

Memory-Centred Cognitive Architectures for Robots Interacting Socially with Humans

Paul Baxter

Centre for Robotics and Neural Systems

The Cognition Institute

Plymouth University, U.K.

paul.baxter@plymouth.ac.uk

Abstract—The Memory-Centred Cognition perspective places an active association substrate at the heart of cognition, rather than as a passive adjunct. Consequently, it places prediction and priming on the basis of prior experience to be inherent and fundamental aspects of processing. Social interaction is taken here to minimally require contingent and co-adaptive behaviours from the interacting parties. In this contribution, I seek to show how the memory-centred cognition approach to cognitive architectures can provide an means of addressing these functions. A number of example implementations are briefly reviewed, particularly focusing on multi-modal alignment as a function of experience-based priming. While there is further refinement required to the theory, and implementations based thereon, this approach provides an interesting alternative perspective on the foundations of cognitive architectures to support robots engage in social interactions with humans.

I. INTRODUCTION

The representation and handling of memory is an important feature of cognitive architectures, with a variety of symbolic and sub-symbolic representation schemes used (generally as passive storage), typically based on assumptions of modularity [1]. As such, memory is generally considered to be structurally separable from the cognitive processing mechanisms, and functions to provide these ‘cognitions’ with the required data.

In the memory-centred cognition perspective, memory is instead considered to be a fundamentally active process that underlies cognitive processing itself rather than being a passive adjunct [2], [3]. Based on evidence and models in neuropsychology, e.g. [4], this approach necessitates a re-examination of the organisation and functions of cognitive architectures, as outlined below (section III).

Previously, I put forward the case for the greater consideration of memory in HRI developments [5]. I argued that memory is pervasive: fundamentally involved in all aspects of social behaviour, beyond mere passive storage of information in data structures. In this brief (and relatively introspective) contribution, I expand on this point, exploring specifically the requirements of social interaction for robots, and consequently what cognitive architectures need to encompass.

II. FACETS OF SOCIAL INTERACTION

Social interaction is a complex phenomena that entails a range of abilities on the part of the interactants; indeed, there are facets of human-human social interaction that are as yet not fully understood, with the neural substrates supporting

these in the individual yet to be characterised. One aspect that is commonly emphasised is the requirement for social signal processing for the individual, where behavioural cues (such as gaze, intonation, gesture, etc) should be interpreted to inform the behaviour of the observer.

One central idea emerging in the behavioural sciences is the notion of ‘social contingency’: the coupling and co-dependency of behaviours between interacting individuals [6]. This explicitly acknowledges the necessary role that the ‘other’ plays to set up the contingent behaviours, and moves away from the emphasis on social signal processing (though not discounting it). Minimal interaction paradigms provide intriguing illustrations of this: even given a low bandwidth interaction environment, there are non-trivial dynamics set up that cannot be explained by observations of an individual [7].

For social interaction generally, and in particular for this latter interacting systems perspective, there is an important role for prediction [8]. When interacting, there is an expectation that the interaction partner is also a social agent, and thus predictable in that context. Infants, for example, can use the gaze behaviour of a robot to infer that the robot is a psychological agent with which they can interact [9]. A previous study has further lent support to the idea that the imposition of expectations of social behaviour (and therefore the arising of socially contingent behaviours, in this case turn-taking) will come about if the interactants view each other as (potentially) social agents [10].

If the interaction partner (whether it is human or robot) is attributed with social agency, initially as a result of anthropomorphism for example [11], then one fundamental characteristic of social interaction between humans that will be seen is the ‘chameleon effect’ [12], or imitation/alignment, e.g. [13], [14], [15]. The presence of this within an interaction, as a type of contingency between the interactants (see above), could be seen as an indicator of sociality.

These phenomena, from attribution of social agency to alignment, illustrate a necessity for social robots (to a certain extent at least) to conform to human cognitive and behavioural features, as well as to their constraints, to enable predictability, consistency and contingency of robot behaviour with respect to the human(s) in the interaction.

III. MEMORY-CENTRED COGNITIVE ARCHITECTURE

From neuropsychology, the Network Memory framework [4] emphasises the central role that distributed associative cortical networks play in the organisation and implementation of cognitive processing in humans. The role of associative networks serves not only as a learning system (through Hebbian-like learning), but also as a substrate for activation dynamics. The reactivation and adaptation of existing networks combine to generate behaviour that is inherently based on prior experience.

The Memory-Centred Cognition perspective, as applied to the domain of cognitive robotics [2], seeks to extend these principles of operation: associative networks supporting activation dynamics that bring prior experience to bear on the current situation. A developmental perspective is necessary in order to do so [16]: the creation (and subsequent updating) of the associative networks must be done through the process of experience in order to form the appropriate associations between information in the present sensory and motor modalities of the robot (or system, in the case of a simulation).

Once an associative structure has been acquired, the principle mechanism at play is *priming* [2]. Priming in a memory-centred system occurs when some sub-set of the system is stimulated (from incoming sensory information for example), which causes activation to flow around the network, in turn causing parts of the network with no external stimulation to become active. Priming in this way fulfils a number of important functions. Firstly, it sets up cross-modal expectations, or the prediction of currently absent stimuli. Secondly, the priming process facilitates an integration of information across different modalities in a way that is explicitly based on prior experience (biased by the weights of the associative network).

A computational implementation of this has been applied to an account of the developmental acquisition of concepts [17]: not only was the system able to complete the task with a high success rate, but also the errors it made were consistent with those made by humans. A similar computational implementation has also been used to demonstrate how word labels for real-world objects can facilitate further cognitive processing [18]. These examples provide a glimpse of the range of cognitive processing (relevant to human cognitive processing) that can be accounted for using the memory-centred perspective.

Regarding social human-robot interaction, and in particular the notion that alignment is a fundamental feature of it (section II), the memory-centred perspective provides an intuitive, and indeed effective, account. Using exactly the same mechanism as for the concept learning study, the structure of an associative network was learned based on human behaviour (across a number of different modalities), which could then be directly used to determine the characteristics of the robot behaviour [14]. Alignment is achieved as a by-product of the way the memory-centred cognitive system operated: the associations were learned through experience, and behaviour was generated from priming (i.e. recall).

IV. ADDRESSING QUESTIONS

From the context outlined above, I now attempt to provide answers to a set of six questions relevant to the notion of social cognitive architectures. I particularly seek to emphasise a principled-basis (as opposed to computational mechanism-basis) for cognitive architectures and for the application to social interaction.

A. Why should you use cognitive architectures - how would they benefit your research as a theoretical framework, a tool and/or a methodology?

The benefit would be in considering cognitive architectures as a set of principles (a theoretical framework), a methodology for assessing these principles, and as a tool for providing robots with autonomous intelligent behaviour.

There are in my view three specific contributions related to scientific development (as opposed to technical implementation) that cognitive architectures can make to HRI research and development, which are centred around the idea of a cognitive architecture being made up of a set of formalised hypotheses.

Firstly, in a principled manner, they allow data and theory from empirical human studies to be integrated into artificial systems. For example, if data from a psychology experiment is to be integrated, a framework for doing so is required (i.e. the architecture enables an interpretation of the data). This first point promotes the idea of a directly human-inspired/constrained architecture. Secondly, treating cognitive architectures as a set of formalised (through implementation) principles, they facilitate a comparison of different architectures at a level abstracted away from the computational systems/algorithms used, enabling a focus on the assumptions. In the presently considered case of social interaction, this is a useful facet given the as yet uncertain nature of what exactly constitutes social interaction (section II). Thirdly, the application of cognitive architectures (in robotic systems for instance) provides a means of evaluating its constituent assumptions and principles. This is related to the first point, but is focused more on the integration of empirical evidence obtained from application/experimentation with the architecture itself.

B. Should cognitive architectures for social interaction be inspired and/or limited by models of human cognition?

Following from the principles of social interaction outlined above, essentially, yes.

Taking the view that social interaction between humans is founded on the intrinsic tendency of humans to expect certain types of behaviour from their interaction partners (see section II), it becomes important to ensure that the robot will not violate expectations. In order not to violate expectation, there must necessarily be some understanding (either on the part of the system designer or learned by the system itself) of what expected human behaviour would be.

In the memory-centred cognition perspective, prior interaction history of the robot with humans would constrain its future behaviour by this experienced behaviour.

C. What are the functional requirements for a cognitive architecture to support social interaction?

The discussion of social interaction (section II) emphasised the importance of contingent behaviour, anticipation/prediction to support this, and adaptation/personalisation. In addition, it is necessary to specify appropriate timing, and embodiment-appropriate responses.

If socially-appropriate behaviour is in the eye of the (human) beholder, then the Keepon robot for example demonstrates the importance of coherence of behaviour and timing [19]. The minimally complex embodiment is convincingly responsive in a social manner, to the extent that it is seen as a communicative partner [20]. Even though it doesn't use language, only uses few degrees of freedom (in contrast to many other robots used in HRI), and is only minimally humanoid in appearance, the effect of apparent sociality is strong.

Integration of sensory and motor modalities in a temporally consistent and responsive manner (i.e. contingency), based on principles of prediction from prior experience (i.e. memory), and coherency with the robot embodiment used (c.f. Keepon example) are therefore fundamental functional requirements for a social cognitive architecture.

D. How would the requirements for social interaction inform your choice of the fundamental computational structures of the architecture (e.g. symbolic, sub-symbolic, hybrid, ...)?

Given the commitment to the memory-centred cognition perspective in this work, there is a natural fit with sub-symbolic computational structures. This provides a number of inherent advantages (section III), such as the integration of predictive behaviour from prior experience, and priming effects (within and between modalities).

However, the nature of applications in human-robot interaction (relying on language for example) means that it is not yet possible to dispense with symbol-processing systems. Nevertheless, there is in principle an effort to push the limits of sub-symbolic processing mechanisms up the processing and representation hierarchy, as revisited below (section V).

E. What is the primary outstanding challenge in developing and/or applying cognitive architectures to social HRI systems?

One of the primary challenges in the application of cognitive architectures to social interaction lies in the general lack of understanding of what is precisely involved in human-human social interaction. To a certain extent it is an attempt to find a solution to a problem that is as yet not fully characterised. This reflects on the requirements for the cognitive architectures that should engage in social interaction: if a commitment to human-like cognition/behaviour is made (see section IV-B), then what precisely are the constraints that need to be incorporated?

A more practical concern that requires further development is the provision of sensory systems for robots that can provide sufficiently complex characterisations of the (social) environment for effective decision making. There is however, in my opinion, no clear distinction between sensory systems and cognitive processing, given the necessity for interpretation of

raw sensory signals (e.g. camera images) at various levels of abstraction.

F. Can you devise a social interaction scenario that current cognitive architectures would likely fail, and why?

The question is whether the application to a single domain can be generalised to other domains, which is where the benefits of cognitive architectures should come (section IV-A). As such, rather than a specific interaction scenario, I would suggest instead that autonomous sociality over variable time-scales poses challenges to current approaches and implementations.

In the short term, the challenge for social robots is to produce behaviour appropriate to the interaction context, informed by prior interaction experience, in a manner consistent with the expectations of the interacting humans. Furthermore, this socially interactive behaviour should adapt to the interaction partner over time, in terms of verbal and non-verbal behaviours for example. The technical challenges to support this in terms of sensory processing are outstanding, but there are also clear challenges in terms of the mechanisms of adaptation required (i.e. the 'cognitive' aspect). The memory-centred approach has ventured an implementation towards this problem, although the account is as yet incomplete.

Over extended periods of time, the challenges are compounded by requirements for stability. This is not just stability in terms of ensuring the system doesn't fail, but also in resolving the apparent trade-off between adaptability to new situations and robustness of the cognitive system. From the perspective of the memory-centred cognition account, the resolution to this question lies in how the formation, maintenance and manipulation of memory is handled in the system in terms of parameters and structures.

V. OUTLOOK

The nature of the discussion above is primarily principled and theoretical rather than focused on specific computational mechanisms. Naturally I believe memory-centred cognition perspective to have a consistency and coherence that merits consideration and further development. However, it is not in its current state able to practically support all aspects of real social interactions with real people.

This is a limitation shared with many 'emergent' cognitive architecture approaches [21]: theoretically interesting and coherent perhaps, but practically limited in terms of what can be done on real systems (use of language and dialogue being good examples of this). This is partly due to an implication of the theoretical perspective: by committing to a holistic approach that emphasises the integration and interplay of many different factors (including, for example, cognition, embodiment, culture, etc), the problem is made more difficult before a computational implementation is even begun. On a practical level, the types of dynamical system (be they neural network-based or other) used are typically not fully understood, or are at least highly complex [22], e.g. in terms of conditions for stability (particularly when adaptation/learning

is incorporated), which does not bode well for social robots that have to be reliable in real interactions with real people.

For these reasons, I do not believe that symbol-based approaches should (or can) be discarded, at least not for the foreseeable future. They provide the means of getting closer to actually achieving the desired behaviours in reality. Having said this, and as noted above (sec. IV-D), I remain intent on pushing the boundary between symbolic and sub-symbolic implementations ‘up’ the abstraction hierarchy, in a manner common with a range of other developmentally-oriented researchers [23], [24].

So, what does a memory-centred cognitive architecture look like if it is to be effectively applied to social interaction? And what does the memory-centred cognitive architecture enable in terms of social robots that would be difficult to achieve with an alternative approach? The functionality of developmental learning of cross-modal associations for prediction and action generation outlined above (section III) provides a technically difficult but in principle effective solution to the issue of learning from a vast array of potential multi-modal information in a way that is useful for action generation. This is not to say that this is the only approach (theoretical or computational) that would be capable of a similar functionality. However, this is where the second aspect, the requirement to fulfill social interaction with humans through conformity with human cognition (section II), becomes a distinguishing characteristic of the memory-centred approach.

In developing the theory, I have applied it to a range of practical systems and applications, as reviewed above (section III). For example using the same mechanism, accounts have been made of concept acquisition [17] and multi-modal robot behaviour alignment to an interaction partner [14]. Other systems using the same principles have been used to demonstrate the development of low-level sensory-motor coordination through experience [16], and the role of words in supporting new cognitive capabilities [18].

Whereas my commitment to the memory-centred cognition perspective for robotics is strong, my commitment to the specific mechanisms used is weak. I must acknowledge that there are a number of weaknesses with the various systems used, notably related to hierarchical structure/representation, and an incomplete account of temporal processing. However, in my view, this does not invalidate the theoretical approach, and merely serves to provide motivation to either find or develop a more appropriate computational implementation that fulfils all of the principles and constraints of the memory-centred cognition perspective.

ACKNOWLEDGEMENT

This work was supported by the EU FP7 project DREAM (grant number 611391, <http://dream2020.eu/>).

REFERENCES

[1] R. Sun, “Desiderata for Cognitive Architectures,” *Philosophical Psychology*, vol. 17, no. 3, pp. 341–373, sep 2004.

[2] P. Baxter, R. Wood, A. Morse, and T. Belpaeme, “Memory-Centred Architectures: Perspectives on Human-level Cognitive Competencies,” in *Proceedings of the AAAI Fall 2011 symposium on Advances in Cognitive Systems*, Arlington, Virginia, U.S.A.: AAAI Press, 2011, pp. 26–33.

[3] R. Wood, P. Baxter, and T. Belpaeme, “A Review of long-term memory in natural and synthetic systems,” *Adaptive Behavior*, vol. 20, no. 2, pp. 81–103, 2012.

[4] J. M. Fuster, “Network Memory,” *Trends in Neurosciences*, vol. 20, no. 10, pp. 451–9, 1997.

[5] P. Baxter and T. Belpaeme, “Pervasive Memory: the Future of Long-Term Social HRI Lies in the Past,” in *Third International Symposium on New Frontiers in Human-Robot Interaction at AISB 2014*, London, UK, 2014.

[6] E. Di Paolo and H. De Jaegher, “The Interactive Brain Hypothesis,” *Frontiers in Human Neuroscience*, vol. 6, no. June, pp. 1–16, 2012.

[7] E. Di Paolo, M. Rohde, and H. Iizuka, “Sensitivity to social contingency or stability of interaction? Modelling the dynamics of perceptual crossing,” *New Ideas in Psychology*, vol. 26, no. 2, pp. 278–294, 2008.

[8] E. C. Brown and M. Brüne, “The role of prediction in social neuroscience,” *Frontiers in Human Neuroscience*, vol. 6, no. May, pp. 1–19, 2012.

[9] A. N. Meltzoff, R. Brooks, A. P. Shon, and R. P. N. Rao, ““Social” robots are psychological agents for infants: a test of gaze following,” *Neural networks*, vol. 23, no. 8-9, pp. 966–72, 2010.

[10] P. Baxter, R. Wood, I. Baroni, J. Kennedy, M. Nalin, and T. Belpaeme, “Emergence of Turn-taking in Unstructured Child-Robot Social Interactions,” in *HRI’13*, Tokyo, Japan: ACM Press, 2013, pp. 77–78.

[11] B. R. Duffy, “Anthropomorphism and the Social Robot,” *Robotics and Autonomous Systems*, vol. 42, pp. 177–190, 2003.

[12] T. L. Chartrand and J. A. Bargh, “The Chameleon Effect: the perception-behavior link and social interaction,” *Journal of Personality and Social Psychology*, vol. 76, no. 6, pp. 893–910, 1999.

[13] K. Dautenhahn and A. Billard, “Studying robot social cognition within a developmental psychology framework,” in *Third European Workshop on Advanced Mobile Robots (Eurobot’99)*, Zurich, Switzerland, 1999, pp. 187–194.

[14] P. E. Baxter, J. de Greeff, and T. Belpaeme, “Cognitive architecture for humanrobot interaction: Towards behavioural alignment,” *Biologically Inspired Cognitive Architectures*, vol. 6, pp. 30–39, 2013.

[15] A.-L. Vollmer, K. J. Rohlfing, B. Wrede, and A. Cangelosi, “Alignment to the Actions of a Robot,” *International Journal of Social Robotics*, vol. 7, no. 2, pp. 241–252, 2015.

[16] P. Baxter and W. Browne, “Memory as the substrate of cognition: a developmental cognitive robotics perspective,” in *Proceedings of the Tenth International Conference on Epigenetic Robotics*, Örenäs Slott, Sweden, 2010, pp. 19–26.

[17] P. Baxter, J. D. Greeff, R. Wood, and T. Belpaeme, ““And what is a Seasnake?”: Modelling the Acquisition of Concept Prototypes in a Developmental Framework,” in *International Conference on Development and Learning and Epigenetic Robotics*. San Diego, USA: IEEE Press, 2012, pp. 1–6.

[18] A. F. Morse, P. Baxter, T. Belpaeme, L. B. Smith, and A. Cangelosi, “The Power of Words,” in *Joint IEEE International Conference on Development and Learning and on Epigenetic Robotics*. Frankfurt am Main, Germany: IEEE Press, 2011, pp. 1–6.

[19] H. Kozima and C. Nakagawa, “Social Robots for Children: Practice in Communication-Care,” in *AMC’06*. Istanbul, Turkey: IEEE Press, 2006, pp. 768–773.

[20] A. Peca, R. Simut, H.-L. Cao, and B. Vanderborght, “Do infants perceive the social robot Keepon as a communicative partner?” *Infant Behavior and Development*, vol. in press, 2015.

[21] D. Vernon, G. Metta, and G. Sandini, “A Survey of Artificial Cognitive Systems: Implications for the Autonomous Development of Mental Capabilities in Computational Agents,” *IEEE Transactions on Evolutionary Computation*, vol. 11, no. 2, pp. 151–180, 2007.

[22] R. D. Beer, “On the Dynamics of Small Continuous-Time Recurrent Neural Networks,” *Adaptive Behavior*, vol. 3, no. 4, pp. 469–509, 1995.

[23] L. B. Smith, “Cognition as a dynamic system: principles from embodiment,” *Developmental Review*, vol. 25, pp. 278–298, 2005.

[24] A. Cangelosi, *et al*, “Integration of Action and Language Knowledge: A Roadmap for Developmental Robotics,” *IEEE Transactions on Autonomous Mental Development*, vol. 2, no. 3, pp. 167–195, 2010.