# Reasoning with Scene Graphs for Robot Planning under Partial Observability

Saeid Amiri, Kishan Chandan, and Shiqi Zhang

Abstract—Robot planning in partially observable domains is difficult, because a robot needs to estimate the current state and plan actions at the same time. When the domain includes many objects, reasoning about the objects and their relationships makes robot planning even more difficult. In this paper, we develop an algorithm called scene analysis for robot planning (SARP) that enables robots to reason with visual contextual information toward achieving long-term goals under uncertainty. SARP constructs scene graphs, a factored representation of objects and their relations, using images captured from different positions, and reasons with them to enable context-aware robot planning under partial observability. Experiments have been conducted using multiple 3D environments in simulation, and a dataset collected by a real robot. In comparison to standard robot planning and scene analysis methods, in a target search domain, SARP improves both efficiency and accuracy in task completion. Supplementary material can be found at https://tinyurl.com/sarp22

Index Terms—Planning under Uncertainty, Probabilistic Inference, Semantic Scene Understanding.

### I. INTRODUCTION

THERE has been great progress in development of service robots in the recent years, e.g., [1], [2]. Those robots are able to conduct everyday tasks in human-inhabited environments over extended periods of time. Robot perception in such domains is partial and unreliable, which brings a major challenge to robot decision making.

Partially Observable Markov Decision Process (POMDP) is a framework that models the uncertainty in both observations and action outcomes [3], and has been used for policy generation in partially observable domains. However, the challenges are two-fold. First, constructing POMDPs requires that the robot has a complete world model, which tends to be infeasible in practice. In particular, real-world environments (say a kitchen) frequently include many objects, making it troublesome to use POMDPs to have a universal representation of all objects. Second, the complexity of reasoning about these objects and their relationships grows exponentially as more objects are considered. In this paper, we aim to develop an approach that reasons with contextual information for scene analysis to enable POMDP-based robot planning.

One of the recent advancements in computer vision has been scene graph generation networks [4]–[10]. Given an

Manuscript received: September, 09, 2021; Revised January, 01, 2022; Accepted January, 30, 2022.

This paper was recommended for publication by Editor Hanna Kurniawati upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by NSF (NRI-1925044), Ford, OPPO, and SUNY RF)

All authors are with the Department of Computer Science, the State University of New York (SUNY) at Binghamton, Binghamton, NY, 13902, USA. Email: {samiri1; kchanda2; zhangs}@binghamton.edu Digital Object Identifier (DOI): see top of this page.

image, scene graph systems generate a graph consisting of detected objects (e.g., a book and a table), their corresponding bounding boxes, and the relationships among the objects (e.g., book on a table). Scene graphs provide a robot with a structured understanding of the world in terms of objects, and their relations. From the robotics perspective, however, current scene graph research has the limitation that the context analysis does not go beyond individual images, even though a robot can easily capture images from different angles and locations for analysis purposes. With the active perception capabilities of robots, we have the objective of developing an approach for domain-wide active scene analysis for mobile robots.

In this work, we develop an algorithm called *scene analysis* for robot planning (SARP) for planning robot actions for context-aware, object-centric scene analysis. SARP uses local scene graphs of single images to build and augment global scene graphs toward context-aware robot planning under partial observability. An overview of SARP is shown in Figure 1. More specifically, a global scene graph is incrementally constructed "on the fly" using local scene graphs generated at different locations when new objects are perceived. Reasoning with this global scene graph produces useful information to help the robot estimate the current world state. This enhanced state estimation enables the robot to improve its performance in goal achievement.

We have evaluated SARP using target search tasks where a robot needs to locate an object in an indoor environment. We use POMDPs to model the robot's perception and actuation skills [3], use Neural Motifs [4] to compute local scene graphs, and use approximate inference methods to build Markov networks computed from large datasets. We have extensively evaluated SARP through comparisons with competitive baselines in simulation. Results show that SARP reduced the overall action costs by 16% compared with a predefined action policy. Also, SARP helps the robot maintain its policy quality in the presence of an increased number of objects, and enables the robot to focus on the areas that are most relevant to the current task.

# II. RELATED WORK

This work aims to enable a robot to represent and reason with contextual information to guide robot planning under uncertainty. Researchers have developed algorithms that reason with contextual knowledge to guide sequential decision making [11]. The contextual knowledge can be in a variety of forms, such as commonsense knowledge [12], action

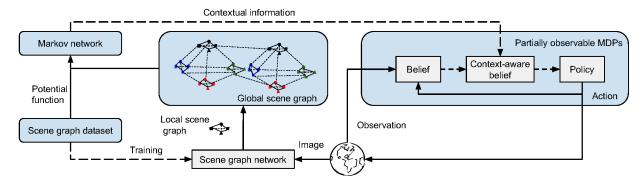


Fig. 1. An overview of SARP, where the robot takes an action using the policy, and receives an observation from the world, updating the belief using the action and observation. In each timestep, a graph is accumulated using the perception sensor to bias the belief. This biased belief provides the contextual information to the policy. Top dashed lines show that the belief biasing not occurring in every iteration. The bottom dashed line indicates that the scene graph network is trained offline.

knowledge [13], [14], and graph-based knowledge [7], [15]. Such contextual knowledge can also be leveraged to guide RL agents. In this section, we examine each of these categories.

a) Rule-based human commonsense knowledge: Researchers have used rule-based commonsense knowledge to guide the robot planning under uncertainty [16]-[18]. In their methods, a robot reasons about human knowledge to compute an informative prior to help a probabilistic planner estimate the current world state, enabling the robot to achieve complex goals with less information-gathering behaviors. Our previous work added a learning component to further improve a robot's reasoning and planning capabilities from its task completion experience [19]. Another example is an algorithm that uses first-order logic to construct decision tree policies for goaloriented factored POMDPs [20]. Others have used action knowledge to build a hierarchical robot planner where the higher level computes a sequence of abstract actions and the lower level implements the higher-level actions using primitive behaviors [21]. Algorithm iCORPP enables a robot to reason with contextual knowledge to compute parameters of a planning agent's reward and transition functions [22]. Others have used commonsense knowledge to guide a classical planner to reason and plan in open worlds [23], [24]. Compared to those methods, SARP uses contextual, object centric information, in the form of a graph, to estimate the current world state and guide robot planning under uncertainty.

b) Hierarchical frameworks: Researchers have been using hierarchies to construct their framework's knowledge-base or planner [21], [25]. In a recent work, researchers used a visual hierarchical planning algorithm for long-horizon manipulation tasks. Their framework integrates neuro-symbolic task planning and graph-based motion generation on graph-based scene representations [25]. Their method has two-level abstractions of a manipulation scene with geometric scene graphs and symbolic scene graphs. To enable a robot to operate in open-world domains, researchers have developed a three-layer hierarchy for reasoning about action knowledge and default knowledge [26]. In particular, the knowledge at higher level was used for correcting lower-level knowledge to guide probabilistic planning. In comparison, SARP (ours)

incrementally builds and reasons about scene graphs to guide robot planning under uncertainty.

c) Graph-based human knowledge: Graph-based representations have been used for reasoning with contextual information to guide robot planning [25], [27], [28]. For instance, researchers have used Conditional Random Fields [29] to construct contextual knowledge bases for robot target search by maintaining a belief over the locations of target objects and landmark objects while exploiting the knowledge of their co-appearances [27].

Other methods were developed to extract prior knowledge from data, e.g., using Long Short-Term Memory (LSTM) networks [30], to guide a planning agent in navigating new environments [28]. In comparison to those methods, SARP leverages contextual information in the form of automatically constructed scene graphs, avoiding domain experts manually developing knowledge bases.

A very recent work leverages scene graph reasoning to disambiguate human utterances when referring to various instances of the same object [31]. Their framework reasons about static domains, whereas our robot leverages scene analysis to plan towards achieving long-term goals.

d) Knowledge-based RL: RL agents are able to leverage contextual information as well. For instance, existing research has shown that an RL agent is able to reason with action knowledge to decompose complex tasks into smaller, tractable subtasks [32]. The concept of reward machines has been introduced for using temporal knowledge toward reusing interaction experience and creating extra feedback for learning purposes [33]. Recent research has shown that action knowledge can be used to guide a model-based RL agent by providing optimistic, artificial interaction experience to speed up the learning process [34], [35]. We assume the availability of the world dynamics, and use planning under uncertainty methods (instead of RL) for robot decision making. To the best of our knowledge, SARP is the first that enables a planning agent to automatically construct, and use scene graphs for context analysis under partial observability.

e) POMDP-based target search: Several reasearchers have used POMDPs for the target search task [36]–[39]. In one work, a robot arm is tasked to search for an object in the

clutter where it should learn the synergies of acting and seeing objects. In this work, the policy learned by a POMDP model could determine when is a good time to look at and detect objects and when it should move objects around in order to reduce the clutter [38]. In another work, a robot searched for multiple objects by specifiying their locations from the user query (e.g., *Find the mugs in the kitchen and books in the library*) [37]. In their framework, a robot can associate the locations to each object class so as to improve its search. To the best of our knowledge, none of these works leverage graph reasoning for target search tasks.

### III. PRELIMINARIES

In this section, we describe the three building blocks of this research, namely scene graphs, and partially observable Markov decision processes.

## A. Scene Graphs

Scene graph G is a representation of the semantic content of an image, consisting of a set of bounding boxes  $B=\{b_1,...,b_n\}$ , a set of objects  $V=\{o_0,o_1,\cdots\}$ , and a set of binary predicates  $E=\{r_0,r_1,\cdots\}$  [40], A triplet of object-predicate-object is called a relationship. Given an image I, the probability distribution of the scene graph pr(G|I) is decomposed into three components:

$$pr(G|I) = pr(B|I)pr(V|B, I)pr(E|V, B, I)$$
(1)

where the bounding box component pr(B|I) generates a set of candidate regions from the input image, pr(V|B,I) predicts the class label for each predicted region, and pr(E|V,B,I) predicts the predicate among objects, which is conditioned on the predicted labels.

# B. Partially Observable MDPs

Markov decision processes (MDPs) can be used for sequential decision making under full observability [41]. Partially observable MDPs (POMDPs) generalize MDPs by assuming current state being partially observable [3]. A POMDP is represented as a tuple  $(S,A,T,R,Z,O,\gamma)$  where S is the state-space, A is the action set, T is the state-transition function, R is the reward function, R is the observation set, R0 is the observation function, and R1 is a discount factor that determines the planning horizon.

A robot maintains a belief state distribution b based on observations  $(z \in Z)$  using the Bayes update rule:

$$b'(s') = \frac{O(s', a, z) \sum_{s \in S} T(s, a, s') b(s)}{pr(z|a, b)}$$
(2)

where s is a state, a is an action, pr(z|a,b) is a normalizer, and z is an observation. Solving a POMDP produces a policy that maps the current belief state distribution to an action toward maximizing long-term utilities.

In this research, scene graphs and POMDPs are used for representing and reasoning about objects (and their relationships), and planning actions toward achieving long-term goals. The main contribution of this paper is the novel interplay between scene graphs and POMDPs for context-aware robot planning under partial observability.

# C. Markov Networks

A Markov network is a graph that consists of variables (nodes)  $\mathcal{X} = \{X_1, \cdots, X_n\}$  and undirected edges connecting pairs of nodes  $(X_i, X_j)$  [15]. Each edge is parameterized using a potential function (factor)  $\phi$  that captures the affinities between the variables. The joint distribution of all variables are:

$$P(X_1, \dots, X_n) = \frac{1}{Z} \prod_{i=1}^m \phi(D_i)$$
 (3)

3

where  $D_i$  is the *i*th edge of the network. A potential function of the edge connecting nodes  $X_i$  and  $X_j$ , consists of  $|X_i||X_j|$  values

In the next section, we describe the algorithmic contribution of this work.

### IV. ALGORITHM

a) Problem Formulation: In this work, we are interested in the problem of target search. A mobile robot receives the task of searching for object Q and can navigate in the environment E in order to find the object. The robot is provided with the environment map and is localized initially. Once the robot navigates in the environment sufficiently, it reports the location of the target object. In order for the robot to better reason about Q, robot requires the scene graph network that is pretrained on a scene graph dataset  $\mathcal{D}$ .

b) POMDP model: In order to solve this target search task, we first define the POMDP model  ${\mathcal M}$  as the tuple  $< S, A, T, R, O, Z, \gamma >$ . Its factored state space set S, is a Cartesian product of two dimensions, and the terminal state.  $S^E$  includes a set of discrete partially observable locations of Q (the target object) and  $S^R$  includes the set of robot's fully observable locations. We define these equidistant, discrete locations manually in the robot's motion planning workspace. SARP's action set A consists of navigation and termination actions. The robot can take navigation action  $qo_i \in A$  to go to the location i. We model the action transition function T(s'|s,a) so that the robot can only go to its closest neighboring locations in the absence of obstacles. Each navigation action has a cost (negative value, R(s,a) < 0), which is proportional to the distance the robot needs to travel to reach that location. The robot can take the termination action, and receive bonus (penalty) if it correctly locates (not find) the target object. Observation set Z is {Detected, NotDetected, *NotApplicable*} where *Detected* happens when the perception detects the query object, NotDetected happens when the query object is not detected, and NotApplicable when the agent takes the termination action. To maximize its planning horizon, we set  $\gamma$  to 0.99. SARP assigns the observation function O = pr(z|s,a) based on the target's object detection accuracy on a test set of  $\mathcal{D}$ .

c) Algorithm Description: After defining the problem and the POMDP model, we present our novel algorithm, called scene analysis for robot planning (SARP), for context-aware robot planning under partial observability. SARP computes a policy that enables a robot to accomplish its task using less action cost using the contextual knowledge. SARP is an object-centric algorithm that bridges the representation gap

between visual scene analysis and robot planning under partial observability.

Algorithm 1 presents SARP, whose input is an object of interest that a mobile robot needs to locate (Q). SARP requires a POMDP solver for policy generation, a pre-trained scene graph generation network, a domain map for navigation, and a dataset  $\mathcal{D}$  that consists of scene graphs. In Line 1, SARP constructs a POMDP according to the object of interest (Q), and computes policy  $\pi$  using the provided POMDP solver. After that, there are a few steps for initializing beliefs (over object locations), and a scene graph (Lines 2-4). It should be noted that we maintain two beliefs b and b' over the location of the target object, where b is updated using POMDP observations, and b' is updated using the current b and available contextual information. This design allows the robot to use different beliefs for decision making and action selection, and avoids possible issues caused by error propagation. Also, this mechanism enables a SARP agent to avoid reusing contextual information in belief updates. This mechanism enables a SARP agent to avoid reusing contextual information in belief updates.

Lines 5-20 form the main control loop of SARP, and it terminates when the current state is a terminal state. In each iteration, the robot uses a captured image to generate  $G^L$ , a local scene graph (Line 6), and uses this local scene graph to update  $G^G$ , the global scene graph (Line 8).

After that, SARP computes a potential function  $\phi$  (Line 9) values for each edge, using the function  $CALC\phi$  (Algorithm 2). This function calculates four values for each relation  $(v,\epsilon,v')$  in  $G^L$  that serve as the potential function.  $CALC\phi$  queries the dataset  $\mathcal{D}$ , to find out the ratio of the times that each object (v,v') or both) in a relation has appeared (not appeared) in  $\mathcal{D}$ , resulting in four values to form the potential function.

We use  $\phi$  to form a Markov network together with global scene graph G in Line 10 where SARP queries the number of times that each pair of nodes of an edge in  $G^G$  has appeared (not appeared) in  $\mathcal{D}$  scene graphs.

In our implementation,  $\mathcal{N}$  is incrementally updated in each iteration, if the robot detects new scene graphs (that were not previously detected), it will add the new nodes and relationships to the existing global scene graph  $G^G$ . We use the robot's localization and the camera depth sensor to approximately localize the detected objects on the map, in order to distinguish different instances of the same objects.

Lines 11-19 correspond to the belief update, and action selection processes of POMDPs. Contextual knowledge in the form of a Markov network ( $\mathcal{N}$ ) is used for biasing belief b only if the object of interest is visually detected in the current image. This is because of avoiding the bias drift that could result from too many biasing at every timestep. We use a belief propagation<sup>1</sup> method [42] to compute the probability of Q being collocated with objects V (Line 15).

$$pr(Q|V) = BeliefPropagation(\mathcal{N}, Q, V)$$
 (4)

where the computations of pr(Q|V) and b are independent. In Line 16, SARP uses pr(Q|G) to compute b', the posterior belief distribution.

### **Algorithm 1 SARP**

Input: Query object Q

**Require:** a POMDP solver, a scene graph generation network, a domain occupancy grid map, robot localization software, scene graph dataset  $\mathcal{D}$ .

- 1: Construct a POMDP based on Q, and compute policy  $\pi$
- 2: Uniformly initialize beliefs b and b' over  $S^{E}$
- 3: Initialize vertices  $V=\{Q,Q'\}$  and edges  $\mathcal{E}=\{\epsilon_{Q-Q'}\}$ , where Q' is a duplicate of Q
- 4: Initialize a scene graph:  $G \leftarrow (V, \mathcal{E})$
- 5: while current state is not terminal do
- 6: Take an image, and generate local scene graph  $G^L = (V', \mathcal{E}')$ , where V' are detected evidence objects
- 7:  $V \leftarrow V \cup V'; \ \mathcal{E} \leftarrow \mathcal{E} \cup \mathcal{E}'$
- 8:  $G^G \leftarrow (V, \mathcal{E})$
- 9: Compute potential function  $\phi = CALC \phi(G^L, \mathcal{D})$
- 10: Form a Markov network:  $\mathcal{N} \leftarrow (G^G, \phi)$
- 11: Select action  $a \leftarrow \pi(b')$ , and execute a
- 12: Make observation z on Q
- 13: Update b based on a and z
- 14: **if** z is Detected **then**
- 14. If z is Detected then
- 15:  $pr(Q|V) \leftarrow BeliefPropagation(\mathcal{N}, Q, V)$
- 16:  $b' \leftarrow \eta \cdot pr(Q|V) \cdot b$
- 17: **else if** z is NotDetected **then**
- 18:  $b' \leftarrow b$
- 19: **end if**
- 20: end while

# Algorithm 2 $CALC_{\phi}$

```
Input:
                  Scene
                                                   dataset
                                                                      \mathcal{D}.
                                                                                 local
                                                                                                                 graph
 1: Initialize the potential function set \phi as empty
 2: for (v, \epsilon, v')^{\hat{i}} in G^L do
           m \leftarrow number of images in \mathcal{D} containing v or v'.
           \phi^{\imath}_{v^1,v'^1} \leftarrow (\text{number of images in } \mathcal{D} \text{ containing } v,v', \text{and } \epsilon)/m
 4:
           \phi^i_{v^0,v'^0} \leftarrow (\text{number of images in } \mathcal{D} \text{ containing } v,v', \text{but not } \epsilon)/m
           \phi^i_{v^0,v'^1} \leftarrow \text{(number of images in } \mathcal{D} \text{ containing } v' \text{ without } v)/m
           \phi^i_{v^1\ v'^0} \leftarrow \text{(number of images in } \mathcal{D} \text{ containing } v \text{ without } v')/m
 7:
           \phi \leftarrow \phi \cup \phi^i
 9: end for
10: return \phi
```

d) CALC function: CALC function that takes as input the scene graph dataset  $\mathcal D$  and the local scene graph  $G^L$ . For each pair of nodes v and v' that are connected in the ith edge  $\epsilon$ , the algorithm calculates four values by querying the number of times each node and their corresponding edge have (not) appeared in the dataset  $\mathcal D$ . We denote  $v^1$  as the node exists in  $\mathcal D$  and  $v^0$  when it does not exist. It returns the computed potential function  $\phi$ .

SARP enables the agent to leverage contextual information towards task completion through building and reasoning with global scene graphs to guide a probabilistic planner. Next, we discuss our experiments for evaluating SARP.

### V. EXPERIMENTS

We conducted two sets of experiments in simulation where the robot is tasked with finding a target object accurately and as quickly as possible. In all the trials, the robot is provided with a domain map, a dataset of scene graphs, and a pretrained network for generating scene graphs. It receives 360-degree images of the environment as the input, where we leveraged

<sup>&</sup>lt;sup>1</sup>Any approximate inference method can be used.

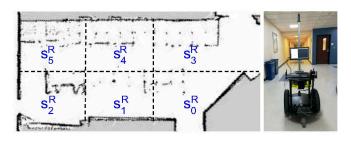


Fig. 2. (**Left**): Map of the environment where the robot can navigate. The discretized locations are shown in blue color. (**Right**): The Segway RMP110 mobile platform, equipped with 360-degree vision, used in this research.

our previous research on 360-degree robot vision [43], and the output is the location of the target object. Our first baseline method is a naive POMDP planner [38] (with uniform prior belief) where the robot action policy solely depends on the model of the world. The second baseline uses a predefined policy where the robot exhaustively visits all discrete positions, updates the belief at each timestep, and reports the object's location based on the argmax of the belief. Our third baseline method is CORPP [16] as another competitive baseline where only the initial belief is biased based on the commonsense rules defined by the human developer in the form of logical probabilistic rules (e.g., a book is likely to be on a desk with 0.8 probability). Similar to SARP, all the baselines maintain a belief of the object's location and update it at each timestep using Bayes update rule. However, none of them use graph reasoning like SARP. We have two evaluation metrics. First, the average action cost that represent the average execution time of all actions taken until the terminal state. Second, is the average success rate in finding the target object's location correctly. By using SARP, we hypothesize that:

- 1) The robot's overall action cost would be less compared to baselines (H-1).
- 2) SARP performs better than the baselines in action cost and success rate in domains with a large number of objects (H-2).

The reward of successfully finding the target object is 100, and the penalty of failure in finding the target object is -100. The reward for all actions  $go_i$  is -10 which is proportional to the time it takes for the robot to execute the action. We solve the POMDP model using an off-the-shelf point-based system [44].

We use Neural Motif [4] for generating local scene graphs where there are a total of 50 predicate classes. The most prevalent objects in this dataset are humans and the most appeared predicates are *in*, *on*, and *belongs*. Given an input image, Neural Motifs produces a scene graph which is a list of objects, their probabilities, relationships and bounding boxes. SARP requires a dataset to assign the Markov network potential function (Line 9 in Algorithm 1). We are using Visual Genome [45] that contains round 108K images, 3.8M objects, 2.8M relationships.

# A. Setup

We used the robot to navigate through the environment and collect images in a hallway in an indoor educational environment (Figure 2) in order to build a dataset. We call it the *hybrid* dataset. We used this dataset to simulate robot's behavior in the experiments. We manually placed multiple objects including *banana*, *laptop*, *human*, *books*, *mug*, *etc*. at different locations.

In addition to the dataset collected by the robot, we collected rendered images using an embedded agent in AI2THOR [46], an open-source interactive environment for embodied AI. AI2THOR provides 30 instances of four types of environments (shown in Figure 4): Kitchens, Living Rooms, Bedrooms, and Bathrooms, totaling 120 environments where an embedded robot is able to take navigation actions. Navigation actions include moving and rotating in orthogonal directions. To make both datasets consistent, we manually create 360 images using the AI2THOR monocular camera. We use the default value of 1.5m for the camera visibility and 90 degrees for the field of view. The actions executions are stochastic with the default Gaussian noise of average 0.001 and standard deviation of 0.005. There are a total of 125 objects in AI2THOR platform<sup>2</sup>, while for each scene in the platform, there are 61 number of objects on average. We randomly select one of those objects for each individual trial in the experiments. For inference on the scene graph, we use pgmpy library [47].

## B. Results

Table I shows the result of the first set of experiments that evaluates the first hypothesis (H-1) where we ran this experiment 1500 times over three batched of 500 experiments. We ran the experiment both using the dataset collected by the real robot and by AI2THOR. We call the first one, the hybrid dataset and the latter one, the rendered dataset. In the hybrid dataset, the average cost of target search for our robot is consistently less than all the baselines while all the methods maintain a high success rate. In the rendered dataset, we categorize the results based on the four different types of indoor environments: kitchen, living room, bedroom and bathroom. We randomly selected five environment from each of the four types, totaling 20 different environments. In each environment, we conducted 100 target searches. Except for the bedroom and bathroom environments where the predefined policy has lower cost, SARP produced the lowest average action cost while maintaining the highest accuracy consistently. The reason is that, in bedroom and bathroom environments, robot has a smaller navigation area. As a result, it takes less action costs using the predefined policy. To evaluate the second hypothesis (H-2), we incrementally added more objects to the POMDP baseline, and to the scene graph of SARP. Figure 5, shows that SARP's average cost remains almost the same with an average of 50.6 s while the POMDP baseline's cost increases as the number of additional objects is increased. With an increase in the number of non-target objects, the uniform POMDP baseline tries to take more actions to find the target

<sup>&</sup>lt;sup>2</sup>https://ai2thor.allenai.org/ithor/documentation/objects/object-types/

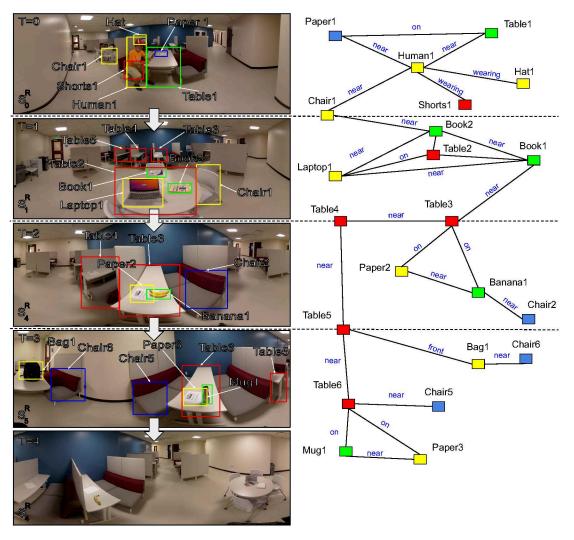


Fig. 3. The global scene graph being constructed incrementally as the robot navigates. Solid lines indicate the relationships between the detected objects, and dashed lines show how the global scene graph is completed from the local scene graphs. At timestep 4, the global scene graph is not augmented, since the robot has not detected new instances of objects.

object with more confidence, however with more than two nontarget objects, it finds that taking more actions is not helpful any more, and therefore is not successful. Also, the POMDP baseline failed to maintain the success rate at a high value due to its poor-quality policy.

### C. Enumerating objects of the same instance

As the robot is navigating the environment, it may perceive different instances of the same objects. Our scene graph network is not able to enumerate these objects. To better differentiate these instances, we use a hashmap to store the labels and location(s) of the detected objects. This facilitates the situations where multiple instances of the same objects need to be distinguished. We approximate the detected objects' locations using the centroid of the bounding boxes outputted by the scene graph. This provides the relative location of the object with respect to the robot. Then, we use coordinates transformation to get the object global scene graph. For the real robot experiments, we use the "Robot Operating System (ROS) [48] transforms" package to transform the locations to

global coordinates. It should be noted that vision-based object association is generally difficult, and introduces errors into our system, which is beyond the scope of this paper.

# VI. DEMONSTRATION

In this demonstration trial, the robot was assigned the task of searching for a banana as shown in Figure 3 while its initial location is in location 0. In Table II, we show an example trial where the robot is tasked with the search of a banana located at  $s_4^E$ . As the robot takes action suggested by the policy, it updates its belief over possible locations of the target object. In this example, the robot follows the trajectory of  $0 \to 1 \to 4 \to 5 \to 4$  until it terminates the trial and reports successfully that banana is in location 4. The robot visits locations 0,1,4, and 5 to build and augment its scene graph. Figure 3 shows the graph generated at each location. At timesteps 1 and 3, robot visits the location that it can easily detect the banana,

### TABLE I

THE EXPERIMENTAL RESULTS FROM THE HYBRID AND RENDERED DATASETS SHOW THAT ON AVERAGE OF 1500 TRIALS, SARP PERFORMS WITH LESS OVERALL ACTION COSTS. BOLD FONT SHOWS THE BEST RESULT. THE AVERAGE NUMBER INDICATES THE AVERAGE NUMBER OF OBJECTS INVOLVED IN A SINGLE EPISODE CONDUCTED IN THAT ENVIRONMENT.

Dataset Type	Average # of objects	SARP (Ours)		CORPP		Uniform POMDP (Li et al.,2016)		Predefined	
		Cost (std.)	Success	Cost (std.)	Success	Cost (std.)	Success	Cost (std.)	Success
Hybrid	15	52.1(13.1)	0.89	82.5 (8.9)	0.83	102.3 (14.6)	0.84	61.1 (0.0)	0.7
Rendered									
Kitchen	70	41.7 (21.6)	0.87	55.5 (18.4)	0.84	59.4 (24.1)	0.82	51.3 (0)	0.79
Living room	37	43.5 (19.1)	0.91	60.1 (33.8)	0.86	69.7 (23.9)	0.73	57.8 (0)	0.79
Bathroom	40	31.2 (16.4)	0.75	45.1 (17.6)	0.73	49.8 (21.7)	0.71	21.2 (0)	0.69
Bedroom	35	23.9 (12.9)	0.74	24.4 (19.1)	0.75	43.7 (22.5)	0.68	<b>19.4</b> (0)	0.64
Overall	45	35.7 (18.1)	0.81	46.3 (21.5)	0.79	55.7 (23.0)	0.73	37.3(0)	0.72



Fig. 4. Different types of environment provided by AI2THOR including kitchen, bathroom, bedroom, and living room.

therefore it biases its belief by inferring the whole scene graph. The overall cost for this trial is 60.

# VII. CONCLUSION & FUTURE WORK

Probabilistic planning methods under partial observability allow the robot to accomplish complex tasks toward maximizing long-term goal. Scene graphs allow the detection of objects and their relationships in images. Aiming at robots capable of accomplishing sophisticated tasks, we design a framework where a robot can use scene understanding information to obtain domain knowledge to provide contextual information to the robot planner where limited perception is a bottleneck. Results show that, by using our approach, the robot can benefit from the contextual information in the form of graph network by reducing action collection cost and avoiding scalability issues.

This work was based on a few assumptions that are sometimes unrealistic. The object detection accuracy and the set of

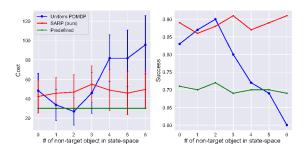


Fig. 5. Results of the experiment conducted on the renderred dataset. The state space consists of three locations, blue curve is the baseline that does not use contextual information but includes non-target objects in its state space incrementally, and the orange one is SARP that adds objects to the graph incrementally.

TABLE II An example trial where the robot is tasked with the search of a  $\it banana$  located in  $s_4^E$  while robot is initially at  $s_3^R$  .Here, we show how the belief is updated at each timestep.

	A: Action	Robot's belief of the target
Step	O: Observe	$[s_0^E, s_1^E, s_2^E, s_3^E, s_4^E, s_5^E]$
	A: go <i>l</i> 1	
	O: No	
0	Update:	[0.18, 0.09, 0.18, 0.18, 0.18, 0.18]
	Bias:No	
	A: go 14	
	O: Yes	
1	Update:	[0.15, 0.07, 0.15, 0.15, 0.33, 0.15]
	Bias:Yes	[0.12, 0.06, 0.11, 0.12, 0.47, 0.12]
	A: go <i>l</i> 5	
	O: No	
2	Update:	[0.13, 0.05, 0.13, 0.14, 0.49, 0.07]
	Bias:No	
	A: go <i>l</i> 4	
	O: Yes	
3	Update:	[0.09, 0.03, 0.07, 0.10, 0.68, 0.05]
	Bias: Yes	[0.07, 0.01, 0.04, 0.06, 0.81, 0.04]
4	A: terminate	

are all limited to the scene graph generation model quality. The scene graph we use, is incapable of face recognition, therefore we assume that humans are stationary (e.g, they are sitting at the time of the robot's task execution). In the future, we intend to further improve this work both in perception capabilities (e.g., enabling face detection) and reasoning capabilities(e.g., considering cases where objects may be replaced during the

task execution).

### ACKNOWLEDGEMENT

This work has taken place at the Autonomous Intelligent Robotics (AIR) Group, SUNY Binghamton. AIR research is supported in part by grants from the National Science Foundation (NRI-1925044), Ford Motor Company (URP Award 2019-2022), OPPO (Faculty Research Award 2020), and SUNY Research Foundation.

### REFERENCES

- [1] P. Khandelwal, S. Zhang, J. Sinapov, M. Leonetti, *et al.*, "Bwibots: A platform for bridging the gap between ai and human–robot interaction research," *IJRR*, vol. 36, no. 5-7, pp. 635–659, 2017.
- [2] N. Hawes, C. Burbridge, F. Jovan, L. Kunze, B. Lacerda, L. Mudrova, J. Young, J. Wyatt, D. Hebesberger, T. Kortner, et al., "The strands project: Long-term autonomy in everyday environments," *IEEE Robotics & Automation Magazine*, vol. 24, no. 3, pp. 146–156, 2017.
- [3] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and acting in partially observable stochastic domains," *Artificial intelligence*, vol. 101, no. 1-2, pp. 99–134, 1998.
- [4] R. Zellers, M. Yatskar, S. Thomson, and Y. Choi, "Neural motifs: Scene graph parsing with global context," in CVPR, 2018.
- [5] Y. Li, W. Ouyang, B. Zhou, K. Wang, and X. Wang, "Scene graph generation from objects, phrases and region captions," in *Proceedings* of the IEEE ICCV, 2017, pp. 1261–1270.
- [6] T. Chen, W. Yu, R. Chen, and L. Lin, "Knowledge-embedded routing network for scene graph generation," in CVPR, 2019, pp. 6163–6171.
- [7] D. Xu, Y. Zhu, C. B. Choy, and L. Fei-Fei, "Scene graph generation by iterative message passing," in *Proceedings of the IEEE Conference on CVPR*, 2017.
- [8] A. Rosinol, A. Gupta, M. Abate, J. Shi, and L. Carlone, "3d dynamic scene graphs: Actionable spatial perception with places, objects, and humans," arXiv preprint arXiv:2002.06289, 2020.
- [9] M. Engelcke, A. R. Kosiorek, O. P. Jones, and I. Posner, "Genesis: Generative scene inference and sampling with object-centric latent representations," arXiv preprint arXiv:1907.13052, 2019.
- [10] Z. Lin, Y.-F. Wu, S. V. Peri, W. Sun, G. Singh, F. Deng, J. Jiang, and S. Ahn, "Space: Unsupervised object-oriented scene representation via spatial attention and decomposition," arXiv preprint arXiv:2001.02407, 2020.
- [11] S. Zhang and M. Sridharan, "A survey of knowledge-based sequential decision making under uncertainty," arXiv preprint arXiv:2008.08548, 2020
- [12] E. Davis and G. Marcus, "Commonsense reasoning and commonsense knowledge in artificial intelligence," *Communications of the ACM*, vol. 58, no. 9, pp. 92–103, 2015.
- [13] D. Nau, M. Ghallab, and P. Traverso, "Automated planning: Theory & practice," 2004.
- [14] P. Haslum, N. Lipovetzky, D. Magazzeni, and C. Muise, "An introduction to the planning domain definition language," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 13, no. 2, pp. 1–187, 2019.
- [15] D. Koller and N. Friedman, Probabilistic graphical models: principles and techniques. MIT press, 2009.
- [16] S. Zhang and P. Stone, "Corpp: Commonsense reasoning and probabilistic planning, as applied to dialog with a mobile robot," in AAAI, 2015.
- [17] D. Lu, S. Zhang, P. Stone, and X. Chen, "Leveraging commonsense reasoning and multimodal perception for robot spoken dialog systems," in *IEEE/RSJ International Conference IROS*, 2017, pp. 6582–6588.
- [18] R. Chitnis, L. P. Kaelbling, and T. Lozano-Pérez, "Integrating humanprovided information into belief state representation using dynamic factorization," in 2018 IEEE/RSJ IROS. IEEE, 2018, pp. 3551–3558.
- [19] S. Amiri, M. S. Shirazi, and S. Zhang, "Learning and reasoning for robot sequential decision making under uncertainty," in *Proceedings of* the AAAI Conference on Artificial Intelligence, vol. 34, no. 03, 2020, pp. 2726–2733.
- [20] B. Juba, "Integrated common sense learning and planning in pomdps," The JMLR, vol. 17, no. 1, pp. 3276–3312, 2016.
- [21] M. Sridharan, M. Gelfond, S. Zhang, and J. Wyatt, "Reba: A refinement-based architecture for knowledge representation and reasoning in robotics," *JAIR*, vol. 65, pp. 87–180, 2019.

- [22] S. Zhang, P. Khandelwal, and P. Stone, "Dynamically constructed (po) mdps for adaptive robot planning." in AAAI, 2017, pp. 3855–3863.
- [23] Y. Jiang, N. Walker, J. Hart, and P. Stone, "Open-world reasoning for service robots," in *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 29, 2019, pp. 725–733.
- [24] G. Cui, W. Shuai, and X. Chen, "Semantic task planning for service robots in open worlds," *Future Internet*, vol. 13, no. 2, p. 49, 2021.
- [25] Y. Zhu, J. Tremblay, S. Birchfield, and Y. Zhu, "Hierarchical planning for long-horizon manipulation with geometric and symbolic scene graphs," arXiv preprint arXiv:2012.07277, 2020.
- [26] M. Hanheide, M. Göbelbecker, G. S. Horn, A. Pronobis, K. Sjöö, et al., "Robot task planning and explanation in open and uncertain worlds," Artificial Intelligence, vol. 247, pp. 119–150, 2017.
- [27] Z. Zeng, A. Röfer, and O. C. Jenkins, "Semantic linking maps for active visual object search," in *ICRA*. IEEE, 2020, pp. 1984–1990.
- [28] Y. Wu, Y. Wu, A. Tamar, S. Russell, G. Gkioxari, and Y. Tian, "Bayesian relational memory for semantic visual navigation," in *Proceedings of the IEEE/CVF ICCV*, 2019, pp. 2769–2779.
- [29] H. M. Wallach, "Conditional random fields: An introduction," *Technical Reports (CIS)*, p. 22, 2004.
- [30] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [31] J. S. K. Yi, Y. Kim, and S. Chernova, "Incremental object grounding using scene graphs," arXiv preprint arXiv:2201.01901, 2022.
- [32] F. Yang, D. Lyu, B. Liu, and S. Gustafson, "Peorl: Integrating symbolic planning and hierarchical reinforcement learning for robust decisionmaking," arXiv preprint arXiv:1804.07779, 2018.
- [33] R. T. Icarte, T. Klassen, R. Valenzano, and S. McIlraith, "Using reward machines for high-level task specification and decomposition in reinforcement learning," in *ICML*, 2018.
- [34] Y. Hayamizu, S. Amiri, K. Chandan, K. Takadama, and S. Zhang, "Guiding robot exploration in reinforcement learning via automated planning," in *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 31, 2021, pp. 625–633.
- [35] M. Leonetti, L. Iocchi, and P. Stone, "A synthesis of automated planning and reinforcement learning for efficient, robust decision-making," *Artif. Intell.*, 2016.
- [36] K. Zheng, D. Bayazit, R. Mathew, E. Pavlick, and S. Tellex, "Spatial language understanding for object search in partially observed city-scale environments," in 2021 30th IEEE International Conference on RO-MAN. IEEE, 2021, pp. 315–322.
- [37] A. Wandzel, Y. Oh, M. Fishman, N. Kumar, L. L. Wong, and S. Tellex, "Multi-object search using object-oriented pomdps," in *ICRA*. IEEE, 2019, pp. 7194–7200.
- [38] J. K. Li, D. Hsu, and W. S. Lee, "Act to see and see to act: Pomdp planning for objects search in clutter," in 2016 IEEE/RSJ International Conference IROS. IEEE, 2016, pp. 5701–5707.
- [39] Y. Xiao, S. Katt, A. ten Pas, S. Chen, and C. Amato, "Online planning for target object search in clutter under partial observability," in 2019 International Conference on Robotics and Automation (ICRA). IEEE, 2019, pp. 8241–8247.
- [40] J. Johnson, R. Krishna, M. Stark, L.-J. Li, D. Shamma, M. Bernstein, and L. Fei-Fei, "Image retrieval using scene graphs," in *Proceedings of the IEEE conference on CVPR*, 2015, pp. 3668–3678.
- [41] M. L. Puterman, Markov decision processes: discrete stochastic dynamic programming. John Wiley & Sons, 2014.
- [42] J. S. Yedidia, W. T. Freeman, Y. Weiss, et al., "Generalized belief propagation," in NIPS, vol. 13, 2000, pp. 689–695.
- [43] K. Chandan, J. Albertson, X. Zhang, X. Zhang, Y. Liu, and S. Zhang, "Learning to guide human attention on mobile telepresence robots with 360 vision," in 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2021, pp. 5297–5304.
- [44] H. Kurniawati, D. Hsu, and W. S. Lee, "Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces." in RSS, vol. 2008. Zurich, Switzerland., 2008.
- [45] 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," *IJCV*, vol. 123, no. 1, pp. 32–73, 2017.
- [46] 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 preprint arXiv:1712.05474, 2017.
- [47] A. Ankan and A. Panda, "pgmpy: Probabilistic graphical models using python," in *Proceedings of SCIPY 2015*. Citeseer, vol. 10. Citeseer, 2015.

[48] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "Ros: an open-source robot operating system," in *ICRA workshop on open source software*, vol. 3, no. 3.2. Kobe, Japan, 2009, p. 5.