

The Radical Interactionism Conceptual Commitment

Olivier L. Georgeon

*Université Lyon 1
LIRIS, CNRS, UMR5205
Villeurbanne F-69622, FRANCE*

OLIVIER.GEORGEON@LIRIS.CNRS.FR

David W. Aha

*Navy Center for Applied Research in AI
Naval Research Laboratory, Code 5514
4555 Overlook Ave., SW
Washington, DC 20375, USA*

DAVID.AHA@NRL.NAVY.MIL

Editor: Haris Dindo, James Marshall, Giovanni Pezzulo

Abstract

We introduce *Radical Interactionism* (RI), which extends Franklin et al.'s (2013) *Cognitive Cycles as Cognitive Atoms* (CCCA) proposal in their discussion on conceptual commitments in cognitive models. Similar to the CCCA commitment, the RI commitment acknowledges the indivisibility of the perception-action cycle. However, it also reifies the perception-action cycle as *sensorimotor interaction* and uses it to replace the traditional notions of *observation* and *action*. This complies with constructivist epistemology, which suggests that knowledge of reality is constructed from regularities observed in sensorimotor experience. We use the LIDA cognitive architecture as an example to examine the implications of RI on cognitive models. We argue that RI permits self-programming and constitutive autonomy, which have been acknowledged as desirable cognitive capabilities in artificial agents.

Keywords: constructivist learning, sensorimotor modeling, self-programming, constitutive autonomy, self-motivation, Enaction.

1. Introduction

In their paper “Conceptual Commitments of the LIDA Model of Cognition”, Stan Franklin, Steve Strain, and Ryan McCall invite Artificial General Intelligence (AGI) researchers to discuss which conceptual commitments are essential for AGI agents. In response to this invitation, we wish to further discuss their fourth commitment: *Cognitive Cycles as Cognitive Atoms* (CCCA). On page 9, they write: “we hesitate to propose it as important for the AGI research in general, since to our knowledge, no other system-level cognitive architecture makes such a commitment”. We make this commitment in our Enactive Cognitive Architecture (ECA, Georgeon, Marshall, and Manzotti, 2013). Moreover, we extend this commitment into a more radical one called Radical Interactionism (RI).



In Section 4.4, Franklin et al. define a cognitive cycle as a cycle that begins by sampling (sense) the environment and ends by selecting an appropriate response (action), traversing various phases including perception, understanding, consciousness, and learning. Modeling such a cognitive cycle as a *cognitive atom* entails considering it as *indivisible*, as opposed to dividing it into a sequence of separate tasks. To our understanding, LIDA acknowledges this indivisibility by implementing cognitive cycles through asynchronous distributed processes. Since LIDA’s design imposes no synchronicity on the distributed processes that generate the cognitive cycle, it is indeed not divided into a predefined sequence. In our view, this approach embodies CCCA in which “atom” denotes a sense of indivisibility. We introduce a more radical view that additionally uses “atom” as a primitive notion for designing the model.

Our Radical Interactionism conceptual commitment invites designers of cognitive models to consider the notion of *sensorimotor interaction* as a primitive, instead of perception and action. This is analogous to a mathematical system’s primitives, which are used in axioms and theorems to define more complex structures, but which are themselves undefined within the system. Traditional cognitive models take perceptions and actions as primitive notions, and derive the notion of sensorimotor interaction from them. The RI conceptual commitment recommends doing the opposite. Although the CCCA commitment takes a first step in the RI direction by considering cognitive cycles as indivisible, Franklin et al. still define a cognitive cycle using the primitive notions of observation and action. The RI commitment eliminates these as primitives and instead frames each cognitive cycle as a sensorimotor interaction.

Intuitively, a sensorimotor interaction fits Franklin et al.’s description of a cognitive cycle: “Each cycle constitutes a unit of sensing, attending and acting. A cognitive cycle can be thought of as a moment of cognition, a cognitive moment” (p4). Within constructivist epistemology, sensorimotor interactions can be viewed as representing a Piagetian (1955) *sensorimotor scheme*, from which the subject constructs knowledge of reality. At a philosophical level, this view relates to phenomenology (e.g., Dreyfus, 2007), which argues that knowledge of the self and of the world derives from regularities in phenomenological experience. Accordingly, a sensorimotor interaction may be understood as a *chunk of phenomenological experience*, making the stream of sensorimotor interactions the primitive material from which the agent constructs all of its knowledge.

2. RI Impact on cognitive modeling

As an example to illustrate the RI commitment, we examine how it would impact the LIDA model. It would imply modifying the left-side part of Figure 1 in Franklin et al.’s paper as shown in our Figure 1 below. Since perception and action do not exist in RI, the *Sensory Memory* and the *Motor Plan Execution* modules would be removed, as well as the *Sensory Stimulus*, *Actuator Execution*, and *Dorsal Stream* connections. Instead, the *Sensory Motor Memory* module would directly connect to the *Internal & External Environment* module through the *Intended Interaction* and the *Enacted Interaction* connections. As a result of these changes, *stimuli* and *actions* would be eliminated from the architecture and replaced by a single type of primitive object: sensorimotor interactions.




Figure 1. A diagrammatic adaptation of LIDA to Radical Interactionism. *Sensory Motor Memory* directly interacts with *Internal & External Environment* through the *Intended Interaction* i_t and the *Enacted Interaction* e_t . The *Other Modules of the Cognitive Model* box encompasses the remaining modules of LIDA. At time t , the agent chooses the sensorimotor interaction i_t that it intends to enact from among the set of sensorimotor interactions I . The attempt to enact i_t may change the environment. The agent then receives the enacted sensorimotor interaction e_t .

The algorithm begins with a predefined set of sensorimotor interactions I , called *primitive interactions*. At a given time t , the agent chooses a primitive interaction i_t that it intends to enact, from among I . The agent ignores this enaction’s meaning; that is, the agent has no rules that would exploit knowledge of how the designer programmed the primitive interactions through actuator movements and sensory feedback (such as: “if a specific interaction was enacted then perform a specific computation”). As a response from the tentative enaction of i_t , the agent receives the *enacted interaction* e_t , which may differ from i_t . The enacted interaction is the only data available to the agent that carries some information about the external world, but the agent ignores the meaning of this information. As an example, the primitive interaction i_t may correspond to actively feeling (through touching) an object in front of the agent, involving both a movement and a sensory feedback. The tentative enaction of i_t may indeed result in feeling an object, in which case $e_t = i_t$, or may result in feeling nothing, if there is no object in front of the agent, in which case the enacted interaction e_t corresponds to a different interaction: moving while feeling nothing. The agent constructs knowledge about its environment and organizes its behavior through regularities observed in the sequences of enacted interactions.

Within the RI commitment, LIDA would not have to construct *action schemes* defined “as an action together with its context and expected result” (Franklin et al., p5, citing Drescher, 1991). Procedural and episodic memories would instead directly process sensorimotor interactions. The learning mechanism would construct higher-level sensorimotor interactions, called *composite interactions*, as pairs of a *context interaction* and an *intention interaction*: $i_{\text{composite}} = \langle i_{\text{context}}, i_{\text{intention}} \rangle$. Such a learning mechanism would be bottom-up, and would generate a hierarchy of composite interactions where each level would contain lower-level composite interactions, except the bottom level, which would be made of primitive interactions. At any time, the agent would represent its current situation as a set of interactions, called *interactional context*. The LIDA behavior selection mechanism would be modified to activate composite interactions $i_{\text{composite}}$ whose context interaction i_{context} belongs to the interactional context. The intention interactions $i_{\text{intention}}$ of the activated composite interactions $i_{\text{composite}}$ would compete to be the one selected as the next intended interaction. Since the intended interaction can be a primitive interaction or a previously learned composite interaction, this mechanism allows asynchrony between the architecture’s behavior-selection mechanism and the primitive interactional mechanism, as illustrated in Figure 2.




Figure 2. Levels of interaction in Radical Interactionism. Over time, the agent constructs composite interactions, which implement a series of primitive interactions. At time t , the behavior selection mechanism may select a previously constructed composite interaction $i_{ct} = \langle i_{p1} \dots i_{pn} \rangle$ to try to enact. The tentative enaction of i_{ct} is delegated to the *Primitive Sensory Motor Memory*, which controls the enaction of the sequence of the n primitive interactions i_{p1} to i_{pn} , and returns the enacted composite interaction e_{ct} to the *Composite Sensory Motor Memory*. The rest of the cognitive architecture sees the primitive loop as an *Internal & External Environment Constructed at Time t* , with which it interacts through higher-level interactions. Since the environment that the cognitive architecture interacts with evolves as the agent develops, the agent can recursively learn increasingly complex behaviors.

The coupling between the cognitive architecture and the environment (represented by the *Internal & External Environment Constructed at Time t* in Figure 2) evolves as the agent learns increasingly sophisticated patterns of behaviors, while simultaneously representing the world in terms of increasingly sophisticated affordances. We refer the reader to other publications for a more thorough description of this mechanism (Georgeon & Ritter, 2012), its implications in a cognitive architecture (Georgeon, Marshall, and Manzotti, 2013), and demonstrations in which primitive interactions consist of active feeling (e.g., Georgeon and Marshall, 2013) and rudimentary active vision (Georgeon, Cohen, and Cordier, 2011).

3. Conclusion

Rather than offering a method to address the traditional perception-action problem, Radical Interactionism presents a profoundly different formulation of this problem. Indeed, many studies define the perception-action problem as learning a mapping between a predefined action space and a predefined observation space. For example, these include studies in reinforcement learning (e.g., Sutton & Barto, 1998), means-end analysis (e.g., Drescher, 1991), and robotics (e.g., Pierce & Kuipers 1997). In contrast, the RI approach begins with a predefined interactional space, and focuses on the problem of generating learning effects that are considered important in AGI agents, namely self-programming (e.g., Thórisson, Nivel, Sanz, and Wang, 2013) and *constitutive autonomy*, which Froese and Ziemke (2009) define as an agent’s ability to “self-constitute its identity”, which they argue is a prerequisite for autonomous sense-making. RI agents realize self-programming because they learn increasingly long sequences of interactions and re-enact these in appropriate contexts. They have constitutive autonomy because their coupling with the environment evolves through their individual experience of interaction.

Moreover, RI allows implementing a type of self-motivation called *interactional motivation* (Georgeon, Marshall, and Gay, 2012). To implement interactional motivation, a designer attaches

a predefined numerical valence to some primitive interactions, and biases the behavior selection mechanism to select sequences of interactions that have the highest total valence. Interactional motivation provides a way to specify inborn preferences (some primitive interactions that the agent innately likes or dislikes) but does not specify the agent's goal; the agent thus engages in open-ended learning to find its own way to fulfill its inborn preferences of interaction. This view implements Glasersfeld's (1984, p29) idea that goals "arise for no other reason than this: a cognitive organism evaluates its experiences, and because it evaluates them, it tends to repeat ones and avoid others".

Associated with interactional motivation, the RI commitment attacks the problem of designing self-motivated agents that learn to master lawful regularities in the flow of sensorimotor contingencies (SMCs, O'Regan and Noë, 2001). Notably, RI agents do not have to learn the first category of SMCs identified by Buhrmann, Di Paolo, and Barandiaran (2013): the *sensorimotor environment*, which they "innately" know through the set of primitive interactions *I*. However, they must construct *intentional actions*—actions that they can choose knowingly of their anticipated effects (e.g., Engel, Maye, Kurthen, & König, 2013). Concurrently, they must learn the existence of possibly persistent entities in the environment that they can observe, displace, transform, or use through intentional actions.

With regard to epistemological theories, RI relates to Glasersfeld's radical constructivism, whose "radical difference [from traditional conceptualizations] concerns the relation of knowledge and reality. Whereas in the traditional view of epistemology, as well as of cognitive psychology, that relation is always seen as a more or less picture-like (iconic) correspondence or match, radical constructivism sees it as an adaptation in the functional sense" (Glasersfeld, 1984, p20). Accordingly, RI does not implement a "picture-like" relation in the input received by the model from the environment. In contrast with observations in traditional models, the input received by an RI model from the environment only consists in the enacted interaction, which does not directly represent the environment. Because the input does not directly represent the environment, RI agents can be designed without ontological assumptions about the environment, which is critical for agents that perform open-ended learning in the real world.

Acknowledgements

This work was supported by the French Agence Nationale de la Recherche (ANR) contract ANR-10-PDOC-007-01, and NRL. The views and opinions contained in this paper are those of the authors and should not be interpreted as representing the official views or policies, either expressed or implied, of NRL or the DoD.

References

- Buhrmann, T.; Di Paolo, E.; and Barandiaran, X. 2013. A dynamical systems account of sensorimotor contingencies. *Frontiers in psychology*. 4:285.
- Drescher, G. L. 1991. *Made-up minds, a constructivist approach to artificial intelligence*. Cambridge, MA: MIT Press.
- Dreyfus, H. 2007. Why Heideggerian AI failed and how fixing it would require making it more Heideggerian. *Philosophical Psychology*. 20(2): 247–268.
- Engel, A.; Maye, A.; Kurthen, M.; König, P. 2013. Where's the action? The pragmatic turn in cognitive science. *Trends in Cognitive Sciences*. 17(5):202–209.

- Froese, T. and Ziemke, T. 2009. Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Artificial Intelligence*. 173(3-4): 466–500.
- Georgeon, O.; Cohen, M.; and Cordier, A. 2011. A model and simulation of early-stage vision as a developmental sensorimotor process. In *proceedings of Artificial Intelligence Application and Innovations (AIAI'2011)*, 11–16. Corfu, Greece.
- Georgeon, O. and Marshall, J. 2013. Demonstrating sensemaking emergence in artificial agents: A method and an example. *International Journal of Machine Consciousness*. 5(2): 131–144.
- Georgeon, O.; Marshall, J.; and Gay, S. 2012. Interactional motivation in artificial systems: between extrinsic and intrinsic motivation. In *proceedings of the Second International Conference on Development and Learning, and on Epigenetic Robotics (EPIROB'2012)*, 1–2. San Diego, CA.
- Georgeon, O.; Marshall, J.; and Manzotti, R. 2013. ECA: An enactivist cognitive architecture based on sensorimotor modeling. *Biologically Inspired Cognitive Architectures*. 6: 46–57.
- Georgeon, O. and Ritter, F. 2012. An intrinsically-motivated schema mechanism to model and simulate emergent cognition. *Cognitive Systems Research*. 15–16: 73–92.
- Glaserfeld, E. V. 1984. An introduction to radical constructivism. In *The invented reality* ed. P. Watzlawick. 17–40. New York, NY: Norton.
- O'Regan, K. and Noë, A. (2001). A sensorimotor account of vision and visual consciousness, *Behavioral and Brain Sciences*. 24(5): 939–1031.
- Piaget, J. 1955. *The construction of reality in the child*. London: Routledge and Kegan Paul.
- Pierce, D. and Kuipers, B. 1997. Map learning with uninterpreted sensors and effectors. *Artificial Intelligence*. 92: 169-227.
- Sutton, R. and Barto, A. 1998. *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Thórisson, K.; Nivel, E.; Sanz, R.; and Wang, P. 2013. Approaches and Assumptions of Self-Programming in Achieving Artificial General Intelligence. *Journal of Artificial General Intelligence*. 3(3): 1–10.