

Learning Hierarchical Interactive Multi-Object Search for Mobile Manipulation

Fabian Schmalstieg*, Daniel Honerkamp*, Tim Welschehold, and Abhinav Valada

Abstract—Existing object-search approaches enable robots to search through free pathways, however, robots operating in unstructured human-centered environments frequently also have to manipulate the environment to their needs. In this work, we introduce a novel interactive multi-object search task in which a robot has to open doors to navigate rooms and search inside cabinets and drawers to find target objects. These new challenges require combining manipulation and navigation skills in unexplored environments. We present HIMOS, a hierarchical reinforcement learning approach that learns to compose exploration, navigation, and manipulation skills. To achieve this, we design an abstract high-level action space around a semantic map memory and leverage the explored environment as instance navigation points. We perform extensive experiments in simulation and the real world that demonstrate that, with accurate perception, the decision making of HIMOS effectively transfers to new environments in a zero-shot manner. It shows robustness to unseen subpolicies, failures in their execution, and different robot kinematics. These capabilities open the door to a wide range of downstream tasks across embodied AI and real-world use cases.

I. INTRODUCTION

Autonomous navigation and exploration in unstructured indoor environments require a large variety of skills and capabilities. Pathways may be blocked and objects of interest may be stored away. Thus far, existing multi-object search tasks and methods have focused on environments that can be freely navigated with openly visible target objects [1]–[3]. However, as we move to human-centered environments such as households, neither of these assumptions hold. We introduce a novel *Interactive Multi-Object Search* task in which target objects may be located inside articulated objects such as drawers and closed doors have to be opened to explore the environment. As a result, only navigation is insufficient to accomplish the task and the robotic agent has to physically interact with the environment to manipulate it to its needs.

Multi-object search tasks pose long-horizon problems with non-trivial optimal policies. Methods such as frontier exploration [4] offer guarantees to explore the entire environment if given enough time. However, they often do not take the context of the environment into account and result in long far from optimal paths while moving from one frontier point to the next. On the other end of the spectrum, learning-based methods can take unstructured observations into account and have been shown to learn good exploration strategies [1], [3], but they struggle with the long-horizon nature of the task. Moreover, since the robot also has to interact with the environment, both the action space

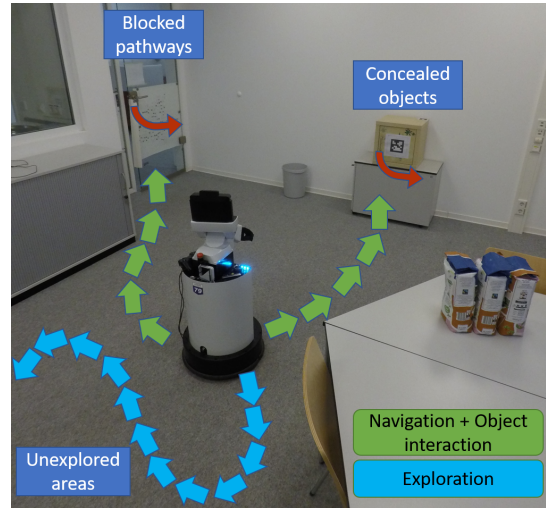


Fig. 1. We introduce the *Interactive Multi-Object Search* task in which an agent has to autonomously search and manipulate the environment to find a set of target objects. To succeed, the agent has to free pathways by opening doors and searching inside articulated objects such as cabinets and drawers.

and task horizon increase even further, and existing exploration methods are insufficient to accomplish the task.

In this work, we propose Hierarchical Interactive Multi-Object Search (HIMOS), a hierarchical, reinforcement learning-based approach to learn both exploration and manipulation skills and to reason at a high level about the required steps. Reinforcement learning has shown to work well in unexplored environments with high-dimensional observation spaces [5], [6]. We combine learned motions for local exploration in continuous action spaces [3] and frontier exploration for long-horizon exploration [4] together with mobile manipulation skills for object interactions [7]. We use semantic maps as the central memory component, which have shown to be an expressive and sample-efficient representation for these tasks [3] and design a high-level action space that exploits the acquired knowledge about the environment. By leveraging explored object instance locations as navigation waypoints, our approach efficiently learns these complex tasks from little data and consistently achieves success rates above 90% even as the number of target objects increases. By equipping all the low-level skills with mobility, we remove the “hand-off” problem in which subpolicies have to terminate in the initial set of the following skill [8], [9]. Lastly, we transfer the trained agent to the real world and demonstrate that it successfully accomplishes these tasks in a real office environment. In particular, we replace the subpolicies from simulation with unseen real-world variations and find that the policy is able to generalize to these unseen subpolicies and is robust to failures in their execution,

*These authors contributed equally.

Department of Computer Science, University of Freiburg, Germany.

This work was funded by the European Union’s Horizon 2020 research and innovation program under grant agreement No 871449-OpenDR. Toyota Motor Europe supported this project with an HSR robot for experiments.

This paper provides a supplementary at <http://arxiv.org/abs/2307.06125>.

making it highly modular and flexible for transfer. Finally, we present ablation studies to evaluate the impact of the main design decisions.

To summarize, the following are the main contributions:

- 1) We propose an interactive multi-object search task that requires physical interactions with articulated objects, opening doors, and searching in cabinets and drawers.
- 2) We present a hierarchical, reinforcement learning-based approach that combines exploration and manipulation skills based on semantic knowledge and instance navigation points to efficiently solve these long-horizon tasks.
- 3) We demonstrate the capabilities of this approach in both simulated and real-world experiments and show that, given accurate perception, its decision making achieves zero-shot transfer to the real world, unseen environments, unseen subpolicies, and is robust to unseen failures.
- 4) We make the code for both the task and models publicly available at <http://himos.cs.uni-freiburg.de>.

II. RELATED WORK

Object search tasks: Exploration and the ability to find items of interest is a key requirement for a wide range of downstream tasks. Previous works have proposed methods to maximize coverage of explored space [10] and to find specific objects of interest based on vision [5], auditory signals [6], [11] or target object categories [12], [13]. In ordered multi-object search tasks, the agent has to find k items in a fixed order in game environments [14] or realistic 3D apartments [1]. In unordered multi-object search, the agent simply has to find the target objects as fast as possible, irrespective of the order [2], [3]. We focus on this unordered task. As we aim to demonstrate our system on a real robot, we follow Schmalstieg et al. [3] and use the full continuous action space. This is in contrast to most previous works which focus on a simplified granular discrete action space. Existing search tasks assume that the desired goals can be freely reached by the agent. The interactive navigation task [15] relaxes this assumption by placing objects that the robot has to push away to reach the target. In contrast, we introduce an interactive search task for mobile robots equipped with a manipulator, that requires interaction with articulated objects to clear the path or reveal concealed objects. This requires integrating navigation and manipulation. Lastly, in contrast to most previous works, we demonstrate that our approach successfully transfers to the real world.

Exploration requires both understanding and memorizing the seen environment and decision-making to explore the remaining space. Previous work has introduced both implicit and explicit memory mechanisms. Implicit memory agents either learn a direct end-to-end mapping from RGB-D images to actions or store embeddings of previous observations and retrieve them with an attention mechanism [2]. Other methods build explicit maps of the environment by projecting the RGB-D inputs into a global map. Commonly, this map is also annotated with semantic labels [1], [3], [14]. Further, combining short- and long-term exploration by learning an auxiliary prediction of the direction to the next closest object has

proven to result in a strong performance in continuous action spaces [3]. We use this approach for low-level exploration. Frontier exploration [4] samples points on the frontier of the explored space and then navigates to these points. Instead of sampling, Ramakrishnan et al. [16] predict a potential function towards a target object. SGoLAM [17] combines mapping and a goal detection module. If no goal object is detected, it explores with frontier exploration and navigates directly to the goal otherwise. This results in strong results without any learning component. We include a similar exploration strategy as an option in our hierarchical approach. Zheng et al. [18], [19] leverage POMDP-solvers over an explicit belief representation to plan next viewpoints.

Articulated object manipulation such as opening doors and drawers requires control of both the base and arm of the robot [20]. Existing approaches often separate both aspects and execute sequential navigation and manipulation. In our evaluation in simulation, we follow this approach and use BiRRT [21] to generate a motion plan for the robot arm. Recent works train an agent to control the base of the robot via reinforcement learning to follow given end-effector motions [7], [22]. We use this method in our real world evaluation as it generalizes across different robots, tasks, and environments.

Hierarchical methods introduce layers of abstraction by decomposing the decision-making into higher and lower-level policies. This shortens the time horizon of the Markov decision process (MDP) for the higher-level policies and enables the agent to combine different modules or skills at lower levels. At the same time, joint training of low- and high-level policies is often unstable and hard to optimize [23]. We focus on combining pretrained subpolicies. Joint finetuning of these policies offers a path to further performance improvements in the future. Pretraining is a common approach to increase the stability of the policy in hierarchical reinforcement learning [24]. While naive skill-chaining of arbitrary skills often results in “hand-off” failures in which the subsequent skill cannot start from the current state [9], we resolve this issue by adding mobility to all the low-level skills in the real world execution, without the need for region-rewards [8]. A common navigation abstraction is to learn to set waypoints [25], [26], however, these often end up as very near points to the agent. In contrast, we propose instance navigation that provides a prior on important locations and action granularity. ASC [27] learns to combine navigation and pick skills for given receptacle locations. In contrast, we learn to search in unexplored environments with objects hidden in articulated objects. Alternatively, behavior trees are often used to decompose tasks hierarchically into a tree structure. These trees can either be fully constructed manually or be used in conjunction with a planner [28]. [29] learn values of tree nodes for cluttered object relocation. In contrast to these approaches, our proposed HIMOS learns high-level decision-making with a two-layer hierarchy.

III. INTERACTIVE MULTI-OBJECT SEARCH

We propose an interactive object search task in which a robotic agent with a mobile base and a manipulator arm is randomly spawned in an unexplored indoor environment. The

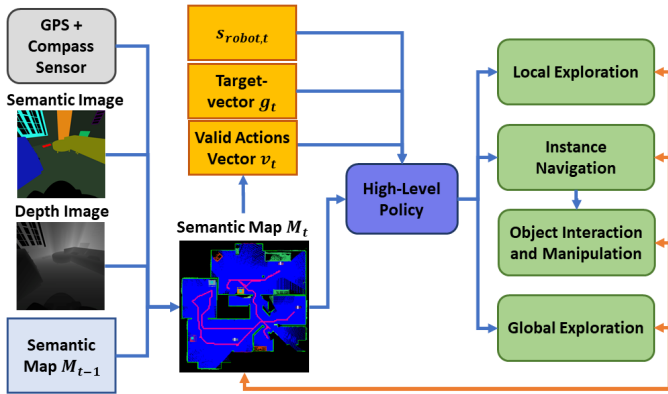


Fig. 2. Schematic overview of HIMOS. A semantic map M_t serves as a central memory component and is used and updated across low- as well as high-level modules. This map is extended to a partial panoptic map with instance IDs of relevant objects. Given the remaining objects g_t to find, the robot state $s_{robot,t}$, and the derived valid actions v_t , the high-level policy acts in an abstract action space consisting of local and global exploration, navigation to mapped object instances, and a mobile manipulation policy.

agent receives a goal vector that indicates k objects out of c categories that it has to find. The episode is considered successful if the agent finds all k objects, where an object is considered found when the agent has seen the object and navigated within a distance of 1.3m of it. The episode is terminated early if the agent exceeds 1,000 timesteps. To succeed, the agent has to explore the space while opening doors that block the way and opening the cabinets that contain the target objects. This results in very long-horizon tasks with complex shortest paths, as in contrast to previous multi-object search tasks [1]–[3], [14], the agent needs to manipulate its environment to achieve its goals.

The agent is acting in a goal-conditional Partially Observable Markov Decision Process (POMDP) $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{O}, T(s'|s, a), P(o|s), R(s, a, g))$, where \mathcal{S} , \mathcal{A} and \mathcal{O} are the state, action, and observation spaces, $T(s'|s, a)$ and $P(o|s)$ describe the transition and observation probabilities and $R(s, a, g)$ is the reward function. The objective is to learn a policy $\pi(a|o, g)$ that maximises the discounted, expected return $\mathbb{E}_\pi[\sum_{t=1}^T \gamma^t R(s_t, a_t, g)]$, where γ is the discount factor. In each step, the agent receives a visual observation o from an RGB-D and semantic camera, together with its current pose in the environment and a binary vector g indicating the objects it has to find. The true state s of the environment is unknown and can only be inferred from its observations.

IV. HIERARCHICAL INTERACTIVE MULTI-OBJECT SEARCH

The challenges introduced by this task require the robot to master an increasing number of different behaviors and skills. To stem this increase in complexity, we propose a hierarchical, reinforcement-based learning approach and design efficient abstractions over states and actions. The model consists of three main components: a mapping module, a set of subpolicies for exploration, navigation, and interaction, and a high-level policy module to select the next low-level action to execute. An overview of the approach is depicted in Figure 2.

Assumptions: The focus of this work is on learning long-horizon decision-making, exploration, and search. We abstract from real-

world perception and assume to have access to the following: (i) a semantic camera that produces accurate semantic object labels, (ii) accurate depth and localization, (iii) an object detection module to perceive object poses in the environment and to detect whether an object interaction was successful. In training and simulation, the object detection module uses ground truth poses from the simulator while in real-world experiments, we rely on AR markers placed on the respective objects. Interaction failure is detected based on the change in pose of these markers after the interaction. (iv) the environment remains static except for the agent’s own interactions with it. For example, objects are not being displaced by other people during the search process.

A. Mapping Module

A semantic map of the environment serves as a central memory component across all policies and modules. To construct this map, we build upon our previously introduced mapping module [3]. We extend this map to a partial panoptic segmentation [30] map by labeling task-specific objects such as doors and cabinets with an instance-specific color, which is randomly assigned whenever a new instance is detected. I.e. an instance’s color changes from episode to episode, but remains consistent within an episode. The robot receives the semantic masks and depth at a resolution of 128×128 pixels, generates a point cloud from this information, and projects the points into a local top-down map, using the top-most point for each cell. With the local map and its current localization, the agent then updates an internal global map which is further annotated with the agent’s trace and encoded into an RGB image. Target objects are mapped as a special color, making the approach agnostic to the underlying classes of the target objects. After the agent has segmented and approached a target object, it updates the object’s annotation with a fixed color coding to mark the corresponding object as “found”. In the simulation, the robot receives the semantic labels from the simulator’s semantic camera. In the real world, we simulate access to semantic labels as follows: we pre-build a map of the environment using Hector SLAM [31] and annotate it with labels for target objects, cabinets, and doors. All the other occupied space is mapped to the wall category. At test time, we then use the robot’s depth camera to build the same local map as in the simulation and overlay it with this pre-annotated semantic map. Localization in the real world is done on the same pre-recorded map.

In each step, the agent then extracts an egocentric map from the global map and passes two representations of this map to the encoder: a coarse map of dimension $224 \times 224 \times 3$ at a resolution of 6.6 cm and a fine-grained map of dimension $84 \times 84 \times 3$ at a resolution of 3.3 cm. I.e., they cover $14.8 \text{ m} \times 14.8 \text{ m}$ and $2.77 \text{ m} \times 2.77 \text{ m}$, respectively.

B. Subpolicy Behaviors

In this section, we describe the different subpolicy behaviors that are available to the high-level policy.

Local Exploration: Object search requires smart movement for local navigation and efficient exploration around corners and

corridors. For this, we use our previously introduced exploration policy [3]. The policy receives the semantic map, robot state, and target objects and predicts the direction to the next closest target object. It then communicates this prediction to a reinforcement learning agent which produces target velocities for the base of the robot. The local exploration policy is pretrained on the multi-object-search task [3] in the same train/test split. However, we change the robot to Fetch, adapt the collision penalty from -0.1 to -0.15 , and include open doors in the scenes. After the exploration policy is trained, it is kept frozen during high-level policy training. We adjust the panoptic labels provided by the mapping module on the fly to match the simpler, instance-unaware, semantic map that this subpolicy was trained with. In particular, doors and cabinets are colored as obstacles for the subpolicy. This shields the subpolicy from information that is not required to solve the downstream task. During training, when selected by the high-level policy, the exploration policy is executed for four time steps, giving the high-level policy control to quickly react to new knowledge of the environment. During the evaluation, we found it beneficial to execute it for a longer period of 20 time steps.

Global Exploration: While the local exploration policy has been shown to produce efficient search behavior, it can struggle to navigate to faraway areas. This, however, is a strength of frontier exploration [4], which samples points at the frontier to (often far away) unexplored areas. While frontier exploration on its own can lead to long inefficient paths, the high-level policy can learn to select the appropriate exploration strategy for the current context. In each iteration, the sampled frontier point is drawn onto the map, allowing the high-level policy to observe where it would navigate to before deciding which subpolicy to execute. Frontier points that lie outside the range of the agent’s egocentric map are projected onto a circle around the agent and marked in a different color, indicating that it is potentially a long-distance navigation. If selected by the high-level policy, the agent uses its navigation policy described below to navigate to the frontier point.

Instance Navigation: Learning navigation at the right level of abstraction can be challenging. Approaches such as setting waypoints are often difficult to optimize or decay to only selecting nearby points, removing the benefits of the abstraction. Instead, we leverage the acquired knowledge about the environment by using object instances as navigation points: the high-level policy can directly navigate to the discovered object instances by selecting their instance ID (for simplicity restricted to target objects, doors, and cabinets). We implement this action space as a one-hot encoding that maps to instance colors on its map (this assumes a maximum number of instances). We furthermore find it beneficial for learning speed [32] to only allow the agent to navigate to doors or cabinets that have not been successfully opened yet (cf. invalid action masking below). While less fine-grained than arbitrary waypoints, this results in an efficient set of navigable points across the map that, as we demonstrate in Section V, is well optimizable and results in a strong final policy.

The respective navigation goals are set to a pose slightly in front, or for the goal objects directly to the detected pose

of the corresponding object. This navigation goal is then fed to an A^* -planner which produces a trajectory at a resolution of 0.5 m. For training speed, we do not execute the full path in simulation, but set the robot’s base pose to the generated waypoints and only collect observations from these points. In the real world, we use the ROS navigation stack to move the robot to the goal. The policy may fail in some situations, for example, due to collision with obstacles or narrow doorways. In this case, the agent returns to the last feasible waypoint, and given the updated semantic map M_t , the high-level policy has to make a new decision.

Object Interaction and Manipulation: If the high-level policy chooses to navigate to a closed door or cabinet, this automatically triggers an interaction action that is executed once the navigation has been successfully completed. For fast simulation, we train the agent with magic actions that either open the object successfully or fail with a probability of 15% and leave the object untouched, in which case the agent has to decide whether to try to open it again. The training with failure cases enables it to learn a re-trial behavior to recover from failed attempts.

At test time, the agent has to physically execute the interactions. In the simulation, we replace the interaction subpolicy with a BiRRT motion planner [21] and inverse kinematics to execute a push-pull motion. The success of these motions depends on the pose of the robot and the object. Implementation details can be found in the supplementary material. In the real world, we replace these subpolicies with the N^2M^2 mobile manipulation policy [7]. Given the pose of the object handle (based on an AR marker) and the object label (door, drawer, or cabinet), it generates end-effector motions learned from demonstrations to open the object together with base commands that ensure that these motions remain kinematically feasible. The model is pretrained without additional retraining or finetuning. The agent returns a success indicator to the high-level policy and in case of failure, the agent again has to decide whether to repeat the interaction.

C. High-level Decision Making

Efficient high-level decision making requires the right level of abstractions of states and actions. We hypothesize that object- and instance-level decision making is such an efficient level of abstraction for embodied search tasks. We design a high-level policy around this idea. In particular: (i) We propose an instance navigation subpolicy that leverages the agent’s accumulated knowledge about the environment. It provides a prior on important places and on the granularity of navigation points, making it data-efficient and well-optimizable. (ii) As objects are discrete instances, the resulting full action space remains discrete, avoiding the complexities of mixed action spaces. At the same time, all the subpolicies still act directly in continuous action spaces, allowing for direct transfer to real robotic systems. (iii) We abstract from reasoning about exact robot placements in the real world by shifting the responsibility of mobility into the subpolicies. This ensures that the subpolicies can start from a large set of initial positions, resolving the “hand-off” problem from naive skill-chaining [9] and strongly simplifies the learning process for the high-level

policy. Furthermore, it enables us to change out the subpolicies to unseen subpolicies in the real world. (iv) We incorporate subpolicy failures into the training process, enabling the high-level policy to learn a retrial behavior if execution fails.

Action Space: The high-level policy acts in a Semi-Markov Decision Process (SMDP) in which the actions model temporarily extended behaviors and act at irregular intervals [33]. It makes a decision whenever the last invoked subpolicy returns, successful or not. The high-level action space consists of (i) invoking the local exploration policy [3], (ii) invoking global frontier exploration [4], (iii) instance navigation with subsequent object interaction (if available), where instance IDs are selected through a one-hot vector mapping to fixed colors. The current task instantiation assumes a maximum of ten instances per episode, resulting in an overall 12-dimensional discrete action space.

Adaptive Discounting: The high-level policy acts at irregular intervals, as the duration of the subpolicies varies largely. We correct for this time bias and accurately reflect the long-term consequences of actions with adaptive discounting.

Invalid Action Masking: The availability of high-level actions varies with the state of the environment, e.g. navigating to an object instance depends on the instance being mapped, opening a cabinet is only possible if it is mapped and closed. We infer a valid actions vector v_t from the agent’s observations and incorporate it into the training by masking out invalid actions as well as including it in the observation space of the agent. As a result, the agent can learn more effectively, speeding up the training process [32]. We implement this by replacing the logits of the invalid actions with a large negative number.

Observation Space and Architecture: The observation space of the high-level policy consists of the coarse semantic map $\in \mathcal{R}^{224 \times 224 \times 3}$, the c -dimensional target object vector g_t and the i -dimensional binary valid actions vector v_t (see invalid action masking), where i is the number of instances, and a 22-dimensional robot state vector $s_{robot,t}$. The robot state consists of linear and angular base velocities, the last low-level actions, sum of collisions over the last ten steps, a current collision flag, and a normalized history over the last 16 high-level actions taken. Following previous work [3], [10], the agent first encodes the coarse semantic map with a ResNet-18 [34]. Then it concatenates these features with the other, structured observations. The high-level policy is trained with Proximal Policy Optimization (PPO) [35]. We report hyperparameters and full network architectures in the supplementary material.

Reward Functions: The high-level policy is trained with the accumulated rewards of the invoked subpolicies. The subpolicies collect the following rewards: (i) A sparse positive reward of +10 for finding a target object, (ii) A sparse positive reward of +3 for opening a door, (iii) A penalty of -0.1 per collision for navigation policies, (iv) A negative traveled distance reward to encourage the high-level policy to find efficient compositions of the subpolicies. As the navigation policy does not get physically executed during training, we set it to -0.05 for each invocation of the local exploration policy

and to $-0.05 \times$ number of waypoints for the navigation policy. This results in a similar penalty per distance traveled.

V. EXPERIMENTAL EVALUATIONS

We extensively evaluate our approach both in simulation and real-world experiments. We aim to answer the following: I) Does the high-level policy learn to make decisions that lead to efficient exploration of the environment, improving over alternative decision rules?

II) What is the impact of the different subpolicies, in particular the local and global exploration policies?

III) Does the learned behavior transfer to the real world and to execution with different subpolicies?

IV) Is the overall system capable of successfully solving extended tasks involving many physical interactions within a single episode in the real world?

A. Experimental Setup

We instantiate the task in the iGibson simulator [36], which builds on the PyBullet physics engine. Each scene contains three cabinets placed randomly across a set of feasible locations. All the doors in the scene are initially in a closed state. We then construct tasks of finding 1-6 target objects, matching the hardest setting in previous work [1], [3]. We randomly place up to three target objects across the free space of the entire apartment and up to three objects inside the cabinets. We use the same eight training scenes as the iGibson challenge and use the remaining seven apartments for evaluation. The embodied agent is a Fetch robot, equipped with a mobile base with a differential drive, a height-adjustable torso, and a 7-DoF arm. We scale the robot’s size by a factor of 0.85 to navigate the narrow corridors of all apartments. Its raw 10-dimensional action space consists of a continuous linear and angular velocity for the base together with the torso and arm-joint velocities. The robot is equipped with an RGB-D camera with a field of view of 79 degrees and a maximum depth of 5.6 m.

B. Baselines

We compare our approach against different high-level decision-making modules and ablations of the action space.

Greedy: A greedy high-level decision-making strategy that immediately drives to any newly mapped task object (door, cabinet, target object) if available and otherwise selects either the local or global exploration policies with equal probability. *SGoLAM+:* SGoLAM [17] combines non-learning based approaches to achieve very strong performance on the CVPR 2021 MultiOn challenge. It explores the map with frontier exploration until it localizes a target object, then switches to a planner to navigate to the target. We reimplement the author’s approach for continuous action spaces and directly use the semantic camera for goal localization which further improves the performance. We then modify the action execution to open doors and cabinets when applicable.

HIMOS: The hierarchical approach presented in Section IV. *w/o frontier* removes global exploration from the subpolicy set. *w/o expl* removes local exploration from the subpolicy set.



Fig. 3. Example trajectories of HIMOS in unseen apartments. Black: unexplored, blue: free space, green: walls, red: agent trace, grey: (found) target objects, other colors: miscellaneous objects. Bottom right: the agent failed to find the last object, marked by the red circle, in the given time.

w/o IAM removes the invalid action masking and instead penalizes the agent with -2.5 for selecting invalid actions.

Metrics: We evaluate the models’ ability to find all the desired objects using the success rate and we evaluate the optimality of the search path with the success-weighted path length (SPL) [37]. In the simulation, we evaluate 25 episodes per scene, the number of target objects, and report the average over three random training seeds. This results in a total of $25 \cdot 6 \cdot 8 \cdot 3 = 3600$ episodes for seen and $25 \cdot 6 \cdot 7 \cdot 3 = 3150$ episodes for unseen apartments for each approach.

C. Simulation Experiments

To test the models’ abilities to learn to complete the tasks, we first evaluate them in the seen apartments for variable numbers of target objects. The results are reported in Table I (left). We find that all the compared models achieve good success rates. All three high-level decision-making variations, greedy, SGoLAM+, HIMOS, are able to make reasonable decisions, demonstrating the benefits of the design of the high-level abstractions discussed in Section IV-C. Furthermore, our proposed method, HIMOS, further improves over the baselines, consistently achieving the highest success rate and the most efficient paths, as measured by the SPL metric.

We then evaluate the models in the unseen apartments. Note that neither the low- nor high-level policies have seen these scenes during training. The results are shown in Table I (right). We find that the models learned to generalize without any clear generalization gap. The performance is even slightly higher than on the seen apartments, this is in accordance with previous observations [3]. This may be due to the validation split containing potentially simpler scenes than the training scenes. The evaluations on unseen scenes confirm the observations from the training scenes: HIMOS consistently achieves the highest success rate and the best SPL across all the numbers of target objects. Finally, we find that our hierarchical approach scales very well to longer scenarios, with a very small drop in success rates as the number of target objects increases. To find all six objects, the agent often has to explore the majority of the apartments and interact with a large number of objects.

D. Ablation Study

Exploration Subpolicies: To evaluate the impact of the exploration policies, we compare HIMOS to *w/o frontier* and *w/o expl*. We find frontier exploration to have a large impact on success rates. Removing this component reduces the success rate to 66.5%. Removing the local exploration policy leads to a smaller, but nonetheless significant drop of 12.0 ppt in average success rates as well as a clear drop in SPL. This indicates the different strengths of the two exploration behaviors, as well as that HIMOS learned to use the local exploration policy to increase search efficiency. Again, we find this effect to be consistent across both seen and unseen apartments.

Invalid Action Masking: We observe that removing the invalid action masking also leads to a drop in both success rates and SPL. Furthermore, we found that masking improves convergence speed by up to 2.5 times. Note that the action masking does not use any further privileged information beyond our perception assumptions (Section IV) of inferring object states.

Qualitatively, we find that the agent learned sensible behaviors for the task at hand. Figure 3 depicts example episodes in the unseen scenes. The agent learned to frequently invoke the local exploration policy while the apartment is still largely unexplored, to then use the global exploration policy to navigate to unexplored corners where target objects could be hidden. Generally, areas further away are being used in order to travel faster to certain areas of interest. The high-level policy also frequently uses frontier-based navigation when the exploration policy is stuck in some area. When a target object lies on the way to another relevant navigation point (a frontier, cabinet, or door), the high-level policy learned to navigate directly to the latter, instead of sequentially navigating to the target object and then proceeding. This saves time and improves efficiency.

E. Real World Experiments

We transfer the trained policy to a *Toyota HSR* robot. The robot has a height-adjustable torso and a 5-DoF arm for environment interactions. It is equipped with an RGB-D camera used for mapping and a 2D lidar employed for localization in the pre-built map and object avoidance by the ROS navigation stack. Both the local exploration policy and the high-level behavior policy are transferred to this real-world setting without any further retraining or fine-tuning. The agent requires only minor adjustments to account for the differences in robot geometry and subpolicies. See Section IV-B and the supplementary material for details on the real-world subpolicies.

The experiments are performed in an office building covering three rooms connected by a hallway with a total of three doors. The operation space covers roughly 180 square meters. We place three articulated objects, two cabinets with a revolute door, and one drawer in the environment. These objects are never seen during training. We evaluate runs with 1-6 target objects in five different scenarios, for a total of 30 episodes. Each scenario defines new positions for the three articulated objects. We randomly chose which doors start in an open or closed state, and start each episode from the room that the last episode terminated in. We cover glass doors to prevent

TABLE I
EVALUATION OF SEEN AND UNSEEN ENVIRONMENTS, REPORTING THE SUCCESS RATE (TOP) AND SPL (BOTTOM).

Model	Seen						Unseen								
	1-obj	2-obj	3-obj	4-obj	5-obj	6-obj	Avg 1-6	1-obj	2-obj	3-obj	4-obj	5-obj	6-obj	Avg 1-6	
Success	Greedy	96.8	94.0	91.0	91.3	89.7	88.0	91.8	96.4	94.9	92.6	93.7	91.4	87.2	92.7
	SGoLAM+	96.0	93.7	90.2	85.8	88.0	85.3	89.8	95.6	94.1	94.7	90.5	91.6	91.6	93.0
	w/o frontier	81.1	68.9	59.1	68.5	71.9	35.7	64.2	82.2	74.2	65.0	61.8	65.3	50.5	66.5
	w/o expl	93.2	88.4	87.2	91.6	87.6	58.4	84.4	93.4	90.2	85.8	81.8	70.8	79.6	83.6
	w/o IAM	95.1	95.1	90.3	89.5	84.0	87.2	90.2	96.5	94.3	92.1	82.7	85.4	78.2	88.2
	HIMOS	96.8	96.5	94.0	94.8	93.5	91.7	94.6	97.5	97.9	95.6	95.6	93.5	93.5	95.6
SPL	Greedy	43.1	45.3	48.6	50.8	54.1	56.9	49.8	45.0	49.5	51.8	57.9	60.6	62.0	54.5
	SGoLAM+	41.5	44.3	47.8	48.9	52.5	55.3	48.4	44.8	49.1	54.2	56.2	60.3	64.2	54.8
	w/o frontier	41.0	41.2	42.9	44.9	58.0	47.4	45.9	32.0	45.2	50.9	46.3	49.6	47.2	45.2
	w/o expl	43.3	45.7	45.4	41.6	48.8	44.9	44.9	40.9	43.1	56.1	54.7	45.7	52.3	48.8
	w/o IAM	46.9	48.5	51.9	38.7	61.0	51.2	49.7	45.9	46.3	45.2	55.0	54.1	52.3	49.8
	HIMOS	49.7	51.5	55.2	57.5	59.5	60.9	55.7	52.2	56.0	57.6	63.5	65.7	68.7	60.6

TABLE II
REAL WORLD EXPERIMENTS ON THE HSR ROBOT.

Model	1-obj	2-obj	3-obj	4-obj	5-obj	6-obj	Total
Success	5	4	4	3	4	3	23
Collision	0	0	1	0	0	1	2
Interaction failure	0	1	0	0	1	1	3
Navigation failure	0	0	0	2	0	0	2
Total Episodes	5	5	5	5	5	5	30




Fig. 4. Top: map of the real world environment. Initial door state and cabinet positions are randomized between episodes. Below: example trajectory in the real world. From top left to bottom right: the agent decides to look inside a target object, then navigates to the hallway, opens a different cabinet and finally opens and drives through a closed door.

the agent from directly looking through them. This evaluation tests (i) the generalization abilities of both low- and high-level behaviors to the real world, given accurate semantic perception, (ii) the generalization abilities of the local exploration policy

and high-level policy to a different robot, and (iii) the high-level policy’s generalization to unseen subpolicies, as we change both the navigation and manipulation modules. This is an important ability for transfers to different robot models and execution requirements. Lastly, (iv) the map representation enables easy transfer to different objects as it only requires a mapping to known semantic and instance colors. Our aim is to evaluate the transfer of decision making ability within our stated assumptions. We leave full system evaluation with integrated perception pipeline for future work. We provide details on how this could be implemented in the supplementary.

The results of the experiments are summarized in Table II and example episodes are shown in Figure 4 as well as in the supplementary video. The agent successfully completes 76.7% of the episodes, requiring long sequences of autonomous navigation and physical interactions. The high-level policy proves robust to failures in the subpolicies. These include navigation failures if the planner does not find a valid path to a frontier point and manipulation failures in which the mobile manipulation skill fails to grasp the handle of an articulated object. In this case, the agent is capable of re-triggering the interactions after detecting the failure. A few irrecoverable failures occurred: reaching a safety limit of the wrist joint during door opening, base collisions, and in two cases repeated failures of the navigation stack.

VI. CONCLUSION

We introduced the interactive multi-object search task in which the agent has to manipulate the environment in order to fully explore it, resembling common household settings. We proposed a novel hierarchical, reinforcement learning-based approach capable of solving this complex task in both simulation and the real world. By combining a high-level policy on abstract action spaces with low-level robot behaviors, we are able to perform long-term reasoning while acting in continuous action spaces. Our approach decouples the perception from decision making which allows a seamless transition to unknown and real-world environments on a differing embodiment. Tighter integration with perception and active perception in the lower level are promising areas

for future work. In extensive experiments, we demonstrated the capabilities of our approach and the importance of the individual components in ablation studies. In future work, we will investigate the benefits of jointly training the high- and low-behavior and integrate more sophisticated mapping modules directly based on the robot sensors. Further, additional low-level behaviors could extend environment interaction options or perform more goal-oriented active perception actions. Detection of more finegrained failure-feedback could further help the high-level decision making. Lastly, the current instance navigation assumes a pre-specified maximum number of possible instances. Regression-based instance selection is a promising avenue to further scale this approach.

REFERENCES

- [1] S. Wani, S. Patel, U. Jain, A. Chang, and M. Savva, "Multion: Benchmarking semantic map memory using multi-object navigation," *Proc. of the Conf. on Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 9700–9712, 2020.
- [2] K. Fang, A. Toshev, L. Fei-Fei, and S. Savarese, "Scene memory transformer for embodied agents in long-horizon tasks," in *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 2019, pp. 538–547.
- [3] F. Schmalstieg, D. Honerkamp, T. Welschehold, and A. Valada, "Learning long-horizon robot exploration strategies for multi-object search in continuous action spaces," *Proceedings of the International Symposium on Robotics Research (ISRR)*, 2022.
- [4] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *Proc. of the IEEE Int. Symp. on Comput. Intell. in Rob. and Aut. (CIRA)*, 1997, pp. 146–151.
- [5] D. S. Chaplot, D. Gandhi, S. Gupta, A. Gupta, and R. Salakhutdinov, "Learning to explore using active neural slam," in *International Conference on Learning Representations*, 2020.
- [6] C. Chen, U. Jain, C. Schissler, S. V. A. Gari, Z. Al-Halah, V. K. Ithapu, P. Robinson, and K. Grauman, "Soundspaces: Audio-visual navigation in 3d environments," in *Proc. of the Europ. Conf. on Computer Vision (ECCV)*, 2020, pp. 17–36.
- [7] D. Honerkamp, T. Welschehold, and A. Valada, " N^2m^2 : Learning navigation for arbitrary mobile manipulation motions in unseen and dynamic environments," *IEEE Transactions on Robotics*, 2023.
- [8] J. Gu, D. S. Chaplot, H. Su, and J. Malik, "Multi-skill mobile manipulation for object rearrangement," in *International Conference on Learning Representations*, 2023.
- [9] A. Szot, A. Clegg, E. Undersander, E. Wijmans, Y. Zhao, J. Turner, N. Maestre, M. Mukadam, D. S. Chaplot, O. Maksymets *et al.*, "Habitat 2.0: Training home assistants to rearrange their habitat," *Proc. of the Conf. on Neural Information Processing Systems (NeurIPS)*, vol. 34, pp. 251–266, 2021.
- [10] T. Chen, S. Gupta, and A. Gupta, "Learning exploration policies for navigation," in *International Conference on Learning Representations*, 2019.
- [11] A. Younes, D. Honerkamp, T. Welschehold, and A. Valada, "Catch me if you hear me: Audio-visual navigation in complex unmapped environments with moving sounds," *IEEE Rob. and Automation Letters*, vol. 8, no. 2, pp. 928–935, 2023.
- [12] Y. Qiu, A. Pal, and H. I. Christensen, "Learning hierarchical relationships for object-goal navigation," in *2020 Conference on Robot Learning (CoRL)*, 2020.
- [13] R. Druon, Y. Yoshidasu, A. Kanezaki, and A. Watt, "Visual object search by learning spatial context," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1279–1286, 2020.
- [14] E. Beeching, J. Debangoye, O. Simonin, and C. Wolf, "Deep reinforcement learning on a budget: 3d control and reasoning without a supercomputer," in *25th International Conference on Pattern Recognition (ICPR)*, 2021, pp. 158–165.
- [15] F. Xia, W. B. Shen, C. Li, P. Kasimbeg, M. E. Tchapmi, A. Toshev, R. Martín-Martín, and S. Savarese, "Interactive gibbon benchmark: A benchmark for interactive navigation in cluttered environments," *IEEE Rob. and Automation Letters*, vol. 5, no. 2, pp. 713–720, 2020.
- [16] S. K. Ramakrishnan, D. S. Chaplot, Z. Al-Halah, J. Malik, and K. Grauman, "Poni: Potential functions for objectgoal navigation with interaction-free learning," in *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 2022, pp. 18 890–18 900.
- [17] J. Kim, E. S. Lee, M. Lee, D. Zhang, and Y. M. Kim, "Sgolam: Simultaneous goal localization and mapping for multi-object goal navigation," *arXiv preprint arXiv:2110.07171*, 2021.
- [18] K. Zheng, R. Chitnis, Y. Sung, G. Konidaris, and S. Tellex, "Towards optimal correlational object search," in *Int. Conf. on Robotics & Automation. IEEE*, 2022, pp. 7313–7319.
- [19] K. Zheng, A. Paul, and S. Tellex, "A system for generalized 3d multi-object search," *Int. Conf. on Robotics & Automation*, 2023.
- [20] A. Röfer, G. Bartels, W. Burgard, A. Valada, and M. Beetz, "Kinverse: A symbolic articulation model framework for model-agnostic mobile manipulation," *IEEE Rob. and Automation Letters*, vol. 7, no. 2, pp. 3372–3379, 2022.
- [21] A. H. Qureshi and Y. Ayaz, "Intelligent bidirectional rapidly-exploring random trees for optimal motion planning in complex cluttered environments," *Robotics and Autonomous Systems*, vol. 68, pp. 1–11, 2015.
- [22] D. Honerkamp, T. Welschehold, and A. Valada, "Learning kinematic feasibility for mobile manipulation through deep reinforcement learning," *IEEE Rob. and Automation Letters*, vol. 6, no. 4, pp. 6289–6296, 2021.
- [23] O. Nachum, S. S. Gu, H. Lee, and S. Levine, "Data-efficient hierarchical reinforcement learning," *Advances in neural information processing systems*, vol. 31, 2018.
- [24] M. Hutsebaut-Buysse, K. Mets, and S. Latré, "Hierarchical reinforcement learning: A survey and open research challenges," *Machine Learning and Knowledge Extraction*, vol. 4, no. 1, pp. 172–221, 2022.
- [25] J. Krantz, A. Gokaslan, D. Batra, S. Lee, and O. Maksymets, "Waypoint models for instruction-guided navigation in continuous environments," in *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 2021, pp. 15 162–15 171.
- [26] C. Chen, S. Majumder, Z. Al-Halah, R. Gao, S. K. Ramakrishnan, and K. Grauman, "Learning to set waypoints for audio-visual navigation," in *International Conference on Learning Representations*, 2020.
- [27] N. Yokoyama, A. W. Clegg, E. Undersander, S. Ha, D. Batra, and A. Rai, "Adaptive skill coordination for robotic mobile manipulation," *arXiv preprint arXiv:2304.00410*, 2023.
- [28] M. Iovino, E. Scukins, J. Styrod, P. Ögren, and C. Smith, "A survey of behavior trees in robotics and ai," *Robotics and Autonomous Systems*, vol. 154, p. 104096, 2022.
- [29] G. Golluccio, D. Di Vito, A. Marino, and G. Antonelli, "Robotic weight-based object relocation in clutter via tree-based q-learning approach using breadth and depth search techniques," in *Int. Conf. on Advanced Robotics. IEEE*, 2021, pp. 676–681.
- [30] R. Mohan and A. Valada, "Amodal panoptic segmentation," in *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 2022, pp. 21 023–21 032.
- [31] S. Kohlbrecher, J. Meyer, O. von Stryk, and U. Klingauf, "A flexible and scalable slam system with full 3d motion estimation," in *Proc. IEEE International Symposium on Safety, Security and Rescue Robotics*, 2011.
- [32] S. Huang and S. Ontañón, "A closer look at invalid action masking in policy gradient algorithms," in *FLAIRS Conference*, 2022.
- [33] R. S. Sutton, D. Precup, and S. Singh, "Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning," *Artificial Intelligence*, vol. 112, no. 1, pp. 181–211, 1999.
- [34] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [35] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [36] C. Li, F. Xia, R. Martín-Martín, M. Lingelbach *et al.*, "igibson 2.0: Object-centric simulation for robot learning of everyday household tasks," in *Conference on Robot Learning*, vol. 164, 2021, pp. 455–465.
- [37] P. Anderson, A. Chang, D. S. Chaplot, A. Dosovitskiy, S. Gupta, V. Koltun, J. Kosecka, J. Malik, R. Mottaghi *et al.*, "On evaluation of embodied navigation agents," *arXiv preprint arXiv:1807.06757*, 2018.