

# Interactive Continual Learning Architecture for Long-Term Personalization of Home Service Robots

Ali Ayub<sup>1\*</sup>, Chrystopher L. Nehaniv<sup>1</sup>, and Kerstin Dautenhahn<sup>1</sup>

**Abstract**— For robots to perform assistive tasks in unstructured home environments, they must learn and reason on the semantic knowledge of the environments. Despite a resurgence in the development of semantic reasoning architectures, these methods assume that all the training data is available a priori. However, each user’s environment is unique and can continue to change over time, which makes these methods unsuitable for personalized home service robots. Although research in continual learning develops methods that can learn and adapt over time, most of these methods are tested in the narrow context of object classification on static image datasets. In this paper, we combine ideas from continual learning, semantic reasoning, and interactive machine learning literature and develop a novel interactive continual learning architecture for continual learning of semantic knowledge in a home environment through human-robot interaction. The architecture builds on core cognitive principles of learning and memory for efficient and real-time learning of new knowledge from humans. We integrate our architecture with a physical mobile manipulator robot and perform extensive system evaluations in a laboratory environment over two months. Our results demonstrate the effectiveness of our architecture to allow a physical robot to continually adapt to the changes in the environment from limited data provided by the users (experimenters), and use the learned knowledge to perform object fetching tasks.

## I. INTRODUCTION

Significant advances in key robotics technologies, such as object recognition and object detection [1], [2], robot navigation and localization [3], [4], robot scene understanding [5], and semantic reasoning [6], have enabled home service robots to perform daily tasks in household environments. To be able to plan and perform tasks in an unstructured home environment, the first step is for the robot to understand the environment i.e. objects (e.g. cup, bottle), contexts (e.g. kitchen), and semantic relations between objects and contexts (e.g. cup is in a cupboard in the kitchen). However, each user’s environment can be unique, and it can continue to change over time, e.g. a user might replace old objects with new ones, or change objects’ locations. In such dynamic environments, developing general-purpose robots might not be the best solution, instead, a personalized approach [7] to developing home service robots might be more appropriate, where these robots can interact with their users to learn about their unique environments. A combination of interactive machine learning [8] and continual learning (CL) [9] is a promising solution to this problem, allowing users to

interactively teach their personal home service robots about their environments through repeated interactions over time.

CL has been the subject of growing interest in recent years, to develop machine learning (ML) models that can learn and adapt over extended periods of time [9], [10], [11]. This research has found valuable applications in the field of robotics, where it enables robots to consistently acquire knowledge about diverse objects, leveraging limited data inputs often provided by their users [12], [13]. Despite significant advancements, most CL methods are developed for object recognition, which is crucial but insufficient to perform home service tasks. Although a few recent CL methods were developed for knowledge graph embeddings for semantic object reasoning, these methods were only developed for symbolic datasets without any application to physical robots where new knowledge might not be provided in constrained setups used for evaluating these methods [14], [15]. Further, common across most CL methods for robotics is that these methods are not developed to learn from real-time interactions with their users. Although interactive ML methods have been developed that can allow robots to interact and learn object categories [8], these methods are only tested in single interaction scenarios and not over long-term interactions for CL.

For understanding unstructured environments beyond object recognition, semantic reasoning architectures [6], [16], [17], [18] have been developed for home service robots. Albeit promising advances in this field, these methods use static knowledge sources that are trained on a large corpus of symbolic datasets without much personalization or any interaction with their users [16], [17], [18], [19]. To the best of our knowledge, we know of no other work on integrating modern CL methods with semantic reasoning architectures that can allow physical robots to learn continually over a long period of time through human-robot interaction (HRI).

In this paper, we combine ideas from CL, interactive ML, and semantic reasoning literature to develop a novel interactive continual learning (ICL) architecture for a personalized home service robot that can interact with, learn from, and assist its users in their dynamic home environments. Our ICL architecture allows users to continually teach a home service robot the semantic knowledge of the environment i.e. objects and contexts in the home environment, and allows the robot to use the dynamic knowledge about the environment to perform object search and fetching tasks. Further, as users might provide limited data and expect quick results, our ICL architecture can learn in real time from only a few training examples of objects and contexts provided by the users.

This research was undertaken, in part, thanks to the Canada 150 Research Chairs Program.

<sup>1</sup>University of Waterloo, Waterloo, ON N2L 3G1, Canada  
{\*a9ayub, cnehaniv, kdautenh}@uwaterloo.ca

We perform extensive proof-of-concept evaluations of the ICL architecture on a Fetch mobile manipulator robot [20] over two months to continually learn 20 common household objects and two home contexts. Our results demonstrate the ability of our architecture to quickly adapt to the environment with limited data without catastrophically forgetting any past knowledge. This paper makes the following contributions:

- We propose a novel ICL architecture that integrates CL, interactive ML, and semantic reasoning to address multiple real-world challenges that have only been partially addressed in prior research, i.e. it can continually learn both objects and contexts through HRI in real-time with limited data, without catastrophically forgetting past knowledge.
- Unlike most prior CL research, we perform extensive CL experiments on a physical robot in a large indoor environment. Our results demonstrate the real-world applicability of our ICL architecture to adapt to its environment and perform object fetching tasks.

## II. RELATED WORK

### A. Continual Learning

Continual learning (CL) models aim to achieve a dual objective: to continually adapt and learn new information over time while simultaneously preserving past knowledge [9], [11], [21], [22]. In the literature, various categories of CL scenarios are defined based on factors such as shifts in input or output distributions and whether inputs and outputs share the same representation space [23]. Among these categories, class-incremental learning (CIL) is the most studied problem, in which an ML model is trained using labeled training data of different classes in each increment and is then evaluated on all the classes it has learned up to that point [9]. One of the primary challenges encountered by CIL models is *catastrophic forgetting*, a phenomenon in which the CIL model forgets the previously learned classes when it is exposed to new classes in an increment [24]. Numerous research avenues have been explored in the quest to mitigate catastrophic forgetting. These include replay-based techniques [9], [11], [25], where data from previously learned classes is stored and replayed during the learning of new classes, regularization techniques [26], [27] that constrain model parameters from changing drastically, and generative replay-based techniques [21], [28] which generate data for old classes using stored class statistics. Common across these CL techniques is the requirement of a large number of training examples per class, making them unsuitable for learning through HRI. In recent years, CL researchers developed models that can learn from only a few labeled training examples per class, a direction called few-shot class incremental learning (FSCIL) [29]. All of the CIL and FSCIL techniques were designed for the narrow context of object classification only, and mostly tested on static datasets in systematically controlled setups and not on physical robots learning in real-time through HRI [29], [10], [12], [30], [31].

### B. Semantic Reasoning Architectures

For real-world applications, robots need to reason on the semantic knowledge of objects (e.g. object categories, relations with other objects, contexts, etc.) to perform everyday tasks, such as object search and manipulation [16], [17], [19], [32], and navigation in dynamic and unknown environments [3]. Extensive research has been conducted in the past on semantic reasoning architectures to learn semantic knowledge of objects and environments, and reason on the mathematical relationships between different semantic properties of an object and between different objects [6], [16], [17], [18]. Most of these architectures use static knowledge sources [6] that are trained and tested on a large corpus of semantic reasoning datasets, which might not be personalized to unique user environments. A few semantic reasoning architectures have been tested on physical robots in constrained setups for object search tasks, however, they were not tested over the long term in dynamic environments.

Although recently a few preliminary research directions have been considered for developing CL methods for graph embeddings [14], [15], these methods simply apply CL methods for object classification on knowledge graphs. Further shortcomings of these methods are: 1) semantic CL methods continually learn knowledge graph embeddings only from symbolic data and do not learn object categories (from image data) simultaneously with the semantic knowledge embeddings, 2) these methods are computationally expensive and tested on static datasets in constrained CL setups without any interaction with the users. In this paper, we develop a novel ICL architecture to continually learn both objects and contexts in a dynamic environment over the long term.

## III. INTERACTIVE CONTINUAL LEARNING ARCHITECTURE

We build an interactive learning framework that includes some of the core abilities of a cognitive architecture, mainly perception, action, learning, adaptation, autonomy, memory, and reasoning. The complete interactive continual learning (ICL) architecture is depicted in Figure 1. All the components in ICL architecture are implemented as ROS nodes, communicating through topic subscriptions.

### A. Perception Layer

The perception layer manages the robot’s perception capabilities to extract information from the visual input from the robot’s RGB-D camera. Given that there can be many objects within a home environment, the first step is to individually detect each object in the robot’s camera view. For a home service robot, an object detection algorithm must be chosen based on real-time constraints, as users would want the robot to learn and perform object fetch tasks efficiently. The most used and appropriate algorithm for real-time object detection is YOLO [33]. We use a pre-trained YOLO for detecting objects in the RGB image taken from the robot’s camera. The pre-trained YOLO can also provide object categories, however, these predicted categories are not useful as the

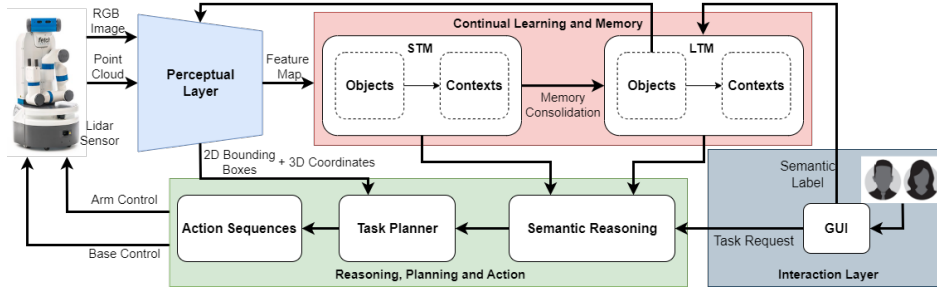


Fig. 1: Overview of the ICL architecture. It is composed of 1) perception, 2) interaction, 3) continual learning and memory, and 4) reasoning, planning, and action.

objects in a user’s home might be different from the pre-trained categories, and even for the same objects a user might want to label objects differently [13]. Given that there will only be a few training examples provided by the users, following the common practice in few-shot learning [34] and FSCIL [10] literature, the detected objects are further processed through a pre-trained CNN on ImageNet [1] for extracting latent features. The latent features can then be used for continual learning of objects. The point cloud data obtained from the depth sensor is combined with the detected objects in the RGB image to determine the 3D location of the object so that the manipulator can reach and pick the objects. A lidar sensor is also used for the robot’s base to navigate in the environment while avoiding obstacles.

### B. Interaction Layer

A key aspect of the ICL architecture is interaction with the users to allow them to teach the robot new knowledge and changes in the environment, and also request the robot to perform object fetch actions. We developed a graphical user interface (GUI) that allows the user to teach the robot objects and contexts in the environment. For teaching an object, the user can place the object in front of the robot’s camera, and type the label of the object in a textbox. Similarly, for teaching a context, say a kitchen, the user can manually drive the robot using a joystick to the context location and then type the name of the context and the location number in a text box. The robot captures the images in its camera view which are then processed by the perception layer (Section III-A). The user can also initiate an object fetch task by typing the name of the object to be fetched in a textbox. Examples of how the GUI is used for teaching and testing are shown in the supplementary video.

Similar to other layers in the architecture, the interactive technology (GUI) in this layer can be replaced. For example, a natural language interaction system, as used by many general-purpose cognitive architectures (e.g. SOAR [35]) can be added. However, to avoid any speech-to-text and natural language understanding errors during learning, we opted to use a GUI.

### C. Continual Learning and Memory

At any instant in time, the user can start a teaching interaction with the robot to either teach a new object, show

more viewpoints of a previously shown object, or teach a part of the context (e.g. an area of a kitchen). The underlying principles for learning object categories and contexts are similar. Therefore, we will first explain the continual learning and memory development procedure for object recognition and then describe some differences in learning contexts.

1) *Continual Learning*: At any time instant  $t$ , the user can teach the robot objects belonging to  $n^t$  different categories. Let’s consider that the robot has learned  $N^{t-1}$  number of object categories before. In this unconstrained setup, new objects can belong to any of the previously learned categories, which is different from the constrained class-incremental learning [9] setting where data for an object category is only available once. To teach the robot a new object belonging to category  $i$ , the user shows  $m_i$  number of object views to the robot. After processing the images through the perception layer, the architecture receives  $m_i$  number of CNN feature maps for the object. For CL, we use a network architecture similar to SUSTAIN [36] that clusters similar stimuli together. This network is inspired by the theories of pattern separation and memory integration in computational neuroscience [37]. For each object belonging to category  $i$  in increment  $t$ , the network begins with a simple solution and initiates a single cluster represented by the cluster centroid  $c_1^i$  computed as a mean of the feature maps in the cluster. Let’s assume that there are  $k_i$  number of clusters already in memory for object  $i$ . For each new  $j$ th feature map  $x_j$  for object  $i$  in increment  $t$ , L1 distance is computed between  $x_j$  and each of the  $k_i$  cluster centroids for object  $i$  over all the feature dimensions. Based on the distance from each cluster, a cluster activation is determined:

$$H_l^{act} = e^{-d_l/w_l} \quad (1)$$

Where  $d_l$  is the L1 distance between the  $l$ th centroid of object category  $i$  and  $x_j$ , and  $w_l$  is the weight attached to cluster  $l$ . In our experiments, we set  $w_l$  to be the total number of feature maps in the cluster. Clusters compete with each other to respond to the input pattern (the feature map  $x_j$ ), and the winning cluster is selected. For the winning cluster, say  $l_w$ th cluster, the output is calculated as,

$$H_{l_w}^{out} = \frac{(H_{l_w}^{act})^\beta}{\sum_{s=1}^{k_i} (H_s^{act})^\beta} H_{l_w}^{act} \quad (2)$$

Where  $\beta$  is an inhibition parameter for the regulation of cluster competition (for our experiments, we set it to 1). Note that based on the above equation, if many clusters are strongly activated, the output of the winning cluster is less, because of inhibition by competing clusters. In this network, only the winning cluster is selected and the output for all the other clusters is set to zero. As mentioned earlier, this network starts with a simple solution including one cluster only but recruits more clusters if needed. The cluster recruitment process is self-supervised. If the winning cluster's output ( $H_{l_w}^{out}$ ) is below a certain threshold ( $\tau$ ) a new cluster is recruited, and the centroid of the cluster is set to the input feature map. Otherwise, if the output is higher than the threshold, the cluster centroid is updated:

$$c_{l_w}^i = \frac{w_{l_w} c_{l_w}^i + x_j^i}{w_{l_w} + 1} \quad (3)$$

The weight  $w_{l_w}$  for the winning cluster is also incremented by 1. Since only the winning cluster is updated for each feature map, this effectively prevents cluster interference, thereby mitigating catastrophic forgetting. Further, as the network requires only a single pass of training data, the learning process is efficient and can be executed in real time, which is a crucial component for learning through HRI.

2) *Long-Term Memory*: The network learned through CL (Section III-C.1) is stored in the long-term memory (LTM) with all the clusters and the associated weights. Although the goal of CL is to avoid catastrophic forgetting, memory fading does happen in humans, especially for events that are not encountered regularly [38], [39]. For a home environment, it is possible that some objects might simply be replaced over time and the previously used objects are no longer in the home. In such cases, it would be useless to store this knowledge in memory. As the robot would not encounter any replaced objects in its environment, we incorporate a memory-fading process in the LTM. We use the predictive performance equation (PPE) model [39], [38] of learning and retention. Based on this model, we represent each weight  $w_i$  for cluster  $i$  in our network as:

$$w_i = w_i^\eta \times T^{-\alpha} \quad (4)$$

Where,  $\eta$  is the learning rate (held at 1 in [38]),  $\alpha$  is the decay rate, and  $T$  is called the *Model Time* in the PPE model, which is a weighted sum of the time since each of the  $n$  previous events  $T = \sum_{j=1}^n h_j t_j$ , where  $t_j$  is the time since the  $j$ th event. In our architecture,  $t_j$  represents the time since cluster  $i$  was activated the  $j$ th time.  $h_j$  is the weight assigned to each event which decreases with time:

$$h_j = \frac{t_j^{-u}}{\sum_{k=1}^n t_k^{-u}} \quad (5)$$

Where  $u$  controls the steepness of weighting, with higher values resulting in larger weights for more recent events. The weights  $h_j$  are normalized and sum to 1. Through this process, the weights of clusters that are not refreshed (details

below in Section III-C.3) are faded over time. We choose the decay rate  $\alpha$  to be small in LTM to allow for slow fading.

3) *Short-Term Memory*: Inspired by the dual-memory system in the mammalian brain [40], [41], our ICL architecture also contains a short-term memory (STM). Although the knowledge learned from the users is stored in LTM, STM stores the daily experiences of the robot encountered through its sensors and processed through the perception layer. STM stores the feature maps of encountered objects (and contexts) while performing tasks in the environment. Note that the clustering network is not used for STM because STM stores specific experiences, in accordance with the dual memory theory [40], whereas the LTM stores abstract knowledge represented as clusters. A weight parameter for each specific object instance (and context) is also stored to keep track of how many times the object was encountered. If the weight parameter exceeds a certain threshold  $\gamma$  (a hyperparameter), the corresponding feature map is consolidated into LTM through the learning process described in Section III-C.1. The memory consolidation process refreshes LTM and prevents relevant object and context memories from fading in LTM. The memory fading process described in Section III-C.2 is also applied to the weights of the feature maps stored in STM. The decay rate  $\alpha$  for STM is chosen to be much higher as STM fades quickly compared to LTM. The data stored in STM can be used for assistive tasks beyond object fetching, for example, helping users find a misplaced object in the home [42] or if they forgot to perform an important task, such as taking their medicine. We will expand our architecture for these applications in the future.

The above-mentioned process (Sections III-C.1 III-C.2 III-C.3) of learning and memory is described for object category learning. For context learning, the user would show a specific view of the context that can have multiple objects. As contexts can be described by the objects present in the scene, the architecture first generates the feature maps of all the objects in the image and then uses the object category network stored in LTM to find object categories in the context image. Following the principle of conceptual spaces [43], we generate a feature map where the dimensions of the feature map represent the object categories present in the context image, and the location of the context. After the feature map is generated, a context cluster network is trained using the CL process described in Section III-C.1. The LTM and STM for the context data also behave in the same way as for objects. Note that context recognition is dependent on object recognition, which means that the user would have to teach the objects first before they are shown in a context. However, a self-supervised learning method can be used to learn a new object in a context if it is not taught by the user.

In summary, the learning and memory mechanisms within the ICL architecture are inspired by cognitive science principles, and they enable efficient real-time learning of objects and contexts with limited data while addressing the challenge of catastrophic forgetting. Nevertheless, similar to the other layers, it's possible to substitute the CL method and memory components with alternative techniques in the literature.

#### D. Reasoning, Planning, and Action

The ICL architecture allows a mobile manipulator robot to perform object fetch tasks in a home environment. For this task, the architecture receives the object to be fetched from the user through the interaction layer (Section III-B). A masked conceptual space feature map is generated using the requested object category and passed through the context learning network in LTM which predicts the most probable context location for the object. The context location is then combined with a metric map (SLAM map [44]) of the environment, and a navigation planner in ROS is used to plan a collision-free path to the context location. The robot then moves to the context location. As there can be multiple objects present at a context location, the architecture localizes the objects using the perception layer (Section III-A) and gets the predicted object labels from the object learning network in LTM. If the object requested by the user is in the list of detected objects at the context location, the architecture plans an arm trajectory to move to the 3D coordinates of the object (from the perception layer). The robot’s arm then moves to the object’s location and picks it up. The navigation planner then plans a path back to the user’s location and the robot’s base moves to that location. Finally, the robot places the object on the table using the arm trajectory planner.

### IV. EXPERIMENTS

In this section, we describe the experimental setup of the ICL architecture on a physical robot and the object fetch experiment to evaluate ICL’s performance. The researchers assume the role of users throughout this section.

#### A. Experimental Setup

We evaluate our architecture on a Fetch mobile manipulator [20] equipped with a 7 DOF arm, a mobile base, a lidar, and an RGB-D camera sensor. For faster object detection and CL, an external workstation equipped with an Nvidia RTX 2060 GPU is used. We set up three mock home contexts (kitchen, home\_office, and dining\_area) in an indoor laboratory environment. We argue that our lab setup is comparable to mock home setups in previous work [45], [16], [19], [17], [18] due to the similar use of table-top environments for object placement. Further, prior works used only a few objects (<10) in physical robot experiments with objects placed individually without many other objects close to each other to avoid detection failures. Unlike these works, in our experimental setup, we use 20 different objects of various sizes and shapes that are commonly found in home environments, and we place them close to each other in the contexts (see Figure 2a for an example). We map the indoor environment using the lidar sensor and store the SLAM map in the robot’s onboard memory (Figure 2b).

The hyperparameters for the CL method and LTM and STM are set as follows: the threshold  $\tau$  in the CL network for recruiting a new cluster is set to 0.0003 for the object recognition network, and it is set to 0.13 for context recognition network. These thresholds were chosen using cross-validation on object and context data in the first increment.

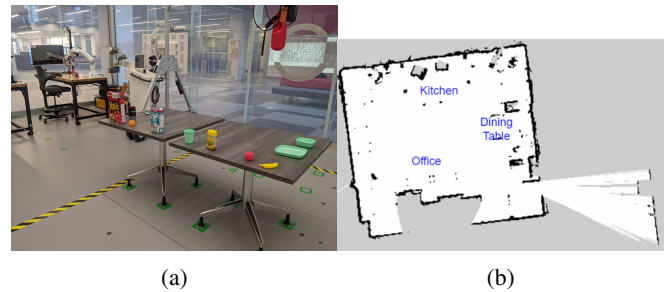


Fig. 2: (a) An example of the kitchen context in our experimental setup. Note the variation in object size, the closeness of objects, and distractions, such as robots and transparent walls behind. See supplementary video for examples of other contexts. (b) SLAM map of the indoor environment.

The decay rate  $\alpha$  is set to 0.2 for LTM, and 15 for STM. These decay rates were chosen such that it would take almost one month for the weights in LTM to decay to 50% their original value ensuring slow fading, and 2 days for the weights in STM to decay to  $\sim 0.003\%$  leading to complete memory fading. As we performed a few runs of the experiment each day, we measured the time since the activation of a cluster in days. The variable  $u$  in equation (5) was set to 0.6 as in [39]. Finally, we use a ResNet18 [46] as the feature extractor in the perceptual layer, as it is used commonly for other CL methods [29], [47], [12].

#### B. Continual Learning Experiment

We performed a long-term CL experiment to teach the robot objects and contexts continually and test it on object fetch tasks. The object fetch tasks can help us evaluate both the object and context learning pipelines given that both are needed to successfully find the object in the environment. In this experiment, we taught the object and context data to the robot over 16 increments. Each increment was usually a day apart, with some days consisting of two increments. For two increment days, we updated the event time in memory fading twice. For robustness, we performed this experiment 5 times randomizing the order of objects and contexts taught to the robot. Out of the 20 objects, 7 were assigned to the office and 13 to the kitchen. In each increment, the robot was either taught some objects only, some context updates only, or both objects and contexts. For the first 10 increments, after each increment, we tested the robot 5 times to fetch some of the objects it had learned so far. By 10 increments all 20 objects were taught to the robot, therefore, for the last 6 increments, only contextual updates were taught to the robot where some objects were randomly moved within contexts or completely removed. In the last 6 increments, we tested the robot 7 times each to fetch objects it had learned so far.

We report the task execution accuracy representing if the robot was able to fetch the correct object in each increment. A successful task execution depends on other components of ICL, therefore we report the accuracy of each of the components. 1) object classification accuracy (%) measures the percentage of correctly classified objects when the robot moves to context locations, 2) context classification accuracy

(%) measures the correct prediction of context location for the requested object, 3) manipulation error measures the number of times the robot was unable to physically pick, place, or move objects, 4) perception error measures the object detection failures in the perception layer, 5) navigation error measures the number of times the robot’s base failed to move to the desired location, and 6) execution time measures the time required by the robot to correctly fetch an object. Given the continual nature of our experiment, and the time required to teach and test the robot, it took approximately 2 months to complete 5 experimental runs (16×5 increments).

Table I shows results for the first 10 increments averaged over 5 runs of the experiment. The task execution accuracy starts at ~80% and decreases to ~56% (~24% decrease) after learning 20 objects and two contexts in 10 increments, therefore we notice some forgetting of past knowledge but no catastrophic forgetting. Diving deeper, we notice that the main source of task accuracy decrease is the object recognition network, where the accuracy starts at ~100% but decreases to ~73% (~27% decrease). The lower object recognition accuracy, however, could be because the network has more classes to classify, and the limited amount of data provided per object. On average, only **5-10 images per object** were shown to the robot for training in the experiment. We also trained the ICL architecture on all the object and context data at once (called joint training (JT)) and performed the object fetch task 15 times. The difference in task execution accuracy of ICL in the 10th increment compared to JT was only 6.1%, which indicates that ICL mitigates catastrophic forgetting effectively. The context recognition network maintained almost perfect accuracy for all increments without any forgetting. The only errors in the context recognition module were because of one misclassified object during the learning phase of the contexts. The object *fork* was predicted as *spoon* by the object recognition network during kitchen context learning. The other sources of errors in task execution were the manipulation and perception errors, both of which were consistently ~10-20% for all increments as these errors did not depend on object or context learning networks. Finally, the navigation stack performed quite reliably, albeit with a few failures in each experimental run.

For the last 6 increments, no new objects were taught but the contexts were updated by moving or removing objects. In these runs, both the object classification and task execution **accuracies remained similar** to the accuracies in the 10th increment. The reason is that our context recognition module had almost zero forgetting and quick adaptation ability, which ensured similar performance for task execution given that all other sources of errors remained consistent. The reason for consistent context recognition accuracy was that as contexts were updated, older memories about those contexts continued to fade which resulted in updated context clusters winning during cluster competition in the context recognition network. In terms of the execution time, on average it took **97 seconds** to fetch an object and place it on the dining table. The reason for the short execution

	Objects	Contexts	Tasks	Execution Time
1	100.0 ± 0	100 ± 0	80.0 ± 0.1	96.6 ± 3.6 s
2	93.1 ± 5.5	96.0 ± 0.1	72.0 ± 0.1	96.0 ± 4.6 s
3	91.8 ± 7.3	96.0 ± 0.1	72.0 ± 0.1	93.0 ± 3.2 s
4	88.7 ± 5.0	100 ± 0	70.0 ± 0.1	98.2 ± 2.3 s
5	85.6 ± 4.8	92.9 ± 0.1	66.0 ± 0.2	99.0 ± 3.5 s
6	83.7 ± 4.3	100 ± 0	64.0 ± 0.2	97.1 ± 3.4 s
7	80.1 ± 3.7	96.0 ± 0.1	62.0 ± 0.2	98.4 ± 2.6 s
8	78.2 ± 3.4	96.0 ± 0.1	62.0 ± 0.2	100.6 ± 3.9 s
9	75.2 ± 2.3	100 ± 0	58.0 ± 0.3	99.4 ± 2.4 s
10	73.3 ± 0.6	100 ± 0	56.0 ± 0.3	97.5 ± 3.1 s
JT	+3.7	+0	+6.1	+0

TABLE I: Incremental task execution, object, and context recognition accuracy (%), and task execution time (seconds) over first 10 increments averaged over 5 runs. JT represents the ICL architecture trained on all the data at once without any continual learning. The row for JT shows the accuracy and time differences compared to the 10th increment.

time was that the robot went to the correct object locations without wandering around in the environment. For the last 6 increments, we also tested the architecture without any memory decay, which resulted in higher task execution times (~130 seconds, **33 seconds** higher) because the robot went to older locations of objects first which wasted some time. Finally, the training and testing time for the CL networks were ~1 second and 0.004 seconds, respectively, ensuring real-time learning and prediction capabilities of the architecture. Examples of teaching and object-fetching tasks are shown in the supplementary video file.

## V. CONCLUSIONS

In this paper we tackle the challenging problem of simultaneous continual learning of object and context knowledge in a dynamic home environment. We developed a novel ICL architecture that builds on the core cognitive principles of learning and memory and allows users to directly teach objects and contexts in their unique home environment. Our results on a mobile manipulator robot demonstrate the ability of ICL to continually adapt to the changes in the environment through limited data provided by the user. We hope that this work can serve as a first step toward developing personalized home service robots that can interact with, learn from, and perform assistive tasks for their users over the long term.

## VI. LIMITATIONS AND FUTURE WORK

Given the long-term nature of the experiments with a physical robot, we did not perform an extensive hyperparameter analysis and compared our proposed approach with JT only. In the future, we hope to perform more experiments to analyze the tradeoff of hyperparameter values and use other SOTA CL models. In this paper, we only considered the supervised learning ability of ICL, however, in the future, we will expand it to perform self-supervised CL. Additionally, the experimental setup did not have overlapping objects in different contexts. In the future, we will analyze how ICL performs when there are overlapping objects in different contexts. Finally, we will conduct long-term HRI studies with human users in home environments to understand the personalization capabilities of ICL in the real world.

## REFERENCES

- [1] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "Imagenet large scale visual recognition challenge," *Int. J. Comput. Vision*, vol. 115, no. 3, pp. 211–252, Dec. 2015.
- [2] R. Girshick, "Fast r-cnn," in *The IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [3] N. Fulda, N. Tibbetts, Z. Brown, and D. Wingate, "Harvesting common-sense navigational knowledge for robotics from uncurated text corpora," in *Conference on Robot Learning*. PMLR, 2017, pp. 525–534.
- [4] J. Wang, V. A. Shim, R. Yan, H. Tang, and F. Sun, "Automatic object searching and behavior learning for mobile robots in unstructured environment by deep belief networks," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 11, no. 3, pp. 395–404, 2018.
- [5] A. Ayub and A. R. Wagner, "Centroid based concept learning for rgb-d indoor scene classification," in *British Machine Vision Conference (BMVC)*, 2020.
- [6] W. Liu, A. Daruna, M. Patel, K. Ramachandruni, and S. Chernova, "A survey of semantic reasoning frameworks for robotic systems," *Robotics and Autonomous Systems*, vol. 159, p. 104294, 2023.
- [7] J. Saunders, D. S. Syrdal, K. L. Koay, N. Burke, and K. Dautenhahn, "Teach Me—Show Me"—End-user personalization of a smart home and companion robot," *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 1, pp. 27–40, 2016.
- [8] A. L. Thomaz and M. Cakmak, "Learning about objects with human teachers," in *2009 4th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2009, pp. 15–22.
- [9] S.-A. Rebuffi, A. Kolesnikov, G. Sperl, and C. H. Lampert, "iCaRL: Incremental classifier and representation learning," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [10] A. Ayub and A. R. Wagner, "Cognitively-inspired model for incremental learning using a few examples," in *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020.
- [11] T. L. Hayes, G. P. Krishnan, M. Bazhenov, H. T. Siegelmann, T. J. Sejnowski, and C. Kanan, "Replay in deep learning: Current approaches and missing biological elements," *Neural computation*, vol. 33, no. 11, pp. 2908–2950, 2021.
- [12] A. Ayub and A. R. Wagner, "Tell me what this is: Few-shot incremental object learning by a robot," *arXiv:2008.00819*, 2020.
- [13] A. Ayub, J. Mehta, Z. De Francesco, P. Holthaus, K. Dautenhahn, and C. L. Nehaniv, "How do human users teach a continual learning robot in repeated interactions?" in *IEEE International Conference on Robot and Human Interactive Communication (RO-MAN) (Accepted)*, 2023.
- [14] A. Daruna, M. Gupta, M. Sridharan, and S. Chernova, "Continual learning of knowledge graph embeddings," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1128–1135, 2021.
- [15] Y. Cui, Y. Wang, Z. Sun, W. Liu, Y. Jiang, K. Han, and W. Hu, "Lifelong embedding learning and transfer for growing knowledge graphs," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 4, 2023, pp. 4217–4224.
- [16] Z. Wang, G. Tian, and X. Shao, "Home service robot task planning using semantic knowledge and probabilistic inference," *Knowledge-Based Systems*, vol. 204, p. 106174, 2020.
- [17] A. Daruna, L. Nair, W. Liu, and S. Chernova, "Towards robust one-shot task execution using knowledge graph embeddings," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 11 118–11 124.
- [18] W. Liu, D. Bansal, A. Daruna, and S. Chernova, "Learning instance-level n-ary semantic knowledge at scale for robots operating in everyday environments," *Autonomous Robots*, pp. 1–19, 2023.
- [19] S. Liu, G. Tian, Y. Zhang, M. Zhang, and S. Liu, "Service planning oriented efficient object search: A knowledge-based framework for home service robot," *Expert Systems with Applications*, vol. 187, p. 115853, 2022.
- [20] M. Wise, M. Ferguson, D. King, E. Diehr, and D. Dymesich, "Fetch and freight: Standard platforms for service robot applications," in *IJCAI, Workshop on Autonomous Mobile Service Robots*, 2016.
- [21] A. Ayub and A. R. Wagner, "EEC: Learning to encode and regenerate images for continual learning," in *International Conference on Learning Representations (ICLR)*, 2021. [Online]. Available: <https://openreview.net/forum?id=1Waz5a9lcFU>
- [22] X. Shi, D. Li, P. Zhao, Q. Tian, Y. Tian, Q. Long, C. Zhu, J. Song, F. Qiao, L. Song, Y. Guo, Z. Wang, Y. Zhang, B. Qin, W. Yang, F. Wang, R. H. M. Chan, and Q. She, "Are we ready for service robots? the openloris-scene datasets for lifelong slam," *arXiv:1911.05603*, 2019.
- [23] V. Lomonaco and D. Maltoni, "Core50: a new dataset and benchmark for continuous object recognition," in *Proceedings of the 1st Annual Conference on Robot Learning*, vol. 78, 2017, pp. 17–26.
- [24] R. M. French, "Dynamically constraining connectionist networks to produce distributed, orthogonal representations to reduce catastrophic interference," *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, pp. 335–340, 2019.
- [25] Y. Wu, Y. Chen, L. Wang, Y. Ye, Z. Liu, Y. Guo, and Y. Fu, "Large scale incremental learning," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [26] J. Kirkpatrick, R. Pascanu, N. C. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell, "Overcoming catastrophic forgetting in neural networks," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 114, no. 13, pp. 3521–3526, 2017.
- [27] Z. Li and D. Hoiem, "Learning without forgetting," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 12, pp. 2935–2947, Dec 2018.
- [28] O. Ostapenko, M. Puscas, T. Klein, P. Jahnichen, and M. Nabi, "Learning to remember: A synaptic plasticity driven framework for continual learning," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019, pp. 11 321–11 329.
- [29] X. Tao, X. Chang, X. Hong, X. Wei, and Y. Gong, "Topology-preserving class-incremental learning," in *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIX*. Springer-Verlag, 2020, p. 254–270.
- [30] M. Dehghan, Z. Zhang, M. Siam, J. Jin, L. Petrich, and M. Jagersand, "Online object and task learning via human robot interaction," in *2019 International Conference on Robotics and Automation (ICRA)*, May 2019, pp. 2132–2138.
- [31] A. Ayub and C. Fendley, "Few-shot continual active learning by a robot," *Advances in Neural Information Processing Systems*, vol. 35, pp. 30 612–30 624, 2022.
- [32] M. Scheutz, E. Krause, B. Oosterveld, T. Frasca, and R. Platt, "Spoken instruction-based one-shot object and action learning in a cognitive robotic architecture," in *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, ser. AAMAS '17. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2017, p. 1378–1386.
- [33] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [34] W.-Y. Chen, Y.-C. Liu, Z. Kira, Y.-C. F. Wang, and J.-B. Huang, "A closer look at few-shot classification," in *International Conference on Learning Representations*, 2019.
- [35] J. E. Laird, A. Newell, and P. S. Rosenbloom, "Soar: An architecture for general intelligence," *Artificial intelligence*, vol. 33, no. 1, pp. 1–64, 1987.
- [36] B. C. Love, D. L. Medin, and T. M. Gureckis, "Sustain: A network model of category learning," *Psychological Review*, vol. 111, no. 2, p. 309–332, 2004.
- [37] M. L. Mack, B. C. Love, and A. R. Preston, "Building concepts one episode at a time: The hippocampus and concept formation," *Neuroscience Letters*, vol. 680, pp. 31–38, 2018.
- [38] M. G. Collins, F. Sense, M. Krusmark, and T. S. Jastrzembski, "Improving predictive accuracy of models of learning and retention through bayesian hierarchical modeling: An exploration with the predictive performance equation," in *CogSci*, 2020.
- [39] M. M. Walsh, K. A. Gluck, G. Gunzelmann, T. Jastrzembski, and M. Krusmark, "Evaluating the theoretic adequacy and applied potential of computational models of the spacing effect," *Cognitive science*, vol. 42, pp. 644–691, 2018.
- [40] J. L. McClelland, B. L. McNaughton, and R. C. O'Reilly, "Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory," *Psychological Review*, vol. 102, no. 3, p. 419–457, 1995.

- [41] T. Kitamura, S. K. Ogawa, D. S. Roy, T. Okuyama, M. D. Morrissey, L. M. Smith, R. L. Redondo, and S. Tonegawa, "Engrams and circuits crucial for systems consolidation of a memory," *Science*, vol. 356, no. 6333, pp. 73–78, 2017.
- [42] J. Shah, A. Ayub, C. L. Nehaniv, and K. Dautenhahn, "Where is my phone? towards developing an episodic memory model for companion robots to track users' salient objects," in *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 621–624.
- [43] P. Gärdenfors, "Conceptual spaces as a framework for knowledge representation," *Mind and Matter*, vol. 2, no. 2, pp. 9–27, 2004.
- [44] M. Csorba, "Simultaneous localisation and map building," Ph.D. dissertation, University of Oxford, 1997.
- [45] 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.
- [46] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [47] C. Zhang, N. Song, G. Lin, Y. Zheng, P. Pan, and Y. Xu, "Few-shot incremental learning with continually evolved classifiers," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 12 455–12 464.