

Context-Aware Planning and Environment-Aware Memory for Instruction Following Embodied Agents

Byeonghwi Kim Jinyeon Kim Yuyeong Kim^{1,*} Cheolhong Min Jonghyun Choi[†]
 Yonsei University ¹Gwangju Institute of Science and Technology

{byeonghwikim, jinyeonkim, cheolhong.min, jc}@yonsei.ac.kr yyeongkim@gm.gist.ac.kr

Abstract

Accomplishing household tasks requires to plan step-by-step actions considering the consequences of previous actions. However, the state-of-the-art embodied agents often make mistakes in navigating the environment and interacting with proper objects due to imperfect learning by imitating experts or algorithmic planners without such knowledge. To improve both visual navigation and object interaction, we propose to consider the consequence of taken actions by **CAPEAM** (Context-Aware Planning and Environment-Aware Memory) that incorporates semantic context (e.g., appropriate objects to interact with) in a sequence of actions, and the changed spatial arrangement and states of interacted objects (e.g., location that the object has been moved to) in inferring the subsequent actions. We empirically show that the agent with the proposed CAPEAM achieves state-of-the-art performance in various metrics using a challenging interactive instruction following benchmark in both seen and unseen environments by large margins (up to +10.70% in unseen env.).

1. Introduction

For decades, the research community has been pursuing the goal of building a robotic assistant that can perform everyday tasks through language directives. Recent advancements in computer vision, natural language processing, and embodied AI have led to the development of several benchmarks aimed at encouraging research on various components of such robotic agents. These benchmarks include navigation [2, 7, 8, 24], object interaction [28, 42], and interactive reasoning [11, 16] in visually rich 3D environments [6, 23, 40]. However, for realistic assistants to be built, active research in interactive instruction following [16, 28, 36, 42] has been in progress. This requires agents to navigate, in-

*Work done while YK was an intern at Yonsei University. †: Corresponding author.

Instruction: Place an apple with a plate in a cabinet.

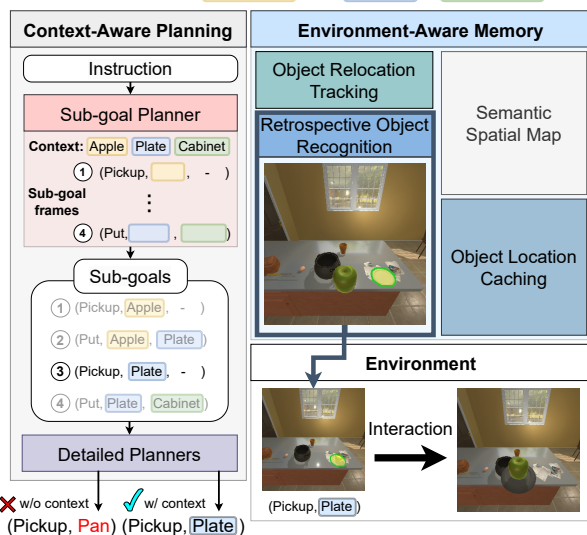


Figure 1: Overview of the proposed ‘Context-Aware Planning (CAP)’ and ‘Environment-Aware Memory (EAM)’. The CAP incorporates ‘context’ (i.e., task-relevant objects) of the task (denoted by ✓ in generating a sequence of sub-goals, compared with the output without the CAP, denoted by ✗). The detailed planners then predict a sequence of agent-executable actions for each respective sub-goal. The agent keeps the state changes of objects and their masks in the EAM and utilizes them when necessary. Even when the agent may not predict the mask of the plate due to occlusion, it can still interact with the plate thanks to the mask remembered in EAM, leading to successful task completion.

teract with objects, and complete long-horizon tasks by following natural language instructions with egocentric vision.

To accomplish a given task, the agent needs to plan a sequence of actions to interact with specific task-relevant objects. However, the agent often plans to interact with irrelevant objects to the task. For instance, for the task “put an apple slice on the table”, after slicing an apple, the agent might plan to pick up a bread slice, which can lead to the failure of the entire task, mainly due to a lack of contextual memory. To address this issue, we first propose a novel approach that

divides the long-horizon planning process into two distinct phases: (1) task-relevant prediction, treated as a *context* prediction, and (2) detailed action planning that considers the contextual memory. We refer to the term ‘context’ as the objects that the command instructs the agent to manipulate. All actions performed by the agent need to focus on these objects, making them the overarching context for the entire plan’s actions. By prioritizing the prediction of the context, we improve the agent’s ability to plan a sequence of actions with less loss of environmental knowledge including objects and their receptacles. We then combine the generated actions with the context to boost the agent’s efficiency in accomplishing long-term objectives by concentrating on interactive objects related to the task.

In addition, changing the object states poses an additional challenge to the agent’s ability to successfully complete tasks that involve object interaction [16, 42]. Failure to track the dynamic object states (*e.g.*, if an object has been already moved or not) can result in unintended interactions and often lead to task failure. For example, for the task “move two apples in the table,” once the agent moves an apple, the agent might try to move the same apple twice if the agent does not know the apple has already been moved and eventually fails at the task.

To address the additional challenge, we further propose to use an environment-aware memory that stores information about the states of objects, as well as their masks for changed visual appearances mainly due to occlusion. This approach allows the agent to interact with objects in their proper states over time. By keeping track of object states and appearances, the agent can ensure interacting with the correct objects and conducting the appropriate actions, ultimately leading to more successful task completion.

For training and evaluation, we use the widely used challenging benchmark for interactive instruction following [36]. We achieve the state-of-the-art success rates and the goal-condition success rates in *seen* and *unseen* environments by large margins (up to +10.70% in unseen SR) and rank the first place in the leaderboard at the moment of submission. Also, CAPEAM with the templated approach with minor engineering won the 1st generalist language grounding agents challenge at the Embodied AI Workshop in CVPR 2023.¹

We summarize our contributions as follows:

- We propose context-aware planning that plans a sub-goal sequence with ‘context’ and conducts respective sub-goals with the corresponding detailed planners.
- We propose environment-aware memory that stores states in spatial memory and object masks for better navigation and interaction with changed object states.

¹See our entry ‘[EAI23] ECLAIR’ in <https://leaderboard.allenai.org/alfred/submissions/public>

- We achieve a state-of-the-art in a challenging interactive instruction following benchmark [36] in all metrics with better generalization to novel environments.

2. Related Work

Action Planning. Embodied AI tasks require agents to reason at multiple levels of abstraction, taking into account the dynamic nature of the environment, their capabilities, and goal formulation [11, 14, 16, 31, 39]. Despite the complexity of tasks, many approaches [30, 32, 37] rely on flat reasoning that directly outputs low-level actions.

Prior arts [5, 10, 12, 16, 41] attempted to address this issue by splitting the layer of actions into two, with the first layer composed of abstract natural language instructions and the second layer of agent-interpretable low actions. However, a significant semantic gap exists between natural language instruction and agent-interpretable actions. Several approaches [24, 25, 32] require large amounts of labeled data or trial-and-error learning to bridge this gap. [4] propose a deep model to reduce the semantic gap issue without requiring excessive labeled data or trial-and-error learning. However, their models often suffer from confusion in predicting the correct objects to interact with.

Meanwhile, the templated approaches [19, 27] have been explored for data-efficient action planning. They rely on pre-designed templates for every task type in the dataset and match each task to the corresponding template. Despite the benefits of efficiency and accuracy, human experts must generate templates whenever a new task type is needed, which is time-consuming and resource-intensive. Additionally, this may not generate optimal plans, as it is restricted to predefined protocols and may not be suitable for tasks outside of the specified protocol. Instead, we propose a method that takes full advantage of the deep learning model to decrease the rigidity of the templated approaches.

For effective planning, it would help agents focus on a task context to parse task-relevant objects from natural language instructions [9, 15]. [9] learn to localize instruction-relevant objects based on implicit language representation. [15] uses the implicit representation of language instructions to guide hallucination of the region outside of the input egocentric map. However, such implicit representation can be affected by language ambiguity (*e.g.*, a red one, the fruit, and apple for ‘Apple’) as the agents are not supervised which they refer to as the same object. In contrast, we explicitly predict task-relevant objects (*e.g.*, Apple, Knife, etc.) that can help the agent focus on the context and thus be robust to variation of language instruction.

Aside from the learning-based approaches for planning, large language models (LLMs) [1, 17, 18, 33] have been actively investigated for their efficacy in reasoning. Despite their significant In [17], LLMs lack physical ground-

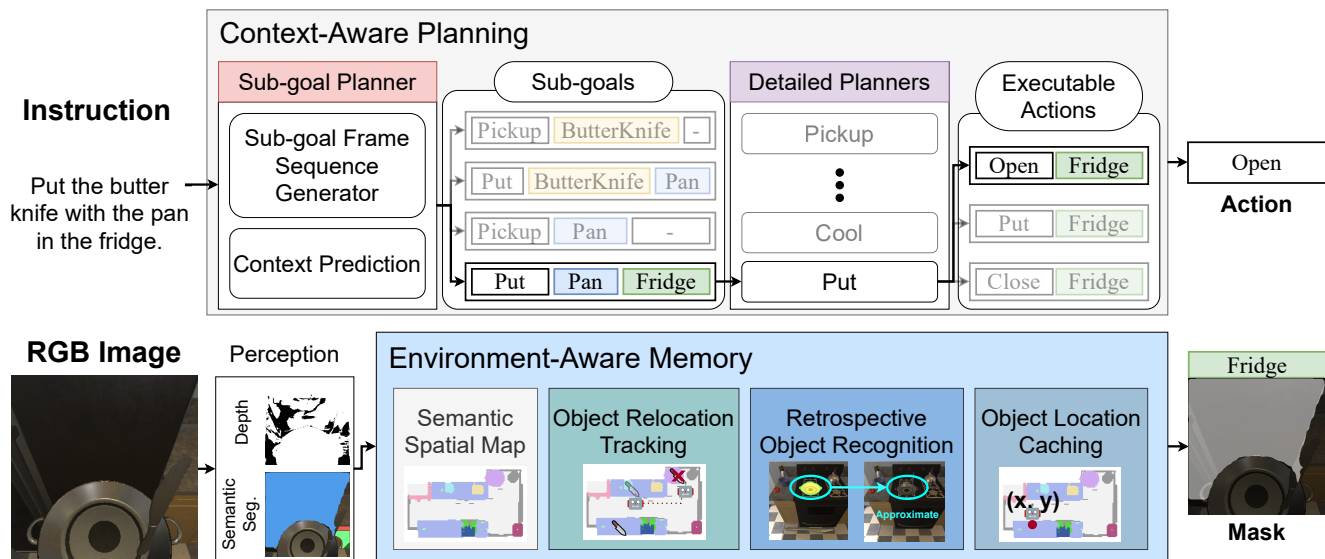


Figure 2: **Model Architecture.** Our agent consists of (1) ‘context-aware planning (CAP)’ and (2) ‘environment-aware memory (EAM)’. Taking the natural language instructions, the sub-goal planner in the CAP predicts ‘context’ (*i.e.*, task-relevant objects) and generates a sequence of ‘sub-goal frames’ that are sub-goals with a predicted action and placeholders for which object should be used with it. Then the objects in the ‘sub-goal frames’ are completed with predicted objects (the context). For each planned sub-goal, a corresponding detailed planner generates a sequence of ‘executable actions.’ In the EAM, the agent maintains the semantic spatial map by integrating the predicted depths and masks into 3D world-coordinates along with the state changes of objects with their masks to utilize them during task completion.

ing due to the absent connection between the physical state of agents and environments (*e.g.*, executable actions, environmental consequences of actions, etc.). This may result in unreasonable interpretations of instructions. For instance, for the task of ‘turn on the lamp while holding a tennis racket,’ the LLM may generate a plausible plan: 1) pick up the tennis racket, 2) plug in the code of the lamp, and 3) turn on the lamp. However, the agent may not support the ‘plug-in’ action, leading to failure at the task. To address the physical grounding issue in LLMs, efforts have been made including reprompting [33], success detectors [18], and skill affordance value functions [1], which implies that sub-goal planning with LLMs still requires further investigation.

Memory for Object Interaction. Utilizing the semantic spatial map makes the process of searching for objects efficient. Previous works [19, 27] record the objects’ positional information on the map when the agent finds key objects during exploration and uses it when the current sub-goal matches one of these records. However, these data do not contain whether the interaction has already been completed or the object should not be moved again. This may cause the agent to re-interact with the object that should not be interacted with. To address this issue, we propose environment-aware memory that keeps track of objects’ information if it has been interacted with, which can reduce misinteraction with unsuitable objects.

In addition to spatial representation encoding, interaction with the same object in multiple time steps is also chal-

lenging as the appearance of an object may change during the interaction (*e.g.*, opened drawer vs. closed drawer) but the agent has to recognize the same object with different appearances. For multiple interactions with the same object, [37] proposes instance association in time (IAT) that introduces a memory to store the previous time step’s mask for an additional mask-selection criterion based on the geological distance of masks to select the object mask closest to the previous time step’s one. However, if the mask generator fails to recognize the object, the IAT then has no mask candidates for selection and therefore the agent is not able to choose the current object’s mask. In contrast, we propose to memorize the previously interacted masks to reduce the effect of the appearance changes mainly due to occlusion.

Semantic Spatial Representation. The earlier strategies to the problem [30, 32, 37] are to directly map visual observations and natural language instructions to a sequence of low-level actions and corresponding masks for object interaction by encoding their history in implicit representation. For instance, [37] proposes to use separate network branches for action and mask prediction and each branch learns to map visual and language features to a sequence of respective actions and masks. [32] proposes a Transformer-based agent that jointly learns to predict actions and masks based on the history of visual observations and language instructions. While they outperform the baseline [36] with large margins, they still lack the ability to complete tasks in unseen environments as observed in the performance gap

from seen environments.

Multiple recent works propose to build an explicit semantic spatial representation such as 2D top-down maps [19, 27], 3D voxel maps [5, 20, 29], or graphs [26]. Such representation enables the agent to accurately perceive 3D environments and plan actions to navigate to and interact with objects on the representation at the expense of additional supervision. Inspired by this, we maintain our agent’s history in a semantic map for room layouts, objects, *etc.*

3. Approach

Generalizing a learned embodied AI agent to the unseen environment is one of the key challenges of building a successful embodied AI agent [5, 19, 27, 29]. To achieve better generalization, recently proposed successful approaches [5, 20, 29] often build an explicit spatial map for the environment and use a hybrid approach of combining learning the environment with well-designed navigation strategies.

Despite being successful, these agents often forget their task contexts and therefore attempt to interact with task-irrelevant objects, eventually leading to task failure. In addition, they might encounter different objects’ states (*e.g.*, appearances, positions, *etc.*) due to object interaction, which may require tracking the states while completing tasks.

To address this issue, we introduce context-aware planning that plans a sequence of actions based on a context (*i.e.*, task-relevant objects) and environment-aware memory that stores states in a spatial memory and object masks for better navigation and interaction with the objects in various states (*e.g.*, an object containing another object inside) in a hybrid model of learning and crafted navigation algorithms. Figure 2 illustrates the architecture of our CAPEAM. We provide details for each component of our method below.

3.1. Context-Aware Planning

Upon receiving a natural language command, an agent needs to interpret and infer the requirements of the given task (*e.g.*, to fetch the object of interest). Once the agent successfully interprets the command, the agent needs to create a plan to achieve the goal. Similar to [4, 5, 41], we propose a novel planning scheme that divides the goal into ‘sub-goals’ and develops each sub-goal into a ‘detailed action sequence’ that the agent can execute.

However, generating a proper sub-goal sequence for the complex task is not trivial as the generated sub-goals often contain irrelevant objects or undesirable actions. Particularly, if there is no proper error-correcting mechanism in place, which is not trivial either, an incorrect sub-goal leads to a failure of the entire task. Even if the agent corrects the error, the corrections take extra steps to complete the task, harming the efficiency of the agent.

For the correct planning with task-relevant objects, we first define ‘context’ as a set of task-relevant objects shared

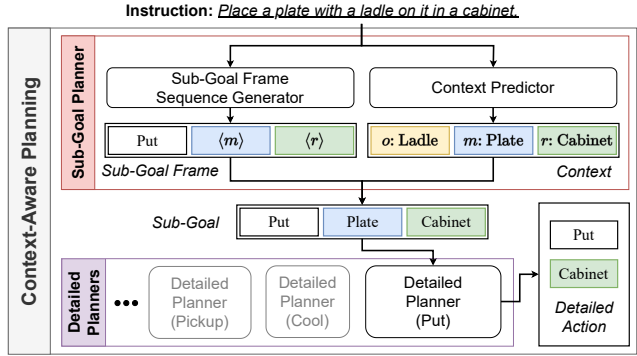


Figure 3: **Context-Aware Planning (CAP)**. It consists of a ‘sub-goal planner’ and a set of ‘detailed planners’ for each sub-goal to generate ‘executable actions.’ The sub-goal planner first predicts a set of objects related to the task, which we call ‘Context.’ Then, the ‘sub-goal frame sequence generator’ in the sub-goal planner generates a sequence of ‘sub-goal frames.’ Finally, the ‘meta-classes’ in each sub-goal frame are replaced with the corresponding objects in the context, resulting in the final sub-goal. A ‘detailed planner’ translates the sub-goal to executable actions.

across sub-goals of a given task. The proposed ‘context-aware planning’ (CAP) divides planning into two phases; 1) a ‘sub-goal planner’ which generates sub-goals, and 2) a ‘detailed planner’ which is responsible for a sequence of detailed actions and objects for interaction for each sub-goal. The sub-goal planner further comprises two sub-modules: the context predictor, which predicts three task-relevant objects, and the sub-goal frame sequence generator, which generates a sequence of sub-goals that do not rely on particular objects, referred to as sub-goal frames.

There are three task-relevant objects referred to as the context. The first object corresponds to the main object to be manipulated. The second object pertains to the container that holds the object. The last object relates to the target object where the object is to be placed in the task. We integrate these predicted task-specific objects into a sequence of sub-goal frames to produce a sub-goal sequence. This allows our agent to plan a sequence of sub-goals conditioned on the task-relevant objects, which helps the agent remember the context, *i.e.*, objects, during action planning.

3.1.1 Sub-Goal Planner

Given a language input l , the sub-goal planner, $f_{sub}(\cdot)$, generates human-interpretable sub-goals. We can write the n^{th} sub-goal as a triplet of action, a small object, and a receptacle that contains the object to be interacted with as:

$$f_{sub}(l) = \{S_n\}_{n=1}^N, \quad (1)$$

$$S_n = (A_n, O_n, R_n),$$

where, A_n denotes a human-interpretable action, such as ‘clean’ or ‘heat’. O_n denotes a small object targeted for ma-

nipulation in the execution of A_n . R_n refers to the location where O_n can be found. N refers to the number of sub-goals in a plan. We have two sub-modules to generate the triplets: the context predictors and the sub-goal frame sequence generator. The context predictors focus on extracting task-relevant objects from natural language instruction. The sub-goal frame generator concentrates on a sequence of interactions that constitute the goal.

We predict task-relevant objects to ensure that all sub-goals in the plan share the same task-relevant object information. By utilizing a ‘meta-class’ to generate sub-goal sequences that identify where these task-relevant objects should be placed. This ensures that all sub-goals share the same task-relevant objects. We refer to a sub-goal filled with the meta-classes as a ‘sub-goal frame’.

Context Prediction. Given a human-described instruction l as an input, we use three context predictors. $f_{ctx}^O(\cdot)$ predicts a primary object being targeted in the task, noted as c_O . $f_{ctx}^M(\cdot)$ predicts a required carrier of the object, noted as c_M . $f_{ctx}^R(\cdot)$ outputs a destination (*i.e.*, receptacle) that contains the target object, referred to as c_R :

$$f_{ctx}^O(l) = c_O, f_{ctx}^M(l) = c_M, f_{ctx}^R(l) = c_R, \quad (2)$$

where c_O is an object assumed to play the main target in a task described as l , c_M , and c_R are a container and destination of c_O , respectively. For instance, if the goal is to “place an apple in a mug on a table”, the agent needs to move the ‘apple’ (c_O) using the container ‘mug’ (c_M) and subsequently place both of them onto the ‘table’ (c_R). To predict the context, we finetune a pretrained language model [13].

Sub-Goal Frame Sequence Generator. After identifying the objects that require manipulation in the task, we determine a series of interactions that will bring objects to the desired goal state. To do this, we generate a sequence of sub-goals *without* specifying the context (*i.e.*, objects) involved. This allows the sub-goals to be later filled with predicted task-relevant objects. This task is accomplished through the sub-goal frame sequence generator.

We introduce a meta-class, which is mapped to one of the corresponding contexts. The use of a meta-class instead of the specific object name allows the sub-goal frame sequence generator to focus more on the object’s role (*i.e.*, a main target, a carrier, and a destination) in the task rather than its own name, as illustrated in Figure 5. If a sub-goal frame includes meta-classes (x_O, x_M, x_R), it is later replaced with the ‘contexts’ (c_O, c_M, c_R) from the context predictors. Note that some contexts may contain ‘None’ indicating that a task does not need them. For instance, for a task, ‘put an apple on the table,’ the carrier (x_M) is not needed and therefore an action sequence is planned based only on the two contexts (*i.e.*, $x_O = \text{Apple}$ and $x_R = \text{Table}$).

The sub-goal frame outputs a sub-goal with a placeholder ($\langle \cdot \rangle$), which is later filled with either an object or a meta-class by the generator. Formally, the sub-goal frame sequence generator $f_{sf}(\cdot)$ takes as input a natural language instruction and generates a sequence of sub-goal frames:

$$\begin{aligned} f_{sf}(l) &= \{F_n\}_{n=1}^N, \\ F_n &= (A_n, \langle O \rangle_n, \langle R \rangle_n), \\ \langle \cdot \rangle &\in E \cup \{x_O, x_M, x_R\}, \end{aligned} \quad (3)$$

where l denotes the instruction, A_n denotes an action of n^{th} sub-goal frame, $\langle O \rangle_n$ and $\langle R \rangle_n$ represent place holder of the small object and receptacle to be interacted with, respectively, E refers to a set of all objects in a given environment, and N refers to the number of sub-goal frames. All meta-class in place holder of sub-goal frames are then replaced with task-relevant objects (c_O, c_M, c_R) from context prediction, resulting in the final output $\{(A_n, O_n, R_n)\}_{n=1}^N$.

3.1.2 Detailed Planners

To execute each sub-goal generated by the sub-goal planner f_{sub} , an agent should render the agent-executable actions which we refer to as ‘detailed actions’ from the inferred ‘sub-goal plan’, $\{S_n\}_{n=1}^N$. We define a ‘detailed planner’ $f_{dp}^g(\cdot)$ for a sub-goal action g to translate a sub-goal with $A_n = g$ into a sequence of the detailed actions as:

$$f_{dp}^g((A_n, O_n, R_n)) = \{(a_t, o_t)\}_{t=1}^{T_n}, \quad (4)$$

where a_t and o_t are the t^{th} detailed action and an object to interact in the given sub-goal. For instance, if a sub-goal is given as (Pickup, Plate, Cabinet), a favorable output of a detailed planner would be (Open, Cabinet), (Pickup, Plate), and (Close, Cabinet). We learn the detailed planner using a self-attention LSTM in a supervised manner [22].

3.2. Environment-Aware Memory

Semantic Spatial Map. One of the challenges in the interactive instruction following task is to accurately perceive the 3D environment from the 2D RGB images captured by the agent for better navigation and interaction with objects. Following [5, 19, 20, 27, 29], we build a semantic spatial map from predicted depth maps and object masks by back-projecting them to 3D world coordinates. In particular, we use depth to maintain the environmental information such as obstacle area, object positions and classes, *etc.*

However, without memory, the agent may not be able to keep track of the configurations of objects in an environment. The absence of such tracking poses an additional challenge for task completion. For instance, if the agent is asked to move multiple apples to a table, the agent should then remember which objects have been moved so far. Here,

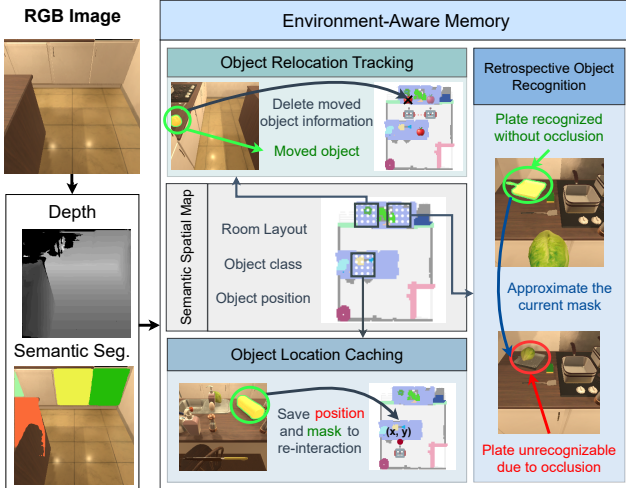


Figure 4: **Environment-Aware Memory (EAM)**. The agent updates the semantic spatial map using predicted depths and object masks for scene information. ‘Retrospective Object Recognition’ preserves the latest object mask to approximate the current object’s mask when mask prediction fails. ‘Object Relocation Tracking’ stores the most recent location of each relocated object and discards it as a future navigation target. ‘Object Location Caching’ remembers the locations and masks of objects whose states change.

we propose to configure memory of past environmental information to remember the past information and predict a proper action sequence during a task. The memory helps the agent remember the current state of the environment and make proper decisions to facilitate task completion.

Retrospective Object Recognition. While completing tasks, the agent often needs to interact with the same object in multiple time steps. For example, to move an apple with a plate, the agent has to interact with the plate twice: 1) put the apple in the plate and 2) pick up the plate with the apple. During multiple object interactions, however, the visual appearance of the object can change due to various reasons such as occlusion. For these reasons, the agent might not be able to recognize the object and would fail at interaction with the intended object, leading to the entire task’s failure.

To address changes in the visual appearance of an object during multiple interactions, we propose to retain the latest segmentation masks of objects and use them as the current object’s mask if the agent is interacting with the same object but fails to recognize it. Exploiting the preserved masks allows the agent to keep interacting with the same object even with visual appearance changes during interactions.

Object Relocation Tracking. In another task completion scenario, the agent often faces scenarios where it needs to relocate multiple objects of the same class to the same location (e.g., “Move two apples on the table.”). While conducting this task, the agent may lose track of which objects

have already been relocated. Consequently, the agent may attempt to navigate to and relocate objects that have already been relocated, which possibly leads to a repeated sequence of unintended object relocation or a task failure.

To circumvent the issues, we propose to maintain the information about the most recent location of each relocated object in a 2D coordinate and exclude it in the semantic map as a future target for our agent’s navigation. The agent can recognize relocated objects among all detected objects by comparing the locations in the memory and the semantic spatial map. This module allows the agent to avoid redundant interaction with already relocated objects, possibly leading to task failure, and therefore successfully navigate to and interact with those that have been not relocated.

Object Location Caching. During another task completion, an agent may need to revisit object locations that it has previously visited to obtain an object with ‘changed states’. Considering the task “put an apple slice on the table” as an example, after the agent has sliced the apple, the agent needs to navigate to the sliced object again at the location where the apple was sliced to interact with the object in its changed state. Without memorizing the locations and masks, however, the agent would need to re-explore the environment to locate the objects again, which could result in inefficient navigation and possible task failure.

To alleviate such an issue, we propose to cache the 2D locations and the segmentation masks in memory for objects whose states change. By preserving the object locations and masks, the agent can efficiently navigate back to the remembered locations and interact with the remembered object masks when necessary. This can reduce the need for the agent to explore the environment again, which can lead to more efficient navigation and interaction and possibly reduce the possibility of navigation and interaction failure.

3.3. Action Policy

To conduct object interaction for task completion, the agent first needs to reach the target objects in close vicinity [31, 34, 39]. Recent approaches [5, 19, 20, 29] plan obstacle-free paths using either deterministic algorithms (e.g., A*, FMM [35], etc.) or learning the path in the discrete navigation space [38], mostly by the imitation learning [3].

Since imitation learning requires a large number of expert trajectories for satisfactory performance, deterministic algorithms currently dominate the literature for significantly better performance, implying the amount of data in the current benchmark with imitation learning might be more limited than necessary. In addition, [19, 27] maintains the obstacle area larger than they actually perceive for safe path planning distant from the obstacles. Inspired by them, we plan navigation paths with the deterministic approach [35] in the discrete space on the expanded obstacle map.

Model	Low Inst.	Tem. Act.	Test Seen		Test Unseen	
			SR	GC	SR	GC
FILM [27]	✗	✓	25.77 (10.39)	36.15 (14.17)	24.46 (9.67)	34.75 (13.13)
Prompter [19]	✗	✓	49.38 (23.47)	55.90 (29.06)	<u>42.64 (19.49)</u>	59.55 (25.00)
CAPEAM*	✗	✓	<u>46.64 (20.81)</u>	<u>55.29 (25.47)</u>	45.72 (20.15)	<u>57.25 (24.73)</u>
FILM [27]	✓	✓	28.83 (11.27)	39.55 (15.59)	27.80 (11.32)	38.52 (15.13)
Prompter [19]	✓	✓	53.23 (25.81)	63.43 (30.72)	<u>45.72 (20.76)</u>	<u>58.76 (26.22)</u>
CAPEAM*	✓	✓	<u>50.62 (22.61)</u>	<u>59.40 (27.49)</u>	49.84 (22.61)	61.10 (27.00)
HL5M [5]	✗	✗	25.11 (6.69)	35.79 (11.53)	16.29 (4.34)	27.24 (8.45)
LGS-RPA [29]	✗	✗	33.01 (16.65)	<u>41.71 (24.49)</u>	27.80 (12.92)	38.55 (20.01)
EPA [26]	✗	✗	<u>39.96 (2.56)</u>	44.14 (3.47)	<u>36.07 (2.92)</u>	<u>39.54 (3.91)</u>
CAPEAM (Ours)	✗	✗	47.36 (19.03)	54.38 (23.78)	43.69 (17.64)	54.66 (22.76)
HL5M [5]	✓	✗	29.94 (8.74)	41.21 (14.58)	20.27 (5.55)	30.31 (9.99)
MAT [20]	✓	✗	33.01 (9.42)	43.65 (14.68)	21.84 (6.13)	32.41 (10.59)
AMSLAM [21]	✓	✗	29.48 (3.28)	40.88 (5.56)	23.48 (2.36)	34.64 (4.63)
LGS-RPA [29]	✓	✗	<u>40.05 (21.28)</u>	<u>48.66 (28.97)</u>	<u>35.41 (22.76)</u>	<u>45.24 (22.76)</u>
CAPEAM (Ours)	✓	✗	51.79 (21.60)	60.50 (25.88)	46.11 (19.45)	57.33 (24.06)
Human	-	-	-	-	91.00 (85.80)	94.50 (87.60)

Table 1: **Comparison with the state of the arts.** The path-length-weighted (PLW) metrics are given in the parentheses for each value. The highest and second-highest values per fold and metric are shown in **bold** and underline, respectively. ‘Low Inst.’ refers to the step-by-step instructions aligned to the respective sub-goals. ‘Tem. Act.’ refers to ‘templated action’ sequences designed in [27]. ✓/✗ denotes the corresponding module is used/not used, respectively. ‘CAPEAM*’ denotes our agents using the templated actions without action planning by our learned model.

#	CAP	EAM	Test Seen		Test Unseen	
			SR	GC	SR	GC
(a)	✓	✓	51.79 (21.60)	60.50 (25.88)	46.11 (19.45)	57.33 (24.06)
(b)	✗	✓	49.45 (20.77)	58.94 (25.43)	42.25 (17.79)	53.93 (22.30)
(c)	✓	✗	46.44 (18.12)	56.37 (23.24)	41.66 (16.03)	52.88 (21.40)
(d)	✗	✗	45.34 (17.55)	56.62 (22.99)	39.83 (15.40)	51.97 (20.15)

Table 2: **Ablation study for each proposed component.** CAP denotes ‘Context-Aware Planning.’ EAM denotes ‘Environmental-aware Memory.’ The ablation of either or both CAP and EAM leads to noticeable performance drops.

4. Experiments

Dataset and Metrics. We employ ALFRED [36] as a challenging interactive instruction following benchmark. There are three splits of environments in ALFRED: ‘train’, ‘validation’, and ‘test’. The validation and test environments are further divided into two folds, *seen* and *unseen*, to assess the generalization capacity. For evaluation, we follow the standard evaluation protocol of the ALFRED benchmark [36]. The primary metric is the success rate, denoted by ‘SR,’ which measures the percentage of completed tasks. Another metric is the goal-condition success rate, denoted by ‘GC,’ which measures the percentage of satisfied goal conditions. Finally, path-length-weighted (PLW) scores penalize SR and GC by the length of the actions that the agent takes. We provide further details of this benchmark in the supplementary for space’s sake.

4.1. Comparison with the State of the Art

We present a quantitative analysis of our method and prior arts in Table 1. For a fair comparison, we compare our method with prior arts that incorporate semantic spatial rep-

#	CAP	Param#	Valid. Seen	Valid. Unseen
(a)	✓	712.60M	83.05	77.34
(b)	✗	651.01M	79.27	74.91
(c)	✗	164.95M	80.12	74.18

Table 3: **Accuracy of the predicted action sequence compared to the ground-truth by the CAP module vs. naïve model size increase.** CAP refers to the ‘Context-Aware Planning.’ Param# denotes the number of parameters to be learned for planning. (b) and (c) share the same LSTM-based architecture of (a) but not use context (*i.e.*, a set of task-relevant objects). CAP noticeably improves the success rate ((a) → (c)) but this is not simply from the model size increase ((c) → (b)).

resentation constructed followed by depth estimation. Following the recent approaches [5, 27], we compare methods that use only a high-level goal statement (*i.e.*, without low-level instructions, denoted by ‘Low Inst.’ by ✗). In addition, we compare the models that generate action sequences using prior knowledge of tasks and environments with the ‘action template’ (‘Tem. Act.’ by ✓) [19, 27].

We first investigate the performance when the hand-designed action sequence templates are combined with our agent (✓ in ‘Tem. Act.’), which is an ablated version of our model. We observe that our agent outperforms all prior arts in novel environments in terms of success rates, which is the main metric of the benchmark. We observe [19] yields better performance in seen environments compared to our agent. We believe that this might be attributed to the strategies to enhance spatial perception such as pose adjustment based on accurate perception models (*e.g.*, depth estimators, *etc.*) that generally perform well in seen environments. Nevertheless, we observe that agent has less performance gap between seen and unseen environments, implying better generalization of our agent to unseen environments.

We then investigate the performance without using the templated action sequences (✗ in ‘Tem. Act.’). We observe that our method outperforms all prior arts by large margins in SR and GC for both seen and unseen environments. As we consistently observe the improvements with and without the low-level instructions, this would imply that our method does not heavily rely on the detailed description of a task. Note that [29] collects an additional dataset with the human-in-the-loop process for interaction failure recovery. With the additional expensive supervision, we observe that [29] achieves better PLW scores by efficiently completing tasks than CAPEAM, and the comparison is not quite fair.

4.2. Ablation Study

To investigate the benefits of each proposed component of our CAPEAM, we conduct a quantitative ablation study and summarize the result in Table 2.

Without Context-Aware Planning. First, we ablate the CAP from our method and the agent therefore learns a

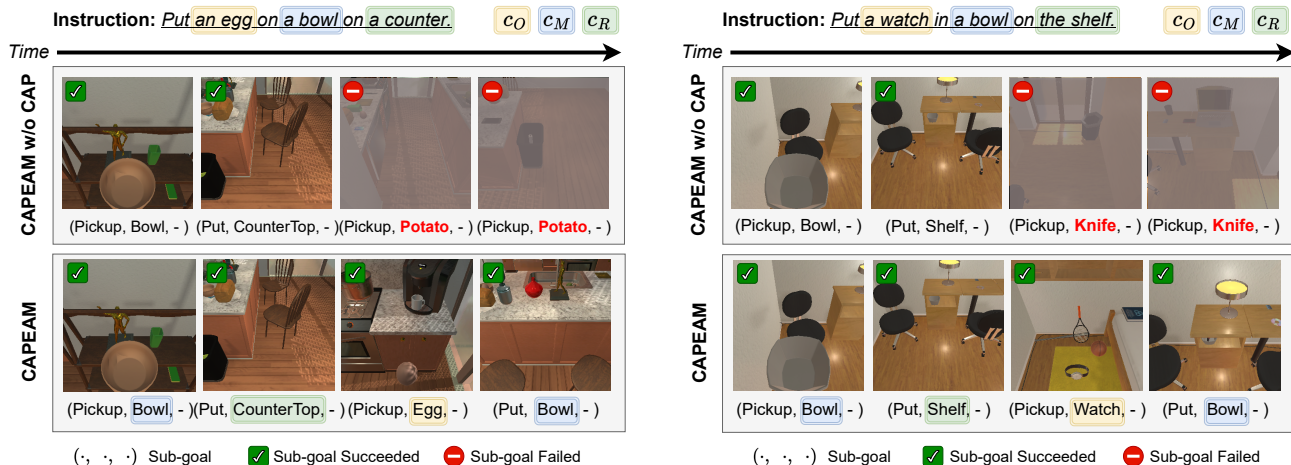


Figure 5: **Benefit of Context-Aware Planning (CAP).** In the two qualitative examples, the ‘contexts’ are denoted by c_O in yellow, c_M in blue, and c_R in green colored boxes. While our CAPEAM plans a sub-goal sequence with task-relevant objects, ‘CAPEAM w/o CAP’ interacts with task-irrelevant objects (*i.e.*, **Potato** or **Knife**) and consequently fails.

monolithic policy that directly maps natural language instructions to a sequence of agent-executable actions. At times, the agent attempts to interact with inappropriate objects to tasks. The result of the task may differ from the intended goal, leading to task failure. This eventually leads to noticeable performance drops (-2.34% , -3.86% in SR in both seen and unseen splits as evidenced in $\#(a)$ vs. $\#(b)$).

Some may argue that the performance gain by the CAP may come from model size increase as we use separate networks for context prediction and sub-goal frame sequence generation. For this, we learn the same policy but with more learnable parameters such as the size of hidden states and embeddings to closely approach to the size of the final model and summarize the result in Table 3. We observe that even with the similar number of parameters (651.01M compared to 712.60M), the planning accuracy of the policy with more parameters remains similar or even slightly drops in the seen split from the one with the smaller network (164.95M). It implies that the performance gain by the CAP may not be attributed to simple model size increase.

Without Environment-Aware Memory. We then ablate the EAM from our agent. Without EAM, the agent has to conduct its actions based on the current state as the agent preserves only limited information about the changed states of objects and their masks. Due to the lack of environmental information, the agent may perform undesired actions (*e.g.*, move an already relocated object) and thus fail. We observe significant performance drops (-5.35% and -4.45% in SR) in both seen and unseen environments ($\#(a)$ vs. $\#(c)$).

Without Both. Without any of the proposed components, the agent may interact with irrelevant objects with limited past environmental information. As expected, our agent

without CAP and EAM achieves the lowest performance among the agents equipped with either or both ($\#(d)$ vs. $\#(a, b, c)$). Moreover, we observe that using both CAP and EAM improves our agent more than using either of them ($\#(d) \rightarrow \#(b, c)$ vs. $\#(d) \rightarrow \#(a)$). This implies that the CAP and the EAM are complementary to each other.

4.3. Qualitative Analysis

Context-Aware Planning. To illustrate the benefit of context-aware planning (CAP), we present two qualitative examples in Figure 5. The left example shows that a sub-goal planner without the CAP may generate sub-goals including irrelevant objects (*i.e.*, **Potato**), even when it is not mentioned in the given instruction. The agent continues searching for a knife for the limited number of steps and fails. But when we use the context predictor, the sub-goal planner correctly infer the task-relevant objects and constructs the sequence with them (Sec. 3.1); the context predictor outputs task-relevant objects from the given human-described instruction such as an ‘egg,’ denoted as o_O , a ‘bowl’, the container to hold the object as o_M , and a ‘counter’, the place to put the object as o_R . The task-relevant objects that is correctly inferred by our model, the sub-goal planner can generate a desirable sub-goal sequence and leads to successful task completion.

In addition, the right example in the Figure 5 shows that our agent without CAP may predict irrelevant objects (*i.e.*, Knife) to the task (*i.e.*, “Put a watch in a bowl on the shelf.”). As the agent without CAP, denoted by ‘CAPEAM w/o CAP,’ tries to find the unintended object (*i.e.*, Knife) which is not present in this room, the agent continues to explore the environment and eventually fails to reach the target object. In contrast, our agent with CAP, denoted by ‘CAPEAM,’ can generate a sequence of executable actions

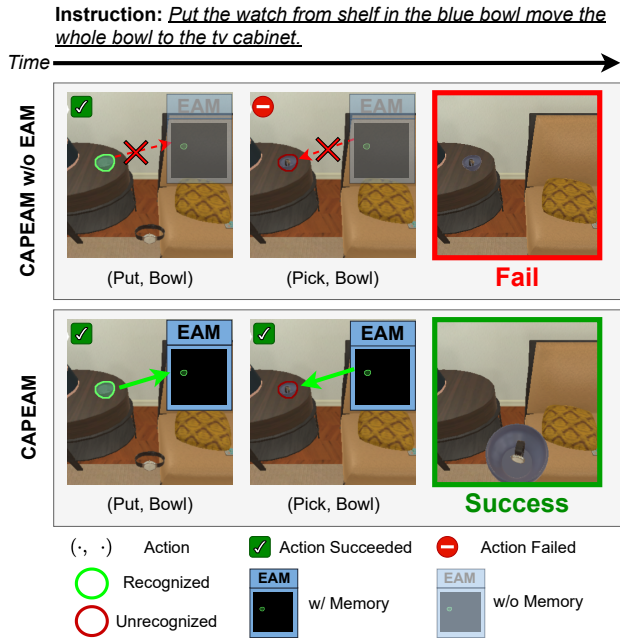


Figure 6: **Benefit of the ‘Retrospective Object Recognition’ in the Environment-Aware Memory (EAM).** → indicates the agent can preserve the object masks in EAM and utilize them while ✗ cannot. ‘CAPEAM w/o EAM’ fails in interaction since it cannot recognize the bowl. In contrast, ‘CAPEAM’ can exploit the preserved bowl’s mask and therefore succeeds in the task.

with relevant objects. By conducting all the actions with the intended objects, the agent finally succeeds in the task.

Environment-Aware Memory. We now conduct a qualitative analysis to assess the impact of retrospective object recognition (Sec. 3.2). It allows the agent to continue interacting by utilizing a previously saved mask even when it cannot recognize the object.

First, we investigate the benefit of the ‘Retrospective Object Recognition’ module in the EAM. Prior arts [4,27] may miss interactions with unrecognized objects because they have limited access to past information and typically rely on current information during the interaction. For instance, in Figure 6, the memoryless model is incapable of detecting the bowl’s mask because of occlusion caused by the watch placed on top of it when it tries to pick up the bowl. Thus, they may initially encounter challenges with interaction, given their incapability to obtain a mask for interaction. On the contrary, CAPEAM with the memory stores the bowl’s mask from the previous action and uses it for interaction. Consequently, even when the agent cannot perceive the bowl, it can interact with it by the saved mask.

Finally, we investigate the benefit of the ‘Object Relocation Tracking’ module in the EAM. Figure 7 illustrates its benefit of keeping track of the relocated objects and preventing re-relocation of the same objects (Sec. 3.2). As observed in ‘CAPEAM’ in the figure, after relocating a tar-

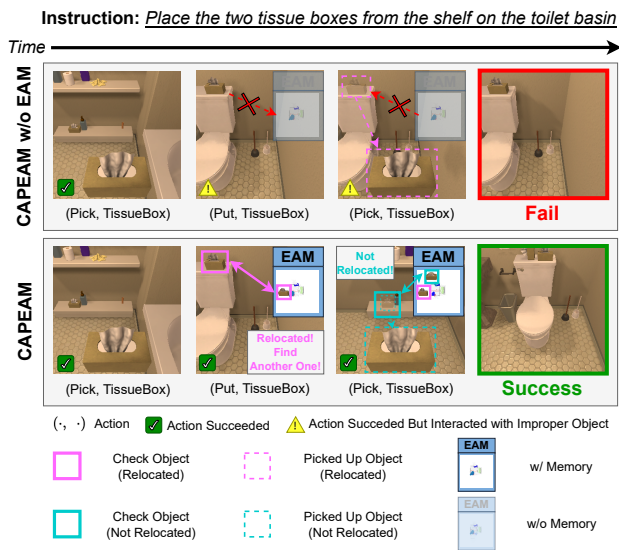


Figure 7: **Benefit of the ‘Object Relocation Tracking’ in the EAM.** While our agent without EAM (‘CAPEAM w/o EAM’) interacts with the already relocated object (‘TissueBox’) as it does not keep track of the relocated location (denoted by the dashed → and ✗), EAM (‘CAPEAM’) allows to avoid interacting with the already relocated one and therefore interact with the intended one.

get object (*i.e.*, ‘TissueBox’), the agent remembers the relocated object’s location and avoids interacting with that already relocated tissue box. On the contrary, as observed in ‘CAPEAM w/o EAM,’ our agent without EAM does not keep track of the relocated object’s location. Thus, it interacts with the already relocated object again and eventually, this leads to task failure.

5. Conclusion

We propose CAPEAM that incorporates *context* in planning and to remember environmental changes in memory for embodied AI agents. It improves navigation and object interaction by avoiding unnecessary exploration and correct planning of actions with appropriate objects to interact with. We empirically validate the benefit of the proposed modules in ALFRED by showing that the proposed method outperforms existing methods, especially in unseen environments even without human-designed plan templates.

Limitation and Future Work. Our context-aware planning fixes the anticipated context during a task execution as an inductive bias. But the context may change even in a single task execution. A prospective avenue for future investigation lies in the modification of context in response to input from environments, thereby enhancing adaptability.

Acknowledgment. This work is partly supported by the NRF grant (No.2022R1A2C4002300) 20%, IITP grants (No.2020-0-01361, AI GS Program (Yonsei University) 5%, No.2021-0-02068, AI Innovation Hub 5%, 2022-0-00077 10%, 2022-0-00113 10%, 2022-0-00959 10%, 2022-0-00871 20%, 2022-0-00951 20%) funded by the Korea government (MSIT).

References

- [1] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Daniel Ho, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Eric Jang, Rosario Jauregui Ruano, Kyle Jeffrey, Sally Jesmonth, Nikhil J Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Kuang-Huei Lee, Sergey Levine, Yao Lu, Linda Luu, Carolina Parada, Peter Pastor, Jornell Quiambao, Kanishka Rao, Jarek Rettinghouse, Diego Reyes, Pierre Sermanet, Nicolas Sievers, Clayton Tan, Alexander Toshev, Vincent Vanhoucke, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Mengyuan Yan, and Andy Zeng. Do as i can, not as i say: Grounding language in robotic affordances. In *CoRL*, 2022. 2, 3
- [2] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *CVPR*, 2018. 1
- [3] Michael Bain and Claude Sammut. A framework for behavioural cloning. In *Mach. Intell.* 15, 1995. 6
- [4] Suvvansh Bhambri, Byeonghwi Kim, and Jonghyun Choi. Multi-level compositional reasoning for interactive instruction following. In *AAAI*, 2023. 2, 4, 9
- [5] Valts Blukis, Chris Paxton, Dieter Fox, Animesh Garg, and Yoav Artzi. A persistent spatial semantic representation for high-level natural language instruction execution. In *CoRL*, 2022. 2, 4, 5, 6, 7
- [6] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. In *3DV*, 2017. 1
- [7] Devendra Singh Chaplot, Kanthashree Mysore Sathyendra, Rama Kumar Pasumarthi, Dheeraj Rajagopal, and Ruslan Salakhutdinov. Gated-attention architectures for task-oriented language grounding. In *AAAI*, 2017. 1
- [8] Howard Chen, Alane Suhr, Dipendra Misra, Noah Snaveley, and Yoav Artzi. Touchdown: Natural language navigation and spatial reasoning in visual street environments. In *CVPR*, 2019. 1
- [9] Peihao Chen, Dongyu Ji, Kunyang Lin, Runhao Zeng, Thomas Li, Mingkui Tan, and Chuang Gan. Weakly-supervised multi-granularity map learning for vision-and-language navigation. *NeurIPS*, 2022. 2
- [10] Rodolfo Corona, Daniel Fried, Coline Devin, Dan Klein, and Trevor Darrell. Modular networks for compositional instruction following. In *NAACL*, 2021. 2
- [11] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied question answering. In *CVPR*, 2018. 1, 2
- [12] Abhishek Das, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Neural modular control for embodied question answering. In *CoRL*, 2018. 2
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019. 5
- [14] Kiana Ehsani, Winson Han, Alvaro Herrasti, Eli VanderBilt, Luca Weihs, Eric Kolve, Aniruddha Kembhavi, and Roozbeh Mottaghi. Manipulathor: A framework for visual object manipulation. In *CVPR*, 2021. 2
- [15] Georgios Georgakis, Karl Schmeckpeper, Karan Wanchoo, Soham Dan, Eleni Miltsakaki, Dan Roth, and Kostas Daniilidis. Cross-modal map learning for vision and language navigation. In *CVPR*, 2022. 2
- [16] Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, and Ali Farhadi. Iqa: Visual question answering in interactive environments. In *CVPR*, 2018. 1, 2
- [17] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *ICML*, 2022. 2
- [18] Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, et al. Inner monologue: Embodied reasoning through planning with language models. In *CoRL*, 2022. 2, 3
- [19] Yuki Inoue and Hiroki Ohashi. Prompter: Utilizing large language model prompting for a data efficient embodied instruction following. *arXiv:2211.03267*, 2022. 2, 3, 4, 5, 6, 7
- [20] Shintaro Ishikawa and Komei Sugiura. Moment-based adversarial training for embodied language comprehension. In *ICPR*, 2022. 4, 5, 6, 7
- [21] Zhiwei Jia, Kaixiang Lin, Yizhou Zhao, Qiaozi Gao, Govind Thattai, and Gaurav S Sukhatme. Learning to act with affordance-aware multimodal neural slam. In *IROS*, 2022. 7
- [22] Alexander Katrompas and Vangelis Metsis. Enhancing lstm models with self-attention and stateful training. In *IntelliSys*, 2022. 5
- [23] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. Ai2-thor: An interactive 3d environment for visual ai. *arXiv:1712.05474*, 2017. 1
- [24] Jacob Krantz, Erik Wijmans, Arjun Majumdar, Dhruv Batra, and Stefan Lee. Beyond the nav-graph: Vision-and-language navigation in continuous environments. In *ECCV*, 2020. 1, 2
- [25] Federico Landi, Lorenzo Baraldi, Massimiliano Corsini, and Rita Cucchiara. Embodied vision-and-language navigation with dynamic convolutional filters. In *BMVC*, 2019. 2
- [26] Xiaotian Liu, Hector Palacios, Muise, and Christian. A planning based neural-symbolic approach for embodied instruction following. In *Embodied AI Workshop @ CVPR*, 2022. 4, 7
- [27] So Yeon Min, Devendra Singh Chaplot, Pradeep Ravikumar, Yonatan Bisk, and Ruslan Salakhutdinov. Film: Following instructions in language with modular methods. In *ICLR*, 2022. 2, 3, 4, 5, 6, 7, 9

- [28] Dipendra Misra, John Langford, and Yoav Artzi. Mapping instructions and visual observations to actions with reinforcement learning. In *EMNLP*, 2017. 1
- [29] Michael Murray and Maya Cakmak. Following natural language instructions for household tasks with landmark guided search and reinforced pose adjustment. In *RA-L*, 2022. 4, 5, 6, 7
- [30] Van-Quang Nguyen, Masanori Suganuma, Okatani, and Takayuki. Look wide and interpret twice: Improving performance on interactive instruction-following tasks. In *IJCAI*, 2021. 2, 3
- [31] Aishwarya Padmakumar, Jesse Thomason, Ayush Srivastava, Patrick Lange, Anjali Narayan-Chen, Spandana Gella, Robinson Piramuthu, Gokhan Tur, and Dilek Hakkani-Tur. Teach: Task-driven embodied agents that chat. In *AAAI*, 2022. 2, 6
- [32] Alexander Pashevich, Cordelia Schmid, Sun, and Chen. Episodic transformer for vision-and-language navigation. In *ICCV*, 2021. 2, 3
- [33] Shreyas Sundara Raman, Vanya Cohen, Eric Rosen, Ifrah Idrees, David Paulius, and Stefanie Tellex. Planning with large language models via corrective re-prompting. In *Foundation Models for Decision Making Workshop NeurIPS*, 2022. 2, 3
- [34] Ram Ramrakhya, Eric Undersander, Dhruv Batra, and Abhishek Das. Habitat-web: Learning embodied object-search strategies from human demonstrations at scale. In *CVPR*, 2022. 6
- [35] James A Sethian. A fast marching level set method for monotonically advancing fronts. In *Proceedings of the National Academy of Sciences*, 1996. 6
- [36] Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. Alfred: A benchmark for interpreting grounded instructions for everyday tasks. In *CVPR*, 2020. 1, 2, 3, 7
- [37] Kunal Pratap Singh, Suvaansh Bhambri, Byeonghwi Kim, Roozbeh Mottaghi, and Jonghyun Choi. Factorizing perception and policy for interactive instruction following. In *ICCV*, 2021. 2, 3
- [38] Aviv Tamar, Yi Wu, Garrett Thomas, Sergey Levine, and Pieter Abbeel. Value iteration networks. In *NeurIPS*, 2016. 6
- [39] Luca Weihs, Matt Deitke, Aniruddha Kembhavi, and Roozbeh Mottaghi. Visual room rearrangement. In *CVPR*, 2021. 2, 6
- [40] Fei Xia, Amir R Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: Real-world perception for embodied agents. In *CVPR*, 2018. 1
- [41] Yichi Zhang and Joyce Chai. Hierarchical task learning from language instructions with unified transformers and self-monitoring. In *Findings of ACL-IJCNLP*, 2021. 2, 4
- [42] Yuke Zhu, Daniel Gordon, Eric Kolve, Dieter Fox, Li Fei-Fei, Abhinav Gupta, Roozbeh Mottaghi, and Ali Farhadi. Visual semantic planning using deep successor representations. In *ICCV*, 2017. 1, 2