

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

A Knowledge-Based System for Prototypical Reasoning

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1550432> since 2017-12-01T17:45:25Z

Published version:

DOI:10.1080/09540091.2014.956292

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)

Preprint version. The final version is available on the Taylor & Francis site:
<http://www.tandfonline.com>.

RESEARCH ARTICLE

A Knowledge-Based System for Prototypical Reasoning

Antonio Lieto^{*,1,2}, Andrea Minieri¹, Alberto Piana¹ and Daniele P. Radicioni^{*,1}

¹*Dipartimento di Informatica – Università degli Studi di Torino
Corso Svizzera 185 10149 – Turin, ITALY*

²*ICAR - CNR, Viale delle Scienze Ed. 11, 90128, Palermo, Italy*

()

In this work we present a knowledge-based system equipped with a hybrid, cognitively inspired architecture for the representation of conceptual information. The proposed system aims at extending the classical representational and reasoning capabilities of the ontology-based frameworks towards the realm of the prototype theory. It is based on a hybrid knowledge base, composed of a classical symbolic component (grounded on a formal ontology) with a typicality based one (grounded on the conceptual spaces framework). The resulting system attempts to reconcile the heterogeneous approach to the concepts in Cognitive Science with the *dual process theories* of reasoning and rationality. 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 extends the representational and reasoning “conceptual” capabilities of standard ontology-based systems.

Keywords: Knowledge Representation; Formal Ontologies; Conceptual Spaces; Common Sense Reasoning; Dual Process Theory; Prototypical Reasoning.

1. Introduction

Representing and reasoning on *common sense* concepts is still an open issue in the field of knowledge representation (KR) 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 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. Following the approach proposed in [13] we introduce a conceptual architecture that, embedded in a larger knowledge-based system, aims at extending the representational and reasoning capabilities available to traditional ontology-based frameworks.

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 [34], preceded by the philosophical analysis of Wittgenstein [40], showed that ordinary concepts do not obey the classical theory (stating that concepts can be defined in terms of sets of necessary and sufficient

*Corresponding author. Email: {lieto,radicion}@di.unito.it

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 [29] and semantic networks [33]– to represent and to conduct reasoning on concepts in “non classical”, prototypical terms [3].

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 KL-ONE [5]. These formalisms are known today as description logics (DLs) [31]. 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 [38, 17, 2, 6], nonetheless various theoretical and practical problems remain unsolved [10].

As a possible way out, we follow the proposal presented in [13], that relies on two main cornerstones: the dual process theory of reasoning and rationality [37, 9, 22], and the heterogeneous approach to the concepts in Cognitive Science [26]. This paper has the following major elements of interest: *i*) we provide the hybrid conceptual architecture envisioned in [13] with a working implementation; *ii*) we show how the system implementing such architecture is able to perform a simple form of non-monotonic categorization that is, *vice versa*, unfeasible by using only formal ontologies.

The paper is structured as follows: in Section 2 we illustrate the general architecture and the main features of the knowledge-based system. 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 the related work (Section 4) and by outlining future work (Section 5).

2. The System

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 knowledge base architecture and of its components and features, based on the conceptual spaces framework [15, 16] and on formal ontologies [18]; *iii*) we elaborate on the inference task, providing the detailed control strategy.

Two cornerstones inspiring the current proposal are the *dual process theory* and the heterogeneous approach to concepts in Cognitive Science. The theoretical framework known as *dual process theory* postulates the co-existence of two

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

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

different types of cognitive systems [37, 9, 22]. 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 each system type can be composed of many sub-systems and processes. According to the reasons presented in [11, 13], the conceptual representation of our system relies on 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.

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.

The second 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 from a representational point of view [26]. By following this approach, we assume that each concept represented in an artificial system can be composed of several bodies of knowledge, each one carrying a specific type of information.¹

A system has been implemented to explore the hypothesis of the hybrid conceptual architecture. To test it, we have considered a basic *inference* task: given an input description in natural language, the system should be able to find, even for typicality based description (that is, most of common sense descriptions), the corresponding concept category by combining ontological inference and typicality based one. We chose this task as a challenging one. In fact, classical queries for concept retrieval based on lists of necessary and sufficient conditions are commonly handled by standard ontology-based systems, and in general by logic oriented systems. Conversely, the answer to typicality based queries –i.e., queries based on prototypical traits–, is almost never addressed by exploiting ontological inference.

2.1. Knowledge Base Architecture

Our system is equipped, then, with a hybrid conceptual architecture based on a classical component and on a typical one. Each component represents a specific conceptual body of knowledge together with the related reasoning procedures 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.

¹In the present implementation we considered two possible types of informational components: the typical one (encoding prototypical knowledge) and the classical one (encoding information in terms of necessary and sufficient conditions). In particular, although in this case we mainly concentrate on representation and reasoning tenets coming from the prototype theory, the typical component can be considered general enough to encode many other forms of representational and reasoning mechanisms related to a wider spectrum of typicality theories such as, for example, the Exemplars theory [30].

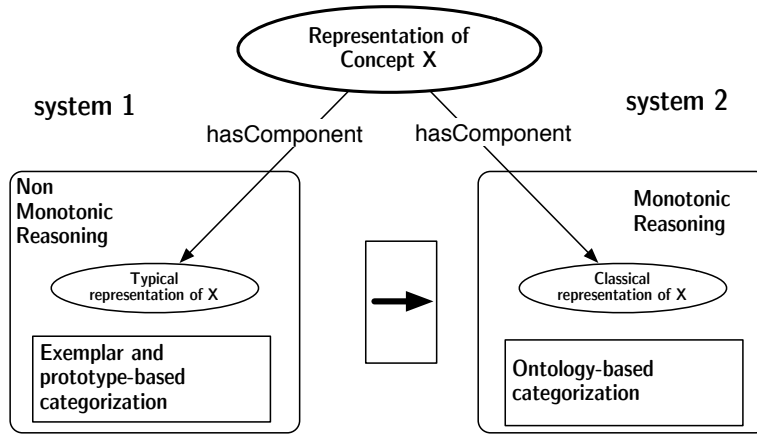


Figure 1. Architecture of the knowledge base.

The proposed architecture (and in particular its prototypical component) can be useful in tasks such as categorization. For example, let us consider the question: In which conditions should we say that someone grasps the concept *water*? In most cases, when we categorize something as *water*, we do not use its classical representation as chemical substance. Rather, we make use of information that *water* is usually liquid, colorless and odorless. The concept *water* is not exceptional from this point of view. Most of everyday or “common sense” concepts behave in the same way. The problem is not different in the case of artificial systems: generally a system that has to categorize something as *water* cannot perform chemical analyses, and it must resort to prototypical evidence. The typicality-based categorization is rooted in the constraints that concern agents with limited access to the relevant knowledge for a given task, and that try to adopt a *heuristic* approach to problem solving [35].²

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 are associated to the Conceptual Spaces framework [15]. 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.¹ Differently from what proposed in [13], 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.

2.1.1. Formalizing conceptual spaces and distance metrics

As mentioned, the aspects related to the typical conceptual component $\mathcal{S}1$ are modeled through the Conceptual Spaces [15]. Conceptual spaces (CS) are an in-

²Therefore, the use of prototypical knowledge in cognitive tasks such as categorization is not a fault of the human mind, as it could be the fact that people are prone to fallacies and reasoning errors (leaving aside the problem of establishing whether recurrent errors in reasoning could have a deeper “rationality” within the general framework of cognition). For the same reason it is also a desired characteristics in the field of intelligent artificial systems.

¹Currently OWL and OWL 2 profiles are not expressive enough to perform the reasoning processes provided by the overall system. However, both language profiles are usable in their DL-safe characterization to exploit taxonomical reasoning. Extending the expressivity of ontological formalisms and languages would be a long-term desideratum in order to enrich the ontological reasoning with more complex inference. To be more expressive and practically usable, a knowledge representation framework should provide an acceptable trade-off in terms of complexity. However, this is an open problem in Fuzzy and Non Monotonic extensions of standard Description Logics.

intermediate level between the symbolic and the sub-symbolic approaches to the knowledge representation relying on geometrical structures, encoded as 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 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 can 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. So, “when natural properties are defined as convex regions of a conceptual space, prototype effects are indeed to be expected” [15, p. 9]. 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 [12, p. 9]).

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 [15, 1]. 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 is the weight associated to the i -th dimension.

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 identifier, the *lemma* of a given concept, information about its typical dimensions, such as color (as the position of the instance on the three-dimensional axes of *brightness*, *hue* and *saturation*) and food.¹ All concepts are mapped onto WordNet synsets: WordNet is a lexical resource whose nodes –the *synsets*– are sets of synonyms, connected through binary relations such as hyponymy/hypernymy and meronymy [28].² Each quality in a domain is associated

¹Typical traits are selected based on statistically relevant information regarding a given concept, as posited by the Prototype Theory [34]. For example, the selection of the information regarding the typical color of a rose (*red*) is given by the fact that roses are often red.

²WordNet information is relevant in our system in that synset identifiers are used by both *S1* and *S2* as a lexical ground to access both the conceptual representations.

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

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

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 [1]), 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 a function of the distance between two regions in a given domain (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 the concepts herein.

2.1.2. Ontology

On the other hand, 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 deductive processes. It is known, in fact, in literature (e.g., by referring to the foundational approach in DOLCE) how to model the fact that “the rose is red”, that is:

- we refer to a given *rose* (*rose#1* in Figure 2);
- it has a certain *color*, expressed via the *inherence* relation, *qtc*: this enables us to specify that *qtc(rose#1)*;
- the particular color of the *rose#1* has a particular *redness* at a certain time t : this is expressed via the *quale*, *ql*, as the relation: *ql(qtc(rose#1, t))*.

However, 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 assume $\#roseP$) nor to describe (and to learn from new cases) the typical




Figure 2. Connecting Concepts to Qualities and Quality Regions in a Foundational Ontology (taken from [27]).

features of the class of prototypical roses. Such aspect has, on the other hand, a natural interpretation by using the conceptual spaces framework.

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, the 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, it is not necessary to specify prototypical values for all individuals in a class: prototypical values can be inherited or overwritten by class individuals, like in structured inheritance networks [4, Chap. 10]. 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 can be summarized as follows (Algorithm 1). 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. In particular, $\mathcal{S}1$ produces in output a list of results, sorted according to a similarity score assessing how similar the input is w.r.t. the available prototypes (Algorithm 1: line 1). All of these elements are then checked by $\mathcal{S}2$ through the cycle at lines 2–14. If the $\mathcal{S}2$ system classifies it as consistent with the ontology, then the classification succeeded and the category provided by $\mathcal{S}2$ (and referred to as 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 subclass of one of those identified by $\mathcal{S}1$ (and referred to as c_i), both cc and c_0 are returned: thus, if $\mathcal{S}2$ provides more specific output, we follow a *specificity* heuristics (Algorithm 1: line 11).

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 even though the $\mathcal{S}1$ output can contain errors, it furnishes approximate answers that cannot be obtained by resorting only to classical ontological inference. On the other hand, the output of

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  $\langle c_0, cc \rangle$ 

```

$\mathcal{S}2$ is always returned with the rationale that it is *safer*,¹ and potentially helpful in correcting the mistakes returned by the $\mathcal{S}1$ process. 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 Thing, the meta class of all the classes in the ontological formalism (Algorithm 1: line 16).

3. Experimentation

We have designed a twofold experimentation on a categorization task, where the system’s results have been studied and compared to state-of-the-art search engines. In the first experiment we tested the system over a restricted domain, and we were mainly interested in assessing the control strategy illustrated in Section 2.2; in the second experiment we tested the system in a broader setting, and we were interested in assessing its robustness and the discriminative features of the $\mathcal{S}1$ component in a multi-domain context.

Two datasets, one composed of 27 “common-sense” linguistic descriptions, and one composed of 56 linguistic descriptions were used, containing two lists of stimuli, including descriptions and their corresponding target categories, such as

$$\underbrace{\text{The big carnivore with yellow and black stripes is the}}_{\text{description } d} \underbrace{\text{... tiger}}_{\text{target } T} .$$

$$\underbrace{\hspace{15em}}_{\text{stimulus } st}$$

The target T is the “prototypically correct” category, and in the following it is referred to as the *expected* result.¹ The set of stimuli was devised by a team of

¹The 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.

¹The *expected* prototypical target category represents a gold standard, since it corresponds to the results provided within a psychological experimentation. In this experimentation 30 subjects were requested to

neuropsychologists and philosophers in the frame of a broader project, aimed at investigating the role of visual load in the 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’) [19].

The $\mathcal{S}1$ - $\mathcal{S}2$ system is part of a larger software pipeline including the extraction of salient information from the input *stimulus*, the access to the hybrid knowledge base, and the retrieval of the corresponding concept (Figure 3). The whole system works as follows: once the linguistic input is given, an internal representation is built by analyzing the stimulus, and by looking for matches between the conceptual spaces and the input. An Information Extraction step was designed that relies on syntactic analysis. The syntactic structure of sentences is computed through the Turin University Parser (TUP) in the dependency format [25]. Dependency formalisms represent syntactic relations by connecting a *dominant* word, the head (e.g., the verb ‘fly’ in the sentence *The eagle flies*) and a *dominated* word, the dependent (e.g., the noun ‘eagle’ in the same sentence). The connection between these two words is represented by using labeled directed edges (e.g., *subject*): the collection of all dependency relations of a sentence forms a tree, rooted in the main verb. In this step, phrases containing nouns and their adjectives are mapped onto the conceptual spaces. Another fundamental step regards the word sense disambiguation [21], necessary to access concepts representation both in the conceptual spaces and in the ontological representation. In the present setting, this Information Extraction step is performed in a supervised fashion.

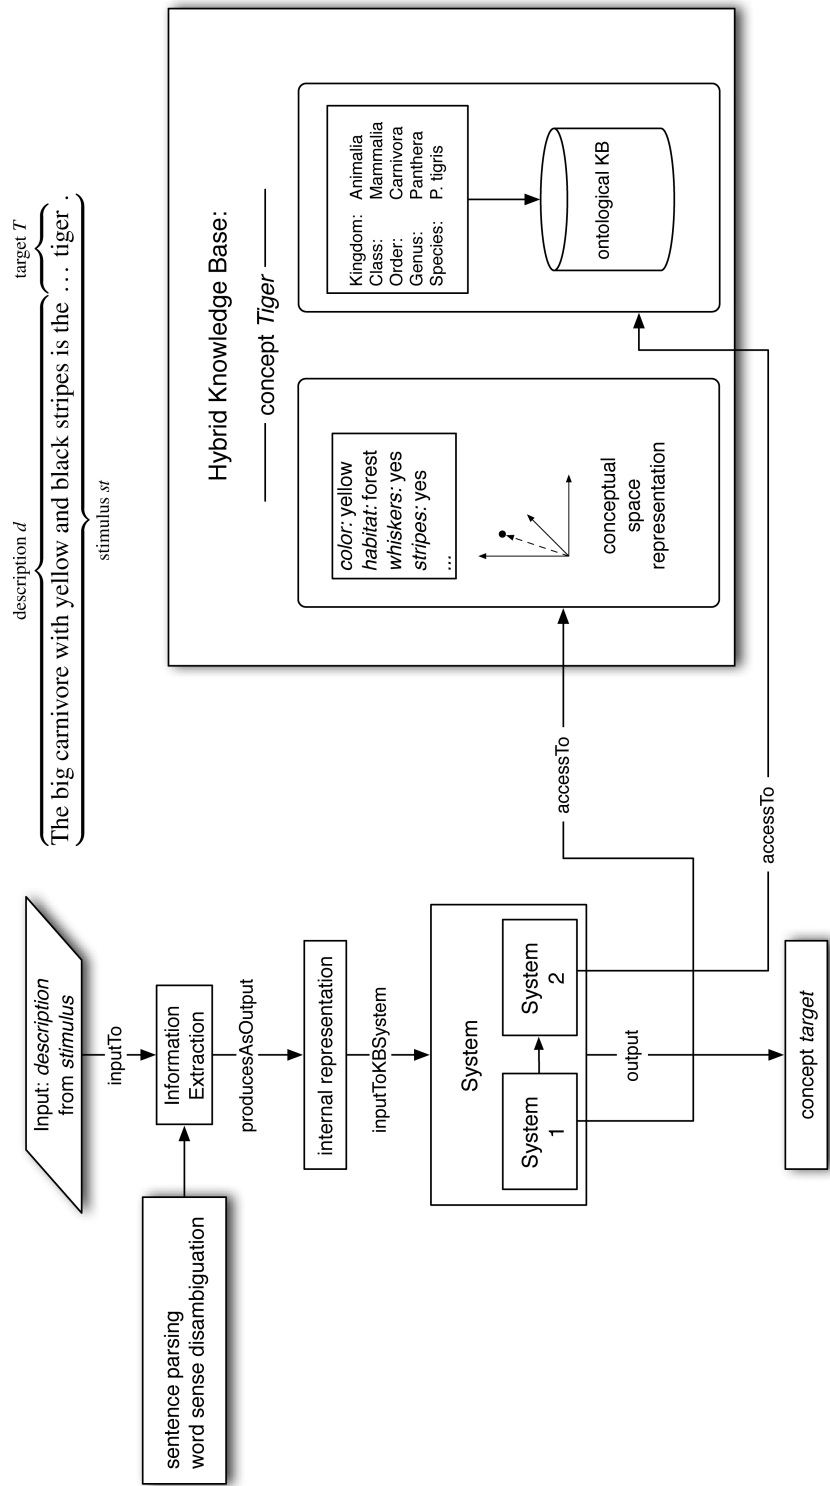


Figure 3. The software pipeline takes in input the linguistic description, queries the hybrid knowledge base and returns the categorized concept.

Experiment 1

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 common-sense descriptions are: “the big carnivore with

black and yellow stripes” denoting the concept of *tiger*, or “the fresh water fish that goes upstream” denoting the concept of *salmon*, and so on.

We devised some metrics to assess both the accuracy of the system, by evaluating it against the *expected* target, and the agreement between $\mathcal{S}1$ and $\mathcal{S}2$. The following information was recorded:

- (1) how often $\mathcal{S}1$ and $\mathcal{S}2$ returned in output the same category. This figure is a measure of the agreement between the two outputs: it scores cases where $\mathcal{S}1$ - $\mathcal{S}2$ output is equal. In this case we do not consider whether the result is the expected category or not;
- (2) the accuracy obtained by $\mathcal{S}1$ alone and by $\mathcal{S}1$ - $\mathcal{S}2$:
 - 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 $\mathcal{S}1$ - $\mathcal{S}2$, that is the overall accuracy of the system also considering, as additional result, the category returned by $\mathcal{S}2$. 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).
- (3) how often Google and Bing —used in a question-answering mode— return pages corresponding to the appropriate concepts, given the same set of definitions and target concepts used to test the proposed system. To these ends, we considered the first 10 results provided by each search engine.¹

In the first experiment 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*.² 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* [36]. 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;³ 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 includes 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*).

The results obtained in the first experimentation are presented in Table 1.

Discussion

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

¹We also tried to extend our evaluation to the well-known semantic question-answering engine Wolfram-Alpha (<https://www.wolframalpha.com>). However, it was not possible to test the descriptions in that it explicitly disregards considering typicality based queries. Namely, the only stimulus correctly categorized is that describing the target cat as “The domestic feline.”

²The ontology is available at the URL http://www.di.unito.it/~radicion/datasets/cs_2014/Naive_animal_ontology.owl

³<http://www.loa-cnr.it/ontologies/DOLCE-Lite.owl>

Table 1. Results of the first experiment. The first column reports the kind of test; the second column reports the number of correctly categorized descriptions, and the third column reports in percentage the same datum as the previous one.

Test cases categorized	descriptions	%
[1.] Cases where $\mathcal{S}1$ and $\mathcal{S}2$ returned the same category	24/27	88.9%
[2a.] Cases where $\mathcal{S}1$ returned the expected category	25/27	92.6%
[2b.] Cases where $\mathcal{S}1$ - $\mathcal{S}2$ returned the expected category	26/27	96.3%
[3.] Google accuracy	19/27	70.4%
[4.] Bing accuracy	18/27	66.7%

concepts (except for one) described with typical features would have been classified through classical ontological inference. It is by 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 \mathcal{C} returned by $\mathcal{S}1$; but because of the categorization performed by $\mathcal{S}2$, 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 expected prototypical answer is *whale*. However, whales properly are mammals, not fishes. In our hybrid system, the $\mathcal{S}1$ 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 $\mathcal{S}1$ is non-monotonic. When then \mathcal{C} (the output of $\mathcal{S}1$) is checked against the ontology, as described by the Algorithm 1 lines 7–13, and an inconsistency is detected,¹ the consistency of the second result in \mathcal{C} (*whale-shark* in this example) is tested against the ontology. Since this answer is an ontologically compliant categorization, then this solution is returned by the $\mathcal{S}2$ component. The final output of the categorization is then the pair $\langle \textit{whale}, \textit{whale-shark} \rangle$: 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 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.

Experiment 2

In order to assess the accuracy of the system in a more demanding experimental setting, we devised a second experimental task, where $\mathcal{S}2$ used the knowledge base OpenCyc.² OpenCyc is one of the largest ontologies publicly available, in that it is a huge attempt at integrating many diverse semantic resources (such as, e.g.,

¹This follows by observing that $c_0 = \textit{whale}$, $cc = \textit{whale-shark}$; and $\textit{whale} \subset \textit{mammal}$, while $\textit{whale-shark} \subset \textit{fish}$; and \textit{mammal} and \textit{fish} are disjoint.

²<http://www.cyc.com/platform/opencyc>.

Table 2. The results of the second experiment.

Test cases categorized	descriptions	%
$\mathcal{S}1$ - $\mathcal{S}2$	44/56	78.57%
Google	37/56	66.07%
Bing	32/56	57.14%

WordNet, DBpedia, Wikicompany, *etc.*). Its coverage and depth were therefore its most attractive features (it contains about 230,000 concepts, 2,090,000 triples and 22,000 predicates). Differently from Experiment 1, we adopted OpenCyc to use a knowledge base independent of our own representational commitments. This was aimed at more effectively assessing the flexibility of the proposed system when using general-purpose, well-known, existing resources, and not only domain-specific ones.

A second dataset of 56 new “common-sense” linguistic descriptions was collected with the same rationale considered for the first experiment.¹

The obtained results are reported in Table 2.

Discussion

While the previous experiment explores the output of both $\mathcal{S}1$ and $\mathcal{S}2$ components, the present one is aimed at assessing it with respect to existing state-of-art search technologies: the main outcome of this experiment is that the trends obtained in the preliminary experiment are confirmed in a broader and more demanding evaluation. Despite being less accurate with respect to the previous experiment, 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 derive from the difficulty of mapping the linguistic structure of stimuli containing very abstract meaning in the representational framework of conceptual spaces. 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*. Similarly, the information extracted from the description “Giving something away for free to someone” could not be mapped onto the features associated to the concept *Gift*. On the other hand, the system shows good performances when dealing with less abstract descriptions based on perceptual features such as *shape*, *color*, *size*, and with some typical information such as *function*.

In this experiment, differently from the previous one (e.g., in the case of *whale*), $\mathcal{S}1$ mostly provided an output coherent with the model in OpenCyc. This datum is of interest, in that although we postulate that the reasoning check performed by $\mathcal{S}2$ is beneficial to ensure a refinement of the categorization process, in this experimentation $\mathcal{S}2$ did not reveal any improvement to the output provided by $\mathcal{S}1$, also when this output was not in accord with the *expected* results. In fact, by analyzing in detail the different answers, we notice that at least one inconsistency should have been detected by $\mathcal{S}2$. This is the case of the description “An intelligent grey fish” associated to the target concept *Dolphin*. In this case, the $\mathcal{S}1$ system returned the expected target, but $\mathcal{S}2$ did not raise the inconsistency since OpenCyc

¹The full list of the second set of stimuli, containing the expected “prototypically correct” category is available at the following URL: http://www.di.unito.it/~radicion/datasets/cs_2014/stimuli.txt.

erroneously represents *Dolphin* as a subclass of *Fish*, rather than a subclass of *Mammal*. Therefore, one of the weaknesses of the overall architecture can be due to ontological misclassifications in the *S2* component.

4. Related work

The presented solution has some analogies with the approach that considers concepts as “semantic pointers” [8, 39], proposed in the field of the computational modeling of brain. In such approach, different informational components are supposed to be attached to a unifying concept identifier. The similarity with their approach is limited to the idea that concepts consist of different types of information. However, the mentioned authors specifically focus on the different modalities of the stimuli contributing to conceptual knowledge, and therefore they identify the different components of concepts according to the different information carriers used to provide the information. Their conceptual components are divided in: sensory, motor, emotional and verbal stimuli, and for each type of carriers a mapping function to a brain area is supposed to be activated. On the other side, our focus is on the type of conceptual information (e.g., classical *vs.* typical information): we do not consider the modality associated to the various sources of information (e.g., visual or verbal, etc.).¹ Rather, we are concerned with the type of information combined in the hybrid conceptual architecture embedded in our *S1-S2* computational system.

In the context of a different field of application, a solution similar to the one adopted here has been proposed in [7]. 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 the 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 [23], developed with the aim of providing a computational architecture of reasoning based on the mental models theory proposed by Philip Johnson-Laird [20]. 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 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 searches for alternative models, if the initial conclusion does not hold or if it is not satisfactory. A second system has been proposed by Larue *et al.*, [24]. The authors adopt an extended version of the dual process approach, which has been described in [37]; it is based on the hypothesis that the system 2 is divided in two further levels, respectively called “algorithmic” and “reflective”. The goal of Laure and colleagues is to build

¹In our view the distinction classical *vs.* prototypical is ‘a-modal’ *per se*, for example both a typical and a classical conceptual information can be accessed and processed through different modalities (that is visual *vs.* auditory, etc.).

a multi-agent and multilevel architecture, to represent the emergence of emotions in a biologically inspired computational environment.

Another system that is close to our present work has been proposed by [32]. 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 akin to the dual process proposal: for example, many hybrid, symbolic-connectionist systems –including cognitive architectures such as, for example, CLARION¹–, 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 [14]). However, to 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 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. 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 system has been tested in a twofold experimentation consisting of a categorization task involving typicality based queries. In the former case we used an *ad-hoc* developed domain ontology; whilst in the latter case we used OpenCyc, a publicly available, huge knowledge base. The obtained results show that in the restricted domain of the animal kingdom the proposed architecture is by far more accurate than general purpose systems such as Google and Bing. This advantage is reduced when dealing with an unrestricted context, but the overall results corroborate the hypothesis that matching in conceptual spaces can be fruitfully paired with ontological inference. However, the ontological inference would not suffice to categorize the common sense stimuli presented in both experiments, with only one exception.

The $\mathcal{S}1$ - $\mathcal{S}2$ system can be applied in a task close to that considered in present experimentation: besides informational queries, the system can be employed in the analysis of search engines web logs to investigate whether and to what extent their results match the actual users' informational needs. A meta-search model can be envisioned, where web engines results are filtered by the system $\mathcal{S}1$ - $\mathcal{S}2$ before being returned to the user, so as to retain only relevant (to the ends of informational queries) results.

In the next future we will complete the automatization of the Information Extraction step from linguistic descriptions, which are currently (in a supervised fashion) mapped onto an internal representation shared by both $\mathcal{S}1$ and $\mathcal{S}2$. In particular, we

¹<http://www.cogsci.rpi.edu/~rsun/clarion.html>

are currently investigating how to exploit information provided by semi-structured sources such as ConceptNet² to compute the values required to fill the conceptual space dimensions. Furthermore, we plan to test the proposed approach in the area of biomedical domain in order to assess disease diagnosis tasks by using in $\mathcal{S}2$ an ontology such as SNOWMED,³ and in $\mathcal{S}1$ conceptual spaces representing the typical symptoms of a given disease. Moreover, the conceptual spaces information can be globally seen as an effort to provide some portions of the WordNet hierarchy with a conceptual-spaces annotation. For the next future, we plan to release the conceptual spaces resource in open format.

Finally, the obtained results show that this approach may be beneficial to a plethora of NLP tasks where a wide-coverage is required, such as summarization (also in a multilingual setting), question answering and information retrieval.

Acknowledgment

This work has been partly supported by the Ateneo-San Paolo project number TO_call03_2012_0046, *The role of visual imagery in lexical processing* (RVILP). The first author's work is also partially supported by the CNR *F.A.C.I.L.E.* project ICT.P08.003.001.

The authors kindly thank Leo Ghignone, for working to an earlier version of the system; Marcello Frixione, for discussions and advices on the theoretical aspects of this approach; the anonymous reviewers, whose valuable suggestions were helpful to improve the work; Manuela Sanguinetti, for her comments on a previous version of the article. We also thank the attendees of the *ConChaMo 4 Workshop*,¹ organized by the University of Helsinki, and the participants of the *Spatial Colloquium Workshop* organized by the Spatial Cognition Center of the University of Bremen² for their comments and insights to initial versions of this work: in particular, we thank David Danks, Christian Freksa, Peter Gärdenfors, Ismo Koponen, and Paul Thagard.

We especially thank Leonardo Lesmo, beloved friend and colleague no longer with us, who strongly encouraged the present line of research.

References

- [1] Benjamin Adams and Martin Raubal. A metric conceptual space algebra. In Kathleen Stewart Hornsby, Christophe Claramunt, Michel Denis, and Gérard Ligozat, editors, *COSIT*, volume 5756 of *Lecture Notes in Computer Science*, pages 51–68. Springer, 2009.
- [2] Piero A Bonatti, Carsten Lutz, and Frank Wolter. Description logics with circumscription. In *KR*, pages 400–410, 2006.
- [3] Ronald J. Brachman and Hector J. Levesque. *Readings in Knowledge Representation*. Morgan Kaufmann Pub, 1985.
- [4] Ronald J. Brachman and Hector J. Levesque. *Knowledge Representation and Reasoning*. Elsevier, 2004.
- [5] Ronald J. Brachmann and James G. Schmolze. An overview of the KL-ONE knowledge representation system. *Cognitive Science*, 9(2):171–202, April 1985.
- [6] Silvia Calegari and Davide Ciucci. Fuzzy ontology, fuzzy description logics and fuzzy-owl. In *Applications of Fuzzy Sets Theory*, pages 118–126. Springer, 2007.
- [7] A. Chella, M. Frixione, and S. Gaglio. A cognitive architecture for artificial vision. *Artificial Intelligence*, 89(1–2):73–111, 1997.
- [8] Chris Eliasmith, Terrence C Stewart, Xuan Choo, Trevor Bekolay, Travis DeWolf, Yichuan Tang, and Daniel Rasmussen. A large-scale model of the functioning brain. *Science*, 338(6111):1202–1205, 2012.

²<http://conceptnet5.media.mit.edu>.

³<http://www.b2international.com/portal/snow-owl>.

¹<http://conceptualchange.it.helsinki.fi/workshops/conchamo4/>.

²<http://www.sfbtr8.uni-bremen.de/en/home/>.

- [9] Jonathan St. B.T. Evans and Keith Frankish, editors. *In two minds: Dual processes and beyond*. Oxford University Press, 2009.
- [10] Marcello Frixione and Antonio Lieto. The computational representation of concepts in formal ontologies-some general considerations. In *KEOD*, pages 396–403, 2010.
- [11] Marcello Frixione and Antonio Lieto. Representing concepts in formal ontologies: Compositionality vs. typicality effects. *Logic and Logical Philosophy*, 21(4):391–414, 2012.
- [12] Marcello Frixione and Antonio Lieto. 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*, pages 171–182, 2013.
- [13] Marcello Frixione and Antonio Lieto. 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.
- [14] Marcello Frixione and Antonio Lieto. Formal Ontologies and Semantic Technologies: A Dual Process Proposal for Concept Representation. *Philosophia Scientiae*, forthcoming.
- [15] Peter Gärdenfors. *Conceptual spaces: The geometry of thought*. MIT press, 2000.
- [16] Peter Gärdenfors. *The Geometry of Meaning: Semantics Based on Conceptual Spaces*. MIT Press, 2014.
- [17] Laura Giordano, Valentina Gliozzi, Nicola Olivetti, and Gian Luca Pozzato. A non-monotonic description logic for reasoning about typicality. *Artificial Intelligence*, 195:165–202, 2013.
- [18] Tom Gruber. Ontology. *Encyclopedia of database systems*, pages 1963–1965, 2009.
- [19] Bernard J. Jansen, Danielle L. Booth, and Amanda Spink. Determining the informational, navigational, and transactional intent of web queries. *Information Processing & Management*, 44(3):1251–1266, 2008.
- [20] Philip N. Johnson-Laird. Mental models in cognitive science. *Cognitive Science*, 4(1):71–115, 1980.
- [21] Daniel Jurafsky and James H. Martin. *Speech and Language Processing: an Introduction to Natural Language Processing, Computational Linguistics and Speech*. Prentice Hall, 2000.
- [22] Daniel Kahneman. *Thinking, fast and slow*. Macmillan, 2011.
- [23] Sangeet Khemlani and PN Johnson-Laird. The processes of inference. *Argument & Computation*, 4(1):4–20, 2013.
- [24] Othalia Larue, Pierre Poirier, and Roger Nkambou. A cognitive architecture based on cognitive/neurological dual-system theories. In *Brain Informatics*, pages 288–299. Springer, Berlin, 2012.
- [25] Leonardo Lesmo. The Rule-Based Parser of the NLP Group of the University of Torino. *Intelligenza Artificiale*, 2(4):46–47, June 2007.
- [26] Edouard Machery. *Doing without concepts*. Oxford University Press Oxford, 2009.
- [27] Claudio Masolo, Stefano Borgo, Aldo Gangemi, Nicola Guarino, and Alessandro Oltramari. WonderWeb deliverable D18 ontology library (final). Technical report, IST Project 2001-33052 WonderWeb: Ontology Infrastructure for the Semantic Web, 2003.
- [28] George A Miller. WordNet: a lexical database for English. *Communications of the ACM*, 38(11):39–41, 1995.
- [29] Marvin Minsky. A framework for representing knowledge. In P. Winston, editor, *The Psychology of Computer Vision*, pages 211–277. McGraw-Hill, New York, 1975.
- [30] Gregory Leo Murphy. *The big book of concepts*. MIT press, Cambridge, Massachusetts, 2002.
- [31] Daniele Nardi and Ronald J. Brachman. An introduction to description logics. In *Description logic handbook*, pages 1–40, 2003.
- [32] Giovanni Pilato, Agnese Augello, and Salvatore Gaglio. A modular system oriented to the design of versatile knowledge bases for chatbots. *ISRN Artificial Intelligence*, 2012, 2012.
- [33] Ross Quillian. *Semantic Memory*, pages 216–270. MIT Press, 1968.
- [34] Eleanor Rosch. Cognitive representations of semantic categories. *Journal of experimental psychology: General*, 104(3):192–233, 1975.
- [35] Herbert A Simon and Allen Newell. Human problem solving: The state of the theory in 1970. *American Psychologist*, 26(2):145, 1971.
- [36] Eliot R. Smith and Nyla R. Branscombe. Category accessibility as implicit memory. *Journal of Experimental Social Psychology*, 24(6):490–504, 1988.
- [37] Keith E Stanovich, Richard F West, et al. Individual differences in reasoning: Implications for the rationality debate? *Behavioral and brain sciences*, 23(5):645–665, 2000.
- [38] Umberto Straccia. Reasoning within fuzzy description logics. *arXiv preprint arXiv:1106.0667*, 2011.
- [39] Paul Thagard and Scott Findlay. *The cognitive science of science: Explanation, discovery, and conceptual change*. MIT Press, 2012.
- [40] Ludwig Wittgenstein. *Philosophische untersuchungen-Philosophical investigations*. B. Blackwell, 1953.