Memory-Centred Architectures: Perspectives on Human-level Cognitive Competencies

Paul Baxter, Rachel Wood, Anthony Morse, Tony Belpaeme

Centre for Robotics and Neural Systems
University of Plymouth, U.K.
{paul.baxter, rachel.wood, anthony.morse, tony.belpaeme}@plymouth.ac.uk

Abstract

In the context of cognitive architectures, memory is typically considered as a passive storage device with the sole purpose of maintaining and retrieving information relevant to ongoing cognitive processing. If memory is instead considered to be a fundamentally active aspect of cognition, as increasingly suggested by empirically-derived neurophysiological theory, this passive role must be reinterpreted. In this perspective, memory is the distributed substrate of cognition, forming the foundation for cross-modal priming, and hence soft cross-modal coordination. This paper seeks to describe what a cognitive architecture based on this perspective must involve, and begins to explore how human-level cognitive competencies (namely episodic memory, word label conjunction learning, and social behaviour) can be accounted for in such a low-level framework. This proposal of a memory-centred cognitive architecture presents new insights into the nature of cognition, with benefits for computational implementations such as generality and robustness that have only begun to be exploited.

Introduction

That memory is important for cognition, and cognitive competencies, is an uncontroversial assertion. Memory is in this context generally considered to be a storage mechanism that is connected to, but distinct from, cognitive processing. This feature is prevalent in the study and implementation of contemporary cognitive architectures, and reflects a perspective on human cognition implicitly based on the 'cognition as computation' metaphor (Bickhard 2009; Miller 2003). However, developments in empirically-based neuropsychological theory suggest that this assumption, and the corrollaries informed by it, do not hold for biological systems; rather, memory should be considered as fundamental and central to cognition. Our research explores the consequences of this perspective on the study of cognitive architecture, and applies the resulting principles to synthetic cognitive systems. These principles, being informed by biological theory and applied in a computational context, are as such 'low-level', emphasising a bottom-up approach to cognitive system design. The purpose of this paper is therefore

to demonstrate how such fundamental principles of cognitive system design and operation allow the examination (in principle) of what are generally considered to be high- or human-level cognitive competencies.

In the most general terms, memory may be described in terms of three mechanisms: the *encoding* and *storage* of previously acquired information, and subsequent *retrieval* of information in the service of some current or ongoing task (Neath and Surprenant 2003). While this definition is very broad, it does incorporate a number of assumptions regarding the nature and structure of memory and cognitive processing, namely, it implies that memory is passive storage, with computation structurally separate and requiring information from this storage device. This paradigm of memory organisation in relation to cognition is prevalent both in general cognitive science and in synthetic cognitive architecture approaches, and can be seen as an extension of the prevalent (but often only implicitly present) mind-as-computer metaphor. However, in terms of neuropsychological evidence, this strict separation does not appear to be fully justified, with an emphasis instead on distributed associative structure (i.e. memory) as the substrate for activation dynamics (i.e. cognition) (Fuster 2000; Postle 2006).

Such a reinterpretation of what memory is, and what its relationship with cognition entails, leads to a necessary re-examination of the role of memory in synthetic cognitive systems. The starting point in the process of designing a cognitive system thus becomes a consideration of what memory is, rather than what is required of memory (Wood, Baxter, and Belpaeme 2010). With such a focus on memory and its mechanisms in their own right, rather than memory as required for a given task environment or architecture, there is no inherent necessity to incorporate application-specific functionality into the theory of operation, leading to a potentially more generalisable computational implementation, and one more robust to changes in application context. In embarking on such a process of cognitive memory reinterpretation, it is necessary to start from low-level principles. Elements of this are indicated by the biologically-derived theory (see next section), which describe fundamental mechanisms of memory: an attractive source of inspiration given that biological systems display the type of robust, general and adaptive behaviour that synthetic cognitive systems are

attempting to attain.

Based on this reinterpretation, we have proposed a number of principles and mechanisms that characterise the role of memory for cognition (Baxter and Browne 2011; Wood, Baxter, and Belpaeme in press; Baxter 2010; Baxter and Browne 2010; Wood, Baxter, and Belpaeme 2010). These are 'low-level' relative to what are normally considered to be cognitive competencies, encompassing the principle of fundamental associationism, and functions such as activation flow, and cross-modal priming (see section 3 for details). The central theme though, is that *cognition should be considered as activation dynamics over the substrate of associative memory*, and that a commitment to this perspective requires a fundamental shift in the manner that synthetic cognitive architecture implementations should be approached. In altering the approach in this low-level manner, it becomes necessary to provide an explanation for how higher-level cognitive competencies can be accounted for. This paper seeks to provide the first steps in answering this question, by proposing that if a number of assumptions regarding the standard interpretation of these higher-level competencies are themselves re-evaluated, and put in the context of the wider cognitive system, then a consistent and coherent account can be formulated based on the memory-centred cognition perspective.

The remainder of this paper is structured as follows. Firstly, we provide an overview of the use of memory in synthetic cognitive systems, and the interpretation of memory structure and function in biological systems (section 2). After a description of the foundations of the memory-centred cognitive architecture perspective, and how it affords lines of enquiry not facilitated by currently prevalent approaches (section 3), three case studies of how such mechanisms can account for functionality approaching that which may be considered human-level are described (section 4). While these case studies report research in varying stages of completion, they provide support for the notion that the 'gap' between low- and high-level cognitive competencies, in terms of both function and structure, may not be as large as it may first appear (section 5).

Memory in Synthetic and Biological Systems

Four characteristics may be distinguished from current memory system implementations. Firstly, synthetic implementations of memory are either modular themselves, or part of an inherently modular architecture. This characteristic is founded in the computational interpretation of mind (where computation and storage are separated), and human memory has been decomposed into functional types (see below). In synthetic cognitive systems, this characteristic is clear, with Soar (Newell 1990) and ACT-R (Anderson et al. 2004) providing notable examples: indeed, this separation of computation and information storage has even been proposed as a theoretical necessity (Sun 2004). The second and third characteristics are related to this point. The second one is that memory is treated as a passive storage structure, in which information may be placed, remaining static during the storage period, for recall by a specific function at some point in the future. The third characteristic is

that the format that the information that is stored in memory takes is typically based on a global (i.e. system-wide) ontology, which ensures that information is in a common format that can be handled by all of the computational (or cognitive) functions. Memory thus contains symbolic information in a state removed/abstracted away from the low-level sensory and/or motor information. These characteristics are a consequence of the physical symbol system hypothesis (Newell and Simon 1976), which states that a physical symbol system has the necessary and sufficient means for intelligent action: cognitive architectures therefore are typically symbol processing systems, see (Goertzel et al. 2010; Langley, Laird, and Rogers 2009; Sun 2004) for general reviews. Finally, the fourth characteristic is that with memory as a purely passive storage structure, specific computational structures have to be defined to handle the information flow to and from memory, with centralised workspaces being a typically used structure, such as CAS (Wyatt and Hawes 2008), ISAC (Kawamura et al. 2004), EPIC (Kieras and Meyer 1997), and Soar (Young and Lewis 1999).

In this fourth characteristic, the influence from psychological models is clear. In addition to the memory/cognition divide, there is a further fractionation of memory into functional sub-systems - in particular a separation of long- from short-term memory, and then a further separation of long-term memory. The distinction between short- and long-term memory is made on purely temporal grounds: short-term memory holds information for a finite length of time (usually commensurate with that of the context to which this information is related), whereas long-term memory holds information indefinitely (with processes of forgetting playing a role here). The characterisation of short-term memory as Working Memory (WM) has been particularly influential (Baddeley and Hitch 1974; Cowan 1999), forming the justification of workspace-like constructs in cognitive architectures. WM, in this context, takes on the role of interface between memory (as passive storage), and cognition (as computation).

In contrast with this characterisation of WM as a functionally, and even structurally, distinct system, a number of neuroscientifically-based theories have arisen emphasising distributed organisation and function (Postle 2006; Fuster 1997; Bar 2007; Glenberg 1997). In proposing an epistemological shift away from 'systems of memory' to 'the memory of systems' (Fuster 1997) proposes that memory function be understood as a distributed property of the same cortical systems that underpin perception and action. This view is founded on the principle that memory consists of the modulation of synaptic contacts within distributed networks of interconnected cortical cells. Memories are formed by selective facilitation and elimination of synaptic links between neuronal aggregates which activate in response to discrete features of their environment (both within and without the body of the agent) and thus are inherently associative: the information they contain consisting of neuronal relationships. Network memories are constructed via Hebbian synchronous convergence wherein co-activation leads to association via summation of temporally coincident inputs and Long-Term Potentiation for reshaping

existing networks. These associative processes produce interconnected, functional units of memory (i.e. Hebbian cell assemblies).

Fuster extends this principle to propose that that the strict dichotomisation of long and short term memory is unnecessary, with working memory realised through temporary, ad hoc activation of perceptual and motor memory networks. These 'active' memory networks can be distinguished from passive (long-term) memory in terms of state of a network rather than cortical distribution. In this network-based view, the formation and reactivation of memory depends on the association and activation of distributed networks located in the posterior cortex. If the content of memory is associated with action then activation propagates, spreading forward to pre-frontal cortex. The central principle of the network memory approach is that memory is a property of the systems it serves and therefore inseparable from them. Thus the sensory-motor dichotomy in memory localisation is useful only to the extent that neural systems for sensing and acting can themselves be separated.

Given this alternative neuroscientifically founded description of memory, it may be seen how many computational implementations of learning mechanisms, for example, may be reinterpreted as being memory in this context. For example, systems in which transient activation dynamics operate on network structures, such as in the cognitive architecture ACT-R (Anderson et al. 2004), in systems such as semantic networks (Rumelhart and Norman 1973; Sowa 1992), and also, more generally, artificial neural networks (see (Haykin 1999) for an overview). In these examples, the notion of binary membership of information to a workspace no longer applies - rather, it is the level of activation that determines the degree to which a given piece of information (be it in discrete or distributed form) participates in the ongoing processing. Given the broad definition of memory discussed above, we propose that these types of system, in the context of cognitive architecture, should also be considered as memory.

The Memory-Centred Architecture Perspective

The four characteristics of memory system implementations in synthetic cognitive systems described above highlight that the function typically defines the mechanism. Given the general acceptance of the encode-store-recall paradigm, the contents of the memory system (the structure of the information stored) are thus defined by the task that it is to be engaged in. In the majority of cases, this results in the use of a symbolic representation scheme, with discrete symbols as 'chunks' of information being manipulated in accordance with the mind-as-computer paradigm (cf. the physical symbol system hypothesis (Newell and Simon 1976)). While there are some exceptions to this characterisation which do begin to approach certain aspects of low-level biological mechanisms, these systems are typically not embedded into general cognitive architectures, and thus can not inform these questions of theory directly: rather, they provide examples of how such mechanisms may be implemented.

1. Memory and Cognition are functionally integrated
2. Memory is an active process, not a passive storage device
3. Memory is fundamentally associative and inherently amodal
4. Memory as an informal coordinator of information across and within modalities

Table 1: A summary of the four proposed properties of Memory.

By instead committing to memory-centred view, memory can be considered without having to define the task context. A focus on mechanism can be applied, enabling a closer integration with the neuroscientific view of memory described above than would otherwise be possible. Four properties of memory can thus be defined, which may be seen as responses to the four characteristics of currently implemented memory systems (Table 1). Firstly, memory is a system that is functionally integrated with cognition, but which is structurally distributed, rather than one which is inherently modular and separates memory and cognition. Secondly, memory is an active process, and not a passive storage device. Closely related to this is the third property: memory is fundamentally associative ad inherently amodal. The explicit storage of symbolic information therefore does not arise, resulting in there being no requirement for a global ontology for information transfer and coordination. Associations are inherently amodal, even if the information associated is modality-specific. As a result of this characteristic, the fourth property (although this borders on the definition of a mechanism) is that memory is an informal coordinator of information both within and across modalities, since it does not provide explicit symbolic storage, by means of activation flow on the substrate of acquired associations. These four properties are task/application independent, since they have been defined without specifying the tasks for which they should be used. They are as a consequence low-level, as they do not specify what the 'behaviour' of the memory system as a whole should exhibit.

On the basis of these four proposed properties, a mechanism may be defined, which describes a fundamental function of the cognitive system in which the memory system operates: priming (be it intra- or cross-modal). Given the proposed lack of a global ontology by which modal-specific information is integrated, the associative structure of memory serves as a coordinator of information across modalities rather than a simple storage device. Further, given that there is no necessity for (symbolic) semantic information to pass through the memory system, cross-modal coordination is proposed to occur through the propagation of activation through the substrate of associative memory. Thus the use of information in one modality (be it recognition in a sensory modality, or the execution of a sequence of actions in

a motor modality) gives rise to activation that may propagate from one modality to another via the associative memory network. The significance of memory as fundamentally associative in structure and mechanism thus becomes clear: associations can only form in terms of a history of activation profiles, which then become the substrate for future cross-modal coordination. Therefore, previously acquired associations (based on a history of perception-action activity) modulate the control of ongoing and future behaviour, i.e. memory as an active system embedded within a wider cognitive system, and fully integrated with behaviour generation.

Towards Human-level Cognitive Competencies: case studies

As described above, the memory-centred architecture perspective is currently focussed primarily on low-level mechanisms, upon which cognitive architectures may be be implemented. However, in order to be considered as a potential account of cognition in synthetic systems, it must be shown how these principles and mechanisms can lead to the sort of behavioural functionality typically associated with cognitive competencies, specifically human-level cognitive competencies. In this section we examine three case studies, each of which describe the application of the memory-centred perspective to higher-level cognitive abilities (if not necessarily human-level): (1) accounting for apparent functional specialisations of memory, in this case episodic memory; (2) cross-modal coordination for social interaction; and (3) the development of capabilities underlying language, namely label use and compositionality. Whilst these case studies are currently the subject on ongoing research, we contend that they justify the applicability of the central principle (memory-centred cognition) to not only low-level functionality, but also to high-level cognitive competencies.

In order to address this divide between low- and high-level functionality, it is necessary to review the current approach to memory implementations. With the separate consideration of what memory does and what memory is, the emphasis typically lies on the replication of memory function, which is in the context of a given task. There are two issues to this approach: (1) there is a difficulty in hand-designing a process without resulting in a brittle and sub-optimal system; and (2) this approach to the definition of memory function is typically based on a task-centred analysis, which does not necessarily reflect the entirety of a given memory function, and which therefore necessitates the integration of data from a wide range of sources. These issues become more acute as the functions are considered in largely behavioural terms, since increasing discontinuities arise in relation to neural mechanisms.

In order to reconcile these two aspects, the approach of synthetic neuroethology may be applied: a methodology that examines the production of adaptive behaviour as resultant from neural mechanisms in the context of embodiment and environmental interaction (Beer 1990; Cliff 1991). This forces a consideration of the context in which memory operates, embedding it within cognition, in the context of perception-action cycles pertaining to an embodied agent.

This expansion of focus regarding the interpretation of memory enables an operationalisation of memory function, as it forces a consideration of what memory actually *is*, in addition to what it *does*. In this manner, the focus on memory rather than task context enables a reinterpretation of higher-level cognitive competencies in terms of how the proposed mechanisms of memory can support these functions, rather than consideration of what properties of memory are required for each task. This is the methodology applied to the following case studies; we show how this results in a closing of the apparent gap between proposed low-level memory operational principles, and higher-level cognitive competencies.

Episodic Memory

One feature of the psychological view of memory is the fractionation of memory into functionally-specific modules, which may be considered as separable systems. One such sub-memory system is Episodic memory, which is in general terms the capacity to recall previous experiences. It is widely held to be a uniquely human capability because it entails the conscious experience of remembering and an element of self-recollection. In current cognitive systems, episodic memory has been treated as a functional sub-system of memory, typically handling information with a given task-centred structure, e.g. (Nuxoll 2007; Kawamura et al. 2004; Ho and Watson 2006; Brom and Lukavsky 2008). Given these assumptions, it would be difficult to provide a low-level, memory-centred mechanistic account of episodic memory. Given the re-evaluation of memory of the type proposed in this paper, it is necessary to provide an account for this apparent functional fractionation on the basis of the proposed properties and mechanisms of memory.
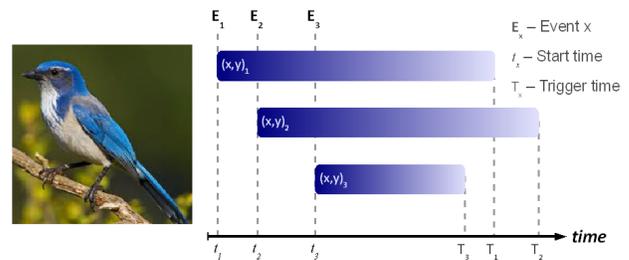

Figure 1: An operational account of episodic memory that does not require a role for conscious recollection, based on the food caching behaviours of corvids: the behaviour entails the maintenance of what/where information in combination with a temporal component, here represented by a simple decay, the time constant of which may have been learned previously. The caching event (E_x) starts a decay function (at time t_x), which when expired, triggers a food retrieval event (at time T_x).

As there are no accepted, non-verbal, behavioural markers for conscious experience in non-human species their capacity for conscious episodic recall cannot be tested. If, how-

ever, the explicit requirement for conscious experience of recall is set aside, then it is apparent that a number of non-human animals exhibit behaviour consistent with the presence of an episodic memory. Tulving's original definition characterises episodic memory as storing temporally organised events in such a way as to also preserve their temporal-spatial relations (Tulving 1972; 2002). So here, episodic memory consists of the recollection of linked what, where and when information, the unique characteristic of this type of memory being the explicit temporal and spatial linking of events, rather than a conscious awareness of their having been personally experienced at some point in the past (Clayton et al. 2001). This functionality can be observed in the food caching behaviour of some birds (Clayton and Dickinson 1998), where items of food with differing decay rates (*what*) are cached over a period of time (*when*) in dispersed locations (*where*) for retrieval at a later time, but before the item has decayed. Thus, by taking an ethological approach we can generate an operational concept of episodic memory (i.e. a systematic notion of what it does) and identify minimal candidate mechanisms for the functions required of what becomes a simple 'what-when-where' system.

The original problem is thus transformed from one requiring specialist functional mechanisms, to one that the low-level mechanisms of the memory-centred approach can account for (Wood, Baxter, and Belpaeme 2010). The application of this principle allows a minimal model of episodic memory as observed in food caching to be formulated (Figure 1). In this case, multiple temporal sequences are a candidate for modelling the item-related decay rate, using neural-inspired mechanisms such as those proposed in (Yamashita and Tani 2008) for example, triggering a retrieval when the temporal sequence has run its course.

Social Interaction

Another cognitive competence that is considered high-level is the capacity for social interaction: whilst there are a number of non-human species that engage in social interaction, the complex, multi-modal and multi-timescale nature of human social interaction is undoubtedly at the pinnacle of cognitive functionality. In terms of synthetic cognitive architecture, an increasingly important research goal has been the development of companion agents (be they robots or simulated avatars) capable of naturalistic social interaction (Fong, Nourbakhsh, and Dautenhahn 2003). It has been proposed that required for this is the adaptive coordination of multiple modalities (based on a history of interactions), such that coherent multi-modal behaviour, and thus engaging interaction, results (Wood, Baxter, and Belpaeme 2010; Baxter et al. 2011).

This competence is thus the subject of active research, not only in terms of human-robot interaction, but also in terms of the more general architectural mechanisms required for the coordination of multiple modalities in the service of coherent system behaviour. In the context of human-robot interaction, this is an issue that has only recently gained attention (Baxter et al. 2011). In cognitive architectures this is typically currently achieved through the use of a global ontology to synchronise and exchange information across

modalities, such as those used in Polyscheme, e.g. (Cassimatis et al. 2004), and CAS, e.g. (Hawes et al. 2007; Wyatt and Hawes 2008). A drawback of such an approach however is the potential brittleness in such symbolic representations with respect to both task environment (i.e. robustness to change in task), and architecture structure (i.e. robustness to change in architectural structure, although this may be more applicable to the design rather than execution stage).
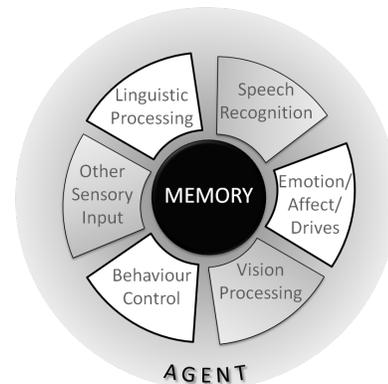

Figure 2: In the context of a cognitive architecture with multiple and discrete cognitive modalities, some main examples of which are shown, the memory system performs cross-modal priming on the basis of associations gained through prior experience, without requiring the explicit processing of symbolic information. Memory may thus mediate between the modalities without directly encoding any modality specific information.

In the context of a memory-centred architecture scheme (Figure 2), cross-modal coordination may be achieved through the activation-based priming mechanism. Since this occurs on the substrate of associations, a necessary part of this process is the formation of these associations between modalities through experience - hence memory (prior experience) in the service of an ongoing task (a current interaction). In this account, there is no requirement for the memory itself to encode symbolic information, rather its formed associations facilitate the priming of information in other modalities: indeed, the memory system itself does not necessarily require any information regarding what information is being associated/coordinated, merely that this process is occurring and ongoing. While this does not provide a complete account of how 'meaningful' behaviour can be generated, it does illustrate the process of cross-modal information itself: the following case study demonstrates how the incorporation of the context of the system as a whole (specifically using the example of embodiment) enables such meaningful behaviour to be produced.

Proto-Language development

Whilst language is a cognitive competence unique to humans[1], what may be described as its cognitive predecessors are frequently overlooked: notably the use of word labels to scaffold cognitive processing (Clowes and Morse 2005). It has been established that the learning of labels for objects is mediated by body posture - in a robotic model (Morse et al. 2010a), and in infants (Smith and Samuelson 2010) - where changes in body posture prime different representations (an example of association formation providing the substrate for cross-modal priming). Based on the proposal that learned labels can be subsequently used to extend cognitive capabilities through scaffolding (Clark 2008), an extension to the label learning setup was proposed, where learned labels enabled an overlapping categorisation task to be completed, which could not be completed if the labels had not been first learned (Morse et al. 2011). Cognition is therefore extended by this use of labels, the learning of which is conducted on a purely associative basis.

The computational model implemented to explore these questions uses Self-Organising Maps (SOMs) for each modality used connected by Hebbian-like associative links to one central SOM representing the posture of the agent (Figure 3). The associative structures are driven by activation from the sensory SOMs, which subsequently forms the substrate of an activation-spreading network (Morse et al. 2010a). On this basis, the system may learn to associate information from multiple modalities, thus enabling for example the learning of labels of objects (which are themselves conjunctions of visual features, such as colour and shape).
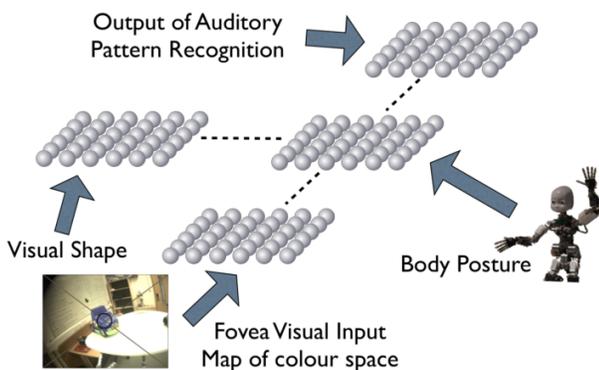

Figure 3: The label learning model architecture: self-organising maps are used to map the colour, shape, word and body spaces, with associative links formed (in a Hebbian-like manner) between peaks of activation in each map.

This work demonstrates a number of aspects relevant to the memory-centred cognitive architecture perspective. Firstly, it shows how fundamentally associative mechanisms may be used in the acquisition and subsequent support of complex cognitive capabilities. Secondly, it demonstrates

---

[1]Displaying the properties of compositionality and complex syntax not present in non-human animal communication schemes.

the necessity for considering memory in the context of the cognitive system as a whole, as indicated by the application of the methods of synthetic neuroethology - specifically embodiment in this case. Finally, it also provides an example of how the low-level mechanisms of association and priming are implemented in a manner that is task independent, since the computational architecture used for this case study has been applied to a range of cognitive tasks (Morse, Lowe, and Ziemke 2009; Morse et al. 2010b).

Discussion

The case studies above provide three examples of higher-level cognition, and indicates how they can be accounted for in a theoretical framework (the memory-centred approach) that is inherently low-level. In the case of episodic memory, it was argued that the observable behaviour corresponding with the functional role of this type of memory can be accounted for with non-specialised, low-level mechanisms. While this case dealt with episodic memory, it is hypothesised (and the subject of ongoing investigation) that other functionally distinct types of memory can be accounted for in a similar manner. In the case of social interaction, and specifically the hypothesised central role for cross-modal coordination, it has been described how the low-level mechanisms that the memory-centred architecture perspective provides can account for this functionality without having to rely on a task- or environment-specific symbolic encoding scheme. Finally, in the case of label use for cognitive scaffolding, it has been shown how an associative substrate (that may be described as memory) with activation dynamics can account for what may be described as complex cognitive processing, without a requirement for central executive control, as a result of embodiment-mediated coordination.

The four general characteristics of current memory system implementations identified above were modularity, passive storage, use of a global ontology, and controlled information flow, and are based upon psychological models, and resulting implicit assumptions regarding the nature of memory. The memory-centred architecture perspective challenges each of these characteristics, resulting in the definition of low-level principles of operation, with corresponding mechanisms. As shown above, the case studies variously support these principles, and so justify the validity of the perspective, and its applicability to the implementation of synthetic cognitive architectures.

It must however be acknowledged that these case studies alone do not provide complete support for the application of the memory-centred architecture perspective to synthetic cognitive architecture, particularly given that they are the subject of ongoing examination. It is still an open question, for example, how this memory-centred account may be reconciled with decision making and planning capabilities (although basic goal-directed behaviours have been demonstrated in this context (Baxter and Browne 2010)). However, they do demonstrate, as a proof of concept, a reconciliation between low-level mechanisms and high-level behavioural competences. These case studies are also indicative of the manner in which this perspective can be applied to issues of human cognition to yield an alternative insight that is not

easily afforded by existing approaches: a task-independent characterisation of memory, and its application to cognitive systems.

There are a number of advantages provided by this memory-centred architecture perspective over the currently prevalent view of memory. Notably among these is that memory is here considered as independent from a specific task: the forming of associations, activation dynamics on this associative substrate, and cross-modal priming will occur irrespective of the task. This affords an active model of memory with a generality of application that is not generally present in existing memory systems. Additionally, because this memory does not deal explicitly with symbolic structures for information coordination, there is an added robustness against task reallocation and potentially also partial system failure.

Finally, the proposal that memory forms the substrate of cognitive processing, and must therefore be considered in the context of the system as a whole (embodiment, the perception-action cycle, etc), means that the memory-centred architecture perspective is readily placed in a developmental framework: the mechanisms of learning and adaptation of the memory system constitute the developmental progress of the system as a whole (Wood, Baxter, and Belpaeme 2010; Baxter and Browne 2010; Baxter 2010). Given that for humans, higher-level cognitive competencies emerge only after a period of prolonged development, the memory-centred perspective can thus readily leverage what is known about human development for further definition and refinement. Indeed, in so doing, there is also the potential for reciprocal knowledge transfer in terms of the further understanding of human cognitive ontogeny (Morse et al. 2010a).

Conclusion

The examination of memory is essential to the study and implementation of cognitive systems. Memory is functionally an active process, structurally the substrate of cognition, and acts in the service of ongoing adaptive behaviour based on previously acquired information. In this paper, we propose that in reinterpreting the role of memory with respect to cognition, it is necessary to also reconsider the role of memory systems in synthetic cognitive architectures. This results in a focus on what memory *is* rather than simply what memory *does*, and leads to an emphasis on low-level mechanism, rather than high-level function. However, it has been shown that high-level cognitive competencies can still be accounted for in this account. While this paper represents only an initial step, this demonstration of an integration between low-level mechanisms and high-level capabilities indicates that there is not such a gap between the two levels of analysis as may be initially apparent.

Acknowledgments

References

Anderson, J. R.; Bothell, D.; Byrne, M. D.; Douglass, S.; Lebiere, C.; and Qin, Y. 2004. An integrated theory of the mind. *Psychological Review* 111(4):1036–1060.

Baddeley, A. D., and Hitch, G. 1974. Working memory. In Bower, G. H., ed., *The psychology of learning and motivation: Advances in research and theory.*, volume 8. Academic Press. 47–89.

Bar, M. 2007. The proactive brain: using analogies and associations to generate predictions. *Trends in Cognitive Sciences* 11(7):280–9.

Baxter, P., and Browne, W. 2010. Memory as the substrate of cognition: a developmental cognitive robotics perspective. In Johansson, B.; Sahin, E.; and Balkenius, C., eds., *Proceedings of the Tenth International Conference on Epigenetic Robotics*, 19–26.

Baxter, P., and Browne, W. 2011. Memory-based cognitive framework: a low-level association approach to cognitive architectures. In Kampis, G.; Karsai, I.; and Szathmry, E., eds., *European Conference on Artificial Life (ECAL 09)*, 402–409. LNCS 5777, Springer-Verlag.

Baxter, P.; Belpaeme, T.; Canamero, L.; Cosi, P.; Demeris, Y.; Enescu, V.; Hiolle, A.; Kruijff-Korbayova, I.; Looije, R.; Nalin, M.; Neerincx, M.; Sommavilla, G.; Tesser, F.; and Wood, R. 2011. Long-term human-robot interaction with young users. In Miyake, N.; Ishiguro, H.; Dautenhahn, K.; and Nomura, T., eds., *Proceedings of the Workshop Robots with Children: Practices for Human-Robot Symbiosis, at the 6th ACM/IEEE International Conference on Human-Robot Interaction.*

Baxter, P. 2010. *Foundations of a Constructivist Memory-Based approach to Cognitive Robotics*. Ph.D. Dissertation, University of Reading.

Beer, R. 1990. *Intelligence as Adaptive Behaviour*. Academic Press.

Bickhard, M. H. 2009. The biological foundations of cognitive science. *New Ideas in Psychology* 27:75–84.

Brom, C., and Lukavsky, J. 2008. *Episodic memory for human-like agents and human-like agents for episodic memory*. 42–47.

Cassimatis, N. L.; Trafton, J. G.; Bugajska, M. D.; and Schultz, A. C. 2004. Integrating cognition, perception and action through mental simulation in robots. *Robotics and Autonomous Systems* 49(1-2):13–23.

Clark, A. 2008. *Supersizing the mind*. Oxford University Press.

Clayton, N. S., and Dickinson, A. 1998. Episodic-like memory during cache recovery by scrub jays. *Nature* 395(September):272–274.

Clayton, N.; Griffiths, D.; Emery, N. J.; and Dickinson, A. 2001. Elements of episodic-like memory in animals. In Baddeley, A.; Conway, M.; and Aggleton., J., eds., *Episodic memory: new directions in research*. Oxford University Press. 232–248.

Cliff, D. 1991. Computational neuroethology: A provisional manifesto. In Meyer, J.-A., and Wilson, S., eds., *From Animals to Animats 3: Proceedings of the First International Conference on the Simulation of Adaptive Behavior*, 29–39.

Clowes, R., and Morse, A. 2005. Scaffolding cognition with words. In *Proceedings of the Fifth International Workshop on Epigenetic Robotics*.

Cowan, N. 1999. An embedded-processes model of working memory. In Miyake, A., and Shah, P., eds., *Models of Working Memory: Mechanisms of active maintenance and executive control*. Cambridge University Press. 62–101.

Fong, T.; Nourbakhsh, I.; and Dautenhahn, K. 2003. A survey of socially interactive robots. *Robotics and Autonomous Systems* 42(3-4):143–166.

Fuster, J. 1997. Network memory. *Trends in Neurosciences* 20(10):451–459.

Fuster, J. 2000. Cortical dynamics of memory. *International Journal of Psychophysiology*. 35(155-164).

Glenberg, a. M. 1997. What memory is for. *The Behavioral and Brain sciences* 20(1):1–19; discussion 19–55.

Goertzel, B.; Lian, R.; Arel, I.; Garis, H. d.; and Chen, S. 2010. World survey of artificial brains, part ii: Biologically inspired cognitive architectures. *Neurocomputing*.

Hawes, N.; Sloman, A.; Wyatt, J.; Zillich, M.; Jacobsson, H.; Kruijff, G.-j. M.; Brenner, M.; Berginc, G.; and Sko?, D. 2007. Towards an integrated robot with multiple cognitive functions. 1548–1553.

Haykin, S. 1999. *Neural Networks*. Prentice-Hall, 2nd edition.

Ho, W. C., and Watson, S. 2006. *Autobiographic Knowledge for Believable Virtual Characters*. 383–394.

Kawamura, K.; Peters, R. A.; Bodenheimer, R. E.; Sarkar, N.; Park, J.; Clifton, C. A.; Spratley, A. W.; and Hambuchen, K. A. 2004. A parallel distributed cognitive control system for a humanoid robot. *International Journal of Humanoid Robotics* 1(1):65–93.

Kieras, D. E., and Meyer, D. E. 1997. An overview of the epic architecture for cognition and performance with application to human-computer interaction. *Human-Computer Interaction* 12:391–438.

Langley, P.; Laird, J. E.; and Rogers, S. 2009. Cognitive architectures: Research issues and challenges. *Cognitive Systems Research* 10(2):141–160.

Miller, G. 2003. The cognitive revolution: a historical perspective. *Trends in Cognitive Sciences* 7(3):141–144.

Morse, A.; Belpaeme, T.; Cangelosi, A.; and Smith, L. 2010a. Thinking with your body: modelling spatial biases in categorization using a real humanoid robot. In *Cognitive Science 2010*.

Morse, A. F.; de Greeff, J.; Belpaeme, T.; and Cangelosi, A. 2010b. Epigenetic robotics architecture. *IEEE Transactions on Autonomous Mental Development* 2(4):325–339.

Morse, A. F.; Baxter, P.; Belpaeme, T.; Smith, L. B.; and Cangelosi, A. 2011. The Power of Words. In *Joint IEEE International Conference on Development and Learning and on Epigenetic Robotics*.

Morse, A.; Lowe, R.; and Ziemke, T. 2009. Manipulating space: Modelling the role of transient dynamics in inattentional blindness. *Connection Science* 21(4):275–296.

Neath, I., and Surprenant, A. M. 2003. *Human Memory: An introduction to research, data and theory*. Wadsworth, 2 edition.

Newell, A., and Simon, H. A. 1976. Computer science as empirical inquiry: symbols and search. *Communications of the ACM* 19(3):113–126.

Newell, A. 1990. *Unified Theories of Cognition*. Harvard University Press, 2nd edition.

Nuxoll, A. M. 2007. *Enhancing Intelligent Agents with Episodic Memory*. Ph.D. Dissertation, University of Michigan.

Postle, B. R. 2006. Working memory as an emergent property of the mind and brain. *Neuroscience* 139:23–38.

Rumelhart, D. E., and Norman, D. A. 1973. *Active semantic networks as a model of human memory*. 450–457.

Smith, L., and Samuelson, L. 2010. *Objects in Space and Mind: From Reaching to Words*. Oxford University Press.

Sowa, J. F. 1992. *Semantic Networks*, volume 5.

Sun, R. 2004. Desiderata for cognitive architectures. *Philosophical Psychology* 17(3):341–373.

Tulving, E. 1972. *Episodic and Semantic memory*. Academic Press. 382–404.

Tulving, E. 2002. Episodic memory: from mind to brain. *Annual review of psychology* 53:1–25.

Wood, R.; Baxter, P.; and Belpaeme, T. 2010. A developmental perspective on memory centered cognition. In B. Johansson, E. S., and Balkenius, C., eds., *Proceedings of the Tenth International Conference on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, volume 149 of *LUCS*. Lund: Lund University Cognitive Studies.

Wood, R.; Baxter, P.; and Belpaeme, T. in press. A review of long-term memory in natural and synthetic systems. *Adaptive Behavior*.

Wyatt, J. L., and Hawes, N. 2008. *Multiple Workspaces as an Architecture for Cognition*. 201–206.

Yamashita, Y., and Tani, J. 2008. Emergence of functional hierarchy in a multiple timescale neural network model: a humanoid robot experiment. *PLoS computational biology* 4(11):1–18.

Young, R. M., and Lewis, R. L. 1999. *The Soar cognitive architecture and human working memory*. Cambridge University Press. 224–256.