

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Typicality-Based Inference by Plugging Conceptual Spaces Into Ontologies

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/146806> since

Publisher:

CEUR Workshop Proceedings

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)



UNIVERSITÀ DEGLI STUDI DI TORINO

This is an author version of the contribution published on:

Leo Ghignone, Antonio Lieto, Daniele P. Radicioni
Typicality-Based Inference by Plugging Conceptual Spaces Into Ontologies
Editor: CEUR
2013

in

Proceedings of the First International Workshop on Artificial Intelligence and
Cognition (AIC 2013)
68 - 79
International Workshop on Artificial Intelligence and Cognition
Turin, Italy
December, 3rd 2013

The definitive version is available at:

<http://www.di.unito.it/~lieto/AIC2013/>

Typicality-Based Inference by Plugging Conceptual Spaces Into Ontologies

Leo Ghignone, Antonio Lieto, and Daniele P. Radicioni

Università di Torino, Dipartimento di Informatica, Italy
{lieto,radicion}@di.unito.it
leo.ghignone@gmail.com

Abstract. In this paper we present a cognitively inspired system for the representation of conceptual information in an ontology-based environment. It builds on the heterogeneous notion of concepts in Cognitive Science and on the so-called *dual process theories* of reasoning and rationality, and it provides a twofold view on the same artificial concept, combining a classical symbolic component (grounded on a formal ontology) with a typicality-based one (grounded on the conceptual spaces framework). The implemented system has been tested in a pilot experimentation regarding the classification task of linguistic stimuli. The results show that this modeling solution extends the representational and reasoning “conceptual” capabilities of standard ontology-based systems.

1 Introduction

Representing and reasoning on *common sense* concepts is still an open issue in the field of knowledge engineering and, more specifically, in that of formal ontologies. In Cognitive Science evidences exist in favor of prototypical concepts, and typicality-based conceptual reasoning has been widely studied. Conversely, in the field of computational models of cognition, most contemporary concept oriented knowledge representation (KR) systems, including formal ontologies, do not allow –for technical convenience– neither the representation of concepts in prototypical terms nor forms of approximate, non monotonic, conceptual reasoning. In this paper we focus on the problem of concept representation in the field of formal ontologies and we introduce, following the approach proposed in [1], a cognitively inspired system to extend the representational and reasoning capabilities of the ontology based systems.

The study of concept representation concerns different research areas, such as Artificial Intelligence, Cognitive Science, Philosophy, etc.. In the field of Cognitive Science, the early work of Rosch [2] 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 [3]– to represent and to conduct reasoning on concepts in “non classical”, prototypical terms [4].

However, these systems lacked in clear formal semantics, and were later sacrificed in favor of a class of formalisms stemmed from structured inheritance semantic networks: the first system in this line of research was the KL-ONE system [5]. These formalisms are known today as description logics (DLs). In this setting, the representation of prototypical information (and therefore the possibility of performing non monotonic reasoning) is not allowed,¹ since the formalisms in this class are primarily intended for deductive, logical inference. Nowadays, DLs are largely adopted in diverse application areas, in particular within the area of ontology representation. For example, OWL and OWL 2 formalisms follow this tradition,² which has been endorsed by the W3C for the development of the Semantic Web. However, under a historical perspective, the choice of preferring classical systems 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, nonetheless various theoretical and practical problems remain unsolved [6].

As a possible way out, we follow the proposal presented in [1], that relies on two main cornerstones: the dual process theory of reasoning and rationality [7,8,9], and the heterogeneous approach to the concepts in Cognitive Science [10]. This paper has the following major elements of interest: *i*) we provided the hybrid architecture envisioned in [1] with a working implementation; *ii*) we show how the resulting system is able to perform a simple form of categorization, that would be unfeasible by using only formal ontologies; *iii*) we propose a novel access strategy (different from that outlined in [1]) to the conceptual information, closer to the tenets of the dual process approach (more about this point later on).

The paper is structured as follows: in Section 2 we illustrate the general architecture and the main features of the implemented system. In Section 3 we provide the results of a preliminary experimentation to test inference in the proposed approach, and, finally, we conclude by presenting the related work (Section 4) and by outlining future work (Section 5).

2 The System

A system has been implemented to explore the hypothesis of the hybrid conceptual architecture. To test it, we have been considering a basic *inference* task: given an input description in natural language, the system should be able to find,

¹ This is the case, for example, of *exceptions* to the inheritance mechanism.

² For the Web Ontology Language, see <http://www.w3.org/TR/owl-features/> and <http://www.w3.org/TR/owl2-overview/>, respectively.

even for typicality based description (that is, most of common sense descriptions), the corresponding concept category by combining ontological inference and typicality based one. To these ends, we developed a domain ontology (the *naive animal ontology*, illustrated below) and a parallel typicality description as a set of domains in a conceptual space framework [11].

In the following, *i)* we first outline the design principles that drove the development of the system; *ii)* we then provide an overview of the system architecture and of its components and features; *iii)* we elaborate on the inference task, providing the detailed control strategy; and finally *iv)* we introduce the domain ontology and the conceptual space used as case study applied over the restricted domain of animals.

2.1 Background and architecture design

The theoretical framework known as *dual process theory* postulates the co-existence of two different types of cognitive systems. The systems³ 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 [7,8,9]. According to the reasons presented in [12,1], the conceptual representation of our systems should be equipped with two major sorts of components, based on:

- type 1 processes, to perform fast and approximate categorization by taking advantage from prototypical information associated to concepts;
- type 2 processes, involved in complex inference tasks and that do not take into account the representation of prototypical knowledge.

Another theoretical framework inspiring our system regards the heterogeneous approach to the concepts in Cognitive Science, according to which concepts do not constitute a unitary element (see [10]).

Our system is equipped, then, with a hybrid conceptual architecture based on a classical component and on a typical component, each encoding a specific reasoning mechanism as in the dual process perspective. Figure 1 shows the general architecture of the hybrid conceptual representation.

The ontological component is based on a classical representation grounded on a DL formalism, and it allows specifying the necessary and/or sufficient conditions for concept definition. For example, if we consider the concept *water*, the classical component will contain the information that *water* is exactly the chemical substance whose formula is H_2O , i.e., 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 its prototypical traits, such as the fact that water occurring in liquid state is usually a colorless, odorless and tasteless fluid.

³ We assume that each system type can be composed by many sub-systems and processes.




Fig. 1: Architecture of the hybrid system.

By adopting the “dual process” notation, in our system the representational and reasoning functions are assigned to the system 1 (executing processes of type 1), and they are associated to the Conceptual Spaces framework [11]. Both from a modeling and from a reasoning point of view, system 1 is compliant with the traits of conceptual typicality. On the other hand, the representational and reasoning functions assigned to the system 2 (executing processes of type 2) are associated to a classical DL-based ontological representation. Differently from what proposed in [1], the access to the information stored and processed in both components is assumed to proceed from the system 1 to the system 2, as suggested by the central arrow in Figure 1.

We now briefly introduce the representational frameworks upon which system 1 (henceforth $\mathcal{S}1$) and system 2 (henceforth $\mathcal{S}2$) have been designed.

As mentioned, the aspects related to the typical conceptual component $\mathcal{S}1$ are modeled through Conceptual Spaces [11]. Conceptual spaces (CS) are a geometrical framework for the representation of knowledge, consisting in a set of *quality dimensions*. In some cases, such dimensions can be directly related to perceptual mechanisms; examples of this kind are temperature, weight, brightness, pitch. In other cases, dimensions can be more abstract in nature. A geometrical (topological or metrical) structure is associated to each quality dimension. The chief idea is that knowledge representation can benefit from the geometrical structure of conceptual spaces: instances are represented as points in a space, and their similarity can be calculated in the terms of their distance according to some suitable distance measure. In this setting, concepts correspond to regions, and regions with different geometrical properties correspond to different kinds of concepts. Conceptual spaces are suitable to represent concepts in “typical” terms, since the regions representing concepts have soft boundaries. In many cases typicality effects can be represented in a straightforward way: for example, in the case of concepts, corresponding to convex regions of a conceptual space, prototypes have a natural geometrical interpretation, in that they correspond to the geometrical centre of the region itself. Given a convex region, we can

provide each point with a certain centrality degree, that can be interpreted as a measure of its typicality. Moreover, single exemplars correspond to single points in the space. This allows us to consider both the exemplar and the prototypical accounts of typicality (further details can be found in [13, p. 9]).

On the other hand, the representation of the classical component $\mathcal{S}2$ has been implemented based on 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 simply contemplate deductive processes.

2.2 Inference in the hybrid system

Categorization (i.e., to classify a given data instance into a predefined set of categories) is one of the classical processes 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. In turn, prototypical values need not be specified for all class individuals, that vice versa can overwrite them: one typical example is the case of birds that (by default) fly, except for special birds, like penguins, that do not fly.

The whole categorization process regarding our system 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. The $\mathcal{S}1$ component takes in input the information extracted from the description d , and produces in output a set of classes $C = \{c_1, c_2, \dots, c_n\}$. This set of results is then checked against cc , the output of $\mathcal{S}2$ (Algorithm 1, line 3): the step is performed by adding to the ontology an individual from the class $c_i \in C$, modified by the information extracted from d , and by checking the consistency of the newly added element with a DL reasoner.

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 top scoring class returned by $\mathcal{S}1$ (Algorithm 1: line 8). If cc –the class computed by $\mathcal{S}2$ – is a superclass or 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*.⁴ If all results in C are

⁴ 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 tries to implement 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.

Algorithm 1 Inference in the hybrid system.

input : textual description d
output : a class assignment, as computed by $S1$ and $S2$
1: $C \leftarrow S1(d)$ /* conceptual spaces output */
2: **for each** $c_i \in C$ **do**
3: $cc \leftarrow S2(\langle d, c_i \rangle)$ /* 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 or superclass of c_i **then**
11: **return** $\langle c_0, cc \rangle$
12: **end if**
13: **end if**
14: **end for**
15: $cc \leftarrow S2(\langle d, \text{Thing} \rangle)$
16: **return** $\langle c_0, cc \rangle$

inconsistent with those computed by $S2$, a pair of classes is returned including c_0 and the output of $S2$ having for actual parameters d and Thing , the meta class of all the classes in the ontological formalism.

2.3 Developing the Ontology

A formal ontology has been developed describing the animal kingdom. It has been devised to meet common sense intuitions, rather than reflecting the precise taxonomic knowledge of ethologists, so we denote it as *naïve animal ontology*.⁵ In particular, the ontology contains the taxonomic distinctions that have an intuitive counterpart in the way human beings categorize the corresponding concepts. Classes are collapsed at a granularity level such that they can be naturally grouped together also based on their *accessibility* [14]. For example, although the category *pachyderm* is no longer in use by ethologists, we created a *pachyderm* class that is superclass to *elephant*, *hippopotamus*, and *rhinoceros*. The underlying *rationale* is that it is still in use by non experts, due to the intuitive resemblances among its subclasses.

The ontology is linked to DOLCE's *Lite* version;⁶ in particular, the tree containing our taxonomy is rooted in the *agentive-physical-object* class, while the body components are set under *biological-physical-object*, and partitioned between the two disjunct classes *head-part* (e.g., for framing horns, antennas, fang, etc.) and *body-part* (e.g., for paws, tails, etc.). The *biological-object* class in-

⁵ The ontology is available at the URL http://www.di.unito.it/~radicion/datasets/aic_13/Naive_animal_ontology.owl

⁶ <http://www.loa-cnr.it/ontologies/DOLCE-Lite.owl>

cludes different sorts of skins (such as *fur*, *plumage*, *scales*), substances produced and eaten by animals (e.g., *milk*, *wool*, *poison* and *fruits*, *leaves* and *seeds*).

2.4 Formalizing conceptual spaces and distance metrics

The conceptual space defines a metric space that can be used to compute the proximity of the input entities to prototypes. To compute the distance between two points p_1, p_2 we apply a distance metrics based on the combination of the Euclidean distance and the angular distance intervening between the points. Namely, we use Euclidean metrics to compute within-domain distance, while for dimensions from different domains we use the Manhattan distance metrics, as suggested in [11,15]. Weights assigned to domain dimensions are affected by the context, too, so the resulting weighted Euclidean distance $dist_E$ is computed as follows

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

where i varies over the n domain dimensions, k is the context, and w_i are dimension weights.

The representation format adopted in conceptual spaces (e.g., for the concept *whale*) includes information such as:

```
02062744n,whale,dimension(x=350,y=350,z=2050),color(B=20,H=20,S=60),food=10
```

that is, the WordNet synset identifier, the *lemma* of the concept in the description, information about its typical dimensions, color (as the position of the instance on the three-dimensional axes of *brightness*, *hue* and *saturation*) and food. Of course, information about typical traits varies according to the species. Three domains with multiple dimensions have been defined:⁷ *size*, *color* and *habitat*. Each quality in a domain is associated to a range of possible values. To avoid that larger ranges affect too much the distance, we have introduced a damping factor to reduce this effect; also, the relative strength of each domain can be parametrized.

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*. Cosine similarity is computed as the cosine of the angle between the considered vectors: 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 and context k is computed as

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

⁷ We defined also further domains with one dimension (e.g., whiskers, wings, paws, fang, and so forth), but for our present concerns they are of less interest. The conceptual space is available at the URL http://www.di.unito.it/~radicion/datasets/aic_13/conceptual_space.txt.

Moreover, to satisfy the triangle inequality is a requirement upon distance in a metric space; unfortunately, cosine similarity does not satisfy triangle inequality, so we adopt a slightly different metrics, the *angular similarity* (\hat{as}), whose values vary over the range $[0, 1]$, and that is defined as

$$\hat{as}(p_1, p_2) = 1 - \frac{2 \cdot \cos^{-1} \cdot \hat{cs}(p_1, p_2, k)}{\pi}.$$

Angular distance allows us to compare the shape of animals disregarding their actual size: for example, it allows us to find that a python is similar to a viper even though it is much bigger.

In the metric space being defined, the distance d between individuals i_a, i_b is computed with the Manhattan distance, enriched with information about context k that indicates the set of weights associated to each domain. Additionally, the relevance of domains with fewer dimensions (that would obtain overly high weights) is counterbalanced by a normalizing factor (based on the work by [15]), so that such distance is computed as:

$$d(i_a, i_b, K) = \sum_{j=1}^m w_j \cdot \sqrt{|D_j|} \cdot \text{dist}_E(p_j(i_a), p_j(i_b), k_j), \quad (1)$$

where K is the whole context, containing domain weights w_j and contexts k_j , and $|D_j|$ is the number of dimensions in each domain.

In this setting, the distance between each two concepts can be computed as the distance between two regions in a given domain, and then to combining them through the Formula 1. Also, we can compute the distance between any two region prototypes, or the minimal distance between their individuals, or we can apply more sophisticated algorithms: in all cases, we have designed a metric space and procedures that allow characterizing and comparing concepts herein. Although angular distance is currently applied to compute similarity in the *size* of the considered individuals, it can be generalized to further dimensions.

3 Experimentation

The evaluation consisted of an inferential task aimed at categorizing a set of linguistic descriptions. Such descriptions contain information related to concepts typical features. Some examples of these common-sense descriptions are: “the big carnivore with black and yellow stripes” denoting the concept of *tiger*, and “the sweet water fish that goes upstream” denoting the concept of *salmon*, and so on. A dataset of 27 “common-sense” linguistic descriptions was built, containing a list of stimuli and their corresponding category: this is the “prototypically correct” category, and in the following is referred to as the *expected* result.⁸ The set of stimuli was devised by a team of neuropsychologists and philosophers in

⁸ The full list is available at the URL http://www.di.unito.it/~radicion/datasets/aic_13/stimuli_en.txt.

Table 1: Results of the preliminary experimentation.

Test cases categorized	27	100.0%
[1.] Cases where $\mathcal{S}1$ and $\mathcal{S}2$ returned the same category	24	88.9%
[2a.] Cases where $\mathcal{S}1$ returned the expected category	25	92.6%
[2b.] Cases where $\mathcal{S}2$ returned the expected category	26	96.3%
Cases where $\mathcal{S}1$ OR $\mathcal{S}2$ returned the expected category	27	100.0%

the frame of a broader project, aimed at investigating the role of visual load in concepts involved in inferential and referential tasks. Such input was used for querying the system as in a typicality based question-answering task. 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. Since it is characterized by uncertain and/or incomplete information, this class of queries is by far the most common and complex to interpret, if compared to queries where users can 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’) [16].

We devised some metrics to assess the accuracy of the system, and namely we recorded the following information:

1. how often $\mathcal{S}1$ and $\mathcal{S}2$ returned in output the same category;
2. in case different outputs were returned, the accuracy obtained by $\mathcal{S}1$ and $\mathcal{S}2$:
 - 2a. the accuracy of $\mathcal{S}1$. This figure is intended to measure how often the top ranked category c_0 returned by $\mathcal{S}1$ is the same as that expected.
 - 2b. the accuracy of $\mathcal{S}2$, that is the second category returned in the output pair $\langle c, cc \rangle$. This figure is intended to measure how often the cc category is the appropriate one w.r.t. the expected result. We remark that cc has not been necessarily computed by starting from c_0 : in principle any $c_i \in C$ might have been used (see also Algorithm 1, lines 3 and 15).

The results obtained in this preliminary experimentation are presented in Table 1. All of the stimuli were categorized, although not all of them were correctly categorized. However, the system was able to correctly categorize a vast majority of the input descriptions: in most cases (92.6%) $\mathcal{S}1$ alone produces the correct output, with considerable saving in terms of computation time and resources. Conversely, none of the concepts (except for one) described with typical features would have been classified through classical ontological inference. It is in virtue of the former access to conceptual spaces that the whole system is able to categorize such descriptions. Let us consider, e.g., the description “The animal that eats bananas”. The ontology encodes knowledge stating that *monkeys* are *omnivore*. However, since the information that *usually* monkeys eat bananas cannot be represented therein, the description would be consistent to all omnivores. The information returned would then be too informative w.r.t. the granularity of the expected answer.

Another interesting result was obtained for the input description “the big herbivore with antlers”. In this case, the correct answer is the third element in the list C returned by $S1$; but thanks to the categorization performed by $S2$, it is returned in the final output pair (see Algorithm 1, line 8).

Finally, the system revealed to be able to categorize stimuli with typical, though ontologically incoherent, descriptions. As an example of such a case we will consider the categorization results obtained with the following stimulus: “The big fish that eats plankton”. In this case the prototypical answer expected is *whale*. However, whales properly are mammals, not fishes. In our hybrid system, $S1$ component returns *whale* by resorting to prototypical knowledge. If further details were added to the input description, the answer would have changed accordingly: in this sense the categorization performed by $S1$ is non monotonic in nature. When then C (the output of $S1$) is checked against the ontology as described by the Algorithm 1 at lines 7–13, and an inconsistency is detected,⁹ the consistency of the second result in C (*shark* in this example) is tested against the ontology. Since this answer is an ontologically compliant categorization, then this solution is returned by the $S2$ component. The final output of the categorization is then the pair $\langle whale, shark \rangle$: the first element, prototypically relevant for the query, would have not been provided by querying a classical ontological representation. Moreover, if the ontology recorded the information that also other fishes do eat plankton, the output of a classical ontological inference would have included them, too, thereby resulting in a too large set of results w.r.t. the intended answer.

4 Related work

In the context of a different field of application, a solution similar to the one adopted here has been proposed in [17]. 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 carry different –though complementary– conceptual information. On the other hand, the previous proposal is mainly used to interpret and ground raw data coming from sensor 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 [18], developed with the aim of providing a computational architecture of reasoning based on the mental models theory proposed by Philip Johnson-Laird [19]. 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. It parses

⁹ This follows by observing that $c_0 = whale$, $cc = shark$; and $whale \subset mammal$, while $shark \subset fish$; and $mammal$ and $fish$ are disjoint.

the premises of an argument by using natural language processing techniques, and it then creates an initial intensional model of them. 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. Another system that is close to our present work has been proposed by [20]. 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 subsymbolic 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 subsymbolic (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 akin to the dual process proposal: for example, many hybrid, symbolic-connectionist systems –including cognitive architectures such as, for example, CLARION (<http://www.cogsci.rpi.edu/~rsun/clarion.html>)–, in which the connectionist component is used to model fast, associative processes, while the symbolic component is responsible for explicit, declarative computations (for a deeper discussion, please refer to [21]). 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.

5 Conclusions and future work

In this paper we presented a cognitively inspired system to extend the representational and reasoning capabilities of classical ontological representations. We tested it in a pilot study concerning a categorization task involving typicality based queries. The results show that the proposed architecture effectively extends the reasoning and representational capabilities of formal ontologies towards the domain of prototype theory.

Next steps will be to complete the implementation of current system: first, we will work to the automatization of the Information Extraction from linguistic descriptions, and then to the automatization of the mapping of the extracted information onto the conceptual representations in $S1$ and $S2$. In near future we will also extend the coverage of the implemented system to further domains.

Yet, we are designing a learning setting to modify weights in conceptual spaces according to experience (thereby qualifying the whole system as a supervised learning one). This line of research will require the contribution of theoretical and experimental psychologists, to provide insightful input to the development of the system, and experimental corroboration to its evolving facets, as well. Future work will also include the evaluation of the system on web data, namely to experiment by using search engine web logs, in order to verify whether

and to what extent the implemented system matches the actual users' informational needs.

References

1. Frixione, M., Lieto, A.: Dealing with Concepts: from Cognitive Psychology to Knowledge Representation. *Frontiers of Psychological and Behavioural Science* **2**(3) (July 2013) 96–106
2. Rosch, E.: Cognitive representations of semantic categories. *Journal of experimental psychology: General* **104**(3) (1975) 192–233
3. Minsky, M.: A framework for representing knowledge. In Winston, P., ed.: *The Psychology of Computer Vision*. McGraw-Hill, New York (1975) 211–277
4. Brachman, R.J., Levesque, H.J.: *Readings in Knowledge Representation*. Morgan Kaufmann Pub (1985)
5. Brachmann, R.J., Schmolze, J.G.: An overview of the KL-ONE knowledge representation system. *Cognitive Science* **9**(2) (April 1985) 171–202
6. Frixione, M., Lieto, A.: The computational representation of concepts in formal ontologies-some general considerations. In: *KEOD*. (2010) 396–403
7. Stanovich, K.E., West, R.F., et al.: Individual differences in reasoning: Implications for the rationality debate? *Behavioral and brain sciences* **23**(5) (2000) 645–665
8. Evans, J.S.B., Frankish, K.E.: *In two minds: Dual processes and beyond*. Oxford University Press (2009)
9. Kahneman, D.: *Thinking, fast and slow*. Macmillan (2011)
10. Machery, E., et al.: *Doing without concepts*. Oxford University Press Oxford (2009)
11. Gärdenfors, P.: *Conceptual Spaces*. MIT Press (2000)
12. Frixione, M., Lieto, A.: Representing concepts in formal ontologies: Compositionality vs. typicality effects. *Logic and Logical Philosophy* **21**(4) (2012) 391–414
13. Frixione, M., Lieto, A.: Representing Non Classical Concepts in Formal Ontologies: Prototypes and Exemplars. In: *New Challenges in Distributed Information Filtering and Retrieval*. Volume 439 of *Studies in Computational Intelligence*. (2013) 171–182
14. Smith, E.R., Branscombe, N.R.: Category accessibility as implicit memory. *Journal of Experimental Social Psychology* **24**(6) (1988) 490–504
15. Adams, B., Raubal, M.: A metric conceptual space algebra. In Hornsby, K.S., Claramunt, C., Denis, M., Ligozat, G., eds.: *COSIT*. Volume 5756 of *Lecture Notes in Computer Science*, Springer (2009) 51–68
16. Jansen, B.J., Booth, D.L., Spink, A.: Determining the informational, navigational, and transactional intent of web queries. *Information Processing & Management* **44**(3) (2008) 1251–1266
17. Chella, A., Frixione, M., Gaglio, S.: A cognitive architecture for artificial vision. *Artificial Intelligence* **89**(1–2) (1997) 73 – 111
18. Khemlani, S., Johnson-Laird, P.: The processes of inference. *Argument & Computation* **4**(1) (2013) 4–20
19. Johnson-Laird, P.: Mental models in cognitive science. *Cognitive Science* **4**(1) (1980) 71–115
20. Pilato, G., Augello, A., Gaglio, S.: A modular system oriented to the design of versatile knowledge bases for chatbots. *ISRN Artificial Intelligence* **2012** (2012)
21. Frixione, M., Lieto, A.: Formal Ontologies and Semantic Technologies: A Dual Process Proposal for Concept Representation. *Philosophia Scientiae* (forthcoming)