

Towards Augmenting Dialogue Strategy Management with Multimodal Sub-Symbolic Context^{*}

Paul Baxter¹, Heriberto Cuayahuitl², Rachel Wood¹,
Ivana Kruijff-Korbayová², and Tony Belpaeme¹

¹ Centre for Robotics and Neural Systems, Plymouth University, U.K.

{paul.baxter,rachel.wood,tony.belpaeme}@plymouth.ac.uk

² Language Technology Lab, DFKI, Saarbrücken, Germany

{heriberto.cuahuitl,ivana.kruijff}@dfki.de

Abstract. A synthetic agent requires the coordinated use of multiple sensory and effector modalities in order to achieve a social human-robot interaction (HRI). While systems in which such a concatenation of multiple modalities exist, the issue of information coordination across modalities to identify relevant context information remains problematic. A system-wide information formalism is typically used to address the issue, which requires a re-encoding of all information into the system ontology. We propose a general approach to this information coordination issue, focussing particularly on a potential application to a dialogue strategy learning and selection system embedded within a wider architecture for social HRI. Rather than making use of a common system ontology, we rather emphasise a sub-symbolic association-driven architecture which has the capacity to influence the ‘internal’ processing of all individual system modalities, without requiring the explicit processing or interpretation of modality-specific information.

Keywords: Context in Cognitive Architecture, Dialogue Strategy Management, Distributed Memory, Multimodal Coordination, Social HRI

1 Introduction

In the application of robotic agents to societal issues, their capacity for extended social interaction with people is a central feature of behaviour that has not yet been resolved [9]. It has been proposed that key to achieving such functionality is the requirement for the coordination of multiple interaction modalities (such as verbal, behavioural, emotional, etc) and the adaptation of the agents’ behaviour to that of the human interactant [13]. Among the various potential interaction modalities, linguistic interaction is clearly of central importance. However, while the joint optimisation of dialogue strategies with other modalities shows improvements over the optimisation of dialogue strategies for individual modalities [4], the combinatorial expansion in the possible number of dialogue context

^{*} This work is funded by the EU FP7 ALIZ-E project (grant 248116).

configurations if encoded symbolically (e.g. [8]) could prove problematic in a real-time social HRI application domain. This variation on the ‘frame problem’ is particularly apparent in those systems where a global symbolic representation scheme (ontology) is used, where an explicit encoding of what information that should be taken into account must be provided. The aim of this contribution is to provide an outline approach (roadmap) to augment existing solutions to this problem by taking a perspective inspired by aspects of biological cognitive systems, emphasising sub-symbolic associative strategies. In so doing, we propose a system that may provide complementary functions to those mechanisms that are used in an existing dialogue strategy management implementation [8].

2 An Association-driven approach to providing ‘internal’ context

Taking inspiration from neuropsychological theories of memory and cognitive processing in biological agents, a framework for the ‘soft’ coordination of multiple modalities within a cognitive architecture has been devised that is based on the hebbian-like associative learning of relationships between information in different modalities [15]. Given that learning of this type may be characterised in a statistical manner, e.g. principle component analysis [12], such a system can extract the statistical relationship between the constituent modalities of a cognitive architecture, over the course of its operation. If the resulting structure is then treated as the substrate for activation dynamics, then the mechanism of priming is instantiated, in which the activity in one modality can influence the processing in another, thereby linking prior experience with ongoing behaviour generation across modalities [2, 15]. A computational system that embodies these characteristics has been implemented, which is informed by Interactive Activation and Competition models (and their adaptive extensions, e.g. [3]). Applied to a classification problem, it is shown that this system both completes the task, and does so in a manner consistent with human behavioural data [6].

In combination, these mechanisms (Hebbian association and activation spreading) enable the embedding of a realtime sub-symbolic statistical relationship component with what may be a largely symbolic processing cognitive architecture. Indeed, this general approach has formed the inspiration for a number of contemporary cognitive robotics implementations, notably in [11] which is used to model a range of human cognitive competencies.

In this perspective, context in a cognitive architecture can be derived on the basis of the associative substrate, and the activation dynamics operating over it: it is an encoding of the activity within the system, resulting from external input and internal processing. From the point of view of a given single cognitive modality (e.g. face recognition or dialogue strategy management), the activation received from other such cognitive modalities constitutes an internal context signal since it is explicitly based on the associations that have been formed, and thus on prior experience. In this way, a symbolic representation of context is not required, thus circumventing combinatorial explosion in symbolic representation

schemes. Furthermore, due to the generalisation (and compact representation) properties inherent to distributed associative networks (e.g. [10]), the proposed mechanism has the capacity to similarly generalise over prior experience.

3 A proposed application to augment Dialogue Strategy Management

One computational approach to dialogue strategy learning and language generation that has recently shown particular promise is that of Hierarchical Reinforcement Learning (HRL) [7, 8]. In this approach, a hierarchy of learning agents optimise a set of policies, rather than a unitary dialogue system policy. Each layer of the hierarchy therefore consists of multiple RL agents optimised to maximise reward for a subset of the state space (which includes context information). Recent developments have relaxed the notion of strict hierarchical organisation by positing a role for informal structure to enhance flexibility [5].

On this basis, it may be proposed that associations can be formed between the currently active dialogue management agent(s) and other cognitive modalities (e.g. face recognition, emotion interpretation, etc, see figure 1). This process would, over experience, encode statistical relationships between the dialogue agents and the state of processing in other cognitive modalities. These multimodal associations thus subsequently act as the substrate for priming effects between the different cognitive modalities and dialogue strategy management, by modulating the transition functions used in the HRL structure. For example, a dialogue agent that has strong positively weighted associations with a set of states in other modalities can be differentially supported in terms of dialogue strategy processing over other dialogue agents should these states re-occur: i.e. the influence of a global ‘context signal’. Thus, though the associative mechanism remains constant, the quantitative effect on action selection changes as a function of experience, and may take effect over relatively short time-periods given the real-time learning rate of the type of associative system in question, e.g. [11]. This scenario is particularly consistent with the loosening of the requirement for strict hierarchy in the HRL schema, as transitions between dialogue strategies in different hierarchy branches can be informed to a greater extent by the information provided by the proposed association mechanism (figure 1).

A number of points can be raised concerning this outline proposal. Firstly, the proposed association mechanism assumes that there is a functioning cognitive architecture upon which the associative mechanisms can operate: the emphasis of the current approach is on augmenting existing processing rather than replacing it. Secondly, there is no explicit semantic/symbolic information transferred between modalities through the proposed system. While this results in the need for an existing architectural mechanism to transfer specific semantic information (e.g. *the ball I see is red rather than green...*), the association and activation spread mechanisms enable multi-modal information to be taken into account in a relatively simplistic manner (architecturally), without requiring translation using a system ontology with the resulting frame-problem issues. Indeed, the

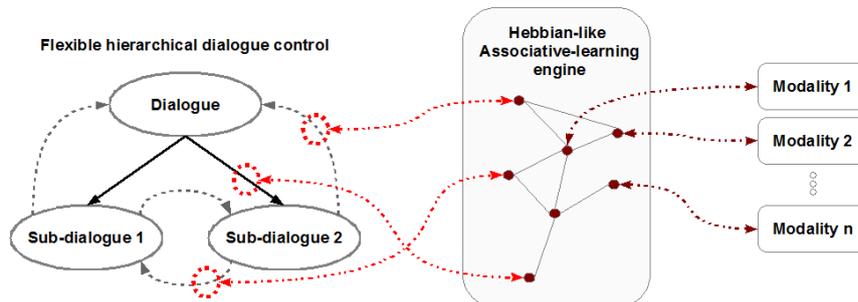


Fig. 1. Schematic of the interaction between a loose hierarchical dialogue strategy management system (left, adapted from [5]), and cognitive modalities (right, such as visual processing, etc), through the sub-symbolic associative learning mechanism (centre, see section 2): encoded statistical relationships can subsequently influence dialogue strategy state transitions (not all possible influences depicted).

combination of sub-symbolic and RL approaches has been previously considered as promising, e.g. [14]. Furthermore, it enables a clear perspective on how the agent can improve its behaviour through interaction experience rather than off-line training (enabling the findings of developmental learning theory to be leveraged [1]), since further associative specification can be progressively gained.

In summary, there are two primary advantages to this type of augmented dialogue strategy management: (1) opportunities for fast learning over the course of interaction with real users; and (2) global integration of modalities for coordinating multimodal conversational behaviours.

4 Perspectives

A roadmap for complementing the functionality of an existing dialogue strategy management implementation has been outlined: making use of statistical multimodal associations learned through experience (on the basis of a hebbian-like learning mechanism) in order to augment the dialogue strategy management process. In this way, there is also the potential for coordination among other modal-specific systems within a cognitive architecture. Even though the flow of explicit information through such a system is sparse (being constituted of scalar activation levels only), it has the capacity to improve performance not through the explicit transfer of information *per se*, but by virtue of the statistical properties of those associations that have been experienced providing a global context. While this paper has only presented a roadmap and is an ongoing programme of research, we believe that this principle of sub-symbolic context signals for augmenting modal-specific processing is generalizable to other cognitive architecture modalities, from perception to action, in the service of achieving social human-robot interaction with autonomous systems.

References

1. Asada, M., MacDorman, K.F., Ishiguro, H., Kuniyoshi, Y.: Cognitive developmental robotics as a new paradigm for the design of humanoid robots. *Robotics and Autonomous Systems*, 37: pp. 185–193, (2001)
2. Baxter, P., Wood, R., Morse, A., Belpaeme, T.: Memory-Centred Architectures: Perspectives on Human-level Cognitive. In *Proceedings of the Advances in Cognitive Systems track at AAAI Fall Symposium 2011, Arlington, USA*, pp. 26–33, (2011)
3. Burton, A.M.: Learning new faces in an interactive activation and competition model. *Visual Cognition*, 1(2): pp. 313–348, (1994)
4. Cuayahuitl, H., Kruijff-Korabayova, I.: Towards Learning Human-Robot Dialogue Policies Combining Speech and Visual Beliefs. In *Proc. of the Workshop on Paralinguistic Information and its Integration in Spoken Dialogue Systems, Granada, Spain*, pp. 133–140, (2011)
5. Cuayahuitl, H., Kruijff-Korabayova, I.: An Interactive Humanoid Robot Exhibiting Flexible Sub-Dialogues. In *Proc. of The 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Demo Session, Montreal, Canada*, (in press)
6. de Greeff, J., Baxter, P., Wood, R., Belpaeme, T.: From Penguins to Parakeets: a Developmental Approach to Modelling Conceptual Prototypes. In *Proc. of the PG Conference on Robotics and Development of Cognition at ICANN 2012, Lausanne, Switzerland*, (in press)
7. Cuayahuitl, H., Dethelfs, N.: Optimizing Situated Dialogue Management in Unknown Environments. In *Proc. of the Annual Conference of the International Speech Communication Association (INTERSPEECH), Florence, Italy*, pp. 1009–1012, (2011)
8. Dethelfs, N., Cuayahuitl, H.: Hierarchical Reinforcement Learning and Hidden Markov Models for Task-Oriented Natural Language Generation. In *Proc. of Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, Portland, USA*, pp. 654–659, (2011)
9. Fong, T., Nourbakhsh, I., Dautenhahn, K.: A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42(3-4): pp. 143–166, (2003)
10. McGarry, K., Wermter, S., MacIntyre, J.: Hybrid Neural Systems: From Simple Coupling to Fully Integrated Neural Networks. *Neural Computing Surveys*, 2: pp. 62–93, (1999)
11. Morse, A., De Greeff, J., Belpaeme, T., Cangelosi, A.: Epigenetic Robotics Architecture (ERA). *IEEE Transactions on Autonomous Mental Development*, 2(4): pp. 325–339, (2010)
12. Oja, E.: A simplified neuron model as a principle component analyzer. *Journal of Mathematical Biology*, 15: pp. 267–273, (1982)
13. Ros, R., Nalin, M., Wood, R., Baxter, P., Looije, R., Demiris, Y., Belpaeme, T., Giusti, A., Pozzi, C.: Child-Robot Interaction in The Wild: Advice to the Aspiring Experimenter. In *Proc. of ACM ICMI, Valencia, Spain*, pp. 335–342, (2011)
14. Santiago, R., Lendaris, G.G.: Reinforcement Learning and the Frame Problem. In *Proc. of IEEE IJCNN*, vol 5: pp. 2971–2976, (2005)
15. Wood, R., Baxter, P., Belpaeme, T.: A Review of long-term memory in natural and synthetic systems. *Adaptive Behaviour*, 20(2): pp. 81–103, (2012)