

Selective visual memory replay with self-evaluation in cognitive robots based on global workspace framework

Wenjie Huang^{a,*,1}, Antonio Chella^{b,c,2}, Angelo Cangelosi^{a,c,3}

^a Manchester Centre for Robotics and AI, University of Manchester, UK

^b University of Palermo, Italy

^c ICAR-CNR, Italy

ARTICLE INFO

Action editor: Serge Thill

Keywords:

Cognitive architecture
Continual learning
Consciousness
Memory
Global Workspace Theory

ABSTRACT

Learning capability for artificial systems is a well-studied topic, with various schemes enabling the system to develop knowledge continually. Methods based on memory replay are commonly adopted in the literature. This work presents a consciousness-based model integrated with a continual learning scheme for class-incremental learning in visual recognition. We suggest a reciprocal relation between memory maintenance and the learning activity of the system based on psychological evidence. The memory capability fits the continual learning problem. In return, a self-evaluation of knowledge mechanism is proposed for the robot to discriminate the important learning data during interactions to alleviate the memory constraint without degrading the distribution representation of abnormal data. The implemented robotic agent autonomously puts more effort into learning novel knowledge without human intervention. The cognitive architecture based on the Global Workspace Theory for the robotic agent is presented, with which the agent can automatically associate information from different modalities. Memory consolidation is implemented to run in parallel to the memory formation process. The work is validated in a class-incremental object recognition experiment on a robotic agent. The results show that the agent automatically balances the memory distribution for learning and maintains a relatively small set of samples during learning.

1. Introduction

Cognitive robotics is still at the infant stage of its development with the increasing interest in the field of Artificial General Intelligence (AGI). Most people would admit that the lack of a well-defined, commonly accepted definition of general cognition for robots is one of the main barriers to achieving it. Attempts to define the framework are made by [Pei \(2008\)](#) and [Samsonovich \(2012\)](#). They define general cognitive intelligence from different aspects. [Samsonovich \(2012\)](#) emphasises different capabilities of human cognition such as memory, emotion, self-consciousness and so on, while [Pei \(2008\)](#) tries to explain different paths to the goal at different levels such as structure, behaviour, capability, function and principles of human cognition. [Sukhobokov, Belousov, Gromozdov, Zenger, and Popov \(2024\)](#) provides the latest review on relevant topics of AGI and cognitive architectures. Amongst various sub-problems in this grand area of intelligent systems, we aim to develop an artificial system based on

human-inspired cognitive architecture, by incrementally introducing key features to cognitive robots.

This work focuses on the learning aspect of cognitive robotics ([Churamani, Kalkan, & Gunes, 2020](#); [Lesort, Lomonaco, Stoian, Maltoni, Filliat, & Díaz-Rodríguez, 2020](#)). Although many studies implement learning algorithms on robotic platforms for automation purposes (e.g., [Berto, Costa, Simões, Gudwin, and Colombini \(2024\)](#), [Costa, Pinho, Botto, and Lima \(2023\)](#), [DeChant, Akinola, and Bauer \(2023\)](#), [Morrison, Corke, and Leitner \(2020\)](#), [Qian, You, Zhou, Xu, and He \(2023\)](#) and [Zhong, Peniak, Tani, Ogata, and Cangelosi \(2019\)](#)), the challenge of continual learning in robotic systems witnesses limited progress. Specifically, continual learning ([Van de Ven & Tolias, 2019](#)) for robots means the agent is supposed to learn gradually via interactions with the environment. [Lesort et al. \(2020\)](#) reviewed existing approaches for continual learning and listed some challenges of continual learning in robotic applications specifically. The work of Churamani

* Corresponding author.

E-mail addresses: wenjie.huang@manchester.ac.uk (W. Huang), antonio.chella@unipa.it (A. Chella), angelo.cangelosi@manchester.ac.uk (A. Cangelosi).

¹ Research Associate.

² Professor of Robotics.

³ Professor of Machine Learning and Robotics.

and colleagues (Churamani et al., 2020) is similar, while they compensate (Lesort et al., 2020) with different challenges of continual learning in robots, especially in interactions with human peers. Both works have discussed the importance of automatic data extraction and memory during the interactive learning process of the robotic agent. In this paper, we implemented an architecture based on the Global Workspace Theory (GWT) (Baars, 1993), which illustrates the automatic information association between different modalities of learning. The continual learning capability of the system is achieved with memory replay. With a maintained memory buffer, the agent could retain knowledge of previous data while learning new information. While most existing works in continual learning emphasise the role of memory in learning, we highlight the reciprocal relationship between memory and learning activities.

As the body of continual AI grows, interesting approaches were proposed, including symbolic, deep learning or hybrid methods. Due to the advantage of interpretability, symbolic and hybrid methods are often studied in cognitive systems. The work of MacLellan and colleagues (MacLellan, Harpstead, Alevan, & Koedinger, 2024) is a recent example of symbolic methods. It makes use of symbolic memory to avoid the forgetting issue in the system. Also, symbolic methods were often adopted in cognitive systems. ICARUS (Langley & Choi, 2006), ACT-R (Ritter, Tehranchi, & Oury, 2019), and Soar (Laird, Newell, & Rosenbloom, 1987) are good examples, deploying symbolic components in the learning module design. However, symbolic and hybrid methods usually require expertise in designing the learning rules or data structure. In addition, they are hard to scale up for large-scale datasets. These led us to the group of research in deep learning approaches. Approaches in this group include regularisation-based (Kirkpatrick et al., 2017), expansion-based (Rusu et al., 2016; Schwarz et al., 2018), parameter-isolation based (Mallya & Lazebnik, 2018; Serra, Suris, Miron, & Karatzoglou, 2018), and data-based (Lopez-Paz & Ranzato, 2017; Rebuffi, Kolesnikov, Sperl, & Lampert, 2017) techniques. Amongst those, this work focuses on the data replay approach. Though we do not claim the implemented agent to be a sufficient cognitive agent yet, we modelled the relation of memory and knowledge accumulation for cognitive agents. Future research would extend this work to approach sufficient cognitive agents. Furthermore, to implement the memory-replay approach for continual learning in real-time applications for robots, we have to cope with the memory growth problem, as mobile systems like robots usually have limited memory allowance and computational limits in the continual learning setting. Thus, selective replay (Hayes, Kafle, Shrestha, Acharya, & Kanan, 2020) comes to the interest of this paper. By adopting selective memory replay, the performance of learning tasks is proven to be improved in many studies (Hayes, Cahill, & Kanan, 2019; Hayes et al., 2020; Rebuffi et al., 2017). Rebuffi et al. (2017) introduced the first approach, which used a fixed allocated memory, which stores samples closest to the feature mean of each class. Though this approach alleviates the memory constraint, it lacks a good representation of samples from the class mean. However, this is not consistent with the human memory study. Recent studies have suggested the novelty-boosted memory effect in human cognition (Lorents, Ruitenbergh, & Schomaker, 2023). Different from forcing a fixed amount of memory, a selective replay approach for the continual learning setting is studied by Hayes and Kanan (2021). In their work, they adopted a selective policy to choose samples from the replay buffer while they held an unlimited memory buffer. They used a computational policy to decide the possibility of replay for each sample in the buffer. However, this approach is challenged by the memory constraint for mobile robots due to the unlimited memory design. For this, we introduce an approach to store a part of the experienced samples without abandoning the abnormal data. The agent can distinguish the importance of samples with its learned knowledge based on previous experience. Another limitation in the learning of robots is automatic association. In Hajizada, Berggold, Iacono, Glover, and Sandamirskaya (2022), the object is allocated to a fixed position, thus, the agent does



Fig. 1. Interactive experiment scenario in the laboratory. The agent is attending to the object manipulated by the human peer.

little on automatic data sampling by itself, as suggested by Lesort et al. (2020). Auddy, Hollenstein, Saveriano, Rodríguez-Sánchez, and Piater (2023) and Ayub, De Francesco, Holthaus, Nehaniv, and Dautenhahn (2023) also both require human assistance to prepare the data. In their experiment, the object is selected by the human peer rather than by spontaneous sampling. Such limitation is overcome in our model by the consciousness mechanism assisted by attention. The architecture implemented in this work enables the agent to automatically pay attention to either visual or auditory stimuli and associate them. To summarise our continual learning scheme, we chose neural network implementation in the learning module, which is achieved by selective visual memory replay with unknown samples only buffered.

Specifically, this work develops the robotic agent with the ability to continually learn visual perceptions via interactions with the human peer. A model based on GWT is implemented, with attention and memory schemes. Based on the data replay approach, the agent is designed to learn a class-incremental object recognition task. The experiment scenario is displayed as in Fig. 1. We used a physical Pepper robot for the human-robot interaction purpose in the lab environment. The Pepper robot is a well-developed humanoid social robot, which demonstrates behaviours with a humanoid appearance. This usually makes the interaction with human peers more natural.

This work makes three novel contributions to the community of Cognitive Robotics. Firstly, this work provides a consciousness-inspired cognitive architecture integrating memory, attention, and learning capabilities for cognitive robotic research. It shows how continual learning behaviour takes place in this consciousness-based model, with the interactions between different cognitive mechanisms implemented in this model via the global workspace. Thus, this work enhances the feasibility of GWT for cognitive system development. Secondly, this work proposes a self-evaluation mechanism, which enables the agent to discriminate important information for learning. By this, we emphasise the autonomy of the agent during learning without supervision on the data selection. Thirdly, the memory consolidation mechanism is introduced into the robotic implementation, which enables the continual learning of the agent in parallel to the memory formation process. This guarantees the agent to keep interacting without interruptions for retraining of the knowledge network.

In Section 2, the theoretical basis of GWT is introduced, and a working definition of consciousness is given for this work. It is followed by Section 3, the methodology part. In this part, the authors first introduce the experiment task on the class-incremental learning task. This is accompanied by the model implementation with explained inspirations from other cognitive system works and supportive studies in neuroscience and psychology. Key concepts of the system, including




Fig. 2. Framework of Global Workspace Theory. Specialist modules run in parallel and compete to enter the Global Workspace, which would broadcast the information received to all other Specialist modules.

cognitive processes, selective memory replay, the mechanism of self-evaluation of knowledge and memory consolidation are explained. After this, in Section 4, the experiment is conducted on a robot agent Pepper (Pandey & Gelin, 2018) in a class-incremental object recognition task via interactions with the human peer. The results are analysed with observations of learning performance, knowledge development and memory growth. In Section 5, this work is concluded with novelties and contributions again, while limitations are recognised and possible solutions are proposed for future work.

2. Global workspace theory and consciousness

In this work, a cognitive architecture is implemented based on GWT, which is consistent with many studies in human cognition (e.g., Baars (2005, 2007) and Dehaene, Sergent, and Changeux (2003)). GWT was proposed to simulate the structure of the conscious brain. It was inspired by psychological and neuro-scientific research and proposed for consciousness (Baars, 1993, 2007). It is derived from the feature of the conscious mind that a person can attend to only one event out of the overwhelming information at a certain time and get this piece of information processed by all brain regions. In GWT, this is attributed to the three postulated features, namely (i) specialised processors, (ii) competition for entering consciousness, and (iii) global availability (Baars, 1993; Shanahan, 2006). This theory provides a general framework for the development of a cognitive architecture. The initial purpose of exploiting this theory in intelligence systems is to create conscious machines by understanding and modelling the consciousness correlates, such as functions and mechanisms. This theory defines the system as a multi-modal agent. A framework of a GWT-based system is demonstrated in Fig. 2.

This framework defines specialist modules operating in parallel and a global workspace for coordinating different modules of the system as a whole. These specialist modules stand for cognitive functionalities in certain modalities, such as hearing, vision, language, etc. These modules keep trying to enter the global workspace node for broadcasting signals to other modules when certain signals need to be consciously processed by the whole brain. For example, a voice signal in the background would only activate the local cognitive area centred on the hearing system in the brain, if it is not consciously perceived by the subject. On the other hand, if the voice is consciously perceived by the subject, he or she would be able to exert other cognitive capabilities on that voice signal, such as visual reference recall, source identification, etc., which need to recruit other specialist modules.

There are also other architectures proposed for cognitive systems, as mentioned in the last section, ICARUS (Langley & Choi, 2006), ACT-R (Ritter et al., 2019), and Soar (Laird et al., 1987). There is no doubt that each of them models specific aspects of human cognition. ICARUS highlights hierarchical skill learning, which is grounded in perception-action loops. ACT-R models rule-based reasoning with mathematical cognitive modelling. For the Soar system, the agent focuses on goal-directed problem-solving via the chunking paradigm. However, most of them and their variations focus on certain cognitive capabilities of the

system. However, the GWT-based framework emphasises the systematic organisation of different sub-systems rather than the importance of certain functionalities. This theory proposes the importance of information integration via “conscious” attention.

Because of the consistency of the GWT in psychology and brain science (Baars, 2017; Deco, Vidaurre, & Kringelbach, 2021; Mashour, Roelfsema, Changeux, & Dehaene, 2020), it becomes a guiding theory for cognitive computing. In this field, researchers implement the Global Workspace framework first but pay little attention to the complexities and capabilities of specialists. Some example attempts of this were made by Dehaene, Kerszberg, and Changeux (1998) and Shanahan (2006). In these works, GWT was exploited to simulate certain cognitive effects related to human consciousness. The works of this group have at least one profound implication. It validates the potential of GWT in computing systems. For instance, Shanahan (2006) combines GWT and internal model theory, enhancing the agent’s cognitive responses to the changing situation. Through this group of work, the viability of the theory for computing modelling of consciousness-related cognition is strengthened, providing later researchers with prototypes. Extensively, instead of taking GWT as the study subject, people are proposing cognitive architectures based on GWT for system development. Two significant examples in this group are LIDA (Franklin, Madl, D’mello, & Sneider, 2013) and IDyOT (Wiggins & Forth, 2014), in which the modules are organised by the formula of GWT to perform cognitive computations. By these, they put forward the work of cognitive architectures, providing concrete models for others. Based on the LIDA architecture, some implemented agents are summarised by Franklin et al. (2016). These implementations of the LIDA architecture repeat the common goal for the discussed second group work on GWT. This work falls into the extensive research community based on GWT to develop cognitive architecture.

3. Methodology

This work studies the reciprocal relationship between memory and learning for cognitive robotics based on GWT. Here, we first introduce the task to demonstrate the relationship between these two cognitive features. Then, the implementation of the cognitive architecture is detailed.

3.1. Task design

An interactive experiment with a human peer is designed. The robot is taught with the name-object pairs in sequence in the learning phase. After this, the agent is exposed to objects in random poses. The agent would articulate the name of a known object. If the shown object is not known, the agent would only visually focus on the object by moving head joints without telling the name. Fig. 1 illustrates the experiment environment. Objects used in this experiment include an apple, a banana and an orange. The experiment consists of teaching and test stages in alternation to incrementally assess the learning outcomes of the agent towards the three objects. Specifically, after the teaching test of each object, the agent is tested with all three objects to evaluate the performance. The total experiment process includes teaching and testing phases of each object in turn. During the experiment, the agent is expected to gradually learn all three object-name pairs, and make use of the learnt knowledge to distinguish important information for memory maintenance. The two phases are explained below. The implementations of involved mechanisms and models are introduced in the architecture implementation part.




Fig. 3. Information flow during teaching stage. 1 The agent perceives both stimuli from visual and hearing modules. 2 When the visual image attends the Global Workspace node, 3 the signal is then broadcast to the hearing module. 4 The associative memory of the object–name pair is formed and stored in a list of associative neural networks, 5 which are used for the consolidation of memory into knowledge.

3.1.1. Teaching phase

In the teaching stage, the agent is taught with the object–name pairs of the three objects in sequence. With the attention mechanism implemented in the model, the human peer moves the object to be taught to attract the attention of the agent. After the agent has focused on the intended object, the name is spoken to the robot. With the speech signal cached in the working memory for a short period, the agent is exposed to the same object with random poses which are manipulated by the human peer. During this process, the agent is observing the object, facilitating the generalisation by learning a limited number of variations of the object from different visual perspectives and poses. Three objects are taught to the agent by sequence. During the teaching process, the cognitive flow of the agent is illustrated in Fig. 3. To be noted, as we defined memory as information transformation, which is further explained in the subsection of the architecture, the associative memory between the image and the speech name is stored as a small neural network. These neural networks simulate the episodic memory of the hearing module. For instance, if an apple image is associated with a specific speech input, the apple image could only directly recall that associated speech input as episodic memory. Though these episodic memory neural networks are resistant to generalised knowledge, they are used to continually train the bigger neural network as her semantic memory which is shown in Fig. 3 as step 5.

3.1.2. Test phase

In the test stage, the agent is only exposed to visual perceptions of experimental objects. Given that the agent has learnt the object–name pairs in the teaching stage, the robotic agent would articulate object names. Otherwise, the agent with self-evaluation, which is explained later in this section, prevents the agent from generating random output. Thus, the agent is silent but only visually attends to the object. Fig. 4 illustrates the cognitive flow when the agent knows the perceived object. This process, different from the learning process, requires the agent to use the continually trained semantic memory network to classify the perceived object representation. If the agent knows the object, the semantic memory neural network generates a response based on the self-evaluation mechanism introduced in sub-section D.

3.2. Cognitive architecture

This paper focuses on the relationship between memory and knowledge. The interest lies especially in their modelling within a cognitive system. Though the definitions of these two concepts are not perfectly pinned down yet (Bolisani, Bratianu, Bolisani, & Bratianu, 2018; Zlotnik




Fig. 4. Information flow during the test stage. 1 The agent is exposed to an object, 2 which is encoded into a low-dimensional representation in the visual module. 3 It enters the global workspace node and gets broadcast to the hearing module, where the image–name knowledge is maintained. 4 When the hearing module receives this broadcast representation, it makes use of the knowledge network to transform it into the modal representation of the name speech of the object. 5 Then the hearing module sends the representation back to the Global Workspace node. 6 This is further forwarded to the speech module by the node, where the corresponding name is articulated by the agent.




Fig. 5. Transformation diagram of concept DOG in Semantic Point Architecture (Blouw, Solodkin, Thagard, & Eliasmith, 2016). The information transformation in the architecture hinges on the conceptual semantic pointer (SP), which is a low-dimensional representation. This SP is connected to other SP in various modalities. Through the SP in perceptual modules, the conceptual SP is associated with high-dimensional representations from the perceptual modules as well.

& Vansintjan, 2019), they are two different aspects of the cognitive system (Tulving, 2002). Following this, they are also defined in our model as two aspects of cognition. Specifically, memory is the store, maintenance, and retrieval processes of experience, and knowledge is the ability of the agent to make predictions based on the learning of experience.

Different from the previous implementation (Huang, Chella, & Cangelosi, 2023a), which implemented memory and knowledge as individual modules, this paper absorbs these two features into the information transformations between modalities. The idea is that both could be regarded as certain transformations of information between different modules, but of different aspects. For instance, the memory of the visual experience is the transformation from a speech cue to the corresponding experience episode. Also, the knowledge of the name of a visual object is the transformation from visual perception to language representation. Though they are both integrated into the transformative processes between different modules, the distinctions between these two cognitive mechanisms are still retained. Memory is only about experience, while knowledge is a generalised employment of experience. This transformative definition is inspired by the work of Semantic Pointer Architecture (SPA) (Blouw et al., 2016; Eliasmith, 2013). The key idea of SPA is the transformation between different modalities and




Fig. 6. The structure of the model. In this model, all modules are bidirectionally connected with the global workspace node. By this, it emphasises that the information transformation between modules is bi-directional. Please refer to Huang, Chella, and Cangelosi (2023b) for more information on the inter-module competition and the role of the Global Workspace node.

levels of cognition. For instance, In Fig. 5, the concept “DOG” could be transformed into other concepts in semantics and sensorimotor representations, which represent the recalling of images, voices, feelings and so on of the agent. On the contrary, sensorimotor representations can be transformed into semantic representations. Thus, three kinds of transformation could be summarised in this architecture, namely compression, decompression and within-semantics. These three processes capture most (if not all) information processing happening in brains. Amongst them, for our interest, the compression and decompression transformations are inspiring as they are quite compatible with our definitions of memory and knowledge, though with differences in implementation from SPA.

As inspired by SPA, in our new model, modules communicate with others via the Global Workspace by transformations. In this case, memory and knowledge are implemented as transformations between different modules, rather than individual modules. In Fig. 6, we focus on the interactions between the modules vision, hearing, motor and speech with the coordination of the Global Workspace. In the hearing module, the learning scheme is integrated for associating consciously perceived visual information to the language inputs. Motor and speech modules are implemented for the robotic agent to engage in physical interactions with human peers. Thus, there are no transformation processes implemented within these two specialist modules yet. Similarly, the vision part is implemented as a receptor and pre-processor of visual stimuli, without a learning process. The focus of our new model is on the learning of associations from visual stimuli to auditory ones. These four modules are explained in detail below.

Vision module receives the perceived image as the input and the output as an encoded representation to compete for entering the Global Workspace. The detailed process is demonstrated in Fig. 7

When an image is perceived, it goes through a detector network Yolov8, resulting in the detected objects. As illustrated in Fig. 7, the array data of an object consists of not only the sub-image but also the position information. The position information is represented by two channels appended to the RGB channels of the image. Also, for computational convenience, the data arrays are reshaped into [64, 64, channelNum]. Thus, each array of data is of [64, 64, 5]. When there is more than one object, these objects would compete with each other based on their salience values. For stimuli from the visual field, the salience values are determined by (1) the difference of pixel values between the sub-image of the certain object *Obj* and its local background *Local*, which is defined as the square area surrounding the object. The size of the area is three times the sub-image in both terms of height and width, as illustrated in Fig. 7. The calculation is defined by Eq. (1); (2) relative sensitivity of retina towards the *Area* of an object, which is illustrated in Fig. 7. The calculation is defined by Eq. (2); (3) relative change of visual field, determined by the difference between current visual episode *Current* and the prior one *Prior* within the *Area* of an object in consecutive perceptions. The calculation is defined by Eq. (3); and (4) context familiarity which is defined as the difference between the perceived object *Obj* and objects in the working memory buffer

BufferedObjs which were perceived by the agent before the current perception. The calculation is defined by Eq. (4).

With all the defined factors, The calculation of salience values for the detected objects is defined by (5).

$$PS = |\sum Local - 2 * \sum Obj| \quad (1)$$

$$RS = \sum Area \quad (2)$$

$$ChS = \sum_{Area} |Current - Prior| \quad (3)$$

$$CtS = \frac{Obj - BufferedObjs_n}{|Obj - BufferedObjs_n|_0} \quad (4)$$

$$Salience = \frac{PS + RS + ChS + CtS}{MaxPixelSalience} \quad (5)$$

where *MaxPixelSalience* is held in *Vision*. It is updated every time if *Salience* is greater than 1.0. As shown in the experiment part, this visual attention mechanism enables the robotic agent to engage with different stimuli in the visual field in a natural way. This is a key factor for the agent to interact with the human peer.

After the calculation of salience values, the object signal with the highest salience wins the within-modality competition. This winner signal is visually attended by the agent, resulting in a representation of that visual stimulus. This is simulating the signal propagation from the retina to the visual cortex. The encoding model is a pre-trained ResNet50, which is directly loaded from the Pytorch library. The activation of the hidden layer of the pre-trained model is taken as the encoded representation. Then, the representation competes with signals from other modalities to enter the Global Workspace for requesting conscious responses across the whole brain.

Hearing is implemented with a learning scheme for associating visual stimuli to language inputs. The processes are illustrated in Fig. 8.

The two processes on the left side are information flow from the external world and the Global Workspace, respectively. The upper process receives language input from the human peer, resulting in an auditory representation of the stimulus. This is assisted by the Whisper model, which is a speech recogniser developed by OpenAI (Radford et al., 2022). The auditory representation would be buffered in the working memory of the hearing module if that signal attends the Global Workspace node, becoming conscious of the agent. The bottom process is broadcast from the Global Workspace node to the hearing module, therefore the broadcast signals are taken as conscious perceptions in our assumption. In this work, we focus on the visual representation. When the hearing module tries to associate the received signal from the Global Workspace node with the buffered signal in its working memory. Notably, the association that happens here is regarded as memory. This memory association has two features. Firstly, both representations are of real experiences from the external world. Secondly, this association has no generalisability, which can only map from one to the other. This memory association forms into short- or long-term memory based on the strength of conscious activity as (Huang et al., 2023a) discussed. In this model, the key factor for this purpose is if the agent knows the association before learning the pair of representations. With all the memory associations stored, the knowledge consolidation process uses them as the source of data replay for continual learning. To sum up, the hearing module has a Whisper model for generating auditory representation, a memory assembly consisting of associations as the memory pool, and a model for transforming broadcast representations from the Global Workspace node to auditory ones, which is updated continually.

Motor functions as an effector for the agent. It has two roles in the interactions with the human peer. Firstly, the motor module coordinates the angles of the robot body based on the position of the attended object in the environment. The movement of the robot head results in a change of visual perception, affecting the salience




Fig. 7. Information processing in vision module. The saliency calculation of objects involves four factors. The input is the image of the visual field, and the output is a low-dimensional representation of the most salient object.




Fig. 8. Processing of information from both the external world and the global workspace in hearing module.

distribution of different areas and objects in the visual field, especially the retina saliency as explained in the vision module. This is a key factor for the dynamics of visual attention. In addition, the motor module is a strong interactive indicator to the human peer, which is important within human–robot interactions.

Speech is implemented as an interface similar to the motor module. When it receives auditory representation from the Global Workspace node, a speech would be articulated by the robot. If there is a visual representation attended by the agent, it results in a predicative response in the hearing module. This is realised by the associative learning discussed above in the hearing module description. The response is also a candidate to enter the Global Workspace node. When it accesses the Global Workspace node and gets broadcast to the speech module, a language output is generated by the agent. This speech response indicates if the agent knows the name of the visual stimulus.

3.3. Learning process in cognitive systems

When it comes to cognitive systems, the continual learning problem goes beyond only techniques for updating the networks. In the normal setting of continual learning, the algorithm or the trained network is fed with pre-defined inputs and target data. On the other hand, with an embodied agent, this is not taken for granted. In this study, the learning process is recognised in two stages. The agent has to identify the knowledge during interactions with human peers first and then convert that into its own. Thus, the continual learning process consists of multiple cognitive processes in systematic implementation instead of a single network model.

For the first step, the agent needs to visually attend to the object which the human peer refers to. This is realised by the visual attention mechanism as discussed in the description of the vision module. Then, the agent hears the language input from the human peer of the object name. With the object name buffered in the working memory of the hearing module, as illustrated in Fig. 8, the memory association is then formed. After this, the second step of learning is implemented as knowledge consolidation. The interactive environment is demonstrated in Fig. 1. During this interactive experiment, the key to success is the attention mechanism implemented in the robotic agent. By the behaviour triggered by the attention mechanism, the human peer could distinguish if the taught object is attended to by the robot properly. Thus, in this teacher–student interaction, not only the robot perceives stimuli guided by the human teacher, but also the human peer needs to evaluate the cognitive behaviours of the robot to achieve the interaction success.

3.4. Self-evaluation based selective replay

3.4.1. Selective replay

Data replay is one of the common approaches for continual learning, as rehearsed data could simply maintain the model's knowledge gained from previous data (Lange et al., 2021). This is inspired by the neuroscientific evidence in human brains. Studies (e.g., Oudiette and Paller (2013) and Van de Ven, Trouche, McNamara, Allen, and Dupret (2016)) have proven the importance of neuronal reactivation in the process of memory consolidation. However, the constraint on the growing body of data storage is the main challenge for this approach in AI research. Relevant works on selective replay (McClelland, McNaughton,

& Lampinen, 2020; Schapiro, McDevitt, Rogers, Mednick, & Norman, 2018) inspire this work on this issue. In this paper, we exploit the self-evaluation of the agent to store only “hard” memory. While emphasising the importance of a subset of previous data is not a novel idea (e.g., Liu et al. (2022)), we implement the agent with self-evaluation to tell if the taught knowledge is novel and important to learn. For this, we assume that the learner should pay more attention to unknown knowledge, which contributes to the knowledge gain. For instance, after the agent learns the object–name pair of the orange, it would pay less effort to learn new associations for the orange name and images. On the other hand, the pair of the apple and the name that it has not learnt will be remembered better (long-term memory). To formalise, the agent maintains a memory buffer according to **Algorithm 1**.

Algorithm 1 Update Memory Buffer with Hard Samples

```

1: Initialize:
2:   Memory buffer  $B \leftarrow \emptyset$ 
3:   Knowledge model  $N$ 
4:   Input set  $(X, Y) = \{(x_1, y_1), \dots, (x_n, y_n)\}$ 

5: for each  $(x_i, y_i) \in (X, Y)$  do
6:   if  $N(x_i) \neq y_i$  then
7:      $B \leftarrow B \cup \{(x_i, y_i)\}$  ▷ Add misclassified sample
8:     Train  $N$  with  $B$ 
9:   else
10:    Continue ▷ Skip correct predictions
11:   end if
12: end for

```

Before implementing this, two questions have to be answered. The first is whether remembering only unknown knowledge items is enough for good learning performance. The second is how the agent could distinguish a novel stimulus, and if the knowledge is unknown.

To validate the idea of replaying only the unknown data to the agent, we conducted comparative studies on four different datasets. They were MNIST (Deng, 2012), a Pepper dataset collected from the robot, iCubWorld 1.0 (Fanello et al., 2013) and the CIFAR 10 (Krizhevsky, Hinton, et al., 2009). All four datasets contain 10 classes of objects. The MNIST dataset contains 60 000 and 10 000 samples for the train and test sets. The Pepper dataset has 477 and 53 samples for the train and test sets. The iCubWorld-10 dataset has 6000 and 1000 samples for the train and test sets. The CIFAR-10 dataset contains 50 000 and 10 000 samples for the train and test sets. Table 1 shows the performance on the MNIST dataset with 1000 training samples and the rest for the test. As shown by the results, continual learning with hard memory replay only needed training on a small subset of the whole training dataset, but achieved close performance to batch learning on the whole dataset across all four datasets. The performance on Pepper data and CIFAR-10 was lower than the other two. This was because the Pepper data collected is limited in samples, and the CIFAR-10 dataset is much more complex in image quality. However, the performance on the CIFAR-10 dataset is comparable to the normal ResNet18 architecture. These results show that replaying only unknown samples to the agent is a feasible idea.

Notably, the buffer size of a complex dataset like CIFAR-10 would still suffer from the challenge of memory management. This indicates the importance of advanced memory management schemes. Other approaches like regularisation (Kirkpatrick et al., 2017) and parameter-isolation (Mallya & Lazebnik, 2018) would be incorporated in our future work with complex data structures.

3.4.2. Self-evaluation

The idea of the ‘hard’ memory replay is validated in the last subsection, enabling the agent the ability to tell if the taught knowledge is known as the second technical problem. This problem involves two

Table 1

Comparison of learning performance with hard-memory-replay continual learning against batch training.

	Batch training	Continual learning	
	Test accuracy	Test accuracy	Buffer size
MNIST	99.0	97.6	1873/60 000
Pepper	85.4	83.3	98/477
iCubWorld	98.9	96.0	915/6000
CIFAR-10	72.9	74.5	16 802/50 000




Fig. 9. Encoder–Decoder–Encoder model. By training this model with samples, the model learns both patterns mapping from input to output and vice versa.

situations. The first one is the easy one, which could be told by simply comparing the target in the associative pair against the classification result of the model. However, the second one is harder. Generative Adversarial Networks (GANs) (Goodfellow et al., 2020) are proposed for this purpose. Conditional GAN (Smith & Smith, 2020) is specifically for the multi-class problem, while the training of GANs is very tricky. One needs to maintain a balance between the generator and the discriminator. In this work, we make use of an encoder–decoder–encoder network for this purpose. Instead of training a discriminator, we train an extra decoder. This decoder guarantees that the agent learns the distribution of each class of objects in the input space. The benefit is that there is no competition between the two sub-networks, but they share the same goal. This saves the effort to keep the tricky balance between them. The model is demonstrated in Fig. 9.

Given the input, it goes through encoding, decoding and another encoding. If the input is known to the agent, the results of the two encoding steps would be the same, otherwise, the input is a new stimulus to the agent. With the **Input** and **Target** given, the loss function for this purpose is as below.

$$Loss = \sum (Input - Input')^2 + \sum (Output' - Target)^2 \quad (6)$$

In Eq. (6), the encoder tries to predict the result for the input, while the decoder tries to recover the input data. In this case, only if both the encoder and decoder are well-trained, the results of both encoding steps would be the same. Notably, the result of the second encoding is not used in the back-propagation. This is because of its impact on the training. With **Output''**, biases of three steps would contribute to the loss equally and independently. However, before the model is well trained, **Input'** is not correct, which would feed the wrong data to the model. This is detrimental to learning. On the other hand, **Output''** does not require learning as it is determined by the first two steps. This means **Output''** would be correct as long as the prior encoding and decoding function well. With this encoder–decoder–encoder model, the agent can tell if the taught data is important during interactions without the intervention of the human peer.

3.5. Consolidation mechanism

Learning in cognitive systems is another story from that in common AI research. In common AI research, the model is updated anytime




Fig. 10. Memory growth of experiment 1. This corresponds to the learning sequence in Table 2. The x-axis is the time step, y-axis is the number of associative name-object pairs in episodic memory.




Fig. 12. Memory growth of experiment 3. This corresponds to the learning sequence in Table 20. The x-axis is the time step, y-axis is the number of associative name-object pairs in episodic memory.




Fig. 11. Memory growth of experiment 2. This corresponds to the learning sequence in Table 11. The x-axis is the time step, y-axis is the number of associative name-object pairs in episodic memory.

there is new data arrives. However, the cognitive system should not be interrupted at any time the data is available. Instead, experienced data is consolidated in the back of online life. The psychological study on the role of sleeping in memory (e.g., [Stickgold \(2005\)](#)) is a good example of the theory of memory consolidation.

As illustrated in Fig. 8, we implement the learning process into two steps. The first step is the formation of memory associations. Because of this, the knowledge model is not updated immediately when the data arrives. On the contrary, the memory associations compose a memory pool. With certain intervals during the cognitive life of the agent, the knowledge model uses the memory pool as the replay source for continually updating its weights to learn.

4. Experiment results

In this part, the results of three repetitions of the interactive task are presented and discussed. For each repetitive experiment, the memory growth of the agent through the task is illustrated (Figs. 10 11 12). In each memory growth plot, the orange line is the growth of long-term memory when the agent does not know the taught object-name pair. The blue dash line is the growth of memory, including both short-term and long-term. Peaks of the green dash mark the regular completion of memory consolidation, which takes all memory at the point. The first table of each experiment result is the learning sequence of three objects, with the corresponding labels of long-term (L) and short-term (S) memory of different samples (Tables 2 11 20). To evaluate the learning performance of the agent in the task, it is tested in 4 stages as explained in the experiment design after each teaching phase. For each stage, there is a table of classification performance. Another table shows

the perceived objects with classified labels, in which the samples and articulated labels (“A” for Apple, “O” for Orange, “B” for Banana, and “-” for Unknown) are given. The knowledge of the agent is measured for the three objects with correctness, incorrectness and unknown rate, the sum of which for a specific object is 1. Correctness measures that the agent articulates the correct name of the shown object, incorrectness measures that the agent articulates the wrong name of the shown object and unknown rate measures that the agent does not articulate anything for the shown object. To avoid coincidental results, the objects are learnt in different orders in the three repetitive experiments. These results are discussed in the next section.

4.1. Experiment 1

The first experiment is teaching the robot three objects in the sequence of [Apple, Orange, Banana, Apple]. [Table 2](#) shows the sequence of learning, and the classification performance was evaluated after learning each category of objects (see [Tables 3–10](#)).

Table 2
Learning sequence of experiment 1, with four teaching phases of the apple, orange, banana and apple again in order.

[illegible]

Table 3
Test performance of the agent for the three objects after learning only the apple.

	Apple	Orange	Banana	Total
Correct	100%	0%	0%	33.3%
Incorrect	0	100%	100%	66.7%
Unknown	0%	0%	0%	0%

Table 4
Classification labels of the agent for samples of the three objects after learning only the apple.

[illegible]

Table 5

Test performance of the agent for the three objects after learning the apple and orange.

	Apple	Orange	Banana	Total Accuracy
Correct	0%	80%	0%	26.7%
Incorrect	40%	20%	60%	40%
Unknown	60%	0%	40%	33.3%

Table 6

Classification labels of the agent for samples of the three objects after learning the apple and orange.

Perception																		
Response	0	-	0	-	-	0	A	0	0	0	0	-	0	0	0	0	-	

Table 7

Test performance of the agent for the three objects after learning all for once.

	Apple	Orange	Banana	Total
Correct	0%	80%	80%	53.3%
Incorrect	40%	0%	0%	13.3%
Unknown	60%	20%	20%	33.3%

Table 8

Classification labels of the agent for samples of the three objects after learning all for once.

Perception																																																																																																																																	
------------	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	--	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	--

Table 9

Test performance of the agent for the three objects after learning all and reviewing the apple.

	Apple	Orange	Banana	Total
Correct	80%	80%	60%	73.3%
Incorrect	0%	0%	0%	0%
Unknown	20%	20%	40%	26.7%

Table 10

Classification labels of the agent for samples of the three objects after learning all and reviewing the apple.

Perception															
Response	B	B	-	-	B	O	O	-	O	O	A	A	A	A	-

4.2. Experiment 2

The second experiment is teaching the robot three objects in the sequence of [Orange, Banana, Apple, Orange]. Table 11 shows the sequence of learning, and the classification performance was evaluated after learning each category of objects (see Tables 12–19).

Table 11

Learning sequence of experiment 2, with four teaching phases of the orange, banana, apple and orange again in order.

[illegible]

Table 12

Test performance of the agent for the three objects after learning only the orange.

	Apple	Orange	Banana	Total
Correct	0%	100%	0%	33.3%
Incorrect	100%	0%	100%	66.7%
Unknown	0%	0%	0%	0%

Table 13

Classification labels of the agent for samples of the three objects after learning only the orange.

[illegible]

Table 14

Test performance of the agent for the three objects after learning the orange and banana.

	Apple	Orange	Banana	Total
Correct	0%	80%	100%	60%
Incorrect	20%	0%	0%	6.7%
Unknown	80%	20%	0%	33.3%

Table 15

Classification labels of the agent for samples of the three objects after learning the orange and banana.

Perception																																																																																																																													
------------	---	---	---	--	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	--

Table 16

Test performance of the agent for the three objects after learning all for once.

	Apple	Orange	Banana	Total
Correct	80%	0%	100%	60%
Incorrect	0%	20%	0%	6.7%
Unknown	20%	80%	0%	33.3%

Table 17

Classification labels of the agent for samples of the three objects after learning all for once.

																		
Response	A	A	-	A	A	-	-	-	-	A	B	B	B	B	B	B	B	B

Table 26

Classification labels of the agent for samples of the three objects after learning all for once.

Response	-	-	A	A	-	O	O	O	O	O	-	-	-	-	-	-

Table 27

Test performance of the agent for the three objects after learning all and reviewing the banana.

	Apple	Orange	Banana	Total
Correct	80%	80%	100%	86.7%
Incorrect	0%	0%	0%	0%
Unknown	20%	20%	0%	13.3%

Table 28

Classification labels of the agent for samples of the three objects after learning all and reviewing the banana.

Response	A	-	A	A	A	O	O	O	O	O	B	B	B	B	B	B

4.4. Result discussion

In the three repetitive experiments, objects are taught to the agent in different orders to avoid possible biases. The order is [apple, orange, banana, apple] for experiment 1, [orange, banana, apple, orange] for experiment 2, and [banana, apple, orange, banana] for experiment 3. For each teaching phase, the agent observed a specific object with 9 different sample positions. For each test phase, the agent is tested with 5 different sample positions of each object.

The memory growth curves and the memory labels of learning sequences show that the long-term memory of the agent throughout the interactions is only a subset of all perceptions. While short-term memory would fade away, long-term memory is retained for learning continually. As demonstrated in Tables 9 10 18 19 27 28 of the final learning performances, the learning outcome is satisfactory, which is consistent with the results of validation experiments. By analysing the results, long-term memory is remarkably growing when there is a new object taught to the agent. However, when the memory is consolidated and learnt as the knowledge of the agent, the storage of long-term associative memory is occasionally increased for an object. Across all three experiments, the memory size barely goes up with only very few new samples after certain objects are learnt. Generally speaking, the agent remembers novel samples for specific objects and pays less effort to familiar perceptions.

As noted, during the experiment, not all objects could be properly learnt by the agent via only one stage of teaching. An interesting finding is that the first object taught in the experiment always requires a review for the agent to learn. The experiment shows that the agent does not learn the first object properly after teaching in the sequence of three objects, though it can properly discriminate different objects. After a review of the first object, the agent eventually learns all three objects with names well. This is attributed to the continual learning scheme adopted. As mentioned above, when only the first object is taught, the agent is overfitting so that almost all perceived objects are recognised as the same. In this case, the agent is over-confident, so it does not

maintain proper memory based on whether it knows the taught object–name pair. The key factor behind this is the balance of memory data, which is a big issue for most AI research (Batista, Prati, & Monard, 2004).

Apart from the above, the difference between knowing and discriminating objects is recognised due to the self-evaluation mechanism. Given an object, the agent can discriminate objects by not only telling their names respectively but also knowing the ignorance of self. In this case, though the agent does not know the name of a certain object, it knows the perceived object is not any of the learned ones. During the learning process, the agent shows an increasing ability to discriminate different stimuli. As demonstrated in the results, the unknown rate for the three objects at the beginning is always 0% due to the over-confidence of the agent in her knowledge of the first taught object. At this time, when the agent is exposed to other objects, the name of the first object would be articulated most of the time. This means that the agent has very limited capability to discriminate objects due to the extremely imbalanced available data. Afterwards, the agent turns to realise her ignorance of different samples with an increased unknown rate. For instance in experiment 1 and experiment 2, after the second object–name pair is taught, the agent can disentangle different stimuli better. Thus, if the agent is then exposed to the third object, the self-evaluation enabled by the encoder–decoder–encoder network of the agent helps discriminate known and unknown stimuli. It shows that the agent gives no speech response to objects it does not recognise. Finally, the unknown rate drops to a low level because of the learning successes of the three objects. With the last object learnt successfully, the agent can discriminate all three objects and correctly articulate their names.

In summary, the agent prioritises different information during learning with the self-evaluation mechanism, which facilitates memory maintenance. In return, those relatively important memory contributes to the continual learning process of the agent for the class-incremental recognition task. In addition, with the ability to distinguish if an object is learnt based on the self-evaluation mechanism, the robot knows whether it knows the name of the perceived object or not, instead of giving a guess. With this, the results show that the agent demonstrated over-confidence at the beginning and eventually showed its awareness of its knowledge after the whole learning process.

5. Conclusion and future work

In this paper, a cognitive architecture based on the GWT is presented. The cognitive agent is implemented with a continual learning scheme based on selective memory replay. The experiment shows the reciprocal relationship between memory and learning activities based on the mechanism of self-evaluation of knowledge. In addition, the agent is implemented with a memory consolidation mechanism, which enables the agent to learn during interactions without interruptions. Compared to the previous implementation of the architecture in Huang et al. (2023a), here memory and knowledge are designed as information transformations between modalities rather than individual modules following the information system design of SPA (Eliasmith, 2013). Secondly, a more sophisticated visual attention mechanism is implemented, resulting in a more natural behaviour of attention for the agent. This facilitates the interactions between the robotic agent and the human peer.

To summarise, the novelties of this work are recognised from three aspects, which contribute to the community of cognitive robotics.

Firstly, this work introduces a cognitive architecture based on consciousness theory. The architecture integrates different cognitive features such as memory, attention, and learning. This work would contribute to cognitive architecture development for cognitive robotics. With the interactions between different cognitive mechanisms implemented in this model via the global workspace, this work enhances the feasibility of GWT for cognitive system development. Secondly, this work studies the reciprocal relationship between memory and learning

processes. The class-incremental learning of the agent is supported by the selective memory replay approach, and in return, the memory maintenance is facilitated by the learnt knowledge with memory. To achieve this, a self-evaluation mechanism is proposed, which enables the agent to discriminate important information for learning. By this, we emphasise the autonomy of the agent during learning without supervision on the data selection. Thirdly, the memory consolidation mechanism is introduced into the robotic implementation, which enables the continual learning of the agent in parallel to the memory formation process. This guarantees the agent to keep interacting without interruptions for retraining of the knowledge network.

However, the limitations of this work are also recognised for future development. Firstly, the design of memory as transformations between different sensorimotor perceptions deprives the temporal patterns of memory. This is a violation of the definition of episodic memory adopted in our previous implementation (Huang et al., 2023a). We make the hypothesis that temporal patterns of episodic memory are the transformations between different perceptual episodes within a modality like hearing or vision. This is an important idea for our future development of the system. Secondly, the memory constraint realised by the hard memory replay approach in this work should not be the best strategy. More continual learning approaches could be explored in the future to refine the implementation. Thirdly, the level of intelligence of the implemented agent is still very low. The cognitive cycles are simple, with the global workspace coordinating the communication between different modules in a naive way. To convince people with the intelligent system, the agent must demonstrate complex communications between different modules (Huang, Cangelosi and Chella, 2023). This requires us to design more complicated experiments, in which the agent would be involved in harder tasks demanding complex collaborations of modules. The learning scheme implemented in this work is only one aspect of the continual learning issue for artificial systems, focusing on class incremental learning. However, an autonomous system is expected to learn different tasks continually. This requires the designer to consider how different task functionalities are organised at the system level.

CRedit authorship contribution statement

Wenjie Huang: Writing – review & editing, Methodology, Formal analysis, Conceptualization, Writing – original draft, Investigation, Data curation. **Antonio Chella:** Writing – review & editing, Methodology, Supervision, Conceptualization. **Angelo Cangelosi:** Writing – review & editing, Methodology, Supervision, Investigation.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Wenjie Huang reports financial support was provided by Innovate UK. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was funded by the Innovate UK (grant number 10089807 for the European Union's Horizon Europe research and innovation programme under Grant Agreement n° 101120727 – PRIMI).

Data availability

Data will be made available on request.

References

- Auddy, S., Hollenstein, J., Saveriano, M., Rodríguez-Sánchez, A., & Piater, J. (2023). Continual learning from demonstration of robotics skills. *Robotics and Autonomous Systems*, 165, Article 104427.
- Ayub, A., De Francesco, Z., Holthaus, P., Nehaniv, C. L., & Dautenhahn, K. (2023). Continual learning through human-robot interaction–human perceptions of a continual learning robot in repeated interactions. *arXiv preprint arXiv:2305.16332*.
- Baars, B. J. (1993). *A cognitive theory of consciousness*. Cambridge University Press.
- Baars, B. J. (2005). Global workspace theory of consciousness: toward a cognitive neuroscience of human experience. *Progress in Brain Research*, 150, 45–53.
- Baars, B. J. (2007). The global workspace theory of consciousness. In *The Blackwell companion to consciousness* (pp. 236–246). Wiley Online Library.
- Baars, B. J. (2017). The global workspace theory of consciousness: Predictions and results. In *The Blackwell companion to consciousness* (pp. 227–242). Wiley Online Library.
- Batista, G. E., Prati, R. C., & Monard, M. C. (2004). A study of the behavior of several methods for balancing machine learning training data. *ACM SIGKDD Explorations Newsletter*, 6(1), 20–29.
- Berto, L., Costa, P., Simões, A., Gudwin, R., & Colombini, E. (2024). A motivational-based learning model for mobile robots. *Cognitive Systems Research*, 88, Article 101278.
- Blouw, P., Solodkin, E., Thagard, P., & Eliasmith, C. (2016). Concepts as semantic pointers: A framework and computational model. *Cognitive Science*, 40(5), 1128–1162.
- Bolisani, E., Bratianu, C., Bolisani, E., & Bratianu, C. (2018). The elusive definition of knowledge. In *Emergent knowledge strategies: Strategic thinking in knowledge management* (pp. 1–22). Springer.
- Churamani, N., Kalkan, S., & Gunes, H. (2020). Continual learning for affective robotics: Why, what and how? In *2020 29th IEEE international conference on robot and human interactive communication RO-MAN*, (pp. 425–431). IEEE.
- Costa, G., Pinho, J., Botto, M. A., & Lima, P. U. (2023). Online learning of MPC for autonomous racing. *Robotics and Autonomous Systems*, Article 104469.
- De Lange, M., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., et al. (2021). A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(7), 3366–3385.
- DeChant, C., Akinola, I., & Bauer, D. (2023). Learning to summarize and answer questions about a virtual robot's past actions. *arXiv preprint arXiv:2306.09922*.
- Deco, G., Vidaurre, D., & Kringelbach, M. L. (2021). Revisiting the global workspace orchestrating the hierarchical organization of the human brain. *Nature Human Behaviour*, 5(4), 497–511.
- Dehaene, S., Kerszberg, M., & Changeux, J.-P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences*, 95(24), 14529–14534.
- Dehaene, S., Sergent, C., & Changeux, J. P. (2003). A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences*, 100(14), 8520–8525.
- Deng, L. (2012). The mnist database of handwritten digit images for machine learning research [best of the web]. *IEEE Signal Processing Magazine*, 29(6), 141–142.
- Eliasmith, C. (2013). *How to build a brain: A neural architecture for biological cognition*. Oxford University Press.
- Fanello, S., Ciliberto, C., Santoro, M., Natale, L., Metta, G., Rosasco, L., et al. (2013). Icube world: Friendly robots help building good vision data-sets. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops* (pp. 700–705).
- Franklin, S., Madl, T., D'mello, S., & Snider, J. (2013). LIDA: A systems-level architecture for cognition, emotion, and learning. *IEEE Transactions on Autonomous Mental Development*, 6(1), 19–41.
- Franklin, S., Madl, T., Strain, S., Faghihi, U., Dong, D., Kugele, S., et al. (2016). A LIDA cognitive model tutorial. *Biologically Inspired Cognitive Architectures*, 16, 105–130.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2020). Generative adversarial networks. *Communications of the ACM*, 63(11), 139–144.
- Hajizada, E., Berggold, P., Iacono, M., Glover, A., & Sandamirskaya, Y. (2022). Interactive continual learning for robots: a neuromorphic approach. In *Proceedings of the international conference on neuromorphic systems 2022* (pp. 1–10).
- Hayes, T. L., Cahill, N. D., & Kanan, C. (2019). Memory efficient experience replay for streaming learning. In *2019 international conference on robotics and automation* (pp. 9769–9776). IEEE.
- Hayes, T. L., Kafle, K., Shrestha, R., Acharya, M., & Kanan, C. (2020). Remind your neural network to prevent catastrophic forgetting. In *European conference on computer vision* (pp. 466–483). Springer.
- Hayes, T. L., & Kanan, C. (2021). Selective replay enhances learning in online continual analogical reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 3502–3512).
- Huang, W., Cangelosi, A., & Chella, A. (2023). An incremental cognitive architecture of consciousness with global workspace theory. Available at SSRN 4462597.
- Huang, W., Chella, A., & Cangelosi, A. (2023a). A cognitive robotics implementation of global workspace theory for episodic memory interaction with consciousness. *IEEE Transactions on Cognitive and Developmental Systems*.

- Huang, W., Chella, A., & Cangelosi, A. (2023b). A cognitive robotics implementation of global workspace theory for episodic memory interaction with consciousness. *IEEE Transactions on Cognitive and Developmental Systems*, 16(1), 266–283.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521–3526.
- Krizhevsky, A., Hinton, G., et al. (2009). Learning multiple layers of features from tiny images.
- Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence*, 33(1), 1–64.
- Langley, P., & Choi, D. (2006). A unified cognitive architecture for physical agents. Vol. 21, In *Proceedings of the national conference on artificial intelligence* (p. 1469). Menlo Park, CA; Cambridge, MA; London: AAAI Press; MIT Press, no. 2. 1999.
- Lesort, T., Lomonaco, V., Stoian, A., Maltoni, D., Filliat, D., & Díaz-Rodríguez, N. (2020). Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges. *Information Fusion*, 58, 52–68.
- Liu, H., Gu, L., Chi, Z., Wang, Y., Yu, Y., Chen, J., et al. (2022). Few-shot class-incremental learning via entropy-regularized data-free replay. In *Computer vision—ECCV 2022: 17th European conference, Tel Aviv, Israel, October 23–27, 2022, proceedings, part XXIV* (pp. 146–162). Springer Nature Switzerland.
- Lopez-Paz, D., & Ranzato, M. A. (2017). Gradient episodic memory for continual learning. *Advances in Neural Information Processing Systems*, 30.
- Lorents, A., Ruitenberg, M., & Schomaker, J. (2023). Novelty-induced memory boosts in humans: The when and how. *Heliyon*.
- MacLellan, C. J., Harpstead, E., Alevan, V., & Koedinger, K. R. (2024). Trestle: a model of concept formation in structured domains. arXiv preprint arXiv:2410.10588.
- Mallya, A., & Lazebnik, S. (2018). Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 7765–7773).
- Mashour, G. A., Roelfsema, P., Changeux, J.-P., & Dehaene, S. (2020). Conscious processing and the global neuronal workspace hypothesis. *Neuron*, 105(5), 776–798.
- McClelland, J. L., McNaughton, B. L., & Lampinen, A. K. (2020). Integration of new information in memory: new insights from a complementary learning systems perspective. *Philosophical Transactions of the Royal Society B*, 375(1799), Article 20190637.
- Morrison, D., Corke, P., & Leitner, J. (2020). Learning robust, real-time, reactive robotic grasping. *The International Journal of Robotics Research*, 39(2–3), 183–201.
- Oudiette, D., & Paller, K. A. (2013). Upgrading the sleeping brain with targeted memory reactivation. *Trends in Cognitive Sciences*, 17(3), 142–149.
- Pandey, A. K., & Gelin, R. (2018). A mass-produced sociable humanoid robot: Pepper: The first machine of its kind. *IEEE Robotics & Automation Magazine*, 25(3), 40–48.
- Pei, W. (2008). What do you mean by “AI”? *Artificial General Intelligence*, 362–373.
- Qian, Z., You, M., Zhou, H., Xu, X., & He, B. (2023). Robot learning from human demonstrations with inconsistent contexts. *Robotics and Autonomous Systems*, 166, Article 104466.
- Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., & Sutskever, I. (2022). Robust speech recognition via large-scale weak supervision. arXiv preprint arXiv:2212.04356.
- Rebuffi, S. A., Kolesnikov, A., Sperl, G., & Lampert, C. H. (2017). Icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2001–2010).
- Ritter, F. E., Tehranchi, F., & Oury, J. D. (2019). ACT-R: A cognitive architecture for modeling cognition. *Wiley Interdisciplinary Reviews: Cognitive Science*, 10(3), Article e1488.
- Rusu, A. A., Rabinowitz, N. C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., et al. (2016). Progressive neural networks. arXiv preprint arXiv:1606.04671.
- Samsonovich, A. V. (2012). On a roadmap for the BICA challenge. *Biologically Inspired Cognitive Architectures*, 1, 100–107.
- Schapiro, A. C., McDevitt, E. A., Rogers, T. T., Mednick, S. C., & Norman, K. A. (2018). Human hippocampal replay during rest prioritizes weakly learned information and predicts memory performance. *Nature Communications*, 9(1), 3920.
- Schwarz, J., Czarnecki, W., Luketina, J., Grabska-Barwinska, A., Teh, Y. W., Pascanu, R., et al. (2018). Progress & compress: A scalable framework for continual learning. In *International conference on machine learning* (pp. 4528–4537).
- Serra, J., Suris, D., Miron, M., & Karatzoglou, A. (2018). Overcoming catastrophic forgetting with hard attention to the task. In *International conference on machine learning* (pp. 4548–4557).
- Shanahan, M. (2006). A cognitive architecture that combines internal simulation with a global workspace. *Consciousness and Cognition*, 15(2), 433–449.
- Smith, K. E., & Smith, A. O. (2020). Conditional GAN for timeseries generation. arXiv preprint arXiv:2006.16477.
- Stickgold, R. (2005). Sleep-dependent memory consolidation. *Nature*, 437(7063), 1272–1278.
- Sukhobokov, A., Belousov, E., Gromozdov, D., Zenger, A., & Popov, I. (2024). A universal knowledge model and cognitive architecture for prototyping AGI. arXiv preprint arXiv:2401.06256.
- Tulving, E. (2002). Episodic memory: From mind to brain. *Annual Review of Psychology*, 53(1), 1–25.
- Van de Ven, G. M., & Tolias, A. S. (2019). Three scenarios for continual learning. arXiv preprint arXiv:1904.07734.
- Van de Ven, G. M., Trouche, S., McNamara, C. G., Allen, K., & Dupret, D. (2016). Hippocampal offline reactivation consolidates recently formed cell assembly patterns during sharp wave-ripples. *Neuron*, 92(5), 968–974.
- Wiggins, G. A., & Forth, J. (2014). IDyOT: a computational theory of creativity as everyday reasoning from learned information. In *Computational creativity research: towards creative machines* (pp. 127–148). Springer.
- Zhong, J., Peniak, M., Tani, J., Ogata, T., & Cangelosi, A. (2019). Sensorimotor input as a language generalisation tool: a neurorobotics model for generation and generalisation of noun-verb combinations with sensorimotor inputs. *Autonomous Robots*, 43, 1271–1290.
- Zlotnik, G., & Vansintjan, A. (2019). Memory: An extended definition. *Frontiers in Psychology*, 10, 2523.