Biologically Inspired Cognitive Architectures (2013) 6, 46–57

Available at www.sciencedirect.com

SciVerse ScienceDirect

journal homepage: www.elsevier.com/locate/bica

RESEARCH ARTICLE

ECA: An enactivist cognitive architecture based on sensorimotor modeling

Olivier L. Georgeon [a,b,*], James B. Marshall [c], Riccardo Manzotti [d]

[a] *Université de Lyon, CNRS, Villeurbanne F-69622, France*
[b] *Université Lyon 1, LIRIS, UMR5205, Villeurbanne F-69622, France*
[c] *Sarah Lawrence College, Bronxville, NY 10708, USA*
[d] *IULM University, Milan, Italy*

Received 15 March 2013; received in revised form 18 April 2013; accepted 14 May 2013

KEYWORDS
Enaction;
Self-motivation;
Cognitive architecture;
Developmental learning

Abstract
A novel way to model an agent interacting with an environment is introduced, called an Enactive Markov Decision Process (EMDP). An EMDP keeps perception and action embedded within sensorimotor schemes rather than dissociated, in compliance with theories of embodied cognition. Rather than seeking a goal associated with a reward, as in reinforcement learning, an EMDP agent learns to master the sensorimotor contingencies offered by its coupling with the environment. In doing so, the agent exhibits a form of intrinsic motivation related to the *autotelic principle* (Steels, 2004), and a value system attached to interactions called interactional motivation. This modeling approach allows the design of agents capable of autonomous self-programming, which provides rudimentary constitutive autonomy—a property that theoreticians of enaction consider necessary for autonomous sense-making (e.g., Froese & Ziemke, 2009). A cognitive architecture is presented that allows the agent to discover, memorize, and exploit spatio-sequential regularities of interaction, called Enactive Cognitive Architecture (ECA). In our experiments, behavioral analysis shows that ECA agents develop active perception and begin to construct their own ontological perspective on the environment.

1. Introduction

In cognitive science, there has been a customary and traditional tripartite division of the mind between perception, the control system, and motor action. This view has been

* Corresponding author at: Université de Lyon, CNRS, France.
E-mail address: olivier.georgeon@liris.cnrs.fr (O.L. Georgeon).

2212-683X/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.bica.2013.05.006

nicely dubbed the ''classic sandwich model'' by Susan Hurley (1998). Many control architectures are built in this way. Since the 1980s there have been many attempts to challenge this traditional picture particularly in the field of robotics (e.g., Brooks, 1991) but also from a more psychological and theoretical perspective (e.g., Hirose, 2002; Shanahan, 2010; Ziemke, 2001). In particular, the idea emerged that it might be a mistake to consider sensation independently from action and that we should design cognitive systems on the basis of *low-level sensorimotor loops* that represent *sensorimotor patterns* of interaction. This intuition gained momentum from other related views such as embodied cognition (e.g., Anderson, 2003; Holland, 2004), ecological psychology (Chemero & Turvey, 2007; Gibson, 1979), sensorimotor theories (O'Regan & Noë, 2001; O'Regan, 2012), morphological robotics (Paul, 2006; Pfeifer & Bongard, 2006; Pfeifer, 1999), developmental robotics (Lungarella, Metta, Pfeifer, & Sandini, 2003), and epigenetic robotics (Berthouze & Ziemke, 2003; Zlatev, 2001). Here, we introduce a modeling approach that goes a step beyond the notion of low-level sensorimotor loops by simply considering sensorimotor patterns—also called *sensorimotor schemes* by Piaget (1951)—as the atomic elements manipulated by our algorithms.

Varela, Thompson, and Rosch (1991) coined the term *enactive perception* to suggest that organism and environment are coupled together. The features of the environment to which an organism responds are singled out by the ongoing activity in the organism. The domain that defines this coupling has been called the *relational domain* (e.g., Froese & Ziemke, 2009). The theory of enaction, initiated by Varela, stresses that the relational domain evolves over the organism's life in a manner that is codetermined by the organism and the environment. The fact that the relational domain is not predefined makes possible the organism's *constitutive autonomy*—the capacity of the organism to ''self-constitute its identity'' (Froese & Ziemke, 2009). These authors argue that constitutive autonomy is an important aspect of organisms because it is a precondition of sense-making and intrinsic teleology, and is thus a property that we should seek to obtain in artificial agents.

Furthermore, the term enaction also incorporates the idea that perception involves physical activity, or action. A model of reference was offered by O'Regan and Noë's (2001) *sensorimotor contingencies* theory. To perceive the world is to master the sensorimotor contingencies between the body and the world. Every sensor modality is characterized by ''the structure of the rules governing the sensory changes produced by various motor actions, that is, what we call the sensorimotor contingencies'' (O'Regan & Noë, 2001, p. 941).

The enactivist approach suggests modeling a cognitive agent on the basis of sensorimotor interactions with the environment. This paper is an attempt in that direction. In the next section, we introduce a new type of algorithm that does not separate perception from action, called an Enactive Markov Decision Process (EMDP). An EMDP provides a useful conceptual framework for designing agents capable of intrinsically-motivated self-programming as they interact with their environment. We qualify such self-programming as *sensorimotor* because it consists of learning a series of sensorimotor schemes that are subsequently executed as programs. We argue that sensorimotor self-programming opens the way to constitutive autonomy.

While acknowledging that EMDP problems are intractable in the general case, we present two instances in which the coupling with the environment allows the agent to learn to master sensorimotor contingencies within a reasonable frame. The first is called a hierarchical sequential EMDP problem. The second is called a Spatial Enactive Markov Decision Process (SEMDP). A SEMDP is intended to model an agent interacting with an environment that has a Euclidian spatial structure, such as the real world. This work leads us to propose a cognitive architecture dedicated to agents confronted with SEMDP problems, called the Enactive Cognitive Architecture (ECA).

2. Formalism for enactive learning problems

The philosophy of an EMDP is that the agent tries to enact an *intended sensorimotor scheme*, and is informed by the environment whether this intended scheme was indeed enacted, or whether another scheme was enacted instead. In the former case, the intended scheme is considered *successfully enacted*; in the latter case, the intended enaction failed and another scheme was *actually enacted* instead. While EMDP problems differ from reinforcement learning problems, we present them using a similar formalism to allow for comparison.

2.1. Enactive Markov Decision Processes

Formally, we define an EMDP as a tuple (S, I, q, v) in which S is the set of environment states; I is the set of *primitive interactions* offered by the coupling between the agent and the environment; q is a probability distribution such that $q(s_{t+1}| s_t, i_t)$ gives the probability that the environment transitions to state $s_{t+1} \in S$ when the agent chooses interaction $i_t \in I$ in state $s_t \in S$; and v is a probability distribution such that $v(e_t| s_t, i_t)$ gives the probability that the agent receives the input $e_t \in I$ after choosing i_t in state s_t. We call i_t the *intended interaction* because it represents the sensorimotor scheme that the agent intends to enact at the beginning of step t; i_t constitutes the agent's output sent to the environment. We call e_t the *enacted interaction* because it represents the sensorimotor scheme that the agent records as actually enacted at the end of step t; e_t constitutes the agent's input received from the environment. If the enacted interaction equals the intended interaction ($e_t = i_t$) then the attempted enaction of i_t is considered a *success*, otherwise, it is considered a *failure*.

As an example, primitive interactions may represent tactile sensorimotor schemes, which consist of the combination of a movement and the sensory stimulation generated by the movement. When the agent tries to enact a tactile interaction, this may result in a success if the agent indeed touched something, or in a failure if the agent touched nothing, in which case the actually enacted interaction represents the sensorimotor scheme of touching nothing. Fig. 1 shows the EMDP cycle and algorithm. A complete EMDP example will be presented later in Fig. 7a and b.

There are four formal differences between an EMDP and a Partially Observable Markov Decision Process (POMDP)