

Bayesian Relational Memory for Semantic Visual Navigation

Yi Wu^{†¶} Yuxin Wu[‡] Aviv Tamar[§] Stuart Russell[†] Georgia Gkioxari[‡] Yuandong Tian[‡]
[†]UC Berkeley [‡]Facebook AI Research [§]Technion [¶]OpenAI

[†]{jxwuyi, russell}@cs.berkeley.edu [‡]{yuxinwu, gkioxari, yuandong}@fb.com [§]avivt@technion.ac.il

Abstract

We introduce a new memory architecture, Bayesian Relational Memory (BRM), to improve the generalization ability for semantic visual navigation agents in unseen environments, where an agent is given a semantic target to navigate towards. BRM takes the form of a probabilistic relation graph over semantic entities (e.g., room types), which allows (1) capturing the layout prior from training environments, i.e., prior knowledge, (2) estimating posterior layout at test time, i.e., memory update, and (3) efficient planning for navigation, altogether. We develop a BRM agent consisting of a BRM module for producing sub-goals and a goal-conditioned locomotion module for control. When testing in unseen environments, the BRM agent outperforms baselines that do not explicitly utilize the probabilistic relational memory structure.

1. Introduction

Memory is a crucial component for intelligent agents to gain extensive reasoning abilities over a long horizon. One such challenge is visual navigation, where an agent is placed to an environment with unknown layouts and room connectivity, acts on visual signal perceived from its surrounding, explores and reaches a goal position efficiently.

Due to partial observability, the agent must memorize its past experiences and react accordingly. Hence, deep learning (DL) models for visual navigation often encodes memory structures in its design. LSTM is initially used to as general-purpose implicit memory [36, 35, 10]. Recently, to improve the performance, explicit and navigation-specific structural memory are used [42, 24, 46, 25].

Two categories exist for navigation-specific memory: the spatial memory [42, 24] and topological memory [25, 46]. The core idea of spatial memory is to extend the 1-dimensional memory in LSTM to a 2-dimensional matrix that represents the spatial structure of the environment, where a particular entry in the matrix corresponds to a 2D location/region in the environment. Due to its regular structure, value iteration [50] can be applied directly for effective planning over the memory matrix.

However, planning on such spatial memory can be computationally expensive for environments with large rooms. To navigate, precise localization of an agent is often not unnecessary. Extensive psychological evidences [46] also show that animals do not rely strongly on metric representations [54, 22]. Instead, humans primarily depend on a landmark-based navigation strategy, which can be supported by qualitative topological knowledge of the environment [22]. Therefore, it is reasonable to represent the memory as a topological graph where the vertices are landmarks in the environment and edges denote short-term reachability between landmarks. During navigation, a localization network is trained to identify the position of the agent and the goal w.r.t. the landmarks in the memory and an efficient graph search can be used for long-term planning.

However, still human navigation shows superior generalization performance which cannot be explained by either spatial or topological memory. For example, first-time home visitors naturally move towards the kitchen (rather than outdoor or toilet) to get a plate; from kitchen to bedroom, they know that living room may be an intermediate waypoint. Although visually different, such *semantic knowledge*, i.e., the “close-by” relations over semantic entities, are naturally shared across environments and can be learned from previous experience to guide future navigation. In comparison, existing approaches of topological memory assume pre-exploration experiences of the environment before navigation starts [25, 46], provide no memory updating operations, and cannot incorporate the prior knowledge of scene layouts and configurations from previously seen environments.

In this work, we propose a new memory design for visual navigation, Bayesian Relational Memory (BRM), which (1) captures the prior knowledge of scene layouts from training environments and (2) allows both efficient planning and updating during test-time exploration. BRM can be viewed as a probabilistic version of a topological memory with semantic abstractions: each node in BRM denotes a semantic concept (e.g., object category, room type, etc), which can be detected via a neural detector, and each edge denotes the relation between two concepts. In each environment, a single relation may be present or not. For each relation (edge), we can average its existences over all training environments

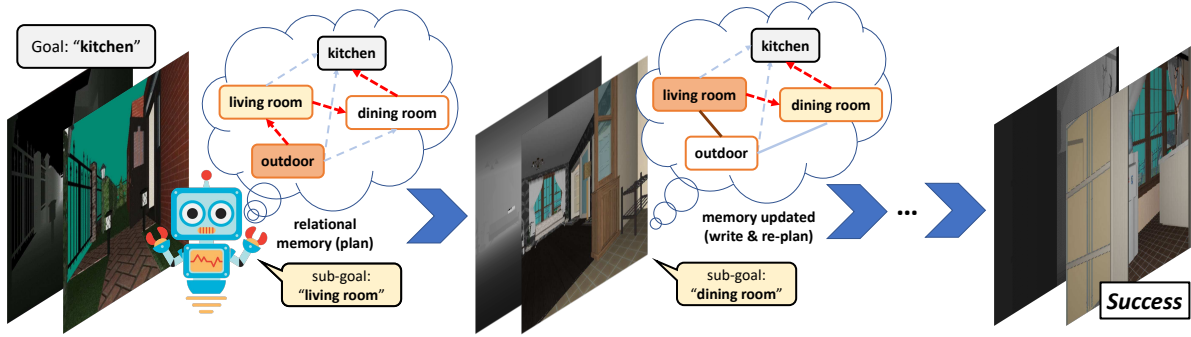


Figure 1. A demonstration of our task and method. The agent perceives visual signals and needs to find the kitchen, which cannot be seen from outside (leftmost). So the agent plans in its memory and conclude that it may first find a likely intermediate waypoint, i.e., living room. Then the agent repeatedly updates its belief of the environment layout, re-plans accordingly and reaches the kitchen in the end.

and learn an existence probability to denote the *prior knowledge* between two semantic concepts. During exploration at test time, we can incrementally observe the existences of relations within that particular test environment. Therefore, we can use these environment specific observations to update the probability of those relations in the memory via the Bayes rule to derive the *posterior knowledge*. Additionally, we train a semantic-goal-conditioned LSTM locomotion policy for control via deep reinforcement learning (DRL), and by planning on the relation graph with posterior probabilities, the agent picks the next semantic sub-goal to navigate towards.

We evaluate our BRM method in a semantic visual navigation task on top of the House3D environment [58, 10], which provides a diverse set of objects, textures and human-designed indoor scenes. The semantic scenes and entities in House3D are fully labeled (with noise), which are natural for our BRM model. In the navigation task, the agent observes first-person visual signals and needs to navigate towards a particular room type. We utilize the room types as the semantic concepts (nodes) in BRM and the "close-by" relations as the edges. We evaluate on unseen environments with random initial locations and compare our learned model against other DRL-based approaches without the BRM representation. Experimental results show that the agent equipped with BRM can achieve the semantic goal with higher success rates and fewer navigation steps.

Our contributions are as follows: (1) we proposed a new memory representation, Bayesian Relational Memory (BRM), in the form of probabilistic relation graphs over semantic concepts; (2) BRM is capable of encoding prior knowledge over training environments as well as efficient planning and updating at test time; (3) by integrating BRM into a DRL locomotion policy, we show that the generalization performances can be significantly improved.

2. Related Work

Navigation is one of the most fundamental problems in mobile robotics. Traditional approaches (like SLAM) build

metric maps via sensory signals, which is subsequently used for planning [17, 4, 51, 16]. More recently, thanks to the advances of deep learning, end-to-end approaches have been applied to tackle navigation in various domains, such as mazes [36, 27], indoor scenes [63, 5, 48, 37, 30], autonomous driving [7, 59, 13], and Google street view [35]. There are also nice summaries of recent progresses [2, 37]. We focus on indoor navigation scenario with House3D environment [58] which contains real-world-consistent relations between semantic entities and provides ground-truth labels of objects and scenes.

There are also works studying visual navigation under natural language guidance, including instruction following [6, 39, 23, 55, 3] and question answering [11, 10, 1]. These tasks require the agent to understand the natural language and reason accordingly in an interactive environment. In our semantic navigation task, the goal instruction is simplified to a single semantic concept and hence we focus more on the reasoning ability of navigation agents.

In our work, the reasoning is performed on the relations over semantic concepts from visual signals. Similar ideas of using semantic knowledge to enhance reasoning have been applied to image classification [34, 56], segmentation [52, 64], situation recognition [33], visual question answering [57, 9, 53, 28, 26], image retrieval [29] and relation detection [62]. Savinov et al. [47] and Kuang et al. [18] also consider extracting visual concepts dynamically by treating every received input frame as an individual concept and storing them in the memory. The most related work to ours is a concurrent one by Wei et al. [60], which considers visual navigation towards an object category and utilizes a knowledge graph as the prior knowledge. Wei et al. [60] use a fixed graph to extract features for the target as an extra input to the locomotion without graph updating or planning. While in our work, the relation graph is used in a Bayesian manner as a representation for the memory, which unites use of prior knowledge, updating and planning altogether.

From a reinforcement learning perspective, our work is related to the model-based approaches [14, 45, 61, 44, 32], in the sense that we model the environment layout via a re-

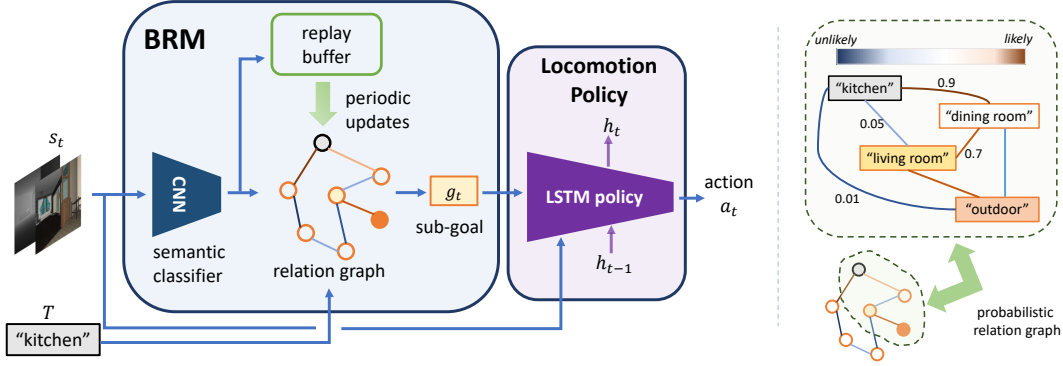


Figure 2. The architecture overview of our proposed navigation agent equipped with the Bayesian Relational Memory (BRM).

lation graph and plan on it. Our work is also related to hierarchical reinforcement learning [49, 12, 31], where the controller (BRM) produces a high-level sub-goal for the sub-policy (locomotion) to pursue. Furthermore, BRM learns from multi-task training and its update operation fast adapts the prior relations to the test environment, which can be also viewed as a form of meta-learning [19, 15, 38].

3. Method

3.1. Task Setup

We consider the semantic visual navigation problem where an agent interacts with an environment with discrete time steps and navigates towards a semantic goal. In House3D, the semantic entities of interest are room types and we assume a fixed number of K categories. In the beginning of an episode, the agent is given a semantic target $T \in \mathcal{T} = \{T_1, \dots, T_K\}$ to navigate towards for a success. At each time step t , the agent receives a visual observation s_t and produces an action $a_t \in \mathcal{A}$ conditioning on s_t and T .

We aim to learn a neural agent that can *generalize* to unseen environments. Hence, we train the agent on a training set $\mathcal{E}_{\text{train}}$, where the ground-truth labels are assumed, and validate on $\mathcal{E}_{\text{valid}}$. Evaluation is performed on another separate set of environments $\mathcal{E}_{\text{test}}$. At test time, the agent only access to the visual signal s_t without any pre-exploration experiences of the test environment.

3.2. Method Overview

The overall architecture of a BRM agent is shown in Fig. 2, which has two modules, the Bayesian Relational Memory (BRM) as well as an LSTM locomotor policy for control (Fig. 2 left). Particularly, the key component of a BRM agent is a probabilistic relation graph (Fig. 2 right), where each node corresponds to a particular semantic target T_i . For semantic target T_i and T_j , the edge between them denotes the “close-by” relation and the probability of that edge implies how likely T_i and T_j are close to each other in the current environment.

At a high level, the locomotion is a semantic-goal-

conditioned policy which takes in both the visual input s_t and the sub-goal $g \in \mathcal{T}$ produced by the BRM module to produce actions towards g . The BRM module takes in the visual observation s_t at each time step, extracts the semantic information via a CNN detector and store them in a replay buffer. We periodically update the posterior probability of each edge in the relation graph and re-plan to produce a new sub-goal. In our work, the graph is updated every fixed number of N steps. Notably, we do not assume existences of all concepts — in case of a missing concept in particular environment, the posterior of its associated edges will approach zero as more experiences gained in an episode.

3.3. Bayesian Relational Memory

The BRM module consists of two parts, a semantic classifier and the most important component, a probabilistic relation graph over semantic concepts, in the form of a probabilistic graphical model allowing efficient planning and posterior updates.

Intuitively, at each time step, the agent detects the surrounding room type and maintains the probability of whether two room types T_i and T_j are “nearby” in the current environment. If the agent starts from room T_i and reaches another room T_j within a few steps, the probability of “nearby” relation between T_i and T_j should be increased¹; otherwise the probability should be decreased. Periodically, the graph is updated and the agent finds the most likely path from the current room towards the target as a navigation guidance.

We introduce these two components in details as well as how to update and plan with BRM as follows.

Semantic classifier: The semantic classifier detects the room type label c_t for the agent’s surrounding region. It can be trained by supervised learning on $\mathcal{E}_{\text{train}}$. Note that for robust room type classification, only using the first-person view image may not be enough. For example, the agent in a bedroom may face towards a wall, where the first-person image is not informative at all for the classifier, but a bed might be just behind. So we take the panoramic view as the

¹In perfect noiseless setting, the relation should with probability 1.





Figure 3. Overview of the planning (Left) and the update operation (Right) with BRM.

classifier input, which consists of 4 images, s_t^1, \dots, s_t^4 with different view angles. We use a 10-layer CNN with batch normalization to extract features $f(s_t^i)$ for each s_t^i and compute the attention weights over these visual features by $l_i = f(s_t^i) W_1^T W_2 [f(s_t^1), \dots, f(s_t^4)]$ $a_i = \text{softmax}(l_i)$ with parameters W_1, W_2 . Then the weighted average of these four features $\sum_i a_i f(s_t^i)$ is used for the final sigmoid predictions for each semantic concept T_i ². This results in a K -dimensional binary vector $c_t \in \{0, 1\}^K$ at time t .

Probabilistic relation graph: We represent the probabilistic graph in the form of a graphical model $P(z, y; \psi)$ with latent variable z , observation variable y and parameter ψ .

Since we have K semantic concepts, there are K^3 nodes and $K(K-1)/2$ relations (edges) in the graph. Each relation is probabilistic (i.e., it may exist or not with probability) and before entering a particular environment, we only hold a prior belief of that relation. Formally, for the relation between T_i and T_j , we adopt a Bernoulli variable $z_{i,j}$ defined by $z_{i,j} \sim \text{Bernoulli}(\psi_{i,j}^{\text{prior}})$, where parameter $\psi_{i,j}^{\text{prior}}$ denotes the prior belief of $z_{i,j}$ existing. During exploration, the agent can noisily observe $z_{i,j}$ and use the noisy observations to estimate the true value of $z_{i,j}$. We define the noisy observation model $P(y_{i,j}|z_{i,j})$ by

$$y_{i,j} \sim \begin{cases} \text{Bernoulli}(\psi_{i,j}^{\text{obs}}) & \text{if } z_{i,j} = 0 \\ \text{Bernoulli}(1 - \psi_{i,j}^{\text{obs}}) & \text{if } z_{i,j} = 1 \end{cases}, \quad (1)$$

where ψ^{obs} is another parameter to learn. At each time step, the agent holds an overall posterior belief $P(z|\mathcal{Y})$ of relation existences within the current environment, based on its experiences \mathcal{Y} , namely the samples of variable y .

Posterior update and planning: A visualization of the procedures is shown in Fig. 3. We assume the agent explores the current environment for a *short* horizon of N steps and stores the recent semantic signals c_t, \dots, c_{t+N} in the replay buffer. Then we compute the bit-OR operation over these binary vectors $B = c_t \text{ OR } \dots \text{ OR } c_{t+N}$. B represents all the visited regions within a short (N -step) exploration period. When two targets appear concurrently in a short trajectory, they are assumed to be “close-by”. For T_i and T_j with $B(T_i) = B(T_j) = 1$, T_i and T_j should




Figure 4. Learning the prior probability of a particular relation by analyzing the existence of that relation in each training environment and then summarizing the overall statistics.

be nearby in the current environment, namely a sample of $y_{i,j} = 1$; otherwise for $B(T_i) \neq B(T_j)$, we get a sample of $y_{i,j} = 0$. With all the history samples of y as \mathcal{Y} , we can perform posterior inference, i.e., compute posterior Bernoulli distribution $P(z_{i,j}|\mathcal{Y}_{i,j})$, for each $z_{i,j}$ by the Bayes rule.

Let $\hat{z}_{i,j} = P(z_{i,j}|\mathcal{Y}_{i,j})$ denote the posterior probability of relation over T_i and T_j . Given the current beliefs \hat{z} , the semantic signal c_t and the target T , we search for an optimal plan $\tau^* = \{\tau_0, \tau_1, \dots, \tau_{m-1}, \tau_m\}$ over the graph, where $\tau_i \in \{1 \dots K\}$ denotes an index of concepts, so that the joint belief along the path from some current position to the goal is maximized: $\tau^* = \arg \max_{\tau} c_t(T_{\tau_0}) \prod_{i=1}^m \hat{z}_{\tau_{i-1}, \tau_i}$.

After obtaining τ^* , we execute the locomotion policy for sub-goal $g_t = T_{\tau_1^*}$, and then periodically update the graph, clear the replay buffer and re-plan every N steps.

Learning the probabilistic graph: The parameter ψ has two parts: ψ^{prior} for the prior of z and ψ^{obs} for the noisy observation y .

For the prior parameter ψ^{prior} , we learn from $\mathcal{E}_{\text{train}}$ with the ground truth labels. A visualization is shown in Fig. 4. For each pair of room types T_i and T_j , we enumerate all training environments and run *random explorations* from some location of room type T_i . If eventually the agent reaches somewhere of room type T_j , we consider T_i and T_j are nearby and therefore a *positive* sample $z_{i,j} = 1$; otherwise a *negative* sample $z_{i,j} = 0$. Suppose \mathcal{Z} denotes all the samples we obtained for z . We run maximum likelihood estimate for ψ^{prior} by maximizing $L_{\text{MLE}}(\psi^{\text{prior}}) = P(\mathcal{Z}|\psi^{\text{prior}})$. We can choose any exploration strategy such that for any “close-by” targets, they appear together in a short exploration trajectory more frequently. Random exploration is the most lightweight option among all.

The noisy observation parameter ψ^{obs} is related to the

²It is a multi-label classification setting. Imagine an open kitchen with both cooking facilities and dining tables could have two labels.

³In fact we have $K+1$ nodes. For clarity purpose, we use K here and explain the details of the extra node later in Sec. 3.5.




Figure 5. The learned prior of the relations, including the most (red) and least (blue) likely nearby rooms for dining room (Left), bedroom (Mid.) and outdoor (Right), with numbers denoting ψ^{prior} , i.e., the prior probability of two room types are nearby.

performance of the locomotion policy $\mu(\theta)$: if $\mu(\theta)$ has a higher navigation success rate, ψ^{obs} should be smaller (i.e., low noise level); when $\mu(\theta)$ is poor, ψ^{obs} should be larger (cf. Eq. (1)). Unfortunately, there is no direct learning supervision for ψ^{obs} . However, we can evaluate the “goodness” of a particular value of ψ^{obs} by evaluating the success rate of the overall BRM agent on $\mathcal{E}_{\text{valid}}$. Therefore, we can simply run grid search to derive the best parameter.

3.4. The Goal Conditioned Policy

We learn an LSTM locomotion policy $\mu(s_t, g; \theta)$ parameterized by θ , which conditions on observation s_t and navigates towards goal g . Following Wu et al. [58], we learn $\mu(s_t, g; \theta)$ by formulating the task as a reinforcement learning problem with shaped reward: when the agent moves towards target room g , it receives a positive reward proportional to the distance decrements; if the agent moves apart or hits an obstacle, a penalty is presented. A success reward of 10 and a time penalty of 0.1 are also assigned. We optimize the policy on $\mathcal{E}_{\text{train}}$ via the actor-critic method [40] with a curriculum learning paradigm by periodically increasing the maximum spawn distance to the target g . Additionally, thanks to a limited set of K targets, we adopt a behavior approach [8] for improved performances: we train a separate policy $\mu_i(s_t; \theta_i)$ for each semantic target T_i and when given the sub-goal g from the BRM module, we directly execute its corresponding behavior network. We empirically observe it performs better than the original gated-attention policy in Wu et al.[58]. Such an architecture is also common technique in other domains such as RL [41, 20] and robotics [21].

3.5. Implementation Details

We introduce those crucial details below and defer the remaining to appendix.

Nodes in the relation graph: In the previous content, we assume K pre-selected semantic concepts as graph nodes. However, it is not rare to reach some situation that cannot be categorized into any existing semantic category. In practice, we treat $c_t = \mathbf{0}$ as a special “unknown” concept. Hence, the BRM module actually contains $K + 1$ nodes. This is conceptually similar to natural language processing: a semantic concept is a word; the set of all concepts can be viewed as the dictionary; and $c_t = \mathbf{0}$ corresponds to the special “out-

of-vocabulary” token.

Smooth temporal classification: Although the semantic classifier achieves high accuracy on validation data, the predictions may not be temporally consistent, which brings extra noise to the BRM module. For temporally smooth prediction at test time, we set a restricted threshold over the sigmoid output and apply a filtering process on top of that: the actual prediction label for room type T_i will be 1 only if the sigmoid output remains at least 0.9 confidence score for consecutively 3 time steps.

Graph learning: For learning ψ^{prior} , we run a random exploration of 300 steps and collect 50 samples for each $z_{i,j}$ per training environment. Also, learning ψ^{prior} does not depend on the locomotion $\mu(s_t, g; \theta)$ and can be reused even with different control policies. ψ^{obs} depends on the locomotion so it must be learned after $\mu(s_t, g; \theta^*)$ is obtained.

4. Experiments

We experiment on the House3D environment and proceed to answer the following two questions: **(1)** Does the BRM agent captures the underlying semantic relations and behave as we expected? **(2)** Does the BRM agent generalize better in test environments than the baseline methods?

We first introduce evaluation preliminaries and baseline methods in Sec. 4.1, 4.2. Then we answer the first question qualitatively in Sec. 4.3. In Sec. 4.4, 4.6, we quantitatively show that our BRM agents generally achieve higher test success rates (i.e., better generalization) and spend fewer steps to reach the targets (i.e., more efficient) than all the baselines. We choose a fixed $N = 10$ for all the BRM agents. Ablation study on the design choices of BRM is presented in Sec. 4.5. More details can be found in appendix.

4.1. Preliminaries

We consider the RoomNav task on the House3D environment [58] where $K = 8$ room types are selected as the semantic goals, including “kitchen”, “living room”, “dining room”, “bedroom”, “bathroom”, “office”, “garage” and “outdoor”. The House3D environment provides a success check for whether the agent has reached a specific room target or not while we also experimented on the setting of the agent predicting termination on its own (Sec. 4.6). House3D provides a training set of 200 houses, a test set of 50 houses and a validation set of 20 houses. At training




Figure 6. Qualitative comparison between BRM (purple), basic DRL policy (pure $\mu(\theta)$, blue) and random policy (red) with horizon $H = 1000$ (Sec. 4.3). Y-axis is the test success rate and x-axis is the distance in meters to the target. When targets become farther way from the starting point, the success rate of BRM stays high while the basic DRL policy quickly degenerates to random.

time, all the approaches adopt the ground-truth semantic labels regardless of the semantic classifier.

We evaluate the generalization performances of different approaches with two metrics, *success rate* and *Success weighted by Path Length (SPL)*, under different horizons. SPL, proposed by Anderson et al.[2], is a function of both success rate and episode length defined by $\frac{1}{C} \sum_i S_i \frac{L_i}{\max(L_i, P_i)}$, where C is total episodes evaluated, S_i indicates whether the episode is success or not, L_i is the ground truth shortest path distance in the episode, P_i is the number of steps the agent actually took. SPL is upper-bounded by success rate and assigns more credits to agents accomplishing their tasks faster.

4.2. Baseline Methods

Random policy: The agent samples a random action per step, denoted by “random”.

Pure DRL agent: This LSTM agent does not have the BRM module and directly executes the policy $\mu(s_t, T; \theta)$ throughout the entire episode, denoted by “pure $\mu(\theta)$ ”. This is in fact the pure locomotion module. As discussed in Sec. 3.4, this is also an improved version of the original policy proposed by Wu et al. [58].

Semantic augmented agent: Comparing to the pure DRL agent, our BRM agents utilizes an extra semantic signals c_t provided by the semantic classifier in addition to the visual input s_t . Hence, we consider a semantic-aware locomotion baseline $\mu_S(\theta_s)$, which is another LSTM DRL agent that takes both s_t and c_t as input (denoted by “aug. $\mu_S(\theta_s)$ ”).

HRL agent with an RNN controller: From a DRL perspective, our BRM agent is a hierarchical reinforcement learning (HRL) agent with the BRM module as a high-level controller producing sub-goals and the locomotion module as a low-level policy for control. Note that update and planning on BRM only depend on (1) the current semantic signal c_t , (2) the target T , and (3) the accumulative bit-OR feature B (see Sec. 3.3). Hence, we adopt the same locomotion $\mu(s_t, g; \theta)$ used by our BRM agent, and train an LSTM con-

troller with 50 hidden units on $\mathcal{E}_{\text{train}}$ that takes all the necessary semantic information and produces a sub-target every N steps as well. The only difference between our BRM agent and this HRL agent is the *representation* of the controller (memory) module. The LSTM controller has access to exactly the same semantic information as BRM and uses a much more complex and generic neural model instead of a relation graph. Thus, we expect it to be a strong baseline and perform competitively to our BRM agent.

4.3. Qualitative Analysis and Case Study

In this section, we qualitatively illustrate that our BRM agent is able to learn reasonable semantic relations and behave in an interpretable manner.

Prior knowledge: Fig. 5 visualizes $P(z|\psi^{\text{prior}})$, the learned prior probability of relations, for 3 room types with their most and least likely nearby (connected) rooms. Darker red means more likely while darker blue implies less likely. The captured knowledge is indeed reasonable: bathroom is likely to connect to a bedroom; kitchen is often near a dining room while garage is typically outdoor.

Effectiveness of planning: The BRM agent can effectively decompose a long-term goal into a sequence of easier sub-goals via graph planning. Fig. 6 visualizes the test success rates of the BRM agent (BRM), random policy (“random”) and the pure locomotion policy (“pure $\mu(\theta)$ ”) for increasingly further targets over a fixed set of 5689 randomly generated test tasks. The x-axis is the shortest distance to the target in meters⁴. As expected, when the target becomes more distant, all methods have lower success rates. However, as opposed to the pure locomotion policy, which quickly degenerates to random as the distance increases, the BRM agent remains a much higher success rate in general.

Case study: Fig. 7 shows a successful trajectory by the BRM agent, where the final target is to get out of the house. We visualize the progression of the episode, describe the plans and show the updated graph during exploration. Note that the final goal is invisible to the agent initially (frame ①) but the agent is able to plan, effectively explore the house (e.g., without ever entering the bottom-right dining room region), and eventually reach the target.

4.4. Quantitative Generalization Performances

We evaluate the generalization performances of different approaches on $\mathcal{E}_{\text{test}}$ with horizons $H = 300$ and $H = 1000$. We set $N = 10$, i.e., memory updated every 10 steps. Tab. 1 reports both success rates (% , percent) and SPL values (% , per mile) over 5689 fixed test tasks. In addition to the overall performances, we also report the results under different *planning distances*, i.e., the shortest sequence of sub-goals on the ground-truth relation graph. For an accurate measurement, we ensure that there are at least 500 test tasks for each planning distance.

⁴Typically one meter in shortest distance requires 3 to 4 actions.

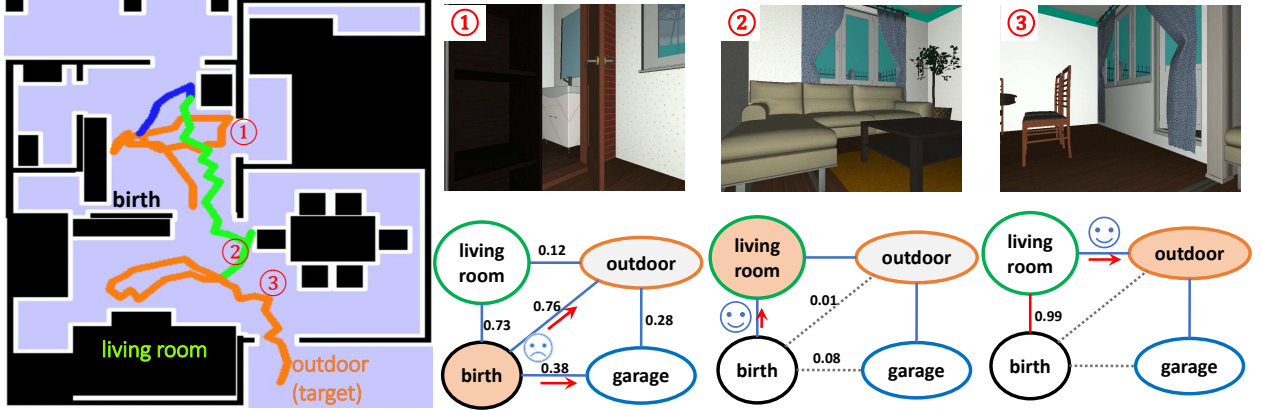


Figure 7. Example of a successful trajectory. The agent is spawned inside the house, targeting “outdoor”. **Left**: the 2D top-down map with goal-conditioned trajectories (“outdoor” – orange; “garage” – blue; “living room” – green); **Right, 1st row**: RGB visual image; **Right, 2nd row**: the posterior of the semantic graph and the proposed sub-goals (red arrow). Initially, the agent starts by executing the locomotion for “outdoor” and then “garage” according to the prior knowledge (**1st graph**), but both fail (top orange and blue trajectories in the map). After updating its belief that garage and outdoor are not nearby (grey edges in the **2nd graph**), it then executes locomotion for “living room” with success (red arrow in the **2nd graph**, green trajectory). Finally, it executes sub-policy for “outdoor” again, explores the living room and reaches the goal (**3rd graph**, bottom orange trajectory).

Our BRM agent has the highest average success rates as well as the best SPL values in *all* the cases. Notably, the margin in SPL is much more significant than that in pure success rate. More importantly, as the horizon increases, i.e., the larger number of planning computations H/N allowed, the overall performance margin (rightmost column) of BRM over the best remaining baseline strictly increases, thanks to the effectiveness of planning.

4.5. Ablation Study

In this section, we show the necessity of all the BRM components and the direction for future improvement.

Benefits of Learned Prior: In BRM, the prior $P(z|\psi^{\text{prior}})$ is learned from training houses. Tab. 2 evaluates BRM with an *uninformative prior* (“unif.”), i.e. $\psi_{i,j}^{\text{prior}} = 0.5$. Generally, the learned prior leads to better success rates and SPL values. Notably, when horizon becomes longer, the gap becomes much smaller, since the graph will converge to the true posterior with more memory updates.

Source of Error: Our approach has two modules, a BRM module for planning and a locomotion module for control. We study the errors caused by each of these components by introducing (1) a hand-designed (imperfect) oracle locomotion, which automatically gets closer to nearby targets, and (2) an optimal planner using House3D labels. The evaluation results are shown in Tab. 3, where the oracle locomotion drastically boosts the results while the BRM performance is close to the optimal planner. This indicates that the error is mainly from locomotion – it is extremely challenging to learn a single neural navigator, which motivates our work to decompose a long-term task into sub-tasks.

Choice of N : The oracle steps for reaching a nearby, i.e., 1-plan-step, target is around 12.27 (top in Tab. 1), so we

choose a slightly smaller value $N = 10$ as the re-planning step size. We investigate other choices of N in Tab. 4. Larger N results in more significant performance drops. Notably, our BRM agent consistently outperforms RNN controller under different parameter choices.

4.6. Evaluation with Terminate Action

In the previous studies, the success of an episode is determined by the House3D environment automatically. It is suggested by Anderson et al. [2] that a real-world agent should be aware of its goal and determine whether to stop by itself. In this section, we evaluate the BRM agent and all previous baselines under this setting: a success will be counted only if the agent terminates the episode correctly in a target room on its own.

There are two ways to include the terminate action: (1) expand the action space with an extra stop action; (2) train a separate termination checker. We observe (2) leads to much better practical performances, which is also reported by Pathak et al. [43]. In our experiments, we simply use the semantic classifier as our termination checker.

The results are summarized in Tab. 5, where we use a long horizon $H = 1000$ to allow the agents to have enough time to self-terminate. Similarly, BRM achieves the best performance in both success rate and SPL metric.

4.7. Discussions

Success rate and SPL: In Tab. 1, the SPL values are typically much smaller than the success rates, namely BRM uses much more steps than the reference shortest path. This is not surprising due to the strong partial observability in the RoomNav task. As a concrete example in Fig. 7, the optimal path from birthplace (near ①) to the outdoor (near

opt. plan-steps	1	2	3	4	5	overall
avg. oracle steps	12.27	42.53	61.09	72.47	63.74	46.86
Horizon $H = 300$						
random	20.5 / 15.9	6.9 / 16.7	3.8 / 10.7	1.6 / 4.2	3.0 / 8.8	7.2 / 13.6
pure $\mu(\theta)$	49.4 / 47.6	11.8 / 27.6	2.0 / 4.8	2.6 / 10.8	4.2 / 13.2	13.1 / 22.9
aug. $\mu_S(\theta)$	47.8 / 45.3	11.4 / 23.1	3.0 / 7.8	3.4 / 8.1	4.4 / 11.2	13.0 / 20.5
RNN control.	55.0 / 49.8	20.0 / 40.8	8.0 / 20.1	5.2 / 15.2	11.0 / 25.2	19.9 / 34.2
BRM	57.8 / 65.4	24.4 / 54.3	10.5 / 28.3	5.8 / 18.6	11.2 / 29.8	23.1 / 45.3
Horizon $H = 1000$						
plan-dist	1	2	3	4	5	avg.
random	24.3 / 17.6	13.5 / 20.3	9.1 / 14.3	8.0 / 9.3	7.0 / 11.5	13.0 / 17.0
pure $\mu(\theta)$	60.8 / 47.6	23.3 / 27.6	7.6 / 4.8	8.2 / 10.8	11.0 / 13.2	22.5 / 22.9
aug. $\mu_S(\theta)$	61.3 / 50.1	23.0 / 26.2	9.4 / 12.0	5.8 / 9.6	9.0 / 13.6	22.4 / 23.8
RNN control.	71.0 / 58.0	39.6 / 51.3	24.1 / 32.7	16.6 / 25.6	23.2 / 39.6	37.0 / 45.2
BRM	73.7 / 74.9	43.6 / 66.0	29.2 / 44.9	20.4 / 27.1	28.4 / 42.5	41.1 / 57.5

Table 1. Metrics of **Success Rate(%) / SPL(%, per mile)** evaluating the generalization performances of BRM and all the baseline approaches (Sec. 4.4). Here $N = 10$. “plan-steps” denotes the shortest planning distance in the ground truth relation graph. “oracle steps” denotes the reference shortest steps required to reach the goal. Our BRM agents have the highest success rates and the best SPL values in all the cases. More importantly, as the horizon increases, which allows more planning, BRM outperforms the baselines more.

BRM ($H=300$)	unif. ($H=300$)	BRM ($H=1k$)	unif. ($H=1k$)
23.1 / 45.3	20.9 / 39.4	41.1 / 57.5	40.4 / 56.6

Table 2. **Success Rate(%) / SPL(%)**: performances of BRM with learned and uninformative prior. $N = 10$, $H = 300, 1000$ (“1k”).

Src. of Err.	LSTM locomotion	oracle locomotion
BRM	41.1 / 57.5	88.6 / N.A.
opt. plan	46.3 / 62.5	96.7 / N.A.

Table 3. **Success Rate(%) / SPL(%)**: performances with an optimal planner and a hand-designed locomotion. $H=1000$, $N=10$.

Choice of N	$N = 10$	$N = 30$	$N = 50$
BRM	41.1 / 57.5	29.7 / 35.2	27.4 / 32.2
RNN control.	37.0 / 45.2	28.2 / 27.7	26.5 / 26.7

Table 4. **Success Rate(%) / SPL(%)**: performances with different choices of N under horizon $H = 1000$.

Horizon $H = 1000$ with Terminate Action				
random	pure $\mu(\theta)$	aug. $\mu_S(\theta_s)$	RNN cont.	BRM
1.8/1.2	8.6/9.0	8.2/8.8	14.5/16.3	17.3/23.0

Table 5. **Success Rate(%) / SPL(%)** with **terminate action** evaluating the generalization performances of BRM and baseline agents with horizon $H = 1000$ and $N = 10$. Our BRM agent achieves the best performances under both metrics.

③) is extremely short if we know the top-down view in advance. However, the outdoor region is out of the agent’s sight (frame ①) and the agent has to *explore* the nearby regions before it sees the door towards the outside (frame ③). Planning and updates on BRM helps guide the agent to explore the unknown house more effectively, which helps lead to higher SPL values. But overall the agent still suffers from the fundamental challenge of partial observability.

Pre-selected concepts: In this work, we focus on the memory representation and simply assume we know all the semantic concepts in advance. It is also possible to generalize to unseen concepts by leveraging the word embedding and knowledge graph from NLP community [60]. It is also feasible to directly discover general semantic concepts from $\mathcal{E}_{\text{train}}$ via unsupervised learning by leveraging the rich semantic information (e.g., object categories, room types, etc) within the visual input. We leave this to our future work.

5. Conclusion

In this work, we proposed a novel design of memory architecture, Bayesian Relation Memory (BRM), for the semantic navigation task. BRM uses a semantic classifier to extract semantic labels from visual input and builds a probabilistic relation graph over the semantic concepts, which allows representing prior reachability knowledge via the edge priors, fast test-time adaptation via edge posteriors and efficient planning via graph search. Our BRM navigation agent uses BRM to produce a sub-goal for the locomotion policy to reach. Experiment results show that the BRM representation is effective and crucial for a visual navigation agent to generalize better in unseen environments.

At a high-level, our approach is general and can be applied to other tasks with semantic context information or state abstractions available to build a graph over, such as robotics manipulations where semantic concepts can be abstract states of robot arms and object locations, or video games where we can plan on semantic signals such as the game status or current resources. In future work, it is also worthwhile to investigate how to extract relations and concepts directly from training environments automatically.

References

- [1] Ankesh Anand, Eugene Belilovsky, Kyle Kastner, Hugo Larochelle, and Aaron Courville. Blindfold baselines for embodied QA. *arXiv preprint arXiv:1811.05013*, 2018.
- [2] 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.
- [3] 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.
- [4] Francisco Bonin-Font, Alberto Ortiz, and Gabriel Oliver. Visual navigation for mobile robots: A survey. *Journal of intelligent and robotic systems*, 53(3):263–296, 2008.
- [5] 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. *International Conference on 3D Vision (3DV)*, 2017.
- [6] Devendra Singh Chaplot, Kanthashree Mysore Sathyendra, Rama Kumar Pasumarthi, Dheeraj Rajagopal, and Ruslan Salakhutdinov. Gated-attention architectures for task-oriented language grounding. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [7] Chenyi Chen, Ari Seff, Alain Kornhauser, and Jianxiong Xiao. Deepdriving: Learning affordance for direct perception in autonomous driving. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [8] Kevin Chen, Juan Pablo de Vicente, Gabriel Sepulveda, Fei Xia, Alvaro Soto, Marynel Vazquez, and Silvio Savarese. A behavioral approach to visual navigation with graph localization networks. *arXiv preprint arXiv:1903.00445*, 2019.
- [9] Xinlei Chen, Li-Jia Li, Li Fei-Fei, and Abhinav Gupta. Iterative visual reasoning beyond convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7239–7248, 2018.
- [10] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied Question Answering. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [11] Abhishek Das, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Neural modular control for embodied question answering. In *Conference on Robot Learning*, pages 53–62, 2018.
- [12] Peter Dayan and Geoffrey E Hinton. Feudal reinforcement learning. In *Advances in neural information processing systems*, pages 271–278, 1993.
- [13] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*, pages 1–16, 2017.
- [14] Kenji Doya, Kazuyuki Samejima, Ken-ichi Katagiri, and Mitsuo Kawato. Multiple model-based reinforcement learning. *Neural computation*, 14(6):1347–1369, 2002.
- [15] Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. RL²: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.
- [16] Hugh Durrant-Whyte and Tim Bailey. Simultaneous localization and mapping: part i. *IEEE robotics & automation magazine*, 13(2):99–110, 2006.
- [17] Alberto Elfes. Sonar-based real-world mapping and navigation. *IEEE Journal on Robotics and Automation*, 3(3):249–265, 1987.
- [18] Kuan Fang, Alexander Toshev, Li Fei-Fei, and Silvio Savarese. Scene memory transformer for embodied agents in long-horizon tasks. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [19] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *International Conference on Machine Learning*, 2017.
- [20] Chelsea Finn, Ian Goodfellow, and Sergey Levine. Unsupervised learning for physical interaction through video prediction. In *Advances in neural information processing systems*, pages 64–72, 2016.
- [21] Chelsea Finn and Sergey Levine. Deep visual foresight for planning robot motion. In *International Conference on Robotics and Automation*, pages 2786–2793. IEEE, 2017.
- [22] Patrick Foo, William H Warren, Andrew Duchon, and Michael J Tarr. Do humans integrate routes into a cognitive map? map-versus landmark-based navigation of novel shortcuts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(2):195, 2005.
- [23] 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 *Advances in Neural Information Processing Systems*, pages 3318–3329, 2018.
- [24] Saurabh Gupta, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [25] Saurabh Gupta, David Fouhey, Sergey Levine, and Jitendra Malik. Unifying map and landmark based representations for visual navigation. *arXiv preprint arXiv:1712.08125*, 2017.
- [26] Ronghang Hu, Marcus Rohrbach, Jacob Andreas, Trevor Darrell, and Kate Saenko. Modeling relationships in referential expressions with compositional modular networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1115–1124, 2017.
- [27] Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. *International Conference on Learning Representation*, 2017.
- [28] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2901–2910, 2017.

- [29] Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. Image retrieval using scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3668–3678, 2015.
- [30] Eric Kolve, Roozbeh Mottaghi, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. AI2-THOR: An interactive 3D environment for visual AI. *arXiv preprint arXiv:1712.05474*, 2017.
- [31] Tejas D Kulkarni, Karthik Narasimhan, Ardavan Saeedi, and Josh Tenenbaum. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. In *Advances in neural information processing systems*, pages 3675–3683, 2016.
- [32] Thanard Kurutach, Aviv Tamar, Ge Yang, Stuart J Russell, and Pieter Abbeel. Learning plannable representations with causal infogam. In *Advances in Neural Information Processing Systems*, pages 8747–8758, 2018.
- [33] Ruiyu Li, Makarand Tapaswi, Renjie Liao, Jiaya Jia, Raquel Urtasun, and Sanja Fidler. Situation recognition with graph neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4173–4182, 2017.
- [34] Kenneth Marino, Ruslan Salakhutdinov, and Abhinav Gupta. The more you know: Using knowledge graphs for image classification. *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [35] Piotr Mirowski, Matt Grimes, Mateusz Malinowski, Karl Moritz Hermann, Keith Anderson, Denis Teplyashin, Karen Simonyan, Andrew Zisserman, Raia Hadsell, et al. Learning to navigate in cities without a map. *Advances in Neural Information Processing Systems*, pages 2419–2430, 2018.
- [36] Piotr Mirowski, Razvan Pascanu, Fabio Viola, Hubert Soyer, Andy Ballard, Andrea Banino, Misha Denil, Ross Goroshin, Laurent Sifre, Koray Kavukcuoglu, et al. Learning to navigate in complex environments. *International Conference on Learning Representation*, 2017.
- [37] Dmytro Mishkin, Alexey Dosovitskiy, and Vladlen Koltun. Benchmarking classic and learned navigation in complex 3D environments. *arXiv preprint arXiv:1901.10915*, 2019.
- [38] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. *International Conference on Learning Representation*, 2018.
- [39] Dipendra Misra, John Langford, and Yoav Artzi. Mapping instructions and visual observations to actions with reinforcement learning. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1004–1015, 2017.
- [40] 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.
- [41] Junhyuk Oh, Xiaoxiao Guo, Honglak Lee, Richard L Lewis, and Satinder Singh. Action-conditional video prediction using deep networks in atari games. In *Advances in Neural Information Processing Systems*, pages 2863–2871, 2015.
- [42] Emilio Parisotto and Ruslan Salakhutdinov. Neural map: Structured memory for deep reinforcement learning. *International Conference on Learning Representation*, 2018.
- [43] Deepak Pathak, Parsa Mahmoudieh, Guanghao Luo, Pulkit Agrawal, Dian Chen, Yide Shentu, Evan Shelhamer, Jitendra Malik, Alexei A. Efros, and Trevor Darrell. Zero-shot visual imitation. In *International Conference on Learning Representation*, 2018.
- [44] Martin Riedmiller, Roland Hafner, Thomas Lampe, Michael Neunert, Jonas Degraeve, Tom Van de Wiele, Volodymyr Mnih, Nicolas Heess, and Jost Tobias Springenberg. Learning by playing-solving sparse reward tasks from scratch. *International Conference on Machine Learning*, 2018.
- [45] Stéphane Ross and Joelle Pineau. Model-based bayesian reinforcement learning in large structured domains. In *Uncertainty in artificial intelligence: proceedings of the... conference. Conference on Uncertainty in Artificial Intelligence*, volume 2008, page 476. NIH Public Access, 2008.
- [46] Nikolay Savinov, Alexey Dosovitskiy, and Vladlen Koltun. Semi-parametric topological memory for navigation. *International Conference on Learning Representation*, 2018.
- [47] Nikolay Savinov, Anton Raichuk, Damien Vincent, Raphael Marinier, Marc Pollefeys, Timothy Lillicrap, and Sylvain Gelly. Episodic curiosity through reachability. In *International Conference on Learning Representations*, 2019.
- [48] Manolis Savva, Angel X Chang, Alexey Dosovitskiy, Thomas Funkhouser, and Vladlen Koltun. Minos: Multi-modal indoor simulator for navigation in complex environments. *arXiv preprint arXiv:1712.03931*, 2017.
- [49] Richard S Sutton, Doina Precup, and Satinder Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211, 1999.
- [50] Aviv Tamar, Yi Wu, Garrett Thomas, Sergey Levine, and Pieter Abbeel. Value iteration networks. In *Advances in Neural Information Processing Systems*, pages 2154–2162, 2016.
- [51] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. MIT press, 2005.
- [52] Antonio Torralba, Kevin P Murphy, William T Freeman, and Mark A Rubin. Context-based vision system for place and object recognition. In *Proceedings of the Ninth IEEE International Conference on Computer Vision-Volume 2*, page 273. IEEE Computer Society, 2003.
- [53] Ramakrishna Vedantam, Karan Desai, Stefan Lee, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. Probabilistic neural-symbolic models for interpretable visual question answering. In *International Conference on Machine Learning*, 2019.
- [54] Ranxiao Frances Wang and Elizabeth S Spelke. Human spatial representation: Insights from animals. *Trends in cognitive sciences*, 6(9):376–382, 2002.
- [55] 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.
- [56] Xiaolong Wang, Yufei Ye, and Abhinav Gupta. Zero-shot recognition via semantic embeddings and knowledge graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6857–6866, 2018.

- [57] Qi Wu, Peng Wang, Chunhua Shen, Anthony Dick, and Anton van den Hengel. Ask me anything: Free-form visual question answering based on knowledge from external sources. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4622–4630, 2016.
- [58] Yi Wu, Yuxin Wu, Georgia Gkioxari, and Yuandong Tian. Building generalizable agents with a realistic and rich 3D environment. *arXiv preprint arXiv:1801.02209*, 2018.
- [59] Huazhe Xu, Yang Gao, Fisher Yu, and Trevor Darrell. End-to-end learning of driving models from large-scale video datasets. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2174–2182, 2017.
- [60] Wei Yang, Xiaolong Wang, Ali Farhadi, Abhinav Gupta, and Roozbeh Mottaghi. Visual semantic navigation using scene priors. In *International Conference on Learning Representation*, 2019.
- [61] Amy Zhang, Adam Lerer, Sainbayar Sukhbaatar, Rob Fergus, and Arthur Szlam. Composable planning with attributes. In *International Conference on Machine Learning*, 2018.
- [62] Hanwang Zhang, Zawlin Kyaw, Shih-Fu Chang, and Tat-Seng Chua. Visual translation embedding network for visual relation detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5532–5540, 2017.
- [63] 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 *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pages 3357–3364. IEEE, 2017.
- [64] Yukun Zhu, Raquel Urtasun, Ruslan Salakhutdinov, and Sanja Fidler. segdeepm: Exploiting segmentation and context in deep neural networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4703–4711, 2015.