

A Memory System of a Robot Cognitive Architecture and its Implementation in ArmarX

Fabian Peller-Konrad, Rainer Kartmann, Christian R. G. Dreher,
Andre Meixner, Fabian Reister, Markus Grotz and Tamim Asfour

*Karlsruhe Institute of Technology (KIT), Institute of Anthropomatics and Robotics (IAR),
High Performance Humanoid Technologies (H²T), Karlsruhe, Germany*

Abstract

Cognitive agents such as humans and robots perceive their environment through an abundance of sensors producing streams of data that need to be processed to generate intelligent behavior. A key question of cognition-enabled and AI-driven robotics is how to organize and manage knowledge efficiently in a cognitive robot control architecture. We argue, that memory is a central active component of such architectures that mediates between semantic and sensorimotor representations, orchestrates the flow of data streams and events between different processes and provides the components of a cognitive architecture with data-driven services for the abstraction of semantics from sensorimotor data, the parametrization of symbolic plans for execution and prediction of action effects.

Based on related work, and the experience gained in developing our ARMAR humanoid robot systems, we identified conceptual and technical requirements of a memory system as central component of cognitive robot control architecture that facilitate the realization of high-level cognitive abilities such as explaining, reasoning, prospection, simulation and augmentation. Conceptually, a memory should be active, support multi-modal data representations, associate knowledge, be introspective, and have an inherently episodic structure. Technically, the memory should support a distributed design, be access-efficient and capable of long-term data storage. We introduce the memory system for our cognitive robot control architecture and its implementation in the robot software framework ArmarX. We evaluate the efficiency of the memory system with respect to transfer speeds, compression, reproduction and prediction capabilities.

Keywords: Humanoid Robotics, Memory-driven Cognitive Architecture, Working Memory, Episodic Memory, Long-term Memory, Knowledge Representation

1. Introduction

The human memory is one of the most astonishing and intricate systems in nature as it receives a huge amount of sensory data and processes it in a highly distributed way [1]. Humanoid robots face similar challenges, thus, developing memory-based cognitive control architectures is a key challenge to endow such robots with intelligent behavior and seamless human-like interaction abilities with the environment. A humanoid robot has multiple sensors producing sensor data and actuators executing motor commands which in turn allow performing actions as depicted in Figure 1. Several processes are needed for filtering, interpreting, and using the data to perform tasks such as object detection, task planning and execution, navigation, motion generation and control, learning and interaction.

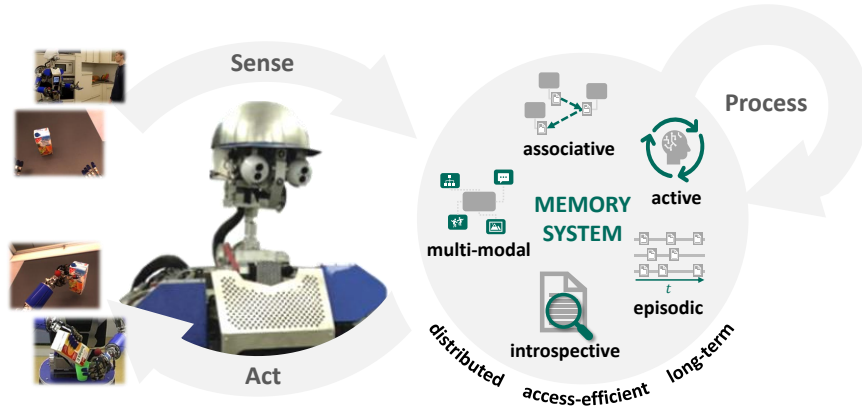


Figure 1: The humanoid robot ARMAR-III [2] perceiving multi-modal data of different sensors such as haptic, visual and audio information. This data has to be stored and processed in an adequate memory system. To optimally support various cognitive processes and action, a memory system must fulfill several characteristics, such as an active design, multi-modality, an inherently episodic structure, associativity and introspectability.

How can we create humanoid robot systems with rich cognitive and sensorimotor capabilities that are comparable to the human's, especially regarding learning and development? The authors in [3, 4] describe cognition as “the process by which an autonomous system perceives its environment, learns from experience, anticipates the outcome of events, acts to pursue goals, and adapts to changing circumstances”. This definition identifies three core components for

robotics. A cognitive system needs (i) *perception* components in order to perceive the environment, (ii) *processing* components in order to learn from experience, anticipate the outcome of actions and adapt to changing circumstances, and (iii) *action execution* components in order to purposely act to achieve goals. In addition, a robot’s cognitive architecture requires a place to hold information acquired from perception and action execution so that the processing components can access and process the data and eventually store the results back in this place – the memory. We argue that a *memory system* (or *memory* in short) is a key element in any cognitive architecture. Such a memory should not only connect system components and store multi-modal information on different levels of abstraction but provide mechanisms and services for abstraction of semantics from sensorimotor data, parametrization of symbolic plans for execution in a given context and prediction of action effects.

To this end, we formulate our first hypothesis: *A memory system in a cognitive robot control architecture mediates between i) high-level abilities, usually represented in a symbolic manner, such as language understanding, scene understanding, planning, plan execution monitoring and reasoning, and ii) low-level abilities, such as sensor data processing, sensorimotor control.* This means that the memory system must be able to process a huge amount of data – no matter if the data is symbolic (e. g. plans, words, relations, etc.) or sub-symbolic (e. g. images, joint configurations, forces, etc.). It must build a bridge between sub-symbolic (sensorimotor) representation and symbolic (semantic) representation, tackling the *signal-to-symbol-gap* by e. g. learning representations of perception-action dependencies in form of Object-Action Complexes (OACs) as proposed by [5].

This observation leads to our second hypothesis: *Multi-modal representations are key.* The ability to store multi-modal information, the efficiency of storage and retrieval or the ability to learn from such data require a meaningful and efficient multi-modal representation of knowledge. This representation must be specific enough to differentiate between symbolic and sub-symbolic information. On the other hand, the representation must also support generalization on sensorimotor and symbolic level. Further, it must support *associations of knowledge* [1] because the system needs an understanding of how perception and action are coupled and which sensations usually occur together. Such multi-modality applies explicitly to information resulting from different cognitive processes within the architecture.

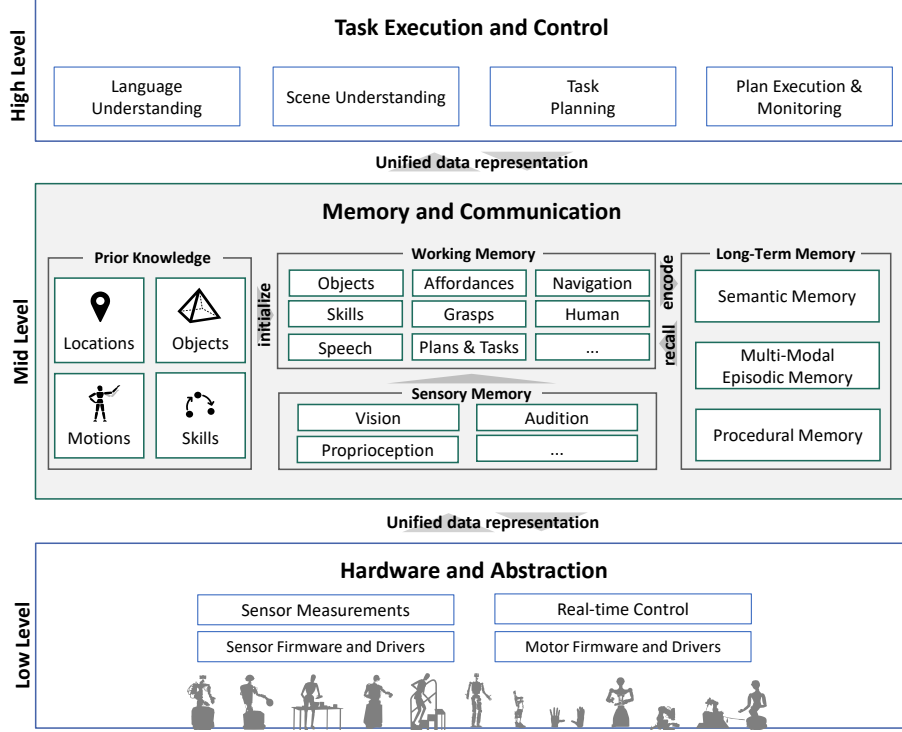


Figure 2: Overview of the memory system as mediator between a symbolic high level and sub-symbolic low level. The memory system is implemented in the robot software framework ArmarX ²

The main contributions of this paper are two-fold: (i) On a conceptual level,
 55 we identified requirements (or characteristics) of memory systems of a cognitive
 control architecture of a humanoid robot to support our two hypotheses. These
 requirements comprise the five concepts *active*, *episodic*, *multi-modal*, *associative*
and introspective, as well as the three technical requirements *distributed*, *access-*
efficient and long-term. We argue that a memory is an active element within the
 60 architecture, i. e. it is not simply a passive storage device but it is able to process
 multi-modal, both symbolic and sub-symbolic data. The ability to inspect
 information and to adapt behaviors based on such information greatly supports

²See <https://armarx.humanoids.kit.edu> and [6].

the active role on the memory. Further, we argue that it is beneficial to manage data episodically, no matter if it is declarative or non-declarative knowledge.

65 Finally, we explain that associations of how knowledge is shared within the system and how such knowledge is extracted from data yield great benefits for data processing, learning and development. The technical requirements are all related to improving the system’s overall efficiency. Further, we promote the need for memory-driven cognitive control architectures for complex robot systems
70 such as humanoid robots with these core characteristics to create systems with advanced cognitive abilities. (ii) The above considerations served as guidelines for the implementation of the memory system with the described core characteristics in our robot software framework ArmarX.

We take inspiration from cognitive science [7] for the realization of the
75 memory system of our cognitive robot control architecture but with special focus on optimizing data flow and making the development and integration of new components easier. The implementation further includes a novel way for representation information in the memory, which allows the inspection of information at run-time. Overall, we provide an overview of the technical
80 implementation of the memory system that meets the requirements of complex robotic systems, such as humanoid robots, including a broad evaluation of its efficiency. We would like to emphasize that with this work we do not describe a complete cognitive architecture, but focus on the aspect of a memory system in such architectures. From a user perspective, we show how our new memory
85 system optimizes data flow, enables mediation between high- and low-level in operation and leverages the integrated key characteristics in multiple use cases.

The remainder of this paper is structured as follows: Section 2 gives an overview of the memory classes known in cognitive psychology, which are often used for the realization of memory-based cognitive architectures for technical
90 systems. In addition, we compare related cognitive architectures with a special focus on their design and implementation of the memory including a comparison of the conceptual characteristics. The conceptual and technical concepts of our memory system are then motivated and described in detail in Section 3. Subsequently, we provide a technical overview of our memory system and its
95 implementation in ArmarX, including the working memory, the data representation format and the long-term memory in section Section 4 A broad quantitative evaluation of our memory system follows in Section 5. Section 6 discusses use cases of the memory system, demonstrating the usage of its new capabilities based on the implemented characteristics. In Section 7 the contributions of this

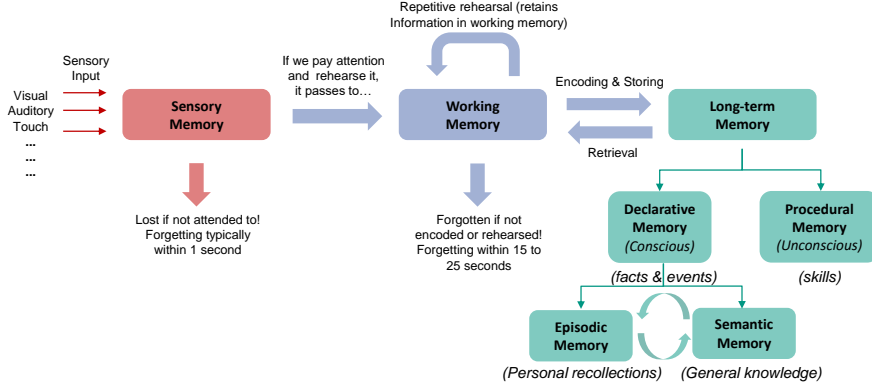


Figure 3: A taxonomy of memory classes. Multi-modal sensory information is passed from a sensory memory over a working memory to a long-term memory. Long-term memory itself is be subdivided into declarative memory and non-declarative (here procedural) memory.

work are summarized and planned extensions of the presented architecture are discussed.

2. Related Work

In this section, we briefly describe the different memory structures as motivated from cognitive psychology. Second, we discuss related cognitive architectures with respect to the identified core characteristics of a memory system.

2.1. Memory

By studying amnesic patients and animals in the second half of the 20th century, researchers found out that different areas of the brain are responsible for different memory tasks, motivating the assumption that memory consists of several subsystems [7]. Findings from the medical history of famous patients such as *Henry Gustav Molaison* or *Clive Wearing* greatly influenced the development of cognitive psychology and theories that attempt to explain the connection between brain function and memory [8].

Based on such theories of memory, [9] proposed the so-called *Multi-Store Model*. This theoretical model of human memory introduced three types of memory: (i) A sensory memory that processes perceptual information, (ii) a long-term memory that holds information for a long duration, and (iii) a short-term memory

that holds information through repetitive rehearsal and which receives information from the sensory memory through attention and from long-term memory through retrieval. Newer works distinguish between short-term memory and working memory or question a clear distinction between different memory types for biological systems altogether [10]. However, such taxonomies provide a basis to structure and classify artificial memory systems into a taxonomy. Figure 3 shows a taxonomy of memory classes. Multi-modal sensory information is passed from a sensory memory over a working memory to a long-term memory. Long-term memory itself is subdivided into declarative memory and non-declarative (here procedural) memory.

2.1.1. Sensory Memory

The sensory memory (SM) holds perceptual information for a short time duration and thus it acts as a cache memory or repository for incoming sensory information. In biological systems, sensory receptors take e.g. visual, auditory or touch information and forward it directly to the nervous system for further processing. It is common to distinguish between different sensory memory types for different modalities such as iconic memory for visual information, echoic memory for auditory information, etc. [11]. The sensory memory is assumed to be a highly volatile storage containing raw, unanalyzed sensory data that is derived from the senses. The data is only stored long enough to be passed to the working memory (WM). To limit the amount of data transferred to the WM, the data is only transferred when it is attended to, and is lost otherwise.

2.1.2. Working Memory

Similar to the SM, the working memory (WM) holds information for a limited time duration. However, this duration is longer and is in the order of seconds [12, 13]. In addition, the WM can consciously be controlled by attention mechanisms and is therefore important for reasoning, learning, problem-solving, and other mental processes. The ability to have certain details ready, even if they are not yet stored in long-term memory, supports a variety of everyday mental processes. Examples include remembering the first part of a sentence or keeping a number in mind while solving a mathematical problem. Studies show that people are able to keep track of several items at the same time [14]. This motivates the assumption that the WM has a relatively small capacity of 7 ± 2 *chunks* [15], when not using exploits like repeating information out loud, regardless of whether the elements are digits, letters, words, or other units. Newer research distinguish

between the modalities, e.g. seven chunks for digits, six for letters, and five for words [16]. According to [17], the total amount of chunks can not be increased –
 155 we can only increase the complexity of the referred information. The WM is not only exclusive to humans as animals have also shown similar abilities such as storing and maintaining several items simultaneously in memory, remembering their order and manipulating them [14].

The working memory is sometimes used interchangeably with short-term
 160 memory (STM) [10], but some consider the two forms of memory to be distinct and argue that working memory allows for the manipulation of stored information, whereas short-term memory refers only to the short-term storage of information [18]. For simplicity, we do not explicitly distinguish between WM and STM in this work.

165 2.1.3. Long-Term Memory

The long-term memory (LTM) is intended for the storage of information over a long period of time. Through the process of repetitive rehearsal and association, memories of WM consolidate to LTM or an existing memory in the LTM gets reinforced [9]. The WM can then retrieve the data from the LTM
 170 when it is needed for processing. During the process of consolidation, data is encoded into a special representation to identify groups and for generalization. However, memories stored in LTM are not saved in a static state. Studies showed that memories in LTM are transformed every single time they are accessed [19]. According to [7], the LTM can be divided into two types. (i) *Declarative memory*,
 175 also known as explicit memory, that contains information such as facts and events and is managed through conscious control, and (ii) *Non-declarative memory* contains implicit knowledge, such as the ability to perform various actions or behavioral control parameters.

180 *Declarative Memory.*

The declarative memory can be further subdivided into two types. The *semantic memory* consists of general knowledge about the internal state and the environment such as facts, ideas and concepts. In the context of cognitive architectures for technical systems, knowledge stored in this part of the memory
 185 is usually assumed to be symbolic [20]. In comparison, the *episodic memory* contains episodes or autobiographical experiences occurring in an explicit spatial and temporal context, i.e. what, where, and when something happened. There is evidence that the knowledge of how information has changed in the past in the

episodic memory as well as the knowledge about facts in the semantic memory
190 form the basis for prediction and explanation. There is a strong coupling between
them, as we derive new concepts from the experiences we have stored in the
episodic memory [21].

Non-Declarative Memory.

195 The non-declarative memory covers all information that is not consciously
accessible. Again, this memory can be further subdivided. *Priming* describes
the ability to strongly accelerate the retrieval of information from long-term
memory by a related stimulus. It requires that knowledge can be associated
heterogeneously, i. e. no matter how the knowledge is represented. Priming can be
200 further subdivided into positive, negative, semantic, perceptual and conceptual
priming. In *classical conditioning*, different stimuli are linked together. A
well-known example of classical conditioning is *Pavlov's dog* [22], which showed
increased saliva production just by ringing the dinner bell. The neutral stimulus
“ringing of the bell” was thus linked with the positive stimulus “there is food”,
205 which triggers a physical reaction. *Non-associative learning* is the simplest
form of learning as it does not require stimuli association [23]. Habituation and
sensitization are the two forms of non-associative learning. Habituation describes
the process of inhibiting a response after repeated exposures to a stimulus. The
degree to which a response is inhibited depends on the repetition rate of the
210 stimulus, its intensity, the duration of the stimulus, and how often the agent is
exposed to the stimulus. On the other hand, a sensitized stimulus has increased
intensity and sensitization does not require repetitive stimuli. Even a single
stimulus may cause a reinforced response. For example, relapse of addiction can
be seen as sensitization. Even a few stimuli, e. g. from drugs or gambling, can
215 trigger a strong physical desire. The *procedural memory* is required for skilled
behaviors and habits. There is less known about how humans store skills and
abilities except that skills are learned and refined through practical training and
that for learning a skill a large variety of areas of the brain are involved.

From the non-declarative memory, artificial systems often only explicitly
220 implement procedural memory [20]. However few approaches focus on how
artificial systems can be conditioned and how this mechanism can be used in
social robots [24].

2.2. Artificial Cognitive Architectures

The development of artificial cognitive architectures is a longstanding and still
225 unsolved problem, e. g. the development of the cognitive architecture *ACT-R* [25]
started in the early 1980s and is still ongoing. In general, artificial cognitive
architectures try to explain and represent the underlying processes of human
cognition, such as perception, attention, action selection, memory, learning,
reasoning, meta-cognition or prospection [26]. There is no implementation that
230 is able to solve all these tasks with similar performance to humans [20]. Thus,
the proposed architectures focus on different aspects of cognition.

[20] estimate the number of artificial cognitive architectures to be around three
hundred, of which about one hundred are being actively developed. Artificial
cognitive architectures can roughly be divided into three classes [20, 27, 3, 28, 29]:
235 (i) *Cognitivist architectures*, where knowledge is usually represented as symbols,
(ii) *emergent architectures* that focus on sub-symbolic processing and self-su-
pervised learning of how knowledge is associated, and (iii) *hybrid architectures*
that combine both types of processing. Not every cognitive architecture is fully
integrated in the sense that it covers all the aforementioned processes of human
240 cognition.

2.2.1. Cognitivist Architectures

In cognitivist approaches, cognition is achieved by computations performed
on internal symbolic representations. Approaches following this paradigm focus
on aspects of cognition that is constant and task independent [26]. This way
245 of representing information is natural, intuitive and often very performant as
irrelevant data can be abstracted away. Symbols are ideal to represent descriptive
problem statements and actions. They can easily be enriched with probabilities
and beliefs. Thus, symbolic representations are usually chosen for complex
high-level abilities such as planning, reasoning and language understanding. On
250 a semantic level, they can be used to learn from experience. Cognitivist agents
that act in the real world require rules to derive symbols from the sub-symbolic
sensory input overcoming the *Signal-to-symbol-gap* [5].

EPIC [30] (Executive-Process/Interactive Control) implements such a cogni-
tivist architecture. One goal of EPIC is to model cognitive executive processes
255 with a focus on detailed timing of human perceptual, cognitive, and motor activ-
ities during multiple tasks. EPIC requires *Sensory* and *Perceptual Processors*
to derive symbolically coded changes in sensory properties. These processors
accept visual, aural and tactile inputs. The WM of EPIC is a collection of

modality-specific items, but also contains perceptually unrelated information
260 such as goals and actions, and is updated periodically. Long-term knowledge is
encoded as production rules.

Related to EPIC, ICARUS [31] is a cognitivist architecture that focuses on
problem-solving. In each recognize-act cycle, perceptual information is converted
into short-term beliefs using categorization and inference and stored in the WM.
265 Such beliefs are used to find an action sequence from the current state to a
goal state. Again, production rules to abstract from sensory input to beliefs are
stored in the LTM, along with current goals and known actions. While EPIC
has a one-way connection from LTM to WM, and thus is not able to learn new
production rules, ICARUS is able to learn the connection between skill execution
270 and goal achievements in the form of new skills. Further, it is able to update
skill constraints from failure executions.

2.2.2. *Emergent Architectures*

Symbolic representations can also have restrictions. Cognitivist architec-
tures that have a fixed predefined set of production rules to generate symbols
275 from perception are less flexible and robust against a constantly changing envi-
ronment [20] as new production rules are not learned during operation³.

Thus, another approach for flexible systems that are able to adapt and learn
from the interaction with the environment focuses directly on sub-symbolic
representations as emergent systems do – typically by exploiting highly parallel
280 models based on connectionism. Ideally, such systems evolve during operation
to a full cognitive state. Emergent approaches are often heavily influenced
by neurology and cognitive psychology studies, especially those of infants in
which innate and learned properties and abilities are discussed and identified.
The architecture itself is usually assumed to be fixed and thus belongs to the
285 innate properties while ontogeny is achieved through learning from interaction
with the environment and experience [27]. Because of the focus on learning
from interaction with the environment, emergent approaches require a physical
body [26].

Unfortunately, implementations of such systems often lack transparency. It
290 is difficult for programmers to implement inference rules or to bring in prior
knowledge because this functionality is not intended by the architecture. The

³If some approach includes learning of new production rules from the sub-symbolic environ-
ment, it would rather be hybrid than cognitivist.

system should learn and adapt by itself.

As an architectural approach, the Self-Aware Self-Effecting (SASE) [32] emergent cognitive architecture is based on the idea that intelligence emerges from the interactions and connections between simple, low-level elements, rather than from a central, pre-programmed control system. In SASE, the basic building blocks of intelligence are simple, self-aware agents that can sense and act in their environment, and that are capable of forming complex networks of interactions and relationships. These agents, working together, are able to generate complex and adaptive behavior, without the need for a centralized control system. Emergence in SASE is achieved by the combination of self-awareness, self-effectuation, and interaction. Self-awareness allows an agent to sense and represent its own state, and thus to adapt its behavior based on that state. Self-effectuation allows an agent to act upon its own state, thereby influencing its own behavior. And interaction allows agents to exchange information and coordinate their behavior, thereby creating the potential for emergent phenomena, such as cooperative and competitive behavior, and the emergence of higher-level structures and patterns.

In our previous work [33], we investigated how an episodic memory can be modeled based on connectionism. In this approach, the memory was tasked to convert visual experiences into a latent representation that facilitates encoding, recalling, and predicting sensorimotor experience using unsupervised learning using Variational Auto-Encoders (VAE) [34]. Through the utilization of two different decoders, the learned representation can be reconstructed and predicted for the next frame. For memory, such a learned representation is beneficial as it generalizes and compresses knowledge. In [35], we extend this approach to support multi-modal information, including visual information, the robot’s configuration and platform pose in global coordinates, actions, plans and goal information, object labels and locations, and recognized speech as text. Using this multi-modal data, the system was tasked to answer user queries given in natural language. A deeper insight into how the memory supports this task is given in Section 6.4.

2.2.3. Hybrid Architectures

Hybrid systems combine the best of both previous approaches. Such implementations often use symbolic representations for high-level cognitive abilities and sub-symbolic representations for learning and development. Today, such cognitive architecture is the most prevalent type [20] and the most suitable for

the realization of robot cognitive architectures.

One of the earliest cognitive architectures that is still maintained and developed is Soar (State, Operator Apply Result) [36, 37]. The goal of the Soar project is to develop an artificial system that has similar cognitive capabilities as humans, i. e. knowledge-intensive reasoning, reactive execution, hierarchical reasoning, planning, and learning from experience, and to find out what computational structures are required to support human-level agents.

Perceptions are converted into a scene-graph-based representation and stored as such in WM along with information about targeted goals. Beyond the WM, Soar manages three different types of long-term memories: (i) A procedural memory that contains skills as if-then-rules, (ii) a semantic memory that contains facts and declarative information about tasks, and (iii) an episodic memory that manages experiences consolidated from the WM in form of snapshots. Next to symbolic representations, Soar includes sub-symbolic processing, i. e. to generate symbols from sub-symbolic percepts or to control symbolic processing [38]. Soar has multiple learning mechanisms for different types of knowledge: Chunking and reinforcement learning acquire procedural knowledge, whereas episodic and semantic learning acquire the corresponding types of declarative knowledge.

The ISAC (Intelligent Soft Arm Control) [39] hybrid cognitive architecture is constructed from an integrated collection of software agents and associated memories. The software agents encapsulate all aspects of perception, cognition and action and operate asynchronously. Perceptual events, encoded through a *Sensory-Ego-Sphere (SES)* [40] are first placed in the STM. An attentional network determines which events are relevant for the current situation and forwards this information to the WM. Additionally, WM temporarily holds information about motivation, goals, actions and internal processes if needed for the current task and encapsulates expectations for the future simulated by a *Central Executive Agent*. The LTM stores procedural, semantic and episodic knowledge (abstractions of SES, enriched with targeted goals, performed actions, outcomes and valuations) in multiple layers, supporting an efficient retrieval of memories. Associations are stored as state transitions in episodic memory.

LIDA (Learning Intelligent Decision Agent) was developed as a biologically inspired cognitive architecture to model all aspects of cognition in form of a global cognitive cycle [41]. It includes a large number of cognitive modules, some of which have short-term or long-term storage capabilities. Its cognitive cycle is divided into three phases: a perception and understanding phase, an attention phase, and an action and learning phase. During the perception and

understanding phase, sub-symbolic data from the environment is analyzed and translated into symbols corresponding to objects, entities or events in the *Perceptual Associative Memory* module. A *Current Situational Model* holds information about an agent’s present situation enriched through recall of experiences from the long-term memory modules by using local associations or similarity measures. Information about the present may decay if not stored in the long-term memory. During the attention phase, the content of the Current Situational Model is surveyed and the most salient information is broadcasted to all modules. During the action and learning phase, the modules use the broadcasted information for learning and execution, e. g. the *Procedural Memory Module* instantiates behaviors that can then be executed.

A comprehensive system is CRAM (Cognitive Robot Abstract Machine) [42] that fuses perceptual information, semantic knowledge taken from a suite of knowledge bases, and execution results of simulated actions in order to carry out vaguely defined goal-directed everyday activities. To abstract execution plans in such tasks, CRAM uses *designators* (i. e. placeholders) which require runtime resolution. As soon as available, these placeholders are filled by knowledge from the internal knowledge base *KnowRob2* [43]. This system contains a large-scale ontology of symbolic information used for reasoning and generalization. To integrate non-symbolic information, computable predicates are used. Episodic knowledge is represented as *NEEMS* (narrative-enabled episodic memories), a first-order time interval logic expression enriched with detailed episodic low-level information, such as perceptual or procedural events and signals. Sub-symbolic information is used through logic interface using computational predicates, inherently grounding resulting representations and ensuring consistency with the environment. The chosen data structure thus allows inspecting information in the memory in order to use it for learning, reasoning, simulating the outcome of actions and evaluating their feasibility. Apart from gathering knowledge in the real world, episodic knowledge can also be generated or refined by simulating actions in an inner-world model consisting of a high-quality virtual reality system and physics engine. Generalization and Specialization are achieved through meta-cognitive induction.

The predecessor memory model of our work was part of a hybrid architecture [44, 6] implemented in ArmarX. A centralized working memory component receives arbitrary data, i. e. symbolic and sub-symbolic information) stores it and eventually retrieves it when needed. The working memory is initialized via prior knowledge, a special sort of long-term knowledge containing only information

provided by programmers, such as object, robot and world model, object features, and pre-learned skills. The perception data is passed to the WM, replacing the last instance of the same type. Thus, the WM always only holds the most recent
405 information. Accessing the WM is done either directly or by listening to memory change events. Snapshots of situations including robot pose, robot configuration and object poses could be stored in the LTM. Skills and actions were stored in the procedural memory in the form of symbols mapping to executable statecharts that can be instantiated and executed if needed.

410 2.3. Comparison of Memory Systems in Artificial Cognitive Architectures

In artificial cognitive architectures, memory is an essential part of the system. Thus, almost every system described in the survey in [20] models some sort of memory. As described above, working memory, procedural memory, and semantic memory are commonly represented, while sensory memory and episodic
415 memory are explicitly modeled by only a few.

However, existing memory systems of cognitive architectures are often criticized for having the wrong focus [1], since in some artificial systems (e.g. [30, 31, 6, 44]), the memory is only assumed to be passive storage. This means that memory has no active state and that new information that enters the
420 system has no influence on how existing information is represented, processed or interpreted. Instead, memory should be an active part of the system that changes during operation. The memory itself must be able to adapt to circumstances and respond to new data. Some architectures partially implement this requirement, e.g. in Soar [36] where the WM does not react to new data while the
425 LTM involves learning from new data. In addition, many cognitive architectures focus on the implementation of different types of memory without considering the interconnection between them [30, 6, 44], associating entities of different modalities. Sometimes, only association between homogeneous data types is possible [39, 36]. Without the possibility to store multi-modal information and
430 to associate multi-modal memories, it is difficult for the system to draw common conclusions from knowledge obtained from different sources. Few systems hide the ability to associate information behind neural networks that learn a connection of data sources [35]. Such associations are however useless for applications outside the neural network, such as reasoning components. [1] also criticized the
435 way how information is represented in artificial memory systems. Physiologically, all areas of memory are the same, i.e. there is no structural difference between WM, episodic memory or semantic memory. Hence, we believe that memory

should support an inherently episodic structure and that the information stored in memory must have a unified representation. Artificial cognitive architectures usually utilize specialized representations and manage modality-specific knowledge in modality-specific containers [30, 31, 41, 42, 39, 6, 44].

In addition to the identified requirements by [1] and closely related to an active memory, we emphasize that introspection is another key component of memory and data representation as it allows the system to adapt its behavior based on the stored information, to learn and use the information for e. g. internal simulation, augmentation and prediction. Thus, introspection is also strongly related to the ability to monitor the system’s internal cognitive processes. Some cognitive architectures only focus on the perception-action coupling [31], however, other architectures explicitly model this capability [32, 36, 42].

In summary, we believe that a memory must be active, inherently episodic and multi-modal, associative and introspective. Only a memory that supports all these characteristics is able to adapt to new situations and respond to new data, have a unified representation of knowledge and support the overall system with data-driven services and data tracing for high-level abilities such as reasoning, explanation, prospection, augmentation or simulation. A more detailed explanation of these characteristics, including technical requirements for a memory system running on a humanoid robot system is given in Section 3.

Table 1 shows a comparison of the cognitive architectures described above. We omitted the architecture SASE [32] as is it more an architectural approach, not an implemented system and thus open to many possibilities. The requirement that the memory adapts to new data and actively changes is fulfilled in particular by emergent and hybrid architectures. For emergent architectures, this property therefore follows directly from their definition. Hybrid architectures usually use the ability to process symbolic and sub-symbolic data for learning and development, also fulfilling the active requirement. Some architectures do implement an episodic memory, but they do not support an inherently episodic structure in all memory structures. On the other hand, in our previous work in [35] we only model episodic memory omitting WM. We strongly believe that an inherently episodic structure is required for all modalities and for all memory structures. Multi-modality is implemented by several systems. [20] gives a more detailed review of which modalities are supported by which cognitive architecture. Associations are not explicitly modeled by every system. Especially associations between different modalities are beneficial as they allow the system to use correlated information during operation. To the best of our knowledge,

Cognitive Architecture	Paradigm	Active	Inherently Episodic	Multi-modal	Associative	Introspective
[30]	cognitivist	✗	✗	✗	✗	✓
[31]	cognitivist	✗	✗	✗	✗	✓
[35]	emergent	✓	(✓)	✓	(✓)	✗
[39]	hybrid	✓	(✓)	✓	✗	✗
[41]	hybrid	✓	(✓)	✗	✗	✗
[42]	hybrid	✓	(✓)	✓	(✓)	✓
[6, 44]	hybrid	✗	✗	✓	✗	✗
This work	hybrid	✓	✓	✓	✓	✓

Table 1: Comparison of the aforementioned cognitive architectures with respect to the identified requirements for a memory system in robotics. (✓) indicates that a requirement is only partially fulfilled.

475 even CRAM does not explicitly link related experiences with each other, but it does link sub-symbolic knowledge of experiences with symbolic derivations.

Finally, the ability to inspect data in the memory is naturally given for cognitivist architectures as data is represented as symbols. Some hybrid architectures claim to have an introspective datatype as well, usually facilitating
480 meta-cognition.

3. Requirements for a Memory System

This section explains the conceptual requirements already outlined and identifies the technical requirements for a memory for a complex robotic system.

3.1. Conceptual Requirements

485 The memory system should centralize knowledge but also structure data flow within cognitive architectures while at the same time being efficient, scalable, flexible and understandable. Instead of various data exchange channels between perception, processing and execution components, general methods and interfaces should allow standardized communication between high-level and low-level.
490 Figure 4 depicts required conceptual characteristics of the memory structure and the underlying knowledge representation. As described in Section 2.3, some




Figure 4: Identified conceptual characteristics for a memory system as motivated in Figure 1.

memory systems of related cognitive architectures already implement those characteristics partially.

In the context of humanoid robots, various sensors perceive the internal and external state. When performing a task such as bimanual mobile manipulation, many processing and execution components are involved simultaneously, including scene interpretation from a static and/or dynamic scene as well as the control of different parts of the humanoid robot such as e. g. hands, arms and platform. Consequently, representing **multi-modal** data in memory systems is mandatory, as is the **association** of knowledge [1], since most data modalities and experiences are highly correlated. In our memory, all knowledge shares a unified representation. This allows us to define general methods for processing multi-modal information. Further, we allow links between entities in the memory, thus fulfilling the requirement of associated knowledge. Those links can be established on a data level, e. g. associating haptic experiences with the performed action, or on an experience level, e. g. associating the task of searching an object with the interaction with the object.

Time-series information is hereby of particular importance. In humans, the episodic memory is concerned with the recollection, organization and retrieval of time-series information [45, 46], i. e. personally experienced events occurring at a particular time, place and context. We believe that this is also true for semantic and procedural knowledge. Seemingly factual knowledge may depend on the temporal context, and is only considered a fact because it changes at a very low frequency. An intuitive example of this is how Pluto was considered a planet for the longest time, until being reclassified as a dwarf planet in 2006. Therefore, all data, either produced periodically in streams or as a consequence of certain conditions, should be considered **episodic** and stored accordingly. Also motivated by the fact that physiologically there is no difference between episodic and working memory [1], we use one unified episodic structure for all parts of our memory, thus making it inherently episodic.

As already stated, the memory system should not be seen as a simple, passive storage device. It is an **active** part of an agent’s cognitive architecture [3, 1] which is highly influenced by the current context. The context not only influences what we store and forget but it also influences how we store information. E. g. highly emotional situations will create much stronger memories which decay much slower than memories taken from a normal situation [47]. In particular, the memory must play a key role in learning how to derive symbols from sub-symbolic multi-modal information by allowing to learn from experience, from interaction with the environment and by trial and error. This also means that the information stored in the memory must be **introspective**, i. e. the memory must be able to analyze the information, adapt its behavior to the given data and possibly discard knowledge if that knowledge is redundant or not necessary for the current situation. In our case, for example, the memory learns predictability from incoming data. Furthermore, the memory adapts to the amount of incoming data.

In order to *reason* about conditions of past events, *explain* why an agent acted in a particular way, to *predict* how the information might change in the future or even to *augment* existing or *simulate* new experiences we conclude that our memory system has to be **multi-modal, associative, episodic, active** and **introspective**.

3.2. Technical Requirements

Complex humanoid robotic systems further pose several technical challenges regarding the implementation of a memory system of their cognitive architecture. In this section, we highlight these challenges, identify the resulting technical requirements, and derive necessary paradigms. The requirements summarized in Figure 5 provide answers to the questions of (i) where to run the memory system, (ii) how to improve the efficiency of access, and (iii) how to deal with space limitations.

3.2.1. Where to Run the Memory

Complex robotic systems such as humanoid robots are data-intensive systems that use several special-purpose computers connected to each other via common interface (usually Ethernet) to process sensory data and control the action execution. To reduce the throughput and response times, in particular in the case of visual task, a cognitive memory architecture must be designed in a **distributed** manner, with memories located directly on the machine where




Figure 5: Technical characteristics of a memory system as part of a robot cognitive control architecture.

the data is produced. This makes extending the system at run-time easier, as additional memories can be enabled/added on demand, or disabled to save resources when not needed anymore.

3.2.2. How to Improve Efficiency of Access

Another possibility to reduce the throughput is to adapt to the production frequency of the data source. In general, we can distinguish between periodically produced and event-driven generated data, i.e. the data is either produced periodically in streams, or as a consequence of certain conditions, thus being event-driven. For systems with limited communication bandwidth, it is desirable to avoid unnecessary requests to the memory, for example, if no new data is yet available. This is especially true for event-driven data sources, where data may be produced at irregular intervals. For this reason, the memory is able to inform clients when new data is available. Clients can then decide whether to request the memory specifically or to wait for the next update, making the memory **access-efficient**.

3.2.3. Assessing What to Store

Humanoid robots are equipped with a variety of sensors that generate large amounts of data, potentially at high frequencies. A memory architecture thus needs to be able to handle such large data volumes, providing **long-term** capabilities due to limited main storage. Since the storage system of a humanoid robot cannot be large and fast enough to simply store everything, the data must be assessed and reduced according to their importance. This can be done by primarily storing data that was produced at or near significant points in time, e.g. keypoints, identified by higher-level cognitive processes, and/or using dimensionality reduction approaches to construct meaningful latent spaces.

Depending on the use case, specific models can also be used to aggregate data into a dedicated representation. Finally, the system must also be able to assess and delete data that has already been stored if it is outdated, irrelevant, or proves to be incorrect. Overall, in terms of assessment, we need to store data that allows better execution of the robots actions, analysis and reasoning on recorded episodes, and the prediction of future states of the robot itself or its environment.

Note that this requirement is strongly linked to *long-term memory* and *forgetting* in biological systems. However, since we look at memory from an engineering perspective, we evaluate this property to be technical. If a system with unbound memory and unbound computational resources were available, then long-term storage and forgetting may not be necessary. The conceptual requirements of a memory motivated in 1 apply to long-term memory as well. With this, we conclude that, in addition to the motivated conceptual requirements, our memory system needs to be **distributed**, **access-efficient**, and **long-term**.

4. The Memory System

In the following, we describe how we integrated the identified requirements in a memory system for our cognitive control architecture shown in Figure 2. Hereby, we rely on a multi-modal data representation to support introspection and fulfill key requirements of a suitable knowledge representation for robot cognitive architectures [48]. Similar to its predecessor described in [6, 44], our novel memory framework consists of:

1. *Sensory Memory (SM)* where data is held for a very short duration until it may be passed to the working memory.
2. *Working Memory (WM)* that holds the current state of the world and the robot’s internal state in a volatile storage. It can be updated by the SM through perceptual processes or by cognitive processes, e. g. recalling an episode from the Long-term Memory or prior knowledge.
3. *Long-Term Memory (LTM)*, which complements the WM, provides persistent storage capabilities and encodes the information into a more graduated representation.

4. *Prior Knowledge (PK)* contains information that is provided by the programmer and thus already known to the robot. During startup, the WM
615 is initialized with this known data such as robot, objects and environment models, pre-defined motion trajectories, etc.

First, we will describe the working memory, which is the part most clients are directly interacting with. Afterwards, we present the long-term memory and its learning capabilities and which adopts the working memory’s principle structure
620 but provides a more permanent storage than the working memory.

From a simplified perspective, the memory system can be viewed on three levels of detail: (i) The *distributed memory system*, which is a collection of memory servers running in their own processes, (ii) a single *memory server*, which stores data in episodically structured segments, and (iii) a single *data*
625 *instance*, which holds data in a general, interpretable format.

4.1. Distributed Memory System

The ArmarX memory system is a distributed system implemented through several *memory servers*. Each memory server is a separate ArmarX component, i. e. a process communicating with and providing interfaces to other components
630 via a middleware. All memory servers offer a common interface for, among others, reading and writing data. A concrete memory server may provide specialized interfaces for its respective modalities (e.g. objects). However, we want to emphasize that all memory servers have the same structure. All memory servers are able to hold arbitrary data – no matter if it’s symbolic or sub-symbolic
635 information. Thus, the memory servers are not modality specific. All memory servers register themselves in the *Memory Name System* (MNS) on startup. The MNS is a central registry, which allows memory clients to get access to the memory server handling a specific modality given a human-readable ID of the modality (e.g. “Object/Instance”). This is similar to and inspired by the Domain
640 Name System (DNS) of the internet, which resolves human-readable URLs to machine-readable IP addresses. Memory servers can be added at any time and distributed to several computers, distributing the load and reducing the network traffic and response times by running memory servers close to related hardware and memory clients such as a camera and our visual perception pipeline, greatly
645 supporting the system’s scalability.

As shown in Figure 6, the distributed memory system can be visualized as a cloud of memory servers, running in separate processes on different machines, each

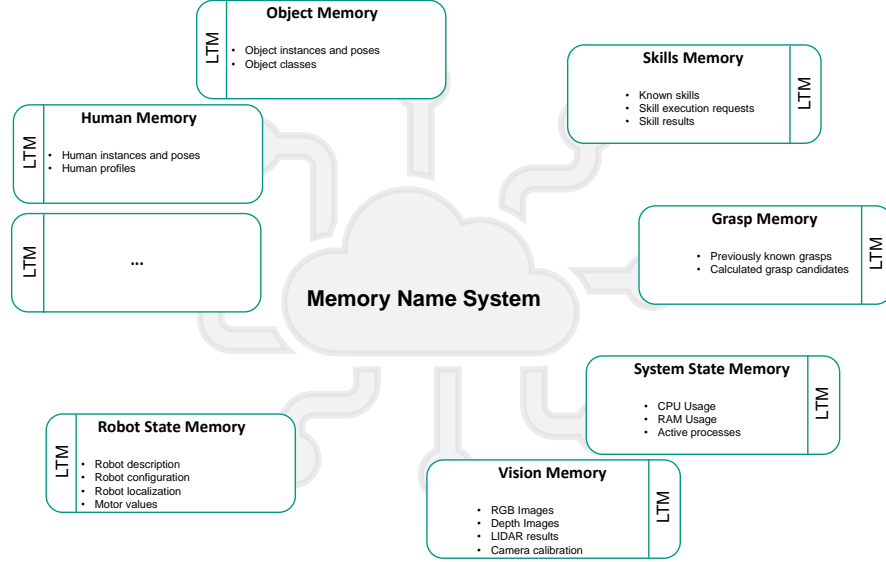


Figure 6: The distributed memory system with some of the memory servers running in ArmarX. Each memory is running in a separate process and may run on a different machines. All memories are connected through the MNS. Each memory server manages its own LTM.

one extending the system with new modalities. The robots' full memory is formed by the union of all distributed memory servers. Each memory server manages its own working memory and a corresponding persistent storage. Working memory and long-term memory of a single memory server share the same structure, i. e. hierarchical segments with a temporal structure with data in a general and introspective format. The working memory part is limited in space while the long-term memory part is not. The working memory can be accessed by clients to read and write data in a hierarchical structure. By querying the memory servers using temporal cues, the user can get access to the stored information. The memory servers accept precise queries to return data for a precise point in time as well as broad queries to return approximate information. Further, the working memory can notify clients as soon as they receive new information thus fulfilling the requirement of access-efficiency. The memory servers are used to distribute and structure the data – usually, each memory server holds information related to a specific modality, sensor or functionality to make the system more understandable for users. E. g. the *Object* memory holds all information related

Level	Description	Example	Key
Memory	Collection of semantically related modalities	Data related to Grasp planning and execution	<i>Grasp</i>
Core Segment	Homogeneous container of a specific modality.	Grasp affordances	<i>Affordance</i>
Provider Segment	Sub-segment containing data of a single provider	Results of a grasp planner	<i>MyGraspPlanner</i>
Entity	Physical thing or concept evolving over time	Grasp affordances of a specific object	<i>blue-cup</i>
Entity Snapshot	State of entity at a specific point in time	Grasp affordances at time t	<i>2022-02-18 13:06:56.492182</i>
Entity Instance	One data instance at a point in time	Second grasp affordance	<i>1</i>

Table 2: Levels of the working memory data structure.

to objects such as object class information (name, parent classes, meshes) or
665 concrete object instances (instance name, pose).

To free up space in the WM, information may be moved to the LTM. The
decision is based on the holding duration, space limitations or internal statistics
whereby associations are considered as well. If entity A has not been queried
recently, but is associated to entity B that is accessed very frequently, A will not
670 be moved to the LTM. Another alternative is to make that decision depending on
the system state. A special *System State* memory server can be used to watch the
CPU and RAM usage and to change the behavior, e. g. to reduce or increase the
maximum size of all memory servers running on this machine. Both, the working
memory and the long-term memory hold the stored information episodically, i. e.
675 they hold a dictionary mapping from timestamps to data instances. Additional
meta-information such as the name of the provider or the time it took to transfer
it to the memory can be used to analyze the data.

The memory servers implement the aforementioned concepts of a distributed
memory system and the inherently required episodic and associative structure.
680 Currently, our working memory holds data in plain text. There is no encoding
as most components require precise information about how the robot is moving
or where objects are located.




Figure 7: Hierarchical memory structure. The distributed memory system comprises several memory servers which are registered in the memory name system (MNS). A memory server has a name (e.g. *Object*) and multiple core segments (e.g. *Class*, *Instance*) specifying the stored modalities. Each core segment can have multiple provider segments (e.g. *PriorKnowledge*, *ObjectLocalizer*). A provider segments stores entities (e.g. *blue-cup*, *bottle*), which represent timelines of a thing or concept evolving over time. Each observation or update of an entity creates a snapshot at a specific point in time. Each snapshot can store multiple instances, which finally contain the payload data.

4.2. Hierarchical Data Structure

The general working memory data structure comprises multiple levels listed in Table 2 and illustrated in Figure 7, starting at a memory server down to a single data instance.

Data structures of level *Memory* as described in Table 2 is a semantic collection of one or many modalities, such as objects, actions, skills, images, locations, relations, the robot itself, and so on. This memory is usually represented by a single memory server and is identified by a memory name, e.g. *Object* or *Navigation*.

Segments are homogeneous containers of specific modalities, for example, object classes, object instances, grasp affordances, grasp actions, images, point clouds etc. The data structure contains two levels of segments: *core segments* and *provider segments*. Core segments are usually defined statically by the containing memory and determine a core data type shared by all data in this segment. Incoming data not fulfilling this data type will be reinterpreted as such type if possible. Provider segments are created by clients writing to the memory, i. e. providing data, at run-time. They serve three purposes: (1) They identify the source of the data, (2) they create a separate namespace for the provider to avoid conflicting names on the next level, and (3) providers are allowed to extend the core segment’s data type, such as pixel-wise labels in addition to bounding boxes in image segmentation, allowing some specialization while still being compatible with the general representation.

705 The elements of a provider segment are *entities*. An entity is a physical thing or concept that exists and evolves over time. Examples are images captured by a camera or object instances reported by an object pose estimation component. Entities are identified by names inside their parent provider segment. As entities are supposed to evolve over time, we represent an entity as a timeline or history
710 states over time. Therefore, entities and thus the whole memory are inherently episodic, fulfilling the aforementioned requirement of an inherently episodic memory. These discrete states are called *entity snapshots*. A snapshot can represent a new sensory observation, the result of a cognitive process, or in general an update to the knowledge of the robot. Finally, each snapshot contains
715 a list of an arbitrary (and potentially varying) number of *entity instances*. In practice, some modalities often only have one instance per snapshot, such as object instances and the robot’s state, but others usually store multiple instances, e. g. two images of a stereo camera or multiple extracted grasp affordances of an object instance.

720 All elements of the hierarchical memory structure have keys, which are unique inside their parent container. The keys of memories, core segments, provider segments, and entities, are textual names. The keys of entity snapshots are time stamps, while a single instance has an integer index. This way, all elements on all levels can be identified by a unique memory ID, which is the path through the
725 hierarchy down to the respective element. For example, the entity instance in Table 2 is identified by *Grasp/Affordance/MyGraspPlanner/blue-cup/2022-02-18 13:06:56.492182/1*, while the mere entity is identified by *Grasp/Affordance/MyGraspPlanner/blue-cup*, and the ID *Grasp/Affordance* refers to the core segment for grasp affordances.

730 These memory IDs allow to create associations and cross-references between instances of different modalities, fulfilling the requirement of an associative memory. Ideally, the output of a processing steps references the respective input. Thus, complete traces of data through the processing pipeline are stored in the memory, which could enable a cognitive agent to reason about the whole process
735 leading to, for example, a failed grasp attempt. Additionally, these traces help programmers to debug and improve processing pipelines.

4.3. Reading and Writing: Queries, Commits and Updates

Among others, memory servers offer interfaces for reading data from and writing data to the working memory. Writing is done via *commits*. A commit
740 is a bundle of *entity updates* which are written to the working memory . Each




Figure 8: Event-driven pipeline scheme. (1) Client C_1 writes data to server S_1 via a commit. (2) Server S_1 broadcasts an update notification, which is received by client C_2 . (3) Client C_2 performs a query to S_1 to read the updated data. (4) Client C_2 processes the data and commits the result to another server S_2 .

entity update adds or updates one snapshot to an entity. Therefore, to create an entity update a client must specify (1) the entity’s ID, (2) the snapshot’s timestamp, and (3) the data of each instance at this point in time.

Each time a memory server receives a commit, it updates its internal data structure and broadcasts a notification message to listening clients. This message contains the IDs of the memory snapshots that have been updated by the commit. When a client receives a notification, it can decide whether it needs to react to the updates; in this case, it can read the updated data from the server and process it. These update notifications allow building data processing pipelines in an event-driven manner as can be seen in Figure 8.

To read from a memory server, a client performs a *query*. The structure of a query follows the memory data structure: For each level of the hierarchy, it specifies which elements should be part of the query result based on their keys. For each of the name-keyed levels (memory, segments and entity), the client can request all entries, a single entry with a specific name, or all elements whose name matches a given regex. For the entity snapshot level, there are many options based on time stamps and logical indices. Commonly used snapshot queries include the single latest snapshot (the current state), the snapshot at a specific point in time, the snapshots in a time window (e. g. in the last 10s), and the latest n snapshots.

Multiple queries at the same level can be combined to create a conjunction, e. g. to request data from two specified core segments. In this case, each core segment query has its own sub-queries for the lower levels. As a result, queries have the same tree-like structure as the memory data structure, and can be flexibly combined to create complex queries.

4.4. Interpretable Data Format

The ability to inspect and analyze data is by default not supported by the used programming language C++. C++ is standard in the context of robotics frameworks in terms of performance. Therefore, we introduce the ArmarX
770 *Interpretable Data Format* (IDF), a special format that allows memory servers to use a common introspectable data representation. In its very essence, this data format implements a recursive variant formulation with special extensions that are frequently used in robotics (such as matrices, orientations, images or point clouds) and minor optimizations for network transfer. All data in the
775 robot’s full memory is a composition of basic data-objects (int, long, float, double, bool, string and multi-dimensional byte array) and recursive data-objects (list and dict). These objects can be inspected at run-time, are performant, are transferable over the network, and offer a strong interface via standard C++ classes.

780 Segments of our memory can be annotated with static type information . That is a mirrored representation consisting of basic type-objects (i. e. int, long, float, double, bool, string, time), special type-objects which get encoded into multi-dimensional byte arrays (i. e. matrix, orientation, image, point cloud) or container type-objects (i. e. list, tuple, dictionary, object, pair). Type annotations
785 do not contain data and only provide additive information for the underlying data objects. Thus, the relation of data-objects and type-objects is not bijective. Type annotations are optional. Even if some segment has no knowledge about the detailed type, it knows which members are present, their name and it can inspect the data. This property is useful when implementing type-agnostic
790 procedures which are applied to all data objects in the memory, no matter what they represent.

Users can specify the static type information through *XML*. Our system uses code generation as an abstraction mechanism to separate the type description from the implementation. Given the type information in XML, IDF automatically
795 generates so called *business objects* for a particular implementation language. It also generates common interfaces to cast IDF data objects to its corresponding business object and vice versa. These business objects are handier to use for clients as the auto-generated C++ classes allow programmers to make use of all programming language features, static type checking and code completion.

800 Separating type and data makes data objects interchangeable. Even if a data object only fulfills a part of some type, it can be casted to the corresponding business object. Unmatched members stay uninitialized.

Another benefit of using code generation mechanisms is that our system is able to generate similar code in different languages while maintaining the network transfer ability of our middleware. Thus, memory servers and memory clients are not required to be implemented in the same language. Currently, IDF supports C++ and Python as target implementation languages.

4.5. Long-term Memory and Deep Episodic Memory

While the working memory parts of memory servers represent a volatile in-place memory that is limited in space, the long-term memory part can hold a large amount of data for a longer period. The exact hold time depends, for example, on the access frequency and importance of the data and on the parameterization of the memory server. In our implementation, the long-term memory will be stored on the hard drive of the machine on which the memory server is running. This reduces the network traffic and increases the response times of the system as data is not transferred via the network to a central storage. To be space-efficient, the long-term memory converts memory instances into a compressed but still generative and predictable format. Further, the long-term memory is able to filter incoming and to forget already stored information.

The full pipeline to filter, store and encode data is shown in Figure 9. In an online phase, data is filtered based on their frequency or using fast similarity measures compared to other entities of the same type. Remaining data gets compressed using standard compression methods (e.g. JPEG [49] or MPEG-4 [50] for images and videos, ZIP [51] for arbitrary snapshots) in order to reduce the final file size. Which online compression technique is suitable depends on the parameterization of the memory server. Users can choose to not filter incoming information and to only use lossless compression, which allows the long-term memory to be used as a general tool for data recording.

In an offline phase, i.e. when the robot is inactive, the memory further compresses and abstracts the episodic information using machine learning techniques into a *deep episodic memory*. We utilize Auto-Encoders (AE) to convert the snapshots into a generalized but still predictable format, similar to [33, 35]. The ability to encode information evolves and improves over time as more training data is available.

Currently, the encoding works for data types that do not contain strings or lists with a varying number of elements as they can not be converted into a well-defined input format for machine learning. Our framework automatically extracts those kinds of information and stores them separately, while the rest

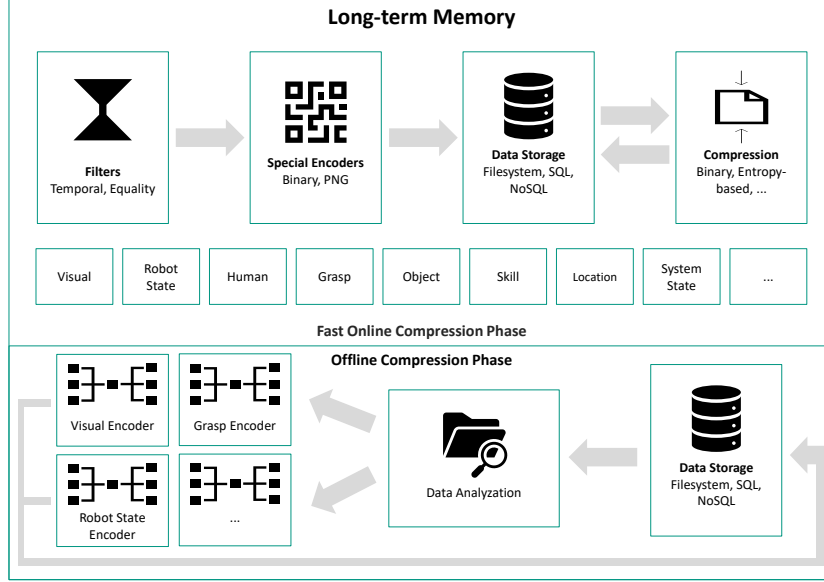


Figure 9: The full pipeline to filter, store and compress data in the long-term memory part of the memory servers.

of the information is passed to the AE. This representation does not affect the efficiency of existence checks of snapshots. If a client queries a snapshot that has already been moved to the LTM and converted into its latent representation, the memory can still access corresponding meta-data. If the memory finds a matching snapshot, it will decode the information, convert it back to IDF and finally return it to the WM. The WM can then forward the result to the client. Thus, at run-time, each long-term memory part of the memory servers manages its data in two representations: (1) recent information, that is filtered and compressed using the online methods and (2) older information, that has already been brought into a latent representation.

We additionally implemented methods to use the information in the LTM for the generation of new knowledge. Internal simulation is a powerful technique to refine the models used for generalization and prediction in our deep episodic memory. Using a high-quality simulation engine [52], we are able to render past experiences in new scenes or even to simulate the interaction with new objects as shown in Figure 10. This technique is only used to simulate visual data, but we plan to use it for physical interactions as well. The simulated data is finally used




Figure 10: Techniques to visually augment experiences using a photo-realistic simulation engine. **(a)** shows the original scene. **(b)** shows the same scene but with a different colored object. **(c)** shows the same scene with a replaced object.

along with real data to train the auto-encoders of the deep episodic memory.

In summary, both the WM and the LTM are active parts of the memory system. The WM can adapt its behavior to the current computer load and refines its models to predict information and the LTM actively searches for suitable data representations in order to compress, reproduce, generalize and predict them.

5. Technical Evaluation

We evaluated the new memory system empirically. Since all components communicate with and through the WM we begin by evaluating the WM. The WM must support large data transfers while providing a high availability for querying information. Second, we evaluate the LTM. Other than the WM, the LTM must be able to store lots of data which, sooner or later, requires the LTM to compress or even forget information.

All measurements were performed on computers equipped with an AMD Ryzen 9 5900X with 12 cores, 64 GB RAM and a NVIDIA GeForce RTX 3060 running Ubuntu 20.04. We define three different data objects with an increasing complexity taken from real-use implementations in ArmarX. First, a simple object only containing a single long value (8 bytes of information). Second, a moderate object containing a long, a string (with a fixed length of five chars) and a memory link. In sum, this moderate data object has 33 bytes of information. Finally, a complex object containing memory links, nested objects, an image with resolution $128 \times 128 \times 3$ and several maps with fixed size. In total, this object contains 49.225 bytes of information.

For each test, we calculated the means and variances over 1000 samples. Further, we performed each test either with one single data entry or with a batch of 20, 50 or 100 entries.




Figure 11: Overview of data transfer to and from the memory from arbitrary components. We evaluated our memory by measuring the times required to commit to and query data from the memory.



Figure 12: Measured times to send a single or multiple instances to the memory including the time to convert the auto-generated business objects to IDF. We calculated the mean values over 1000 individual samples per data type and per batch size. Note that the first two charts have a logarithmic scale on the y-axis.

5.1. Evaluation of the WM

We measure the time a producer needs to convert business objects and send IDF data-objects to the memory and how long a consumer needs to receive and convert the information back after being notified by the memory. Since all components run in different processes and maybe on different machines, the chosen middleware ZeroC Ice [53] requires time to serialize the information, access the network stack, transfer the information to the remote machine and finally deserialize the information. We treat the middleware as a black box system, thus, we have no influence on the times required for each of those steps. Because our memory is written independent of the chosen middleware, it could also be replaced by e.g. ROS, which has lower mean latency values for the transmission of data [54]. We will compare the measured results against direct peer-to-peer connections and publish/subscribe channels.

Figure 11 provides an overview of how the measured times affect the overall data flow to and from memory. The times to commit information to the memory

are shown in Figure 12. The diagram “*Full Commit*” depicts the times needed to commit IDF data-objects to the memory, including the network overhead. The diagram “*Data Storage*” however only shows the time needed to insert the
900 received information into the memory structure.

With increasing complexity the times required to convert or commit large batches of objects increase linearly with their data size. However, when transferring small amounts of data over the network we observe minimum costs for accessing the network stack. Even a transfer of small objects takes at least 100 μ s.
905 We can also observe, that the times to convert and transfer batches of moderate data are closer to the ones of complex data although its data size is closer to the simple type. One has to keep in mind that the biggest part of the complex type comes from an image (49.152 B) which will be converted into a single IDF NDAarray which can be converted and transferred much more efficiently. Apart
910 from the image, the complex type is approximately two times bigger than the moderate type, which can also be observed when comparing the transfer duration of both experiments.

The time to convert a business object to its IDF representation not only depends on the data size – it also depends on the amount of members and on
915 the batch size. This is why the conversion of a moderate type is approximately 15 times slower than the simple one. The moderate and complex objects almost have a similar number of members. The times to arrange information in the memory only depend on the data size. In our experiments, we used a single memory server with already 1000 entries to ensure that all insertion times are
920 comparable. Further, a memory with already existing entries is closer to real-use scenarios. While we can observe a small static overhead to access the memory structure and to insert commits of experiments, the measured times for the batched experiments prove the aforementioned dependency on the data size.

Once the memory receives new data, it will broadcast a notification to all
925 subscribers. The subscribers may then query the information to receive the new data. The times required to query information are shown in Figure 13. The chart “*Full Query*” shows the times it takes to query and receive data from a memory server without converting the data from IDF to its business object representation and “*Memory Lookup*” shows the time needed to search for a
930 specific entry in the memory.

Similar to committing information via the network we observe that querying information depends on the data size and the batch size while there is a static overhead for accessing the network stack. We find out that the time required to




Figure 13: Measured times to query a single or multiple instances from the memory including the time to convert the transferred data in IDF to the auto-generated classes. Again, we calculated the mean values over 1000 individual samples per data type and per batch size. Similar to 12, the first and the second chart have a logarithmic scale on the y-axis.




Figure 14: Averaged age of data as difference of timestamps between data reception and production, either using the memory, peer-to-peer (P2P) connections or publish/subscribe (PS) channels provided by the middleware ZeroC Ice.

search for instances in the memory is almost similar over all experiments with the same batch size because the memory only returns references to the stored information. During a query, the data is automatically converted back to its business object representation. Similar to the opposite conversion we find a similar dependency on the number of members and the data size.

To conclude the results of our first experiment, we compared the full times from producing information to receiving that information in the consumer component with direct peer-to-peer (P2P) connections and publish/subscribe (PS) channels managed through a centralized service as shown in Figure 14. P2P times can be seen as a theoretical lower bound for sending data in a distributed system. A PS service is more comparable to the memory as it receives data from producer components and forwards it to an arbitrary number of subscribers. However, it lacks the ability to access a history of data or to actively request data.

The memory performs worse of the three, but it offers important features for a cognitive system, such as being able to query information from the past or to provide general data-driven services to clients. Further, this evaluation should be seen as a worst-case scenario. Usually, producer, memory and consumer processes are not implemented as independent processes or running on different machines. Producers often implement a memory server, so that data does not have to be transferred over the network. The times measured here can therefore be halved in many cases and by a well-chosen distribution of components. Nonetheless, even latency times measured here for moderate data types allow data to be retrieved from the memory at a maximum frequency of over 160 Hz, which is enough for state-of-the-art model predictive control algorithms such as in [55]. For the few components in our software framework that require accessing the data more frequently, such as self-collision checks and the emergency stop, side channels (P2P and PS channels) are still possible and valid.

5.2. Evaluation of the Deep Episodic Memory

In the long-term memory, it is important to find models with a good compression rate, generalization ability and prediction ability which we will briefly evaluate in this section. For this evaluation, we recorded five experiments on the real humanoid robot ARMAR-III. Similar to [35], the recordings include stereo camera RGB images (with a resolution of 640×480 pixels at a frequency of 20 Hz), the robots pose and configuration including the positions and velocities of 43 DoF, action information, plans, object instances and their locations. The robot was instructed to perform simple pick-and-place tasks in a kitchen environment.

If unfiltered, the LTM receives 35.665 MB of data per second. Images are the largest part of the data. The online compression used for the recordings reduced the amount of required disk space to 0.213 MB/s (reduction by 99.4%). First, the input data was downsampled to a maximum frequency of 5 Hz. Further, images were compressed using PNG compression and all data was converted into a binary format using ZIP compression. Offline compression using Auto Encoders further reduced the required disk size to 0.003 83 MB/s (which represents an additional reduction by 98.2%).

For the quantitative evaluation of the network, a leave-one-out cross-validation was performed to obtain 5 folds of training and testing sets. To compare the results to our previous work [33], we will concentrate on visual information and calculate the PSNR [56] as shown in Table 3 for the prediction and for the




Figure 15: Reconstruction and predictions of the next few frames returned by our deep episodic memory. The first original image is encoded into a latent space. A decoder model reconstructs the image or predicts the next frames.

reconstruction. Because our deep episodic memory automatically instantiates
 985 one AE per entity, other modalities can be reconstructed and predicted in
 the same way. The results show, that our multimodal model outperforms the
 previous work in terms of reconstruction and prediction although the original
 work was pre-trained using large datasets. Still, the performance to predict
 frames decreases with increasing horizon. This is natural, as the uncertainty
 990 increases. Examples of reconstructed and predicted visual experiences are shown
 in 15.

With this work, we want to briefly introduce the intermediate results of our
 research regarding representation learning in the robot’s long-term memory. We
 are still working on a final and all-encompassing evaluation of the presented
 995 system with many different modalities and real application examples.

Iteration	Reconstruction	Prediction of next frame	Prediction of over next frame
1	46.828 591 2	33.675 896 98	28.308 418 52
2	48.637 354 92	32.267 090 11	34.612 041 75
3	44.220 407 39	32.056 580 62	30.045 160 65
4	52.723 655 49	35.525 423 89	30.982 891 46
5	48.212 150 29	33.116 856 94	28.998 586 07

Table 3: PSNR values (averaged) of the test set during a leave-one-out-cross-validation of the generalization network used to compress visual information in the long-term memory. The generalized latent representation offers the possibility of reproduction, production and finding of similar instances as described in [33].

6. Case studies

In this section, we will showcase different use cases of varying complexity, showing how different components are supported by the memory system and how processing pipelines can structure their data flow to make the best use of it. The following examples should be seen as different views on one and the same robot’s cognitive architecture, e. g. there is only one *Robot State* memory managing information related to the robot pose and configuration which can be used by any other component.

In our architecture, executable robot skills and behaviors are managed through a *Skill* memory. It connects executable codelets with memory. In the following, we describe this memory-driven skill framework using the example of how a high-level speech command is being stored and processed within the memory architecture resulting in the desired robot’s low-level behavior. Here, we focus on two different aspects (1) high-level symbolic skill representation and (2) sub-symbolic skill execution. The former will explain how arbitrary skills can be registered to the memory. The latter will explain how the memory system listens to certain conditions to instantiate and execute skills.

6.1. Skill Representation




Figure 16: The connection of executable skills to the memory system. Skills can be triggered by e. g. speech.

Skill Providers register information related to executable codelets, so called *skills*, to the skill memory as shown in Figure 16. Skills that are available during start-up are taken from the prior knowledge (PK) of the robots’ memory. However, skills can be added, removed, enabled or disabled at any time. An instance of a skill in the memory incorporates a symbolic description that itself

links to executable code. There can be a variety of skill providers which register
1020 skills in this memory. Skill providers thus group corresponding and related skills.
Depending on the definition of the skill, it can be parameterized both using
symbolic (e.g. "Go to the table") and sub-symbolic data (e.g. "Go to position
(10,10)"). A skill may be linked to conditions that trigger its execution. The
simplest condition may be that the execution of a skill is directly requested.
1025 However, one can also map complex conditions, such as in response to a certain
events such as a recognized object in the case of grasping. Once all conditions
of a skill are fulfilled, the memory covers the instantiation and execution of the
executable codelet of the skill. An event that the skill is executed is stored in
the skill memory. The result of the execution (e.g. whether it was successful,
1030 stopped or unsuccessful) is stored in the memory as well. Intermediate results
are usually associated with the execution of the skill. In the following, we will
further describe how a skill internally uses the memory for its operation using
the examples of navigation, object localization and knowledge verbalization.

6.2. Navigation

1035 Robot navigation is realized through navigation skills. Once the user requests
the robot to "Go to the table" such a navigation skill gets executed. It will
access the navigator service as depicted in Figure 17. This service resolves
the Cartesian position of the location, plans a collision-free path, executes the
movement and reacts to dynamic changes in the environment. Once the planning
1040 phase ends, the result will be stored in a specific *Navigation* memory. During
planning, accessed data can be multi-modal such as named locations and cost
maps needed for path planning, as well as the robot's global pose. In addition,
components such as the cost map provider ensure that the navigation memory
actively updates itself, as a change to the object memory can trigger an update to
1045 the distance cost map. The planning result is associated with the corresponding
skill execution request. This association may be used for question answering
(e.g. "Why did you move to the table") or for debugging purposes. Afterwards,
the navigator service continues with the execution phase. The skill waits for the
robot to reach the desired location but can also react to unforeseen events such
1050 as the robot getting trapped due to surrounding humans.

As for any memory server, clients subscribe to receive updates on memory
segments. Thus, other components can also listen to the *Navigation/Events* core
segment and react to events such as the robot starting to move or there is a
planning failure.

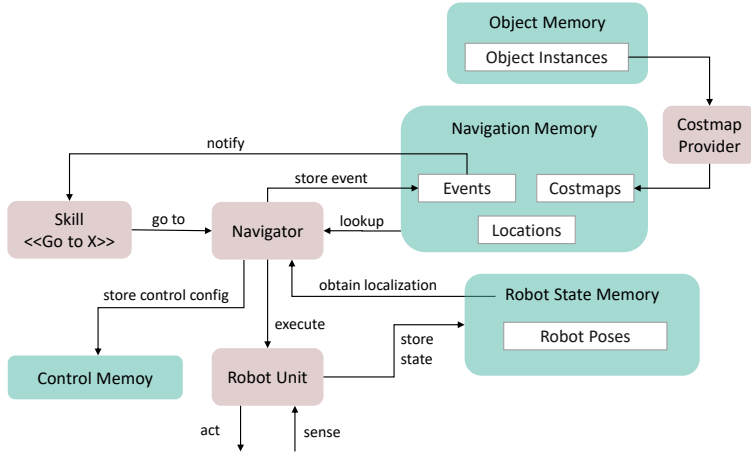


Figure 17: Initial planning phase during the navigation skill execution.

1055 All information in the memory system is fully introspective for the system
but also for the user. As an example, the state of the memory system after the
initial planning phase is shown in Figure 18. On the left, the GUI shows the
content of the *navigation memory* as a tree view. Each entity instance can be
investigated as shown on the right. In this example, the view shows the robot's
1060 platform location associated with the table. On the bottom right, the content of
the memory is visualized in 3D.

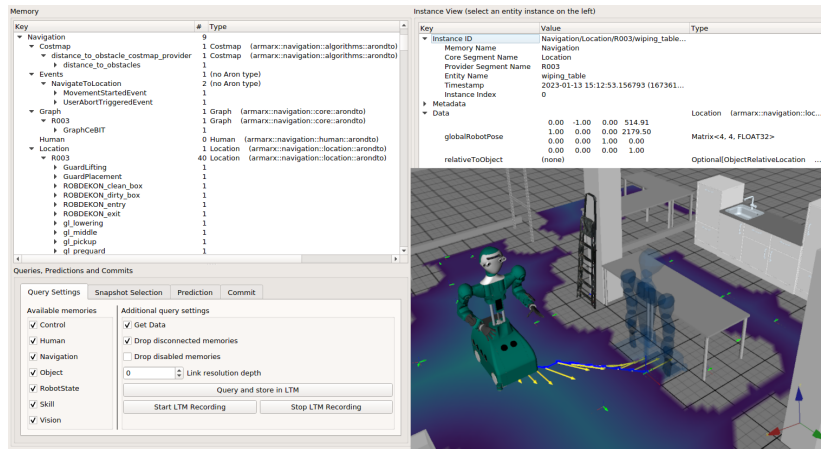


Figure 18: The memory GUI and the 3D visualization.

6.3. Object Detection and Localization

Another important ability of a robot is to detect and localize known or unknown objects in its environment. In this example, we consider the task of
1065 6D localization of known, textured objects based on RGB or RGB-D images. Consider the processing pipeline in Figure 19. A camera driver component reads data from a hardware camera. In each frame, it commits the RGB and depth images to two separate segments in the *Vision* memory. If enabled by the user (e.g. by running an object localization skill) an object localization component
1070 listens to updates of the *Vision/RGB* and *Vision/RGB-D* segments to get notified about new images. When notified, it queries the latest images from both segments. In addition, it uses multi-modal related prior object information such as 3D models, likely positions or pre-extracted visual features from the *Object/Class* segment which itself is initialized from the prior knowledge. The localization
1075 component then applies its internal method to detect objects and estimate their poses. Finally, it commits the new observations to the *Object/Instance* segment, which holds the current state of objects in the scene. Object poses are associated with the images they are calculated from. Additionally, they are linked to the current robot pose. The latter allows for analyzing how the object was placed
1080 relative to the robot and how the robot interacted with the object retrospectively. For efficiency, camera drivers and the vision memory are running on a separate machine, equipped with special-purpose hardware for computer vision. Prior knowledge is equally available on all machines.

Some 6D localization components benefit from an initial guess e.g. regarding
1085 the pose of a moving object. Our memory supports those components by providing methods to predict information. This functionality is not limited to object poses – if a memory server only contains data of similar type it can predict the numeric elements. Prediction is realized on two levels: (i) fast methods only using information stored in the WM with updates each time new data enters the
1090 WM, and (ii) more complex systems using the information stored in the LTM as described in 4.5. Requesting a prediction from the memory is done by querying a future timestamp. Depending on the memory server’s configuration and the amount of received data, the memory decides to either use the fast methods of (i) or more complex systems of (ii) to generate a prediction. Refining the models
1095 based on an evaluation of how good this estimation has been is not part of this work.

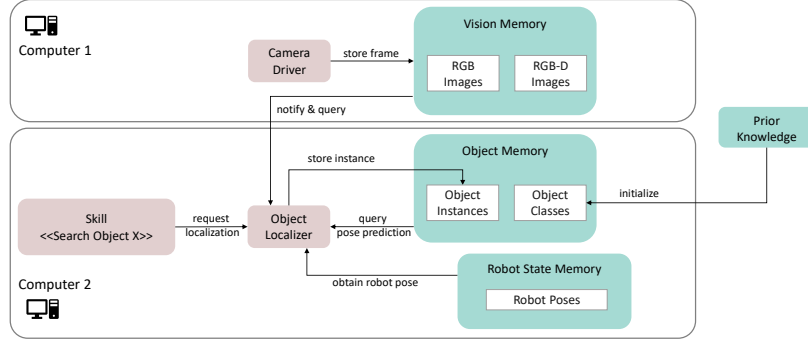


Figure 19: Integrating an object localization component into the distributed memory system.

6.4. Verbalization of Episodic Knowledge

Verbalization of episodic knowledge is fundamental for natural human-robot interaction. This ability requires the robot to understand the user’s query, access its memory and finally generate a response. In our previous work [35] we presented a system that is able to answer user questions about the past, such as “Where have you seen the red cup?” or “What did you yesterday?” by using the knowledge stored in the robots deep episodic memory. Further, the robot is able to verbalize important events such as successes or failures of actions, e. g. “I tried to put the milk in the fridge but I failed to open the door”.

Technically, a special verbalization skill listens to the *Speech* memory server. This memory server contains recognized spoken text as simple strings. Every time a spoken text is recognized and this update is broadcasted by the memory, the verbalization skill queries the latest recognized text and interprets this text as a verbalization command. Thus, we assume that the skill has been activated so that the robot is in verbalization mode.

The user query is passed into a neural network which internally applies attention [57] on the latent instances stored in the robot’s deep episodic memory. Plain text knowledge of the WM is used for verbalization. The latent representations are learned unsupervised and beforehand as described in section 4.5. As the robot should be able to verbalize symbolic and sub-symbolic knowledge our verbalization component requires access to multi-modal data sources. In our experiments, we granted access to executed plans and skills , object information , robot configuration and positions and visual information. A decoder network finally converts the latent representation of the user query and the latent

representations of the attended episodes into a response from the robot. The generated response is again stored in a separate segment of the speech memory of the robot. The deep episodic memory depends on its content, as its weights may be updated with each new instance which gets consolidated from WM to LTM. Thus, this part of our memory cannot be seen as a simple passive storage device, fulfilling the requirement of an active memory. Further, learning from data instances requires an understanding of the data. This would not be possible without the aforementioned introspectability of our data representation.

7. Conclusion and Future Work

In this work, we identified several characteristics of a memory system for a robot cognitive control architecture and revealed details about the functionality and representation of structures in the memory. We also described the implementation of the memory system in our robot software framework ArmarX. Our new memory system fulfills these requirements in the following way:

- It is an *active* memory. The working memory (WM) adapts its behavior based on the current computational load and learns from incoming data and the long-term memory (LTM) adapts its ability to predict and generalize to the given data.
- It is *multi-modal*. The memory has no constraints against the input data. Everything can be stored and encoded, no matter if it is symbolic or sub-symbolic information.
- It is inherently *episodic*. All information is stored episodically all over the memory – even semantic information.
- It supports an *associative* structure. Entities can be linked to other entities of the same or even of different modalities supporting reasoning and explainability. We believe, that the knowledge of how information is connected is as important as the information itself.
- The proposed data representation is *introspective*. This allows the memory to investigate the data and to use it for learning and development, to check constraints and to ensure consistency or even to change information for simulation or augmentation.

- Due to the fact that our robots have several special-purpose computers, our memory follows a *distributed* design approach. All subsystems of the memory (memory servers) may run on different machines leading to reduced network traffic and increased response times. A special centralized component, the *Memory Name Service*, manages the connections to all memory servers.
- Programmers can choose to either listen to memory updates or to actively poll the memory. Both access types have benefits and drawbacks which depend on the production frequency of data sources. Enabling the memory to support both access types makes the memory more *access-efficient*.
- Finally, as robot systems are not online all the time and they have limited storage, the memory must be able to efficiently store information in the *long-term*.

We showed how the data is structured in distributed memory servers, how the data can be accessed and passed to client components. Each memory server contains a long-term memory back end, which allows the memory server to store knowledge persistently. Whether and when data is moved to the LTM is determined by internal statistics, or by the user. The LTM has the task of storing the data as efficiently as possible. For this purpose, it makes use of various compression methods. In an offline compression step, data is converted into a generalized format that allows the prediction and comparison (in terms of distance measuring) of snapshots using machine learning. An introspective yet efficient data representation allows inspecting data instances and eases communication with the memory through code generation.

The performed evaluation shows that this memory implementation increases transfer times of data from producer to consumer components in comparison to peer-to-peer transfer due to the new features. Nonetheless, due to optimizations in knowledge representation of large data types, even those increased times are acceptable for robotic applications.

Further, we presented three use cases explaining how the implementation of these requirements influences the way we develop and evaluate our software. Nonetheless, the system is by far not complete as the development of a novel cognitive architecture (from which the memory is the key element) usually requires years. So far the memory accepts only a temporal structure and temporal requests, however, human episodic memory manages knowledge spatially and

contextually [58]. Allowing spatio-temporal and contextual access to information will allow more efficient retrieval of data. Regarding the efficient management of data, our LTM currently only provides rudimentary tools to filter incoming data and remove existing instances. As a design choice, the WM data is not filtered or encoded at all. During our experiments, we did not reach the limits of storage capacity but we believe that we need a wider range of filter and removal algorithms to manage the data in the LTM more efficiently, especially when considering lifelong learning and 24/7 operating modes. Thus, we are also working on extending the internal statistics (e.g., using saliency models) to provide a larger groundwork for filtering and forgetting algorithms. Moreover, it is currently not possible to incrementally refine the models found for learning representations in the LTM. We either generate a new model only for the experiences of the current work cycle or we remove the models we have, concatenate the decoded experiences and train a new model with this dataset extended by the current work cycle. The latter approach is less space-consuming but also carries the decoding error into the next model. Incremental approaches are desirable.

Last but not least, there is a lack of methods to evaluate and learn from the stored data. In particular, it would be interesting to compare the memories of several robots and make them available to each other (comparable to what is done in [59]). Shared memories, with other robots or with humans, also open up questions about security and privacy which should already be addressed at memory level. For this reason, we recently started working towards privacy-aware memory systems and cognitive architecture [60].

The planned work mentioned above represents only a small subset of the further development of the memory system of our cognitive architecture. We believe that the architecture we have chosen, based on the identified requirements of a memory system, is general and at the same time specific enough to support the development of cognition-enabled and AI-driven robotic applications on complex robot systems such as humanoid robots.

8. Acknowledgements

The research leading to these results has received funding from the Baden-Württemberg Ministry of Science, Research and the Arts (MWK) as part of the state’s ”digital@bw” digitization strategy in the context of the Real-World Lab ”Robotics AI” and from the European Union’s Horizon 2020 research and innovation programme under the grant agreement No 730994 (TERRINet).

References

- 1225 [1] R. Wood, P. Baxter, T. Belpaeme, A review of long-term memory in
natural and synthetic systems, *Adaptive Behavior* 20 (2012). doi:10.1177/
1059712311421219.
- [2] T. Asfour, K. Regenstein, P. Azad, J. Schroder, A. Bierbaum,
N. Vahrenkamp, R. Dillmann, ARMAR-III: An Integrated Humanoid Plat-
form for Sensory-Motor Control, in: *Proceedings of the IEEE-RAS Interna-
tional Conference on Humanoid Robots (HUMANOIDS)*, 2006, pp. 169–175.
1230 doi:10.1109/ICHR.2006.321380.
- [3] D. Vernon, *Artificial Cognitive Systems: A Primer*, MIT Press, 2014.
- [4] D. Vernon, Cognitive System, in: *Computer Vision: A Reference Guide*,
2014, pp. 100–106.
- 1235 [5] N. Krüger, C. Geib, J. Piater, R. Petrick, M. Steedman, F. Wörgötter,
A. Ude, T. Asfour, D. Kraft, D. Omrčen, A. Agostini, R. Dillmann, Ob-
ject–Action Complexes: Grounded abstractions of sensory–motor pro-
cesses, *Robotics and Autonomous Systems* 59 (2011) 740–757. doi:
10.1016/j.robot.2011.05.009.
- 1240 [6] N. Vahrenkamp, M. Wächter, M. Kröhnert, K. Welke, T. Asfour, The
robot software framework ArmarX, *it - Information Technology* 57 (2015).
doi:10.1515/itit-2014-1066.
- [7] L. R. Squire, Memory systems of the brain: A brief history and current
perspective, *Neurobiology of Learning and Memory* 82 (2004). doi:10.
1016/j.nlm.2004.06.005.
- 1245 [8] A. J. McComas, Clive Wearing and Henry Molaison Reconsidered, in:
Aranzio’s Seahorse and the Search for Memory and Consciousness, Ox-
ford University Press, 2022, Ch. 38, pp. 261–266. doi:10.1093/oso/
9780192868244.003.0040.
- 1250 [9] R. C. Atkinson, R. M. Shiffrin, Human memory: A proposed system and its
control processes., *The Psychology of learning and motivation: Advances in
research and theory* 2 (1968) pp. 89–105.

- [10] B. Aben, S. Stapert, A. Blokland, About the distinction between working memory and short-term memory, *Frontiers in Psychology* 3 (2012). doi: 10.3389/fpsyg.2012.00301.
- 1255 [11] C. Wan, P. Cai, M. Wang, Y. Qian, W. Huang, X. Chen, Artificial Sensory Memory, *Advanced Materials* 32 (2020).
- [12] N. Cowan, T. Alloway, The development of working memory, in: *The development of memory in childhood*, Psychology Press, 1997, Ch. 7, pp. 163–199.
- 1260 [13] N. Cowan, The magical number 4 in short-term memory: A reconsideration of mental storage capacity, *Behavioral and Brain Sciences* 24 (2001) pp. 87–114. doi:10.1017/S0140525X01003922.
- [14] O. Barak, M. Tsodyks, Working models of working memory, *Current Opinion in Neurobiology* 25 (2014) 20–24. doi:10.1016/j.conb.2013.10.008.
- 1265 [15] G. A. Miller, The magical number seven, plus or minus two: Some limits on our capacity for processing information, *The Psychological Review* 63 (1956) pp. 81–97. doi:10.1037/h0043158.
- [16] E. Service, The Effect of Word Length on Immediate Serial Recall Depends on Phonological Complexity, Not Articulatory Duration, *The Quarterly Journal of Experimental Psychology* 51 (1998) 283–304. doi:10.1080/713755759.
- 1270 [17] F. R. Gobet, Some shortcomings of long-term working memory, *British journal of psychology* 91 (2000) pp. 551–70.
- [18] A. D. Baddeley, G. Hitch, Working memory, in: *Psychology of learning and motivation*, Vol. 8, Academic Press, 1974, pp. 47–89.
- 1275 [19] J. C. Taylor, A dynamic model of memory for research on human information processing, *Instructional Science* 12 (1983) pp. 367–374.
- [20] I. Kotseruba, J. K. Tsotsos, 40 years of cognitive architectures: core cognitive abilities and practical applications, *Artificial Intelligence Review* (2018). doi:10.1007/s10462-018-9646-y.
- 1280 [21] K. S. Graham, J. S. Simons, K. H. Pratt, K. Patterson, J. R. Hodges, Insights from semantic dementia on the relationship between episodic and

- semantic memory, *Neuropsychologia* 38 (2000) pp. 313–324. doi:10.1016/S0028-3932(99)00073-1.
- 1285 [22] D. Coon, J. Mitterer, T. Martini, *Introduction to Psychology: Gateways to Mind and Behavior*, Cengage Learning, 2018.
- [23] A. Ioannou, X. Anastassiou-Hadjicharalambous, Non-associative Learning, in: *Encyclopedia of Evolutionary Psychological Science*, Springer International Publishing, Cham, 2018, p. 5419–5432. doi:10.1007/978-3-319-16999-6_1027-1.
- 1290 [24] R. Novianto, M.-A. Williams, P. Gärdenfors, G. Wightwick, Classical Conditioning in Social Robots, in: *International Conference on Social Robotics*, 2014, pp. 279–289. doi:10.1007/978-3-319-11973-1_29.
- [25] J. Anderson, D. Bothell, M. Byrne, S. Douglass, C. Lebiere, Y. Qin, An integrated theory of the mind, *Psychological Review* 111 (2004) pp. 1036–1060. doi:10.1037/0033-295X.111.4.1036.
- 1295 [26] D. Vernon, *Cognitive Architectures*, in: *Cognitive Robotics*, MIT Press, 2022, Ch. 10, pp. 191–212. doi:10.7551/mitpress/13780.003.0015.
- [27] D. Vernon, C. Von Hofsten, L. Fadiga, *A Roadmap for Cognitive Development in Humanoid Robots*, Vol. 11, Springer Science & Business Media, 2011.
- 1300 [28] R. Sun, Desiderata for cognitive architectures, *Philosophical Psychology* 17 (2004). doi:10.1080/0951508042000286721.
- [29] W. Duch, R. J. Oentaryo, M. Pasquier, Cognitive architectures: Where do we go from here?, in: *Artificial General Intelligence*, Vol. 171, 2008, pp. 122–136.
- 1305 [30] D. E. Kieras, D. E. Meyer, An Overview of the EPIC Architecture for Cognition and Performance With Application to Human-Computer Interaction, *Human-Computer Interaction* 12 (1997) pp. 391–438. doi:10.1207/s15327051hci1204_4.
- 1310 [31] P. Langley, D. Choi, A Unified Cognitive Architecture for Physical Agents, in: *Proceedings of the National Conference on Artificial Intelligence*, Vol. 21, Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2006, pp. 1469–1474.

- 1315 [32] J. Weng, A Theory for Mentally Developing Robots, in: Proceedings 2nd International Conference on Development and Learning, IEEE, 2002, pp. 131–140. doi:10.1109/DEVLRN.2002.1011821.
- [33] J. Rothfuss, F. Ferreira, E. E. Aksoy, Y. Zhou, T. Asfour, Deep Episodic Memory: Encoding, Recalling, and Predicting Episodic Experiences for Robot Action Execution, IEEE Robotics and Automation Letters 3 (2018) pp. 4007–4014. doi:10.1109/LRA.2018.2860057.
- 1320 [34] D. P. Kingma, M. Welling, Auto-Encoding Variational Bayes (2013). doi:10.48550/ARXIV.1312.6114.
- [35] L. Bärmann, F. Peller-Konrad, S. Constantin, T. Asfour, A. Waibel, Deep Episodic Memory for Verbalization of Robot Experience, IEEE Robotics and Automation Letters 6 (2021) pp. 5808–5815. doi:10.1109/LRA.2021.3085166.
- 1325 [36] J. E. Laird, A. Newell, P. S. Rosenbloom, SOAR: An Architecture for General Intelligence, Artificial Intelligence 33 (1987). doi:10.1016/0004-3702(87)90050-6.
- 1330 [37] J. E. Laird, The Soar Cognitive Architecture, The MIT Press, 2012.
- [38] J. E. Laird, Extending the Soar Cognitive Architecture, Frontiers in Artificial Intelligence and Applications 171 (2008) pp. 224–235.
- [39] K. Kawamura, S. M. Gordon, P. Ratanaswasd, E. Erdemir, J. F. Hall, Implementation of Cognitive Control for a Humanoid Robot, International Journal of Humanoid Robotics 5 (2008) pp. 547–586. doi:10.1142/S0219843608001558.
- 1335 [40] R. A. Peters, K. E. Hambuchen, K. Kawamura, D. M. Wilkes, The Sensory Ego-Sphere as a Short-Term Memory for Humanoids, in: Proceedings of the IEEE-RAS International Conference on Humanoid Robots (HUMANOIDS), 2001, pp. 451–459.
- 1340 [41] S. Franklin, T. Madl, S. D’mello, J. Snider, LIDA: A Systems-level Architecture for Cognition, Emotion, and Learning, IEEE Transactions on Autonomous Mental Development 6 (2013) pp. 19–41. doi:10.1109/TAMD.2013.2277589.
- 1345

- [42] M. Beetz, L. Mösenlechner, M. Tenorth, CRAM—A Cognitive Robot Abstract Machine for everyday manipulation in human environments, in: IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE, 2010, pp. 1012–1017. doi:10.1109/IR0S.2010.5650146.
- 1350 [43] M. Beetz, D. Beßler, A. Haidu, M. Pomarlan, A. K. Bozcuoğlu, G. Bartels, Know Rob 2.0 — A 2nd Generation Knowledge Processing Framework for Cognition-Enabled Robotic Agents, in: IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2018, pp. 512–519. doi:10.1109/ICRA.2018.8460964.
- 1355 [44] M. Wächter, E. Ovchinnikova, V. Wittenbeck, P. Kaiser, S. Szedmak, W. Mustafa, D. Kraft, N. Krüger, J. Piater, T. Asfour, Integrating multi-purpose natural language understanding, robot’s memory, and symbolic planning for task execution in humanoid robots, *Robotics and Autonomous Systems* 99 (2018) pp. 148–165. doi:https://doi.org/10.1016/j.robot.2017.10.012.
- 1360 [45] E. Tulving, Episodic and semantic memory, in: *Organization of Memory*, Academic Press, 1972, pp. 381–403.
- [46] E. Tulving, Précis of Elements of episodic memory, *Behavioral and Brain Sciences* 7 (1984) pp. 223–238. doi:10.1017/S0140525X0004440X.
- 1365 [47] A. V. Samsonovich, Emotional biologically inspired cognitive architecture, *Biologically Inspired Cognitive Architectures* 6 (2013) pp. 109–125. doi:10.1016/j.bica.2013.07.009.
- [48] D. Paulius, Y. Sun, A Survey of Knowledge Representation in Service Robotics, *Robotics and Autonomous Systems* 118 (2019). doi:10.1016/j.robot.2019.03.005.
- 1370 [49] M. Al-Ani, F. Awad, The JPEG image compression algorithm, *International Journal of Advances in Engineering & Technology* 6 (2013) pp. 1055–1062.
- [50] F. C. Pereira, F. M. B. Pereira, F. C. Pereira, F. Pereira, T. Ebrahimi, *The MPEG-4 book*, Prentice Hall Professional, 2002.
- 1375 [51] D. Huffman, A Method for the Construction of Minimum-Redundancy Codes, *Resonance* 11 (2006) pp. 91–99. doi:10.1007/BF02837279.

- [52] Blender Online Community, Blender - a 3D modelling and rendering package (2022).
- [53] M. Henning, A New Approach to Object-Oriented Middleware, *IEEE Internet Computing* 8 (2004) pp. 66–75. doi:10.1109/MIC.2004.1260706.
- [54] S. Macenski, T. Foote, B. Gerkey, C. Lalancette, W. Woodall, Robot Operating System 2: Design, architecture, and uses in the wild, *Science Robotics* 7 (2022). doi:10.1126/scirobotics.abm6074.
- [55] M. Bhardwaj, B. Sundaralingam, A. Mousavian, N. D. Ratliff, D. Fox, F. Ramos, B. Boots, STORM: An Integrated Framework for Fast Joint-Space Model-Predictive Control for Reactive Manipulation, in: *Conference on Robot Learning*, PMLR, 2022, pp. 750–759. doi:10.48550/arXiv.2104.13542.
- [56] M. Mathieu, C. Couprie, Y. LeCun, Deep multi-scale video prediction beyond mean square error, *arXiv preprint* (2015). doi:10.48550/arXiv.1511.05440.
- [57] L. Xue, N. Constant, A. Roberts, M. Kale, R. Al-Rfou, A. Siddhant, A. Barua, C. Raffel, mT5: A massively multilingual pre-trained text-to-text transformer, *arXiv preprint* (2020). doi:10.48550/arXiv.2010.11934.
- [58] E. Dere, M. D. S. Silva, J. P. Huston, Higher Order Memories for Objects Encountered in Different Spatio-temporal Contexts in Mice: Evidence for Episodic Memory, *Reviews in the Neurosciences* 15 (2004) pp. 231–240. doi:10.1515/REVNEURO.2004.15.4.231.
- [59] M. Waibel, M. Beetz, J. Civera, R. d’Andrea, J. Elfring, D. Galvez-Lopez, K. Häussermann, R. Janssen, J. Montiel, A. Perzylo, et al., RoboEarth, *IEEE Robotics & Automation Magazine* 18 (2011) pp. 69–82. doi:10.1109/MRA.2011.941632.
- [60] S. Bayreuther, F. Jacob, M. Grotz, R. Kartmann, F. Peller-Konrad, F. Paus, H. Hartenstein, T. Asfour, Combining Task Planning and Activity-Centric Access Control for Assistive Humanoid Robots, in: *Symposium on Access Control Models and Technologies (SACMAT)*, 2022, pp. 185–194. doi:10.1145/3532105.3535018.