Towards a Unifying Framework for Conceptual Represention and Reasoning in Cognitive Systems

Undefined 1 (2016) 1–5
IOS Press

Towards a Unifying Framework for Conceptual Represention and Reasoning in Cognitive Systems

Antonio Lieto [a,b,*], Daniele Radicioni [a], Valentina Rho [a] and Enrico Mensa [a]
[a] *Dipartimento di Informatica, Università di Torino, Italy*
E-mail: {lieto,mensa,radicion,rho}@di.unito.it
[b] *ICAR-CNR, Palermo, Italy*

Abstract.
In this paper we present the rationale adopted for the integration of the knowledge level of DUAL-PECCS, a cognitive system for conceptual representation and categorization, with two different cognitive architectures: SOAR and LIDA. In previous works we already showed how the representational and reasoning framework adopted in DUAL-PECCS was integrable with diverse cognitive architectures, i.e. ACT-R and CLARION, making different representational assumptions and adopting diverse knowledge processing mechanisms. The additional integrations presented here suggest that the underlying knowledge representation and reasoning structure adopted in DUAL-PECCS can be used as a unifying framework for the knowledge level of agents endowed with different cognitive architectures. The current version of the system has been experimentally assessed in a task of conceptual categorization where a target concept illustrated by a simple common-sense linguistic description had to be identified by resorting to a mix of categorization strategies. The output has then been compared to human and artificial responses. The novel integration allowed us to extend our previous evaluation.

Keywords: Knowledge Representation, Categorization, Conceptual Spaces, Cognitive Architectures, SOAR, LIDA, Heterogeneous Proxytypes, Prototypes, Exemplars, Common-Sense Reasoning, Dual Process Theory.

1. Introduction

In this work we present a novel version of the DUAL-PECCS system [40], integrated with two diverse Cognitive Architectures: SOAR [26] and LIDA [16]. The result of such novel integrations suggests that the underlying knowledge representation and processing mechanisms adopted in DUAL-PECCS can be considered as a plausible candidate for providing a unifying representation and reasoning framework for agents endowed with different cognitive architectures.

The DUAL-PECCS conceptual architecture is based on a cognitively-inspired categorization system able to perform, in an integrated way, two well-known types of common-sense conceptual reasoning in human cognition: *prototypical* and *exemplars-based* reasoning. The system relies on a representational and reasoning framework designed and implemented according to insights and experimental evidences coming from Cognitive Science. From a representational perspective it relies on the hypothesis of conceptual structures represented as *heterogeneous proxytypes*, proposed and developed in the area of BICA (Biologically Inspired Cognitive Architectures) [31]. From a reasoning perspective, it integrates both types of the above mentioned human-like common-sense reasoning (i.e., prototypical and exemplars-based reasoning) with standard monotonic categorization procedures according to the tenets coming from the dual process theory of reasoning [25].

*Corresponding authors: Antonio Lieto, Daniele P.Radicioni and Valentina Rho. E-mail: lieto,radicion,rho@di.unito.it

This work has the following main strengths: it extends our previous work [39,40] where the integration of the common-sense representation and reasoning system, based on the above mentioned assumptions, was first attempted into the ACT-R [2] and CLARION [59] cognitive architectures. Specifically, that system is now integrated and tested into two additional cognitive architectures, SOAR and LIDA, implementing different assumptions on the underlying representational and reasoning structures governing human (and artificial) cognition [28]. Notably, despite the type of the provided integration is the same of our previous work (i.e., we provide an extension of the representations provided in the Declarative Memories of each architecture), the two newly integrated CAs are also very different (from a representational perspective) w.r.t. the previously integrated ones[1]. This aspect represents, therefore, an additional advancement that allows us to fully collocate the employed representational framework adopted in DUAL-PECCS within the recent debate concerning the identification of a Standard Model of Mind [27]. In the following, the integration details are analysed, and the two implementations are assessed in a conceptual categorization test. Finally, the whole workflow has been significantly improved with richer information extraction procedures to automatically acquire the knowledge base dealing with common-sense representation and reasoning tasks.

This work is organized as follows: in the first sections we sketch (by referring to previous work for additional details) the main elements inspiring our system and its theoretical bases as well as its overall architecture; in Section 5 we describe the improvements related to the automatic generation of common-sense knowledge; in Section 7 we show how our hybrid system for conceptual categorization was integrated into SOAR and LIDA, and finally we describe the evaluation experiments and elaborate on the future works.

[1]It is worth-noting that, in this setting, the provided integration should not be seen as a mere implementation fact. On the contrary, the implemented integration of our framework with the software instantiations of the involved cognitive architecture follows a previous, and more complex, integration developed between the abstract models of cognition assumed by each architecture and our system. As a consequence, the resulting integration is provided by respecting the architectural constraints of each of these systems. This means, in other words, that our work provides evidence that the adopted representational framework is cognitively compliant with such diverse agent architectures.

2. Types of Conceptual Representations

In Cognitive Science different theories about how humans organise, reason and retrieve conceptual information have been proposed. The oldest one, known as "classical", states that concepts can be simply represented in terms of sets of necessary and sufficient conditions. In the mid '70s of the last Century, however, Rosch's experimental results demonstrated its inadequacy for ordinary –or *common sense*– concepts, that cannot be described in terms of necessary and sufficient traits [57].

In particular, Rosch's results showed that ordinary concepts are organized in our mind in terms of *prototypes*. Since then different theories of concepts have been proposed to explain different representational and reasoning aspects concerning the problem of *typicality*. We recall here the prototype theory and the exemplars theory.[2] According to the *prototype* view, knowledge about categories is stored in terms of prototypes, i.e., in terms of some representation of the most typical instance of the category. In this view, the concept *bird* should coincide with the representation of a typical bird (e.g., a robin). According to the *exemplar* view, a given category is mentally represented as a set of specific exemplars explicitly stored in memory: the mental representation of the concept *bird* is a set containing the representations of (some of) the birds we encountered during our past experience. Although these approaches have been largely considered as competing ones, several results (starting from the work of Malt [42]) suggested that human subjects may use, in different occasions, different representations to retrieve and categorize concepts. Such experimental evidences led to the development of the so called "heterogeneous hypothesis" about the nature of concepts, hypothesizing that different types of conceptual representations coexist at the same time: prototypes, exemplars, classical representations, and so on [41].

This hypothesis has been recently extended within the field of knowledge representation applied to biologically inspired cognitive architectures: a novel approach to concept representation has been proposed, considering concepts as "heterogeneous proxytypes" [31]. In this view, conceptual structures in natural and artificial cognitive systems and architectures are assumed to be composed by heterogeneous representations (or *bodies of knowledge*) referring to the

[2]A review of all the typicality-theories is available in [41].

same conceptual entity. Each body of knowledge provides specific types of information and specific access and reasoning procedures to the concept it is referred to. Such heterogeneous representations are *proxytypes* [56], in the sense that they can be contextually activated by external stimuli, coming from the environment, and 'go proxy' in working memory (a sort of temporary buffer available in human and artificial memory structures), for their reference category. The proxyfication may be then the result of activities such as concept identification, recognition, retrieval, and so forth. The different types of conceptual representations hypothesized to coexist in the heterogeneous proxytypes approach are typicality-based representations of a given concept (i.e., prototypes and exemplars-based representations), as well as representations in terms of necessary and/or sufficient conditions. As an example of such type of representational hypothesis let us consider the ordinary concept of *water*. The classical component will contain the information that water is exactly the chemical substance whose formula is H_2O, that is the substance whose molecules have two hydrogen atoms with a covalent bond to the single oxygen atom. On the other hand, the prototypical facet of the concept will grasp that water usually occurs in liquid state, and is a colourless, odourless and tasteless fluid. The exemplar-based representations grasp information on individuals, such as a given instance of *water* presenting, for example, an unusual water color (e.g., red-water). According to the heterogeneous proxytypes approach, the activation in working memory of such conceptual structures is context-dependent: if an agent perceives an instance of *water* in a typical scenario – for instance, in a bottle– the only type of conceptual knowledge that will be activated and proxyfied will be the prototypical knowledge associated to that concept (and not, for example, the classical or exemplar-based information associated to the same conceptual entity).

3. A Dual Process Conceptual Categorization

From a reasoning perspective the heterogeneous hypothesis assumes that the retrieval of the above mentioned representations is driven by different process types. In particular, prototype and exemplar-based retrieval is based on a fast and approximate kind of categorization, and benefits from common-sense information associated to concepts.[3] On the other hand, the retrieval of classical representation of concepts is featured by explicit rule following, and makes no use of common-sense information. These two differing categorization strategies have been widely studied in psychology of reasoning in the frame of the dual process theory, that postulates the co-existence of two different types of cognitive systems [15,25]. The systems of the first type (*type 1*) are automatic, associative, parallel and fast. The systems of the second type (*type 2*) are more recent, conscious, sequential and slow, and featured by explicit rule following. We assume that both systems can be composed in turn by many sub-systems and processes.

4. The System

The DUAL-PECCS system relies on both the heterogeneous proxytypes approach and on the dual process theory of reasoning. The heterogeneous conceptual representation has been implemented as a hybrid knowledge base composed of heterogeneous representations of the same conceptual entities: that is, the hybrid knowledge base includes both common-sense representation (prototypes and exemplars) and classical representations for the same concept. Then, such representations are associated with different sorts of processes: *Type 1* processes have been designed to deal with prototypes- and exemplar-based retrieval and categorization, while *Type 2* processes have been designed to deal with deductive inference.

Following the hypotheses in [18,37] the two sorts of reasoning processes interact, since Type 1 processes are executed first and their results are then refined by Type 2 processes. In the implemented system the typical representational and reasoning functions are assigned to the System 1 (hereafter $\mathcal{S}1$), which executes processes of Type 1, and is associated to the Conceptual Spaces framework [21]. The reasoning functions herein are implemented as similarity calculations in a

[3] If we have to categorize a stimulus with the following features: "it has fur, woofs and wags its tail", the result of a *prototype-based categorization* would be *dog*, since these cues are associated to the prototype of *dog*. Prototype-based reasoning, however, is not the only type of reasoning based on typicality. In fact, if an exemplar corresponding to the stimulus being categorized is available too, it is acknowledged that humans use to classify it by evaluating its similarity w.r.t. the exemplar, rather than w.r.t. the prototype associated to the underlying concepts [19]. This type of common sense categorization is known in literature as *exemplars-based categorization*.

metric space. On the other hand, the classical representational and reasoning functions are assigned to the System 2 (hereafter $\mathcal{S}2$) to execute processes of Type 2, and are associated to a standard symbolic based ontological representation (in our case the OpenCyc ontology [29] was used).

In Section 6 we briefly describe the categorization pipeline of the system by presenting the dynamics of the interaction between $\mathcal{S}1$ and $\mathcal{S}2$ processes. In the following we introduce the two representational and reasoning frameworks used in our system, by focusing on *i)* how typicality information (including both prototypes and exemplars) and their corresponding non monotonic reasoning procedures can be encoded through conceptual spaces; and on *ii)* how classical information can be naturally encoded in terms of formal ontologies.

Conceptual spaces (CSs) are a representational framework where knowledge is represented as a set of *quality dimensions*, and where a geometrical or topological structure is associated to each quality dimension. Instances 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, or standard cosine similarity).[4] In this setting, concepts correspond to convex regions, and regions with different geometrical properties correspond to different sorts of concepts [20,21]. Prototypes have a natural geometrical interpretation in conceptual spaces, in that they correspond to the geometrical centre of a convex region. This can be thought of as a centroid, that is the mean position of all the points in all dimensions. This representation also allows us, given a convex region, to associate each point to a certain centrality degree, that can be interpreted as a measure of its typicality [22]. This framework has been used also to encode the exemplars, represented as points in the multidimensional space. Conceptual spaces can be also used to compute the proximity between any two entities, and between entities and prototypes. Concepts, in this framework, are characterized in terms of *domains* [20,21]. Typical domain examples are color, size, shape, texture. In turn, domain information can be specified along some dimensions, e.g., regarding *color* domain, relevant dimensions are hue, chromaticity, and brightness.[5]

On the other hand, the representation of the classical information related to a given concept is demanded to classical ontological formalizations. In this setting, formal ontologies provide the characterization of concepts in terms of necessary and sufficient conditions (if these conditions exists: as mentioned, most common sense concepts cannot be characterized in these terms). Additionally, the ontological representations are used by the $\mathcal{S}2$ component (as mentioned, in our implementation it is grounded on the OpenCyc ontology).

Figure 1 shows an example of the heterogeneous representation for the concept *dog*. In this example, the exemplar and prototype-based representations make use of non classical (or typical) information and, as mentioned, are represented by using the framework of the conceptual spaces. Namely, the prototypical representation grasps information such as that dogs are usually conceptualized as domestic animals, with typically four legs, a tail *etc.*; the exemplar-based representations grasp information on individuals. For example, in Fig. 1 it is represented the individual of *Lessie*, which is a particular exemplar of *dog* with white and brown fur and with a less domestic attitude w.r.t. the prototypical dog (e.g. its typical location is lawn). Both sorts of representations activate Type 1 processes. On the other hand, the classical body of knowledge is filled with necessary and sufficient information to characterize the concept (representing, for example, the taxonomic information that a dog is a mammal and a carnivore), and activates Type 2 processes. This body of knowledge is represented with standard ontological formalisms and is grounded on OpenCyc). For the sake of readability the information in Figure 1 is visualized with a uniform format, even though the different representations are actually encoded in different formalisms.

With respect to the previous versions of the system, one of the main advances of the current version of DUAL-PECCS is the automatic acquisition of the portion of the knowledge base encoded as conceptual spaces. Most existing approaches try to induce conceptual spaces based on distributional semantics by directly accessing huge amounts of textual documents to extract the multidimensional feature vectors that de-

[4] A full account of the semantic similarity calculated in the conceptual spaces is out of the scope of this contribution; in the present setting, distances are computed in a multi-dimensional space that can be thought of as a vectorial model.

[5] It has been shown that such framework presents some advantages if introduced, in general cognitive architectures, as intermediate representational level between the symbolyc and the subsymbolic one (for more details on this issue we remind to [33].

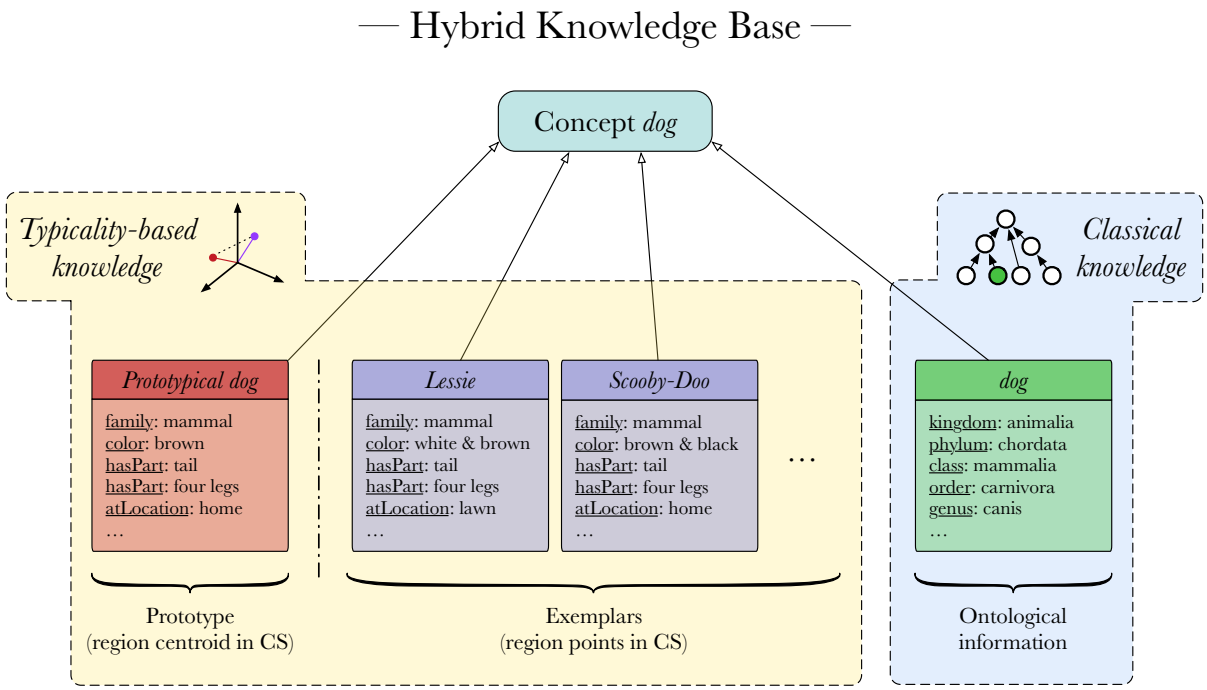

Fig. 1. Heterogeneous representation of the *dog* concept in the hybrid knowledge base.

scribe the conceptual spaces [13]. Conversely, we use a resource-driven approach that exploits: BabelNet, a multilingual encyclopedic resource built on WordNet and Wikipedia [49]; the lexical vectors in NASARI [9], a vectorial representation of BabelNet; and ConceptNet, a large scale semantic network aimed at expressing common-sense knowledge [23].[6] In order to automatically fill the conceptual spaces dimensions (which have been devised beforehand to represent mainly the animals and physical objects) with their values, we have designed a two-steps procedure, including a *semantic extraction* phase and a *semantic match* phase.

5. Building a KB for CSs Processing

In this Section we provide an overview of the methods used for the automatic population of conceptual spaces starting from linguistic resources. A complete description of this module can be found in [35]. The portion of the system performing this task is called TTCS (so named after 'Terms to Conceptual Spaces'). The TTCS system takes in input a pair $\langle t, c^t \rangle$, where t is a term and c^t is the BabelSynset ID representing a sense underlying t, [7] and produces as output a set of

[6] All mentioned resources are either built on top or directly linked to WordNet (WN), which is a lexical database for the English language [45]. Rather than organizing terms alphabetically (like ordinary dictionaries, where senses are possibly scattered) WN groups terms into synonyms sets called *synsets*, that are equipped with short definitions and usage examples. Such sets are represented as the nodes of a large semantic network, where the intervening edges represent a number of semantic relations among synset elements (such as hyponymy, hypernymy, antonymy, meronymy, holonymy).

[7] The pairs term-concept fed to the TTCS are obtained through the CLOSEST algorithm, that takes in input a term and provides in output its most common senses (BabelSynset IDs). The CLOSEST algorithm solves a well-known problem afflicting most lexical resources (such as WordNet and BabelNet): that is, too fine-grained semantic distinctions may be unnecessary and even detrimental in many tasks [55]. The CLOSEST algorithm thus produces more coarse-grained sense inventories based on a simple heuristics. It is largely acknowledged that, in human lexical and conceptual systems, the main senses of a given term gained more room than its marginal senses. This phenomenon determines the *availability* and *salience* of words and phrases [60]. These aspects are also reflected in automatically built encyclopedic resources, where the main senses are assumed to be more richly represented, in accordance with their use in spoken and written language. Building on these insights, the CLOSEST algorithm takes in input a term t and returns $\{c_1^t, c_2^t, \ldots, c_n^t\}$, the list containing the most relevant senses for t. Full details on this module can be found in [36].

ISA	PARTOF	MEMBEROF	HASA
CAPABLEOF	ATLOCATION	HASPROPERTY	ATTRIBUTE
MADEOF	SYMBOLOF	USEDFOR	INSTANCEOF

Table 1
The list of the considered ConceptNet relations.

attribute-value pairs, in the form

$$\bigcup_{d \in \mathcal{D}} \{\langle ID_d, \{v_1, \cdots, v_n\}\rangle\} \quad (1)$$

where ID_d is the identifier of the d-th quality dimension, and $\{v_1, \cdots, v_n\}$ is the set of values chosen for d. Such values will be used as fillers for Conceptual Space dimensions $d \in \mathcal{D}$. The output of the system is then a Conceptual Space representation for the input term t.

The control strategy implemented by the TTCS includes two main steps, semantic extraction and semantic matching. In the *semantic extraction* step we explore the ConceptNet associations regarding the term t and select only those involving t in the sense intended by c^t, thus producing a bag-of-concepts C semantically related to the seed term. In the *semantic matching* step, a new empty exemplar is created, corresponding to an empty vector in the CSs; we then use the bag-of-concepts extracted in the previous step to identify the values appropriate as fillers for the Conceptual Space quality dimensions.

Semantic Extraction. In the semantic extraction step, we access the ConceptNet node associated with t and scan its incoming and outgoing edges: in so doing we retrieve the related terms. The list of 12 relations that are presently considered –out of the 57 relations available in ConceptNet– is provided in Table 1. Since ConceptNet does not provide any anchoring mechanism to associate its terms to meaning identifiers (BabelSynset IDs), it is necessary to determine which edges are relevant for the concept associated to t. In other words, when we access the ConceptNet page for t, we find not only the edges regarding t with a given sense, but all the edges regarding t in any possible meaning. To select only (and possibly all) the edges that concern the sense c^t, we introduce the notion of *relevance*. To give an intuition of this process, terms found in ConceptNet are relevant (and thus retained) either if they exhibit a heavy weight in the NASARI vector corresponding to the considered concept, or if they share at least some terms with the NASARI vector. Finally, relevant terms are disambiguated and added to the bag of concepts C [35, pp. 438–41]. For example, given in input the pair $\langle bank, c^{bank}\rangle$, where $c^{bank} = \texttt{00008364n}$ is the ID corresponding to the sense 'A bank is a financial institution that creates credit by lending money to a borrower', we inspect the edges of the ConceptNet node 'bank' and thanks to the relevance notion we get rid of sentences such as 'bank ISA flight maneuver' since the term 'flight maneuver' is not present in the vector associated to concept c^{bank}; conversely, we'll accept sentences such as 'bank HASA branch'. Finally, 'branch' will be identified as a concept and its BabelSynset ID will be added to C.

Semantic Matching. The semantic matching step consists in generating a new exemplar *ex* in the CS representation, and in filling it with the information previously extracted. An exemplar is basically a list of sets of BabelSynset IDs, where each set corresponds to a quality dimension; it is named and identified in accordance with the seed term t and its meaning c^t.

The system adopts a set of quality dimensions that has been designed to meet the representational requirements posed by a set of sentences and cross-domain concepts collected in the frame of an interdisciplinary research project, aimed at investigating the neuroscientific bases of lexical processing. The selected quality dimensions aim at representing perceptually salient features (such as *size*, *shape* and *color*), with commonsense knowledge (e.g., *partOf*, *hasPart*, *function*) and taxonomic information (*class*, *family*). Some of these dimensions were borrowed from (the most frequent ones in) ConceptNet, whilst in other cases it was necessary to undertake an *ad-hoc* approach, as illustrated below. Without loss of generality, the considered set of dimensions can be extended or refined to describe some specific domain, for example by devising further dimensions to represent physical dimensions.

The process of assigning a certain value to a quality dimension is called *dimension anchoring*, and its implementation differs according to the way quality dimensions are filled: every quality dimension can be filled either based on ConceptNet or on a dictionary. In the former case (ConceptNet-driven approach) the process of extracting values to fill d leverages the set of edges; in the latter case (dictionary-driven approach) we exploit the dictionary associated with the quality dimension d. Additionally, every quality dimension can be metric or not (the whole picture is provided in Table 2). For metric quality dimensions we devised a set of translation maps (e.g., in the case of *color*, we di-

Name	BabelSynset ID	Metric	DD	CND
class	00016733n	no	-	ISA
family	00032896n	no	✓	-
shape	00021751n	no	✓	-
color	00020726n	yes	✓	-
locationEnv	00057017n	yes	✓	-
atLocation	00051760n	no	-	ATLOCATION
feeding	00029546n	yes	✓	-
hasPart	00021395n	no	-	HASA
partOf	00021394n	no	-	PARTOF
locomotion	00051798n	yes	✓	-
symbol	00075653n	no	-	SYMBOLOF
function	00036822n	no	-	USEDFOR

Table 2

List of the considered quality dimensions; the last two columns indicate respectively whether each dimension is filled in a dictionary-driven (DD) or in a ConceptNet-driven (CND) way.

rectly translate the *red* color into its L*a*b color space: $\langle 53, 80, 67 \rangle$).

In the conceptual spaces framework metrical values are fundamental to be able to compute forms of common-sense reasoning by exploiting the distances between exemplars in the resulting geometrical framework. After the new exemplar ex is filled with the values extracted through the above mentioned procedure, we translate the values of the *metric* quality dimensions by exploiting the related translation maps. Translation maps have been devised to map the extracted values onto the corresponding set of metric values in the conceptual space. For example, the *locomotion* dimension is used to account for the type of movement (1:swim, 2:dig, 3:crawl, 4:walk, 5:run, 6:roll, 7:jump, 8:fly). In the conceptual space representation the above mentioned values are translated into a numerically ordered scale such that the distance between indexes of values mirrors the semantic distance between the different types of locomotion: e.g., in this setting "dig" and "crawl" are assumed to be closer than "swim" and "fly" [7].

5.1. Building the CSs representation

In order to build an actual conceptual space, the TTCS took in input a set of 593 cross-domain pairs term-concept. To briefly account for the computational effort required in the KB building process, the TTCS handled over $2.8M$ NASARI vectors (restricting to consider the first 100 features); these were at the base of the relevance computation. In the extraction step, the TTCS accessed around $10M$ of ConceptNet assertions, linking about $3M$ nodes; specifically, $28K$ ConceptNet nodes (on average 47.6 per input term) were extracted and $2.3K$ of them were selected as relevant and finally disambiguated.

The semantic extraction step ended up with 516 success cases, where the bag-of-concepts C contained at least one extracted concept. The 76 failures were caused by the lack of the ConceptNet node for the input term, rather than by the extraction of irrelevant concepts. Further 30 input terms were dropped because their final bag-of-concepts did not contain any suitable value to fill the exemplars. This led to a total of 486 correctly extracted exemplars, and to filling overall $2,388$ dimensions (on average 4.9 per exemplar).

6. Categorization Pipeline of the System

In this Section we briefly recall, for the sake of self-containedness, how the whole categorization pipeline of DUAL-PECCS works. The overall details can be found in [39,40]. The input to the system is a simple linguistic description, like "The feline with mane", and the expected output is a given category evoked by the description (e.g., the category *lion* in this case). After an Information Extraction (IE) step[8], the input information is encoded into an internal format devised to store conceptual spaces information, which is then used as input in the categorization task. The system answers rely on the the output of $\mathcal{S}1$ and $\mathcal{S}2$, respectively. In particular: according to the linguistic stimulus being categorized DUAL-PECCS chooses, based on a similarity calculation between the stimulus and the typical representations available in $\mathcal{S}1$ knowledge base, whether to select an exemplar of a prototype (we refer to this process as $\mathcal{S}1$ categorization). By following a preference that has been experimentally observed in human cognition [44], our algorithm favors the results of the exemplars-based categorization if the knowledge-base stores any exemplars similar to the input being categorized. Once the result of $\mathcal{S}1$

[8] A shallow IE approach has been devised, where the morphological information computed from input sentences has been used to devise a simple finite-state automaton describing the input sentences' structure (more on the input descriptions in Section 8). This approach would not scale to handle more complex sentences. We defer to future work the adoption of richer language models. Despite these limitation, however, it allowed us to complete the automatization of the software pipeline going all throughout from the simple linguistic input description used for the evaluation (that will be described later) to its final conceptual categorization.

is selected (i.e., either a prototype or an exemplar is *proxyfied* in working memory), such approximate categorization result is then checked with the ontological knowledge base of $\mathcal{S}2$ (we refer to this process as $\mathcal{S}1$-$\mathcal{S}2$ categorization). This check is implemented by type 2 processes, and it is therefore based on deductive inference. If the categorization result provided by $\mathcal{S}1$ is consistent with the ontology, then the categorization succeeded and the category provided by $\mathcal{S}2$ is returned along with the top scoring class returned by $\mathcal{S}1$. Otherwise the system evaluates a fixed amount of $\mathcal{S}1$ candidates, meantime keeping track of the inconsistent elements: in case all such candidates are inconsistent w.r.t. the ontology in $\mathcal{S}2$, the output of $\mathcal{S}2$, computed independently of $\mathcal{S}1$, is provided along with the first result initially. The control strategy implements a tradeoff between ontological inference and the output of $\mathcal{S}1$, which is more informative but also formally less reliable.

6.1. Mapping Conceptual Spaces and Ontological Representation

A relevant issue we face is aligning knowledge resources based on different sorts of representational formalisms. Under an architectural perspective, $\mathcal{S}1$ and $\mathcal{S}2$ rely on knowledge bases encoded in different ways, which need to be connected and mapped onto a shared and uniform representation of meaning in order to allow the SOAR and LIDA layers to operate them (see Section 7).

The heterogeneous proxytypes approach, in particular, requires the existence of co-referring representational structures to account for conceptual knowledge. Such co-reference relies on the fact that all the different bodies of knowledge are assumed to semantically point to the main reference conceptual container. In our system, such a container has been automatically provided with a WordNet synset ID. In addition, also the pointing representations, containing the different types of conceptual knowledge, have been equipped with the same WordNet synset ID referred to their corresponding concept. The anchoring mechanism between such heterogeneous representations follows two different ways. The first one corresponds to the mapping between the concept and its related ontological component. This mapping is provided in the OpenCyc ontology, which is sometimes equipped with the information regarding the corresponding WordNet synset ID. On the other hand, the anchoring between the conceptual space representations and the corresponding general concept is obtained thanks to the connection obtained with the BabelNet linguistic resource via the TTCS subsystem. In particular, since most of the lexical items in BabelNet used by the TTCS are equipped with a WordNet synset ID, once the overall output of the TTCS is obtained, the typical Knowledge Base represented as conceptual spaces results to be automatically equipped with WordNet synsets. Once this mapping is provided, such representation is linked to the corresponding identifier of the general concept it belongs to. The Figure 2 shows a pictorial representations of the resources involved to build our co-referring representational structure.

7. Integration in SOAR and LIDA

The current system has been integrated with two additional widely known cognitive architectures[9]: SOAR [26] and LIDA [16]. SOAR is one of the most mature cognitive architectures and has been used by many researchers worldwide during the last 30-years. 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 fired in parallel to produce reasoning cycles. From a representational perspective, SOAR exploits symbolic representations of knowledge (called chunks) and uses pattern matching and spreading activation to select relevant knowledge elements. The LIDA architecture, on the other hand, is a partial implementation of the LIDA cognitive model [17] and employs a variety of modules that are designed using a variety of computational mechanisms drawn from AI. A peculiarity of this architecture resides in its detailed workflow determining the interactions between automatic (sub-conscious) processes and controlled (conscious) ones in the memory systems of its agent. Such model is explicitly grounded on the Global Workspace Theory [4,5].

The rationale underlying such integration efforts was to investigate whether the outlined approach is compatible with architectures implementing different cognitive theories of mind, so to be able to argue that

[9]The term "cognitive architecture" was introduced by Allen Newell and his colleagues in their work on unified theories of cognition [50]. One of the main reasons justifying the introduction of such systems in the AI and Computational Cognitive Science fields was the goal of reaching human level intelligence in a general setting (on the role of CAs for general intelligent systems see also [32]).

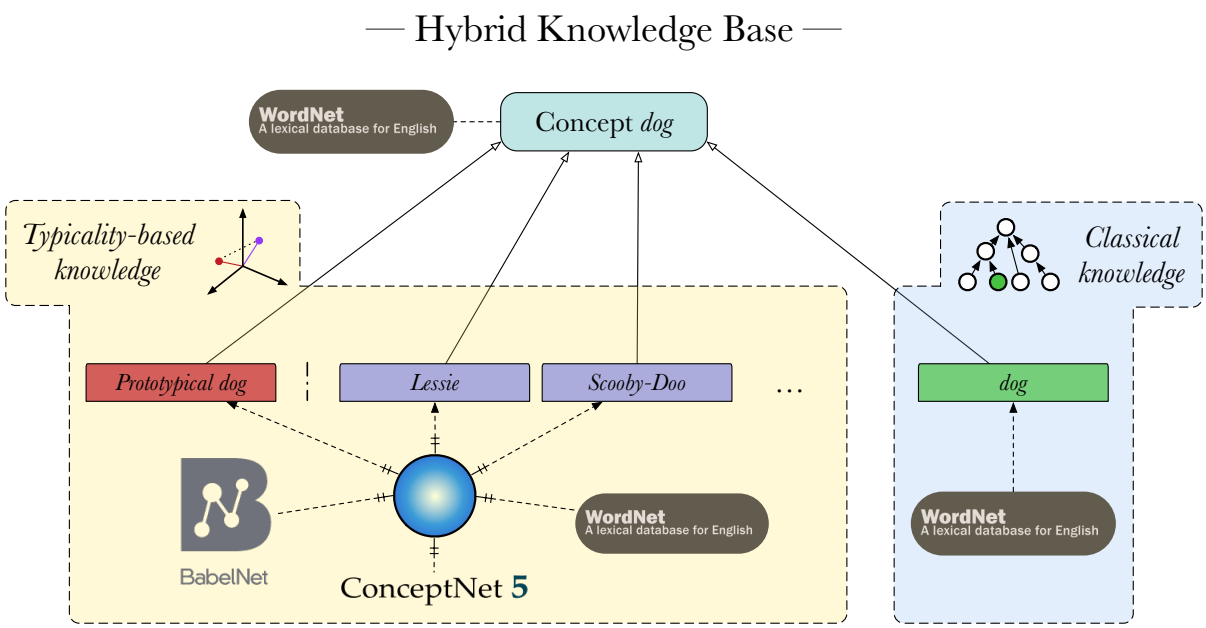

Fig. 2. Resources involved for the grounding of heterogeneous representations via WordNet.

it can be considered as a framework general enough for representation and reasoning on conceptual information.

One main difference between the two architectures is that LIDA is considered a hybrid architecture, while SOAR, on the other hand, is entirely symbolic. In particular, LIDA employs both symbolic and subsymbolic representational elements, and, more specifically, it employs the subsymbolic activation of symbolic representational chunks; SOAR adheres to the Newell and Simon's physical symbol system hypothesis [51] which states that symbolic processing is a necessary and sufficient condition for intelligent behavior [50]. Both architectures, however, are not natively dual-process based. Therefore in both cases, the dual mechanisms of reasoning needed to be explicitly designed and instantiated within an existing general framework. In addition, none of the architectures addresses the problem concerning the representation of (and the reasoning on) common-sense knowledge components, such as prototypes and exemplars (and the related reasoning strategies). In SOAR this problem arises despite the fact that the chunks can be represented as a sort of frame-like structures containing some common-sense (e.g., prototypical) information. In fact, the main problem of this architecture w.r.t. the heterogeneity assumption, relies on the fact that it does not specify how the typical conceptual components (that can eventually be represented in terms of frame-like slots) and the corresponding non monotonic-reasoning strategy, can interact with a possibly conflicting representational and reasoning procedures characterizing a different conceptualisation of the same conceptual entity [10]. In short it assumes, like most of the symbolic-oriented CAs, the availability of a monolithic conceptual structure (e.g., a frame-like prototype or a classical concept) without specifying how such information can be integrated and harmonized with other knowledge components to form the whole knowledge spectrum characterizing a given concept. In LIDA, on the other hand, despite many kinds of approximate comparisons and similarity-based reasoning (e.g., in tasks such as categorization), are, in theory, possible to execute, the peculiarity concerning prototype or exemplars based representations (along with the the design of the interaction between their different reasoning strategies) is not provided. In this sense, the integration of DUAL-PECCS with such architectures yields

[10] A classical example of the described situation is the following: let us think to the case of WHALE. A prototypical conceptualization would classify whales as FISH (since a whale share many typical traits with fishes). On the other hand, a classical conceptualization would classify whales as MAMMAL.

as profit a significant improvement in their representational and inferential power.

As for the previous integrations [39,40], for both architectures we focused on the Declarative Memory (also named Semantic Memory), and Working Memory modules, and on the corresponding retrieval mechanisms. Besides, the dual process strategies of concept categorization have been integrated into SOAR and LIDA processes and connected to the retrieval request executed in the Working Memory.

Figure 3 gives a very general overview of the rationale behind the described integration of DUAL-PECCS with different cognitive architectures. The DUAL-PECCS knowledge base in this diagram is to be intended as a sort of extension of the declarative memory modules available in the considered architectures, and it can interact with the different working memory areas to perform reasoning.

It is worth-noting that, in the literature, there are efforts proposed to extend the Declarative Memories of CAs in order to enable more complex and cognitively inspired forms of reasoning. To this class of works belongs that one proposed by [53] aiming at extending the knowledge layer of ACT-R with external ontological content related to the event modelling or that one by Salvucci [58], enriching the knowledge model of the Declarative Memory (DM) 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 [6] presenting an integration of the ACT-R Declarative and Procedural Memory with the Cyc ontology [29]. The main problematic aspect concerning the extension of the DM with such wide-coverage integrated ontological resources, however, is that the underlying formalisms of such frameworks only allow to represent conceptual information in terms of symbolic structures. As a consequence, they encounter the standard problems affecting this class of representations in dealing with the representation of common-sense knowledge components, mostly absent in such resources (for a more detailed and extended discussion we refer the interested reader to [34]). In this sense, they provide an integration only with the $\mathcal{S}2$ classical conceptual component. This aspect represents a problem since all these integrated symbolic systems do not represent at all the typical information associated to a given concept. As we will see in more detail in Section 8, this phenomenon prevents the type of common-sense conceptual retrieval based on typical traits that, on the other hand, represent one of the main contributions resulting by our integrations.

7.1. SOAR Integration

Concepts are represented in SOAR as empty chunks (that is, chunks having no associated information, except for the *WordNet synset ID* and a human readable *name*), referred to by the external bodies of knowledge (prototypes and exemplars) acting like semantic pointers. Here we have integrated the hybrid knowledge base directly into the *semantic memory* or SMEM of SOAR, that is equivalent to the Declarative Memory in LIDA and ACT-R. This sort of memory is accessed through two dedicated working memory channels, called `^command` and `^result`. In particular, `^command` is the branch of the working memory buffer where the informational query being posed to the semantic memory is provided. In our implementation this piece of information is filled with the information extracted from the input description, which is automatically converted into a SOAR chunk request. This conversion allows us to consider both the query and its result in terms of the SOAR representational language. However, instead of querying the standard Semantic Memory in SOAR, we have introduced some modification to the architecture by working on the SOAR Kernel and by creating novel RHS (Right Hand Side) functions to query the external $\mathcal{S}1$ knowledge base in order to take advantage of the similarity based reasoning possible in the conceptual spaces. The result of the first categorization process (based on common-sense information) produces in output the exemplar or prototype-based representation that is closer to the linguistic stimulus in input. Such result is stored in the `^result` channel, the branch of the SOAR working memory buffer devoted to acquiring the output from the external modules. Once the result of $\mathcal{S}1$ is "proxyfied", it is then checked, by executing a novel set of RHS functions, on a second extension of the SOAR Semantic Memory corresponding to the Cyc Ontology. The final response of the second categorization step, based on $\mathcal{S}2$ procedures, is stored in the `^result` channel to check $\mathcal{S}1$'s result consistency with the ontological model. In case an inconsistency is detected, the following best results of $\mathcal{S}1$ proxyfication is tested until a concept is returned that is compliant with the ontology in $\mathcal{S}2$.

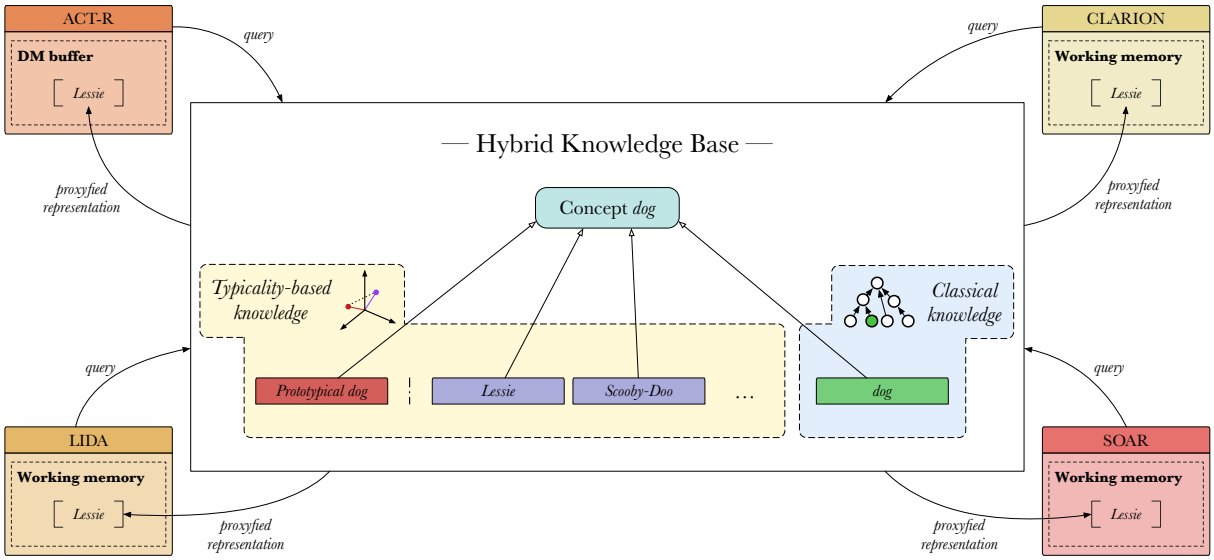

Fig. 3. General overview of the DUAL-PECCS integration within different cognitive architectures.

7.2. LIDA Integration

The integration at the representational and reasoning level in LIDA followed the same rationale as indicated for SOAR. In particular, we adapted the following modules of the architecture to integrate the structure assumed in DUAL-PECCS:

- the *Sensory Module*, receiving input from the external environment and implementing the IE step;
- the *Perceptual Associative Memory* (PAM), that receives the encoded linguistic stimulus by the Sensory Module and sends its instantiation to a working memory buffer called Workspace;
- the *Declarative Memory*, that is queried via a cue-based retrieval by the Workspace.

The dual process based categorization mechanisms have been implemented based on the following procedure: every request is encoded in the Workspace working memory as a particular type of instance (*instance chunk*). The dimensions and values of every instance chunk are filled by the PAM module with the information extracted from the linguistic description. Such process is arranged as a series of *rounds*, each producing a query, by using the cue retrieval mechanism provided by the architecture, to the implicit $\mathcal{S}1$ component and to the explicit $\mathcal{S}2$ module. As indicated in Section 6, these mechanisms require to handle different types of responses returned by the dual systems. Such responses involve different parts of the memory structure of LIDA. In particular, once the the chunk request is built, a retrieval request is executed on the $\mathcal{S}1$ knowledge base, with the aim at retrieving an exemplar or a prototype-based representation. The obtained $\mathcal{S}1$ result is then *proxyfied* and temporarily stored in a buffer of the LIDA working memory (the so called Workspace). Afterward, a second request is sent to the Declarative Memory in order to check, as previously illustrated, the results of the $\mathcal{S}1$ with the external $\mathcal{S}2$ knowledge base represented by the Cyc ontology.

An important aspect that is modelled in LIDA regards the fact that, while the consistency check of the second request to $\mathcal{S}2$ is performed, the temporary result obtained by the $\mathcal{S}1$ categorization process is broadcasted to the entire system. By using the LIDA's architectural terminology, this means that the $\mathcal{S}1$ result stored in the working memory buffer (and waiting for the $\mathcal{S}2$ response) is also sent to the to the Global Workspace module of the architecture. This additional step allows the system to make immediately available the $\mathcal{S}1$ output to the remaining modules of the architectures without waiting for the slow $\mathcal{S}2$ result. This process enables the LIDA agents to perform other tasks in real time with the available information. In case the response of $\mathcal{S}2$ results inconsistent with the previous $\mathcal{S}1$ result (and if the task for which the categorization was requested is still in focus), then the architecture will be available to broadcast to the Global Workspace a revised answer.

8. Evaluation

By following the suggestions presented in [52] we tested our integrated categorization system in a conceptual categorization task very similar to the psychological test known as "Word Reasoning". For human subjects, the Word Reasoning task consists in identifying a concept based on one to three clues. The testee might be told "You can see through it" as a first clue; "It is square and you can open it", and so on. The processing required by a Word Reasoning items goes beyond retrieval because the testee has to integrate the clues and choose among alternative hypotheses. In addition, such task can be seen as a commonsense reasoning one,known to be still on of the grand challenges of AI [46], since the answer to ths kind of queries require to resort to a "common knowledge about the world that is possessed by every schoolchild that has the methods for making obvious inferences from this knowledge" [11]. Unfortunately, as reported by [52], the standard specific questions provided for this task in the Wechsler Preschool and Primary Scale of Intelligence are proprietary. Nonetheless, the general structure of each sentence is public, so that we have re-used a dataset composed of 111 linguistic descriptions (corresponding to very simple riddles) designed by a team of linguists and neuroscientist in the frame of a research project investigating neural correlates of lexical processing. Such descriptions exhibit a structure similar to that of the Word Reasoning task: on average, no more than 3 cues are present in each riddle. The descriptions were given in input to the implemented system. An example of such descriptions is "The mice hunter with whiskers and long tail", where the expected category to be retrieved was *cat*, and in particular its representation corresponding to the "prototype of *cat*"; conversely, a description such as "The felin mice hunter without fur" was expected to lead as answer to "exemplar of *canadian-sphynx*". The *expected* categorical targets represent a gold standard, since they correspond to the results provided by 45 human subjects in a psychological experimentation. The present experimentation extends that presented by [39] in several ways: firstly, the simulation of the categorization processes is now performed on two additional, integrated, cognitive architectures; secondly, it considers an extended dataset; finally, more than half of the $\mathcal{S}1$ knowledge base has been automatically extracted, as described in Section 5.

We designed a twofold experimentation where our system results were compared with both the human responses, and with the results obtained by other systems: Wolfram Alpha, and two general-purpose search engines (Google and Bing) used in question-answering mode [24]. In particular, for the search engines we compared our system results with the first 10 answers they returned for each riddle/query.[11] We then manually evaluated the content of each page in order to assess whether the resulting document was associated to the expected category (i.e., *expected* w.r.t. the human answers). We tested the whole pipeline, where the salient information is extracted by starting from the linguistic description, the corresponding representation is retrieved, proxyfied and loaded in working memory, in both the architectures, according to the dual process approach. The information extraction of the linguistic input is not implemented in the two cognitive architectures, but it relies on the CoreNLP Stanford Parser [43], which is used to convert the textual description into a chunk request. Our evaluation records two distinct metrics:

- *Concept-categorization accuracy* (CC-ACC metric); this metric is intended to evaluate the final categorization, that is the accuracy in retrieving the expected concept (in this case, the wrong proxyfication did not count as error).
- *Proxyfication accuracy* (P-ACC metric); this metric is intended to evaluate whether given in input a description evoking a given concept, the expected proxyfied representation was retrieved. In this case the confusion between prototype and exemplar (or between exemplars) is scored as an error even if the expected category is returned.

8.1. Results and Discussion

The integrated system shows good results for the detection of the expected concepts, compared to both hu-

[11] We considered the first 10 results returned by the search engines, and not only the first one, since we are aware that these systems are not *standard* Question-Answering Systems. However, the underlying purpose of such systems can be plausibly ascribed to the area of Question-Answering, that is, the search engines try to answer to queries, that express a question/information need and looks for an answer. This fact is corroborated by evidence in literature, showing that over more than 80% of Web queries are informational in nature [24], and by recent works in QA and in Semantic Textual Similarity [1,48]. In particular, in these settings, a relevant trend is based on adopting IR techniques that —rather than focussing on the generation of direct answers— are aimed, as the search engines, at finding text excerpts that contain the answer within large collections of documents [47].

Table 3
The accuracy results (Table 3-a) and the analysis of the proxyfication errors (Table 3-b).

a. Accuracy rates obtained for the conceptual categorization accuracy (CC-ACC) and proxyfication accuracy (P-ACC) metrics.

test	CC-ACC	P-ACC
Integrated System/Humans	76.6% (85/111)	67.1% (57/85)
Google/Humans	66.7% (74/111)	71.6% (53/74)
Bing/Humans	60.4% (67/111)	77.7% (52/67)
W. Alpha/Humans	2.07% (3/111)	100% (3/3)

b. Analysis of the errors in the proxyfication (P-ACC metrics).

test	Proxyfication error		
	Ex-Proto	Proto-Ex	Ex-Ex
System	28.2% (24/85)	0.0% (0/85)	4.7% (4/85)
Google	6.8% (5/74)	18.9% (14/74)	2.7% (2/74)
Bing	7.4% (5/67)	1.4% (1/67)	13.4% (9/67)
W.Alpha	0.0% (0/3)	0.0% (0/3)	0.0% (0/3)

man and artificial systems responses. These figures are reported in Table 3-a. Table 3-b reports the detailed errors committed in the proxyfication phase. Provided that *proxyfication* errors occur only when the concept has been correctly categorized, three kinds of errors were recorded: an exemplar returned in place of an expected prototype (column *Ex-Proto*); a prototype returned in place of an expected exemplar (column *Proto-Ex*), or a wrong exemplar retrieved in place of a correct one (e.g., an individual of *polar_bear* in place of an individual of *asian_black_bear*, column *Ex-Ex*).

Compared to human response the vast majority of errors of our system are due to the confusion between exemplars and prototypes. In particular, in the 28.2% of the considered stimuli an exemplar-proxyfied representation has been returned by the system in spite of the expected prototype. This unexpected error is due to the heuristics —proper to human cognition [44] and implemented by the categorization algorithm — that favors the results of the exemplars-based categorization w.r.t. their prototypical counterpart. While this heuristics is helpful in most cases, the analysis of such proxyfication errors points out that the interaction of such common-sense reasoning mechanisms deserves additional clarification in the theoretical and experimental psychological literature. Such heuristics, in fact, resulted counterintuitive and not efficacious for the correct categorization of general descriptions.

The other systems suffer from different problems: in particular, Google committed the majority (18.9%) of errors in favoring prototypes over expected exemplars, whilst the predominant error in Bing was due to the retrieval of a wrong exemplar (*Ex-Ex* error). The results of Wolfram Alpha, finally, were not surprising since this system is able to answer classical scientific-oriented queries, while it is not yet equipped to deal with common-sense knowledge and reasoning. This aspect, in fact, still represents one of the more challenging aspects in the current knowledge-based AI systems [30] affecting, as reported in [12], also systems such as IBM Watson.[12]

Additional empirical results supporting the feasibility of the overall dual-process based framework are provided in [34]. Such results, that we mention here for the sake of completeness, show that the $\mathcal{S}1$-$\mathcal{S}2$ component adopted in DUAL-PECCS presents by far better results with respect to those obtained by considering only an $\mathcal{S}2$ "ontological" knowledge component in isolation (this test has been executed by using, as $\mathcal{S}2$, different large ontological knowledge bases such as Cyc [29] and DBpedia [3]). This finding is consistent with the hypothesis that the typicality based knowledge component, in our environment associated to the $\mathcal{S}1$ module, is crucial for common-sense question answering.

9. Conclusions

This paper has proposed, as a unifying representational and reasoning framework for artificial agents, a cognitively-inspired conceptual architecture implemented in a system which has now been integrated into SOAR and LIDA. The proposed framework has shown a good deal of compatibility with such general

[12] A comparison with the Watson system for common-sense queries represents a mid-term evaluation goal. We recently (in Spring/Summer 2017) inspected the possibility of querying the Watson system by using its Open Services `https://www.ibm.com/watson/products-services/`. Unfortunately, the available services do not provide the possibility of querying the Watson Knowledge Base. Currently, in fact, every user can create his/her own KB collecting a set of documents and then answer queries on the internal representation extracted from that documents. Despite this is an interesting feature, in our opinion a comparison with an *ad-hoc* built Knowledge Base would be unfair and rather difficult to actually implement.

cognitive architectures making different assumptions about the structures and processes of human cognition (and of human-like cognition in AI systems). Although there is room for both refining the framework and tuning the implemented system, the results obtained in a task of common-sense conceptual categorization are encouraging when compared to both human and artificial responses. As above mentioned, we stress that, with respect to our previous work, the extensions described in this paper enable us to fully collocate the employed representational framework adopted in DUAL-PECCS within the recent debate concerning the identification of a Standard Model of Mind [27] and, in particular, in the debate concerning the 'Memory and Content' issues of the Declarative Memories of current Cognitive Architectures. At the current stage of development, in fact, the employed representational framework has been proved to be versatile enough to be integrated with both fully symbolic and hybrid architectures endowed with a strong sub-symbolic component (e.g., CLARION and LIDA). In our opinion this is due to the adoption of the Conceptual Spaces framework as $\mathcal{S}1$ component.[13]

As a mid term goal, we plan to integrate the proposed representational and reasoning framework into further general cognitive architectures (e.g., in LEABRA [54] or SPAUN [8,14]), based on still different representational and reasoning assumptions w.r.t. SOAR and LIDA. Should also these integrations be feasible, we would be allowed to reinforce our claim that the proposed representational and reasoning framework can be used as a reference for the knowledge level of different cognitive agents, thus providing a sort of *interlingua* for heterogeneous architectures. In particular, in a multi-agent setting, the provided framework can be seen as a communication layer. Such a layer is suited *i*) to extend the knowledge stored in the Long Term Memory of the individual agents; and *ii*) to provide a more advanced —shared across the architectures— set of reasoning procedures to query, retrieve and reason on conceptual knowledge coupling standard and common-sense reasoning procedures. Such procedures contribute to fill the gap between the existing cognitive architectures and the categorization heuristics used by human cognition and not previously available (or only partially available) in those systems [38].

Another strength of the proposed approach regards the possibility, for diverse cognitive agents, to interpret and process the shared knowledge level, meantime maintaining the specificities and the constraints proper to each architecture.

[13]For the reasons explaining why the framework of the Conceptual Spaces can be considered in Cognitive Architectures an efficacious, intermediate, representational level connecting symbolic and sub-symbolic representations and what are the advantages offered by such framework with respect to some of the main problems affecting the representational level of CAs we refer to [10,33].

Acknowledgements

We thank Francesco Ghigo and Francesco Caronte for their help with the implementation and the anonymous reviewers for their insightful comments.

References

[1] E. Agirre, C. Banea, D. M. Cer, M. T. Diab, A. Gonzalez-Agirre, R. Mihalcea, G. Rigau, and J. Wiebe. Semeval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation. In *SemEval@ NAACL-HLT*, pages 497–511, 2016.

[2] J. R. Anderson, D. Bothell, M. D. Byrne, S. Douglass, C. Lebiere, and Y. Qin. An integrated theory of the mind. *Psychological review*, 111(4):1036, 2004.

[3] S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak, and Z. Ives. Dbpedia: A nucleus for a web of open data. *The semantic web*, pages 722–735, 2007.

[4] B. J. Baars. In the theatre of consciousness. global workspace theory, a rigorous scientific theory of consciousness. *Journal of Consciousness Studies*, 4(4):292–309, 1997.

[5] B. J. Baars and S. Franklin. An architectural model of conscious and unconscious brain functions: Global workspace theory and ida. *Neural Networks*, 20(9):955–961, 2007.

[6] J. Ball, S. Rodgers, and K. Gluck. Integrating act-r and cyc in a large-scale model of language comprehension for use in intelligent agents. In *AAAI workshop*, pages 19–25, 2004.

[7] A. Bejan and J. H. Marden. Constructing animal locomotion from new thermadynamics theory: Although running, flying and swimming appear to be distinctly different types of movement, they may have underlying physics in common. *American Scientist*, 94(4):342–349, 2006.

[8] P. Blouw, E. Solodkin, P. Thagard, and C. Eliasmith. Concepts as semantic pointers: a framework and computational model. *Cognitive science*, 2015.

[9] J. Camacho-Collados, M. T. Pilehvar, and R. Navigli. Nasari: a novel approach to a semantically-aware representation of items. In *Proceedings of NAACL*, pages 567–577, 2015.

[10] A. Chella, M. Frixione, and A. Lieto. Representational issues in the debate on the standard model of the mind. In *2017 AAAI Fall Symposium Series "A Standard Model of the Mind"*, 2017.

[11] E. Davis. *Representations of commonsense knowledge*. Morgan Kaufmann, 2014.

[12] E. Davis and G. Marcus. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Communications of the ACM*, 58(9):92–103, 2015.

[13] J. Derrac and S. Schockaert. Inducing semantic relations from conceptual spaces: a data-driven approach to plausible reasoning. *Artificial Intelligence*, 228:66–94, 2015.

[14] C. Eliasmith, T. C. Stewart, X. Choo, T. Bekolay, T. DeWolf, Y. Tang, and D. Rasmussen. A large-scale model of the functioning brain. *Science*, 338(6111):1202–1205, 2012.

[15] J. S. B. Evans and K. E. Frankish. *In two minds: Dual processes and beyond.* Oxford University Press, 2009.

[16] S. Franklin and F. Patterson Jr. The lida architecture: Adding new modes of learning to an intelligent, autonomous, software agent. *pat*, 703:764–1004, 2006.

[17] S. Franklin, S. Strain, J. Snaider, R. McCall, and U. Faghihi. Global workspace theory, its lida model and the underlying neuroscience. *Biologically Inspired Cognitive Architectures*, 1:32–43, 2012.

[18] M. Frixione and A. Lieto. Representing concepts in formal ontologies: Compositionality vs. typicality effects. *Logic and Logical Philosophy*, 21(4):391–414, 2012.

[19] M. Frixione and A. Lieto. Dealing with Concepts: from Cognitive Psychology to Knowledge Representation. *Frontiers in Psychological and Behavioural Science*, 2(3):96–106, July 2013.

[20] P. Gärdenfors. *Conceptual spaces: The geometry of thought*. MIT press, 2000.

[21] P. Gärdenfors. *The Geometry of Meaning: Semantics Based on Conceptual Spaces*. MIT Press, 2014.

[22] P. Gärdenfors and M.-A. Williams. Reasoning about categories in conceptual spaces. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence - Volume 1*, IJCAI'01, pages 385–392, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.

[23] C. Havasi, R. Speer, and J. Alonso. Conceptnet: A lexical resource for common sense knowledge. *Recent advances in natural language processing V: selected papers from RANLP*, 309:269, 2007.

[24] B. J. Jansen, D. L. Booth, and A. Spink. Determining the informational, navigational, and transactional intent of web queries. *Information Processing & Management*, 44(3):1251–1266, 2008.

[25] D. Kahneman. *Thinking, fast and slow*. Macmillan, 2011.

[26] J. Laird. *The Soar cognitive architecture*. MIT Press, 2012.

[27] J. E. Laird, C. Lebiere, and P. S. Rosenbloom. A standard model of the mind: Toward a common computational framework across artificial intelligence, cognitive science, neuroscience, and robotics. *AI Magazine*, pages 1–19, in press.

[28] P. Langley, J. E. Laird, and S. Rogers. Cognitive architectures: Research issues and challenges. *Cognitive Systems Research*, 10(2):141–160, 2009.

[29] D. B. Lenat. Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11):33–38, 1995.

[30] H. J. Levesque. *Common Sense, the Turing Test, and the Quest for Real AI: Reflections on Natural and Artificial Intelligence*. MIT Press, 2017.

[31] A. Lieto. A computational framework for concept representation in cognitive systems and architectures: Concepts as heterogeneous proxytypes. *Procedia Computer Science*, 41:6–14, 2014.

[32] A. Lieto, M. Bhatt, A. Oltramari, and D. Vernon. The role of cognitive architectures in general artificial intelligence. *Cognitive Systems Research*, 2017.

[33] A. Lieto, A. Chella, and M. Frixione. Conceptual spaces for cognitive architectures: A lingua franca for different levels of representation. *Biologically Inspired Cognitive Architectures (BICA)*, 19:1–9, 2017.

[34] A. Lieto, C. Lebiere, and A. Oltramari. The knowledge level in cognitive architectures: Current limitations and possible developments. *Cognitive Systems Research*, 2017.

[35] A. Lieto, E. Mensa, and D. P. Radicioni. A Resource-Driven Approach for Anchoring Linguistic Resources to Conceptual Spaces. In *XVth International Conference of the Italian Association for Artificial Intelligence, Genova, Italy, November 29-December 1, 2016, Proceedings*, volume 10037 of *Lecture Notes in Artificial Intelligence*, pages 435–449. Springer, 2016.

[36] A. Lieto, E. Mensa, and D. P. Radicioni. Taming sense sparsity: a common-sense approach. In *Proceedings of Third Italian Conference on Computational Linguistics (CLiC-it 2016) & Fifth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian. Final Workshop (EVALITA 2016), Napoli, Italy, December 5-7, 2016.*, 2016.

[37] A. Lieto, A. Minieri, A. Piana, and D. P. Radicioni. A knowledge-based system for prototypical reasoning. *Connection Science*, 27(2):137–152, 2015.

[38] A. Lieto and D. P. Radicioni. From human to artificial cognition and back: New perspectives on cognitively inspired ai systems. *Cognitive Systems Research*, 39:1–3, 2016.

[39] A. Lieto, D. P. Radicioni, and V. Rho. A Common-Sense Conceptual Categorization System Integrating Heterogeneous Proxytypes and the Dual Process of Reasoning. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 875–881, Buenos Aires, July 2015. AAAI Press.

[40] A. Lieto, D. P. Radicioni, and V. Rho. Dual PECCS: A Cognitive System for Conceptual Representation and Categorization. *Journal of Experimental & Theoretical Artificial Intelligence*, 29(2):433–452, 2017.

[41] E. Machery. *Doing without concepts*. OUP, 2009.

[42] B. C. Malt. An on-line investigation of prototype and exemplar strategies in classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15(4):539, 1989.

[43] C. D. Manning, M. Surdeanu, J. Bauer, J. Finkel, S. J. Bethard, and D. McClosky. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, 2014.

[44] D. L. Medin and M. M. Schaffer. Context theory of classification learning. *Psychological review*, 85(3):207, 1978.

[45] G. A. Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41, 1995.

[46] M. Minsky. *The emotion machine: Commonsense thinking, artificial intelligence, and the future of the human mind*. Simon and Schuster, 2007.

[47] D. Mollá and J. L. Vicedo. Question answering in restricted domains: An overview. *Computational Linguistics*, 33(1):41–61, 2007.

[48] P. Nakov, L. MÃărquez, W. Magdy, A. Moschitti, J. Glass, and B. Randeree. Semeval-2015 task 3: Answer selection in community question answering. In D. M. Cer, D. Jurgens, P. Nakov, and T. Zesch, editors, *SemEval@NAACL-HLT*, pages 269–281. The Association for Computer Linguistics, 2015.

[49] R. Navigli and S. P. Ponzetto. Babelnet: The automatic construction, evaluation and application of a wide-coverage multi-

lingual semantic network. *Artificial Intelligence*, 193:217–250, 2012.

[50] A. Newell. *Unified theories of cognition*. Harvard University Press, 1994.

[51] A. Newell and H. A. Simon. Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19(3):113–126, 1976.

[52] S. Ohlsson, R. H. Sloan, G. Turán, D. Uber, and A. Urasky. An approach to evaluate ai commonsense reasoning systems. In *FLAIRS Conference*, 2012.

[53] A. Oltramari and C. Lebiere. Pursuing artificial general intelligence by leveraging the knowledge capabilities of act-r. In *Artificial General Intelligence*, pages 199–208. Springer, 2012.

[54] R. C. O'Reilly and Y. Munakata. *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. MIT press, 2000.

[55] M. Palmer, O. Babko-Malaya, and H. T. Dang. Different Sense Granularities for Different Applications. In *Proceedings of Workshop on Scalable Natural Language Understanding*, 2004.

[56] J. J. Prinz. *Furnishing the mind: Concepts and their perceptual basis*. MIT press, 2002.

[57] E. Rosch. Cognitive representations of semantic categories. *J. Exp. Psychol. Gen.*, 104(3):192–233, 1975.

[58] D. D. Salvucci. Endowing a cognitive architecture with world knowledge. In *Procs. of the 36th Annual Meeting of the Cognitive Science Society*, 2014.

[59] R. Sun. The CLARION cognitive architecture: Extending cognitive modeling to social simulation. *Cognition and multi-agent interaction*, pages 79–99, 2006.

[60] P. Vossen and C. Fellbaum. *Multilingual FrameNets in Computational Lexicography: Methods and Applications*, chapter Universals and idiosyncrasies in multilingual WordNets. Trends in linguistics / Studies and monographs: Studies and monographs. Mouton de Gruyter, 2009.