

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

## Extending Ontological Categorization Through a Dual Process Conceptual Architecture

### **This is the author's manuscript**

*Original Citation:*

*Availability:*

This version is available <http://hdl.handle.net/2318/1562401> since 2016-05-24T17:36:41Z

*Publisher:*

Springer International Publishing AG

*Published version:*

DOI:10.1007/978-3-319-25840-9

*Terms of use:*

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

# Extending Ontological Categorization through a Dual Process Conceptual Architecture

Antonio Lieto, Daniele P. Radicioni,  
Marcello Frixione and Valentina Rho  
{lieto,radicion}@di.unito.it,  
frix@dist.unige.it, rho.valentina@gmail.com

## Abstract

In this work we present a hybrid knowledge representation system aiming at extending the representational and reasoning capabilities of classical ontologies by taking into account the theories of typicality in conceptual processing. The system adopts a categorization process inspired to the *dual process theories* and, from a representational perspective, is equipped with a heterogeneous knowledge base that couples conceptual spaces and ontological formalisms. The system has been experimentally assessed in a conceptual categorization task where common sense linguistic descriptions were given in input, and the corresponding target concepts had to be identified. The results show that the proposed solution substantially improves the representational and reasoning “conceptual” capabilities of standard ontology-based systems.

## 1 Introduction

One of the main open problems in the field of ontology engineering is that formal ontologies do not allow –for technical convenience– neither the representation of concepts in prototypical terms nor forms of approximate, non monotonic, conceptual reasoning. Conversely, in Cognitive Science evidences exist in favor of prototypical representation of concepts, and typicality-based conceptual reasoning has been widely investigated in the field of human cognition. The early work of Rosch [29] showed that ordinary concepts do not obey the classical theory (stating that concepts can be defined in terms of sets of necessary and sufficient conditions). Rather, they exhibit *prototypical* traits: e.g., some members of a category are considered *better instances* than other ones; more *central* instances share certain typical features –such as the ability of flying for birds– that, in general, cannot be thought of as necessary nor sufficient conditions. These results influenced pioneering KR research, where some efforts were invested in trying to take into account the suggestions coming from Cognitive Psychology: artificial systems were designed –e.g., frames [26] and early semantic networks– to represent concepts in “non classical”, prototypical terms.

However, these systems were later sacrificed in favor of a class of formalisms stemmed from structured inheritance semantic networks and based in a more rigorous semantics: the first system in this line of research was the KL-ONE system [3]. These formalisms are known today as description logics (DLs) [27].

In this setting, the representation of prototypical information (and therefore the possibility of performing non monotonic reasoning) is not allowed,<sup>1</sup> since the formalisms in this class are primarily intended for deductive, logical inference.

Under a historical perspective, the choice of preferring classical systems, which are based on a well defined –Tarskian-like– semantics left unsolved the problem of representing concepts in prototypical terms. Although in the field of logic oriented KR various fuzzy and non-monotonic extensions of DL formalisms have been designed to deal with some aspects of “non-classical” concepts [30, 15], nonetheless various theoretical and practical problems remain unsolved [6].

In this paper a conceptual architecture is presented that, embedded in a larger knowledge-based system, aims at extending the representational and reasoning capabilities available to traditional ontology-based frameworks. The whole system implementing the proposed conceptual architecture is part of a larger software pipeline; it includes the extraction of salient information from the input *stimulus*, the access to the hybrid knowledge base, and the retrieval of the corresponding concept (Figure 1). The paper is structured as follows: in Section 2 we illustrate the theoretical motivations inspiring the proposed system, its general architecture and the main features of the knowledge-base. In Section 3 we provide the results of a twofold experimentation to assess the accuracy of the system in a categorization task. Finally, we conclude by presenting some related works (Section 4) and outlining future developments (Section 5).

## 2 The System

Two cornerstones inspiring the current proposal are the *dual process theories* of mind and the heterogeneous approach to concepts in Cognitive Science. The theoretical framework known as *dual process theory* postulates the co-existence of two different types of cognitive systems [5, 18]. The systems of the first type (*type 1*) are phylogenetically older, unconscious, 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 their turn by many subsystems and processes. According to the hypotheses in [6, 10], the conceptual representation of our system includes two main sorts of components, based on two sorts of processes. *Type 1* processes are used to perform fast and approximate categorization, and benefit from prototypical information associated to concepts. *Type 2* processes, used in classical inference tasks, do not take advantage from prototypical knowledge. The two sorts of system processes are assumed to interact, since type 1 processes are executed first and their results are then refined by type 2 processes. Another source of inspiration for our work has been the heterogeneous approach to the concepts in Cognitive Science and Philosophy [23]. According to this perspective, concepts do not constitute a unitary phenomenon; rather, concepts can consist of several bodies of knowledge, each one conveying a specific type of information.

Before starting with the description of the system, we introduce the theoretical framework inspiring our present work.

---

<sup>1</sup>This is the case, for example, of *exceptions* to the inheritance mechanism.

## 2.1 Theoretical Framework

According to the aforementioned hypotheses (*dual process* theory and *heterogeneous approach* to the concepts), we designed a hybrid conceptual architecture that builds on a classical ontological component, and on a typicality based one. Both components represent a specific conceptual body of knowledge, together with the related reasoning procedures, in a dual process perspective. A DL formalism is the base of the ontological component, which permits to express necessary and/or sufficient conditions to define concepts. For example, if we consider the concept *water*, the classical representation should contain the information that *water* is a natural substance, whose chemical formula is  $H_2O$ . On the other hand, the prototypical traits include information about the fact that water usually occurs in liquid state, and it is mostly a tasteless, odorless and colorless fluid.

According to our dual process view, in the implemented system the typical representational and reasoning functions are assigned to the system 1 (hereafter  $\mathcal{S}1$ ), that executes processes of type 1, and are associated to the Conceptual Spaces framework [11]. 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 DL-based ontological representation. Different from what proposed in [7], the access to the information stored and processed in both components is assumed to proceed from the  $\mathcal{S}1$  to the  $\mathcal{S}2$ . In the following we introduce the two representational and reasoning frameworks used in our system.

## 2.2 Conceptual Spaces

Conceptual spaces (CS) are a geometrical representational framework where knowledge is represented as a set of *quality dimensions* [11]. A geometrical structure is associated to each quality dimension. In this framework instances may be represented as points in this multidimensional space, and their similarity can be computed as the intervening distance, based on some suitable metrics. In this setting, concepts correspond to regions, and regions with different geometrical properties correspond to different sorts of concepts.

Conceptual spaces are suitable to represent concepts in “typical” terms, since the regions representing concepts can have soft boundaries. Prototypes have a natural geometrical interpretation, in that they correspond to the geometrical centre of a convex region; conversely, given a convex region we can associate to each point a certain centrality degree, that can be interpreted as a measure of its typicality. Although other forms of typicality-based representation (e.g. the exemplars) are not presently accounted for by our system, this framework can be extended to consider both the exemplar and the prototypical accounts of typicality [9].

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*; a domain is “a set of integral dimensions that are separable from all other dimensions” [12]. Typical domain examples are color, size, shape, texture. In turn, domain information can be specified along some dimensions: e.g., in the case of the *color* domain, relevant dimensions are hue, chromaticity, and brightness. In order to compute the distance between two

points  $p_1, p_2$  we use Euclidean metrics to calculate within-domain distance, while for dimensions from different domains we use the Manhattan distance metrics, as suggested in [11, 1]. The weighted Euclidean distance  $dist_E$  is computed as follows

$$dist_E(p_1, p_2) = \sqrt{\sum_{i=1}^n w_i (p_{1,i} - p_{2,i})^2},$$

where  $i$  varies over the  $n$  domain dimensions and  $w_i$  are dimension weights.

In our implementation of Conceptual Spaces, we represent points as vectors (with as many dimensions as required by the considered domain), whose components correspond to the point coordinates, so that a natural metrics to compute the similarity between them is *cosine similarity*. In this perspective two vectors with same orientation have a cosine similarity 1, while two orthogonal vectors have cosine similarity 0. The normalized version of cosine similarity ( $\hat{cs}$ ), also accounting for the above weights  $w_i$  is computed as

$$\hat{cs}(p_1, p_2) = \frac{\sum_{i=1}^n w_i (p_{1,i} \times p_{2,i})}{\sqrt{\sum_{i=1}^n w_i (p_{1,i})^2} \times \sqrt{\sum_{i=1}^n w_i (p_{2,i})^2}}.$$

In the metric space we have defined, the distance between an individual and prototypes is computed with the Manhattan distance. The distance between two concepts can be computed as the distance between two regions: namely, we can compute the distance between their prototypes, or the minimal distance between their *individuals*<sup>2</sup>, or we can apply more sophisticated algorithms. Further details about technical issues can be found in [14].

Optionally, a context  $k$  can be defined as a set of weights, to grade the relative relevance of the considered dimensions –thus resulting in the following formulas:  $dist_E(p_1, p_2, k)$  and  $\hat{cs}(p_1, p_2, k)$ –, and to adapt the computation to a variety of settings, such as, e.g., default values *vs.* known values, explicitly asserted values *vs.* computed values and/or inherited, etc..

We stress that inference in conceptual spaces can be performed on incomplete and/or noisy information: that is, it is frequent the case that only partial information is available to categorize a given input individual, and for some individuals the values of one or more dimensions can be undefined. Conceptual spaces are robust to this sort of lack of information, which is conversely problematic in the context of formal ontologies. In such cases we restrict to considering domains that contain points in the input: if the description for a given individual does not contain values for some domains, the distance for those domains is set to a default value.

### 2.3 S1: Modeling Domains and Dimensions in Conceptual Spaces

A processing format for the modelling of Conceptual Spaces has been provided and proposed in [25]. Although the format has been developed by attempting to keep it as general as possible (to extend its usage to further domains), the present implementation has been devised based on specific representational

<sup>2</sup>Individuals can be thought of as *exemplars*, that is elements that belong to the given concept by sharing properties and their related sets of values.

needs described in more detail in Section 3. The basic format structure processed by the system is named `genericDescription`; it encodes the salient aspects of the entities being considered. A `genericDescription` is a super-domain that hosts information about physical and non physical features arranged into nine domains: `size`, `shape`, `color`, `location`, `feeding`, `locomotion`, `hasPart`, `partOf`, `manRelationship`.

The size of entities is expressed through the three Euclidean dimensions; the `shape` allows expressing that an object has circular, square, spherical, cubic, etc., shape. The `color` space maps object's features onto the  $L^*a^*b^*$  color space.  $L^*$  ( $0 \leq L \leq 100$ ) is the correlate of lightness,  $a^*$  ( $-128 \leq a \leq 127$ ) is the chromaticity axis ranging from green to red, and  $b^*$  ( $-128 \leq b \leq 127$ ) is the chromaticity axis ranging from blue to yellow. The `location` space indicates the place where the object being modeled can be typically found. It actually results from the combination of five dimensions, and namely: `humidity`, indicated as a percentage; `temperature`, ranging in  $[-40^\circ, 50^\circ]$ ; `altitude`, ranging in  $[-11000, 8848]$ ; `vegetation`, ranging in  $[0, 100]$ ; `time`. In turn, `time` contains a partitioning of the hours of the day into `sunrise` (4–6 AM), `morning` (6–12 AM), `afternoon` (12–5 PM), `evening` (5–10 PM) and `night` (10 PM–4 AM). The domain `feeding` is currently specific to animals, and it allows mapping an element over two dimensions, `typeOfFood` and `amountOfFood`. The `typeOfFood` is associated to an integer indicating 1: herbivore, 2: lectivore, 3: detritivore, 4: necrophage, 5: carnivore. The underlying rationale is that close elements (e.g., necrophage and carnivore, that are one step apart in the proposed scale) are represented as close in this space due to their proximity under an ethological viewpoint, whilst different categories (e.g., herbivore and carnivore) are featured by larger distances in the considered scale [13]. Similar to the previous one, also the `locomotion` domain combines two dimensions: the former 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), and the latter one is used to account for the speed, expressed in km/h [2]. The `hasPart` and `partOf` domains are used to complement the analogous ontological properties: in particular, we collected information about the following dimensions: `name`, `number`, and `partSize`, `partColor` that are intended to specialize the above illustrated spaces. Finally, the `manRelationship` space is used to grasp entities as related to man by function (both a train and a horse can be used as 'transport'), product (chicken produce 'eggs', and 'chicken' themselves are a food product), symbol ('lion' can be used as a symbol for 'strength' and 'royalty').

A simplified example of the *lion* prototype information is reported below.

```
<object name="lion">
  <genericPhysicalDescription>
    <size name="lion_size">
      <x>70</x>
      <y>120</y>
      <z>200</z>
    </size>
    <color name="beige">
      <l>63</l>
      <a>13</a>
      <b>32</b>
    </color>
    <location name="savanna">
      <humidity>50</humidity>
```

```

    <temperature>40</temperature>
    <altitude>100</altitude>
    <vegetation>50</vegetation>
  </location>
  <locomotion name="walk">
    <movement>4</movement>
    <speed>10</speed>
  </locomotion>
  <locomotion name="run">
    <movement>5</movement>
    <speed>40</speed>
  </locomotion>
</genericPhysicalDescription>
</object>

```

## 2.4 $\mathcal{S}2$ : Ontological Component of the Knowledge Base

The representation of the classical component  $\mathcal{S}2$  is implemented through a formal ontology. As already pointed out, the standard ontological formalisms leave unsolved the problem of representing prototypical information. Furthermore, it is not possible to execute non monotonic inference, since classical ontology-based reasoning mechanisms contemplate exclusively deductive processes.

In this setting we cannot represent even simple prototypical information, such as ‘A typical rose is red’. This is due to the fact that being red is neither a necessary nor a sufficient condition for being a rose, and therefore it is not possible neither to represent and to automatically identify a prototypical rose (let us call it  $\#roseP$ ) nor to describe (and to learn from new cases) the typical features of the prototypical roses. Such aspect have, on the other hand, a natural interpretation in terms of the Conceptual Spaces framework. The ontological component of a given concept, therefore, mainly represents taxonomical necessary information and, as will be described below, is used as control module w.r.t. the output provided by the inferences performed within the conceptual spaces.

## 2.5 Inference in the hybrid $\mathcal{S}1$ - $\mathcal{S}2$ system

Categorization (i.e., the process of assigning a given instance to a certain category) is one of the classical inferences automatically performed both by symbolic and sub-symbolic artificial systems. In our system categorization is based on a two-step process involving both the typical and the classical component of the conceptual representation. These components account for different types of categorization: approximate or non monotonic (performed on the conceptual spaces), and classical or monotonic (performed on the ontology). Different from classical ontological inference, in fact, categorization in conceptual spaces proceeds from *prototypical* values. Prototypical values need not be specified for all class members, that vice versa can overwrite them: one typical example is the case of birds that (by default) fly, except for atypical birds, like penguins, that do not fly.

The whole categorization process can be summarized as follows. The system takes in input a textual description  $d$  and produces in output a pair of categories  $\langle c_0, cc \rangle$ , the output of  $\mathcal{S}1$  and  $\mathcal{S}2$ , respectively (see Algorithm 1). If the  $\mathcal{S}2$  system classifies it as consistent with the ontology, then the classification succeeded and the category provided by  $\mathcal{S}2$  ( $cc$ ) is returned along with  $c_0$ , the

---

**Algorithm 1** Inference in the hybrid system.

---

**input** : textual description  $d$   
**output** : a class assignment, as computed by  $\mathcal{S}1$  and  $\mathcal{S}2$

```
1:  $C \leftarrow \mathcal{S}1(d)$  /* conceptual spaces output */
2: for each  $c_i \in C$  do
3:    $cc \leftarrow \mathcal{S}2(\langle d, c_i \rangle)$  /* ontology based output */
4:   if  $cc == \text{NULL}$  then
5:     continue /* inconsistency detected */
6:   end if
7:   if  $cc$  equals  $c_i$  then
8:     return  $\langle c_0, cc \rangle$ 
9:   else
10:    if  $cc$  is subclass of  $c_i$  then
11:      return  $\langle c_0, cc \rangle$ 
12:    end if
13:  end if
14: end for
15:  $cc \leftarrow \mathcal{S}2(\langle d, \text{Thing} \rangle)$ 
16: return  $\langle c_0, cc \rangle$ 
```

---

top scoring class returned by  $\mathcal{S}1$  (Algorithm 1: line 8). If  $cc$  –the class computed by  $\mathcal{S}2$ – is a subclass of one of those identified by  $\mathcal{S}1$  ( $c_i$ ), both  $cc$  and  $c_0$  are returned (Algorithm 1: line 11). Thus, if  $\mathcal{S}2$  provides more specific output, we follow a *specificity* heuristics; otherwise, the output of  $\mathcal{S}2$  is returned, following the rationale that it is *safer*.<sup>3</sup> A pair of results is always returned, including both the output of  $\mathcal{S}1$  and the output of  $\mathcal{S}2$ , thereby providing typically valid answers (through  $\mathcal{S}1$ ) that are checked against a logically valid reasoning conducted on the ontological knowledge base (through  $\mathcal{S}2$ ). In so doing, we follow the rationale that despite the  $\mathcal{S}1$  output can contain errors, it furnishes approximate answers that cannot be obtained by resorting only to classical ontological inference.

If all results in  $C$  are inconsistent with those computed by  $\mathcal{S}2$ , a pair of classes is returned including  $c_0$  and the output of  $\mathcal{S}2$  having for actual parameters  $d$  and  $\text{Thing}$ , the meta-class of all the classes in the ontological formalism.

An important function provided by  $\mathcal{S}2$  regards the explanation of the detected inconsistencies. This function is obtained by recurring to standard DL reasoners.<sup>4</sup> One main problem encountered in the explanation of inconsistencies regards the fact that reasoners' output is usually quite verbose, since it provides the whole chain of all the possible reasons explaining why a given model is not consistent w.r.t. the represented assertions. For example, let us suppose that the ontological KB is provided with an assertion about the fact that *whale isA fish*. Since whales are not fishes (they are in the order of *cetacea* and cetacea are *mammalia*) an inconsistency is detected. The initial results obtained by

---

<sup>3</sup>The output of  $\mathcal{S}2$  cannot be wrong on a purely logical perspective, in that it is the result of a deductive process. The control strategy implements a tradeoff between ontological inference and the output of  $\mathcal{S}1$ , which is more informative but also less reliable from a formal point of view. However, in next future we plan to explore different conciliation mechanisms to ground the overall control strategy.

<sup>4</sup>To actually access the KBs we used the Jena framework, <https://jena.apache.org>.



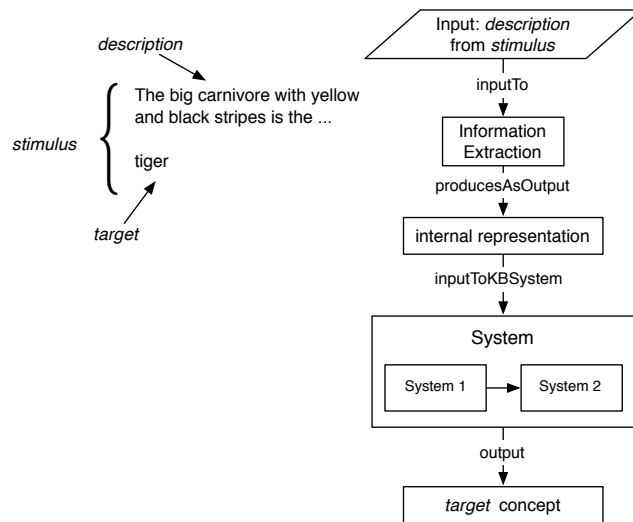


Figure 1: The software pipeline takes in input the linguistic description, queries the hybrid knowledge base and returns the categorized concept.

the reasoner will report all the clauses (i.e. all the models) provoking such an inconsistency. Not only it will report that —since *whale isA mammal* and *mammal* and *cetacea* are disjoint— whales cannot be fishes; it will also provide a huge amount of information about the facts that such an inconsistency has generated. For example, it will provide information about the fact that *whales* cannot be *reptiles*, *birds* and so on. Although factually correct and complete, this explanation is quite long (in an ontological KB with good coverage each class contains many subclasses) and is it not very informative for the punctual explanation of the raised inconsistency. The only relevant information, in this case, regards the fact that the tested class *whale* cannot be classified as *fish* because *mammal* and *fish* are mutually disjoint classes. Therefore we designed and implemented a software layer that runs on top of Jena explanation utilities to extract a *laconic explanation* from the longer one: the main focus of the laconic explanation is to make apparent the cause of the inconsistency. In so doing, we adopted a simple heuristic according to which the only explanation reported is that focused on the tuples of classes under investigation.

## 2.6 Categorization Pipeline of the Dual Process Architecture

The whole system embedding the proposed conceptual architecture works as follows. The input to the system is a simple linguistic description, like ‘The animal that eats bananas’, and the expected output is a given category referred to the description (e.g. the category *monkey* in this case). After the Information Extraction (IE) step, an internal representation is fed into the structure of the hybrid knowledge representation system which is then concerned with the categorization task by adopting the strategies described above.

A shallow IE approach has been devised: the morphological information computed from input sentences has been used to devise a finite-state automaton

to describe the input sentences’ structure. In this setting, the POS (Part-Of-Speech) information has been computed through the Stanford POS Tagger [31]. Then POS information has been used to encode the automaton states and transitions, that allow individuating salient information. In this automaton, states contain some kind of salient information required in the internal representation (e.g., the *place* where animals live; the description of their *skin* or *fur*; the *function* of an artifact, etc.; see below, Section 2.3). On the other hand, transitions between any two states encode connectives and prepositions that contain modifications to the the noun phrase they are referred to (like in ‘the big carnivore *with* [yellow and black stripes]’). This approach makes no use of the sentence dependency structure, and has many known limitations determined from merely using morphological information, and also inherent in finite-state machines (e.g., they cannot deal with parenthetical clauses, like relative clauses). It would not scale to handle more complex sorts of language. We defer to future work the adoption of richer language models; in particular, we will extend to the present context a deep semantic approach developed to perform IE from legal texts [20]. Despite these limitations, this approach allowed us to complete the automatization of the software pipeline going all throughout from the linguistic input description to its final conceptual categorization, thus improving the evaluation of the whole implemented system. In the following we describe the new experimentation and compare current and past results (where the IE step was performed in supervised fashion).

### 3 Experimentation

We have designed an experimentation on a categorization task to the ends of assessing the overall system. In the past experiments we tested the system over a restricted domain (the animal kingdom domain) [14]. Additionally, in [22] we tested the system in a broader context, since we were interested in assessing its robustness and the discriminative features of the  $\mathcal{S}1$  component in a multi-domain setting. Finally, in the present experimentation we tested the whole software pipeline, including the Information Extraction step (which, on the other hand, was simply performed in a supervised fashion in experiments 1 and 2).

The dataset used for the new experiment is composed of 87 “common-sense” linguistic descriptions.<sup>5</sup> Each stimulus  $st = \langle d, T \rangle$  is a pair of description and target, such as ⟨‘The big carnivore with yellow and black stripes’, ‘tiger’⟩ (please refer to Figure 1). The target  $T$  is the “prototypically correct” category, and in the following it is referred to as the *expected* result. The set of stimuli was devised by a team of neuropsychologists and philosophers in the frame of a broader project aimed at investigating the role of visual load in concepts involved in inferential and referential tasks (further details on neural correlates of lexical processing in [24]). The expected prototypical target category represents a gold standard, since it corresponds to the results provided within a psychological experimentation, where 30 subjects were requested to provide the corresponding target concept for each description. Such input was then used for querying our system as in a typicality based question-answering task.

<sup>5</sup>The full list of the stimuli is available at the URL: [http://www.di.unito.it/~radicion/datasets/CCIS\\_2014/stimuli.txt](http://www.di.unito.it/~radicion/datasets/CCIS_2014/stimuli.txt).

In Information Retrieval such queries are known to belong to the class of “informational queries”, i.e., queries where the user intends to obtain information regarding a specific information need [16]. Furthermore, this class of queries is characterized by uncertain and/or incomplete information, thereby resulting in the most complex to interpret, if compared to queries where users search for the URL of a given site (‘navigational queries’), or look for sites where some task can be performed, like buying music files (‘transactional queries’). Additionally, informational queries are by far the most common ones –based on the analysis of 1.5M user web logs [16]–, and they are therefore of the utmost applicative interest.

### 3.1 Summary of previous experiments

A first experiment was made by using as  $\mathcal{S}2$  component publicly available common sense domain ontologies. Namely, we selected the Animal in Context Ontology (ACO) developed by the Veterinary Medical Informatics Laboratory at the Virginia-Maryland Regional College, and the BBC WildLife Ontology.<sup>6</sup> They were both retrieved by using a mixed search over Sindice and Swoogle, and they were selected as guaranteeing a granularity of information and a coverage adequate for describing the stimuli being categorized (belonging to the animal domain). In this case we recorded a categorization accuracy of the  $\mathcal{S}1$ - $\mathcal{S}2$  about 95% with  $\mathcal{S}2$  using the ACO ontology, and over 92% with  $\mathcal{S}2$  using the BBC ontology. The accuracy was determined by comparing the results provided by the system with that ones provided by the human subjects in the above mentioned psychological experiment. Full details are provided in [25].

In order to deeply assess the accuracy of the system in a more demanding experimental setting, we devised a second experiment where we used as  $\mathcal{S}2$  the knowledge base OpenCyc.<sup>7</sup> OpenCyc is one of the largest ontologies publicly available, in that it is an enormous attempt at integrating many diverse semantic resources (such as, e.g., WordNet, DBpedia, Wikicompany, *etc.*). In this case we compared the results obtained by our  $\mathcal{S}1$ - $\mathcal{S}2$  systems with the results obtained by the Google and Bing search engines for the same queries [22]. Differently from the previous experiment, we decided to make use of an encyclopedic source of knowledge in order to investigate the differences with results obtained by plugging into  $\mathcal{S}2$  domain-specific knowledge bases. In this case the hybrid knowledge based  $\mathcal{S}1$ - $\mathcal{S}2$  system was able to categorize and retrieve most of the new typicality-based stimuli provided as input and still showed a better performance w.r.t. the general purpose search engines Google and Bing used in question-answering mode. The major problems encountered in this experiment derived from the difficulty of mapping the linguistic structure of stimuli containing very abstract meaning in the representational framework of  $\mathcal{S}1$  as they are actually encoded according to the conceptual space. For example, it was impossible to map the information contained in the description “the place where kings, princes and princesses live in fairy tales” onto the features used to characterize the prototypical representation of the concept *Castle*. On the other hand, the system shows good performances when dealing with less abstract descriptions based on perceptual and metric-reducible features such as

---

<sup>6</sup> Available at the URLs: <http://vts1.vetmed.vt.edu/aco/Ontology/aco.zip> and <http://www.bbc.co.uk/ontologies/wo>.

<sup>7</sup> <http://www.cyc.com/platform/opencyc>.

Table 1: Analysis of the POS-tagger and automaton failures and analysis of the correct results, in categorizing artifacts, plants, animals. Overall 36 descriptions of artifacts, 6 of plants and 45 of animals were considered.

	POS-tagger errors	automaton errors	successes in the IE step	$\Sigma$
# stimuli (%)	11 (12.6% of 87)	22 (25.3% of 87)	54 (62.1% of 87)	87
S1 successes	3 (27.3% of 11)	6 (27.3% of 22)	50 (92.6% of 54)	59

*shape, color, size*. The final categorization accuracy obtained by the system was around 77%, while Google and Bing successfully categorized 65% and 57.5% of linguistic descriptions, respectively.

### 3.2 New Experimentation

A new experiment to assess the whole pipeline has been devised: specifically, in this case the information available in the descriptions has been extracted through the automatic Information Extraction (IE) process and then used to query the hybrid knowledge base. The  $\mathcal{S}2$  system has been equipped with the OpenCyc ontology, like in the previous experiment. The dataset included 87 stimuli, 36 related to the artifacts domain, 6 to plants and 45 about animals.

Overall, 59 targets were correctly categorized out of the 87 input descriptions, thereby attaining 67.8% of correct responses. If we further examine the results, we can *i*) disaggregate cases where failures in the IE step were determined by the POS-tagger, or by the automaton that did not match the input descriptions structure; *ii*) compare failures to success cases. Interestingly, this analysis unveils that in some cases also in presence of noisy or lacking information the  $\mathcal{S}1$  component is able to retrieve the expected category. Detailed figures are reported in Table 1. The first two cells in the first row report the errors committed in the two stages preceding the access to the hybrid knowledge base. The second row illustrates how  $\mathcal{S}1$  recovers from partial or lacking information provided by the POS tagger and by the automaton:  $\mathcal{S}1$  obtained a correct categorization in 27% of cases when wrong and/or incomplete information was received by the POS tagger and by the automaton. On the other hand, the second row shows that even when all the linguistic information is correctly extracted and mapped onto the internal  $\mathcal{S}1$  format, it does not suffice to provide the expected categorization for all stimuli (the final accuracy for these cases was 92.6%). This datum is in line with previous results [14, 25]; although encouraging, it suggests that the current stage of development of the hybrid system is not yet sufficient to solve the prototypical categorization problem.

Differently from what emerged in past experimentation —where  $\mathcal{S}2$  detected some inconsistencies in the categories returned by  $\mathcal{S}1$ — in this case  $\mathcal{S}2$  did not detect an inconsistency of  $\mathcal{S}1$ ’s output as expected. In particular, given the description “An intelligent grey fish” (associated to the target concept *Dolphin*), the  $\mathcal{S}1$  system returned *Dolphin*, but  $\mathcal{S}2$  did not raise the inconsistency since OpenCyc erroneously represents *Dolphin* as a subclass of *Fish*, rather than a subclass of *Mammal*. Of course, by correcting the information available in the ontological knowledge base, we would obtain a refined, and ontologically correct,

result.

A concluding remark is about the advantages of the  $\mathcal{S}1$ - $\mathcal{S}2$  system w.r.t. adopting solely the  $\mathcal{S}2$  for the categorization task. As we have experimentally observed,  $\mathcal{S}2$  alone is not able to categorize this kind of descriptions [14], which proves to be a challenging task even for state-of-the-art search engines such as Google and Bing [22].

## 4 Related work

In the context of a different field of application, a solution similar to the one adopted here has been proposed in [4]. The main difference with their proposal concerns the underlying assumption on which the integration between symbolic and sub-symbolic system is based. In our system the conceptual spaces and the classical component are integrated at the level of the representation of concepts, and such components are assumed to convey different –though complementary– conceptual information. On the other hand, the previous proposal is mainly used to interpret and ground raw data coming from sensors in a high level symbolic system through the mediation of conceptual spaces.

In other respects, our system is also akin to that ones developed in the field of the computational approach to the above mentioned dual process theories. A first example of such “dual based systems” is the *mReasoner* model [19], developed with the aim of providing a computational architecture of reasoning based on the mental models theory proposed by Philip Johnson-Laird [17]. The *mReasoner* architecture is based on three components: a system 0, a system 1 and a system 2. The last two systems correspond to those hypothesized by the dual process approach. System 0 operates at the level of linguistic pre-processing. System 1 uses this intensional representation to build an extensional model, and uses heuristics to provide rapid reasoning conclusions; finally, system 2 carries out more demanding processes to search for alternative models, if the initial conclusion does not hold or if it is not satisfactory.

A second system that is close to our present work has been proposed by [28]. The authors do not explicitly mention the dual process approach; however, they build a system for conversational agents (chatbots) where agents’ background knowledge is represented using both a symbolic and a sub-symbolic approach. They also associate different sorts of representation to different types of reasoning. Namely, deterministic reasoning is associated to symbolic (system 2) representations, and associative reasoning is accounted for by the sub-symbolic (system 1) component. Differently from our system, however, the authors do not make any claim about the sequence of activation and the conciliation strategy of the two representational and reasoning processes. It is worth noting that other examples of this type of systems can be considered that are in some sense analogous to the dual process proposal: for example, many hybrid, symbolic-connectionist systems –including cognitive architectures such as, for example, CLARION<sup>8</sup>–, in which the connectionist component is used to model fast, associative processes, while the symbolic component is responsible for explicit, declarative computations. However, at the best of our knowledge, our system is the only one that considers this hybridization with a granularity at the level of individual conceptual representations.

---

<sup>8</sup><http://www.cogsci.rpi.edu/~rsun/clarion.html>

## 5 Conclusions

In this work we have presented a knowledge-based system relying upon a cognitively inspired architecture for the representation of conceptual knowledge. The system is grounded on a hybrid framework coupling classical and prototypical representation and reasoning, and it aims at extending the representational and reasoning capabilities of classical ontological-based systems towards more realistic and cognitively grounded scenarios, such as those envisioned by the prototype theory. The results obtained in new experimentation reveal that in a broader and composite domain including artifacts, animals and plants (w.r.t. past experimentation, mainly considering the animal kingdom) the proposed architecture provides encouraging results in tasks of prototype-based conceptual categorization.

In next future we plan to test the proposed approach in the area of biomedical domain to assess disease diagnosis tasks by grounding  $\mathcal{S}2$  on SNOWMED,<sup>9</sup> and  $\mathcal{S}1$  on conceptual spaces representing the typical symptoms of a given disease. A further development of the current work consists in extending the hybrid knowledge base with an additional layer of typicality-based representation (and reasoning): the exemplars. In particular, exemplar-based representations can be implemented by adopting the conceptual spaces [8] and therefore the proposed approach seems to be useful for the development of a system endowed with an integrated suite of categorization capacities. Such integrated prototype-exemplar based categorization could also be plausibly adopted in the area of cognitive architectures [21].

## Acknowledgments

This work has been partly supported by the Project *the Role of the Visual Imagery in Lexical Processing*, grant TO-call03-2012-0046 funded by Università degli Studi di Torino and Compagnia di San Paolo. The authors wish to thank Leo Ghignone, Andrea Minieri and Alberto Piana for working to the early stages of the project.

## References

- [1] Adams, B., Raubal, M.: A metric conceptual space algebra. In: Hornsby, K., et al. (eds.) COSIT. LNCS, vol. 5756, pp. 51–68. Springer (2009), <http://dblp.uni-trier.de/db/conf/cosit/cosit2009.html#AdamsR09>
- [2] Bejan, A., Marden, J.H.: Constructing animal locomotion from new thermodynamics theory. *American Scientist* 94(4), 342 (2006)
- [3] Brachmann, R.J., Schmolze, J.G.: An overview of the KL-ONE knowledge representation system. *Cognitive Science* 9(2), 171–202 (April 1985)
- [4] Chella, A., Frixione, M., Gaglio, S.: A cognitive architecture for artificial vision. *Artificial Intelligence* 89(1–2), 73 – 111 (1997), <http://www.sciencedirect.com/science/article/pii/S0004370296000392>

---

<sup>9</sup><http://www.b2international.com/portal/snow-owl>.

- [5] Evans, J.S.B., Frankish, K.E.: In two minds: Dual processes and beyond. Oxford University Press (2009)
- [6] Frixione, M., Lieto, A.: Representing concepts in formal ontologies: Compositionality vs. typicality effects. *Logic and Logical Philosophy* 21(4), 391–414 (2012)
- [7] Frixione, M., Lieto, A.: Dealing with Concepts: from Cognitive Psychology to Knowledge Representation. *Frontiers of Psychological and Behavioural Science* 2(3), 96–106 (July 2013)
- [8] Frixione, M., Lieto, A.: Exemplars, prototypes and conceptual spaces. In: *Biologically Inspired Cognitive Architectures 2012*, pp. 131–136. Springer (2013)
- [9] Frixione, M., Lieto, A.: Representing Non Classical Concepts in Formal Ontologies: Prototypes and Exemplars. In: *New Challenges in Distributed Information Filtering and Retrieval. Studies in Computational Intelligence*, vol. 439, pp. 171–182 (2013)
- [10] 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), 257–276 (March 2014)
- [11] Gärdenfors, P.: *Conceptual spaces: The geometry of thought*. MIT press (2000)
- [12] Gärdenfors, P.: *The Geometry of Meaning: Semantics Based on Conceptual Spaces*. MIT Press (2014)
- [13] Getz, W.M.: Biomass transformation webs provide a unified approach to consumer–resource modelling. *Ecology letters* 14(2), 113–124 (2011)
- [14] Ghignone, L., Lieto, A., Radicioni, D.P.: Typicality-Based Inference by Plugging Conceptual Spaces Into Ontologies. In: Lieto, A., Cruciani, M. (eds.) *Proceedings of the International Workshop on Artificial Intelligence and Cognition*. CEUR (2013)
- [15] Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: A non-monotonic description logic for reasoning about typicality. *Artificial Intelligence* 195, 165–202 (2013)
- [16] Jansen, B.J., Booth, D.L., Spink, A.: Determining the informational, navigational, and transactional intent of web queries. *Information Processing & Management* 44(3), 1251–1266 (2008)
- [17] Johnson-Laird, P.: Mental models in cognitive science. *Cognitive Science* 4(1), 71–115 (1980)
- [18] Kahneman, D.: *Thinking, fast and slow*. Macmillan (2011)
- [19] Khemlani, S., Johnson-Laird, P.: The processes of inference. *Argument & Computation* 4(1), 4–20 (2013)

- [20] Lesmo, L., Mazzei, A., Palmirani, M., Radicioni, D.P.: TULSI: an NLP System for Extracting Legal Modificatory Provisions. *Artificial Intelligence and Law* 12(4), 1–34 (2012)
- [21] Lieto, A.: A computational framework for concept representation in cognitive systems and architectures: Concepts as heterogeneous proxytypes. *Procedia Computer Science* 41(0), 6–14 (2014), <http://www.sciencedirect.com/science/article/pii/S1877050914015233>, 5th Annual International Conference on Biologically Inspired Cognitive Architectures (BICA)
- [22] Lieto, A., Minieri, A., Piana, A., Radicioni, D.P.: A knowledge-based system for prototypical reasoning. *Connection Science* (In press, 2014)
- [23] Machery, E.: *Doing without concepts*. OUP (2009)
- [24] Marconi, D., Manenti, R., Catricalà, E., Della Rosa, P.A., Siri, S., Cappa, S.F.: The neural substrates of inferential and referential semantic processing. *Cortex* 49(8), 2055–2066 (2013)
- [25] Minieri, A., Lieto, A., Piana, A., Radicioni, D.P.: A Dual Process Architecture for Ontology-Based Systems. In: Filipe, J., Dietz, J., Aveiro, D. (eds.) *Proceedings of KEOD 2014, 6th International Conference on Knowledge Engineering and Ontology Development*. pp. 48–55. SCITEPRESS – Science and Technology Publications (2014)
- [26] Minsky, M.: A framework for representing knowledge. In: Winston, P. (ed.) *The Psychology of Computer Vision*, pp. 211–277. McGraw-Hill, New York (1975), [ftp://publications.ai.mit.edu/ai-publications/\discretionary-{}{}{}pdf/\discretionary-{}{}{}AIM-306.pdf](ftp://publications.ai.mit.edu/ai-publications/\discretionary-{}{}pdf/\discretionary-{}{}{}AIM-306.pdf)
- [27] Nardi, D., Brachman, R.J.: An introduction to description logics. In: *Description logic handbook*. pp. 1–40 (2003)
- [28] Pilato, G., Augello, A., Gaglio, S.: A modular system oriented to the design of versatile knowledge bases for chatbots. *ISRN Artificial Intelligence* 2012 (2012)
- [29] Rosch, E.: Cognitive representations of semantic categories. *J. Exp. Psychol. Gen.* 104(3), 192–233 (1975)
- [30] Straccia, U.: Reasoning within fuzzy description logics. arXiv preprint arXiv:1106.0667 (2011)
- [31] Toutanova, K., Klein, D., Manning, C.D., Singer, Y.: Feature-rich part-of-speech tagging with a cyclic dependency network. In: *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology-Volume 1*. pp. 173–180. Association for Computational Linguistics (2003)