

Heterogeneous Embodied Multi-Agent Collaboration

Xinzhu Liu, Di Guo, and Huaping Liu

Abstract—Multi-agent embodied tasks have recently been studied in complex indoor visual environments. Collaboration among multiple agents can improve work efficiency and has significant practical value. However, most of the existing research focuses on homogeneous multi-agent tasks. Compared with homogeneous agents, heterogeneous agents can leverage their different capabilities to allocate corresponding sub-tasks and cooperate to complete complex tasks. Heterogeneous multi-agent tasks are common in real-world scenarios, and the collaboration strategy among heterogeneous agents is a challenging and important problem to be solved. To study collaboration among heterogeneous agents, we propose the heterogeneous multi-agent tidying-up task, in which multiple heterogeneous agents with different capabilities collaborate with each other to detect misplaced objects and place them in reasonable locations. This is a demanding task since it requires agents to make the best use of their different capabilities to conduct reasonable task planning and complete the whole task. To solve this task, we build a heterogeneous multi-agent tidying-up benchmark dataset in a large number of houses with multiple rooms based on ProcTHOR-10K. We propose the hierarchical decision model based on misplaced object detection, reasonable receptacle prediction, as well as the handshake-based group communication mechanism. Extensive experiments are conducted to demonstrate the effectiveness of the proposed model. The project’s website and videos of experiments can be found at <https://hetercol.github.io/>.

Index Terms—Embodied Multi-Agent Collaboration, Heterogeneous Agents, Embodied Visual Tasks, Group-based Communication, Tidying-up Task.

1 INTRODUCTION

THE embodied tasks are a series of comprehensive tasks integrating the process of environment perception, scene understanding as well as action execution, which are of great significance in the field of intelligent robotics. Recently, a number of embodied tasks have been studied to learn the relationship between visual perception and actions including visual navigation [1], [2], [3], instruction following [4], room rearrangement [5] and so on. However, these tasks are mainly studied in the single-agent scenarios, the work efficiency is likely to be low and the system is of poor fault tolerance. To solve the weakness of the single-agent setting in executing tasks, there are already some works investigating how to collaborate between multiple agents. Some simple disembodied environments such as the grid-world or 2D game scene are firstly used to learn the cooperation strategy [6], [7], [8], [9]. Recently, studies on multi-agent collaboration in the embodied visual environment have been conducted. *FurnLift* [10] and *FurnMove* [11] tasks are proposed to learn the tight collaboration in the action level between two agents. Then multi-agent exploration [12], multi-agent visual navigation [13], [14], multi-agent EQA [15] and multi-agent embodied task planning [16] have been developed to study the cooperation strategy between multiple agents. But most of the multi-agent embodied tasks only consider the homogeneous agents, that is, all agents have the same ability. In practical tasks, different agents are likely to have different capabilities, and heterogeneous agents could leverage




Fig. 1. The overview of the heterogeneous multi-agent tidying-up task. *Agent 1* and *Agent 2* can only navigate, while *Agent 3* is capable of both navigation and manipulation. The height of *Agent 1* is low, while the heights of *Agent 2* and *Agent 3* are high. During exploration, *Agent 1* and *Agent 2* find a misplaced *Spoon* and a misplaced *Pillow* respectively but cannot pick them up, so they communicate with *Agent 3*. Then *Agent 3* navigates to pick up the *Spoon* and *Pillow*, and put them in reasonable places.

their different capabilities to deal with some complex problems that homogeneous agents cannot. The heterogeneous multi-agent task is common in real-world scenarios and has important practical applications. While the current research on heterogeneous agents is still limited to relatively simple experimental setups [17], [18]. In this paper, we concentrate on the collaboration strategy of heterogeneous agents in the embodied visual task.

To further study heterogeneous embodied multi-agent collaboration, verify and evaluate the proposed collaboration strategy, we propose the heterogeneous multi-agent tidying-up task, in which multiple agents need to find all misplaced objects and

- Xinzhu Liu and Huaping Liu are with the Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China, and also with the State Key Laboratory of Intelligent Technology and Systems, Beijing National Research Center for Information Science and Technology, Tsinghua University, Beijing 100084, China.
- Di Guo is with the School of Artificial Intelligence, Beijing University of Posts and Telecommunications, Beijing, China.
- Corresponding author: Huaping Liu (hpliu@tsinghua.edu.cn).

TABLE 1
Comparison Between Our Work and Other Related Embodied Tasks

Task	Multi-Agent	Heterogeneous Agents	Multi-Room Scene	Navigation	Manipulation	Reasoning
MQA [20]	×	×	×	×	✓	×
Rearrangement [5]	×	×	×	✓	✓	✓
TIDEE [21]	×	×	×	✓	✓	✓
Housekeep [22]	×	×	✓	✓	✓	✓
MA-EQA [15]	✓	×	×	✓	×	×
FurnLift [10] & FurnMove [11]	✓	×	×	✓	✓	×
CollaVN [14]	✓	×	✓	✓	×	×
MA-Exploration [12]	✓	×	✓	✓	×	×
MA-VSN [13]	✓	×	×	✓	×	×
MA-TaskPlannig [16]	✓	×	✓	✓	✓	✓
*CH-MARL [19]	✓	✓	✓	✓	✓	×
Ours	✓	✓	✓	✓	✓	✓

*CH-MARL studies collaboration between two agents, the humanoid and drone, in indoor scenes, and agents can clearly know the target object and target receptacle without reasoning. Our work considers collaboration among multiple agents moving on the ground with or without the manipulator, and agents need to reason the reasonable target receptacle.

place them in reasonable locations. An overview of the proposed task is demonstrated in Fig. 1, in which multiple heterogeneous agents with different capabilities work together to detect the misplaced objects in the house and put them back to reasonable locations. The proposed task is quite challenging as it requires heterogeneous agents to make the best use of their own advantages and the complementary abilities of manipulation and navigation to conduct reasonable task planning as well as collaborate with each other to achieve the tidying-up task. Besides, this task also requires a deeper understanding of the scene. Agents need to analyze whether each object and its surroundings are placed reasonably and infer the reasonable receptacle to put the misplaced object based on visual observations and common sense reasoning without extra guidance. The most relevant work with ours is *CH-MARL* in [19], but its experimental setup is relatively simple, since it only considers the collaboration among two agents, the humanoid and the drone, and agents can clearly know the target objects and receptacles without reasoning. The task has been completed in only 7 scenes. Our task studies collaboration among a larger number of agents moving on the ground in a larger number of indoor houses, and agents need to infer target receptacles themselves.

In this paper, We propose a heterogeneous embodied multi-agent collaboration framework based on the hierarchical decision strategy and the handshake-based group communication mechanism. We evaluate the effectiveness of our model in the proposed heterogeneous multi-agent tidying-up task. We also build a benchmark dataset to complete the evaluation experiments in the heterogeneous multi-agent setting in multi-room scenes based on ProcTHOR-10K [23]. The comparisons of our benchmark with other relevant embodied tasks are illustrated in Table 1. The main contributions are summarized as follows:

- We propose a heterogeneous embodied multi-agent collaboration framework, which makes full use of the different capabilities of heterogeneous agents to collaboratively complete the whole task.
- We develop a handshake-based group communication mechanism for heterogeneous multi-agent collaboration. A hierarchical decision strategy based on scene understanding, scene reasoning with prior knowledge, and communication mechanism is proposed to solve the heteroge-

neous multi-agent collaboration task.

- We propose the heterogeneous multi-agent tidying-up task to evaluate the proposed heterogeneous collaboration strategy, in which multiple heterogeneous agents need to find all misplaced objects and place them to reasonable locations. We build a novel benchmark dataset to evaluate our model, which includes *Single-Room* and *Cross-Room* settings. We also provide the method to generate the task data in indoor houses for the tidying-up task. Experimental results demonstrate the effectiveness of the proposed heterogeneous multi-agent collaboration framework.

This paper is organized as follows. In Section 2, the related existing researches are summarized. In Section 3, the heterogeneous multi-agent collaboration framework and heterogeneous multi-agent tidying-up task are formulated. In Section 4, we demonstrate the newly built dataset. The proposed heterogeneous multi-agent model is described in Section 5. Then the experiments are shown in Section 6. Finally, we come to the conclusion in Section 7.

2 RELATED WORK

2.1 Embodied Tasks

In recent years, embodied visual tasks have attracted a large number of researches. Researchers have proposed and studied the visual navigation task including visual semantic navigation [1], [2], [24] and visual-and-language navigation [3], [25], [26], [27], [28], [29], visual exploration [30], remote embodied visual referring expression [31], embodied question answering (EQA) [32], [33], [34], [35], [36] and instruction following [4], [37]. Recently proposed visual room rearrangement [5] and tidying-up tasks [21], [22] require the agent to learn the abilities of deep scene understanding and scene reasoning in long-horizon tasks. However, the above tasks are mainly studied in single-agent scenarios, which may lead to low efficiency and poor fault tolerance. Researchers have further studied multi-agent embodied tasks, including collaborative object movement [10], [11], collaborative exploration [12], multi-agent visual navigation [13], [14], collaborative EQA [15] and collaborative task planning [16]. All these works consider homogeneous agents with exactly the same capabilities.

2.2 Heterogeneous Multi-Agent Tasks

Different from homogeneous agents, heterogeneous agents can leverage their different capabilities to allocate corresponding sub-tasks and cooperate to complete complex tasks. Heterogeneous multi-agent tasks are common in real-world scenarios and have important practical applications, such as the air-ground cooperative system, collaborative manipulators with different functions in the factory, etc. Recently, researchers have proposed heterogeneous multi-agent reinforcement learning (MARL) methods to solve different tasks including known environment mapping [17], cooperative cache in cognitive radio networks [38], autonomous separation assurance [39]. Then inductive heterogeneous graph MARL [40], heterogeneous MARL with mean field control [18], and meta RL-based approach [41] have been proposed to improve the method. In this paper, we consider the collaboration and communication among heterogeneous agents in embodied tasks, in which all heterogeneous agents simultaneously perceive the environment and collaboratively perform actions to complete the task. It is different from *CoMON* [42], in which one oracle agent only provides information without executing actions and only one agent executes the task. *CH-MARL* considers the task of heterogeneous agents in the embodied environment with two different robots including humanoid and drone [19], but the experimental setup is relatively simple, in which agents only need to move an object to a specific receptacle without the requirement of scene reasoning. Besides, *TarMac* [43] and *When2com* [44] are also related to our work on multi-agent communication, but they only study the communication of homogeneous agents without considering different abilities of heterogeneous agents. *TarMac* broadcasts the signature and value vectors to others, *When2com* utilizes a handshake-based method to complete the intra-group communication, while our model develops the handshake-based group communication mechanism with both intra-group and inter-group communication, which is effective in collaboration of heterogeneous agents.

3 PROBLEM FORMULATION

For the heterogeneous multi-agent collaboration tasks, we consider that there exist N heterogeneous agents $A^{(1)}, A^{(2)}, \dots, A^{(N)}$. Each agent $A^{(i)}$ has a capability representation $B^{(i)}$, which consists of three different dimensions, namely the perception ability $P^{(i)}$, the action ability $E^{(i)}$ and the agent morphology $H^{(i)}$, $B^{(i)} = (P^{(i)}, E^{(i)}, H^{(i)})$. Specifically, in the heterogeneous multi-agent tidying-up task, each agent has the same visual perception ability. The action ability $E^{(i)}$ can be indicated as $E^{(i)} = (Nav^{(i)}, Mani^{(i)})$, where $Nav^{(i)}, Mani^{(i)}$ denote whether the agent has the ability of navigation and manipulation respectively. If the agent has the corresponding ability, its value is 1, otherwise, the value is 0. As for the agent morphology $H^{(i)}$, we consider different heights of the agents and represent it as $H^{(i)} = (Hei^{(i)})$, where $Hei^{(i)} = 1$ denotes the height of the agent is high, roughly 0.9 meters, and $Hei^{(i)} = 0$ denotes the height of the agent is low, roughly 0.22 meters. N heterogeneous agents have different $(E^{(i)}, H^{(i)})$, and different action spaces $I^{(i)}$. Specifically, if $Nav^{(i)} = 1$ and $Mani^{(i)} = 0$, $I^{(i)} = I_{nav}$. If $Nav^{(i)} = 1$ and $Mani^{(i)} = 1$, $I^{(i)} = I_{nav} \cup I_{mani}$, where $I_{nav} = \{MoveAhead, MoveRight, MoveLeft, RotateRight, RotateLeft, LookUp, LookDown, Stop\}$ denotes the navigation action set and $I_{mani} = \{PickUp, PutDown, Drop\}$ represents the manipulation actions.

There exist M objects in the house which are denoted as $\mathcal{O} = \{o_1, o_2, \dots, o_M\}$. Among these M objects, the subset $\mathcal{O}_{pick} \subseteq \mathcal{O}$ denotes the set of objects which can be picked up by the agent. The subset $\mathcal{O}_{recep} \subseteq \mathcal{O}$ represents the set of objects which can be used as a receptacle and other objects can be placed on their surface or inside. The location of the j -th object can be represented as (o_j, p_j, r_j) , where $p_j \in \mathcal{O}_{recep}$ denotes the receptacle that o_j is placed in, and r_j is the room where o_j and p_j are in. The environment can provide a judgment of whether the current object is in a reasonable location with a discriminator D . If $D(o_j, p_j, r_j) = True$, o_j is in a reasonable location, otherwise, o_j is in an unreasonable location. There exist k objects which are in unreasonable locations in the house, namely misplaced objects, and agents do not know the value of k . At the t -th step, the agent $A^{(i)}$ obtains the perception observation $e_t^{(i)}$ and the communication information $c_t^{(i)}$ from other agents. Each agent $A^{(i)}$ needs to learn a action strategy $\pi^{(i)}$ to generate the action $a_t^{(i)}$ from its action space $I^{(i)}$, $a_t^{(i)} = \pi^{(i)}(e_t^{(i)}, c_t^{(i)})$. The objective of this tidying-up task is to learn the optimal collaboration action strategy $\pi^{(i)}$ for each agent such that they can collaboratively find all misplaced objects and place them in reasonable locations to make $D(o_j, p_j^*, r_j^*) = True, j = 1, 2, \dots, M$ with as few steps as possible, where (o_j, p_j^*, r_j^*) denotes the location of objects after all agents complete their actions and stop.

4 DATASET

To evaluate the proposed heterogeneous multi-agent collaboration model, we create the novel dataset for the heterogeneous multi-agent tidying-up tasks based on the ProcTHOR-10K [23] which provides a large number of houses with multiple rooms. In each task, several objects are placed in unreasonable locations, and multiple heterogeneous agents are required to find these misplaced objects and put them to the reasonable locations. The newly built dataset contains the *Single-Room (Single)* and *Cross-Room (Cross)* tasks to train and evaluate the proposed model.

ProcTHOR-10K is a generated dataset of houses with multiple rooms [23] built based on AI2-THOR [45], including 10,000 training houses, 1000 validation houses, and 1000 testing houses, and each room belongs to one type of *Kitchen, LivingRoom, Bedroom* and *Bathroom*. We propose a method to generate multi-agent tidying-up tasks in ProcTHOR-10K. Since every single room in ProcTHOR-10K is built based on AI2-THOR, we regard the object placement relationships in AI2-THOR as the reasonable placement constraints in ProcTHOR-10K. We obtain the properties of objects and their reasonable locations from ProcTHOR-10K. Specifically, we obtain the objects which have the properties *Pickupable* and *Receptacle* to form the set \mathcal{O}_{pick} and \mathcal{O}_{recep} respectively. For each object type in \mathcal{O}_{pick} , we generate the triples of reasonable candidate locations $Re_i = \{(o_i, p_j, r_j)\}$ from the metadata, where $o_i \in \mathcal{O}_{pick}, p_j \in \mathcal{O}_{recep}, r_j \in \{Kitchen, LivingRoom, Bedroom, Bathroom\}$, o_i can be put on p_j in room r_j . We generate 10 meta-tasks for each house in ProcTHOR-10K, and the process is shown in Fig.2. For each tidying-up meta-task, we randomly select k objects to change their current location $(o_i, p_i^{cur}, r_i^{cur})$ to an unreasonable location $(o_i, p_i^{new}, r_i^{new})$, and k is randomly chosen from $\{1, 2, 3, 4, 5\}$. We use the action *PutObject* in ProcTHOR-10K to place the misplaced object to an unreasonable receptacle or use the action *DropHandObject, ThrowObject* to drop the object to a random location. For each misplaced object, p_i^{new} and r_i^{cur} may




Fig. 2. The process of generating the tidying-up task from the original scene. For each tidying-up task, we randomly select k objects from the original house in ProcTHOR-10K and then change their current locations to unreasonable receptacles or unreasonable rooms to generate the task data.

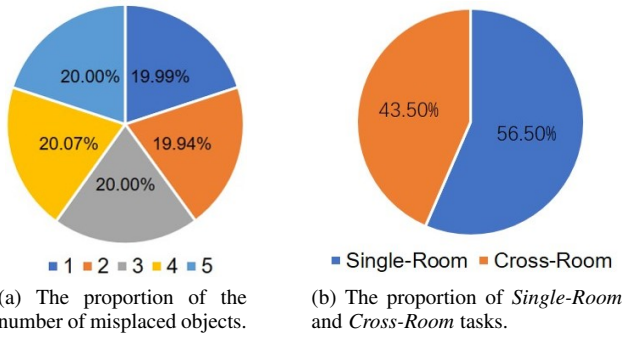


Fig. 3. The statistics of the built dataset.

be the same or different. If $r_i^{new} = r_i^{cur}$, the misplaced object may be placed in an unreasonable location in the same room. If $r_i^{new} \neq r_i^{cur}$, the misplaced object may be placed to the location in another unreasonable room. All meta-tasks are divided into *Single* and *Cross* tasks. In *Single* tasks, all k misplaced objects are placed in the same room as their initial positions, and in *Cross* tasks, at least one misplaced object is placed in a different room from its initial position. The statistics of the dataset are shown in Fig. 3. The proportion of the selected number of misplaced objects in the generated dataset is shown in Fig. 3a, and Fig. 3b illustrated the proportion of *Single-Room* tasks and *Cross-Room* tasks.

For each meta-task, we randomly generate 5 initial positions for multiple heterogeneous agents. We also generate the expert demonstration with heuristic shortest path methods for imitation learning. Specifically, we obtain the metadata and scene information of each task, and collect labels for the sub-task and sub-goal predictors of the action strategy module. We generate oracle sub-task lists for each agent from task information and find viewpoints that agents can interact with misplaced objects as target locations for manipulation. We generate the sub-task allocation results for each agent at each step according to the distance between each agent and misplaced objects. Then, we generate the heuristic shortest path in the sub-goal level, generate 5 reasonable trajectories per training task, and obtain pairs of state features and sub-goal actions $(\Delta x_t, \Delta y_t, \Delta rot_t, ope_t, stop_t)$ from them, which are mentioned in Section 5. Then we use sub-goal actions as labels for imitation learning.

We notice that our dataset is different from that in [21]. The tasks in our dataset are more difficult than [21], and our tasks require deeper scene understanding and reasoning. Firstly, the

tidying-up tasks in our dataset are built on large houses with multi-rooms and misplaced objects are likely to be placed in another unreasonable room, while the dataset in [21] are only built in the single room in AI2-THOR. Secondly, we consider the relationship among pickupable objects, receptacles and rooms, which means that the misplaced objects may be placed in an unreasonable receptacle, or in a reasonable receptacle but an unreasonable room. But in [21], misplaced objects are generated only by being thrown away on the ground in the same room, and it does not need to consider the relationship between objects and rooms. It does not consider the relationship between objects and receptacles either. Besides, our dataset is used to solve the multi-agent tasks and each task has initial positions for multiple agents, while the dataset in [21] is designed for the single-agent task. Furthermore, the major differences between our dataset and *Housekeep* in [22] are that our dataset has a larger scale, can be automatically generated through the proposed program, and is designed for heterogeneous multi-agent collaboration. We build a benchmark dataset based on ProcTHOR-10K which contains 10K houses with multiple rooms, while dataset in [22] only has 14 scenes. We also provide the program to generate task data and our dataset can be automatically generated through the program, while the generation of *Housekeep* needs a large number of efforts from human annotations. Our program can also be easily extended to any other indoor simulation platform with a large number of houses to generate the corresponding task data. Meanwhile, our dataset is designed for heterogeneous multi-agent task, while *Housekeep* is designed for the single-agent task. Misplaced objects in our dataset are likely to be placed in the receptacle or thrown on the floor, which are easier to be observed by agents with high or low height respectively, making it more suitable for the study of heterogeneous multi-agent collaboration.

5 METHODOLOGY

The proposed model that solves the multi-agent tidying-up task consists of four main modules: the misplaced object detector, the reasonable receptacle predictor, the communication module, and the hierarchical decision. The overview of this model is shown in Fig. 4. The misplaced object detector judges whether there exists an object placed in an unreasonable location. The reasonable receptacle predictor generates a reasonable receptacle and room type to place the misplaced objects. The communication module transmits the communication information to other heterogeneous agents. The hierarchical decision module predicts the next sub-task, sub-goal, and next actions for each agent to execute.




Fig. 4. An overview of the structure of the proposed heterogeneous multi-agent collaboration model. The model consists of the misplaced object detector, the reasonable receptacle predictor, the communication module, and the hierarchical decision. The misplaced object detector judges whether there exists an object placed in an unreasonable location. The reasonable receptacle predictor generates a reasonable location to place the misplaced objects. The communication module shares communication messages among heterogeneous agents with the handshake-based group communication strategy. The hierarchical decision module predicts the next sub-task, sub-goal, and next low-level actions for each agent.

5.1 Misplaced Object Detector

Each agent builds a top-down semantic map with the input RGB and depth using the method similar to the semantic mapping in [13]. At each sub-goal step, the agent obtains a local semantic map of size $G \times G \times (K_{total} + 2)$, where G indicates the size of the local map, and K_{total} is the number of object categories. Then agent $A^{(i)}$ generates the map embedding $sm_t^{(i)}$ with a pre-trained scene encoder similar to that in [16] consisting of multiple convolutional layers.

We filter out spatial relationships among objects, receptacles and room types existing in our dataset from Visual Genome [46] to form a graph structure. We regard these relationships as the commonsense prior knowledge. We use the Glove embedding [47] to encode the class label, concatenate with the visual embedding as node features, and use GCN [48] to obtain the embedding of the commonsense $sq_t^{(i)}$. ResNet [49] is utilized to extract the visual embedding $sv_t^{(i)}$ of the visual observation. Meanwhile, the pre-trained Mask-RCNN [50] is used to detect the existing objects and obtain the feature $sl_t^{(i)}$. We construct a room classifier to predict the room type where the agent is currently located based on detected objects and their placement relationships. The classifier extracts the room embedding $sr_t^{(i)}$ of the current state. The misplaced object detector F_d , a binary classifier consisting of two linear layers, is utilized to fuse embeddings and detect whether the object is in a reasonable location. $det_t^{(i)} = F_d(sq_t^{(i)}, sv_t^{(i)}, sl_t^{(i)}, sr_t^{(i)})$. $det_t^{(i)} = 1$ denotes the object is misplaced in an unreasonable location and $det_t^{(i)} = 0$ denotes the object is in a reasonable location.

5.2 Reasonable Receptacle Predictor

The detection feature $sl_t^{(i)}$ and the commonsense embedding $sq_t^{(i)}$ are fed into an attention layer to obtain the object-commonsense attention $atll_t^{(i)}$. The room embedding $sr_t^{(i)}$ and $sq_t^{(i)}$ are also fed into an attention layer to generate the room-commonsense attention $attr_t^{(i)}$. The reasonable receptacle predictor F_p including linear layers is built to fuse the visual embedding $sv_t^{(i)}$ with the attention embedding and predicts the reasonable lo-

cation including receptacle p_t and room type r_t for the misplaced object o_t , then we have $(p_t, r_t) = F_p(sv_t^{(i)}, atll_t^{(i)}, attr_t^{(i)})$.

5.3 Communication

Before communication, each agent generates a characteristic vector $va^{(i)}$ based on their capability and property, $va^{(i)} = (Nav^{(i)}, Mani^{(i)}, Hei^{(i)})$, and obtains a pose vector $pose_t^{(i)} = (x_t^{(i)}, y_t^{(i)}, rot_t^{(i)})$. The state feature extractor F_s including linear layers and the LSTM [51] layer is utilized to obtain the fused state feature $sf_t^{(i)} = F_s(va^{(i)}, pose_t^{(i)}, fd_t^{(i)}, fp_t^{(i)})$, where $fd_t^{(i)}, fp_t^{(i)}$ are the extracted features from the second-from-last layers of F_d and F_p (the layer before the last classification layer) respectively. To solve the collaboration among heterogeneous agents, we propose the *Handshake-based Group Communication (HanGrCom)*. In this module, the agent generates the query vector $qry_t^{(i)}$, key vector $key_t^{(i)}$, value vector $val_t^{(i)}$ and inter-group information vector $inn_t^{(i)}$ based on its state feature $[qry_t^{(i)}, key_t^{(i)}, val_t^{(i)}, inn_t^{(i)}] = [\theta^q, \theta^k, \theta^v, \theta^e](sf_t^{(i)})$, where $\theta^q, \theta^k, \theta^v, \theta^e$ are corresponding vector generators consisting of linear layers respectively. At the t -th sub-goal step, each agent sends $qry_t^{(i)}$ to others and calculates the scaled inner product attention of received query vectors and its key vector $attc_{ij} = \frac{qry_t^{(j)} key_t^{(i)}}{\sqrt{d}}$, where d is the dimension of $qry_t^{(j)}$ and $key_t^{(i)}$. Then a communication matrix $T_t = \sigma([attc_{ij}]_{N \times N})$ is obtained, where $attc_{ij}$ indicates the effectiveness of the information send from $A^{(j)}$ to $A^{(i)}$, and σ denotes row-wise softmax function. When communicating, we set two thresholds δ and μ . If $attc_{ii} < \delta$, $A^{(i)}$ needs to receive messages from other agents, since its own information is not efficient. If $attc_{ij} > \mu (i \neq j)$, $A^{(i)}$ can receive messages from $A^{(j)}$. Based on T_t , agents can be implicitly divided into several groups. Intra-group communication exchanges the state information from agents in the same group to complete short-term sub-tasks, and $A^{(i)}$ obtain the intra-group communication information $inn_t^{(i)} = \sum_{(j \neq i, attc_{ij} > \mu)} attc_{ij} \cdot val_t^{(j)}$. Inter-group communication conveys higher-level information on the task allocation across different groups which is beneficial for agents to make future decisions in the long-term task completion, and $A^{(i)}$ obtains the inter-group communication information

TABLE 2
The Details of the Model Structure

Module	Layers
Room Classifier	$Linear(256, 256) \rightarrow Linear(256, 4)$ $[r_t^{(i)}]$
Misplaced Object Detector	$Linear(256, 256) \rightarrow Linear(256, 2)$ $[det_t^{(i)}]$
Receptacle Predictor	$Linear(256, 256) \rightarrow Linear(256, 43)$ $[p_t^{(i)}]$ $\searrow Linear(256, 4)$ $[r_t^{(i)}]$
Sub-Task Predictor	$\nearrow Linear(256, 2)$ $[Explore\ or\ Place]$ $\nearrow Linear(256, 118)$ $[o_t^{(i)}]$ $LSTM(512) \rightarrow Linear(512, 256) \rightarrow Linear(256, 43)$ $[p_t^{(i)}]$ $\searrow Linear(256, 4)$ $[r_t^{(i)}]$
Sub-Goal Predictor	$\nearrow Linear(256, 9)$ $[\Delta x_t^{(i)}]$ $\nearrow Linear(256, 9)$ $[\Delta y_t^{(i)}]$ $LSTM(512) \rightarrow Linear(512, 256) \rightarrow Linear(256, 5)$ $[\Delta rot_t^{(i)}]$ $\searrow Linear(256, 4)$ $[ope_t^{(i)}]$ $\searrow Linear(256, 2)$ $[stop_t^{(i)}]$

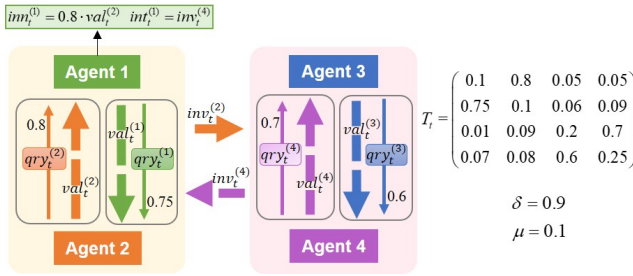


Fig. 5. An example of the communication process. Based on T_t , at that step, *Agent 1* and *Agent 2* would form a communication group, *Agent 3* and *Agent 4* would form a group. The intra-group and inter-group communication information of $A^{(1)}$ is obtained based on T_t , δ and μ .

$int_t^{(i)} = inv_t^{(j)}$, $j = \text{argmax}(\text{avg}(\text{att}_{*j}))$, where avg is the average operation and $*$ denotes the ids of agents that belong to the same group as $A^{(i)}$. $inn_t^{(i)}$ and $int_t^{(i)}$ are utilized to generate the sub-task. An example of this mechanism is shown in Fig. 5.

In the training process, in order to ensure the generation of communication information is differentiable, it is assumed that agents can obtain state features of other agents. At each sub-goal step, $A^{(i)}$ obtains the intra-group communication information $inn_t^{(i)} = \sum_{j=1}^N att_{ij} \cdot val_t^{(j)}$ and inter-group information $int_t^{(i)} = \sum_{j=1}^N (1 - att_{ij}) \cdot inv_t^{(j)}$. In the inference process after training, agents can only obtain their own state features and complete the communication process according to the latent divided communication groups under the threshold δ, μ as mentioned above. We set $\mu = 0.2, \delta = 0.8$ in our experiments.

5.4 Hierarchical Decision

The hierarchical decision module generates the sub-task for each agent, predicts the sub-goal for the corresponding sub-task, and chooses the low-level actions to execute.

5.4.1 Sub-Task Planning

The sub-task planning part generates the sub-task for each agent to complete. In this tidying-up task, two types of sub-tasks can be executed, *Explore* and *Place*. In the sub-task *Explore*, the agent needs to explore the scene from its current position and search for misplaced objects. The sub-task *Place* contains three parameters, which are denoted as $(Place, o_t, p_t, r_t)$, meaning to

put the misplaced object o_t to the receptacle p_t in the room type r_t . A sub-task predictor F_a including the LSTM layer and linear layers is utilized to predict the type and parameters of the executed sub-task with $sf_t^{(i)}$, $inn_t^{(i)}$ and $int_t^{(i)}$. We use the output of the layer before the last prediction layer to obtain the embedding $se_t^{(i)}$ of the predicted sub-task $se_t^{(i)} = F_a(sf_t^{(i)}, inn_t^{(i)}, int_t^{(i)})$.

5.4.2 Action Decision

Each agent encodes the $pose_t^{(i)}$ to get the pose embedding $sp_t^{(i)}$. A sub-goal predictor F_g including the LSTM layer and linear layers is utilized to predict the next sub-goal $sg_t^{(i)}$ with $se_t^{(i)}$, $sm_t^{(i)}$, $sp_t^{(i)}$ and the visual embedding $sv_t^{(i)}$. $sg_t^{(i)} = (\Delta x_t^{(i)}, \Delta y_t^{(i)}, \Delta rot_t^{(i)}, ope_t^{(i)}, stop_t^{(i)}) = F_g(se_t^{(i)}, sm_t^{(i)}, sp_t^{(i)}, sv_t^{(i)})$, where $\Delta x_t^{(i)}$, $\Delta y_t^{(i)}$, $\Delta rot_t^{(i)}$ are the movement in egocentric x and y axis and the rotation angle of the agent respectively. $ope_t^{(i)}$ denotes whether performing action *PickUp*, *PutDown*, *Drop* or *NoAction* after reaching the sub-goal. $stop_t^{(i)}$ predicts whether the agent stops or not after this sub-goal step. Then, a low-level action strategy based on the shortest path algorithm is utilized to generate specific actions from $I^{(i)}$ to reach the current sub-goal.

5.5 Model Training

We demonstrate the details of our model structure in Table 2. $Linear(d_{in}, d_{out})$ indicates the linear layer, and d_{in} , d_{out} denote the dimensions of input and output of the layer respectively. $LSTM(d_{hid})$ indicates the LSTM layer, and d_{hid} denotes the hidden size. $[*]$ represents the predicted contents. Specifically, $r_t^{(i)}$ indicates the room types which include 4 categories, $det_t^{(i)}$ indicates whether the object is misplaced or not, $p_t^{(i)}$ denotes the receptacle types which include 43 categories, $o_t^{(i)}$ denotes the object types which include 118 categories. The agent would predict the egocentric sub-goal within 1 meter in the x and y axis respectively, and the agent moves 0.25 meters each time corresponding to one grid after the discretization of the environment. The values of $\Delta x_t^{(i)}$ and $\Delta y_t^{(i)}$ are $\{-4, -3, -2, -1, 0, 1, 2, 3, 4\}$ respectively, where the negative values indicate moving in the negative direction of x and y axis, the positive values indicate moving in the positive direction of x and y axis. $rot_t^{(i)}$ can take the value of 0, 90, 180, or 270.




Fig. 6. Different methods of communication. The thickness of the lines represents the communication amount. The thicker the line, the greater the amount. The solid line indicates the information is transmitted at each communication step, and the dotted line indicates the information is transmitted only when certain conditions are met. The rectangle on the line represents the delivered specific vectors.

Since the function of each module and their training process are relatively independent, we train them separately. The misplaced object detector and reasonable receptacle predictor are trained with supervised learning, and the action decision part of the hierarchical decision module is trained with imitation learning.

We first train the room classifier in the misplaced object detection module with the true labels of the rooms where the specific object is located obtained from the houses in ProcTHOR-10K. During the training process, the parameters of the pre-trained ResNet and Mask-RCNN are frozen. After the training of the room classifier, we fix its parameters and utilize it to extract the room embedding. Then we select the pickupable objects in a number of training houses in ProcTHOR-10K and place them to reasonable or unreasonable locations in the houses to get the training samples for the misplaced object detector and the reasonable receptacle predictor. We generate the true labels of whether the corresponding objects are misplaced and extract the reasonable locations of misplaced objects based on the metadata of ProcTHOR-10K to obtain the ground truth of training samples for the misplaced object detector and the reasonable receptacle predictor respectively. We train the misplaced object detector and the reasonable receptacle predictor with the parameters of the room classifier fixed. The loss function of this process is defined as follows,

$$\begin{aligned} Loss_{cla} &= Loss_{type} \\ Loss_{det} &= \alpha Loss_{mis} + \beta Loss_{rec1} + \lambda Loss_{room1} \end{aligned} \quad (1)$$

where $Loss_{cla}$ denotes the loss function of the room classifier, and $Loss_{type}$ indicates the cross entropy loss between the predicted and the true type of the room where the specific object is located. $Loss_{det}$ denotes the loss function of the misplaced object detector and the reasonable receptacle predictor. $Loss_{mis}$ indicates the binary cross entropy of whether the object is misplaced. $Loss_{rec1}$, $Loss_{room1}$ denotes the multi-label classification loss between the predicted and true reasonable types of the receptacle and the room where the object can be placed. α, β, λ are the hyper-parameters. In our implementation, we set α, β, λ to 1.

The scene encoder consists of convolutional layers and is pre-trained to extract the map embedding of the semantic map. We design the query to predict whether a type of object exists in a sub-region of the top-down map. We generate a series of queries, which consist of a sub-area and a specific type of object, and extract the true answer of the query from the metadata of the houses. Then we pre-train the scene encoder with supervised learning, and the loss function is denoted as the entropy loss between the predicted answers and the true answers.

The hierarchical decision module predicts the next sub-task and sub-goal for each agent. The sub-task planning part predicts

the type of sub-task (*Explore* or *Place*) and the parameters of the sub-task, which contain the misplaced object type $o_t^{(i)}$, the reasonable receptacle type $p_t^{(i)}$ and the reasonable room type $r_t^{(i)}$. The action decision part predicts the next sub-goal specified by the grid of 0.25 meters. It independently predicts the number of movement grids of the sub-goal within 1 meter in the x and y axis $\Delta x_t^{(i)}, \Delta y_t^{(i)}$ respectively, the rotation angle $\Delta rot_t^{(i)}$, the operation action $ope_t^{(i)}$ and the probability of performing the action $Stop_t^{(i)}$. $rot_t^{(i)}$ takes the value of 0, 90, 180, or 270. $ope_t^{(i)}$ takes the value of *PickUp*, *PutDown*, *Drop* or *NoAction*. Since agents are heterogeneous and their abilities are different, the hierarchical decision modules of each agent are trained separately, and they do not share weights. We use behavior cloning to train the hierarchical decision module with the generated expert demonstrations of allocated sub-task and sub-goal actions. The loss function of training each sub-task planning part and action decision part with imitation learning is defined as follows,

$$\begin{aligned} Loss_{subtask} &= Loss_{task} + \gamma_1 Loss_{obj} + \delta_1 Loss_{rec2} \\ &\quad + \theta_1 Loss_{room2} \\ Loss_{subgoal} &= Loss_{xloc} + Loss_{yloc} + \gamma_2 Loss_{rot} \\ &\quad + \delta_2 Loss_{ope} + \theta_2 Loss_{stop} \end{aligned} \quad (2)$$

where $Loss_{task}$ indicates the binary cross entropy of the sub-task type. $Loss_{obj}$, $Loss_{rec2}$, $Loss_{room2}$ denotes the cross entropy loss between the predicted $o_t^{(i)}, p_t^{(i)}, r_t^{(i)}$ and those in demonstrations respectively. $Loss_{xloc}$, $Loss_{yloc}$, $Loss_{rot}$, $Loss_{ope}$ denotes the cross entropy loss between the predicted results and those in demonstrations respectively, $Loss_{stop}$ denotes the binary cross entropy of whether the agent stops or not. $\gamma_1, \gamma_2, \delta_1, \delta_2, \theta_1, \theta_2$ are the hyper-parameters that can control the training process. In our implementation, we set $\gamma_1, \gamma_2, \delta_1, \delta_2, \theta_1, \theta_2$ to be 1.

6 EXPERIMENTS

6.1 Experiment Setting

We focus on heterogeneous multi-agent collaboration, and we evaluate the proposed collaboration strategy in the tidying-up task. To effectively and efficiently train and verify our model, we select 120 scenes from the newly built dataset, 80 houses for training, 20 for validation, and 20 for testing. The number of scenes in our experiment is the same as that in other embodied tasks in AI2-THOR, which is widely used in different embodied tasks. Each house contains 10 different meta-tasks, and each meta-tasks contains 5 different initial positions for heterogeneous agents.

We consider two different heterogeneous settings with different numbers of heterogeneous agents and different ability settings. In *Setting 1*, there are three agents with the same visual

TABLE 3
Quantitative Results in Heterogeneous Multi-agent Tidying-up Task in *Setting I*

Methods	<i>Success</i> (↑)			%PS(↑)			%FM(↑)			#PL(↓)			ACm(↓)			CES(↑)		
	Single	Cross	All	Single	Cross	All	Single	Cross	All	Single	Cross	All	Single	Cross	All	Single	Cross	All
SA	0.045	0.030	0.038	0.096	0.081	0.089	0.525	0.483	0.507	278.6	294.5	285.5	-	-	-	-	-	-
SA(Oracle)	0.090	0.068	0.080	0.181	0.163	0.173	0.703	0.681	0.693	184.8	197.1	190.1	-	-	-	-	-	-
Random	0	0	0	0	0	0	0.003	0	0.002	300.0	300.0	300.0	0	0	0	-	-	-
QMIX	0.032	0.015	0.025	0.082	0.069	0.076	0.472	0.421	0.450	293.1	300.0	296.1	512.0	512.0	512.0	0	0	0
CondComm.	0.081	0.066	0.075	0.176	0.151	0.165	0.712	0.709	0.710	192.3	263.1	223.1	523.1	501.3	513.6	0.7	0.7	0.7
CmprComm.	0.067	0.053	0.061	0.167	0.131	0.151	0.695	0.677	0.687	188.3	253.3	216.6	220.0	220.0	220.0	1.0	1.0	1.0
IntenComm.	0.047	0.032	0.040	0.136	0.101	0.121	0.645	0.609	0.629	214.6	267.5	237.6	20.0	20.0	20.0	1.0	1.0	1.0
Ours	0.121	0.079	0.103	0.194	0.160	0.179	0.739	0.726	0.733	186.1	252.6	215.0	363.1	366.5	364.6	2.1	1.4	1.8
BroadComm.	0.130	0.089	0.112	0.201	0.184	0.194	0.742	0.731	0.737	183.3	245.9	210.5	820.0	820.0	820.0	1.0	0.7	0.9
CentralComm.	0.131	0.092	0.114	0.205	0.166	0.188	0.755	0.739	0.748	183.2	243.1	209.3	820.0	820.0	820.0	1.0	0.8	0.9

TABLE 4
Quantitative Results in Heterogeneous Multi-agent Tidying-up Task in *Setting II*

Methods	<i>Success</i> (↑)			%PS(↑)			%FM(↑)			#PL(↓)			ACm(↓)			CES(↑)		
	Single	Cross	All	Single	Cross	All	Single	Cross	All	Single	Cross	All	Single	Cross	All	Single	Cross	All
SA	0.045	0.030	0.038	0.096	0.081	0.089	0.525	0.483	0.507	278.6	294.5	285.5	-	-	-	-	-	-
SA(Oracle)	0.090	0.068	0.080	0.181	0.163	0.173	0.703	0.681	0.693	184.8	197.1	190.1	-	-	-	-	-	-
Random	0	0	0	0	0	0	0.005	0	0.003	300.0	300.0	300.0	0	0	0	-	-	-
QMIX	0.035	0.021	0.029	0.091	0.076	0.084	0.483	0.425	0.458	290.5	296.3	292.9	512.0	512.0	512.0	0	0	0
CondComm.	0.090	0.072	0.082	0.179	0.156	0.169	0.720	0.711	0.716	188.5	230.7	206.9	620.6	613.1	617.3	0.7	0.7	0.7
CmprComm.	0.073	0.059	0.067	0.173	0.139	0.158	0.703	0.683	0.694	183.1	226.9	202.1	330.0	330.0	330.0	0.8	0.9	0.8
IntenComm.	0.048	0.033	0.041	0.139	0.112	0.127	0.661	0.620	0.643	210.1	255.3	229.8	30.0	30.0	30.0	1.0	1.0	1.0
Ours	0.129	0.090	0.112	0.203	0.169	0.188	0.745	0.739	0.743	166.9	221.3	190.6	472.5	478.3	475.0	1.8	1.3	1.6
BroadComm.	0.136	0.094	0.118	0.209	0.189	0.200	0.748	0.741	0.745	161.1	213.3	183.8	1230.0	1230.0	1230.0	0.7	0.5	0.6
CentralComm.	0.138	0.097	0.120	0.211	0.173	0.194	0.760	0.745	0.753	160.6	213.1	183.4	1230.0	1230.0	1230.0	0.8	0.5	0.7

perception ability but different action abilities and morphological characteristics. Specifically, $A^{(1)}$ only has the navigation ability with the low height, $Nav^{(1)} = 1, Mani^{(1)} = 0, Hei^{(1)} = 0$. $A^{(2)}$ only has the navigation ability and its height is high, $Nav^{(2)} = 1, Mani^{(2)} = 0, Hei^{(2)} = 1$. $A^{(3)}$ has both the navigation and manipulation abilities and its height is high, $Nav^{(3)} = 1, Mani^{(3)} = 1, Hei^{(3)} = 1$. In *Setting II*, there are four heterogeneous agents. The settings of $A^{(1)}$, $A^{(2)}$ and $A^{(3)}$ are the same as *Setting I*, and $A^{(4)}$ has the same abilities with $A^{(3)}$. The reasons that we select the *Setting I* and *Setting II* are that these two settings can represent the general heterogeneous collaboration scenarios, and the experimental results with these two settings can demonstrate the generalization of our model across different numbers and heterogeneous settings of agents. In *Setting I*, $A^{(3)}$ can pick up objects which are pickupable from the correct interaction locations. It is easier for $A^{(1)}$ with the lower field of view to find the objects thrown on the floor, and it can share its detected information to $A^{(3)}$ to help it pick up thrown objects. $A^{(2)}$ with the higher field of view can observe some areas that $A^{(1)}$ cannot, and its perspective can complement that of $A^{(1)}$ to assisting $A^{(3)}$ complete the tidying-up task. The visual perception of $A^{(1)}$ and $A^{(2)}$ with perspectives at different heights can complement each other during exploration, and $A^{(3)}$ can operate the misplaced objects with the assistance of $A^{(1)}$ and $A^{(2)}$ at different heights to complete the task. In *Setting II*, both $A^{(3)}$ and $A^{(4)}$ can pick up misplaced objects, and they can also cooperate with each other to improve the efficiency of putting objects to the reasonable locations. These two settings can solve different situations of this task. $A^{(3)}$ ($A^{(4)}$) can pick up misplaced objects detected but cannot be picked up by $A^{(1)}$ and $A^{(2)}$. The

different perspectives of $A^{(1)}$, $A^{(2)}$ can complement each other to easier find objects thrown on the floor. For each setting, there are a total of 4000 training tasks ($80 \times 10 \times 5$), 1000 validation tasks ($20 \times 10 \times 5$), and 1000 testing tasks. In testing tasks, there are 565 *Single* tasks and 435 *Cross* tasks.

6.2 Evaluation Metrics

We propose six metrics to evaluate the performance of the model, and N_{epi} denotes the number of tasks:

1) *Success (Suc)*: $Suc = \frac{1}{N_{epi}} \sum_{i=1}^{N_{epi}} R_i$, where $R_i = 1$ if all misplaced objects are re-placed in reasonable locations, otherwise $R_i = 0$. It denotes the strict success rate.

2) *%PatialSuccess (%PS)*: $\%PS = \frac{1}{N_{epi}} \sum_{i=1}^{N_{epi}} \frac{K_i^{suc}}{K_i}$, where K_i denotes the number of misplaced objects in the i -th task, and K_i^{suc} is the number of misplaced objects that are re-placed to reasonable locations by agents in this task. It can reflect the proportion of misplaced objects that can be successfully tidied in each task.

3) *%FindMisObj (%FM)*: $\%FM = \frac{1}{N_{epi}} \sum_{i=1}^{N_{epi}} \frac{K_i^{det}}{K_i}$, where K_i^{det} is the number of misplaced objects that are detected and picked up by agents in this task. It can reflect the effect of misplaced objects detector to some extent.

4) *#PathLength (#PL)*: $\#PL = \frac{1}{N_{epi}} \sum_{i=1}^{N_{epi}} Len_i$, where Len_i is the number of steps in the multi-agent trajectory when synchronously completing the i -th task, that is, the maximum number of steps of three agents. It evaluates the efficiency of completing the task. The fewer steps, the higher efficiency of the multi-agent collaboration.

TABLE 5
Ablation Experiments Results in Heterogeneous Multi-agent Tidying-up Task in *Setting I*

Methods	<i>Success</i> (↑)			%PS(↑)			%FM(↑)			#PL(↓)			<i>ACm</i> (↓)			<i>CES</i> (↑)		
	Single	Cross	All	Single	Cross	All	Single	Cross	All	Single	Cross	All	Single	Cross	All	Single	Cross	All
Ours w/o Know.	0.042	0.030	0.037	0.087	0.080	0.084	0.425	0.366	0.399	235.1	267.3	249.1	363.5	366.6	364.8	0	0	0
Ours w/o MisObjDec.	0.051	0.046	0.049	0.109	0.091	0.101	0.515	0.471	0.496	234.6	264.5	247.6	363.3	366.6	364.7	0.2	0.4	0.3
Ours w/o ReaRecPre.	0.063	0.058	0.061	0.120	0.098	0.110	0.551	0.509	0.533	231.1	260.3	243.8	363.3	366.9	364.9	0.5	0.8	0.6
Ours w/o Comm.	0.045	0.030	0.038	0.096	0.081	0.089	0.525	0.483	0.507	278.6	294.5	285.5	0	0	0	-	-	-
Ours w/o HierDec.	0.055	0.049	0.052	0.106	0.093	0.100	0.573	0.532	0.555	266.3	284.5	274.2	366.1	367.2	366.6	0.3	0.5	0.4
Ours	0.121	0.079	0.103	0.194	0.160	0.179	0.739	0.726	0.733	186.1	252.6	215.0	363.1	366.5	364.6	2.1	1.4	1.8

TABLE 6
Ablation Experiments Results in Heterogeneous Multi-agent Tidying-up Task in *Setting II*

Methods	<i>Success</i> (↑)			%PS(↑)			%FM(↑)			#PL(↓)			<i>ACm</i> (↓)			<i>CES</i> (↑)		
	Single	Cross	All	Single	Cross	All	Single	Cross	All	Single	Cross	All	Single	Cross	All	Single	Cross	All
Ours w/o Know.	0.049	0.033	0.042	0.103	0.096	0.010	0.452	0.375	0.419	227.0	258.5	240.7	473.3	479.2	475.9	0.09	0.06	0.08
Ours w/o MisObjDec.	0.058	0.052	0.055	0.118	0.103	0.111	0.523	0.480	0.504	225.5	256.1	238.8	473.1	479.0	475.7	0.3	0.5	0.4
Ours w/o ReaRecPre.	0.075	0.069	0.072	0.128	0.105	0.118	0.560	0.515	0.540	220.6	251.9	234.2	473.3	478.9	475.8	0.6	0.8	0.7
Ours w/o Comm.	0.049	0.033	0.042	0.105	0.091	0.099	0.535	0.496	0.518	273.1	280.6	276.4	0	0	0	-	-	-
Ours w/o HierDec.	0.063	0.055	0.060	0.111	0.99	0.106	0.581	0.537	0.562	260.1	273.6	266.2	475.3	478.8	477.3	0.4	0.5	0.4
Ours	0.129	0.090	0.112	0.203	0.169	0.188	0.745	0.739	0.743	166.9	221.3	190.6	472.5	478.3	475.0	1.8	1.3	1.6

5) *AvgCom* (*ACm*): $ACm = \frac{1}{N_{epi}} \sum_{i=1}^{N_{epi}} \frac{Total_i}{Len_i \cdot N}$, where $Total_i$ is the total dimensions of communication messages among agents. It evaluates the average communication amounts per agent per step.

6) *CommunicationEfficiencyScore* (*CES*): $CES = \max(0, \frac{10000 \cdot (Suc - SucSA)}{ACm})$, where $SucSA$ is the success rate of the single-agent model. It evaluates the proportion of efficiency improvement of different communication methods. The higher *CES*, the higher efficiency of the communication mechanism.

6.3 Quantitative Analysis

We evaluate the proposed model with different action decision strategies and communication methods in testing tasks. For different action strategies, we compare the results with *SA*, *SA(Oracle)*, *Random*, and *QMIX* model. *SA* only utilizes one agent with both navigation and manipulation abilities to complete the task, and the other parts are the same as the proposed model. In *SA(Oracle)*, the objects' locations are known to the agent, and the other parts are the same as *SA*. *Random* model randomly generates actions for each agent from their action spaces. *QMIX* is a MARL method based on [52], in which agents obtain all state features from others, fuse with their own state features, and directly predict the next low-level actions based on fused state features. For different communication methods, we consider several methods that save bandwidth. In *CondComm.*, only when an agent detects the misplaced object in the current view, it would broadcast its state and map to others. *CmprComm.* utilizes an encoder-decoder structure to compress the communication messages, and decode received messages from other agents, similar to that in [13]. In *IntenComm.*, agents exchange state and sub-task information, infer other agent's next sub-goal, and then decide their next sub-task, which is referred to as the idea in [53]. We also pay attention to the central and broadcast mechanisms, which require large communication amounts. *CentralComm.* adopts a central node to obtain the messages from all agents and generates the sub-task for each agent. In *BroadComm.*, each agent broadcasts its messages

to all agents and determines the next actions with its state and received messages from others. All these communication baselines are demonstrated in Fig. 6.

We show the quantitative results of *Single* and *Cross* tasks in *Setting I*, *Setting II* in Table 3 and Table 4 respectively. The best performance except for central and broadcast methods are bolded. The maximum number of steps per task is 300. The results show that the heterogeneous multi-agent tidying-up task is challenging. %*FM* is relatively high, but the success rate of the whole task is relatively low, which indicates that moving the detected misplaced objects to the reasonable location is challenging, since agents may not reason the proper receptacle or not generate the valid manipulation actions to re-place the misplaced object to predicted locations. The results in *Setting I* and *Setting II* indicate that our model can generalize to different numbers and heterogeneous settings of agents. Compared with multi-agent models, the success rate of task completion of *SA* and *SA(Oracle)* is lower, which demonstrates that collaboration and communication among heterogeneous multi-agents can help agents improve the accuracy and efficiency of tasks. Although agents in *QMIX* can obtain state features of other agents, its performance is relatively low, indicating that this MARL method is not effective enough to deal with visual information, and it is difficult to directly predict low-level actions.

As for different communication methods, the performances of *Ours*, *CondComm.*, *CmprComm.* and *IntenComm.* in task completion (*Suc*, %*PS*, %*FM*, #*PL*) are lower than *CentralComm.* and *BroadComm.*, because agents share partial information with others for decision-making. Since *CentralComm.* utilizes a central node to process information, its task completion performance can be regarded as the upper bound of multi-agent models. Our model performs better than other communication methods that share partial messages with agents in task completion. In our model, the agent with manipulation ability can obtain the state and map information from other agents with only the navigation ability, which is beneficial for the agent to put the object to a reasonable location, and its performance is better than the other three models. In *CondComm.*, agents exchange their state only when a misplaced

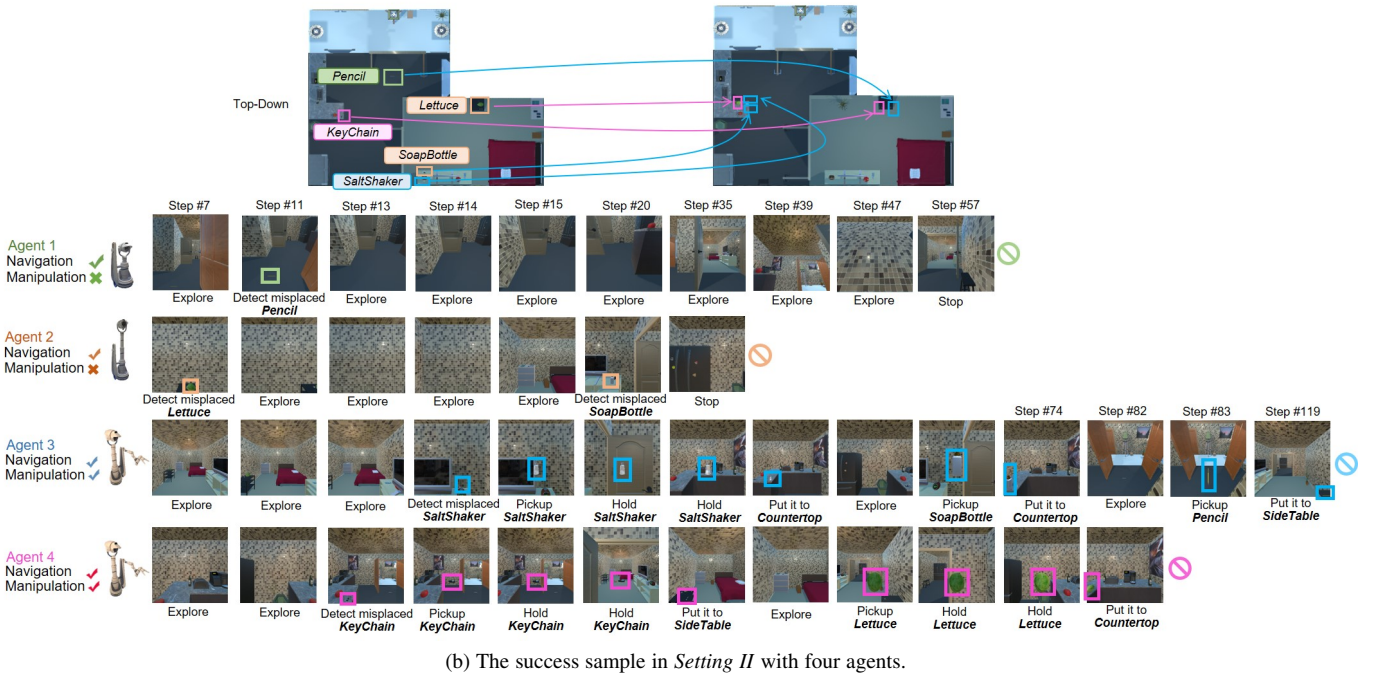
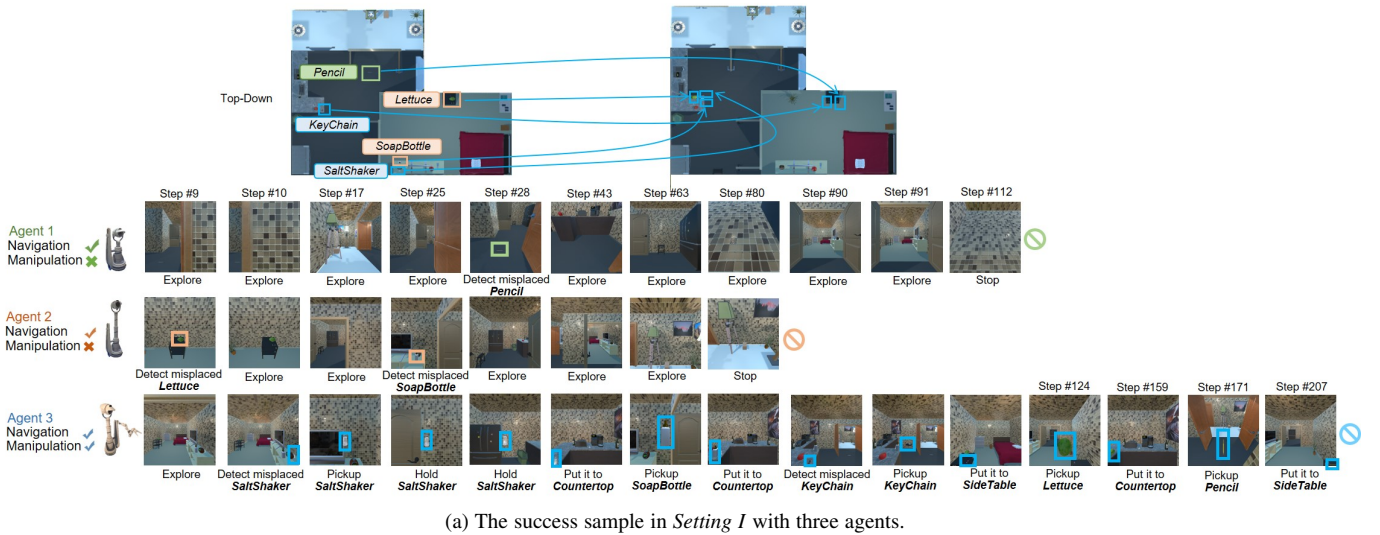


Fig. 7. The successful samples in the same house in *Setting I* and *Setting II*. In the top-down view, the color of bounding boxes in the left indicates which agent detects the specific misplaced objects, and the color of arrows indicates which agent puts the misplaced objects to reasonable locations.

object is detected, and the incomplete shared map information makes its performance lower than our model. The performance of *CmprComm.* indicates that decoding the compressed communication message would lose some effective information though the communication bandwidth can be reduced. The performance of *IntenComm.* is relatively low due to the improper reasoning results of other agents' decisions and incomplete exchanging of information. Furthermore, in *QMIX*, each agent receives the fused 512-dimensional state features from other agents, so ACm of *QMIX* is always 512. The size of local map embedding $sm_t^{(i)}$ is 20×20 . The transmitted state information contains the 3-dimensional $va^{(i)}$, 3-dimensional pose $pose_t^{(i)}$ and 4-dimensional sub-task information ($Place/Explore, o_t, p_t, r_t$), so the number of total dimensions of the state information from one agent is 10. Since *IntenComm.* only transmits the state information from other

agents to each agent, so ACm in *Setting I* is 20 and ACm in *Setting II* is 30. In *CmprComm.*, each agent compresses the dimensions of the map embedding to be 100, and transmits the compressed map embedding as well as the state information to other agents, so ACm in *Setting I* is $2 \times (100 + 10) = 220$, ACm in *Setting II* is $3 \times (100 + 10) = 330$. In *BroadComm.* and *CentralComm.*, each agent transmits their state information and map embeddings to all other agents, so ACm in *Setting I* is $2 \times (400 + 10) = 820$ and ACm in *Setting II* is $3 \times (400 + 10) = 1230$. ACm of *CentralComm.* and *BroadComm.* are the largest, and with the increased number of heterogeneous agents, their ACm have increased greatly, but the declines in SCE are relatively larger, which indicates that they are not cost-effective enough. Although *IntenComm.* uses the smallest ACm , its CES is not high, since its Suc is not good. Our method can save communication amounts, achieve

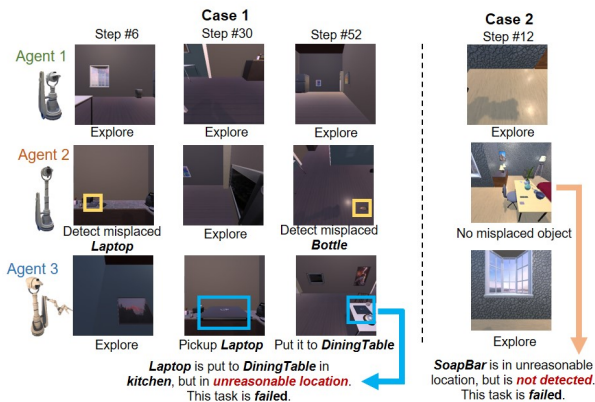


Fig. 8. The failed cases in *Setting I* of the task. In *Case 1*, Agent 3 place Laptop to the Kitchen, an unreasonable location. In *Case 2*, Agent 2 fail to detect SoapBar in the DiningTable in Bedroom.

good performance of task completion, and obtain the highest communication efficiency *CES*, demonstrating the effectiveness of the proposed communication mechanism.

6.4 Qualitative Analysis

We demonstrate successful samples with three agents in *Setting I* and four agents in *Setting II* respectively in Fig. 7. The tidying-up tasks in 7a and 7b are the same. Agents complete the same task more quickly in *Setting II*, since Agent 3 and Agent 4 with both the navigation and manipulation abilities can collaborate with each other to pick up misplaced objects and put them to reasonable locations, which demonstrates that the proposed model can make effective use of different abilities of heterogeneous to improve efficiency. We also show two failed cases in Fig. 8, where the agent places the misplaced object in an unreasonable location or agents fail to detect all the misplaced objects.

6.5 Ablation Experiments

To evaluate the role and effectiveness of different modules of our model, we conduct ablation experiments by removing the specific part of the model and comparing the performance. *Ours w/o Know.* removes the commonsense prior knowledge, which judge the misplaced objects and infer the reasonable locations directly from visual features. *Ours w/o MisObjDec.* removes the misplaced object detector, which extracts visual features and relationship features from the current observation as the subsequent input of the model. *Ours w/o ReaRecPre.* removes the reasonable receptacle predictor and directly extracts the receptacle features with the linear layer to input to the subsequent model to generate next actions. We remove the communication module to obtain *Ours w/o Comm.*, in which agents do not communicate with others and each agent performs actions independently. *Ours w/o HierDec.* removes the hierarchical decision, and agents do not predict the executing sub-tasks and sub-goals, but directly generate low-level actions with their state features. The results of the ablation experiments for *Single* and *Cross* tasks in *Setting I*, *Setting II* in Table 5 and Table 6 respectively.

The performance of *Ours w/o Know.* drops obviously, indicating that the commonsense knowledge is quite important for scene reasoning in this task. Since the misplaced object detector and the reasonable receptacle predictor can provide a clear target object to

be picked up and a clear receptacle to put the object for subsequent sub-tasks generating, the performance of *Ours w/o MisObjDec.* and *Ours w/o ReaRecPre.* would decrease. In *Ours w/o Comm.*, each agent has no communication with others, since in *Setting I* only one agent has the manipulation capability, the performance is the same as *SA* in *Setting I*. *Ours w/o HierDec.* directly generates the low-level actions with the detected misplaced object, the reasoned receptacle, and the state features. Because it is difficult to directly learn the relations between state features and executed low-level actions, this model would generate some invalid actions for the executing task, which leads to worse performance. Since the misplaced object detector and reasonable receptacle predictor are retained in *Ours w/o HierDec.*, its %*FM* is still higher than that of *Ours w/o MisObjDec.* and *Ours w/o ReaRecPre.*. The results demonstrate that each module of the proposed model is effective for heterogeneous multi-agent task completion.

7 CONCLUSION

In this paper, we propose the heterogeneous multi-agent collaborative framework based on the handshake-based group communication strategy and hierarchical decision model. To evaluate the effectiveness of the framework, we propose the heterogeneous multi-agent tidying-up task and generate a benchmark dataset for this task in houses with multiple rooms. The results demonstrate the effectiveness of each module of our model. In the future, we will continue to study the communication mechanism in more complex tasks and larger multi-agent systems with both homogeneous and heterogeneous agents.

ACKNOWLEDGMENTS

This work was supported in part by the National Natural Science Foundation of China under Grants 62025304.

REFERENCES

- [1] M. Wortsman, K. Ehsani, M. Rastegari, A. Farhadi, and R. Mottaghi, "Learning to learn how to learn: Self-adaptive visual navigation using meta-learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 6750–6759.
- [2] W. Yang, X. Wang, A. Farhadi, A. Gupta, and R. Mottaghi, "Visual semantic navigation using scene priors," *arXiv preprint arXiv:1810.06543*, 2018.
- [3] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, I. Reid, S. Gould, and A. Van Den Hengel, "Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3674–3683.
- [4] M. Shridhar, J. Thomason, D. Gordon, Y. Bisk, W. Han, R. Mottaghi, L. Zettlemoyer, and D. Fox, "Alfred: A benchmark for interpreting grounded instructions for everyday tasks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 10 740–10 749.
- [5] L. Weihs, M. Deitke, A. Kembhavi, and R. Mottaghi, "Visual room rearrangement," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
- [6] S. Omidshafiei, J. Pazis, C. Amato, J. P. How, and J. Vian, "Deep decentralized multi-task multi-agent reinforcement learning under partial observability," in *International Conference on Machine Learning*. PMLR, 2017, pp. 2681–2690.
- [7] J. K. Gupta, M. Egorov, and M. Kochenderfer, "Cooperative multi-agent control using deep reinforcement learning," in *International Conference on Autonomous Agents and Multiagent Systems*. Springer, 2017, pp. 66–83.
- [8] S. Iqbal and F. Sha, "Actor-attention-critic for multi-agent reinforcement learning," in *International conference on machine learning*. PMLR, 2019, pp. 2961–2970.

- [9] M. Zhou, Z. Liu, P. Sui, Y. Li, and Y. Y. Chung, "Learning implicit credit assignment for cooperative multi-agent reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 11 853–11 864, 2020.
- [10] U. Jain, L. Weihs, E. Kolve, M. Rastegari, S. Lazebnik, A. Farhadi, A. G. Schwing, and A. Kembhavi, "Two body problem: Collaborative visual task completion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 6689–6699.
- [11] U. Jain, L. Weihs, E. Kolve, A. Farhadi, S. Lazebnik, A. Kembhavi, and A. Schwing, "A cordial sync: Going beyond marginal policies for multi-agent embodied tasks," in *European Conference on Computer Vision*. Springer, 2020, pp. 471–490.
- [12] C. Yu, X. Yang, J. Gao, H. Yang, Y. Wang, and Y. Wu, "Learning efficient multi-agent cooperative visual exploration," in *European Conference on Computer Vision*. Springer, 2022, pp. 497–515.
- [13] X. Liu, D. Guo, H. Liu, and F. Sun, "Multi-agent embodied visual semantic navigation with scene prior knowledge," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3154–3161, 2022.
- [14] H. Wang, W. Wang, X. Zhu, J. Dai, and L. Wang, "Collaborative visual navigation," *arXiv preprint arXiv:2107.01151*, 2021.
- [15] S. Tan, W. Xiang, H. Liu, D. Guo, and F. Sun, "Multi-agent embodied question answering in interactive environments," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*. Springer, 2020, pp. 663–678.
- [16] X. Liu, X. Li, D. Guo, S. Tan, H. Liu, and F. Sun, "Embodied multi-agent task planning from ambiguous instruction," *Robotics: Science and Systems*, 2022.
- [17] C. Wakilpoor, P. J. Martin, C. Rebhuhn, and A. Vu, "Heterogeneous multi-agent reinforcement learning for unknown environment mapping," *arXiv preprint arXiv:2010.02663*, 2020.
- [18] W. U. Mondal, M. Agarwal, V. Aggarwal, and S. V. Ukkusuri, "On the approximation of cooperative heterogeneous multi-agent reinforcement learning (marl) using mean field control (mfc)," *Journal of Machine Learning Research*, vol. 23, no. 129, pp. 1–46, 2022.
- [19] V. Sharma, P. Goyal, K. Lin, G. Thattai, Q. Gao, and G. S. Sukhatme, "Ch-marl: A multimodal benchmark for cooperative, heterogeneous multi-agent reinforcement learning," *arXiv preprint arXiv:2208.13626*, 2022.
- [20] Y. Deng, D. Guo, X. Guo, N. Zhang, H. Liu, and F. Sun, "Mqa: Answering the question via robotic manipulation," *Robotics: Science and Systems*, 2021.
- [21] G. Sarch, Z. Fang, A. W. Harley, P. Schydlow, M. J. Tarr, S. Gupta, and K. Fragkiadaki, "Tidee: Tidying up novel rooms using visuo-semantic commonsense priors," in *European Conference on Computer Vision*, 2022.
- [22] Y. Kant, A. Ramachandran, S. Yenamandra, I. Gilitschenski, D. Batra, A. Szot, and H. Agrawal, "Housekeep: Tidying virtual households using commonsense reasoning," 2022.
- [23] M. Deitke, E. VanderBilt, A. Herrasti, L. Weihs, J. Salvador, K. Ehsani, W. Han, E. Kolve, A. Farhadi, A. Kembhavi, and R. Mottaghi, "Procthor: Large-Scale Embodied AI Using Procedural Generation," 2022.
- [24] A. Pal, Y. Qiu, and H. Christensen, "Learning hierarchical relationships for object-goal navigation," in *Conference on Robot Learning*. PMLR, 2021, pp. 517–528.
- [25] D. Fried, R. Hu, V. Cirik, A. Rohrbach, J. Andreas, L.-P. Morency, T. Berg-Kirkpatrick, K. Saenko, D. Klein, and T. Darrell, "Speaker-follower models for vision-and-language navigation," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [26] X. Wang, Q. Huang, A. Celikyilmaz, J. Gao, D. Shen, Y.-F. Wang, W. Y. Wang, and L. Zhang, "Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 6629–6638.
- [27] —, "Vision-language navigation policy learning and adaptation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 12, pp. 4205–4216, 2021.
- [28] B. Lin, Y. Zhu, Y. Long, X. Liang, Q. Ye, and L. Lin, "Adversarial reinforced instruction attacker for robust vision-language navigation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 10, pp. 7175–7189, 2022.
- [29] Y. Qiao, Y. Qi, Y. Hong, Z. Yu, P. Wang, and Q. Wu, "Hop+: History-enhanced and order-aware pre-training for vision-and-language navigation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 7, pp. 8524–8537, 2023.
- [30] D. S. Chaplot, D. Gandhi, S. Gupta, A. Gupta, and R. Salakhutdinov, "Learning to explore using active neural slam," *arXiv preprint arXiv:2004.05155*, 2020.
- [31] Y. Qi, Q. Wu, P. Anderson, X. Wang, W. Y. Wang, C. Shen, and A. v. d. Hengel, "Reverie: Remote embodied visual referring expression in real indoor environments," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9982–9991.
- [32] A. Das, S. Datta, G. Gkioxari, S. Lee, D. Parikh, and D. Batra, "Embodied question answering," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1–10.
- [33] D. Gordon, A. Kembhavi, M. Rastegari, J. Redmon, D. Fox, and A. Farhadi, "Iqa: Visual question answering in interactive environments," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4089–4098.
- [34] E. Wijmans, S. Datta, O. Maksymets, A. Das, G. Gkioxari, S. Lee, I. Essa, D. Parikh, and D. Batra, "Embodied question answering in photorealistic environments with point cloud perception," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 6659–6668.
- [35] H. Luo, G. Lin, Y. Yao, F. Liu, Z. Liu, and Z. Tang, "Depth and video segmentation based visual attention for embodied question answering," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 6, pp. 6807–6819, 2023.
- [36] S. Tan, M. Ge, D. Guo, H. Liu, and F. Sun, "Knowledge-based embodied question answering," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–13, 2023.
- [37] Y. Zhang and J. Chai, "Hierarchical task learning from language instructions with unified transformers and self-monitoring," *arXiv preprint arXiv:2106.03427*, 2021.
- [38] A. Gao, H. Liu, Y. Hu, W. Liang, and S. X. Ng, "Cooperative cache in cognitive radio networks: A heterogeneous multi-agent learning approach," *IEEE Communications Letters*, vol. 26, no. 5, pp. 1032–1036, 2022.
- [39] M. Brittain and P. Wei, "Scalable autonomous separation assurance with heterogeneous multi-agent reinforcement learning," *IEEE Transactions on Automation Science and Engineering*, 2022.
- [40] S. Yang, B. Yang, Z. Kang, and L. Deng, "Ihg-ma: Inductive heterogeneous graph multi-agent reinforcement learning for multi-intersection traffic signal control," *Neural networks*, vol. 139, pp. 265–277, 2021.
- [41] J. Chen, Y. Gao, J. Hu, F. Deng, and T. L. Lam, "Meta reinforcement learning based sensor scanning in 3d uncertain environments for heterogeneous multi-robot systems," *arXiv preprint arXiv:2109.13617*, 2021.
- [42] S. Patel, S. Wani, U. Jain, A. G. Schwing, S. Lazebnik, M. Savva, and A. X. Chang, "Interpretation of emergent communication in heterogeneous collaborative embodied agents," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15953–15963.
- [43] A. Das, T. Gervet, J. Romoff, D. Batra, D. Parikh, M. Rabbat, and J. Pineau, "Tarmac: Targeted multi-agent communication," in *International Conference on Machine Learning*. PMLR, 2019, pp. 1538–1546.
- [44] Y.-C. Liu, J. Tian, N. Glaser, and Z. Kira, "When2com: Multi-agent perception via communication graph grouping," in *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 2020, pp. 4106–4115.
- [45] E. Kolve, R. Mottaghi, W. Han, E. VanderBilt, L. Weihs, A. Herrasti, D. Gordon, Y. Zhu, A. Gupta, and A. Farhadi, "AI2-THOR: An Interactive 3D Environment for Visual AI," *arXiv*, 2017.
- [46] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, *et al.*, "Visual genome: Connecting language and vision using crowdsourced dense image annotations," *International journal of computer vision*, vol. 123, no. 1, pp. 32–73, 2017.
- [47] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.
- [48] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Conference Track Proceedings*, 2017.
- [49] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [50] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2961–2969.
- [51] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [52] T. Rashid, M. Samvelyan, C. S. De Witt, G. Farquhar, J. Foerster, and S. Whiteson, "Monotonic value function factorisation for deep multi-agent reinforcement learning," *The Journal of Machine Learning Research*, vol. 21, no. 1, pp. 7234–7284, 2020.

- [53] Y. Wang, F. Zhong, J. Xu, and Y. Wang, "Tom2c: Target-oriented multi-agent communication and cooperation with theory of mind," *arXiv preprint arXiv:2111.09189*, 2021.