statistical knowledge to forecast the unexplored portion of the environment (Fig. 2(a-e)). For instance, starting from a waypoint v in \mathcal{G} , the world model uses SS to imagine the future (i.e., $\hat{Y}_{p_{v'}}$) if the agent navigates to a previously unvisited waypoint v' of v . With the synthesized observation $\hat{Y}_{p_{v'}}$ at waypoint v' , the world model can make further future prediction. Notably, our world model, or more precisely, its parametric part – SS, is trained independently, which lifts the navigation policy from learning inherent structures of environments.

3.2. Mental Planning: Forecasting the Future in the World Model

Powered by the world model, DREAMWALKER gains the ability of anticipating the consequences of its actions, for an extended period into the future. This is because the world model can simulate how the environment changes in response to agent’s actions. DREAMWALKER can thus make advanced planning, through perform mental experiments in the simulated world. Basically, at every navigation-decision making step, DREAMWALKER first uses the world model to synthesize many future navigation trajectories, and then selects the best one for execution. Here we adopt Monte Carlo Tree Search (MCTS) [41, 21], a powerful approach for decision problems, to achieve world model based online planning.

MCTS based Mental Planning. As a best-first tree search algorithm, MCTS represents the space of candidates into a

tree and finds an optimal solution in an iterative manner. In our case, each tree node s represents a possible world state: the root node s_0 is the world state \mathcal{G} at current navigation step, while other nodes indicate future world states, synthesized by the world model. The edge from node s to s' denotes the action a taken in s to reach s' , and is identified by pair (s, a) .

The core idea of MCTS based planning is to gradually expand the search tree and evaluate rewards, *i.e.*, create many plans and assess the outcomes in mind. This is achieved by an iterative tree search process. Each search iteration starts from the root state s_0 , and sequentially samples states and actions, based on the simulation of four phases (Fig. 3):

- i) *Selection*: Actions/edges are selected according to a *tree policy*. A commonly used policy is to select greedily with respect to Upper Confidence bounds for Trees (UCT) [7]:

$$\text{UCT}(s, a) = Q(s, a) + C \sqrt{\frac{\log N(s)}{N(s, a)}}, \quad (4)$$

where $Q(s, a)$ is the average accumulated reward of taking action a ; $N(s, a)$ and $N(s)$ return the visiting times of edge a and state s , respectively; C is a scalar exploration constant. Starting from the root node, the action is selected as:

$$a^* = \underset{a \in \mathcal{A}(s)}{\text{argmax}} \text{UCT}(s, a), \quad (5)$$

where $\mathcal{A}(s)$ is the global action space at state s , which enumerates all the possible waypoints. The selection is performed recursively until an unexpanded edge is selected.

- ii) *Expansion*: With the finally selected unexpanded edge, a new leaf node \dot{s} is further appended, *i.e.*, $N(\dot{s}) = 1$.
- iii) *Rollout*: A quick rollout is performed to predict further multiple steps into the future, according to a certain *rollout policy*. Based on a reward function R , the value of the new leaf node \dot{s} , *i.e.*, $V(\dot{s})$, is obtained.
- iv) *Back-up*: The statistics of nodes and edges are updated bottom-up through the tree from \dot{s} until reaching the root node.

$$\begin{aligned} N(s, a) &= N(s'), \\ Q(s, a) &= R(s, a) + \gamma V(s'), \\ N(s) &= 1 + \sum_{a \in \mathcal{A}(s)} N(s, a), \\ V(s) &= \frac{1}{N(s)} \sum_{a \in \mathcal{A}(s)} (N(s, a) Q(s, a)). \end{aligned} \quad (6)$$

Here $R(s, a)$ gives the reward for taking action a at state s ; $s' = T(s, a)$, where $T(s, a)$ is the deterministic transition function. The design of our reward function R and the rollout policy will be detailed later.

As such, in each iteration, DREAMWALKER creates and executes a plan in the world model by expanding the search tree, and anticipates the consequence by estimating the rewards. After several times of iterative simulation, the agent

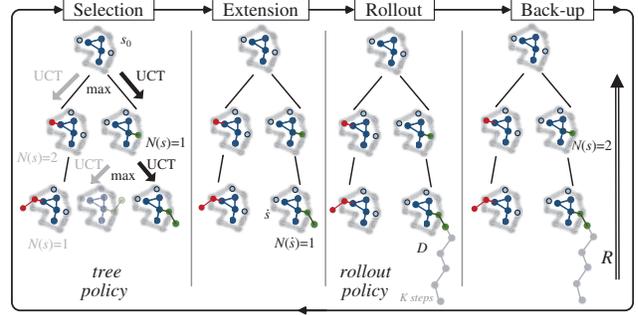


Figure 3: MCTS based Mental planning. Each node in the research tree refers to a possible world state, corresponding to a future plan.

conducts a lot of mental experiments, and the best action is selected among the edges starting from the root node:

$$a^* = \underset{a \in \mathcal{A}(s_0)}{\text{argmax}} Q(s_0, a). \quad (7)$$

Here a^* corresponds to a certain waypoint observed at current world state s_0 . According to EG G , the agent can reach this waypoint by taking a sequence of low-level actions. If a^* is never visited before, EG is updated after its first visit, followed by new-round mental planning and next-step action. When a^* is a visited waypoint, the agent will move to this waypoint and choose **stop** to terminate navigation. In this way, our agent predicts future in the discrete world model, and navigates in the continuous environment.

Reward. The immediate reward $R(s, a)$ is defined according to the change of distance to goal after taking action a . Let $\mathcal{G}_s = (\mathcal{V}_s, \mathcal{E}_s)$ denote EG corresponding to state s , where $\mathcal{V}_s - \mathcal{V}$ refers to those waypoints visited in dream world state s , the distance to the target location in state s is defined as:

$$D(s) = \min_{v \in \mathcal{V}_s} F_d(v, \mathcal{G}_s, X), \quad (8)$$

where F_d is a learnable distance function that predicts the distance from waypoint v to the target location, conditioned on EG \mathcal{G}_s and instruction X . More specifically, $D(s)$ is built upon a graph attention (GAT) network [78]:

$$\begin{aligned} F_d(v, \mathcal{G}_s, X) &= \text{MLP}(\tilde{v}), \\ \tilde{V} = \{\tilde{v}\}_v &= \text{GAT}(\{[\mathbf{v}, \mathbf{X}]\}_v, \{[e_{u,v}, \mathbf{u}, \mathbf{X}]\}_{u,v}), \end{aligned} \quad (9)$$

where \mathbf{v} (\mathbf{u}) is the initial embedding of node v (u), $e_{u,v}$ is the edge embedding of $e_{u,v} \in \mathcal{E}$. \mathbf{X} denotes the textual embedding of the instruction X , and $[\cdot, \cdot]$ refers to concatenation. After fusing node and edge embeddings with the textual context, GAT outputs the collection of improved node embeddings \tilde{V} . After that, a small multilayer perceptron (MLP) is applied per node for distance regression.

Given the deterministic transition $s' = T(s, a)$, we have $a = \mathcal{V}_{s'} - \mathcal{V}_s$. If action a is a previously visited waypoint, *i.e.*, $a \in \mathcal{V}$, the immediate reward $R(s, a)$ is given as:

$$R(s, a) = \begin{cases} +5, & F_d(v, \mathcal{G}_s, X) \leq 3, \\ -5, & F_d(v, \mathcal{G}_s, X) > 3. \end{cases} \quad (10)$$

Here $a \in \mathcal{V}$ means the agent chooses **stop** at waypoint a . We estimate if its distance to the target is within the success criteria, *i.e.*, 3m, and assign a constant positive/negative reward (+5/-5) accordingly. If waypoint a is never visited before, we define:

$$R(s, a) = D(s') - D(s). \quad (11)$$

That is to say, the reward is defined as the distance that a can bring the agent closer to the target than before.

Rollout Policy. In the rollout phase, a rollout policy is adopted to guide the fast ployout starting from the new expanded leaf state \dot{s} . This can be intuitively viewed as further imagining the future of state \dot{s} with several steps through fast sampling. For the sake of Monte Carlo property and simplicity, we treat the distance function $D(\cdot)$ (*cf.* Eq. 8) as the rollout policy. For example, at the first rollout step, the action distribution \mathbf{p} over possible waypoints $\mathcal{A}(\dot{s})$ at state \dot{s} is given as:

$$\mathbf{p}[a] = \text{softmax}_{a \in \mathcal{A}(\dot{s})} F_d(a, \mathcal{G}_{\dot{s}}, X). \quad (12)$$

The rollout stops when a previously visited waypoint is selected or a maximum rollout depth K is reached. The rewards $\{R_1, R_2, \dots, R_K\}$ along the rollout record are collected to compute the accumulated discounted reward as the value of \dot{s} :

$$V(\dot{s}) = \sum_{k=1}^K \gamma^{k-1} R_k. \quad (13)$$

4. Experiments

After stating our experimental setup (§4.1) and implementation details (§4.2), we provide performance comparison results with VLN-CE state-of-the-arts (§4.3). Then we identify there is still room for improvement in the aspects of world model (§4.4) and distance estimation (§4.6), revealing possible future directions. We further verify the contribution of our world model based mental planning (§4.5). Finally, we evaluate the impacts of core hyper-parameters (§4.7) and offer visual analyses (§4.8).

4.1. Experimental Setup

Dataset. We conduct experiments on VLN-CE dataset [45]. The dataset has 16, 844 trajectory-instruction pairs across 90 Matterport3D [8] scenes, and is divided into four sets, *i.e.*, `train` (10, 819 pairs, 61 scenes), `val seen` (778 pairs, 53 scenes), `val unseen` (1, 839 pairs, 11 scenes), and `test` (3, 408 pairs, 18 scenes). As the scenes in `val unseen` and `test` are not exposed in `train`, the performances on `val unseen` and `test` are more important than `val seen`.

Evaluation Metric. Following [45, 43, 44], we use five metrics for evaluation, *i.e.*, Navigation Error (NE), Trajectory Length (TL), Success Rate (SR), Oracle success Rate (OR), and Success rate weighted by Path Length (SPL), where SR is the priority. Please see [5, 6] for full details on metrics.

4.2. Implementation Details

Network Architecture. As in [43, 44, 45, 29], the RGB and depth channels of the panoramic observation Y are respec-

tively encoded by ImageNet [68] pre-trained ResNet-18 [27] and PointGoal [90] pre-trained ResNet-50. The language instruction X is encoded through a bi-LSTM with GLoVe [33] word embeddings. *Scene_Synthesizer* is built as a two-stage generator, following [42]. *Waypoint_Predictor* is the one used in [29]. For the distance function F_d , GAT is implemented as standard [78], and MLP has 1, 024 hidden neurons.

Network Training. The training of *Waypoint_Predictor* and *Scene_Synthesizer* are scheduled as in [29] and [42] respectively. For our distance function F_d , we progressively construct EGs along the ground-truth trajectory of each training episode. F_d is trained by minimizing the L2 loss between the predicted distance and ground-truth distance for each waypoint for each EG. More specifically, for robust distance prediction, for a ground-truth trajectory with L waypoints, we construct L EGs where the l -th EG is constructed by adding the l -th waypoint of the trajectory as well as up to 5 random sampled, accessible waypoints into the $(l-1)$ -th EG. In addition, we randomly replace ground-truth waypoints with the ones created by *Scene_Synthesizer*, making F_d better adapted to the working mode of mental planning.

Reproducibility. Our hyper-parameters are set as follows. The discount factor γ is 0.98 (Eq. 6 and 13). The horizon of mental planning, *i.e.*, the maximum step of imagining the future trajectory, is 4. The number of search iterations for each time of mental planning is 50. Our model is implemented in PyTorch with Habitat [70] simulator, and trained on one NVIDIA RTX 3090 GPU. Our code is released at <https://github.com/hanqingwangai/Dreamwalker>.

4.3. Comparison with VLN-CE State-of-the-Arts

We compare our agent, DREAMWALKER, with five previously published VLN-CE models [45, 43, 67, 29, 44]. As illustrated in Table 1, our agent outperforms all the competitors across all the splits, in terms of SR. In particular, DREAMWALKER surpasses BridgingGap [29], the current top-leading VLN-CE model, by 7%, 5%, and 7%, on `val seen`, `val unseen`, and `test` splits, respectively. Since all the involved competitors are model-free approaches, our promising results demonstrate the superiority of our world model based algorithmic design, which relieves the burden on the agent to learn the knowledge about the environments during navigation policy training. In addition, we can find that DREAMWALKER boosts SR score greatly without introducing much extra travel expense. This evidences that, compared to greedily following a sophisticated navigation policy, planning ahead in the mental world enables advanced decision-making with little expense of physical execution.

4.4. How Far We Are From a Perfect World Model?

The world model allows our DREAMWALKER to employ ‘mental imagery’ to perform mental experiments, so as to plan ahead before taking action. The previous experiments by our world model with the corresponding actual

| Model | val seen | | | | | val unseen | | | | | test | | | | |
|--------------------------------------|-------------|------|-----------|-----------|-----------|-------------|------|-----------|-----------|-----------|-------------|------|-----------|-----------|-----------|
| | NE↓ | TL | SR↑ | OR↑ | SPL↑ | NE↓ | TL | SR↑ | OR↑ | SPL↑ | NE↓ | TL | SR↑ | OR↑ | SPL↑ |
| CMA [45] _[ECCV20] | 7.21 | 9.06 | 34 | 44 | 32 | 7.60 | 8.27 | 29 | 36 | 27 | 7.91 | 8.85 | 28 | 36 | 25 |
| Waypoint [43] _[ICCV21] | 5.48 | 8.54 | 46 | 53 | 43 | 6.31 | 7.62 | 36 | 40 | 34 | 6.65 | 8.02 | 32 | 37 | 30 |
| LAW [67] _[EMNLP21] | 6.35 | 9.34 | 40 | 49 | 37 | 6.83 | 8.89 | 35 | 44 | 31 | 7.69 | 9.67 | 28 | 38 | 25 |
| BridgingGap [29] _[CVPR22] | 5.02 | 12.5 | 50 | 59 | 44 | 5.74 | 12.2 | 44 | 53 | 39 | 5.89 | 13.3 | 42 | 51 | 36 |
| Sim2Sim [44] _[ECCV22] | 4.67 | 11.2 | 52 | 61 | 44 | 6.07 | 10.7 | 43 | 52 | 36 | 6.17 | 11.4 | 44 | 52 | 37 |
| DREAMWALKER (Ours) | 4.09 | 11.6 | 59 | 66 | 48 | 5.53 | 11.3 | 49 | 59 | 44 | 5.48 | 11.8 | 49 | 57 | 44 |

Table 1: Impacts of core method components on VLN-CE dataset [45] (§4.4-§4.5).

(*cf.* §4.3 and Table 1) primarily demonstrated the power of the world model in strategic navigation planning, through the comparison with existing model-free VLN-CE agents. To help highlight how far we are from a perfect world model, we derive a variant – “perfect imagination” – from our algorithm, by replacing the future scenarios forecast-ahead observations from the environment. For completeness, we also provide a “lazy” world model – “copy memory” – which simply memorizes all past observations and generates future predictions by simply copying the nearest memory. From Table 2 we can find that, compared with DREAMWALKER, “perfect imagination” yields solid performance boost. This is because the agent coupled with a perfect world model can make perfect prediction of the future. This also suggests the upperbound of our performance w.r.t. world model, and is consistent with prior studies [84, 42] which find that allowing agent to look ahead at actual observations from environments can facilitate decision-making in the abstract setting. At the other extreme, “copy memory”, not surprisingly, gives the worst results, since the agent does not make any imagination of the future. This verifies again the benefit of world model in navigation planning.

4.5. Advanced Planning or Greedy Selection?

It is also interesting to quantify the contribution of our world model based mental planning. To this end, in Table 2, we report a baseline - “greedy selection”: at each decision-making step, the agent greedily selects the waypoint in the view of field for navigation, according to the estimated distance to the target location. This decision-making strategy is adopted by all current VLN-CE agents, yet lacking explicit planning. Our world model enables sampling-based planning — first assessing the consequences of possible navigation solutions in the mind then taking actual action. From Table 2 we can find that, mental planning exhibit significantly better performance compared with the greedy action selection strategy. This evidences our primary hypothesis that imagination of possible futures in the internal world abstract is necessary for strategic navigation planning.

4.6. What Is the Impact of Distance Function?

During mental planning, DREAMWALKER makes use of the distance function F_d (*cf.* Eq. 9) to calculate the immediate reward R (*cf.* Eq. 11) for a certain action. In other words,

| Variant | val seen | | | | | val unseen | | | | |
|---------------------|----------|------|-----|-----|------|------------|------|-----|-----|------|
| | NE↓ | TL | SR↑ | OR↑ | SPL↑ | NE↓ | TL | SR↑ | OR↑ | SPL↑ |
| DREAMWALKER (Ours) | 4.09 | 11.6 | 59 | 66 | 48 | 5.53 | 11.3 | 49 | 59 | 44 |
| Perfect Imagination | 3.75 | 10.8 | 64 | 69 | 60 | 4.88 | 11.1 | 54 | 63 | 49 |
| Copy Memory | 7.10 | 13.5 | 35 | 44 | 31 | 7.76 | 13.8 | 27 | 35 | 24 |
| Greedy Selection | 5.22 | 10.5 | 47 | 56 | 43 | 5.93 | 10.9 | 42 | 53 | 36 |

Table 2: Impacts of core method components on VLN-CE dataset [45] (§4.4-§4.5).

DREAMWALKER assesses the outcome of mental plans by means of the distance function. We therefore conduct experiments to study the influence of the distance function. Specifically, we report the performance by randomly replacing the distance estimates with the ground-truth in different probabilities. The performance are plotted in curves in Fig. 4. We can observe that, when the distance estimation becomes more accurate, DREAMWALKER performs better. With perfect distance estimation – the agent is allowed to access accurate distance between any waypoint and the target goal, 100% SR can be reached. We also find that, even when the replacement probability is as low as 20%, the agent can achieve a rather high 70 SR. It demonstrates the central role of the distance function in our algorithm and, also, suggests a feasible direction for further improvement.

4.7. Hyper-Parameter Study

In this section, we examine our hyper-parameter setup.

Search Iterations of MCTS. We first investigate the influence of searching iterations in MCTS based mental planning. Intuitively, with more searching rounds, the search tree has a higher probability to reach a good terminal state and can better estimate the outcome of possible actions, yet, at the expense of larger simulation cost. Therefore, in addition to estimating the navigation performance with different searching rounds (*i.e.*, 10, 30, 50, 70), we also report statistics for the runtime. As shown in Table 3, the runtime grows linearly as the number of searching rounds increases. However, when the number of searching rounds exceeds 50, the performance gain becomes marginal. Thus we finally set the number of searching rounds to 50 to save the unproductive computation cost. We also stress that our mental planning is very fast which typically responds within 1.5s.

Horizons of Mental Planning. We report the performance with different planning horizons — imaging future trajectories with a maximum length of 0, 2, 4, or 6 forward steps. From Table 4 we can find that, both the performance and the runtime generally grows as the horizons increased from

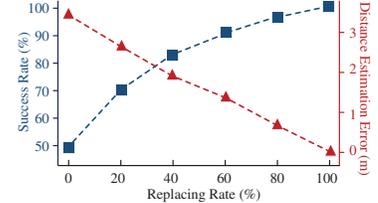


Figure 4: Curves of success rate and distance estimation error (§4.6).

| # | Searching Iteration | val seen | | | | | val unseen | | | | | Runtime (s/step) ↓ |
|---|---------------------|----------|------|------|------|-------|------------|------|------|------|-------|--------------------|
| | | NE ↓ | TL | SR ↑ | OR ↑ | SPL ↑ | NE ↓ | TL | SR ↑ | OR ↑ | SPL ↑ | |
| 1 | 10 | 4.43 | 12.1 | 55 | 63 | 45 | 5.76 | 12.5 | 44 | 55 | 38 | 0.43 |
| 2 | 30 | 4.29 | 10.9 | 57 | 65 | 47 | 5.62 | 11.1 | 46 | 57 | 42 | 1.08 |
| 3 | 50 | 4.09 | 11.6 | 59 | 66 | 48 | 5.53 | 11.3 | 49 | 59 | 44 | 1.43 |
| 4 | 70 | 4.02 | 12.6 | 59 | 67 | 48 | 5.49 | 12.9 | 50 | 60 | 44 | 1.74 |

Table 3: Impact of sampling iteration for navigation plan searching (§4.7).

| # | Planning Horizon | val seen | | | | | val unseen | | | | | Runtime (s/step) ↓ |
|---|------------------|----------|------|------|------|-------|------------|------|------|------|-------|--------------------|
| | | NE ↓ | TL | SR ↑ | OR ↑ | SPL ↑ | NE ↓ | TL | SR ↑ | OR ↑ | SPL ↑ | |
| 1 | 0 | 5.22 | 10.5 | 47 | 56 | 43 | 5.93 | 10.9 | 42 | 53 | 36 | 0.09 |
| 2 | 2 | 4.21 | 11.8 | 56 | 64 | 45 | 5.66 | 12.2 | 47 | 57 | 41 | 1.15 |
| 3 | 4 | 4.09 | 11.6 | 59 | 66 | 48 | 5.53 | 11.3 | 49 | 59 | 44 | 1.43 |
| 4 | 6 | 4.18 | 12.5 | 57 | 65 | 44 | 5.59 | 12.9 | 48 | 58 | 41 | 2.05 |

Table 4: Impact of planning horizon (§4.7). The maximum searching round is 50.

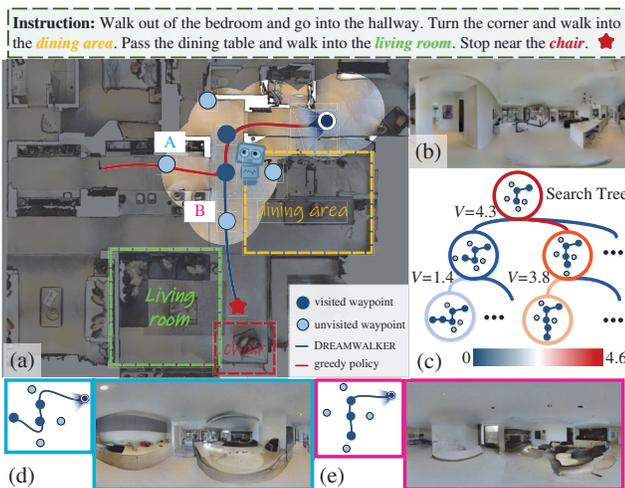


Figure 6: (a) Trajectories of an episode navigated by a greedy policy (red) and DREAMWALKER (blue). (b) Current panoramic observation. (c) The search tree rooted by the world state of (a). The nodes and edges are painted according to their $V(s)$ and $Q(s, a)$ respectively. (d) Imagined view at waypoint A. (e) Imagined view at waypoint B. See §4.8 for more details.

0 to 4 steps. Comparing #3 and #4 rows, we can find that the performance gain becomes marginal or even negative. This happens probably due to that the capacity of the world model is overloaded. As in [42], we assess the fidelity of our synthesized panoramic RGBD views by computing Fréchet Inception Distance (FID) [28] to the ground-truth scenes. As shown in Fig. 5, the error between simulated views and corresponding actual observations from real environments accumulates as the trajectory rolls out. Hence the scenarios forecasted over long horizons might be useless or even detrimental to mental planning.

4.8. Qualitative Result

To better understand the superior performance of our method and demonstrate the interpretability of the mental planning mechanism, we analyse a challenging qualitative

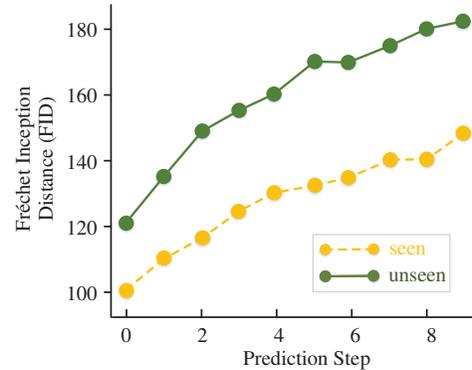


Figure 5: FID curve of synthesized panoramic view w.r.t. prediction step (§4.7).

case in Fig. 6. In this case, we visualize the navigation trajectory performed by a greedy policy (*i.e.*, greedily selecting the best waypoint predicted by the distance function) and our DREAMWALKER. As seen, given a complicated instruction “Walk out \dots to outside.”, the agent with the greedy policy soon gets lost after it enters the dining room, while our DREAMWALKER manages to reach the target location. For an intuitive comprehension of the planning procedure, here we visualize a part of the search tree. The nodes and edges in the search tree are colored according to their corresponding $V(s)$ and $Q(s, a)$ values, *i.e.*, red color indicates high value and blue color indicates low value. We can clearly observe that the search tree is split into two branches when the planning proceeds to “walk into the dining area”, *i.e.*, DREAMWALKER starts to imagine the outcomes of two possible actions, and the branch of the correct action finally wins with a large margin as the branch of the wrong action receives rather low rewards (deep blue). We visualize the imagined future waypoints of the correct action and the wrong action respectively and find that the correct action leads to a room which looks more like a “living room” mentioned in the instruction, while the wrong action leads to a corridor. This study demonstrates that our DREAMWALKER can provide strong interpretability of decision making by conducting tractable planning and presenting intuitive visualization for imagined observations.

5. Conclusion and Discussion

In this article, we devise DREAMWALKER, a world model based VLN-CE agent. Our world model is built as a discrete and structured abstraction of the continuous environment, allowing DREAMWALKER to synthesize, assess, and interpret possible plans in the mind before taking actual actions. We empirically confirm the superiority of DREAMWALKER over existing model-free agents, in terms of performance and interpretability. Building more expressive and versatile world models, and exploring world model based VLN-CE policy learning will be the focus of our future work.

Acknowledgements

This project is partially sponsored by the National Natural Science Foundation of China (NSFC) under Grant No.62172043, and CCF-Tencent Open Fund.

References

- [1] Michal Adamkiewicz, Timothy Chen, Adam Caccavale, Rachel Gardner, Preston Culbertson, Jeannette Bohg, and Mac Schwager. Vision-only robot navigation in a neural radiance world. *IEEE Robotics and Automation Letters*, 2022. 3
- [2] Dong An, Yuankai Qi, Yan Huang, Qi Wu, Liang Wang, and Tieniu Tan. Neighbor-view enhanced model for vision and language navigation. In *ACMMM*, 2021. 2
- [3] Dong An, Yuankai Qi, Yangguang Li, Yan Huang, Liang Wang, Tieniu Tan, and Jing Shao. Bevbort: Topo-metric map pre-training for language-guided navigation. *arXiv preprint arXiv:2212.04385*, 2022. 2
- [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
- [5] Peter Anderson, Angel Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, et al. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*, 2018. 6
- [6] 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*, 2018. 1, 2, 6
- [7] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 2002. 5
- [8] 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*, 2017. 2, 6
- [9] Jianyu Chen, Shengbo Eben Li, and Masayoshi Tomizuka. Interpretable end-to-end urban autonomous driving with latent deep reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 2021. 2
- [10] Shizhe Chen, Pierre-Louis Guhur, Cordelia Schmid, and Ivan Laptev. History aware multimodal transformer for vision-and-language navigation. In *NeurIPS*, 2021. 2
- [11] 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*, 2022. 2
- [12] Silvia Chiappa, Sébastien Racaniere, Daan Wierstra, and Shakir Mohamed. Recurrent environment simulators. In *ICLR*, 2017. 3
- [13] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *NeurIPS*, 2018. 3
- [14] Antonia Creswell, Rishabh Kabra, Chris Burgess, and Murray Shanahan. Unsupervised object-based transition models for 3d partially observable environments. In *NeurIPS*, 2021. 3
- [15] Zhiwei Deng, Karthik Narasimhan, and Olga Russakovsky. Evolving graphical planner: Contextual global planning for vision-and-language navigation. In *NeurIPS*, 2020. 2
- [16] Chelsea Finn and Sergey Levine. Deep visual foresight for planning robot motion. In *ICRA*, 2017. 3
- [17] Jay W Forrester. Counterintuitive behavior of social systems. *Theory and Decision*, 2(2):109–140, 1971. 1
- [18] 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
- [19] Karol Gregor, Danilo Jimenez Rezende, Frederic Besse, Yan Wu, Hamza Merzic, and Aaron van den Oord. Shaping belief states with generative environment models for rl. In *NeurIPS*, 2019. 3
- [20] Karol Gregor, George Papamakarios, Frederic Besse, Lars Buesing, and Theophane Weber. Temporal difference variational auto-encoder. In *ICLR*, 2018. 3
- [21] Xiaoxiao Guo, Satinder Singh, Honglak Lee, Richard L Lewis, and Xiaoshi Wang. Deep learning for real-time atari game play using offline monte-carlo tree search planning. In *NeurIPS*, 2014. 3, 4
- [22] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In *NeurIPS*, 2018. 3
- [23] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. In *ICLR*, 2020. 2, 3
- [24] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In *ICML*, 2019. 3
- [25] Danijar Hafner, Timothy P Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models. In *ICLR*, 2020. 3
- [26] Weituo Hao, Chunyuan Li, Xiujuan Li, Lawrence Carin, and Jianfeng Gao. Towards learning a generic agent for vision-and-language navigation via pre-training. In *CVPR*, 2020. 2
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [28] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017. 8
- [29] 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. 2, 3, 6, 7
- [30] Yicong Hong, Qi Wu, Yuankai Qi, Cristian Rodriguez-Opazo, and Stephen Gould. A recurrent vision-and-language bert for navigation. In *CVPR*, 2021. 2
- [31] Haoshuo Huang, Vihan Jain, Harsh Mehta, Alexander Ku, Gabriel Magalhaes, Jason Baldrige, and Eugene Ie. Transferable representation learning in vision-and-language navigation.

- gation. In *ICCV*, 2019. 2
- [32] Vihan Jain, Gabriel Magalhaes, Alexander Ku, Ashish Vaswani, Eugene Ie, and Jason Baldridge. Stay on the path: Instruction fidelity in vision-and-language navigation. In *ACL*, 2019. 2
- [33] Richard Socher, Jeffrey Pennington, and Christopher D Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014. 6
- [34] Philip Nicholas Johnson-Laird. *Mental models: Towards a cognitive science of language, inference, and consciousness*. Harvard University Press, 1983. 1, 3
- [35] Philip N Johnson-Laird. Mental models and human reasoning. *Proceedings of the National Academy of Sciences*, 2010. 1
- [36] Łukasz Kaiser, Mohammad Babaeizadeh, Piotr Miłoś, Błażej Osiniński, Roy H Campbell, Konrad Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey Levine, et al. Model based reinforcement learning for atari. In *ICLR*, 2019. 3
- [37] Maximilian Karl, Maximilian Soelch, Justin Bayer, and Patrick Van der Smagt. Deep variational bayes filters: Un-supervised learning of state space models from raw data. In *ICLR*, 2017. 3
- [38] Liyiming Ke, Xiujuan 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 *CVPR*, 2019. 2
- [39] Nan Rosemary Ke, Amanpreet Singh, Ahmed Touati, Anirudh Goyal, Yoshua Bengio, Devi Parikh, and Dhruv Batra. Learning dynamics model in reinforcement learning by incorporating the long term future. In *ICLR*, 2019. 3
- [40] Kuno Kim, Megumi Sano, Julian De Freitas, Nick Haber, and Daniel Yamins. Active world model learning with progress curiosity. In *ICML*, 2020. 3
- [41] Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In *ECML*, 2006. 2, 3, 4
- [42] Jing Yu Koh, Honglak Lee, Yinfei Yang, Jason Baldridge, and Peter Anderson. Pathdreamer: A world model for indoor navigation. In *CVPR*, 2021. 3, 4, 6, 7, 8
- [43] Jacob Krantz, Aaron Gokaslan, Dhruv Batra, Stefan Lee, and Oleksandr Maksymets. Waypoint models for instruction-guided navigation in continuous environments. In *ICCV*, 2021. 2, 3, 6, 7
- [44] Jacob Krantz and Stefan Lee. Sim-2-sim transfer for vision-and-language navigation in continuous environments. In *ECCV*, 2022. 2, 3, 6, 7
- [45] 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, 7
- [46] Alexander Ku, Peter Anderson, Roma Patel, Eugene Ie, and Jason Baldridge. Room-Across-Room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In *EMNLP*, 2020. 2
- [47] Alex X Lee, Anusha Nagabandi, Pieter Abbeel, and Sergey Levine. Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model. In *NeurIPS*, 2020. 3
- [48] Kuang-Huei Lee, Ian Fischer, Anthony Liu, Yijie Guo, Honglak Lee, John Canny, and Sergio Guadarrama. Predictive information accelerates learning in rl. In *NeurIPS*, 2020. 3
- [49] Juncheng Li, Siliang Tang, Fei Wu, and Yueting Zhuang. Walking with mind: Mental imagery enhanced embodied qa. In *ACMMM*, 2019. 3
- [50] Zhixuan Lin, Yi-Fu Wu, Skand Peri, Bofeng Fu, Jindong Jiang, and Sungjin Ahn. Improving generative imagination in object-centric world models. In *ICML*, 2020. 3
- [51] Rui Liu, Xiaohan Wang, Wenguan Wang, and Yi Yang. Bird’s-eye-view scene graph for vision-language navigation. In *ICCV*, 2023. 2
- [52] Chih-Yao Ma, Jiasen Lu, Zuxuan Wu, Ghassan AlRegib, Zolt Kira, Richard Socher, and Caiming Xiong. Self-monitoring navigation agent via auxiliary progress estimation. In *ICLR*, 2019. 2
- [53] Chih-Yao Ma, Zuxuan Wu, Ghassan AlRegib, Caiming Xiong, and Zolt Kira. The regretful agent: Heuristic-aided navigation through progress estimation. In *CVPR*, 2019. 2
- [54] Arjun Majumdar, Ayush Shrivastava, Stefan Lee, Peter Anderson, Devi Parikh, and Dhruv Batra. Improving vision-and-language navigation with image-text pairs from the web. In *ECCV*, 2020. 2
- [55] Arun Mallya, Ting-Chun Wang, Karan Sapra, and Ming-Yu Liu. World-consistent video-to-video synthesis. In *ECCV*, 2020. 4
- [56] Anusha Nagabandi, Gregory Kahn, Ronald S Fearing, and Sergey Levine. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *ICRA*, 2018. 3
- [57] Ashvin V Nair, Vitchyr Pong, Murtaza Dalal, Shikhar Bahl, Steven Lin, and Sergey Levine. Visual reinforcement learning with imagined goals. In *NeurIPS*, 2018. 3
- [58] Suraj Nair, Silvio Savarese, and Chelsea Finn. Goal-aware prediction: Learning to model what matters. In *ICML*, 2020. 2, 3
- [59] Junhyuk Oh, Xiaoxiao Guo, Honglak Lee, Richard L Lewis, and Satinder Singh. Action-conditional video prediction using deep networks in atari games. In *NeurIPS*, 2015. 3
- [60] Alexander Pashevich, Cordelia Schmid, and Chen Sun. Episodic transformer for vision-and-language navigation. In *ICCV*, 2021. 2
- [61] Alexandre Piché, Valentin Thomas, Cyril Ibrahim, Yoshua Bengio, and Chris Pal. Probabilistic planning with sequential monte carlo methods. In *ICLR*, 2018. 3
- [62] AJ Piergiovanni, Alan Wu, and Michael S Ryoo. Learning real-world robot policies by dreaming. In *IROS*, 2019. 2
- [63] 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*, 2020. 2
- [64] Yanyuan Qiao, Yuankai Qi, Yicong Hong, Zheng Yu, Peng Wang, and Qi Wu. Hop: history-and-order aware pre-training for vision-and-language navigation. In *CVPR*, 2022. 2
- [65] Yanyuan Qiao, Yuankai Qi, Yicong Hong, Zheng Yu, Peng Wang, and Qi Wu. Hop+: History-enhanced and order-aware pre-training for vision-and-language navigation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 2
- [66] Sébastien Racanière, Théophane Weber, David Reichert,

- Lars Buesing, Arthur Guez, Danilo Jimenez Rezende, Adrià Puigdomènech Badia, Oriol Vinyals, Nicolas Heess, Yujia Li, et al. Imagination-augmented agents for deep reinforcement learning. In *NeurIPS*, 2017. 3
- [67] Sonia Raychaudhuri, Saim Wani, Shivansh Patel, Unnat Jain, and Angel Chang. Language-aligned waypoint (law) supervision for vision-and-language navigation in continuous environments. In *EMNLP*, 2021. 2, 3, 6, 7
- [68] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115(3):211–252, 2015. 6
- [69] Oleh Rybkin, Karl Pertsch, Konstantinos G Derpanis, Kostas Daniilidis, and Andrew Jaegle. Learning what you can do before doing anything. In *ICLR*, 2018. 2, 3
- [70] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, and J. Malik. Habitat: A platform for embodied ai research. In *ICCV*, 2019. 6
- [71] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020. 3
- [72] Ramanan Sekar, Oleh Rybkin, Kostas Daniilidis, Pieter Abbeel, Danijar Hafner, and Deepak Pathak. Planning to explore via self-supervised world models. In *ICML*, 2020. 3
- [73] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 2016. 2, 3
- [74] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 2017. 3
- [75] Aravind Srinivas, Allan Jabri, Pieter Abbeel, Sergey Levine, and Chelsea Finn. Universal planning networks: Learning generalizable representations for visuomotor control. In *ICML*, 2018. 3
- [76] Adam Stooke, Kimin Lee, Pieter Abbeel, and Michael Laskin. Decoupling representation learning from reinforcement learning. In *ICML*, 2021. 3
- [77] Hao Tan, Licheng Yu, and Mohit Bansal. Learning to navigate unseen environments: Back translation with environmental dropout. In *NAACL*, 2019. 2
- [78] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. In *ICLR*, 2018. 5, 6
- [79] Huanjie Wang, Hongbo Gao, Shihua Yuan, Hongfei Zhao, Kelong Wang, Xiulai Wang, Keqiang Li, and Deyi Li. Interpretable decision-making for autonomous vehicles at highway on-ramps with latent space reinforcement learning. *IEEE Transactions on Vehicular Technology*, 2021. 2
- [80] Hanqing Wang, Wei Liang, Luc V Gool, and Wenguan Wang. Towards versatile embodied navigation. In *NeurIPS*, 2022. 2
- [81] Hanqing Wang, Wei Liang, Jianbing Shen, Luc Van Gool, and Wenguan Wang. Counterfactual cycle-consistent learning for instruction following and generation in vision-language navigation. In *CVPR*, 2022. 2
- [82] Hanqing Wang, Wenguan Wang, Wei Liang, Steven CH Hoi, Jianbing Shen, and Luc Van Gool. Active perception for visual-language navigation. *International Journal of Computer Vision*, 2023. 2
- [83] Hanqing Wang, Wenguan Wang, Wei Liang, Caiming Xiong, and Jianbing Shen. Structured scene memory for vision-language navigation. In *CVPR*, 2021. 2
- [84] Hanqing Wang, Wenguan Wang, Tianmin Shu, Wei Liang, and Jianbing Shen. Active visual information gathering for vision-language navigation. In *ECCV*, 2020. 2, 7
- [85] Tingwu Wang and Jimmy Ba. Exploring model-based planning with policy networks. In *ICLR*, 2020. 3
- [86] 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*, 2019. 2
- [87] Xiaohan Wang, Wenguan Wang, Jiayi Shao, and Yi Yang. Lana: A language-capable navigator for instruction following and generation. In *CVPR*, 2023. 2
- [88] Xin Wang, Wenhan Xiong, Hongmin Wang, and William Yang Wang. Look before you leap: Bridging model-free and model-based reinforcement learning for planned-ahead vision-and-language navigation. In *ECCV*, 2018. 2
- [89] Manuel Watter, Jost Springenberg, Joschka Boedecker, and Martin Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. In *NeurIPS*, 2015. 3
- [90] Erik Wijmans, Abhishek Kadian, Ari Morcos, Stefan Lee, Irfan Essa, Devi Parikh, Manolis Savva, and Dhruv Batra. Dd-ppo: Learning near-perfect pointgoal navigators from 2.5 billion frames. In *ICLR*, 2019. 6
- [91] Lin Yen-Chen, Maria Bauza, and Phillip Isola. Experience-embedded visual foresight. In *CoRL*, 2020. 2, 3
- [92] Yusheng Zhao, Jinyu Chen, Chen Gao, Wenguan Wang, Lirong Yang, Haibing Ren, Huaxia Xia, and Si Liu. Target-driven structured transformer planner for vision-language navigation. In *ACMMM*, 2022. 2
- [93] Fengda Zhu, Yi Zhu, Xiaojun Chang, and Xiaodan Liang. Vision-language navigation with self-supervised auxiliary reasoning tasks. In *CVPR*, 2020. 2
- [94] Wanrong Zhu, Yuankai Qi, Pradyumna Narayana, Kazoo Sone, Sugato Basu, Xin Eric Wang, Qi Wu, Miguel Eckstein, and William Yang Wang. Diagnosing vision-and-language navigation: What really matters. In *NAACL*, 2022. 2