

Topological Planning with Transformers for Vision-and-Language Navigation

Kevin Chen¹ Junshen K. Chen¹ Jo Chuang¹ Marynel Vázquez² Silvio Savarese¹

¹Stanford University

²Yale University

Abstract

Conventional approaches to vision-and-language navigation (VLN) are trained end-to-end but struggle to perform well in freely traversable environments. Inspired by the robotics community, we propose a modular approach to VLN using topological maps. Given a natural language instruction and topological map, our approach leverages attention mechanisms to predict a navigation plan in the map. The plan is then executed with low-level actions (e.g. FORWARD, ROTATE) using a robust controller. Experiments show that our method outperforms previous end-to-end approaches, generates interpretable navigation plans, and exhibits intelligent behaviors such as backtracking.

1. Introduction

Enabling robots to understand natural language and carry out communicated tasks has long been desired. A critical step towards this goal in the context of mobile robotics is vision-and-language navigation (VLN) [4]. In VLN, the agent is provided with a navigation instruction such as: “Exit the bedroom, walk to the end of the hall, and enter the kitchen on your right.” The agent is then expected to follow the instruction and move to the specified destination.

The majority of VLN systems [16, 20, 24, 29, 30, 43, 48, 55] are end-to-end deep learning models and utilize unstructured memory such as LSTM [22]. These methods work well when the movement is constrained to pre-defined locations, but performance drops significantly when the agent is allowed to move freely [26]. Moreover, learning to perform navigation, including mapping, planning, and control, in a fully end-to-end manner can be difficult and expensive. Such approaches often require millions of frames of experience [9, 47, 50], and yet performance substantially degrades without ground truth odometry [9, 50].

To address the aforementioned issues, recent visual robot navigation literature has explored using structured memory (e.g. metric maps, topological memory) and using a modular approach, where the algorithm is explicitly divided into

relevant subcomponents such as mapping, planning, and control [6, 7, 8, 15, 17, 23, 33, 39]. These approaches have been demonstrated to work well for tasks such as target image navigation and environment exploration. However, they have not been well studied in the context of VLN.

In this work, we employ a modular approach and leverage topological maps for VLN. Topological maps, inspired in part by cognitive science, typically represent environments as graphs where nodes correspond to places and edges denote environment connectivity or reachability. Compared with metric maps, topological maps eliminate the need for meticulous map construction. They promote efficient planning, interpretable navigation plans, and navigation robustness using cheaper sensors [33, 44]. In particular, the symbolic nature of topological maps lends them suitable for navigation with language [32, 44], as the space discretization provided by the maps can facilitate learning a relationship between instructions and spatial locations.

Importantly, using topological maps synergizes well with sequence prediction models. Predicting a navigation plan in a topological map bears many similarities with predicting sequences for language tasks such as language modeling and neural machine translation (NMT). By drawing the parallel between navigation planning and language sequence prediction, we can leverage powerful attention mechanisms [46] that have enabled significant breakthroughs in language tasks for the navigation problem. Recently, these attention-based models have even been demonstrated to achieve comparable performance to convolutional neural networks on image recognition tasks [14].

We propose using a cross-modal attention-based transformer to compute navigation plans in topological maps based on language instructions. Whereas a language model predicts a word at a time, our transformer predicts one topological map node in the navigation plan at a time. Structuring our model in this manner allows the agent to attend to relevant portions of the navigation instruction and relevant spatial regions of the navigation environment during the navigation planning process. For example, this enables our model to relate the word “bedroom” in a navigation instruction to the physical room in the environment.

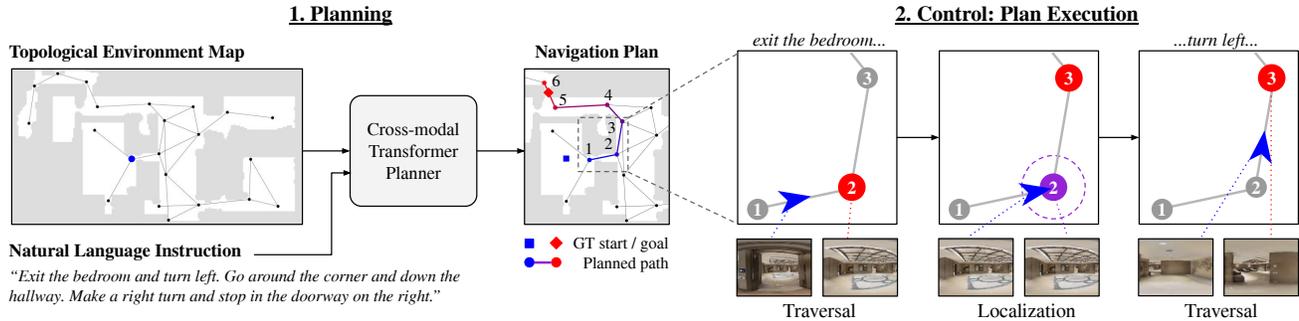


Figure 1: The agent uses the natural language instruction to generate a navigation plan in the topological map (left). A controller then executes the predicted plan by sequentially traversing to each subgoal node in the plan (right). As the agent approaches the subgoal node, it consumes that active node, and the subsequent node in the plan becomes the new active node.

Altogether, we propose a full navigation system for VLN. Unlike much of the prior work in VLN, we use a more challenging setup, allowing the agent to freely traverse the environment using low-level discrete actions. To this end, we define a topological map representation that can be constructed by the agent after it has freely explored the environment. The maps are used as part of a modular navigation framework which decomposes the problem into planning and control (Fig. 1). For each navigation episode, our agent first uses the cross-modal transformer to compute a global navigation plan from the navigation instruction and topological map. This navigation plan is executed by a robust local controller that outputs low-level discrete actions.

We evaluate our approach using the VLN-CE dataset [26]. Our experiments show that cross-modal attention-based planning is effective, and that our modular approach enables learning a robust controller capable of correcting for navigation mistakes like moving in the wrong direction.

2. Related Work

Language modeling and pretraining. Recent language model work has benefited from attention-based transformer architectures [46] and generalized pre-training of language models for fine-tuning on specific tasks [12]. These advances tie into the increasing interest in more effective methods for image recognition [13] and language grounding with visual cues [28, 41, 42]. In addition to capturing semantics that can only be visually demonstrated, these models allow for a more diverse range of applications than language-only approaches including VLN.

Vision-and-language navigation (VLN). Much of the progress in VLN has been achieved using end-to-end models [11, 16, 20, 24, 29, 30, 43, 49, 48, 55] trained using the Matterport3D Simulator [4] in which the agent assumes a panoramic action space [16] that allows it teleport to and from a fixed set of pre-defined locations in the environment. This setup promotes fast iteration and evaluation of different VLN approaches. However, it ignores the problem of

associating action directions in the panorama with the pre-defined positions in the environment as well as the motion feasibility problem of moving from one location to another. As pointed out with the end-to-end approach by Krantz *et al.* [26], the VLN problem becomes significantly harder when the agent is allowed to freely traverse the environment. In contrast, in our work the agent builds a topological map of the environment from ground up using exploration trajectories. This is done via pre-exploration similar to prior work [48, 55]. At test time, the agent predicts a navigation path in the map and uses the local controller to execute the predicted path with low-level discrete actions, bringing the task of VLN one step closer to reality.

Although they operate in the simplified action space, Hao *et al.* [20] and Majumdar *et al.* [31] propose attention-based transformer approaches that perform self-supervised pre-training on a cross-modal network using visual language pairings, either from the web or from demonstrations of a typical VLN trajectory. By fine-tuning these pre-trained models on specific VLN tasks, they achieve better generalized performance across different variants of navigation tasks. Their approaches operate on top of an existing navigation graph and perform cross-modal attention between image regions and language instead of spatial regions and language. Additionally, Majumdar *et al.* propose a trajectory scoring mechanism and relies on other methods to generate candidate routes. On the other hand, our approach predicts routes and executes them with low-level actions.

Memory for visual navigation. Early learning-based approaches to visual navigation were reactive [14, 56] or based on unstructured memory [34, 35]. Later works investigate explicit representations of environments, such as metric map-based representations [6, 17, 18, 19, 37, 54]. For example, [2] use metric maps for VLN, and [15] use transformers for explorative tasks by storing observations and poses as memory. Topological maps have also been demonstrated across different navigation problems [7, 8, 33, 39]. For language navigation, [32] and [53] use topological maps

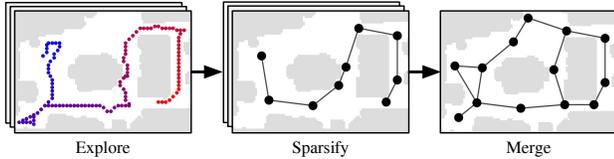


Figure 2: **Topological map construction.** The agent explores the environment multiple times, sparsifies the generated graph from each trajectory, and merges them together.

but do not use RGBD cameras as in our work.

3. Method

Problem setup. Our setup allows the agent to explore the environments using *exploration trajectories* (Fig. 2) prior to executing the VLN tasks. This mimics settings closer to the real world in which the agent has to build an understanding of the indoor environment on its own rather than being handed a pre-defined map. After the exploration phase, the agent is expected to perform VLN and is provided with (1) the instruction text and (2) the current RGBD panorama observation. Our agent has orientation information (or heading) but not positional information during navigation time.

Overview of our approach. From the exploration phase, the agent builds a topological map of the environment (Sec. 3.1). The agent then uses this topological understanding of the environment along with the instruction text and current observation to execute the VLN task.

For executing the VLN task, we take inspiration from the robotics community and use a modular approach. Specifically, we separate our approach into planning and control. Using a modular approach has been proven to work well [36, 45] and has numerous advantages including interpretability, robustness, and flexibility [23].

In the planning stage, prior to taking any action steps in the environment, the agent first uses the navigation instruction to generate an *interpretable* global navigation plan through the environment (Sec. 3.2). The robot then tries to follow the navigation plan using a repeating cycle of localization and control until it reaches the destination (Sec. 3.3). We use a hierarchical controller that first predicts waypoint subgoals. Then the waypoints are translated into low-level actions like FORWARD ($0.25m$) or ROTATE (15°).

3.1. Topological Map Representation

Before planning can take place, the agent must first construct a topological representation of the environment. For VLN, it would be desirable that the map 1) covers a large proportion of traversable area; 2) contains rich information in nodes and edges to enable localization as well as planning from language; and 3) consists of connected nodes that are reachable with high probability when executing the plan.

Similar to prior work [7, 33, 39], we represent the environment as a graph in which each node is an observation and each edge represents the connectivity (or reachability) between two nodes. As the agent explores the environment, it places a node over its current location and connects this node to the previous one.

In prior work [33, 39] the agent may end up placing multiple nodes in the same location. For example, two observations from the same position can look very different if they are oriented 180° apart. This may result in dense topological maps and exacerbate difficulties in localization, planning, and control as each node only contains partial information about each location. To reduce redundancy and mitigate these issues, we represent the nodes as 360° panoramas oriented in a consistent fixed global orientation.

An overview of the topological map construction is presented in Fig. 2. We run multiple pre-defined exploration trajectories per environment. The graph from each exploration trajectory is sparsified by a reachability estimator as proposed by Meng *et al.* [33] and then merged into a single graph that can be used for localization and planning. Further details on the exploration trajectories and map construction can be found in the supplementary material.

We append coarse geometric information to each directed edge in the map. To this end, we introduce a quantized polar coordinate system with 8 directions and 3 distances ($0-2m$, $2-5m$, $> 5m$). Such information is useful for resolving ambiguities (e.g. symmetric environments) and may facilitate language/spatial reasoning. Since mapping is not our main focus, we use ground truth odometry to compute the edge categories during map construction. However, this assumption is relaxed during navigation test time and we instead use a neural network for localization (Sec. 3.3).

3.2. Cross-Modal Planning

As stated earlier, we modularize our approach into planning and control. For planning, the agent uses the constructed topological map and the navigation instruction to formulate a global navigation plan as a path in the map. This path, represented as a sequence of nodes, is then passed to the controller at the end of the planning stage.

As depicted in Fig. 3, our planner has two main components: a graph neural network (GNN) and a cross-modal transformer (comprised of a map, language, and cross-modal encoder). The GNN computes representations of the environment which capture visual appearance and environment connectivity. These representations are passed along with the navigation instruction to the cross-modal transformer which selects the next node in the plan in a sequential, auto-regressive fashion. This process repeats until the planner classifies that the end of the plan has been reached.

Relationship to language modeling and translation. Notably, our problem setup and approach bear similari-

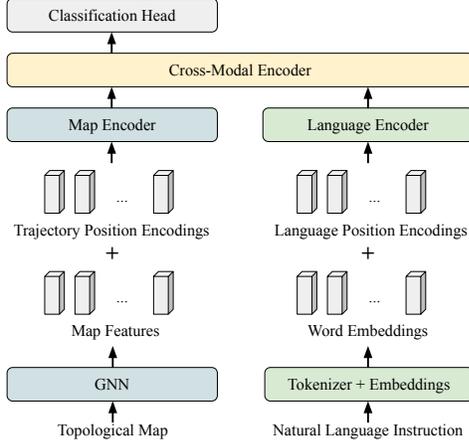


Figure 3: **Planner.** The planner processes the topological map and language instruction separately. The information is fused with a cross-modal transformer (map, language, and cross-modal encoder) to classify the next step in the plan.

ties to language modeling and translation. In neural machine translation (NMT) specifically, the model conditions on the source sentence x in addition to its prior predictions y_1, \dots, y_{t-1} in the predicted sentence.

$$p(y|x) = \prod_{t=1}^T p(y_t|y_1, \dots, y_{t-1}, x) \quad (1)$$

Analogously, in our setup the agent predicts each step of the navigation plan y_t conditioned on its previous predictions y_1, \dots, y_{t-1} as well as the topological map G and the navigation instruction L (i.e., $x = (G, L)$). This important insight allows us to approach the VLN planning problem in a manner similar to NMT and language modeling.

The following subsections illustrate the architectural setup of our approach (Fig. 3). The planner has two branches for encoding the map and the language instruction. In the map branch, we use a graph neural network (GNN) to process the topological map and generate map features (Sec. 3.2.1). In the language branch, the instruction is mapped to word embeddings. Finally, the map features and word embeddings are passed into the cross-modal transformer to produce a navigation plan (Sec. 3.2.2).

3.2.1 Learning Environment Representations using Graph Neural Networks (GNNs)

To facilitate learning a mapping between language and physical space, the map features passed to the transformer should encapsulate visual appearance and environment connectivity. In our work, these map features are learned via a graph neural network (GNN), which carries strong relational inductive biases appropriate for topological maps.

The GNN is composed of sequential graph network (GN) blocks [5]. Each GN block takes as input a graph $\tilde{G} = (\tilde{\mathbf{u}}, \tilde{V}, \tilde{E})$ and produces an updated graph $\tilde{G}' = (\tilde{\mathbf{u}}', \tilde{V}', \tilde{E}')$ where the output features are computed according to the graph structure. To do this, the GN block is comprised of update functions $\phi^v(\cdot)$, $\phi^e(\cdot)$, $\phi^u(\cdot)$ and aggregation functions $\rho^{e \rightarrow v}(\cdot)$, $\rho^{e \rightarrow u}(\cdot)$, $\rho^{v \rightarrow u}(\cdot)$. We implement the update functions as multi-layer perceptrons (MLP) and use summation for the aggregation functions. For more details on GNNs, we refer the reader to Battaglia *et al.* [5].

Input graph representation. The input to the GNN is the topological map encoded as a directed graph $G = (\mathbf{u}, V, E)$. To capture visual appearance and semantics, we encode each vertex as a ResNet152 feature [21] extracted from the corresponding fixed orientation RGB panorama. To encode the relative geometry between nodes, we map each edge category (Sec. 3.1) to a learnable embedding. Each edge thus captures information about the relative orientation and distance between a pair of connecting nodes. Lastly, the global feature \mathbf{u} is also a learned embedding.

Output graph representation. At the final layer of the GNN, we extract the output node features $\hat{V} = \{\hat{v}_i\}_{i=1:n}$ as the environment map features. Due to the message passing nature of the GNN, this set of map features not only captures visual appearance but also environment connectivity. The features are passed along with the tokenized instruction to the cross-modal transformer described in the next section.

3.2.2 Cross-Modal Transformer

At the core of our planner is a cross-modal transformer which takes as input the topological map encoded as map features (modality 1; Sec. 3.2.1) and the navigation instruction (modality 2). The start node is also provided. Each forward pass computes one step of the navigation plan. Our transformer is based on LXMERT [42] which has been demonstrated to be effective for vision and language tasks.

Attention. A fundamental component of this model is the self-attention layer [46]. Specifically, an input sequence $x \in \mathbb{R}^{n \times d}$ of length n and dimension d is linearly projected into a set of keys $K \in \mathbb{R}^{n \times d_k}$, queries $Q \in \mathbb{R}^{n \times d_k}$, and values $V \in \mathbb{R}^{n \times d_v}$. The queries and keys are then used to compute a weighted sum over the values.

$$\text{Attn}(Q, K, V) = \text{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right)V \quad (2)$$

Single modal encoder. Each modality (map and language) is encoded via its corresponding single-modality encoder. This encoder follows a standard transformer model comprised of self-attention layers and feedforward layers. As depicted in Fig. 3, the input to the map encoder is the sum of the map features (output from the GNN) with trajectory positional encodings. The language branch pro-

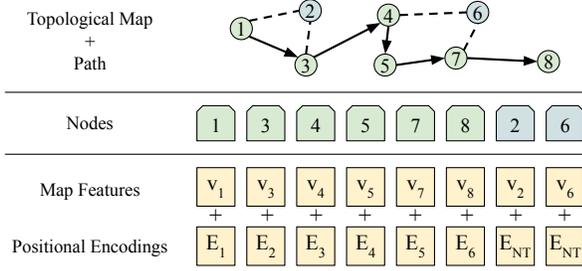


Figure 4: **Trajectory position encoding.** Nodes in the predicted trajectory (green) utilize a positional encoding for the corresponding position, whereas non-trajectory (NT) nodes (blue) utilize the same positional embedding E_{NT} .

cesses word embeddings from the natural language instruction summed with learned positional encodings.

Cross modal encoder. The cross-modal encoder [42] exchanges information between the two modalities. Each layer of the cross-modal encoder is composed of cross-modal attention, self-attention, and feed forward layers. In the cross-modal attention layer, the query comes from one modality, and the keys and values come from the other modality. The calculation for two input sequences x_A, x_B corresponding with modalities A and B is in Eq. 3.

$$\text{CrossAttn}(x_A, x_B) = \text{Attn}(Q_A, K_B, V_B) \quad (3)$$

The node features produced by the cross-modal encoder are passed to a classification head, implemented as an MLP, to compute the next step in the plan.

Trajectory position encoding. Since the transformer predicts a single node token at a time, it is essential that it keeps track of what has been predicted so far. To do this, we add positional embeddings to the node sequence as illustrated in Fig. 4. While sinusoidal positional encodings are commonly used [46], we found that learned positional embeddings yielded higher planning accuracy.

Stop action. We use a [STOP] token to indicate the stop action. It is appended to the map features and classification is performed over the nodes and the [STOP] token.

Training. The planner (GNN and transformer) is trained end-to-end with cross entropy loss. We use the AdamW optimizer [25, 27] with a linear warmup schedule. Further architecture and implementation details are described in the supplementary material.

3.3. Controller: Executing the Navigation Plan

The controller is responsible for converting the plan (a topological path) into a series of low-level actions that take the agent to the goal. The inputs of the controller are the RGBD panorama of each planned node in the sequence, the current RGBD panoramic observation at each time step, and heading information. The output action space of the con-

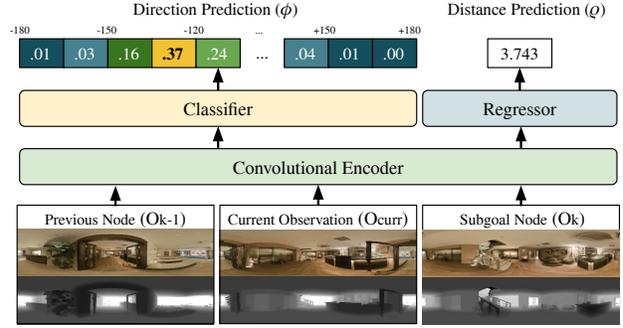


Figure 5: **High level controller.** C_{high} uses the agent’s observation and the previous and current subgoal node to predict a waypoint relative to the agent in polar coordinates.

troller is defined as a set of parameterized low-level actions: FORWARD (0.25m), ROTATE LEFT/RIGHT (15°).

The controller produces actions which move the agent from one node to another. To do this, it must also perform localization in order to determine when the current subgoal node has been reached, at which point the agent can move on to the next node in the plan. We abstract our controller C to two layers: a low level controller C_{low} that moves the agent towards a geometric waypoint, and a high level controller C_{high} that predicts such a waypoint. This high-level controller is used for localization.

High level controller. The high-level controller uses three panoramas, including the current and subgoal observations, to predict a waypoint to move towards in order to reach the subgoal. We use a quantized polar coordinate system to describe the waypoint relative to the agent. The polar coordinate is partitioned to 24 sectors, encompassing 15 degrees each. A sector index and a numeric distance defines the waypoint position. The predicted distance is used for determining whether the subgoal has been reached (localization).

Fig. 5 illustrates the design of the high level controller C_{high} . The partially observable agent state at any time step is represented by a tuple of RGBD panoramas (o_{k-1}, o_{curr}, o_k) where k is the index of the current subgoal. The observations correspond with the previous subgoal node, the current position, and the current subgoal node, respectively. All observations are oriented in the same heading direction using the heading information. To optimize for panoramic observations, we use circular convolution layers [40] to encode the observations into features. The features are then used to predict a classification vector ϕ (direction of the subgoal node) and a scalar ρ (distance between the agent and the subgoal node).

Both outputs are optimized simultaneously by imitation learning using dataset aggregation (Dagger) [38], supervised by A* expert waypoints of a fixed lookahead distance. The controller is trained separately from the planner.

Handling Oscillation. Because the agent’s state is defined by three panoramic observations only, there exists the possibility of perceptual aliasing in which the agent’s observation is similar in multiple directions (e.g. a symmetric-looking hallway with an anchor node at each end). In this case, it is possible for the agent to move towards either direction in alternating time steps, resulting in no progress.

To alleviate this issue, we introduce “stubbornness” to the predicted heading probabilities such that the predicted direction is biased towards the prediction of the previous time step. We define a bias function $B(\phi'_{t-1}, \phi_t; \sigma^2)$ to reweight the polar prediction at time t by multiplying each element by a Gaussian distribution centered at the predicted direction at the previous time step:

$$\phi'_{t_i} = \phi_{t_i} * N(i - \arg \max_j (\phi'_{t-1_j}), \sigma^2) \quad (4)$$

where i, j are indices into ϕ , and σ^2 is the variance of the Gaussian such that the lower the variance, the more biased the waypoint is towards the previous predicted direction.

Mapping to low-level actions. After C_{high} produces a waypoint, a robot-specific low-level controller translates it to a robot action. For our problem setup, we define C_{low} as a simple policy that maps ϕ'_t to quantized low-level actions: FORWARD if agent is facing the direction of $\arg \max_i \phi'_{t_i}$; otherwise ROTATE towards that direction. This can be replaced by more sophisticated robot-specific controllers.

Localization and trajectory following. When the agent traverses within a threshold distance d to the current subgoal node such that $\rho \leq d$, it consumes that node. The subsequent node in the plan becomes the new subgoal. This process is repeated until no more nodes remain in the plan.

4. Experiments

We evaluate our approach using the Interactive Gibson simulator (iGibson) [51, 52]. The agent is equipped with a 360° panoramic RGBD camera and ground truth heading information. Prior to navigation, the agent explores each environment via 10 pre-defined trajectories and constructs topological maps denoted as agent-generated (AG) graphs. For evaluation, we use the VLN-CE dataset [26] and evaluate on environments which were seen (val-seen) and unseen (val-unseen) during training. We use the following metrics: success rate (SR), oracle success rate (OS), and navigation error from goal in meters (NE). Oracle success rate uses an oracle to determine when to stop. A navigation episode is successful if the agent stops within $3m$ of the goal position. Details on these metrics can be found in [1, 4].

4.1. Planner Evaluation

In this section, we evaluate the performance of the planner in isolation from the controller. In addition to the agent-generated maps, we also compare planning performance on

		Val-Seen			Val-Unseen		
		SR ↑	OS ↑	NE ↓	SR ↑	OS ↑	NE ↓
AG	GNN	30.6	45.6	8.04	23.5	40.7	8.15
	CMTP	37.9	54.0	6.54	26.0	39.1	7.62
	CMTP-BCE	33.6	48.1	7.11	27.3	42.9	8.13
	CMTP-R	37.4	52.6	6.28	26.5	41.3	7.66
	CMTP-NS	-	66.2	-	-	64.0	-
R2R	GNN	27.3	51.2	9.62	20.7	49.5	11.21
	CMTP	38.6	49.1	6.76	28.7	37.6	7.74
	CMTP-BCE	37.1	51.6	6.97	30.2	42.2	8.36
	CMTP-R	38.4	49.8	6.91	26.8	36.7	8.04
	CMTP-NS	-	61.4	-	-	60.5	-

Table 1: **Planner performance.** Our CMTP models outperform the baseline GNN and perform better on R2R graphs. CMTP-NS is provided as an ablation. See text for details.

the pre-defined navigation graphs from Room2Room (R2R) [4]. We compare our approach with a graph neural network (GNN) baseline as well as variants of our cross-modal transformer planner (CMTP) with different stop mechanisms:

- **Graph neural network (GNN):** A GNN that directly predicts the final navigation goal in the topological map when given the map, start node, and instruction. The instruction is encoded using a transformer and passed as the global feature to the GNN. The GNN then performs classification over the nodes. To extract a navigation plan, we compute a shortest path from the start node to the predicted goal. This baseline is similar in spirit to [11].
- **Cross-Modal Transformer Planner (CMTP):** Our planner which predicts the stop action by performing classification over the [STOP] token and node tokens.
- **CMTP Binary Cross Entropy (CMTP-BCE):** Uses a separate (binary) cross-entropy loss for the [STOP] token, trained jointly with the node classification loss.
- **CMTP Repeated (CMTP-R):** Instead of using a [STOP] token, this CMTP is trained to indicate the stop action by predicting the same node consecutively.
- **CMTP No Stop (CMTP-NS):** Never produces a stop command and is instead used as an ablation. It provides insight into planning performance when correctly timing the stop action is taken out of the picture. The episode terminates after 20 planning steps.

Quantitative results are presented in Table 1. The transformer planner achieves higher success rates and lower navigation errors compared to the GNN. Across all models, the performance on seen environments is significantly higher than unseen environments. This overfitting is most prominent in the CMTP model. On the other hand, the CMTP-

Instruction: Head forward past the kitchen and take the first left at the dining table. Wait in the hallway between the living room and dining room.

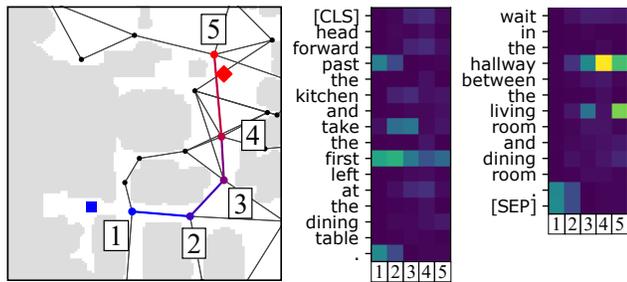


Figure 6: **Cross modal attention** (right) from nodes onto instructions (top). Each column represents a node in the predicted plan (left), going from blue to red.

BCE model has less overfitting at the cost of lower performance on seen environments.

We observe that the oracle success rate of CMTP-NS is significantly higher than all models. In other words, the transformer planner often navigates to the correct locations but struggles to stop at the right time. These results highlight the challenge of accurately stopping at the right moment while illustrating the capability of our model for choosing the correct places to navigate towards.

Different topological maps also affect performance. For example, the transformer planners achieve higher performance on the pre-defined R2R graphs due to the better positioned nodes and edges compared to the AG maps, which sometimes has awkwardly placed nodes near walls or corners. This indicates the importance of (1) developing models which are able to navigate consistently well under different topological maps and (2) exploring new memory representations which support effective and robust navigation.

Cross modal attention. The cross modal attention design of our planner facilitates informational flow between the topological map and language domains, which is intended to ground correlated visual node features to their instructional counterparts. An example is shown in Fig. 6, which visualizes the first layer cross modal attention from predicted nodes to the input instruction. We see that words which encode the first half of the instruction correspond with earlier nodes in the trajectory, whereas words such as “hallway” and “living room” in the latter half of the sentence correspond to the last three nodes of the trajectory.

We note that the model will sometimes focus attention on articles or punctuation. This behavior is in line with previous observations [42] that the cross modal layer tends to focus on articles over single noun references (e.g. attention on “the” or “living” instead of “room”), using them as anchor words while likely relying on the language encoder to connect the articles to the specific objects they reference.

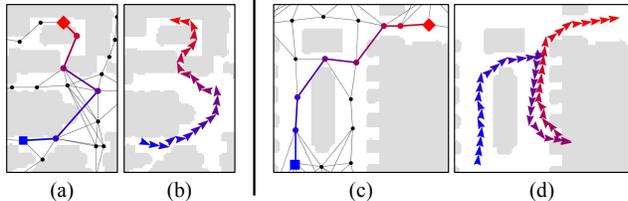


Figure 7: **Plan execution.** The controller is able to follow the navigation plans (a, c) as shown by the trajectories (b, d). It can also correct for mistakes via backtracking (d).

	Unbiased	$\sigma^2 = 5$	$\sigma^2 = 3$	$\sigma^2 = 2$
Pointgoal	99.8 / 99.5	99.9 / 99.4	99.9 / 99.7	99.6 / 99.2
AG	83.3 / 82.5	89.4 / 87.2	93.1 / 89.6	77.7 / 75.9
R2R	87.2 / 82.0	93.3 / 89.1	92.8 / 92.6	84.3 / 81.8

Table 2: **Controller success rate** in [seen / unseen] environments for random point goals ($\leq 5m$ away) and for randomly sampled plans from R2R or AG graphs (≤ 12 nodes away). Adding the bias greatly improves performance.

4.2. Controller Evaluation

We evaluate the controller separate from our planner by executing episodes of random plans in the simulated environment. A navigation success threshold distance of $1m$ is used when evaluating the controller alone.

Fig. 7 shows qualitative results of the controller executing navigation plans. The controller is able to navigate while avoiding collision, despite graph edges cutting through obstacles in the planned path. With the inclusion of the previous node into the controller input, the agent is able to learn intelligent backtracking behavior in which the agent backtracks to the previously known position when it makes a mistake and deviates from the path.

The quantitative results reported in Tab. 2 support our observations that the controller works well. In navigation to point goals (row 1), it achieves high success rates even in unseen environments. Executing multi-node plans (rows 2-3) also produces high success rates, with better performance on R2R graphs than AG graphs due to their better positioned nodes providing more informative observations.

4.3. VLN Evaluation

Finally, we evaluate our full navigation system. We compare with the cross-modal approach in VLN-CE [26] since it is the only prior work, to our knowledge, that considers low-level actions like our work. VLN-CE is trained end-to-end and utilizes a gated recurrent unit (GRU) [10] to encode the agent’s history. In addition to prior metrics, we report success weighted by path length (SPL) [1], and planner success (PS) and controller success (CS) for modular approaches. Our reported VLN-CE results in iGibson differ

		Val-seen					Val-unseen						
		PS \uparrow	CS \uparrow	SR \uparrow	OS \uparrow	NE \downarrow	SPL \uparrow	PS \uparrow	CS \uparrow	SR \uparrow	OS \uparrow	NE \downarrow	SPL \uparrow
	VLN-CE [26]	-	-	22.2	29.8	7.6	16.5	-	-	19.7	27.2	7.8	14.1
AG	GNN	30.6	91.7	29.6	46.0	8.2	26.3	23.5	84.8	22.1	40.0	8.4	17.2
	CMTP	37.9	92.2	35.9	56.2	6.6	30.5	26.0	86.6	23.1	39.2	7.9	19.1
	CMTP-BCE	33.6	89.4	32.6	49.6	7.2	27.5	27.3	82.6	25.2	42.2	8.4	20.2
	CMTP-R	37.4	93.1	36.3	53.9	6.5	31.3	26.5	87.8	25.3	42.6	7.9	20.3
R2R	GNN	27.3	89.1	26.6	51.9	9.4	21.2	20.7	88.6	20.3	48.3	10.9	16.4
	CMTP	38.6	91.7	36.1	45.4	7.1	31.2	28.7	90.0	26.4	38.0	7.9	22.7
	CMTP-BCE	37.1	89.6	34.4	49.3	7.4	29.9	30.2	89.6	28.9	40.7	8.4	24.1
	CMTP-R	38.4	89.6	35.1	48.2	7.5	31.2	26.8	90.0	25.2	36.0	8.3	20.9

Table 3: **Integrated system performance.** The modular approaches (GNN, CMTP) outperform the end-to-end VLN-CE model in success rate (SR), with our CMTP models achieving the best performance. Exploration steps not included in SPL.

slightly from [26] – see supplemental for details.

Table 3 shows that our CMTP models achieve the highest VLN success rates and SPL compared to the GNN and VLN-CE baselines. Moreover, the methods which use topological maps outperform the end-to-end VLN-CE baseline. While these results may be expected given that the former methods pre-explore the environments, the gap in performance is also due to reducing the search space through the topological maps. By combining our cross-modal planner and controller into an integrated system, we get the best of both worlds: the planner achieves high planning success rates while the controller is able to bring those plans to fruition with low-level discrete actions.

We also ask: how do exploration trajectories influence navigation performance? We collected 5 sets of AG maps using different exploration trajectories, averaging 3.6m traveled per m^2 of traversable area in the environment and about $300\times$ longer than the test-time navigation trajectories. Using our CMTP model, we found consistent performance across all maps with a standard deviation of only 1.1% in both full navigation SR and SPL. While this highlights the consistency of navigation performance under our AG maps, further investigation into more effective topological representations is left for future work.

Lastly, Fig. 8 shows exciting qualitative results of our full navigation system. As illustrated by the agent’s predicted plan and executed trajectory, the modular approach makes it easy to interpret the agent’s intents and evaluate its ability to realize those intents. For instance, in Fig. 8 (top) the planner predicts a path between the pool table and couch as specified by the instruction. The controller executes the plan by converting it into a smooth trajectory which traverses between the two objects rather than around them.

5. Conclusion

We introduced a novel approach for VLN with topological maps. By using a modular system in conjunction

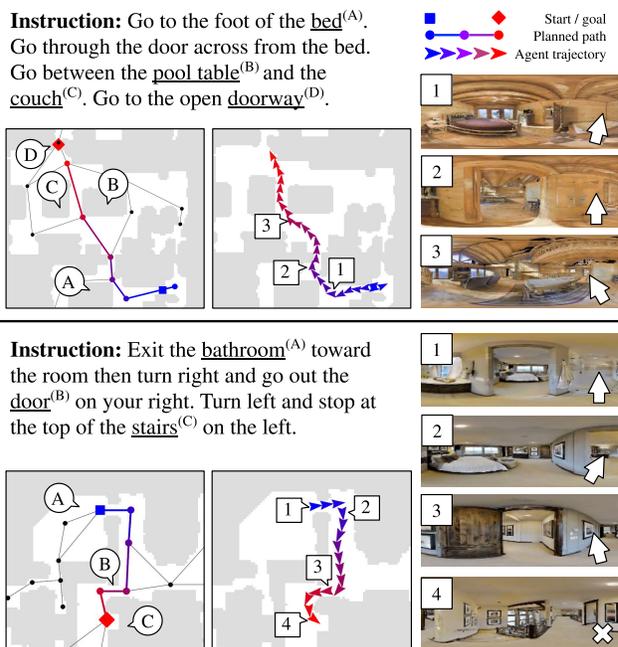


Figure 8: **Our full navigation system.** We show the navigation instruction (top-left), the panorama observations (right), and the predicted navigation plan (bottom-left) and execution (bottom-middle). The arrows in each panorama indicate the predicted action at the given panorama.

with topological maps, we saw a boost in navigation performance compared to an end-to-end baseline. Potential extensions include demonstrating VLN on a real robot [3] and using topological maps constructed at test time [7]. We hope our work inspires further research in building robot systems which communicate effectively and operate robustly.

Acknowledgments. This work was funded through the support of ONR MURI Award #W911NF-15-1-0479. We also thank anonymous reviewers for their feedback.

References

- [1] 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, 7
- [2] Peter Anderson, Ayush Shrivastava, Devi Parikh, Dhruv Batra, and Stefan Lee. Chasing ghosts: Instruction following as bayesian state tracking. In *Advances in Neural Information Processing Systems*, pages 371–381, 2019. 2
- [3] Peter Anderson, Ayush Shrivastava, Joanne Truong, Arjun Majumdar, Devi Parikh, Dhruv Batra, and Stefan Lee. Sim-to-real transfer for vision-and-language navigation. *arXiv preprint arXiv:2011.03807*, 2020.
- [4] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3674–3683, 2018. 1, 2, 6
- [5] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018. 4
- [6] Devendra Singh Chaplot, Dhiraj Gandhi, Saurabh Gupta, Abhinav Gupta, and Ruslan Salakhutdinov. Learning to explore using active neural slam. In *International Conference on Learning Representations*, 2019. 1, 2
- [7] Devendra Singh Chaplot, Ruslan Salakhutdinov, Abhinav Gupta, and Saurabh Gupta. Neural topological slam for visual navigation. In *CVPR*, 2020. 1, 2, 3,
- [8] Kevin Chen, Juan Pablo de Vicente, Gabriel Sepulveda, Fei Xia, Alvaro Soto, Marynel Vázquez, and Silvio Savarese. A behavioral approach to visual navigation with graph localization networks. In *Proceedings of Robotics: Science and Systems*, Freiburg/Breisgau, Germany, June 2019. 1, 2
- [9] Tao Chen, Saurabh Gupta, and Abhinav Gupta. Learning exploration policies for navigation. *arXiv preprint arXiv:1903.01959*, 2019. 1
- [10] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. 7
- [11] Zhiwei Deng, Karthik Narasimhan, and Olga Russakovsky. Evolving graphical planner: Contextual global planning for vision-and-language navigation. *arXiv preprint arXiv:2007.05655*, 2020. 2, 6
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [14] Alexey Dosovitskiy and Vladlen Koltun. Learning to act by predicting the future. In *International Conference on Learning Representations (ICLR)*, 2017. 1, 2
- [15] Kuan Fang, Alexander Toshev, Li Fei-Fei, and Silvio Savarese. Scene memory transformer for embodied agents in long-horizon tasks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 538–547, 2019. 1, 2
- [16] 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 S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 3314–3325. Curran Associates, Inc., 2018. 1, 2
- [17] Saurabh Gupta, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2616–2625, 2017. 1, 2
- [18] Saurabh Gupta, David Fouhey, Sergey Levine, and Jitendra Malik. Unifying map and landmark based representations for visual navigation. *arXiv preprint arXiv:1712.08125*, 2017. 2
- [19] Saurabh Gupta, Varun Tolani, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. *International Journal of Computer Vision*, 2019. 2
- [20] Weituo Hao, Chunyuan Li, Xiujun Li, Lawrence Carin, and Jianfeng Gao. Towards learning a generic agent for vision-and-language navigation via pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13137–13146, 2020. 1, 2
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [22] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 1
- [23] Peter Karkus, Xiao Ma, David Hsu, Leslie Pack Kaelbling, Wee Sun Lee, and Tomás Lozano-Pérez. Differentiable algorithm networks for composable robot learning. In *Proceedings of Robotics: Science and Systems*, 2019. 1, 3
- [24] 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 Conference on Computer Vision and Pattern Recognition*, pages 6741–6749, 2019. 1, 2
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5

- [26] Jacob Krantz, Erik Wijmans, Arjun Majumdar, Dhruv Batra, and Stefan Lee. Beyond the nav-graph: Vision-and-language navigation in continuous environments. 2020. [1](#), [2](#), [6](#), [7](#),
- [27] Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam. 2018. [5](#)
- [28] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. VILBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*, pages 13–23, 2019. [2](#)
- [29] Chih-Yao Ma, Jiasen Lu, Zuxuan Wu, Ghassan AlRegib, Zsolt Kira, Richard Socher, and Caiming Xiong. Self-monitoring navigation agent via auxiliary progress estimation. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019. [1](#), [2](#)
- [30] Chih-Yao Ma, Zuxuan Wu, Ghassan AlRegib, Caiming Xiong, and Zsolt Kira. The regretful agent: Heuristic-aided navigation through progress estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6732–6740, 2019. [1](#), [2](#)
- [31] 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, 2020. [2](#)
- [32] Cynthia Matuszek, Dieter Fox, and Karl Koscher. Following directions using statistical machine translation. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 251–258. IEEE, 2010. [1](#), [2](#)
- [33] Xiangyun Meng, Nathan Ratliff, Yu Xiang, and Dieter Fox. Scaling local control to large-scale topological navigation. *arXiv preprint arXiv:1909.12329*, 2019. [1](#), [2](#), [3](#)
- [34] Piotr Mirowski, Razvan Pascanu, Fabio Viola, Hubert Soyer, Andrew J Ballard, Andrea Banino, Misha Denil, Ross Goroshin, Laurent Sifre, Koray Kavukcuoglu, et al. Learning to navigate in complex environments. *arXiv preprint arXiv:1611.03673*, 2016. [2](#)
- [35] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937, 2016. [2](#)
- [36] Nils J Nilsson. Shakey the robot. Technical report, SRI INTERNATIONAL MENLO PARK CA, 1984. [3](#)
- [37] Emilio Parisotto and Ruslan Salakhutdinov. Neural map: Structured memory for deep reinforcement learning. In *International Conference on Learning Representations*, 2018. [2](#)
- [38] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635, 2011. [5](#)
- [39] Nikolay Savinov, Alexey Dosovitskiy, and Vladlen Koltun. Semi-parametric topological memory for navigation. In *International Conference on Learning Representations (ICLR)*, 2018. [1](#), [2](#), [3](#)
- [40] S. Schubert, P. Neubert, J. Pöschmann, and P. Protzel. Circular convolutional neural networks for panoramic images and laser data. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 653–660, 2019. [5](#)
- [41] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. VI-bert: Pre-training of generic visiolinguistic representations. In *International Conference on Learning Representations*, 2020. [2](#)
- [42] Hao Tan and Mohit Bansal. LXMERT: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5100–5111, Hong Kong, China, Nov. 2019. Association for Computational Linguistics. [2](#), [4](#), [5](#), [7](#)
- [43] Hao Tan, Licheng Yu, and Mohit Bansal. Learning to navigate unseen environments: Back translation with environmental dropout. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2610–2621, 2019. [1](#), [2](#)
- [44] Sebastian Thrun. Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, 99(1):21–71, 1998. [1](#)
- [45] Sebastian Thrun, Mike Montemerlo, Hendrik Dahlkamp, David Stavens, Andrei Aron, James Diebel, Philip Fong, John Gale, Morgan Halpenny, Gabriel Hoffmann, et al. Stanley: The robot that won the darpa grand challenge. *Journal of field Robotics*, 23(9):661–692, 2006. [3](#)
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [1](#), [2](#), [4](#), [5](#)
- [47] Ayzaan Wahid, Austin Stone, Kevin Chen, Brian Ichter, and Alexander Toshev. Learning object-conditioned exploration using distributed soft actor critic. *arXiv preprint arXiv:2007.14545*, 2020. [1](#)
- [48] Xin Wang, Qiuyuan Huang, Asli Celikyilmaz, Jianfeng Gao, Dinghan Shen, Yuan-Fang Wang, William Yang Wang, and Lei Zhang. Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6629–6638, 2019. [1](#), [2](#)
- [49] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 37–53, 2018. [2](#)
- [50] 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. *arXiv preprint arXiv:1911.00357*, 2019. [1](#)
- [51] Fei Xia, Chengshu Li, Kevin Chen, William B Shen, Roberto Martín-Martín, Noriaki Hirose, Amir R Zamir, and Li Fei-Fei Silvio Savarese. Gibson env v2: Embodied simulation environments for interactive navigation. 2019. [6](#)
- [52] Fei Xia, William B Shen, Chengshu Li, Priya Kasimbeg, Micael Edmond Tchammi, Alexander Toshev, Roberto Martín-

- Martín, and Silvio Savarese. Interactive gibbon benchmark: A benchmark for interactive navigation in cluttered environments. *IEEE Robotics and Automation Letters*, 5(2):713–720, 2020. [6](#)
- [53] Xiaoxue Zang, Ashwini Pokle, Marynel Vázquez, Kevin Chen, Juan Carlos Niebles, Alvaro Soto, and Silvio Savarese. Translating navigation instructions in natural language to a high-level plan for behavioral robot navigation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2657–2666, 2018. [2](#)
- [54] Jingwei Zhang, Lei Tai, Joschka Boedecker, Wolfram Burgard, and Ming Liu. Neural slam: Learning to explore with external memory. *arXiv preprint arXiv:1706.09520*, 2017. [2](#)
- [55] Fengda Zhu, Yi Zhu, Xiaojun Chang, and Xiaodan Liang. Vision-language navigation with self-supervised auxiliary reasoning tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10012–10022, 2020. [1](#), [2](#)
- [56] Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 3357–3364. IEEE, 2017. [2](#)