Reconciling Autonomy with Utility: A Roadmap and Architecture for Cognitive Development

David VERNON [1]
Circular Dynamics, Ireland

Abstract.
This paper focusses on a general issue which arises when one attempts to design a cognitive architecture for autonomous cognitive systems: the incompatibility of autonomy with external control and the consequent problem of getting these systems to perform prescribed tasks. In addressing this issue, we consider a specific endeavour to design and implement a biologically-inspired cognitive architecture for the iCub humanoid robot. This architecture is based on 43 guidelines arising from an extended study of the requirements imposed by developmental psychology, neuroscience, enactive cognition, and existing computational models.

Keywords. Autonomy, Utility, Development, Cognitive Architecture, Phylogeny, Ontogeny, Coaching, Imitative Learning.

Introduction

Natural cognitive systems are autonomous. Ideally, so too are biologically-inspired artificial ones. However, autonomy implies a self-determination that precludes explicit control by external agencies, including human users. This creates a problem: how can an autonomous cognitive system be designed so that it can exhibit the behaviours and functionality that its users require of it, such as the ability to perform prescribed tasks? We argue that the answer to this question has two components: one phylogentic and one ontogenetic. Both are linked to the developmental characteristic of autonomous cognitive systems [1] whereby the system acquires new capabilities and skills over time and, in the process, constructs its own understanding of the world around it through its interactions [2,3,4,5].

The developmental process is driven by various task non-specific motives which modulate the affective state of the system and thereby the actions in which it enagages. In turn, these motives reflect an innate value system: as Edelman notes, “the brain must ... establish regularities of behaviour under constraints of inherited value systems and of idiosyncratic perceptual and memorial events. In humans, such systems and events necessarily involve emotions and biases” [6]. If we wish these behaviours to exhibit some desired utility, we must embed in the system’s phylogeny — in its cognitive architec-

[1] Corresponding author e-mail: david@vernon.eu.

ture — the pre-disposition to acquire these behaviours and to fulfil this function. That is, the system phylogeny must have the appropriate value system and the associated motives. The subsequent ontongeny must then be structured to allow the cognitive system to develop the requisite skills and bring about the required behaviour. We argue that the most effective way to do this is through a process of human-robot coaching and imitative learning [7,8,9,10]. We illustrate this argument by making detailed reference to the iCub cognitive architecture, a biologically-inspired architecture that has been modelled on several aspects of human neuroscience and developmental psychology.

1. Autonomy

Autonomy is the self-maintaining organizational characteristic of living systems that enables them to use their own capacities to manage their interactions with the world in order to remain viable [11,3]. An autonomous system is entirely self-governing and self-regulating: it is not controlled by any outside agency and this allows it to stand apart from the rest of the environment, establishing a dynamic but self-sustaining identity. Characteristically, autonomy entails a capacity to contribute to its own persistence [12]. While the system may be dependent on the environment for material or energy, its operational identity is independent and any influences exerted on the system are brought about only by mutual interactions that do not threaten the autonomous operation of the system or control the system in any causal way. And therein lies the problem: if an external agent can't exert a causal influence on an autonomous cognitive system, how can one get it to do something useful? There are two aspects to the answer to this question, one dealing with the phylogeny of the cognitive system and the other with its ontongeny.

2. Phylogeny

What is the relationship between autonomy and cognition? One position is that cognition is the process by which an autonomous self-governing agent acts effectively in the world in which it is embedded [14]. As such, the dual purpose of cognition is to increase the agent's repertoire of effective actions and its power to anticipate the need for future actions and their outcomes [13]. Futhermore, development plays an essential role in the realization of these cognitive capabilities.

For an autonomous cognitive system to exhibit some utility from the perspective of its user or owner, it must have the same goals as its user or owner. So how does one specify the goals of an autonomous agent? The answer to this question seems to be that one can't, at least not directly. The best that one can do is embed a suitable set of motives in the system's phylogeny and subsequently use these to influence the system's goals.

Based on a study of enactive cognitive science, developmental psychology, neurophysiology, and computational modelling, the roadmap for cognitive development in humanoid robots in [13] identifies 43 guidelines for the design of a cognitive architecture and its deployment in a humanoid robot (see Table 1). In the context of the current discussion, this roadmap makes four important observations.

First, a cognitive system's actions are guided by prospection, directed by goals, and triggered by affective motives [15]. They are initially constrained in their numbers of

Embodiment
Rich array of physical sensory and motor interfaces Humanoid morphology Morphology integral to the model of cognition
Perception
Attention fixated on the goal of an action Perception of objecthood Discrimination & addition of small numbers; groups of large numbers Attraction to people (faces, their sounds, movements, and features) Preferential attention to biological motion Recognition of people, expression, and action Prolonged attention when a person engages in mutual gaze Perceive & communicate emotions by facial gesture and engage in turn-taking Involvement of the motor system in discrimination between percepts Mechanism to learn hierarchical representations Pre-motor theory of attention —spatial attention Pre-motor theory of attention —selective attention
Action
Movements organized as actions Early movements constrained to reduce the number of degrees of freedom Navigation based on dynamic ego-centric path integration Re-orientation based on local landmarks Action selection modulated by affective motivation mechanisms Hierarchically-structured representations of action-sequence skills
Anticipation
Internal simulation to predict, explain, & imagine events, and scaffold knowledge
Adaptation
Self-modification to expand actions and improve prediction Autonomous generative model construction Learning affordances Grounding internal simulations in actions Learn from experience the motor skills associated with actions Transient and generalized episodic memories of past experiences Procedural memory of actions and outcomes associated with episodic memories
Motivation
Social and explorative motives Affective drives associated with autonomy-preserving processes of homeostasis
Autonomy
Autonomy-preserving processes of homeostasis Encode space in motor & goal specific manner Minimal set of innate behaviours for exploration and survival Separate representations associated with each component / sub-system Concurrent competitive operation of components and subsystems

Table 1. Guidelines for the Phylogeny of a Developmental Cognitive System (from [13]).

freedom and the motor-programs that constitute them are learned. Second, because cognitive systems are anticipatory and prospective, they must have a mechanism to rehearse hypothetical scenarios through some process of internal simulation in order to predict, explain, and imagine events, and they must have a mechanism to use this outcome to modulate the behaviour of the system [16,17,18,19,20]. These processes should incorporate transient and generalized episodic memories of events and a procedural memory that links actions to perceptions [21]. Third, a developmental cognitive architecture must be capable of adaptation and self-modification, both in the sense of parameter adjustment of phylogenetic skills through learning and through the modification of the structure and organization of the system itself so that it is capable of altering its system dynamics based on experience in order to expand its repertoire of actions and enhance its prospective capabilities. Fourth, development should be driven by both explorative and social motives, the first concerned with both the discovery of novel regularities in the world and the potential of the system's own actions, the second with inter-agent interaction, shared activities, and mutually-constructed patterns of shared behaviour.

These observations, among others, and the guidelines set out in Table 1 formed the basis for the design of the iCub cognitive architecture which is summarized in Section 4. Before looking at this, we first consider the second aspect of how autonomy can be reconciled with utility: the ontogeny of the cognitive system.

3. Ontogeny

Development arises due to changes in the central nervous system as a result of dynamic interaction with the environment. Development is manifested by the emergence of new forms of action and the acquisition of predictive control of these actions. Mastery of action relies critically on prospection, i.e. the perception and knowledge of upcoming events. Repetitive practice of new actions is not focused on establishing fixed patterns of movement but on establishing the possibilities for prospective control in the context of these actions [22]. This highlights the importance of internal simulation in accelerating the scaffolding of early developmentally-acquired sensorimotor knowledge. Sharing this knowledge with other cognitive systems is only possible if they have a common history of experiences and if they have a similar phylogeny and a compatible ontogeny.

Development depends crucially on motivations which define the goals of actions. The two most important motives that drive actions and development are social and explorative. There are at least two exploratory motives: the discovery of novelty and regularities in the world and the discovery of the potential of the infant's own actions. The social motives are manifest as a fixation on social stimuli, imitation of basic gestures, and engagement in social interaction (e.g. joint attention and turn taking [23,24,25]).

The primary focus of the early stages of ontogenesis is to develop manipulative action based on visuo-motor mapping, learning to decouple motor synergies (e.g. grasping and reaching), anticipation of goal states, learning affordances, interaction with other agents through social motives, and imitative learning.

Imitiative learning in neonates is innate in their phyogeny and crucial to their developement, but the right ontogenetic circumstances must still be provided. In the same way, if the phylogeny of a biologically-inspired artificial autonomous cognitive systems is designed correctly, specifically by adhering to the guidelines set out in [13] and summarized in the previous section and in Table 1, then imitative learning can be used through a process of coaching to get the system to perform tasks and fulfil required functions, despite being autonomous. The imitative learning and associated coaching plays on the innate values and motives of the system in the following ways.

First, imitiation assists in the decomposition of actions into movements: it shows how goal-oriented goal-directed action is decomposed into or is constituted by component movements. Seond, goals are perceived directly as the desired outcomes of actions. Third, imitation is a restricted and focussed form of social engagement and is driven in

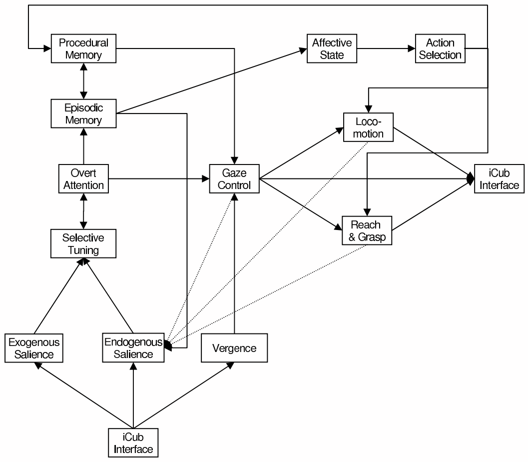

Figure 1. The iCub cognitive architecture.

part by a desire to belong to the social group and manifest by mutually-consistent action. This is important because, in terms of the two primary purposes of cognition to expand a repertoire or space of possible actions and to extend the prospective capability of the agent, acting together in a consistent or conformant manner reinforces the beliefs (or models) that are constructed by the system. Coaching and imitative learning not only provides an efficient and effective way acquiring new skills without the need to explore inappropriate or ineffectual alternatives, it also positively reinforces the system's predictive model. The repetitive nature of goal-directed action imitiative learning further enhances this effect.

Thus, imitative learning results from a dynamic balance between the exploratory motives and the social motives, both of which target the expansion of action capabilities and predictive power of the cognitive system. In this, imitiative learning through coaching brings together and consolidates the other early developmental tendencies.

4. The iCub Cognitive Architecture

The iCub cognitive architecture follows a significant subset of the roadmap guidelines summarized in Table 1, focussing on several key capabilities. Gaze control, reaching, and locomotion constitute the initial simple goal-directed actions. Episodic and procedural memories are included to effect a simplified version of internal simulation in order to provide capabilities for prediction and reconstruction, as well as generative model construction bootstrapped by learned affordances. In addition, motivations encapsulated in the system's affective state are made explicit so that they address curiosity and experimentation, both explorative motives, triggered by exogenous and endogenous factors, respectively. This distinction between the exogenous and the endogenous is reflected by the need to include an attention system to incorporate both factors. A simple process of homeostatic self-regulation governed by the affective state provides elementary action selection. Finally, all the various components of the cognitive architecture operate concurrently so that a sequence of states representing cognitive behaviour emerges from the interaction of many separate parallel processes rather than being dictated by some state-machine as in the case of most cognitive architectures. The iCub cognitive architecture at present comprises thirteen components (please refer to Fig. 1; more complete details can be found in [13]).

Together, the exogenous salience [26], endogenous salience, selective tuning [27], and overt attention [28] components comprise the iCub's perception system. Similarly, gaze control [29], vergence, reach & grasp, locomotion comprise the iCub's actions system. The episodic memory and the procedural memory together provide the iCub's principle mechanism for anticipation and adaptation. The affective state component effects the iCub motivations which together with the action selection component provide a very simple homeostatic process which regulates the autonomous behaviour of the iCub. The iCub interface component completes the architecture and reflect the embodiment of the iCub from an architecture point of view.

5. Conclusion

The primary assertion being put forward in this paper is that a suitably-designed cognitive system, i.e. one with a well-configured cognitive architecture, will exhibit a natural tendency to behave in a manner suggested by an instructing coach. Despite being autonomous, such cognitive systems are inherently trainable. The ease with which they can be trained depends on the balance of exploration and social motivation. This leaves unanswered three questions. First, what is the mechanism by which that part of the cognitive architecture concerned the association of movement and perception in goal-directed prospectively-guided action grows and develops?[2] For example, what metric guides or governs the reinforcement learning of perception-movement associations and the construction of action pathways? Second, in a related manner, how is the right balance between exploratory and social motives achieved? Third, how are episodic memories consolidated into semantic memories so that the resultant experiential learning is transferrable to new situations? Answering these questions poses an important research challenge.

Acknowledgements

Aspects of this work were supported by the European Commission, Project IST-004370, under Strategic Objective 2.3.2.4: Cognitive Systems. Discussions with Gordon Cheng and Marcia Riley, Institute of Cognitive Systems, Technical University of Munich, were very valuable in developing the ideas on coaching and imitative learning.

References

[1] D. Vernon, G. Sandini, and G. Metta. The icub cognitive architecture: Interactive development in a humanoid robot. In *Proceedings of IEEE International Conference on Development and Learning (ICDL)*, Imperial College, London, 2007.

[2] E. Thompson. *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press, Boston, 2007.

[3] T. Froese and T. Ziemke. Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Artificial Intelligence*, 173:466–500, 2009.

[2] In the iCub cognitive architecture, this comprises the procedural memory, episodic memory, affective state, and action selection components.

[4] D. Vernon. Enaction as a conceptual framework for development in cognitive robotics. *Paladyn Journal of Behavioral Robotics*, 1(2):89–98, 2010.
[5] J. Stewart, O. Gapenne, and E. A. Di Paolo. *Enaction: Toward a New Paradigm for Cognitive Science*. MIT Press, 2011.
[6] G. M. Edelman. *Second Nature: Brain Science and Human Knowledge*. Yale University Press, New Haven and London, 2006.
[7] A. Billard, S. Calinon, R. Dillmann, and S. Schaal. Robot programming by demonstration. In *Springer Handbook of Robotics*, pages 1371–1394. 2008.
[8] R. Dillmann, T. Asfour, M. Do, R. Jäkel, A. Kasper, P. Azad, A. Ude, S. Schmidt-Rohr, and M. Lösch. Advances in robot programming by demonstration. *KI*, 24(4):295–303, 2010.
[9] M. Riley, A. Ude, C. Atkeson, and G. Cheng. Coaching: An approach to efficiently and intuitively create humanoid robot behaviors. In *IEEE-RAS Conference on Humanoid Robotics*, pages 567–574, 2006.
[10] S. Schaal, A. Ijspeert, and A. Billard. Computational approaches to motor learning by imitation. 2003.
[11] W. D. Christensen and C. A. Hooker. An interactivist-constructivist approach to intelligence: self-directed anticipative learning. *Philosophical Psychology*, 13(1):5–45, 2000.
[12] M. H. Bickhard. Autonomy, function, and representation. *Artificial Intelligence, Special Issue on Communication and Cognition*, 17(3-4):111–131, 2000.
[13] D. Vernon, C. von Hofsten, and L. Fadiga. *A Roadmap for Cognitive Development in Humanoid Robots*, volume 11 of *Cognitive Systems Monographs (COSMOS)*. Springer, Berlin, 2010.
[14] H. Maturana and F. Varela. *The Tree of Knowledge – The Biological Roots of Human Understanding*. New Science Library, Boston & London, 1987.
[15] R. Núñez and W. J. Freeman. *Reclaiming Cognition — The Primacy of Action, Intention and Emotion*. Imprint Academic, Thorverton, UK, 1999.
[16] A. Berthoz. *The Brain's Sense of Movement*. Harvard University Press, Cambridge, MA, 2000.
[17] G. Hesslow. Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Sciences*, 6(6):242–247, 2002.
[18] R. Grush. The emulation theory of representation: motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27:377–442, 2004.
[19] M. P. Shanahan. Cognition, action selection, and inner rehearsal. In *Proceedings IJCAI Workshop on Modelling Natural Action Selection*, pages 92–99, 2005.
[20] M. P. Shanahan. A cognitive architecture that combines internal simulation with a global workspace. *Consciousness and Cognition*, 15:433–449, 2006.
[21] P. Langley, J. E. Laird, and S. Rogers. Cognitive architectures: Research issues and challenges. *Cognitive Systems Research*, 10(2):141–160, 2009.
[22] E. S. Reed. *Encountering the world: towards an ecological psychology*. Oxford University Press, New York, 1996.
[23] J. Nadel, C. Guerini, A. Peze, and C. Rivet. The evolving nature of imitation as a format for communication. In J. Nadel and G. Butterworth, editors, *Imitation in Infancy*, pages 209–234. Cambridge University Press, Cambridge, 1999.
[24] G. S. Speidel. Imitation: a bootstrap for learning to speak. In G. E. Speidel and K. E. Nelson, editors, *The many faces of imitation in language learning*, pages 151–180. Springer Verlag, 1989.
[25] C. Trevarthen, T. Kokkinaki, and G. A. Fiamenghi Jr. What infants' imitations communicate: with mothers, with fathers and with peers. In J. Nadel and G. Butterworth, editors, *Imitation in Infancy*, pages 61–124. Cambridge University Press, Cambridge, 1999.
[26] J. Ruesch, M. Lopes, J. Hornstein, J. Santos-Victor, and R. Pfeifer. Multimodal saliency-based bottom-up attention - a framework for the humanoid robot icub. In *Proc. International Conference on Robotics and Automation*, pages 962–967, Pasadena, CA, USA, May 19-23 2008.
[27] J. K. Tsotsos, S. Culhane, W. Wai, Y. Lai, N. David, and F. Nufb. Modeling visual attention via selective tuning. *Artificial Intelligence*, 78:507–547, 1995.
[28] A. Zaharescu, A. L. Rothenstein, and J. K. Tsotsos. Towards a biologically plausible active visual search model. In L. Paletta, J. K. Tsotsos, E. Rome, and G. Humphreys, editors, *Proceedings of the Second International Workshop on Attention and Performance in Computational Vision, WAPCV*, volume LNCS 3368, pages 133–147. Springer, 2004.
[29] M. Lopes, A. Bernardino, J. Santos-Victor, C. von Hofsten, and K. Rosander. Biomimetic eye-neck coordination. In *IEEE International Conference on Development and Learning*, Shanghai, China, 2009.