ARTICLE IN PRESS

Available online at www.sciencedirect.com

ScienceDirect

Cognitive Systems Research xxx (2017) xxx–xxx

Cognitive Systems RESEARCH

www.elsevier.com/locate/cogsys

The knowledge level in cognitive architectures: Current limitations and possible developments

Antonio Lieto [a,b,*], Christian Lebiere [c], Alessandro Oltramari [d]

[a] *University of Turin, Dip. di Informatica, Torino, Italy*
[b] *ICAR-CNR, Palermo, Italy*
[c] *Carnegie Mellon University, Department of Psychology, Pittsburgh, USA*
[d] *Bosch Research and Technology Center, Pittsburgh, USA*

Received 15 November 2016; received in revised form 7 May 2017; accepted 7 May 2017

Abstract

In this paper we identify and characterize an analysis of two problematic aspects affecting the representational level of cognitive architectures (CAs), namely: the limited *size* and the homogeneous *typology* of the encoded and processed knowledge. We argue that such aspects may constitute not only a technological problem that, in our opinion, should be addressed in order to build artificial agents able to exhibit intelligent behaviors in general scenarios, but also an epistemological one, since they limit the plausibility of the comparison of the CAs' knowledge representation and processing mechanisms with those executed by humans in their everyday activities. In the final part of the paper further directions of research will be explored, trying to address current limitations and future challenges.

Keywords: Knowledge representation; Cognitive architectures; Knowledge heterogeneity; Knowledge processing

1. Introduction

Handling a considerable amount of knowledge, and selectively retrieving it according to the needs emerging in different situational scenarios, is an important aspect of human intelligence. For this task, in fact, humans adopt a wide range of heuristics (Gigerenzer & Todd, 1999) due to their bounded rationality (Simon, 1955). In this perspective, one of the requirements that should be considered for the design, the realization and the evaluation of intelligent cognitively-inspired systems should consist in their ability to heuristically identify, retrieve, and process, from the general knowledge stored in their artificial Long Term Memory (LTM), that one which is synthetically and contextually relevant. This requirement, however, is often neglected. Currently, artificial cognitive systems and architectures are not able, de facto, to deal with complex knowledge structures that can be even slightly comparable to the knowledge heuristically managed by humans. In this paper we will argue that this is not only a technological problem but also, in the light of the distinction between functionalist and structuralist models of cognition, an epistemological one. The rest of the paper is organized as follows: Section 2 introduces the two main problematic aspects concerning the knowledge level in cognitive architectures, namely the *size* and the *homogeneous typology* of the encoded knowledge. Section 3 provides a focused review of the Knowledge Level of four of the most well known and widely used cognitive architectures (namely SOAR, ACT-R, CLARION and Vector-LIDA) by pointing out the respective

* Corresponding author at: University of Turin, Dip. di Informatica, Torino, Italy.
E-mail address: lieto.antonio@gmail.com (A. Lieto).

http://dx.doi.org/10.1016/j.cogsys.2017.05.001
1389-0417/

differences and, in the light of our axis of analysis, their problematic issues.[1] In doing so we will illustrate the main attempts that have been proposed to address such problems and we will highlight the current limitations of such proposals. In the final sections, we present an overview of three different alternative approaches that can provide a possible solution for dealing with, jointly, both the size and the knowledge homogeneity problems: namely the Semantic Pointer Perspective (Section 4), the idea of Conceptual Space as intermediate level of representation connecting connectionist and symbolic approaches (Section 5) and the novel versions of the Hybrid Neuro Symbolic Approaches currently developed in the field of CAs (Section 6). Interestingly all such proposals converge in suggesting that the neural level of representation can be considered irrelevant for attacking the above mentioned problems, and suggest to address these issues by operating at more transparent and abstract levels of representation. Section 7, finally, considers the dual process hypothesis as a possible reference framework for the integration of different types of knowledge processing mechanisms assumed to cooperate in a CA adopting a heterogeneous representational perspective. As we will show, the advantages provided by the adoption of this approach are still not completely clear and deserve further investigations.

2. Open issues: Knowledge size and knowledge homogeneity

Current cognitive artificial systems and architectures are not equipped with knowledge bases comparable with the conceptual knowledge that humans possess and use in the everyday life. From an epistemological perspective this shortcoming represents a problem: in fact, endowing cognitive agents with more realistic knowledge bases, in terms of both the size and the type of information encoded, would allow, at least in principle, to test the artificial systems in situations closer to those encountered by humans in real life. This problem becomes more relevant if we take into account the knowledge level[2] of Cognitive Architectures (Newell, 1982; Newell, 1994). While cognitively-inspired systems, in fact, could be designed to deal with only domain-specific information (e.g. a computer simulation of a poker player), Cognitive Architectures (CA), on the other hand, have also the goal and the general objective of testing - computationally - the general models of mind they implement. Therefore, if such architectures only process a simplistic amount (and a limited typology) of knowledge, the structural mechanisms that they implement concerning knowledge processing tasks (e.g., retrieval, learning, reasoning, etc.) can be only loosely evaluated, and compared, w.r.t. those used by humans in similar knowledge-intensive situations. In other words, from an epistemological perspective, the explanatory power of their computational simulation is strongly affected (on these aspects see (Cordeschi, 2002; Miłkowski, 2013; Lieto & Radicioni, 2016)). This aspect is problematic since this class of systems, designed according to the "cognition in the loop" approach, aims both at (i) detecting novel and hidden aspects of the cognitive theories by building properly designed computational models of cognition and (ii) at providing technological advancement in the area of Artificial Intelligence (AI) of cognitive inspiration. In this perspective, purely *functionalist* models (Putnam, 1960), based on a weak equivalence (i.e. the equivalence in terms of functional organization) between cognitive processes and AI procedures, are not considered as having a good explanatory power w.r.t. the target cognitive system taken as source of inspiration. Conversely, the development of plausible "structural" models of our cognition (based on a more constrained equivalence between AI procedures and their corresponding cognitive processes) are assumed to be the way to follow in order to build artificial cognitive models able to play both an explanatory role about the theories they implement and to provide advancements in the field of the artificial intelligence research.

By following this line of argument, therefore, we claim that computational cognitive architectures aiming at providing a knowledge level based on the "structuralist" assumption should address, at their representational level, both the problems concerning the limited "size" and "homogeneity" of the encoded knowledge. Let us explore in more details the nature of such aspects: while the size problem is intuitively easy to understand (i.e. it concerns the dimension of the knowledge base available to the agents), the problem concerning the homogeneous typology of the encoded knowledge needs some additional clarification and context. In particular, this problem relies on the theoretical and experimental results coming from Cognitive Science. In this field, different theories about how humans organize, reason and retrieve conceptual information have been proposed. The oldest one, known as "classical" or Aristotelian theory, states that concepts - the building blocks of our knowledge infrastructure - can be simply represented in terms of sets of necessary and sufficient conditions (and this is completely true, for example, for mathematical concepts: e.g. an EQUILATERAL

[1] In the present paper we will leave aside many other aspects (e.g. those related to the knowledge acquisition problems) which are related to, and also affect, the problems in focus.

[2] The description of the 'knowledge problems' affecting the current Cognitive Architectures (i.e. the knowledge size and the knowledge homogeneity, see below) is provided at the knowledge level in the sense intended by Newell (i.e. we point out that, given the current state of affairs, the rational behavior of a cognitive artificial agent adopting such architectures can be predicted as a limited one simply on the basis of the analysis of the content of its available representations, its limited knowledge of its goals, etc.). On the other hand, the possible solutions proposed for dealing with these problems, sketched in the final part of the paper, are focused on what Newell calls the Symbol Level, since they concern the actual information-processing mechanisms that the system uses in order to reach its goal, given the knowledge that possesses. The close relations between these two levels is explained in [3]. According to Newell, this hierarchy of levels (that includes further levels involving the hardware implementation), characterizes the Physical Symbol System architecture (Newell, 1980).