Topological Semantic Graph Memory for Image-Goal Navigation

Nuri Kim, Obin Kwon, Hwiyeon Yoo, Yunho Choi, Jeongho Park, and Songhwai Oh

Department of Electrical and Computer Engineering, ASRI, Seoul National University {firstname.secondname}@rllab.snu.ac.kr, songhwai@snu.ac.kr

Webpage: https://bareblackfoot.github.io/TopologicalSemanticGraphMemory

Abstract: A novel framework is proposed to incrementally collect landmarkbased graph memory and use the collected memory for image goal navigation. Given a target image to search, an embodied robot utilizes semantic memory to find the target in an unknown environment. In this paper, we present a topological semantic graph memory (TSGM), which consists of (1) a graph builder that takes the observed RGB-D image to construct a topological semantic graph, (2) a cross graph mixer module that takes the collected nodes to get contextual information, and (3) a memory decoder that takes the contextual memory as an input to find an action to the target. On the task of an image goal navigation, TSGM significantly outperforms competitive baselines by +5.0-9.0% on the success rate and +7.0-23.5% on SPL, which means that the TSGM finds efficient paths. Additionally, we demonstrate our method on a mobile robot in real-world image goal scenarios.

Keywords: Landmark-Based Navigation, Incremental Topological Memory

1 Introduction

Navigation with rich visual observations has been a critical issue in a variety of embodied agent tasks, such as exploration, image goal navigation, and object goal navigation [1–16]. A crucial ingredient for successful visual navigation is to construct a *memory*, which can represent the structure of the environment along with compact visual features for representing high-dimensional visual inputs. A metric-map memory [5, 17] created with SLAM, and a graph memory [8, 9, 18–20] with nodes and edges are the two standard memory construction approaches for navigation algorithms. Even though navigation systems that use metric maps produce powerful results with exact localization and mapping, it is not practical because the navigation agent is susceptible to sensory noises. The topological map, which represents geometric properties and spatial relations of places in the form of a graph, is proposed to construct a map without accurate mapping. Previous visual navigation methods [9, 18] with topological map exploit image features as nodes and edges connecting the nodes in proximity. Since a node indicates a location, the robot's position can be estimated by the nodes in the topological map. Therefore, the graph can be used to navigate successfully even in a real noisy environment.

Unfortunately, using semantic information in topological graph memory presents several difficulties. The first challenge is to incorporate landmarks into a topological graph. As demonstrated by studies [21–23] that show animals navigate utilizing contextual cues from landmarks, the problem of incorporating landmarks, such as object compositions, is a critical issue in visual navigation. A recent research [15] addresses the problem by making an object graph, which connects objects in a field of view of a directional camera to leverage object features. Although it guides an agent to effective action to discover the target utilizing object relationships, it often misses 3D object context information since it only connects objects observed from the same viewpoint. Imagine a goal is given to find a cooking pot, and you know that the cooking pot is usually kept adjacent to an oven. If the oven is nearby but not visible, the agent may miss this crucial information and pick an inefficient path. The more difficult problem is inferring contextual information from objects using geometrically arranged landmarks. There are various advantages to obtaining contextual features from topological graphs. By defining an object through neighboring objects, the contextual representation helps to eliminate the ambiguity of similar but different objects. As shown in Figure 1(a),



Figure 1: **Importance of semantic contextual information.** (a) When two similar items, e.g., cups, are in different locations, they can be identified individually. (b) A kitchen can be recognized when refrigerator, oven, and dining table are detected.

by describing cups in terms of their neighboring objects, it is possible to distinguish two similar cups. For example, a cup in the kitchen can be perceived as one next to a chair and snack box, while a cup in the bathroom can be shown as one that is near to a toothbrush and washstand. Moreover, a place can be better described through objects. A kitchen, for example, can be defined by the presence of a refrigerator, oven, and dining table (see Figure 1(b)).

To this end, this paper addresses the challenges described above using topological semantic graph memory (TSGM), which has two types of nodes: image nodes and object nodes. Image nodes represent different positions that include image representations, while each object node represents a unique object with its visual representation regardless of viewpoint. If a previously undiscovered image node is identified, it is connected to the previously visited image node, producing a spatially meaningful graph. For each image input, objects are detected using MaskRCNN [24] and then connected to image nodes where the object can be seen. A cross graph mixer, a learnable message-passing network that exchanges object and image information, is then used to make the memory contextual. Using this contextual memory, the agent determines the best strategy for finding the goal. We applied our proposed method to image goal navigation in the Gibson [25] to validate our method. As a result, the proposed method significantly outperforms competitive baselines by +5.0-9.0% on the success rate and +7.0-23.5% on SPL, which means that the TSGM finds efficient paths. On a mobile jackal robot, we demonstrate our method in a real-world environment. The experiment is carried out using episodes where the goal location is not visible and is 5 to 10m away. Despite that the episodes are challenging for image goal navigation, TSGM demonstrates successful navigation.

2 Related Work

There has been a lot of research using memory for visual navigation. There are three types of memory formation methods: implicit memory, metric memory, and topological memory.

Implicit memory. Using a recurrent neural network (RNN) as a policy network is a simple method to make an implicit memory [26]. TargetDriven [26] has a vanilla RL policy with a CNN backbone followed by an LSTM which is implicit memory. Since RNN has difficulty backpropagating for a long sequence, RNN is replaced with an explicit memory structure [2, 5, 8, 9, 17, 20].

Metric memory. Active Neural SLAM [17] has a hierarchical structure to explore an environment: global and local policy. The global policy constructs a top-down 2D map and estimates a global goal. It consists of a neural network for flexible output on the input modalities. Given the global goal from the global policy module, a local policy module plans a path to the goal using the simple local navigation algorithm.Exp4nav [5] tackles the exploration problem for navigation. It builds a global metric map, combining egocentric metric maps. Then, to estimate the action, it embeds images with CNN, and a recurrent policy takes the embedding of current observation and the target to output an action.

Topological memory. In [2], LSTM is a policy network that finds the currently located node and derives an action when a topological map is provided. The goal of NeuralPlanner [20] is primarily to explore the environment and maximize coverage. To this end, a topological map is generated from



Figure 2: **Overview of the proposed method.** TSGM has a spatially meaningful structure that is generated and updated online. An image node V_{im} is represented by a blue circle, and an object node V_{ob} is represented by a pink triangle. The current observation x_t is shown in green, while the goal x_g is highlighted in red. A cross graph mixer module is used to update the constructed graph memory to encode the context of features. Then, the attention module chooses the updated memory by querying the current observation and target node. The action policy network determines the action a_t based on the selected memory. Note that the node position is only used for visualization purposes.

an exploratory policy during the rollout. Then, a pretrained neural planner calculates the path to the node most similar to the target image. Semi-parametric topological memory (SPTM) [9] forms a topological memory through exploration before training an agent. Then, an agent navigates to the destination based on the topological memory using the Dijkstra algorithm that plans a path to reach the waypoint. Scene memory transformer (SMT) [8] stacks all the visual features of the past observations as a navigation memory. It uses transformer network [6] to process this memory in the context of the current and target observations. In other words, the authors completely replaced the RNN with an attention mechanism, which shows good performance in long-term tasks. Visual graph memory (VGM) [18] constructs a topological visual memory to navigate an environment while not utilizing the landmark information of the scene. No RL no simulator (NRNS) [27] uses models trained using image input without interaction with the simulator. It creates a node for an unexplored area and finds the goal by estimating the geodesic distance to the goal.

Our method builds upon prior topological memory methods to handle the problem of creating and exploiting the object-based semantic graph to find the most efficient path. Once placed in a new environment, our method explores while incrementally constructing a topological graph, inspired by prior work [8, 18, 27]. In contrast to these previous methods, our method employs spatially meaningful landmarks without using exact positions.

3 TSGM: Topological Semantic Graph Memory

3.1 Problem Statement

We consider the problem of goal-directed exploration for efficient navigation. Given a current observation and a target image to search, an embodied robot aims to find the target (see Figure 2). The current observation contains RGB-D sensory input x_t and detected objects $\{z_1, ..., z_K\} \in \mathbf{z}_t$.

3.2 Graph Memory Construction

Let us consider an agent equipped with an RGB-D camera that was dropped into a novel environment that it has never been in before. We want to build and exploit a graph memory that enables this agent to find the goal in a new environment efficiently. To this end, a semantically structured graph memory, named topological semantic graph memory (TSGM), is built online while an agent traverses the new environment. For example, as shown in Figure 2, when a new node (green circle) is found, the new node is connected to the previous image node (the blue circle connected to the green circle), and the object nodes found at the new node (the green triangles). Since the graph is built with landmarks, recalling the landmarks the agent has seen before will help the agent to discover efficient paths. The TSGM contains two types of nodes and edges, i.e., $\mathcal{G} = \{\mathcal{V}_{im}, \mathcal{V}_{ob}, \mathcal{E}_{im}, \mathcal{E}_c\}$ with image nodes $\mathcal{V}_{im} = \{x_i\}_{i=1}^N$ and its edges \mathcal{E}_{im} with the affinity matrix $A_{im} \in \mathbb{R}^{N \times N}$, object nodes $\mathcal{V}_{ob} = \{z_i\}_{i=1}^M$ where $z_i \in \mathbb{R}^{1 \times D}$, and edges connecting image and object nodes \mathcal{E}_c with the affinity matrix $A_c \in \mathbb{R}^{N \times M}$. N is the number of image nodes, M is the number of object nodes, and D is the dimension of the object representation. The graph expands continuously as the agent explores the environment. The algorithm for constructing TSGM is provided in Appendix B.

Object graph construction. Objects are retrieved using MaskRCNN [24] and added to the graph memory as nodes. When the object is already in memory, the object node is updated with a node with a higher object detection score. If it is not determined that the object is the same, a new object node (z_k) is added to the memory, connecting with the corresponding current image node (x_t) using the affinity matrix (A_c) . To connect objects in proximity in the graph memory, we calculate the affinity matrix for object nodes to be connected to objects in the neighbors of the current image node using the image affinity and image-object affinity. Here, since we want a single object node for an object, the trained object encoder using Supervised contrastive learning [28] learns whether the objects are the same object even when the images are shot from various points of view.

Image graph construction. To construct an image graph, a pretrained similarity encoder [29] learns image similarity in an unsupervised manner and is then used to determine whether or not the node is already in memory. If an observed node is new, i.e., the node feature is not similar to the existing memory, it is added to the graph memory; otherwise, it is utilized to update the corresponding memory node. While connecting image nodes, a new technique based on objects is utilized, assuming that the likelihood of being at the same place is significantly low if there are few co-visible objects between two image nodes. Therefore, the similarity between the object nodes collected at the current location and the object graph nodes is calculated. We consider two image nodes to represent different locations when there are no similar items between them. If the similarity is low, the observed image node is added to memory; otherwise, the image node is updated to the most recent image representation. When it turns out that the image node is new, it is connected to the previous image node (x_l) , making the image affinity matrix (A_{im}) between x_l and x_t connected.

3.3 Cross Graph Mixer

The cross graph mixer is developed from MPNNs [30] to encode scene contexts by combining information from the image and object graphs. The upper row of Figure 3 depicts the image graph, while the lower row represents the object graph. The message passing phase runs for L steps and is defined of message functions M^l and vertex update functions U^l , where $l \in \{1, ..., L\}$. The message function M^l is composed of two functions, the self-update function (S^l) and the cross-update function (C^l) . To begin, image and object nodes self-update to obtain contextual representations of nearby locations or objects, respectively:

$$\hat{mi}_{v}^{l} = \sum_{w \in \mathcal{N}_{i}(v)} S_{i}^{l}(hi_{w}^{l}, A_{im}, g), \ \hat{mo}_{v}^{l} = \sum_{k \in \mathcal{N}_{o}(v)} S_{o}^{l}(ho_{k}^{l}, A_{ob}, g),$$
(1)

where $\mathcal{N}_i(v)$ and $\mathcal{N}_o(v)$ denote the image/object neighbors of the vth node. The node features h are concatenated with the goal feature g, and it is embedded into a feature using two-layered neural networks. Then, the self-update message function S makes \hat{mi}_v^l and \hat{mo}_v^l by aggregating connected nodes with the message passing method, which is illustrated in the left column of Figure 3. After self-update, each graph exchanges information with the others to make a complete the messages,

$$mi_{v}^{l} = \sum_{k \in \mathcal{N}_{o}(v)} C_{i}^{l}(\hat{m}i_{v}^{l}, \hat{m}o_{k}^{l}, A_{c}), \ mo_{v}^{l} = \sum_{w \in \mathcal{N}_{i}(v)} C_{o}^{l}(\hat{m}o_{v}^{l}, \hat{m}i_{w}^{l}, A_{c}),$$
(2)

where C_i and C_o are cross updating functions that produce m_v^l and m_v^l , which are cross-updated messages. The cross-update module's central concept is that images are updated through connected



Figure 3: Cross graph mixer. The upper row shows the image graph update while the lower row shows the object graph update.

objects, and objects are updated using connected images. To cross-update the graphs, the selfupdated image/object features aggregate the connected object/image node features. Then, the update function U transfers messages from object nodes to image nodes and vice versa,

$$hi_v^{l+1} = U_i^l(hi_v^l, mi_v^l), \ ho_v^{l+1} = U_o^l(ho_v^l, mo_v^l).$$
(3)

where U is a two-layered neural network that maps information of an object/image to the connected objects/images. The image-object update iterates for L steps to create semantic contextual node features. After L steps of the image-object update iterations, the cross-mixed memory for image graph **mi** and object graph **mo** are generated for each node. The details of the module structure are described in the Appendix C.

3.4 Memory Attention Module

TSGM is composed of visited image nodes and object nodes that have been observed. For this, an attention network is used to discover the node closest to the goal among the memory. Since this module utilizes attention, interpolation between visited nodes enables goal features to be extracted from unexplored nodes. Additionally, to find the relative path, the current node is better to be searched. The goal feature x_g is given as a query q when extracting a memory-conditioned goal feature, \mathbf{Ci}_g , \mathbf{Co}_g , and the current observation feature x_t is given as an input when selecting the memory-conditioned current feature, \mathbf{Ci}_t , \mathbf{Co}_t . Using the decoder module of the transformer network [6], current contextual feature \mathbf{Ci}_t , \mathbf{Co}_t and goal context feature \mathbf{Ci}_g , \mathbf{Co}_g are collected,

$$\mathbf{Ci} = \sigma \Big(\frac{(\mathbf{W}_q q)(\mathbf{W}_k \{\mathbf{mi}_1, ..., \mathbf{mi}_n\}^T)}{\sqrt{d}} \Big) (\mathbf{W}_v q), \mathbf{Co} = \sigma \Big(\frac{(\mathbf{W}_q q)(\mathbf{W}_k \{\mathbf{mo}_1, ..., \mathbf{mo}_n\}^T)}{\sqrt{d}} \Big) (\mathbf{W}_v q),$$
(4)

where \mathbf{W}_v , \mathbf{W}_q , and \mathbf{W}_k are matrix parameters and the set $\{\mathbf{mi}_1, ..., \mathbf{mi}_N\}$ is the image memory, the set $\{\mathbf{mo}_1, ..., \mathbf{mo}_M\}$ is the object memory, and σ is a softmax function.

3.5 Action Policy Network

Given the contextual feature for goal and current nodes, the action policy network finds an action to reach the goal. For *t*th time step, we encode the sequential information, $\mathbf{h}_a^t = G(\mathbf{h}_c^t, \mathbf{h}_g^t, \mathbf{h}_a^{t-1}, a_{t-1})$, where G is a recurrent neural network and \mathbf{h}_a^t is a *t*th hidden state for action, and a_{t-1} is the previous action of an agent. Then, the action a_t is sampled from the Categorical distribution, i.e. $a_t \sim \text{Categorical}(\mathbf{W}\mathbf{h}_a^t + \mathbf{b})$, where W and b are matrix parameters.

3.6 Optimization

Learning action. Proximal policy optimization (PPO) [31] is used to learn a policy for picture goal navigation. The policy is trained to maximize the expected return, which is defined as the total reward $\mathbb{E}_{\tau}[R(s_t, a_t)]$ over time trajectories $\tau = (s_t, a_t)_{t=1}^H$ of the policy's time horizon H. When an agent takes action a_t at state s_t , $\mathcal{L}_{ppo} = -\mathbb{E}_t \left[\nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \hat{A}_t \right]$, where and \hat{A}_t is an estimator of the advantage function at time step t, and *pitheta* is a stochastic policy. We formulated it as a negative log likelihood to the oracle action, $\mathcal{L}_{act} = -\mathbb{E}_{\tau \sim \mathbb{D}} \left[\sum_{t=1}^{T_{\tau}} a_t^* \log(a_t | x_t) \right]$, where $\tau = (x_t, a_t^*)$ for $t \in \{1, 2, ..., T_{\tau}\}$, and a_t^* is the oracle action at time step t. Since PPO can better estimate return after creating successful episodes, the policy is pretrained with imitation learning and then finetuned with reinforcement learning.

Auxiliary metrics. Knowing when to stop is a significant issue since an episode is considered successful if an agent calls the stop action upon seeing the goal. Fortunately, it is well known that auxiliary metrics, such as a progress monitor and goal sensor, can assist an agent in determining when to stop [32]. The progress monitor assesses progress based on the current observation and a target to determine when to press the stop button, which is obtained by calculating the geodesic distance between an agent and the goal. The goal sensor indicates how close the current position is to the goal. If the target is within the success criteria for image goal navigation, the goal sensor produces one. The progress monitor takes the current observation and a target as input and estimates the progress, \hat{p}_t . The target sensor takes the same input of the progress monitor and estimates that the current observation containing a target, \hat{s}_t . Then, two estimations, \hat{p}_t and \hat{s}_t , are optimized with $L_2 \text{ loss}$, $\mathcal{L}_{aux} = \mathbb{E}_{\tau \sim \mathbb{D}} \left[\sum_{t=1}^{T_{\tau}} ||s_t^* - \hat{s}_t||_2 + ||p_t^* - \hat{p}_t||_2 \right]$, where s_t^* and p_t^* are the ground truths. Since the auxiliary loss is simply added to the action loss, the final loss function for training an agent with imitation learning is $\mathcal{L}_{bc} = \mathcal{L}_{act} + \lambda \mathcal{L}_{aux}$, where λ is a balancing parameter. For reinforcement learning, the combined loss $\mathcal{L}_{rl} = \mathcal{L}_{ppo} + \lambda \mathcal{L}_{aux}$, is used.

4 Experimental Evaluation

4.1 Baselines

We compare our method with baselines with various types of memory.**RGBD + RL** [26] has a vanilla RL policy with a CNN backbone followed by an LSTM adapted from [26]. Active Neural **SLAM** [17] has a metric memory for exploration. To adapt the algorithm to the image goal navigation, we set the output of the global policy to be the relative position to the target when the target is detected using a pretrained target pose estimator. **Exp4nav** [5] tackles the exploration problem.To finetune the image goal navigation task model, we change the coverage reward to the image goal reward. **Neural Planner** [20] is a model adapted the method [20] for image goal navigation from exploration task. **SPTM** [9] creates a topological graph, and then the Dijkstra algorithm then creates a path to a waypoint. **SMT** [8] stacks all the visual features of the past observations and the pose information as a navigation memory. **VGM** [18] builds a topological memory to navigate an environment without using landmark information. **NRNS** [27] trained agent without interaction with the simulator. In order to compare the method fairly, we adapted the method to use a panoramic camera.

4.2 Experiment Settings

We evaluate TSGM in the Gibson environment with a habitat simulator, which is photo-realistic. For image observation, we used a panoramic image, the same setting with [18], inspired by the human neuron not recognizing the heading for localizing [33]. The ground truth objects from the Gibson dataset and detected objects from the detector are used for training. For testing, a detector [24] pretrained with a COCO dataset is used. We use a discrete action space defined as $\mathcal{A} = \{go \text{ forward}, turn \text{ left}, turn \text{ right}, stop\}$, which is a common choice for navigation problems. The step size of the forward action is 0.25m, and the rotation angle is set to 10° in all experiments. The reward function R is defined using the progress of an agent, using a geodesic distance between an agent and the target, and +10 is given when the agent reaches the goal. Further implementation details of the proposed method are provided in Appendix A.

Method	fethod Memory No Pose		Object	Easy		Medium		Hard		Overall	
			~~ J	Success	SPL	Success	SPL	Success	SPL	Success	SPL
RGBD + RL [26]	implicit	×	×	72.5	69.5	53.1	48.6	22.3	17.7	49.3	45.3
Active Neural SLAM [17]	metric	×	×	74.2	20.5	68.4	22.9	29.9	11.0	57.5	18.1
Exp4nav [5]	metric	×	×	70.2	61.8	60.6	52.4	46.9	38.5	59.2	50.9
SMT [8]	graph	×	×	81.9	77.4	65.6	52.2	55.6	39.7	67.7	56.4
Neural Planner [20]	graph	×	×	71.7	41.3	64.7	38.5	42.0	27.0	59.5	35.6
SPTM [9]	graph	 ✓ 	×	66.5	40.6	64.2	38.5	42.1	25.4	57.6	34.8
VGM [18]	graph	 	×	86.1	79.6	81.2	68.2	60.9	45.6	76.1	64.5
TSGM (Ours)	graph	 ✓ 	v	91.1	83.5	82.0	68.1	70.3	50.0	81.1	67.2

Table 1: Comparison of TSGM with memory-based baselines on image goal navigation on Gibson.

Table 2: Comparison of TSGM with image goal navigation baselines on straight/curved episodes on Gibson.

Path Type	Method	Easy		Medium		Hard		Overall	
- uni -5 pe		Success	SPL	Success	SPL	Success	SPL	Success	SPL
	NRNS [27]	67.1	57.8	52.4	41.2	32.6	22.4	50.7	40.5
Straight	VGM [18]	81.0	54.4	82.0	69.9	67.3	54.4	76.7	59.6
	TSGM (Ours)	94.4	92.1	92.6	84.3	70.3	62.8	85.7	79.7
	NRNS [27]	31.7	13.0	29.0	13.6	19.2	10.4	26.6	12.3
Curved	VGM [18]	81.0	45.5	78.8	59.5	62.2	46.9	74.0	50.6
	TSGM (Ours)	93.6	91.0	89.7	77.8	64.2	55.0	82.5	74.1

Evaluation metrics. We use **Success** and **SPL** (success rate over path length) to evaluate the navigation tasks. **Success** is calculated by dividing the total success number by the number of test episodes. We set the success criterion as 1m to the goal, a common criterion for image goal navigation. **SPL** multiplies the success rate and the ratio of the shortest path length and traveled path.

Episode settings. We divide test episodes into three difficulty levels for image goal navigation: easy, medium, and hard. The difficulty is determined by the geodesic length of an episode. Following [17, 18], we set the length as follows: easy (1.5 - 3m), medium (3 - 5m), and hard (5 - 10m).

4.3 Results

TSGM outperforms baselines. Table 1 shows the performance of our TSGM and relevant memory-based baselines on test splits of the Gibson dataset. Our TSGM algorithm outperforms the baseline methods in terms of Success and SPL @ 1m. TSGM improves upon the implicit memory model [26] across splits of Gibson by 31.8% on the success rate. Compared to the VGM [18], our method improved by +5% on the success rate the 4.1% on SPL. To measure performance in various situations, the test episodes are divided according to whether the path is curved or not in Table 2. Our TSGM shows surprising performance improvement in both curved and straight situations. In particular, for an easy episode, the SPL increased quite a surprising amount, 45.5% to 91.0% (100%) for curved episodes and 54.4% to 92.1% (69.3%) for straight episodes compared to the last state-of-the-art result of VGM [18], which can be seen that if an object existing at the target position is present in the field of view, the agent can find the goal position very efficiently. In averaged results in Table 2, TSGM outperforms competitive baselines by 8.5-9.0% on the success rate and 7.0-23.5% improvement on SPL.

TSGM finds efficient paths. We checked that the SPL, which indicates the efficiency of the path, is improved a lot. To investigate how the object graph assists the agent in choosing efficient paths, we visualized the paths in Figure 4. When an object graph is not provided, an agent passes the goal without realizing it is a goal, while an agent with an object graph quickly understands that the goal has been achieved and terminates the search (see Figure 4(a)). In Figure 4(b), the agent enters a deadlock when the object graph is not provided. In the case of a TSGM with an object graph, on the other hand, the goal is successfully found with a more efficient path.



Figure 4: **Visualization of path changes when object graph is not given.** Our TSGM is compared with the method that does not have the object graph. The red flag highlights the image goals. The lines with an arrow represent trajectories of the agent with the color changes with time. (a) shows that an inefficient path is created when the object graph is not given. (b) The agent is trapped in the deadlock states without the object graph.



Figure 5: **Visualization of the robot path on real environment.** The color gradation represents the flow of time. The blue arrow symbolizes the starting point, and the red flag indicates the destination. An image node is represented by a blue circle, and an object node is represented by a pink triangle.

Cross update module is effective. As shown in Table 3, the ablation study on the graph update was performed, where we only used imitation learning results on hard episodes for the ablation experiments. The object update version indicates the update is done only to the image node to the object node, and the visual update version shows the results when the update is the opposite. Using the semantic graph update shows a 4.5% improvement in success rate and +5.3% on SPL. Surprisingly, the updating object node with visual node shows better

Table	3:	Update rules.	
-------	----	---------------	--

Update	Success	SPL
No	0.533	0.393
Visual	0.578	0.446
Object	0.613	0.458
Cross	0.627	0.471

results than updating graph with opposite direction, showing +8.0% on the success rate and +6.5% on SPL compared to the No update version. Finally, the cross update shows the most powerful results, +9.4% on the success rate and +7.8% on SPL. The results indicate that cross updating improves performance by making image and object features informative.

TSGM can handle the real noisy world. Figure 5 shows paths and topological semantic graphs from navigation experiments using a Jackal robot. The episodes examined in the real world are sampled from *hard* episodes with curves. Nonetheless, our agent demonstrated successful navigation by stopping at the correct goal location. The supplementary video shows more paths, with the graph being incrementally built over time.

5 Conclusion and Limitation

The proposed method explores while incrementally generating a semantic topological graph using landmarks such as objects. The core idea of our method is derived from animal behavior, which uses landmarks as a navigational cue. Contextual information is taken from the constructed graph to discover an efficient path to the goal. We demonstrate that TSGM performs efficiently and effectively, with significant performance improvements, particularly in SPL. The proposed method is demonstrated using a Jackal mobile robot to show its effectiveness for practical visual navigation in the real-world.

While it is expected to be suitable for finding objects since object contexts are successfully trained in TSGM, a limitation of the proposed method is that no experiments on object-goal navigation are conducted, due to time constraint. We may adapt TSGM to object-goal navigation in the future.

Acknowledgments

This work was supported by the Institute of Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No. 2019-0-01309, Development of AI Technology for Guidance of a Mobile Robot to its Goal with Uncertain Maps in Indoor/Outdoor Environments)

References

- P. Mirowski, M. Grimes, M. Malinowski, K. M. Hermann, K. Anderson, D. Teplyashin, K. Simonyan, A. Zisserman, R. Hadsell, et al. Learning to Navigate in Cities Without a Map. In *Neural Information Processing Systems (NeurIPS)*, 2018.
- [2] K. Chen, J. P. de Vicente, G. Sepulveda, F. Xia, A. Soto, M. Vázquez, and S. Savarese. A Behavioral Approach to Visual Navigation with Graph Localization Networks. *Robotics: Science* and Systems (RSS), 2019.
- [3] D. Gordon, A. Kadian, D. Parikh, J. Hoffman, and D. Batra. SplitNet: Sim2Sim and Task2Task Transfer for Embodied Visual Navigation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [4] J. Yang, Z. Ren, M. Xu, X. Chen, D. Crandall, D. Parikh, and D. Batra. Embodied Visual Recognition. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2019.
- [5] T. Chen, S. Gupta, and A. Gupta. Learning Exploration Policies for Navigation. International Conference on Learning Representations (ICLR), 2019.
- [6] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention Is All You Need. In *Neural Information Processing Systems (NeurIPS)*, 2017.
- [7] A. Kumar, S. Gupta, D. Fouhey, S. Levine, and J. Malik. Visual Memory for Robust Path Following. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [8] K. Fang, A. Toshev, L. Fei-Fei, and S. Savarese. Scene Memory Transformer for Embodied Agents in Long-Horizon Tasks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [9] N. Savinov, A. Dosovitskiy, and V. Koltun. Semi-Parametric Topological Memory for Navigation. In International Conference on Learning Representations (ICLR), 2018.
- [10] S. Gupta, J. Davidson, S. Levine, R. Sukthankar, and J. Malik. Cognitive mapping and planning for visual navigation. In *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2017.
- [11] E. Parisotto and R. Salakhutdinov. Neural Map: Structured Memory for Deep Reinforcement Learning. *International Conference on Learning Representations (ICLR)*, 2018.
- [12] J. F. Henriques and A. Vedaldi. MapNet: An Allocentric Spatial Memory for Mapping Environments. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018.
- [13] G. Avraham, Y. Zuo, T. Dharmasiri, and T. Drummond. EMPNet: Neural Localisation and Mapping using Embedded Memory Points. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [14] Y. Lv, N. Xie, Y. Shi, Z. Wang, and H. T. Shen. Improving Target-driven Visual Navigation with Attention on 3D Spatial Relationships. arXiv preprint arXiv:2005.02153, 2020.
- [15] H. Du, X. Yu, and L. Zheng. Learning Object Relation Graph and Tentative Policy for Visual Navigation. In European Conference on Computer Vision (ECCV), 2020.
- [16] Y. Qiu, A. Pal, and H. I. Christensen. Target driven visual navigation exploiting object relationships. arXiv preprint arXiv:2003.06749, 2020.

- [17] D. S. Chaplot, D. Gandhi, S. Gupta, A. Gupta, and R. Salakhutdinov. Learning to Explore using Active Neural SLAM. *International Conference on Learning Representations (ICLR)*, 2020.
- [18] O. Kwon, N. Kim, Y. Choi, H. Yoo, J. Park, and S. Oh. Visual Graph Memory with Unsupervised Representation for Visual Navigation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [19] D. S. Chaplot, R. Salakhutdinov, A. Gupta, and S. Gupta. Neural Topological SLAM for Visual Navigation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [20] E. Beeching, J. Dibangoye, O. Simonin, and C. Wolf. Learning to Plan with Uncertain Topological Maps. In *European Conference on Computer Vision (ECCV)*, 2020.
- [21] J. A. Pérez-Escobar, O. Kornienko, P. Latuske, L. Kohler, and K. Allen. Visual landmarks sharpen grid cell metric and confer context specificity to neurons of the medial entorhinal cortex. *Elife*, 5:e16937, 2016.
- [22] E. Chan, O. Baumann, M. A. Bellgrove, and J. B. Mattingley. From objects to landmarks: the function of visual location information in spatial navigation. *Frontiers in psychology*, 3:304, 2012.
- [23] T. S. Collett and P. Graham. Animal navigation: path integration, visual landmarks and cognitive maps. *Current Biology*, 14(12):R475–R477, 2004.
- [24] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask R-CNN. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2017.
- [25] F. Xia, A. R. Zamir, Z. He, A. Sax, J. Malik, and S. Savarese. Gibson Env: Real-World Perception for Embodied Agents. In *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2018.
- [26] Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi. Target-driven Visual Navigation in Indoor Scenes using Deep Reinforcement Learning. In *IEEE international conference on robotics and automation (ICRA)*, 2017.
- [27] M. Hahn, D. S. Chaplot, S. Tulsiani, M. Mukadam, J. M. Rehg, and A. Gupta. No RL, No Simulation: Learning to Navigate without Navigating. In *Neural Information Processing Systems* (*NeurIPS*), 2021.
- [28] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan. Supervised contrastive learning. In *Neural Information Processing Systems (NeurIPS)*, 2020.
- [29] J. Li, P. Zhou, C. Xiong, and S. C. Hoi. Prototypical Contrastive Learning of Unsupervised Representations. *International Conference on Learning Representations (ICLR)*, 2021.
- [30] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural Message Passing for Quantum Chemistry. In *International Conference on Machine Learning (ICML)*, pages 1263–1272. PMLR, 2017.
- [31] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal Policy Optimization Algorithms. arXiv preprint arXiv:1707.06347, 2017.
- [32] J. Ye, D. Batra, A. Das, and E. Wijmans. Auxiliary Tasks and Exploration Enable Objectgoal Navigation. In *IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [33] S. N. Weber and H. Sprekeler. Learning place cells, grid cells and invariances with excitatory and inhibitory plasticity. *Elife*, 7:e34560, 2018.
- [34] N. Zioulis, A. Karakottas, D. Zarpalas, and P. Daras. OmniDepth: Dense Depth Estimation for Indoors Spherical Panoramas. In *European Conference on Computer Vision (ECCV)*, 2018.

- [35] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang. Matterport3D: Learning from RGB-D data in indoor environments. *International Conference on 3D Vision (3DV)*, 2017.
- [36] H. Du, X. Yu, and L. Zheng. VTNet: Visual Transformer Network for Object Goal Navigation. In *International Conference on Learning Representations (ICLR)*, 2021.
- [37] R. Fukushima, K. Ota, A. Kanezaki, Y. Sasaki, and Y. Yoshiyasu. Object Memory Transformer for Object Goal Navigation. In *IEEE international conference on robotics and automation (ICRA)*, 2022.

Appendices

Appendix A Training Details and Experimental Settings

Training process. The training process is as follows: 1. The agent builds a graph map, which can be incrementally updated based on the observation. 2. Extract latent contextual representation from memory generated from a graph map using the attention method. 3. Extract the action by putting the attended memory, current observation, and prior action into the action policy function. 4. The action is first optimized with imitation learning with oracle action, and then the agent learns more suited for the environment through reinforcement learning.

Implementation details. Baselines and the proposed method, excluding the methods that do not use training [27], are trained using imitation learning with 14,400 episodes, 200 episodes per scene for 72 training scenes. Then, reinforcement learning is applied for 10M frames to fix erroneous behaviors by interacting with simulation environments. The cross-update is carried out twice for both graphs since we wanted to update nodes using updated nodes while utilizing the least computational costs.

Real robot experiment settings. We implement TSGM on a Clearpath Jackal platform for a real-world setting. Jackal is equipped with a single Ricoh Theta V camera, a 360° RGB sensor. Since this camera collects images with equirectangular projections, the equirectangular image is converted into a panoramic image of the 12 pinhole cameras having 30° FoV which is used for training. Since it was found that getting the exact panoramic depth is challenging, depth was extracted from the pretrained depth estimator, using the Omni-Depth [34] algorithm for the real-world experiments. Omni-Depth [34] estimates depth from panoramic RGB images gathered from indoor datasets such as Matterport dataset [35]. We utilized the same test simulator environment for training and testing the depth estimator since we assumed a very accurate depth input was given. To control the Jackal, we use a PD controller.

Episode settings. Gibson dataset is divided into 72 train scenes and 14 test scenes for image goal navigation. The test episodes in the Gibson dataset are derived from 14 test scenes that do not overlap the training scene. The test episodes for Table 1 of the paper are proposed in VGM [18]. The number of episodes for each difficulty



Figure 6: Clearpath Jackal

is 1,007, totaling 3,021 episodes. Since episodes with similar beginning and end points might be chosen as duplicates if the scene is tiny, the number of test episodes was chosen in proportion to the scene's size.

The test episodes used in Table 2 of the paper are proposed in NRNS [27], which are divided into two types: straight and curved. The ratio of shortest path geodesic-distance to euclidean-distance between the starting and target locations in straight episodes is 1:2, and the rotational difference between the start position and destination is 45° . All other start-goal location pairs are labeled as curved episodes. There are 2,806 straight episodes and 3,000 curved episodes in total. Similar to the episode configuration used in Table 1, there are around 1,000 items divided according to the difficulty level.

Similarity thresholds. Two pretrained encoders are used to determine whether to add or update the graph with a new observation when building TSGM. We adopt PCL [29] to obtain image similarity and Supervised contrastive learning [28] d to train object similarity. We apply the pretrained PCL [29] model and set the image similarity threshold (τ_i) to 0.75. We trained the model [28] for object similarity utilizing provided identity values for objects on Gibson train environments as labels. After training, the object threshold is chosen to distinguish objects with the greatest margin of

error when tested on the Gibson test scene. To evaluate the trained object encoder, we apply clustering metrics, normalized mutual information (NMI) and purity. NMI is calculated by taking the mutual information between the true object identity distribution and the estimated distribution and then normalizing the value by the self-information of the two distributions. Furthermore, purity is defined as the number of majority positives divided by the total number of samples. With 0.839 NMI and 0.878 purity, the trained object encoder performed well in predicting object identifications on the clustering measures. As a result, the object threshold (τ_o) is set as 0.8 for all experiments.

Appendix B Graph Memory Construction

When an agent is placed in an unknown environment, the agent explores the environment to acquire a graph memory. The graph memory is constructed using two pretrained encoders: an image encoder enc_{im} and object encoder enc_{obj} . To improve the graph with common objects, we add a rule that prevents adding a new image node if the ratio of the number of common objects of the two image nodes is larger than 10%,

$$CO(x_m, x_t) = \frac{\#(O(x_m) \cap O(x_t))}{\#(O(x_m) \cup O(x_t))} > 10\%,$$
(5)

where $O(x_m)$ represents the objects connected to the image node x_m . The detailed algorithm for building a topological semantic graph is described in Algorithm 1.

Object node. Object features are retrieved using ROI-align [24] based on the position of the object bounding box. The object features are extracted from the output of the conv4 layer of ResNet18. The object category from the detector is encoded with two-layered neural networks and then concatenated to the object features to make an object node state $z_k = \langle \mathcal{F}_k^o, c_k, r_k \rangle$ for kth object. \mathcal{F}_k^o is the extracted object feature, c_k represents the category of the object, and r_k is the detection score of the object. The number of categories is 80 since we employ a detector that is pretrained on the COCO dataset.

Object node updating rules. The graph builder compares the object's detection scores when the same object node is found. If a new input object's feature is similar to an object in memory and the category is the same, the graph builder determines that the two objects are identical. The graph builder compares the detection scores to increase the detected features' quality. Only when the newly discovered object has a higher score, the graph builder updates the node feature.

Graph connections. While building the graph, the affinity matrices can be calculated based on the edge information,

$$A_{im}[x_v, x_w] = 1 \ \forall (x_v, x_w) \in \mathcal{E}_{im}, A_c[x_v, z_k] = 1 \ \forall (x_v, z_k) \in \mathcal{E}_c.$$
(6)

To make objects in proximity connected in the graph memory, we calculate the affinity matrix for object nodes to be connected to objects in the neighbors of the current image node using the image affinity A_{im} and image-object affinity. The object graph affinity matrix can be calculated using the above matrices,

$$A_{ob} = A_c^T (A_{im} + \mathbf{I}) A_c, \tag{7}$$

where I is an identity matrix, which connects the object nodes that share the same image node.

Appendix C Cross Graph Mixer

The cross graph mixer module runs for L steps and is defined in terms of a message function M and vertex update function U. The message function can be seen as a composition of two functions, $M = C \circ S$, where S is the self-update function, and C is the cross-update function. At first, the image memory state x and object memory state z are the inputs of the network,

$$hi_{v}^{1} = x_{v}, \forall v \in \mathcal{V}_{im}$$

$$ho_{k}^{1} = z_{k}, \forall k \in \mathcal{V}_{ob}.$$
(8)

Algorithm 1: Build Topological Semantic Graph Memory **Input:** Sequence of the input images, \mathcal{B} **Output:** Topological semantic graph, \mathcal{G} 1 Initialize empty vertex set, $\mathcal{V}_{im} \leftarrow \varnothing$, $\mathcal{V}_{ob} \leftarrow \varnothing$ 2 Detect objects from the input image, $\mathbf{z}_1 \leftarrow \text{Detector}(I_1)$ 3 Initialize the last localized image, $I_l \leftarrow I_1$ 4 Initialize the neighboring objects $O_l \leftarrow \mathbf{z}_1$ 5 Let $\mathcal{F}_{l}^{i} \leftarrow enc_{im}(I_{l}); x_{l} \leftarrow \langle \mathcal{F}_{l}^{i} \rangle$ 6 for $I_t \in \mathcal{B}$, each observation in the batch, do Encode image, $\mathcal{F}_t^i \leftarrow enc_{im}(I_t); x_t \leftarrow \langle \mathcal{F}_t^i \rangle$ 7 if current place is where an agent already explored, i.e., $\exists x_m \in \mathcal{V}_{im}$ s.t. 8 $cos_{-}dist(\mathcal{F}_m^i, \mathcal{F}_t^i) \geq \tau_i$ and $CO(x_m, x_t) > 10\%$, then Update the most similar image node in the memory using new feature: 9 10 $k \leftarrow \operatorname{argmax}_m cos_dist(\mathcal{F}_m^i, \mathcal{F}_t^i)$ $\mathcal{F}_k^i \leftarrow \mathcal{F}_t^i$ 11 Connect the kth node to the last localized node: 12 $\mathcal{E}_{im} \leftarrow \mathcal{E}_{im} \cup (x_k, x_l)$ 13 end 14 else if I_t is new, i.e., $cos_dist(\mathcal{F}_l^i, \mathcal{F}_t^i) \leq \tau_i$ or t=0, then 15 Add x_t , which is $\langle \mathcal{F}_t^i \rangle$ to the graph \mathcal{G} : 16 $\mathcal{V}_{im} \leftarrow \mathcal{V}_{im} \cup \{x_t\}$ 17 $\mathcal{E}_{im} \leftarrow \mathcal{E}_{im} \cup (x_t, x_l)$ 18 Update the last localized image node: 19 $I_l \leftarrow I_t$ 20 Detect objects from the input image, $\mathbf{z}_t \leftarrow \text{Detector}(I_t)$ 21 Sort objects by detection scores. 22 for $\forall b_k \in \mathbf{z}_t$, each object in the current observation, do 23 Encode objects, $\mathcal{F}_k^o \leftarrow enc_{obj}(b_k)$ 24 if b_k exists in the neighboring objects O_l , i.e., $\exists z_m \in O_l$ s.t. $cos_dist(\mathcal{F}_m^o, \mathcal{F}_k^o)$ 25 $\geq \tau_o$ and $c_k = c_m$ and $r_k > r_m$ then Update the most similar object node in the O_l using new feature: 26 $k \leftarrow \operatorname{argmax}_m cos_dist(\mathcal{F}_m^o, \mathcal{F}_t^o)$ 27 $\mathcal{F}_k^o \leftarrow \mathcal{F}_t^o$ 28 end 29 else if Input object is new according to the neighboring objects, i.e., 30 $cos_dist(\mathcal{F}_m^o, \mathcal{F}_k^o) \leq \tau_o, \forall b_m \in \hat{\mathbf{b}}_n \text{ or } c_k \neq c_m, \text{ then }$ Add the object state z_k , which is $\langle \mathcal{F}_k^o, c_k, r_k \rangle$ to the graph \mathcal{G} : 31 $\mathcal{V}_{ob} \leftarrow \mathcal{V}_{ob} \cup \{z_k\}$ 32 $\mathcal{E}_c \leftarrow \mathcal{E}_c \cup (x_l, z_k)$ 33 34 end end 35 $O_l \leftarrow \langle \mathbf{z}_t, \text{Neighbor}(\mathbf{z}_t) \rangle$ 36 end 37 38 end 39 $\mathcal{G} \leftarrow \{\mathcal{V}_{im}, \mathcal{V}_{ob}, \mathcal{E}_{im}, \mathcal{E}_c\}$ 40 return G

For *l*th step, where $l \in \{1, ..., L\}$, image and object node states are self-updated to get contextual representations between nearby locations or objects,

$$\hat{mi}_{v}^{l} = \sum_{w \in \mathcal{N}_{i}(v)} S_{i}^{l}(hi_{w}^{l}, A_{im}, g),$$

$$\hat{mo}_{v}^{l} = \sum_{k \in \mathcal{N}_{o}(v)} S_{o}^{l}(ho_{k}^{l}, A_{ob}, g),$$
(9)

where $\mathcal{N}_i(v)$ and $\mathcal{N}_o(v)$ denote the image/object neighbors of the vth node, respectively. The selfupdate message function S makes a d dimensional vector by aggregating connected nodes, i.e.,

$$S_{i}^{l}(hi_{w}, A_{im}, g) = \mathbf{A}_{im}[v, w] \mathbf{B}_{i}^{l}(e_{vw}) f_{i}(hi_{w}^{l} || g)$$

$$S_{o}^{l}(ho_{k}, A_{ob}, g) = \mathbf{A}_{ob}[v, k] \mathbf{B}_{o}^{l}(e_{vk}) f_{o}(ho_{k}^{l} || g),$$
(10)

where || is a concatenation operation and f_i and f_o are two-layered neural networks. Here, $B(e_{vw})$ is a learned connection relationship between vth and wth nodes, i.e., the edge vectors between the nodes, which maps the edge vector e_{vw} to a $d_i \times d_i$ matrix,

$$\mathbf{B}_{i}^{l}(e_{vw}) = \mathbf{v}_{i}^{T} \operatorname{concat}(\mathbf{W}_{i}hi_{v}^{l}, \mathbf{W}_{i}hi_{v}^{l}) \\
\mathbf{B}_{o}^{l}(e_{vk}) = \mathbf{v}_{o}^{T} \operatorname{concat}(\mathbf{W}_{i}ho_{v}^{l}, \mathbf{W}_{o}ho_{k}^{l}),$$
(11)

where $\mathbf{W}_i \in \mathbb{R}^{C \times d_i}$, $\mathbf{v}_i \in \mathbb{R}^{2d_i \times d_i}$, $\mathbf{W}_o \in \mathbb{R}^{C \times d_o}$, $\mathbf{v}_o \in \mathbb{R}^{2d_o \times d_o \times d_o}$ are the matrix parameters, which are applied for all latent nodes.

Then, the nodes aggregate the different types of nodes during the message passing phase to make a complete message mi_v^l and mo_v^l . For this, nodes cross update is done for image nodes to aggregate messages from object nodes, while object nodes use image nodes to create messages,

$$mi_{v}^{l} = \sum_{k \in \mathcal{N}_{o}(v)} C_{i}^{l}(\hat{m}i_{v}^{l}, \hat{m}o_{k}^{l}, A_{c}),$$

$$mo_{v}^{l} = \sum_{w \in \mathcal{N}_{i}(v)} C_{o}^{l}(\hat{m}o_{v}^{l}, \hat{m}i_{w}^{l}, A_{c}),$$

(12)

where

$$C_{i}^{l}(\hat{m}i_{v}^{l}, \hat{m}o_{k}^{l}, A_{c}) = \hat{m}i_{v}^{l} + f_{i}'(\mathbf{A}_{c}[v, k]\mathbf{B}_{i}^{l}(e_{vk})f_{i}^{''}(\hat{m}o_{k}^{l}||g)),$$

$$C_{o}^{l}(\hat{m}o_{v}^{l}, \hat{m}i_{w}^{l}, A_{c}) = \hat{m}o_{v}^{l} + f_{o}'(\mathbf{A}_{c}[w, v]\mathbf{B}_{o}^{l}(e_{wv})f_{o}^{''}(\hat{m}i_{w}^{l}||g)),$$
(13)

where f'_{o}, f'_{i}, f''_{o} and f''_{i} are two-layered neural networks.

Finally, the update function U transfers messages from object nodes to image nodes,

$$hi_{v}^{l+1} = U_{i}^{l}(hi_{v}^{l}, mo_{v}^{l}), ho_{v}^{l+1} = U_{o}^{l}(ho_{v}^{l}, mi_{v}^{l}),$$
(14)

where,

$$U_{i}^{l}(hi_{v}^{l}, mo_{v}^{l+1}) = hi_{v}^{l} + g_{i}(A_{c}[\cdot, v]mo_{v}^{l+1}),$$

$$U_{o}^{l}(ho_{v}^{l}, mi_{v}^{l+1}) = ho_{v}^{l} + g_{o}(A_{c}[v, \cdot]mi_{v}^{l+1}),$$
(15)

where g_i is a neural network that maps hidden states of an object to hidden states of the connected images and g_o is a neural network that maps hidden states of an image to hidden states of the connected objects.

After L iterations, the output of the network become a contextual memory. The output of the image stage is,

$$\mathbf{mi}_{v} = hi_{v}^{L} \ \forall v \in \{1, ..., N\},$$

$$\mathbf{mo}_{k} = ho_{k}^{L} \ \forall k \in \{1, ..., M\}.$$

(16)

Note that $\mathbf{mi} = {\mathbf{mi}_1, ..., \mathbf{mi}_N}$ is the contextual image memory and $\mathbf{mo} = {\mathbf{mo}_1, ..., \mathbf{mo}_M}$ is the contextual object memory.

Appendix D Object-Based Methods

We compared TSGM to navigational methods [36, 37] that employ object information. In VT-Net [36], an image and objects observed by an agent are combined to produce a fused representation of the place. VTNet [36] does not have any explicit memory and only has an implicit RNN memory.



Figure 7: Correlation between the object score and image score. It indicates that the object nodes help image nodes to localize to the current node.

TSGM, on the other hand, can explicitly employ previous knowledge gathered while navigating the environment. VTNet [36] connects object information when they are detected in the same image. Since it combines objects detected in the same image, VTNet [36] can be seen as a method that only connects objects detected in the same image. On the other hand, TSGM can connect neighboring object nodes even if the objects are not detected in the same image, resulting in contextually solid representations.

In Object memory transformer [37], image and object representation are saved in the explicit memory every time step. Since TSGM only puts a new node into a graph memory based on the similarity between memory and current observations for both image and object graphs, it has less redundancy than [37]. [37] utilizes only the preceding T chunks of data are utilized. TSGM, on the other hand, makes use of all graph memory information derived from past exploration of the environment. This is achievable due to TSGM's low redundancy.



Figure 8: Attention scores of the image and object nodes.



Figure 9: Comparison with the method without objects. The stop button can be pressed more precisely through the object context.



Figure 10: Comparison with the method without objects. The agent can utilize object configurations to get the semantic knowledge how to go from the bathroom to living room.

Appendix E Impact of Landmarks

Object helps localization. To demonstrate the influence of object information on performance improvement, we conducted additional experiments using attention values from the memory attention module (Section 3.4). The agent draws the best image memory in the memory attention module given the current observation. We assume that localization is successful if the selected image node, which has the highest attention score, is the closest to the agent location. The success rate of localization is computed using the 1,007 hard episodes, which are used in Table 1 of the manuscript. As a result, localization performance is 38.5% without an object node whereas it is 68.2% with an object node, an improvement of 29.7%.

Object and image are correlated. To demonstrate that incorporating objects assists in localization, we investigate the correlation between image and object nodes. We averaged the obtained attention values across the objects in an image since it contains many objects. The vertical axis in Figure 7 reflects the average attention score of the object nodes connected with the corresponding ranking's image node. Since the attention score tends to diminish as the number of image nodes rises, the image node ranking was employed as the horizontal axis by sorting in descending order using the attention value. As a consequence, the image node selected as the current node demonstrated a substantial association with a higher object node score (Figure 7). The likelihood of being in the adjacent image node increases as the object node score increases, suggesting that the inclusion of object nodes benefited with agent localization. As a consequence, the image node selected as the current node demonstrated a substantial association with a higher object node score. In summary, it is claimed that the utilization of object nodes enables the agent to search the path more efficiently while avoiding getting lost by properly localizing the agent.

Visualization of the attention score. We showed the attention scores of image and object nodes to demonstrate the significance of object features. To use a ground truth detector, we finetuned TSGM using PPO for 2M frames in the Gibson tiny dataset for this experiment. For the comparison experiments, a model without objects was also finetuned in the same setting, using PPO for 2M frames. The categories in the Gibson tiny dataset are the same as those in the COCO dataset; we displayed 24 of them. Besides the illustration, we show the legend that indicates the color of each object category. The image node is blue, whereas the object node is colored differently for each category. The attention values are multiplied by 100 and scaled to a score ranging from 0 to 100 for the best view. In Figure 8, the attention scores are depicted in Figure 8(a) at step 51 and Figure 8(b) at step 59. When the *bench* node is not visible, the nearest image node has a score of 29.5. After the bench node is recognized, the current node's attention score climbs to 41.5. The score of the currently located *bench* and *potted plant* is high, indicating that the current image node's localization accuracy has improved. A model in Figure 9(a) is trained using an object, whereas a model in Figure 9(b) was trained without an object. The initial paths taken by Figure 9(a) and Figure 9(b) are similar, but Figure 9(b) failed since it pressed the stop button at the incorrect place. The model in Figure 9(b) was unsuccessful since it could not localize at the goal. On the other hand, Figure 9(a) successfully hit the stop button precisely at the target location after realizing that the place at time step 28 was not the target position by using the object configuration information. Figure 10 shows the result of the model trained without an object (Figure 10(b)) and the version trained with an object (Figure 10(a)). The goal is successfully reached by Figure 10(a) in 75 steps, but Figure 10(b) demonstrates that it does not exit the bathroom. This issue appears since the agent failed to acquire the semantic relevant knowledge to go from the bathroom or room to the living room.

Appendix F Path and Graph Visualizations

The experimental results on real world are visualized in Figure 11,12,13,14. All episode goals are sampled from 6m to 8m. Since the locations of nodes are not accurate in real world, we collected the node locations for the best view.

Robot paths on the simulator are illustrated in Figure 15,16,17,18,19,20,21,22,23,24,25. In the Figures, the color gradation represents the flow of time. The blue arrow symbolizes the starting point, and the red flag indicates the destination. An image node is represented by a blue circle, and an object

node is represented by a pink triangle. The paths are from test episodes proposed in VGM [18], also used in Table 1. The paths are drawn in two columns in pairs. The path generated by the proposed algorithm is in the left column, and the path obtained by the comparison algorithm is in the right column. The routes provided are more effective and efficient than the comparison algorithm [18] and reach towards the target.



Figure 11: Real robot experimental results from the laboratory environment.



Figure 12: Real robot experimental results from the laboratory environment.



Figure 13: Real robot experimental results from the laboratory environment.



Figure 14: Real robot experimental results from the home environment.



Figure 15: Examples from Gibson's Cantwell environment.



Figure 16: Examples from Gibson's Cantwell environment.



Figure 17: Examples from Gibson's Eastville environment.



Figure 18: Examples from Gibson's Denmark (bottom), Greigville (top left), Ribera (top right) environment.



Figure 19: Examples from Gibson's Mosquito environment.



Figure 20: Examples from Gibson's Mosquito environment.



Figure 21: Examples from Gibson's Mosquito environment.



Figure 22: Examples from Gibson's Pablo (top) and Mosquito (bottom) environment.



Figure 23: Examples from Gibson's Scioto environment.



Figure 24: Examples from Gibson's Scioto environment.



Figure 25: Examples from Gibson's Scioto environment.