

Representational Limits in Cognitive Architectures

Antonio Lieto

University of Turin, Department of Computer Science, Italy

ICAR-CNR, Palermo, Italy

lieto.antonio@gmail.com

<http://www.antoniolieto.net>

Abstract—This paper proposes a focused analysis on some problematic aspects concerning the knowledge level in General Cognitive Architectures (CAs). In particular, it addresses the problems regarding both the limited size and the homogeneous typology of the encoded (and processed) conceptual knowledge. As a possible way out to face, jointly, these problems, this contribution discusses the possibility of integrating external, but architecturally compliant, cognitive systems into the knowledge representation and processing mechanisms of the CAs.

Keywords—*cognitive architectures, knowledge representation, knowledge level, common-sense reasoning.*

I. INTRODUCTION

The research on Cognitive Architectures (CAs) is a wide and active area involving a plethora of disciplines such as Cognitive Science, Artificial Intelligence, Robotics and, more recently, the area of Computational Neuroscience. CAs have been historically introduced i) “to capture, at the computational level, the invariant mechanisms of human cognition, including those underlying the functions of control, learning, memory, adaptivity, perception and action” [1] and ii) to reach human level intelligence, also called General Artificial Intelligence, by means of the realization of artificial artifacts built upon them. During the last decades many cognitive architectures have been realized, - such as ACT-R [2], SOAR [3] etc. - and have been widely tested in several cognitive tasks involving learning, reasoning, selective attention, multimodal perception, recognition etc. Despite the recent developments, however, in the last decades the importance of the “knowledge level” [4] has been historically and systematically downsized by this research area, whose interests have been mainly based on the analysis and the development of mechanisms and the processes governing human and (artificial) cognition. The knowledge level in CAs, however, presents several problems that may affect the overall heuristic and epistemological value of such artificial general systems and therefore deserves more attention.

II. TWO PROBLEMS FOR THE KNOWLEDGE LEVEL IN CAs

Handling a huge amount of knowledge, and selectively retrieve it according to the needs emerging in different situational scenarios, represents an important aspect of human intelligence. For this task humans adopt a wide range of heuristics [5] due to their “bounded rationality” [6]. Currently, however, the Cognitive 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 other terms: CAs are general structures without a

general content. This means that the knowledge embedded and processed in such architectures is usually very limited, ad-hoc built, domain specific, or based on the specific tasks they have to deal with. Thus, every evaluation of the artificial systems relying upon them, is necessarily task-specific and do not involve not even the minimum part of the full spectrum of processes involved in the human cognition when the “knowledge” comes to play a role. As a consequence, the structural mechanisms that the CAs implement concerning knowledge processing tasks (e.g. that ones of retrieval, learning, reasoning etc.) can be only loosely evaluated, and compared w.r.t. that ones 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 [7,8]. Such knowledge limitation, in our opinion, does not allow to obtain significant advancements in the cognitive science research about how the humans heuristically select and deal with the huge amount of knowledge that possess when they have to make decisions, reason about a given situation or, more in general, solve a particular cognitive task involving several dimensions of analysis. This problem, as a consequence, also limits the advancement of the research in the area of General Artificial Intelligence of cognitive inspiration.

The “content” limit of the cognitive architectures has been recently pointed out in literature [1] and some technical solutions for filling this “knowledge gap” have been proposed [9]. In particular the use of ontologies and of semantic formalisms and resources (such as DBpedia) has been seen as a possible solution for providing effective content to the structural knowledge modules of the cognitive architectures. Some initial efforts have been done in this sense but cover only part of the “knowledge problem” in CAs (i.e. the one concerning the limited “size” of the adopted knowledge bases). However, also these solutions, do not address another relevant aspect affecting the knowledge level of CAs: namely, the problem concerning the “knowledge homogeneity” issue. In other terms: the type of knowledge represented and manipulated by most CAs (including those provided with extended knowledge modules) is usually homogeneous in nature. It mainly covers, in fact, only the so called “classical” part of conceptual information (that one representing concepts in terms of necessary and sufficient information and compliant with ontological semantics (see [10]) on these aspects). On the other hand, the so called “common-sense” conceptual components of our knowledge (i.e. those that, based on the results from the cognitive science, allow to characterize concepts in terms of “prototypes”, “exemplars” or “theories”)

is largely absent in such computational frameworks. The possibility of representing and handling, in an integrated way, an heterogeneous amount of common sense conceptual representations (and the related reasoning mechanisms), in fact, is not sufficiently addressed both by the symbolic-based "chunk-structures" adopted by the most common general CAs (e.g. SOAR) and by fully connectionist architectures (e.g. LEABRA). This aspect is problematic also in the hybrid solutions adopted by CAs such as CLARION [11] or ACT-R (the different reasons leading to a non satisfactory treatment of this aspect are detailed in [12]). This type of knowledge, however, is exactly the type of "cognitive information" crucially used by humans for heuristic reasoning and decision making. This paper presents an analysis of the current situation by proposing a comparison of the representational level of SOAR, ACT-R, CLARION and Vector-LIDA. Finally, we suggest that a possible way out to deal with this problem could be represented by the integration of external cognitive systems into the knowledge representation and processing mechanisms of general cognitive architectures. Some initial efforts in this direction, have been proposed (see e.g. [13, 14]) and will be presented and discussed.

III. KNOWLEDGE REPRESENTATION IN CAs

In the following we provide a short overview of: SOAR [3], ACT-R [2], CLARION [11] and LIDA [15] (in its novel version known as Vector-LIDA [16]). The choice of these architecture has been based on the fact that they represent some of the most widely used systems (adopted in scenarios ranging from robotics to video-games) and their representational structures present some relevant differentiations that are interesting to investigate in the light of the issues raised in this paper. By analyzing, in brief, such architectures we will exclusively focus on the description of their representational frameworks since a more comprehensive review of their whole mechanisms is out of the scope of the present contribution (detailed reviews of their mechanisms are described in [17]; and [18]). We will show how all of them are affected, at different levels of granularity, by both the size and the knowledge homogeneity problems.

A. SOAR

SOAR is one of the oldest cognitive architectures. This system was considered by Newell a candidate for a Unified Theory of Cognition [19]. One of the main themes in SOAR is that all cognitive tasks can be represented by problem spaces that are searched by production rules grouped into operators. These production rules are red in parallel to produce reasoning cycles. From a representational perspective, SOAR exploits symbolic representations of knowledge (called chunks) and use pattern matching to select relevant knowledge elements. Basically, when a production match the contents of declarative (working) memory the rule fires and then the content from the declarative memory (called Semantic Memory in SOAR) is retrieved. This system adheres strictly to the Newell and Simon's physical symbol system hypothesis which assumes that symbolic processing is a necessary and sufficient condition for intelligent behavior. The SOAR system encounter, in general, the standard problems affecting symbolic formalisms at the representational level: it is not well equipped to deal with common-sense knowledge representation and reasoning (since approximate comparisons are hard and computationally

intensive to implement with graph-like representations), and, as a consequence, the typology of encoded knowledge is biased towards the "classical" (but unsatisfactory) representation of concepts in terms of necessary and sufficient conditions [10]. This characterization, however, is problematic for modelling real world concepts and, on the other hand, the so called common-sense knowledge components (i.e. those that, allow to characterize and process conceptual information in terms of *typicality* and involving, for example, prototypical and exemplar based representations and reasoning mechanisms) is largely absent. This problem arises despite the fact that the chunks in SOAR can be represented as a sort of frame-like structures containing some common-sense (e.g. prototypical) information [12]. W.r.t. to the size problem, the SOAR knowledge level is also problematic. SOAR agents, in fact, are not endowed with general knowledge and only process ad-hoc built (or task-specific learned) symbolic knowledge structures.

B. ACT-R

ACT-R is a cognitive architectures explicitly inspired by theories and experimental results coming from human cognition. Here the cognitive mechanisms concerning the knowledge level emerge from the interaction of two types of knowledge: declarative knowledge, that encodes explicit facts that the system knows, and procedural knowledge, that encodes rules for processing declarative knowledge. In particular, the declarative module is used to store and retrieve pieces of information (called chunks, featured by a type and a set of attribute-value pairs, similar to frame slots) in the declarative memory. ACT-R employs a wide range of sub-symbolic processes for the activation of symbolic conceptual chunks representing the encoded knowledge. Finally, the central production system connects these modules by using a set of IF-THEN production rules using a set of IF-THEN production rules. Differently from SOAR, ACT-R allows to represent the information in terms of prototypes and exemplars and allow to perform, selectively, either prototype or exemplar-based categorization. This means that this architecture allows the modeller to manually specify which kind of categorization strategy to employ according to his specific needs. Such architecture, however, only partially addresses the homogeneity problem since it does not allow to represent, jointly, these different types of common-sense representations for the same conceptual entity (i.e. it does not assume a heterogeneous perspective). As a consequence, it is also not able to autonomously decide which of the corresponding reasoning procedures to activate (e.g. prototypes or exemplars) and to provide a framework able to manage the interaction of such different reasoning strategies (however its overall architectural environment provides, at least in principle, the possibility of implementing cascade reasoning processes triggering one another). Even if, in such architecture, some attempts exist concerning the design of harmonization strategies between different types of common-sense conceptual categorizations (e.g. exemplars-based and rule based, see [20]) however they do not handle the problem concerning the interaction of the prototype or exemplars-based processes according to the results coming from the experimental cognitive science (for example: the *old item effect*, privileging exemplars w.r.t. prototypes is not modelled. See again [12] for a detailed analysis of this aspect). Summing up: w.r.t. the knowledge homogeneity problem, the components needed to

fully reconcile the Heterogeneity approach with ACT-R are present, however they have not been fully exploited yet. Regarding the size problem: as for SOAR, ACT-R agents are usually equipped with task-specific knowledge and not with general cross-domain knowledge. In this respect some relevant attempts to overcome this limitation have been recently done by extending the Declarative Memory of the architecture. They will be discussed in section *E* along with their current implications.

C. CLARION

CLARION is a hybrid cognitive architecture based on the dual-process theory of mind. From a representational perspective, processes are mainly subject to the activity of two sub-systems, the Action Centered Sub-system (ACS) and the Non-Action Centered Sub-system (NACS). Both sub-systems store information using a two-layered architecture, i.e., they both include an explicit and an implicit level of representation. Each top-level chunk node is represented by a set of (micro)features in the bottom level (i.e., a distributed representation). The (micro)features (in the bottom level) are connected to the chunk nodes (in the top level) so that they can be activated together through bottom-up or top-down activation. Therefore, in general, a chunk is represented by both levels: using a chunk node at the top level and distributed feature representation at the bottom level. W.r.t. to the knowledge size and homogeneity problems, CLARION, encounter problems with both these aspects since i) there are no available attempts aiming at endowing such architecture with a general and cross-domain knowledge ii) the dual-layered conceptual information does not provide the possibility of encoding (manually or automatically via learning cycles) the information in terms of the heterogeneous classes of representations presented in the section 2. In particular: the main problematic aspect concerns the representation of the common-sense knowledge components. As for SOAR and ACT-R, also in CLARION the possible co-existence of typical representations in terms of prototypes, exemplars and theories (and the interaction among them) is not treated. In terms of reasoning strategies, notwithstanding that the implicit knowledge layer based on neural network representations can provide forms of non monotonic reasoning (e.g. based on similarity), such kind of similarity-based reasoning is currently not grounded on the mechanisms guiding the decision choices followed, for example, by prototype or exemplars-based reasoning.

D. Vector-LIDA

Vector LIDA is a cognitive architecture employing, at the representational level, high-dimensional vectors and reduced descriptions. High-dimensional vector spaces have interesting properties that make them attractive for representations in cognitive models. The distribution of the distances between vectors in these spaces, and the huge number of possible vectors, allow noise-robust representations where the distance between vectors can be used to measure the similarity (or dissimilarity) of the concepts they represent. Moreover, these high-dimensional vectors can be used to represent complex structures, where each vector denotes an element in the structure. However, a single vector can also represent one of these same complex structures in its entirety by implementing a reduced description, a mechanism to encode complex

hierarchical structures in vectors or connectionist models. These reduced description vectors can be expanded to obtain the whole structure, and can be used directly for complex calculations and procedures, such as making analogies, logical inference, or structural comparison. Vectors in this framework are treated as symbol-like representations, thus enabling different kind of operations executed on them (e.g. simple forms of compositionality via vectors blending). Vector-LIDA, encounters the same limitations of the other CAs since i) its agents are not equipped with a general cross-domain knowledge and therefore can be only used in very narrow tasks (their knowledge structure is either ad hoc build or ad hoc learned). Additionally, this architecture does not address the problem concerning the heterogeneity of the knowledge typologies. In particular its knowledge level does not represent the common-sense knowledge components such as prototypes and exemplars (and the related reasoning strategies). In fact, as for CLARION, despite vector-representations allow to perform many kind of approximate comparisons and similarity-based reasoning (e.g. in tasks such as categorization), the peculiarity concerning prototype or exemplars based representations (along with the the design of the interaction between their different reasoning strategies) are not provided. In this respect, however an element that is worth-noting is represented by the fact that the Vector-LIDA representational structures are very close to the framework of Conceptual Spaces. Conceptual Spaces are a geometric knowledge representation framework proposed by Peter Gärdenfors [21]. They can be thought as a particular class of vector representations where knowledge is represented as a set of quality dimensions, and where a geometrical structure is associated to each quality dimension. They are discussed in more detail in section 5. The convergence of the Vector-LIDA representation towards Conceptual Spaces could enable, in such architecture, the possibility of dealing with at least prototype and exemplars-based representations and reasoning, thus overcoming the knowledge homogeneity problem.

E. Attempts to Overcome the Knowledge Limits

As mentioned above, some initial efforts to deal with the limited knowledge availability for agents endowed with cognitive architecture have been done. In particular, within Mind'sEye program (a DARPA founded project), the knowledge layers of ACT-R architecture have been semantically extended with an external ontological content coming from three integrated semantic resources composed by the lexical databases WordNet [22], FrameNet [23] and by a branch of the top level ontology DOLCE [24] related to the event modelling. In this case, the amount of semantic knowledge selected for the realization of the Cognitive Engine (one of the systems developed within the MindEye Program) and for its evaluation, despite by far larger w.r.t. the standard ad-hoc solutions, was tailored on the specific needs of the system itself. It, in fact, was aimed at solving a precise task of event recognition through a video-surveillance intelligent machinery; therefore only the ontological knowledge about the events was selectively embedded in it. While this is a reasonable approach in an applicative context, still does not allow to test the general cognitive mechanisms of a Cognitive Architecture on a general, multi faceted and multi-domain, knowledge. Therefore it does not allow to evaluate *strictu sensu* to what extent the designed heuristics allowing to retrieve and process, from a massive and composite knowledge

base, conceptual knowledge can be considered *satisfying* w.r.t. the human performances. More recent works have tried to completely overcome at least the size problem of the knowledge level. To this class of works belongs that one proposed by Salvucci [9] aiming at enriching the knowledge model of the Declarative Memory of ACT-R with a world-level knowledge base such as DBpedia (i.e. the semantic version of Wikipedia represented in terms of ontological formalisms) and a previous one proposed in [25] presenting an integration of the ACT-R Declarative and Procedural Memory with the Cyc ontology [26] (one of the widest ontological resources currently available containing more than 230,000 concepts). Both the wide-coverage integrated ontological resources, however, represents conceptual information in terms of symbolic structures and encounter the standard problems affecting this class of formalisms and discussed above. Some of these limitations can be, in principle, partially overcome by such works, since the integration of such wide-coverage ontological knowledge bases with the ACT-R Declarative Memory allows to preserve the possibility of using the common-sense conceptual processing mechanisms available in that architecture (e.g. prototype and exemplars based). Therefore, in principle, dealing with the size problem also allows to address some aspects concerning the heterogeneity problem. Still, however, remains the problem concerning the lack of the representation of common-sense information to which such common-sense architectural processes can be applied: e.g. a conceptual retrieval based on prototypical traits (i.e. a prototype-based categorization) cannot be performed on such integrated ontological knowledge bases since these symbolic systems do not represent at all the *typical* information associated to a given concept ([12] presents an experiment on this aspect). In addition, as already mentioned, it remains not yet addressed the problem concerning the interaction, in a general and principled way, of the different types of common-sense processes involving different representations of the same conceptual entity. In the light of the arguments presented above it can be argued, therefore, that the current proposed solutions for dealing with the knowledge problems in CAs are not completely satisfactory. In particular, the integrations with huge world-level ontological knowledge bases can be considered a necessary solution for solving size problem. It is, however, insufficient for dealing with the knowledge homogeneity problem and with the integration of the common-sense conceptual mechanisms activated on heterogeneous bodies of knowledge, as assumed in the heterogeneous representational perspective. In the next sections we outline a possible alternative solution that, despite being not yet fully developed is, in perspective, suitable to account for both for the heterogeneous aspects in conceptualization and for the size problems.

IV. INTEGRATING EXTERNAL COGNITIVE SYSTEMS IN CA

Recently some available conceptual categorization systems, explicitly assuming the heterogeneous representational hypothesis and integrated with wide-coverage knowledge bases (such as Cyc) have been developed and integrated with the knowledge level of available CAs. For our purposes, we will consider here the DUAL PECCS system [13, 14]. We will not discuss the results obtained by such system in tasks of conceptual categorization, since they have been already presented elsewhere [14]. We shall briefly focus, in the following, on the representational level of the system.

The knowledge level of DUAL PECCS is heterogeneous in nature since it is explicitly based and designed on the assumption that concepts are “heterogeneous proxytypes” [27] and, as such, they are composed by heterogeneous knowledge components selectively and contextually activated in working memory. In particular, by following the proposal presented in [28, 29], the representational level of DUAL PECCS couples Conceptual Spaces representations and ontological knowledge (consisting in the Cyc ontology) for the same conceptual entity. Conceptual Spaces [21] is used to represent and process the common-sense conceptual information. In such framework, to each quality dimension is associated a geometrical (topological or metrical) structure. In some cases, such dimensions can be directly related to perceptual mechanisms; examples of this kind are temperature, weight, brightness, pitch. In other cases, dimensions can be more abstract in nature. In this setting, concepts correspond to convex regions, and regions with different geometrical properties correspond to different sorts of concepts [21]. Here, prototypes and prototypical reasoning have a natural geometrical interpretation: prototypes correspond to the geometrical centre of a convex region (the centroid). Also exemplars-based representation can be represented as points in a multidimensional space, and their similarity can be computed as the intervening distance between each two points, based on some suitable metrics (such as Euclidean and Manhattan distance etc.). The ontological component, on the other hand, is used to provide and process the “classical” knowledge component for the same conceptual entity.

The representational level of DUAL PECCS (and the corresponding knowledge processing mechanisms) has been successfully integrated with the representational counterpart of some available CAs [14, 30] by extending, *de facto*, the knowledge representation and processing capabilities of cognitive architectures based on diverse representational assumptions. One of the main novelties introduced by DUAL PECCS (and therefore one of the main advantages obtained by the CAs extended with such external cognitive system) consists in the fact that it is explicitly designed the flow of interaction between common-sense categorization processes (based on prototypes and exemplars and operating on conceptual spaces representations) and the standard deductive processes (operating on the ontological conceptual component). The harmonization regarding such different classes of mechanisms has been devised based on the tenets coming from the dual process theory of reasoning [31, 32]. Additionally, in DUAL PECCS, also the interaction of the categorization processes occurring *within* the class of non monotonic categorization mechanisms (i.e. prototypes and exemplars-based categorization) has been devised and is dealt with at the Conceptual Spaces level. This latter aspect is of particular interest in the light of the multifaceted problem concerning the heterogeneity of the encoded knowledge. In fact, since the design of the interaction of the the different processes operating with heterogeneous representations still represents, as seen before, a largely unaddressed problem in current CAs, this system shows the relative easiness that its knowledge framework (and, in particular, the Conceptual Spaces component) provides to naturally model the dynamics between prototype and exemplars-based processes. For what concerns the size problem, finally, the possible grounding of the Conceptual Spaces representational component with symbolic structures enables the integration with wide-coverage

knowledge bases such as Cyc. Thus, the solution adopted in DUAL PECCS is, in principle, able to deal with both the size and the knowledge homogeneity problems affecting the CAs. In particular, the extension of the Declarative Memories of the current CAs with this external cognitive system allowed to empower the knowledge processing and categorization capabilities of such general architectures (an important role, in this respect, is played by the Conceptual Spaces component). Despite there is still room of improvements and further investigations, this seems a promising way to deal with the both the knowledge problems discussed in this paper.

ACKNOWLEDGMENTS

I thank the reviewers of the EuCognition Conference for their useful comments. The arguments presented in this paper have been discussed in different occasions with Christian Lebiere, Alessandro Oltramari, Antonio Chella, Marcello Frixione, Peter Gärdenfors, Valentina Rho and Daniele Radicioni. I would like to thank them for their feedback.

REFERENCES

1. Oltramari A., Lebiere C., Pursuing Artificial General Intelligence By Leveraging the Knowledge Capabilities Of ACT-R, AGI 2012 (5th International Conference on "Artificial General Intelligence"), Oxford, 2012.
2. Anderson, J.R., Bothell, D., Byrne, M.D., Douglass, S., Lebiere, C., & Qin, Y., An integrated theory of the mind, *Psychological Review*, 111(4), 1036-1060, <http://dx.doi.org/10.1037/0033-295X.111.4.1036>, 2004.
3. Laird, John E., *The Soar Cognitive Architecture*, MIT Press, 2012.
4. Newell, A. The knowledge level, *Artificial intelligence* 18 (1), 87-127, 1982.
5. Gigerenzer, G., Todd, P., *Simple Heuristics that make us smart*", Oxford University Press, 1999.
6. Simon, H., *A Behavioral Model of Rational Choice*, in *Mathematical Essays on Rational Human Behaviour in Social Setting*. NY: Wiley, 1957.
7. Minkowski M., *Explaining the Computational Mind*, MIT Press, 2013.
8. Lieto, A., Radicioni, D.P. From Human to Artificial Cognition and back: Challenges and Perspectives of Cognitively-Inspired AI systems, *Cognitive Systems Research*, 39 (2), pp. 1-3, 2016.
9. Salvucci, D., Endowing a Cognitive Architecture with World Knowledge, *Proceedings of the 36th Annual Meeting of the CogSci Soc.*, 2014.
10. Frixione, M., Lieto A., Representing concepts in formal ontologies: Compositionality vs. typicality effects, *Logic and Logical Philosophy* 21 (4) (2012) 391-414, 2012.
11. Sun, R. The CLARION cognitive architecture: Extending cognitive modeling to social simulation. *Cognition and multi-agent interaction* pp. 79-99, 2006.
12. Lieto, A., Lebiere C., Oltramari, A. The Knowledge Level in Cognitive Architectures: Current Limitations and Possible Developments, *Cognitive Systems Research*, *Cognitive Systems Research*, forthcoming.
13. Lieto, A., Radicioni, D.P., Rho, V., A Common Sense Conceptual Categorization System Integrating Heterogeneous Proxytypes and the Dual Process of Reasoning. In *Proc. of IJCAI 2015*, AAAI Press, 2015.
14. Lieto, A., D.P. Radicioni, V. Rho, Dual PECCS: A Cognitive System for Conceptual Representation and Categorization *Journal of Experimental and Theoretical Artificial Intelligence*, Taylor and Francis. doi: <http://dx.doi.org/10.1080/0952813X.2016.1198934>, 29 (2), 2017.
15. Franklin, S., F. Patterson, F., The Lida architecture: Adding new modes of learning to an intelligent, autonomous, software agent, 764-1004, 2006.
16. Snaider, S. Franklin, Vector Lida, *Procedia Computer Science* 4,188-203, 2014.
17. Vernon, D, Metta, G., Sandini, G., A survey of artificial cognitive systems: Implications for the autonomous development of mental capabilities in computational agents, *IEEE Transactions on Evolutionary Computation* 11 (2), 2007.
18. Langley, P., Laird, J., Rogers, S., *Cognitive architectures: Research issues and challenges*, *Cognitive Systems Research* 10 (2), 141-160, 2009.
19. Newell, A. *Unified theories of cognition*. Cambridge, MA: Harvard University Press, 1990.
20. Anderson, J., Betz, J., A hybrid model of categorization, *Psychonomic Bulletin & Review* 8 (4) 629-647, 2001.
21. Gärdenfors, *Conceptual spaces: The geometry of thought*, MIT press, 2000.
22. Fellbaum, C. (ed.): *WordNet - An Electronic Lexical Database*. Cambridge, Massachusetts. MIT Press, 1998.
23. Fillmore, C.J., *The case for case*. Bach, E., Harms, T. eds. *Universals in Linguistic Theory*. New York: Rinehart and Wiston, 1968.
24. Masolo, C., Borgo, S., Gangemi, A., Guarino, N., & Oltramari, A., *Wonderweb deliverable d18, ontology library (final)*. ICT project, 33052, 2003.
25. Ball, J., Rodgers, S., Gluck, K., Integrating act-r and cyc in a large-scale model of language comprehension for use in intelligent agents, in: *AAAI workshop*, pp. 19-25, 2004.
26. Lenat, D., Cyc: A large-scale investment in knowledge infrastructure, *Communications of the ACM* 38 (11), 33-38, 1995.
27. Lieto A., A Computational Framework for Concept Representation in Cognitive Systems and Architectures: Concepts as Heterogeneous Proxytypes, *Procedia Computer Science*, 41, 6-14, <http://dx.doi.org/10.1016/j.procs.2014.11.078>, 2014.
28. Frixione M, Lieto A., Towards an Extended Model of Conceptual Representations in Formal Ontologies: A Typicality-Based Proposal, *Journal of Universal Computer Science* 20 (3) (2014) 257-276, 2014.
29. Lieto, A., Chella, A., Frixione, M., *Conceptual Spaces for Cognitive Architectures: A Lingua Franca for Different Levels of Representation*, In *Biologically Inspired Cognitive Architectures*, 19 (2), 1-9, 2017.
30. Lieto, A., Radicioni, D.P., Rho, V., Mensa, E., *Towards a Unifying Framework for Conceptual Representation and Reasoning in Cognitive Systems*, in *Intelligenza Artificiale*, forthcoming.
31. Evans, J., Frankish, K., *In two minds: Dual processes and beyond*, Oxford University Press, 2009.
32. Stanovich, K., West, R., *Advancing the rationality debate*, *Behavioral and brain sciences*, 23 (05), 701-717, 2000.