

Enhancing Object Search in Indoor Spaces via Personalized Object-factored Ontologies

Akash Chikhalikar, Ankit A. Ravankar, Jose Victorio Salazar Luces and Yasuhisa Hirata

Abstract—Personalization is critical for the advancement of service robots. Robots need to develop tailored understandings of the environments they are put in. Moreover, they need to be aware of changes in the environment to facilitate long-term deployment. Long-term understanding as well as personalization is necessary to execute complex tasks like *prepare dinner table* or *tidy my room*. A precursor to such tasks is that of *Object Search*. Consequently, this paper focuses on locating and searching multiple objects in indoor environments. In this paper, we propose two crucial novelties. Firstly, we propose a novel framework that can enable robots to deduce *Personalized Ontologies* of indoor environments. Our framework consists of a personalization schema that enables the robot to tune its understanding of ontologies. Secondly, we propose an *Adaptive Inferencing* strategy. We integrate *Dynamic Belief Updates* into our approach which improves performance in multi-object search tasks. The cumulative effect of personalization and adaptive inferencing is an improved capability in long-term object search. This framework is implemented on top of a multi-layered semantic map. We conduct experiments in real environments and compare our results against various state-of-the-art (SOTA) methods to demonstrate the effectiveness of our approach. Additionally, we show that personalization can act as a catalyst to enhance the performance of SOTAs. Video Link: <https://bit.ly/3WHk9i9>

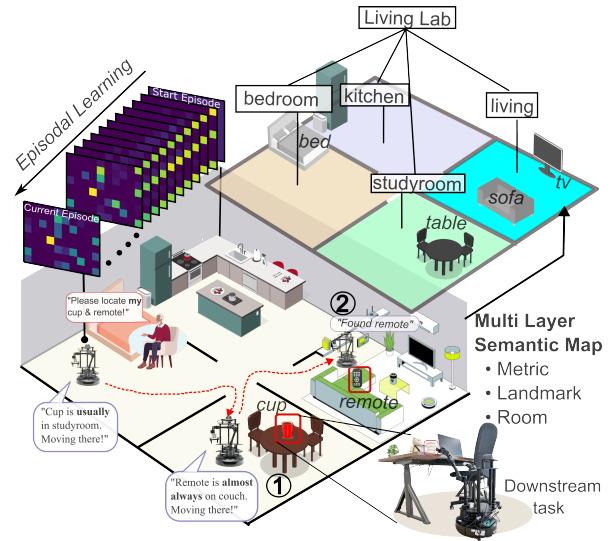


Fig. 1: Our personalized object search framework. The robot uses adaptive inferencing and dynamic belief updates in object search.

I. INTRODUCTION

Personalization is essential for facilitating multiple long-horizon tasks. Service robots, in indoor environments, commanded with *tidy up* tasks like ‘tidy up my living room’ need to be aware of which objects are in place or out of place. Information of such nature will be specific to each house. Another long-horizon task belongs to the *prepare* category, wherein robots can be asked to ‘prepare table for dinner’ or ‘prepare bag for gym’. Humans, who perform such tasks daily, subconsciously optimize their performance using memory, foresight and holistic decision-making. In order to assimilate into human-centric environments, robots will need a personalized understanding of user preferences and environments. The concept of personalization, however, goes beyond these task execution capabilities. In the future, robots will develop personalized relationships with their owners by providing enjoyment, companionship as well as care.

Object search tasks are primarily carried out in indoor environments. A typical indoor environment, apart from humans, consists of different rooms, various pieces of furniture, electronics (e.g., TVs, fridge), smaller items (e.g., tableware, phones) and so on. We categorize the larger, relatively immovable objects as ‘landmarks’. Our environment is thus a network of ‘object-landmark-rooms’. Searching for

landmarks is a simple task due to their fixed position in the long-term. On the other hand, searching for smaller, daily-use items is a challenging task due to their dynamic and uncertain positions. These items are the ‘objects’ for our search task. Object search is crucial for high-level downstream tasks like *tidy up* (Fig. 1).

For an effective object search, the core components are the inferencing mechanism and the robot’s internal belief. A convenient way to represent this belief is via an *Ontology*. Ontologies are a formal representation of concepts (e.g., rooms, landmarks) within a domain (house), including the relationships between them (e.g., in, on, and near). An instance of such ontology would be: ‘laptop-on-table-in-living room’. Ontologies have been used for efficient representations of indoor environments [1]. These ontologies would be different for every environment and hence require a personalized understanding. On the other hand, the inferencing strategy should consider trade-offs between minimizing search distance and maximizing search probability. The mechanism needs to adapt to its current observations and prior beliefs and accordingly generate a plan with an appropriate horizon for the robot to follow.

To this end, the main contributions of this paper include:

- A novel framework for *personalization* of ontologies of various object-landmark-room relationships using an episode-based learning strategy.
- A novel *adaptive inferencing* approach that adapts the

The authors are with the Department of Robotics, Graduate School of Engineering, Tohoku University, Sendai 980-8579, Japan. (e-mail: {a.k.chikhalikar, ankit, j.salazar, hirata}@srd.mech.tohoku.ac.jp)

search horizon based on the ontology priors. The approach includes *dynamic belief updates* that improves performance in multi-object search tasks.

- Quantitative results obtained in real environments and comparisons with various SOTAs to demonstrate significant improvement in object search tasks.

The remainder of the paper is structured as follows. We review related work and our position in the domain in Section II. Multi-layer semantic mapping is presented in Section III. Section IV details our novel framework for personalization of ontologies and is the crux of this submission. In Section V we discuss our inferencing strategy and dynamic belief updates along with some experimental results. A majority of experimental results and comparisons with SOTAs can be found in Section VI. Finally, Section VII concludes the paper with a discussion on future works.

II. RELATED WORK

Object search and retrieval is a mandatory precursor task for service robots. The problem has received increased attention since the realization of maturity in robot-based object detection, semantic segmentation and indoor navigation. We review previous research in the fields of personalization, and object search below.

A. Personalization

Personalization has been explored in different domains for improving the performance of robots. Learning from user-generated pre-arranged scenes, a Graph Neural Network (GNN) is used to obtain user preference vectors and perform personalized arrangements in [2]. LLMs have also shown potential as few-shot learners of personalization as demonstrated in [3]. A framework for personalization of risk assessment was presented in [4]. Spatio-temporal models for proactive robot assistance have been discussed in [5]. Our work, to the best of our knowledge, is a first in presenting mechanisms for deducing *personalized ontologies* without explicit supervision or human robot interactions. Due to this our method can be beneficial in real-world applications where the training data may be scarce or non-existent.

B. Object Search

Object search has been targeted by different researchers until now. Apart from service robotics, the problem is also interesting in the context of probabilistic inferencing and planning under uncertainty. A majority of previous works in solving the problem of object search fall into two categories: Partially Observable Markov Decision Process (POMDP) and Next Best View (NBV) methods.

POMDPs: Object search being a problem of decision-making under uncertainty can be inherently modelled as POMDP. The first work with this approach used belief road maps to reduce search spaces [6]. Due to their computational intractability in large domains, a hierarchical resolution of the environment was used in [7]. A multi-resolution POMDP for optimizing joint rotations to maximize visual area coverage was presented in [8]. In order to optimize the POMDP-based object search, a spatial correlational model was introduced in [9].

NBV: Seemingly the most used, NBV methods treat object search as an explicit inferencing problem and different strategies are used to determine the next location. A popular approach is [10], wherein the authors determined candidate points and selected them according to belief maximization. In [11], authors prepared heat maps according to the time of day to incorporate temporal dimension in object search. On the contrary, [12] proposes a temporal decay for improving search efficiency. They further improved upon their method with a heuristic that considers room sizes and distances [13]. Environment modelling has shown to improve priors for NBV [14]. Recently, graph structures are being used for modelling the environment and planning for NBV [15], [16]. NBV approaches can also be used for an open-vocabulary object search [17].

Another approach is [18] wherein Gaussian Mixture Models are used to prepare a costmap encoding probabilistic information and then an Informative Path Planning is carried out for search. In the absence of prior knowledge, the authors in [19] first determined the object category (e.g., toys, electronics) and planned robot trajectories using prior knowledge about the category.

Our work can be broadly categorized as an NBV approach. In this paper, we present a first attempt at personalization of ontologies to significantly improve performance in object search tasks. Additionally, we propose an adaptive inferencing approach for balancing long and short-horizon planning with dynamic belief updates for multi-object search. In the domain of object search, we summarize our contribution with respect to the state of the arts in Table I below.

Reference	Scope		Ontologies			Inferencing		
	Single object	Multi-object	Object-Landmark	Object-Landmark-Room	Personalized	Dynamic Belief Update	Fixed-horizon	Adaptive
[9]	✓	-	✓	-	-	✓	✓	-
[10]	✓	-	✓	-	-	✓	✓	-
[11]	✓	-	✓	-	-	-	✓	-
[12]	✓	✓	✓	✓	-	✓	✓	-
[15]	✓	✓	✓	-	-	✓	✓	-
Ours	✓	✓	✓	✓	✓	✓	-	✓

TABLE I: Positioning with respect to related works of object search

III. MULTI-LAYER SEMANTIC MAPPING

Object search is a complex task in human-centric environments. A semantic map functions as an important preliminary for object search. The type of semantic map the robot has access to, will significantly influence the level of personalization the robot can achieve, and the sequence of locations it would visit when assigned a task. This section briefly describes each layer of our semantic map.

Metric Layer: A 2-dimensional occupancy grid layer for which, we fuse data from RGB-D sensor, 2D Lidar and wheel odometry using visual SLAM techniques.

Semantic Layer: The semantic layer consists of ubiquitous landmarks such as sofa, bed, etc. We use YOLOv7 [20] for object detection. The corresponding point cloud data is randomly sampled, and the mean coordinates are computed for localization of the landmark on the map.

After localization, persistent tracking of landmarks is necessary to distinguish between new and prior observations of the objects. For such persistent tracking, we maintain a record of K prior observations as $\mathbf{P} = \{P_0, \dots, P_K\}$ and new observations in the latest frame as $\mathbf{C} = \{C_0, \dots, C_L\}$. We compute an association matrix ($\mathbf{D}_{k,l}$) using Euclidean distance as shown below:

$$\forall (k, l) \in (K, L) : \mathbf{D}_{k,l} = \sqrt{(\mathbf{P}_k - \mathbf{C}_l)^T (\mathbf{P}_k - \mathbf{C}_l)} \quad (1)$$

After the association matrix is computed, the association between new observations (i.e., set \mathbf{C}) and prior observations (i.e., set \mathbf{P}) is determined using the Hungarian algorithm [21]. A Kalman filter is then used to localize the landmark based on the associated new observation. For more details related to this implementation, please refer [22].

Segmentation and Classification Layer: Additional information regarding room regions is necessary for inferencing based on object-landmark-room ontologies. To achieve this, we segment the metric map using Voronoi-graphs [23], and then assign room labels to each region based on empirical data. In a segmented region R , containing a set of landmarks \mathbf{L} , with $\|\mathbf{L}\| = n$, the room label is computed as:

$$P(R_j | L_i) = \frac{F_{ij}}{\sum_{j=1}^{N_r} F_{ij}} \quad (2)$$

$$P(R_j | L_1, \dots, L_n) = \prod_{i=1}^n P(R_j | L_i)$$

Where, F_{ij} accounts for occurrences of the object category ‘i’ in the room category ‘j’ in the Places365 [24] dataset. The number of room categories N_r is set to 4 [Bedroom, Living room, Bathroom, Study room]. Finally, a Multi-layered Semantic Map is obtained (Fig. 2).



Fig. 2: Multi-layered Semantic Map

IV. PERSONALIZATION FRAMEWORK

We define personalization as the strategy by which a robot can understand preferences of a user pertaining to the object locations within an environment. We interpret the preferences as ‘personalized ontologies’. For example, a general ontology would be: cups-on-tables. However, if there are multiple tables in different rooms, such ontology wouldn’t be informative. If a user spends more time in the living room than in the study room, then ontologies should prioritize the presence of a cup on the living room table over the study

room table. In principle, a landmark in the study room would be weighted differently from a landmark in the living room. Hence, a personalized ontology would have different weights for ‘cup-on-table-in-living room’ vs ‘cup-on-table-in-dining room’. It is precisely these weights that we intend our robot to learn. Our strategy combines adaptive episode-based belief updates with a reliable termination criterion to obtain this personalized knowledge.

A. Ontology Personalization Schema

We use an episodic process to develop personalized ontologies between different objects, landmarks and rooms present in the environment. The true probabilistic relationships between object-landmarks-rooms, for these episodes, are shown in Fig. 3d. For each episode, we sample the locations of objects based on a multinomial distribution of the true values. Then the robot navigates through the environment and updates its belief. After each episode, the locations are changed as per the sampling and the process is reiterated. In this context, the term episode refers to the entire sequence where the robot navigates through the environment once, collects observations, and updates its belief. With this setup, the robot learns the personalized probabilistic relationships. We consider a uniform initial estimate, wherein, all objects have equal probability at all landmarks, irrespective of rooms. We represent the landmark-room tuple as $T = (L, R)$. Later (in Sec. VI-A), we analyze the results with different initial estimates.

After each episode, $P(O|T)$ changes (Fig. 3a-3c). To infer the probabilistic relationships for the current episode, a discounted sum of the previous episodes is considered as follows:

$$P_{\tau+1}(O|T) = \sum_{t=0}^H \gamma^t P_{\tau-t}(O|T) \quad (3)$$

Where, $\gamma (= 0.9)$ is the discounting factor and $H (= 5)$ is the finite horizon. During each episode, the landmarks are ranked based on the aggregate probabilities of all objects and visited according to a probabilistic greedy strategy.

B. Multi-resolution P-learning

Assume that during any episode an object O was observed at landmark L in Room R (i.e., $T = (L, R)$). To develop personalized ontologies, firstly, the robot needs to update the beliefs at both the room and landmark level. Secondly, updating the beliefs at all other locations (T') uniformly is also not appropriate. A landmark of the same category in a different room should be updated differently from a different landmark in the same room. Hence, a multi-resolution update is more suitable than mere frequency-based updates. Considering such conflicting conditions, we propose the equations for Multi-resolution P-learning as follows:

$$P_{t+1}(O|T) = P_t(O|T) + \frac{\beta}{\sqrt{N}} [1 - P_t(O|T)] \quad (4)$$

$$P_{t+1}(O|T') = P_t(O|T') - \frac{\beta S(T, T')}{\sqrt{N}} [1 - P_t(O|T')] \quad (5)$$

Eq. 4 updates the probability at the observed location while Eq. 5 updates for all other unobserved locations. $S(\cdot)$ is

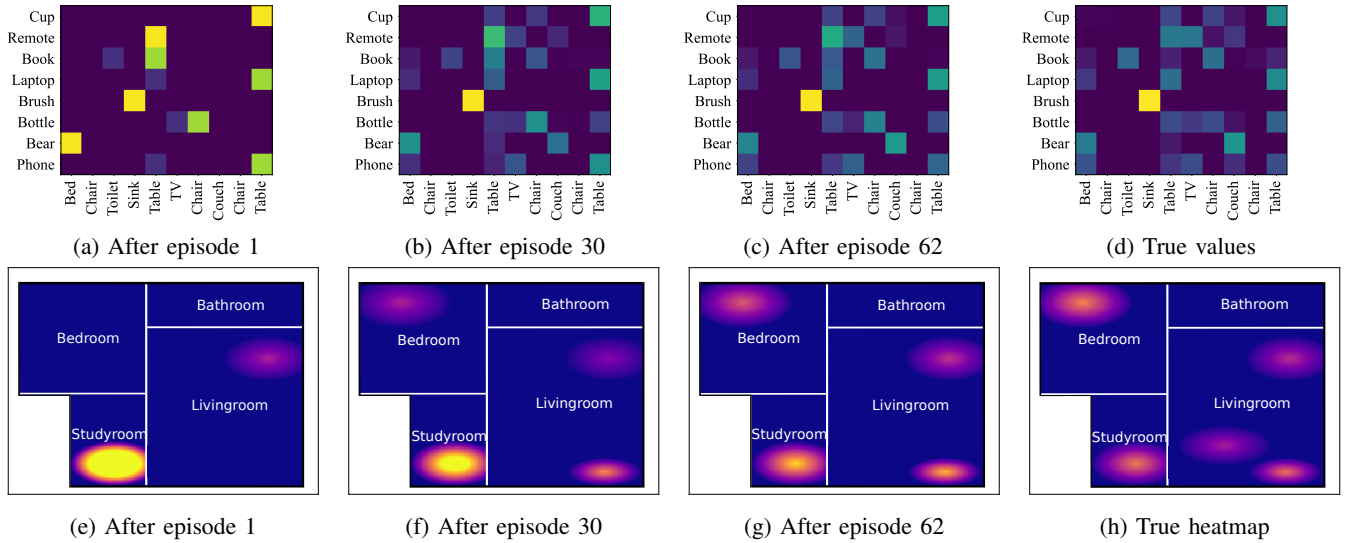


Fig. 3: In {a-c}, based on the uniform initial estimate, the probabilistic relationship between objects and landmarks is shown for different episodes. After every episode, the robot updates its belief (that is, $P(O|T)$), which starts resembling the true values {d}. {e-g} shows the heatmap progression for the object: Phone. The heatmap is populated based on observations (normalized) during each episode. {h} represents the heatmap based on the true *phone-landmark-room* relationship.

the ‘similarity kernel’ which helps in appropriate scaling of the probability updates at unobserved locations. It is defined as:

$$S(T, T') = 2 - [\delta(L, L') + \delta(R, R')] \quad (6)$$

Where, $\delta(\cdot)$ is the Kronecker-delta function. $\beta (= 0.1)$ is the learning rate and N is the episode number. The probabilities are re-normalized after each update step.

C. Confidence based termination

While the aforementioned episodal learning strategy will converge to true values by the law of large numbers, it is necessary to establish a quicker convergence criterion.

The presence of objects at different landmarks and rooms can be modelled as a multinomial distribution. However, determining confidence interval (C.I) for multinomial distributions requires a large number of samples. In contrast, whether a room contains an object or not can be modelled as a binomial distribution. We can thus compute C.Is for all rooms and leverage the intervals to establish a termination criterion.

Since the robot learns the probabilistic relationships between objects, landmarks and rooms, we can find the converse as follows:

$$P(R|O) = \frac{P(O|R)P(R)}{\sum_{\forall R} P(O|R)P(R)} \quad (7)$$

The next step is to calculate the C.I for $P(R|O)$. We use Wilson C.I as it provides better results when the distribution is skewed in either direction and is calculated as follows:

$$C.I = \frac{zN}{z^2 + N} \cdot \sqrt{\sigma^2 + \frac{z^2}{4N^2}} \quad (8)$$

Where, $z = 1.96$ corresponds to 95% quantile and σ is Standard Deviation for the binomial distribution. When the

95% C.I is within 0.05 for all objects, we terminate the episodes.

In Fig. 4, the progression of probability $P(O|R)$ against episode number, starting with the uniform initial estimate, is shown for two objects: laptop and phone. The corresponding heatmap progression for phone is given in Fig. 3e-3g. In our scenario, the laptop is much more likely to be in the study room than the bedroom whereas, the phone is equally likely to be in the living or the study. Consequently, the personalization results in $P(\text{Phone}|\text{Living}) \approx P(\text{Phone}|\text{Study})$ and $P(\text{Laptop}|\text{Study}) \gg P(\text{Laptop}|\text{Bedroom})$. It is important to note that the framework is robust to any sort of personalization as per user preferences.




Fig. 4: Probability of object given room

V. LONG-TERM OBJECT SEARCH

We propose an adaptive inferencing strategy with dynamic belief updates for enabling object search in the long-term. It is important to note that our inferencing approach is decoupled from the personalization strategy since it only depends on the ontology prior, irrespective of whether it is personalized.

A. Inferencing

The task of searching discrete locations to minimize the time to detection is NP-hard [25]. Hence, sub-optimal

solutions using approximations or heuristic-based methods are necessary. A short-horizon strategy yields better results if the beliefs are skewed towards a particular landmark. For e.g., visiting a location T , where $P(O|T) > 0.5$, will lead to a quicker search more than 50% of the time, irrespective of the distances. On the other hand, for a more uniform probabilistic distribution, a long-horizon strategy that considers a trade-off between probability and distance will lead to better results. Essentially, the strategy should adapt the search horizon according to the distribution. Thus, we use an inferencing strategy that is a hybrid between short-horizon and long-horizon planners. The authors in [26] suggest an Optimal Lookahead Search: a form of receding horizon control transposed to solve the search path-planning problem. Their method consists of a hyper-parameter ‘w’ which is the *lookahead window*. We propose a formal approach to determine this window based on the Cumulative Distribution Function (CDF) of the object (see Line 8, Algorithm 1). After, determining the lookahead window, we calculate the utility of each tuple in it as follows:

$$U_T = P(O|T) + \frac{\alpha}{Dist(T, A)} \quad (9)$$

Where, $Dist(T, A)$ represents the distance between the landmark and agent (robot) obtained via the A -star planner and α is the scaling parameter. With the utilities computed, object search is transposed as a Travelling Salesperson Problem (TSP) with known start and arbitrary ends. We use Google-OR TSP solver CP-SAT to solve this problem and determine the next landmark. The process begins again if the object is not found. We summarize our strategy in Algorithm 1.

B. Dynamic Belief Update

In long-term deployments, the robot may be repeatedly tasked with searching for objects, during which the state of the environment may not change much. In such scenarios, it is necessary for the inferencing strategy to adapt to information gained while the search is ongoing. When tasked with searching for an object(s), this manifests as the capability of tracking other objects during the search. We use a dynamic belief update (DBU) mechanism to do the same.

Consider that a robot has visited a set of landmarks $\mathbf{L}_s \subset \mathbf{L}$ before time t . Next, at time t , the robot approached landmark $L \in \mathbf{L} - \mathbf{L}_s$. In such cases, \forall Objects (O), the belief for the next time step is updated as follows:

If $O \in$ Found at L , $\forall L \in \mathbf{L} - \mathbf{L}_s - \{L\}$:

$$\begin{aligned} P_{t+1}(O|T) &= 1 \\ P_{t+1}(O|\tilde{T}) &= 0 \end{aligned} \quad (10)$$

Else:

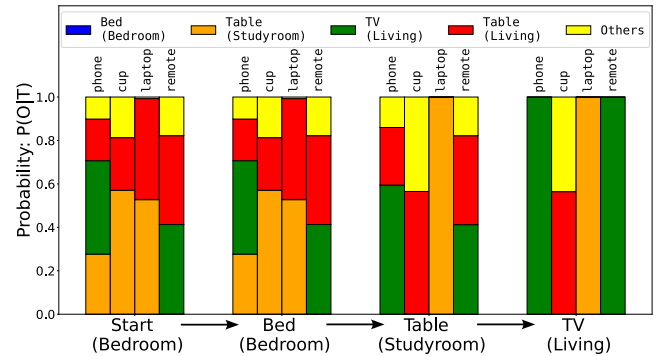
$$\begin{aligned} P_{t+1}(O|T) &= 0 \\ P_{t+1}(O|\tilde{T}) &= \frac{P_t(O|\tilde{T})}{[1 - P_t(O|T)]} \end{aligned} \quad (11)$$

The motivation behind DBU is to improve performance in multi-object search. When assigned with multiple tasks, the robot uses its personalized beliefs for the first search. For the next search, the updated beliefs due to our DBU mechanism are used. This reduces the overall time and distance travelled during the search. The qualitative impact of DBU is shown


Algorithm 1 Algorithm for inferencing

Require: Map : \mathcal{M} , $T = (L, R)$, Belief : $P(O|T)$
1: **Input:** Search Object(O^*)
2: **while** Not_Found **do**
3: Sort $[P_t(O^*|T)]$
4: **if** Max $[P_t(O^*|T)] > 0.5$ **then**
5: $T_{next} \leftarrow \text{argmax}_T P(O^*|T)$
6: $A = \text{Check}(T_{next}|O^*)$
7: **else**
8: WindowSize = Size[CDF $\{P_t(O^*|T)\} > 0.5]$
9: **for** i in WindowSize **do**
10: $U_T \leftarrow P_t(O^*|T) + \alpha/D(T, Agent)$
11: **end for**
12: $T_{next} \leftarrow TSPSolver(U_T, X)$
13: $A = \text{Check}(T_{next}|O^*)$
14: **end if**
15: $P_{t+1}(O|T) \leftarrow DBU(P_t(O|T), T) \quad \triangleright \forall Objects$
16: **end while**

in Fig. 5. The robot is tasked with searching for the Phone but simultaneously keeps track of all other objects during the search. Fig. 5 shows the internal belief propagation of the robot for other objects (i.e., Cup, Remote, and Laptop) as it visits different landmarks. After the Phone is found, the robot returns to the start position. If the robot is then tasked with searching for either of the above items, it can search for them faster. Quantitative results for the impact of DBU are given in the ablation study (Sec. VI-B).



(a) Belief propagation: $P(O|T)$



(b) Heatmap propagation for object: Bottle

Fig. 5: The robot is tasked with searching for the phone, which is under the TV. The landmarks visited are: Start (Bedroom) \rightarrow Bed (Bedroom) \rightarrow Table (Studyroom) \rightarrow TV (Living). During this search, the robot finds the laptop on the studyroom table and the remote under the TV.

VI. EXPERIMENTAL STUDIES

We present the results of our experiments in this section. We study the influence of initial estimates, personalization, and adaptive Inferencing through the experiments. All experiments were conducted in the Living Lab at Tohoku University. The Living Lab is designed to accelerate research

in service robots and human-robot interactions [27]. The facility resembles a house to the extent possible. We use the Turtlebot as our mobility platform and the ROS as our communication framework. We use YOLOv7, TEB local planner [28] and Spatio-Temporal Voxel Layer [29] for navigation.

A. Initial Estimate Comparison

Initial estimate is an important aspect of any episode-based approach. With our framework (Sec. IV), we compared the effectiveness of three different initial estimates in personalization. We sampled object locations from the multinomial distribution and the robot progressively personalized its belief over each episode from each initial estimate. The estimates are as follows:

- Uniform estimate: All objects have equal probability at all landmarks. This estimate is void of any bias.
- In-situ estimate: The robot navigates through the environment for five episodes and builds an initial estimate based on the frequency of observations.
- Empirical estimate: We mine the Places365 and COCO datasets to obtain probabilistic relationships.

$$P(O|T) = \frac{\mathbf{N}(O \cap T) + \lambda}{\mathbf{N}(O)} + \lambda n \quad (12)$$

Where, n represents the total size of classes in the dataset, $\mathbf{N}(\cdot)$ represents the tuples observed and $\lambda (= 0.5)$ accounts for unobserved tuples and is set according to Jeffrey-Perk’s law.

The number of episodes required for termination for each object are given in Table II. Five runs were conducted and the summary is presented below.


Object	Initial Estimate								
	Uniform			In-Situ			Empirical		
	Min.	Max.	Avg.	Min.	Max.	Avg.	Min.	Max.	Avg.
Cup	32	62	51	33	99	59	32	92	57
Remote	7	15	12	9	29	21	7	18	11
Book	23	49	39	28	91	57	14	63	27
Laptop	22	33	26	25	56	39	25	58	41
Toothbrush	7	7	7	7	7	7	7	7	7
Bottle	19	33	25	14	58	39	34	58	48
Teddy Bear	11	17	14	10	26	16	10	23	13
Phone	23	36	28	24	41	33	37	63	50

TABLE II: Episodes required for different initial estimates.

For each initial estimate, we also compute the KL Divergence statistic for object-landmark-room relationships for the episodes. This is shown in Fig. 6. For an object, the KL Divergence with smoothing ($\epsilon = 0.0001$) is computed as:

$$D_{KL}^{epi}(P||P_{true}) = \sum_{\forall T} P(O|T) \log \left[\frac{P(O|T) + \epsilon}{P_{true}(O|T) + \epsilon} \right] \quad (13)$$

Table II and Fig. 6 compares the performance of all three estimates. The performance of the In-situ estimate is dependent on the multinomial sampling during the initial episodes. The empirical estimate is more accurate, but its performance declines if the true values deviate further from the estimate. Due to these reasons, from here on, we consider all probabilistic relationships based on the uniform initial estimate.



(a) Object: Bottle

(b) Object: Cup

Fig. 6: KL Divergence for object-landmark-room relationships with respect to true values

B. Ablation Study

We consider personalization and DBU as separate modules for the ablation study. Subsequently, the baseline only consists of empirical estimate of priors (see Eq. 12) along with our inferencing strategy. We perform the study on a multi-object search. For this, we task the robot to search for two objects sequentially. The second object is only revealed after the robot completes the first task. The distance for first task (D_1), cumulative distance (D_T), and total number of landmarks visited (Ldmk.) during the two tasks are shown in Table III.

Search Tasks	A			A + B			A + B + C		
	D_1 (m)	D_T (m)	Ldmk.	D_1 (m)	D_T (m)	Ldmk.	D_1 (m)	D_T (m)	Ldmk.
Phone+Laptop	21.14	46.61	15	17.71	36.23	6	16.81	21.73	4
Cup+Remote	9.89	26.20	6	9.01	26.44	5	9.13	22.99	4
Phone+Book	23.42	55.63	18	14.92	33.92	6	15.21	33.93	6
Laptop+Bottle	28.89	41.67	11	8.82	20.11	4	8.91	14.64	3
Cup+Book	11.7	48.69	14	6.43	26.1	5	6.22	24.75	4
Average	19.01	43.76	12.8	11.38	29.69	5.4	11.26	21.64	4

Module A: Baseline Module B: Personalization Module C: DBU

TABLE III: Ablation study with personalization and DBU.

The effects of including personalization and DBU for multi-object search are evident. We see that inclusion of personalization improved the D_1 metric by 68% and reduced the overall distance by 47% as compared to the Baseline. DBU doesn’t affect the first task. Due to this, the average D_1 value remains the same compared to including personalization. However, the D_T metric reduces by 37% with the inclusion of DBU.

C. Comparisons with State of the Art and Baselines

From the comparison, we intend to quantify the impact due to our two proposed novelties, viz. ontology personalization and adaptive inferencing. During the comparative experiments, the robot uses the same path planning algorithm and configuration of the navigation stack for maintaining consistency. The differences exist in the probabilistic associations and inferencing strategies which are the core value propositions of our approach. We compare against the following methods:

- PKS: PKS is a probabilistic greedy approach where the next landmark to navigate to is the one with the highest probability of finding the object. This is a baseline method for benchmarking.
- LTOS: The authors in [15] minimize the cost of different traversal routes according to the target. Each route is weighted (WPL) as follows:

Object	PKS			LTOS			HSKOS			HOS			Ours (Ep=15)			Ours (Ep=30)			Ours (Ep=62)		
	D(m)	Ldmk.	T(s)	D(m)	Ldmk.	T(s)	D(m)	Ldmk.	T(s)	D(m)	Ldmk.	T(s)	D(m)	Ldmk.	T(s)	D(m)	Ldmk.	T(s)	D(m)	Ldmk.	T(s)
Cup	16.1	3	72.9	11.01	4.3	60.5	8.9	2.3	37.7	7.1	-	43.6	7.7	2	35.5	7.5	2	35.9	6.1	2	30.9
Book	29.01	7.3	126.9	19.3	6.3	101.9	24.6	8	150.2	11.1	-	60.1	20.8	4	86.7	17.03	3.7	81.9	14.02	2.3	69.2
Teddy	18.5	3.3	76.1	11.1	3.7	68.4	9.7	3.7	46.8	6.3	-	24.7	12.13	2	58.3	7.2	2	27.9	6.97	2	28.2
Bottle	14.8	3	61.6	13.3	5	71.6	15.7	3.7	71.4	10.1	-	61.3	12.8	3	61.4	9.4	2	48.7	9.4	2	50.4
Phone	16.9	4	86.2	17.7	6.3	88.8	19.2	5.3	95.6	10.5	-	57.2	9.2	2.3	43.8	10.01	2	41.4	9.8	2	35.1
Mean	19.1	4.1	84.7	14.5	5.1	78.2	15.6	4.6	80.3	8.2	-	49.4	12.7	2.7	57.1	10.2	2.4	47.2	9.3	2.1	42.8

TABLE IV: Comparison of the overall framework with state of the art and baseline methods.

$$\min WPL = \sum_{i=0}^n \frac{Dist(L_i, A)}{(1 + \alpha * P(O|L_i)) * 2^i} \quad (14)$$

- HSKOS: In [13], the authors consider the relative room size R_s , the distance between the agent and the room $Dist(R, A)$ and the probability of finding the object in the room $P(R|O)$ with the following Heuristic:

$$\arg \max_R G(R) = \frac{P(R|O)}{R_s + Dist(R, A)} \quad (15)$$

- Human Operated Search (HOS): Humans are considered the gold standard for executing service robotic tasks. The human teleoperates the robot to search for the object. The human is given a preview tour of the environment.
- Ours: We compare the SOTAs and baselines against our strategy. We consider three different levels of personalization. This corresponds to the beliefs after episode 15, 30, and 62.

During the experiments, the robot (or operator) is tasked with a single object search, starting from the same location. We place the target object at its three most likely locations, which are defined as the locations with the highest probability according to the values visualized in Fig. 3d. The distance (D) and time (T) to search are recorded. Averages of the data recorded from our experiments are shown in Table IV.

The results in Table IV show that our method outperforms both the SOTA as well as baseline and is close to human performance. LTOS has a better distance-to-landmark ratio since it is a long-horizon planner and HSKOS has fewer room transitions due to its hierarchical strategy. Despite this, our method outperforms HSKOS by around 33% in terms of time and distance and LTOS by 41% in time and 30% in distance. Sample trajectories are shown in Fig. 7 for reference.

We compared the frameworks as a whole in Table IV above. Now, we compare the influence of our adaptive inferencing alone. We consider the strategies of PKS, LTOS and HSKOS for the comparison. PKS and HSKOS are greedy planners while LTOS is a long-horizon planner. However, all the methods have a fixed horizon unlike ours which adapts the horizon based on the ontology priors and distances. In this comparison, all methods use the empirical estimate for the ontology prior. Other conditions are same as for Table IV and results are given in Fig. 8.

We see that only LTOS has comparable search distance (Fig. 8a) and only PKS has comparable landmark visits (Fig. 8b). However, a nuanced look amplifies the improvements due to our approach. While PKS may have comparable landmark visits, the distance travelled is much greater ($\approx 61\%$)




Fig. 7: Trajectories followed to search for Object: Teddy in the environment. PKS is omitted to improve clarity.

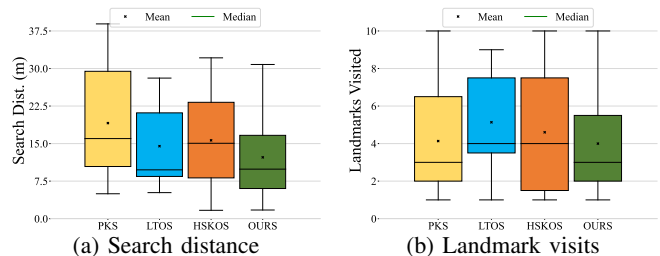


Fig. 8: Comparisons with respect to inferencing strategy.

since it does not consider distances while planning. On the contrary, LTOS may have comparable search distance but is worse ($\approx 33\%$) on landmark visits since it optimizes for worst-case scenarios. HSKOS has a more rounded performance but is worse on both the metrics. Additionally, HSKOS has a much larger inter-quartile range ($\approx 53\%$) for landmark visits than ours representing higher variability in performance.

D. Enhancement Study

We conduct an enhancement study to understand the standalone influence of personalization. We claim that other SOTA methods can also benefit from personalization of ontologies. In this context, we arm the SOTAs with our personalized relationships and use their respective planners. Other conditions are same as for Table IV and the results are given below.

The performance of LTOS and HSKOS improved by around 25% and 35% respectively, in terms of average distance. Correspondingly, HSKOS as well as LTOS required two less landmark visits on average. Thus, our personalization framework can be used as a catalyst for improving the

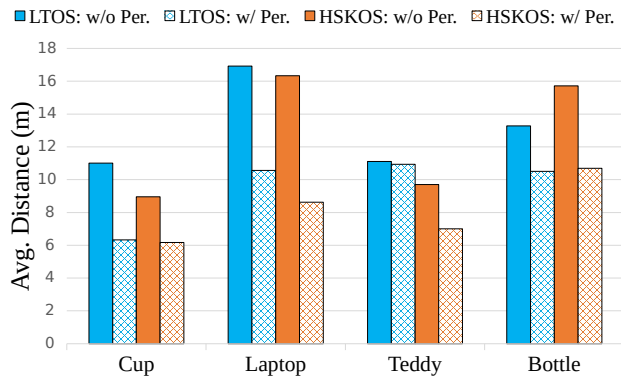


Fig. 9: Enhancement of SOTA with personalization.

performance of other SOTA methods.

VII. CONCLUSIONS AND DISCUSSIONS

We presented a novel approach for the *personalization* of ontologies of *object-landmark-room* relationships. Our methodology drastically improves efficiency of object search as the robot learns more about the environment. We tested our framework with three different initial estimates and determined that a uniform initial estimate fares better in personalization as compared to others. We also proposed an *adaptive inferencing* strategy with a formal approach for deciding the lookahead window. The ablation study confirmed that the inclusion of ontology personalization and DBU improved the performance compared to the baseline for multi-object search. Our framework was able to outperform multiple SOTAs in real environments. Our adaptive inferencing strategy fared better than the SOTA strategies for the same priors. We showed that an adaptive horizon provides better performance than fixed horizon during planning. Additionally, the use of our personalization framework as a catalyst improved the performance of SOTAs. We thus developed a framework that not only provides better results than SOTAs but is also able to bolster their performance.

An important assumption in our work, however, is the fixed position of landmarks. In its truest sense, long-term would mean that the locations of landmarks can change, or even new landmarks can be added or removed. Secondly, the object categories are limited, whereas an open-vocabulary object search would be more desirable for long-term applications.

REFERENCES

- [1] D. Fernandez-Chaves, J.-R. Ruiz-Sarmiento, N. Petkov, and J. Gonzalez-Jimenez, "Vimantic, a distributed robotic architecture for semantic mapping in indoor environments," *Knowledge-Based Systems*, vol. 232, p. 107440, 2021.
- [2] I. Kapelyukh and E. Johns, "My house, my rules: Learning tidying preferences with graph neural networks," in *Proceedings of the 5th Conference on Robot Learning*. PMLR, 2022, pp. 740–749.
- [3] J. Wu, R. Antonova, A. Kan, M. Lepert, A. Zeng, S. Song, J. Bohg, S. Rusinkiewicz, and T. Funkhouser, "Tidybot: personalized robot assistance with large language models," *Autonomous Robots*, vol. 47, no. 8, p. 1087–1102, Nov. 2023.
- [4] I. Sena, A. Chikhalikar, A. A. Ravankar, J. V. S. Luces, and Y. Hirata, "Context-aware risk estimation in home environments: A probabilistic framework for service robots," in *2025 IEEE International Conference on Robot and Human Interactive Communication (ROMAN)*, 2025.
- [5] M. Patel and S. Chernova, "Proactive robot assistance via spatio-temporal object modeling," in *Proceedings of The 6th Conference on Robot Learning*, vol. 205, 2023, pp. 881–891.

- [6] C. Wang, J. Cheng, J. Wang, X. Li, and M. Q.-H. Meng, "Efficient object search with belief road map using mobile robot," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3081–3088, 2018.
- [7] L. Holzherr, J. Förster, M. Breyer, J. Nieto, R. Siegwart, and J. J. Chung, "Efficient multi-scale pomdps for robotic object search and delivery," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6585–6591.
- [8] K. Zheng, Y. Sung, G. Konidaris, and S. Tellex, "Multi-resolution pomdp planning for multi-object search in 3d," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 2022–2029.
- [9] K. Zheng, R. Chitnis, Y. Sung, G. Konidaris, and S. Tellex, "Towards optimal correlational object search," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 7313–7319.
- [10] A. C. Hernandez, E. Derner, C. Gomez, R. Barber, and R. Babuška, "Efficient object search through probability-based viewpoint selection," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 6172–6179.
- [11] M. Mantelli, F. M. Noori, D. Pittol, R. Maffei, J. Torresen, and M. Kolberg, "Semantic temporal object search system based on heat maps," *Journal of Intelligent & Robotic Systems*, vol. 106, no. 4, 2022.
- [12] Y. Zhang, G. Tian, X. Shao, M. Zhang, and S. Liu, "Semantic grounding for long-term autonomy of mobile robots toward dynamic object search in home environments," *IEEE Transactions on Industrial Electronics*, vol. 70, no. 2, pp. 1655–1665, 2023.
- [13] M. Zhang, G. Tian, Y. Cui, Y. Zhang, and Z. Xia, "Hierarchical semantic knowledge-based object search method for household robots," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 8, no. 1, pp. 930–941, 2024.
- [14] A. Chikhalikar, A. A. Ravankar, J. V. S. Luces, and Y. Hirata, "Semantic-Based Multi-Object Search Optimization in Service Robots Using Probabilistic and Contextual Priors," *IEEE Access*, vol. 12, pp. 113 151–113 164, 2024.
- [15] F. Zhou, H. Liu, H. Zhao, and L. Liang, "Long-term object search using incremental scene graph updating," *Robotica*, vol. 41, 2023.
- [16] I. Remy, A. Gupta, and K. Leung, "Semantically-driven object search using partially observed 3d scene graphs," in *NeurIPS 2023 Foundation Models for Decision Making Workshop*, 2023.
- [17] A. Chikhalikar, A. A. Ravankar, J. V. S. Luces, and Y. Hirata, "Open vocabulary object search utilizing large language models and fuzzy inferencing," in *2025 IEEE/SICE International Symposium on System Integration (SII)*, 2025, pp. 345–351.
- [18] C. Wang, J. Cheng, W. Chi, T. Yan, and M. Q.-H. Meng, "Semantic-aware informative path planning for efficient object search using mobile robot," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 51, no. 8, pp. 5230–5243, 2021.
- [19] S. Hasegawa, A. Taniguchi, Y. Hagiwara, L. El Hafi, and T. Taniguchi, "Inferring place-object relationships by integrating probabilistic logic and multimodal spatial concepts," in *2023 IEEE/SICE International Symposium on System Integration (SII)*, 2023, pp. 1–8.
- [20] C.-Y. Wang, A. Bochkovskiy, and H.-Y. M. Liao, "Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors," *arXiv preprint arXiv:2207.02696*, 2022.
- [21] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.
- [22] A. Chikhalikar, A. A. Ravankar, J. V. S. Luces, S. A. Tafrihi, and Y. Hirata, "An object-oriented navigation strategy for service robots leveraging semantic information," in *2023 IEEE/SICE International Symposium on System Integration (SII)*, 2023, pp. 1–6.
- [23] R. Bormann, F. Jordan, W. Li, J. Hampp, and M. Hägele, "Room segmentation: Survey, implementation, and analysis," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016.
- [24] B. Zhou, A. Lapedriza, A. Khosla, A. Oliva, and A. Torralba, "Places: A 10 million image database for scene recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [25] K. Trummel and J. Weisinger, "The complexity of the optimal searcher path problem," *Operations Research*, vol. 34, no. 2, pp. 324–327, 1986.
- [26] T. H. Chung and J. W. Burdick, "A decision-making framework for control strategies in probabilistic search," in *2007 IEEE International Conference on Robotics and Automation*, 2007, pp. 4386–4393.
- [27] A. A. Ravankar, S. A. Tafrihi, J. V. S. Luces, F. Seto, and Y. Hirata, "Care: Cooperation of ai-robot enablers to create a vibrant society," *IEEE Robotics and Automation Magazine*, 2023.
- [28] R. Rösmann, F. Hoffmann, and T. Bertram, "Integrated online trajectory planning and optimization in distinctive topologies," *Robotics and Autonomous Systems*, vol. 88, pp. 142–153, 2017.
- [29] S. Macenski, D. Tsai, and M. Feinberg, "Spatio-temporal voxel layer: A view on robot perception for the dynamic world," *International Journal of Advanced Robotic Systems*, vol. 17, no. 2, 2020.