Journal of Universal Computer Science, vol. 20, no. 3 (2014), 257-276
submitted: 24/3/13, accepted: 4/2/14, appeared: 1/3/14 © *J.UCS*

Towards an Extended Model of Conceptual Representations in Formal Ontologies: A Typicality-Based Proposal

Marcello Frixione
(University of Genoa, Genova, Italy
marcello.frixione@gmail.com)

Antonio Lieto
(University of Turin, Torino, Italy
ICAR – C.N.R., Palermo, Italy
lieto@di.unito.it)

Abstract: In this paper we propose a possible solution for the problem of the computational representation of non-classical concepts (i.e. concepts that cannot be characterized in terms of necessary and sufficient conditions) in the field of formal ontologies. In particular, taking into account empirical evidences coming from cognitive psychology, according to which concept representation is not a unitary phenomenon, we suggest that a similar approach to the representation of conceptual knowledge could be useful also in the field of ontology based technologies. Finally we propose, in a linked open data perspective, conceptual spaces as a suitable framework for developing some aspects of the presented proposal.

Keywords: Concept Representation, Formal Ontologies, Concepts Theories, Conceptual Spaces, Prototypes, Hybrid Conceptual Representation
Categories: I.2.0, I.2.4

1 Introduction

This article deals with the problem of representing non-classical concepts, with particular attention to artificial systems, such as formal ontologies[1]. According to our approach, concept representation in artificial systems can take great advantage from the theoretical and empirical results of cognitive sciences. By non-classical concepts we mean concepts that cannot be represented in terms of sets of necessary and/or sufficient conditions (as maintained by the so called "classical" theory of the concepts). After introducing the problem (sect. 2), we review some empirical evidence from cognitive psychology (sect. 3). In section 4 we individuate some possible suggestions coming from different aspects of cognitive research: the distinction between two different types of reasoning processes developed within the context of the so-called "dual process" accounts of reasoning; the proposal to keep typicality effects separate from classical representation of concepts; and, the

[1] According to [Gruber, 1993] formal ontologies are "a specification of a conceptualization". Therefore they allow to explicitly and precisely represent certain conceptualizations usually by mean of formal languages.

possibility of developing hybrid conceptual representations by combining, respectively, classical and typical (or non-classical) components and prototype and exemplar-based representations. In particular, prototype and exemplar based models of non-classical concepts are both plausible, and can account for different aspects of human abilities (sect. 5). We argue that these results could suggest the adoption of a hybrid approach also in the fields of knowledge representation and formal ontologies. In section 7 we present the proposal of a possible architecture for concept representation based on the integration of a classical component (e.g. based on Description Logics) with a typical one (e.g. prototypes and exemplars based) connected via rules (the role of rules as "connectors" of the conceptual components is anticipated in sect. 6). Then we introduce conceptual spaces (sect. 8) as a suitable framework for the development of some aspects of our proposal (sect. 9). Some conclusions (sect. 10) follow.

2 Representing Non Classical Concepts

The representation of common sense concepts is still an open problem in ontology engineering and, generally, in Knowledge Representation (KR) [Frixione and Lieto, 2012a]. Cognitive Science showed the empirical inadequacy of the so-called "classical" theory of concepts, according to which concepts should be defined in terms of sets of necessary and sufficient conditions. Rather, Eleanor Rosch's experiments [Rosch, 1975] showed that ordinary concepts can be characterized in terms of prototypical information.

These results influenced the early researchers in knowledge representation: the KR practitioners initially tried to keep into account the suggestions coming from cognitive psychology, and designed artificial systems – such as frames [Minsky, 1975] and early semantic networks [Quillian, 1968] – able to represent concepts in "non-classical" (typicality-based) terms (for early KR developments, see also the papers collected in Brachman and Levesque, [Brachman and Levesque, 1985]).

However, these early systems lacked a clear formal semantics and a satisfactory meta-theoretic account, and were later sacrificed in favor of a class of formalisms stemmed from the so-called structured inheritance semantic networks and the KL-ONE system [7]. These formalisms are known today as *description logics* (DLs) [Baader et al. 2010] and do not allow for exceptions to inheritance, and for the possibility to represent concepts in typical terms. From this point of view, therefore, such formalisms can be seen as a revival of the classical theory of concepts. As far as typical information is concerned, such formalisms offer only two possibilities: representing it by resorting to tricks or ad hoc solutions[2], or, alternatively, ignoring it.

[2] A typical example of such a misuse of DLs is the inclusion, in the representation of concepts, of information that is not necessary nor sufficient for the definition of the corresponding category – e.g., the information that DOGs have exactly four legs. In fact, despite the fact that the representation of concepts in ontologies is dependent on the purpose of the ontology itself, in the case of dog, the intended interpretation of such representation is that ∀(x) if DOG(x) then hasFourLegs(x). This leads to the need of recurring to ad-hoc solutions for the representation of atypical dogs. An example is that one of specializing different classes of DOGs based on the number of their legs (e.g. the ThreeLeggedDogs, the EightLeggedDogs and so on).

For obvious reasons, the first solution cannot be take into account: it would have disastrous consequences for the soundness of the knowledge base and for the performances of the entire system. The second choice severely reduces the expressive power of the representation.

Nowadays, DLs are widely adopted within many fields of application, in particular within the area of ontology representation. For example, OWL (Ontology Web Language, see [OWL] is a formalism in this tradition, which has been endorsed by the World Wide Web Consortium for the development of the Semantic Web. However, DL formalisms leave unsolved the problems of representing concepts in typical terms.

Within the field of logic oriented knowledge representation, rigorous approaches have been proposed in order to allow the representation of exceptions. Such approaches are, therefore, at least in principle, suitable for dealing with (some aspects of) "non-classical" concepts. Examples are fuzzy and non-monotonic extensions of DL formalisms. Nevertheless, such approaches pose various theoretical and practical problems, which in part remain unsolved (for a more detailed account on this aspect, see [Frixione and Lieto 2010]).

As a possible way out, we outline here a tentative proposal that goes in a different direction and that is based on some suggestions coming from empirical cognitive science research.

3 The Advantage of a Cognitive Approach to Typicality in Artificial Systems

As anticipated before, typicality effects in categorisation and, in general, in category representation, are not only crucial for the empirical study of human cognition. They are also of the greatest importance in representing concepts in artificial systems. Let us first consider human cognition. Under what conditions should we say that somebody *knows* the concept DOG (or, in other terms, that she possesses an adequate mental representation of it)? It is not easy to say. However, if a person does not know that, for example, dogs usually bark, that they typically have four legs and that their body is covered with fur, that in most cases they have a tail and that they wag it when they are happy, then we probably should conclude that this person does not grasp the concept DOG. Nevertheless, all these pieces of information are neither necessary nor sufficient conditions for being a dog. In fact, they are traits that characterise dogs in typical (or prototypical) cases. The problem is exactly the same if we want to represent knowledge in an artificial system. Let us suppose that we want to provide a computer program with a satisfactory representation of DOG. Then we probably also want to represent the kind of information mentioned above: for many applications, a representation of DOG that does not include the information that dogs usually bark is a bad representation also from a technological point of view. Therefore, if a system does not allow to represent information in typical terms (as is the case of standard description logics), then it is not adequate in this respect. With standard DLs, the only way to face this problem should be the recourse to tricks or ad hoc solutions (as often happens in many applications). The concept DOG is not exceptional from this point of view. The majority of everyday concepts behave in this way. For most concepts, a

classical definition in terms of necessary and sufficient conditions is not available (or, even if it is available, it is unknown to the agent). On the other hand, it can happen that we know the classical definition of a concept, but typical knowledge still plays a central role in many cognitive tasks. Therefore, the use of typical knowledge in cognitive tasks such as categorization can be seen as a powerful heuristic approach to problem solving that concern every finite agent that has a limited access to the relevant knowledge for a given task. This is the case of both natural and artificial cognitive systems.

4 Some Suggestions from Cognitive Science

Some recent trends of psychological research favour the hypothesis that reasoning is not a unitary cognitive phenomenon. At the same time, empirical data on concepts seem to suggest that prototypical effects could stem from different representation mechanisms. In this spirit, we individuate some hints that, in our opinion, could be useful for the development of artificial representation systems, namely: (i) the distinction between two different types of reasoning processes, which has been developed within the context of the so-called "dual process" accounts of reasoning [Evans and Frankish, 2008] (sect. 4.1 below); (ii) the proposal to keep typical effects separate from classical representation of concepts (sect. 4.2); and (iii) the possibility to develop hybrid representations of concepts coupling both classical and typical representations and, within the latter, 'prototype- and exemplar-based' models (sect. 4.3).

4.1 A Dual Process Approach

Cognitive research about concepts [Murphy, 2002; Machery, 2009] seems to suggest that concept representation does not constitute a unitary phenomenon from the cognitive point of view. In this perspective, a possible solution should be inspired by the experimental results of empirical psychology, in particular by the so-called dual process theories of reasoning and rationality [Evans and Frankish 2008]. In such theories, the existence of two different types of cognitive systems is assumed. The systems of the first type (type 1) are phylogenetically older, unconscious, automatic, associative, parallel and fast. The systems of the type 2 are more recent, conscious, sequential and slow, and are based on explicit rule following. In our opinion, there are good *prima facie* reasons to believe that, in human subjects, many monotonic forms of reasoning which are defined on semantic networks, and which are typical of DL systems, are likely to be type 2 tasks (they are difficult, slow, sequential). On the contrary, exceptions play an important role in processes such as categorization and inheritance, which are more likely to be tasks of the type 1: they are fast, automatic, usually do not require particular conscious effort, and so on.

Therefore, a reasonable hypothesis is that a concept representation system should include different "modules": a monotonic module of type 2, involved in "difficult" logical tasks, and a non-monotonic module involved in categorization, which takes advantage from the management of exceptions. This last module should be a "weak" system, able to perform only some simple forms of non-monotonic inferences (mainly related to categorization and to exceptions inheritance).

4.2 A Pseudo-Fodorian Proposal

According to Fodor [Fodor, 1987], concepts cannot be prototypical representations, since concepts must be compositional, and prototypes do not compose. On the other hand, in virtue of the criticisms to "classical" theory, concepts cannot be definitions. Therefore, Fodor argues that (most) concepts are atoms, i.e., are symbols with no internal structure. Their content is determined by their relation to the world, and not by their internal structure and/or by their relations with other concepts. Of course, Fodor acknowledges the existence of prototypical effects. However, he claims that prototypical representations are not part of concepts.

We borrow from Fodor the hypothesis that compositional representations and prototypical effects are demanded to different components of the representational architecture. We assume that there is a compositional component of representations, which does not admit exceptions and exhibits no prototypical effects, and which can be represented, for example, in the terms of some classical DL knowledge base. In addition, a typical representation of categories is responsible for such processes as 'prototype- and exemplar-based' categorisation, but it does not affect the inferential behaviour of the compositional component.

4.3 Multiple Concept Theories

Within the field of cognitive psychology, different positions and theories on the nature of concepts are available. Usually, they are grouped in three main classes, namely: prototype views, exemplar views and theory-theories (see e.g. [Murphy, 2002] and [Machery, 2009]). All of them are assumed to account for (some aspects of) typicality effects in conceptualization and conceptual reasoning.

According to the prototype view, the knowledge about categories is stored in terms of prototypes, where a prototype is a representation of the "best" instance of a category. For example, the mental representation of the concept CAT should coincide with a representation of a prototypical cat, where a prototypical cat is a cat whose body is covered with fur that has four legs and retractile claws, that meows and purrs, and so on. In the simpler versions of this approach, prototypes are represented as (possibly weighted) lists of features, such *has fur*, *has retractile claws*, *meows*. The weights are numeric values that express the relevance of each feature.

According to the exemplar view, categories are not mentally represented as specific, local structures such as prototypes. Rather, a category is represented as a set of specific exemplars explicitly stored within memory. For example, the mental representation of the concept CAT is the set of the representations of (some of) the cats we encountered during our lifetime.

Theory-theory approaches adopt some form of holistic point of view about concepts. According to some versions of the theory-theories, concepts are analogous to theoretical terms in a scientific theory. For example, the concept CAT is individuated by the role it plays in our mental theory of zoology. In other versions of the approach, concepts themselves are identified with micro-theories of some sort. For example, the concept CAT should be identified with a mentally represented micro-theory about cats.

These approaches are not mutually exclusive. Rather, they seem to succeed in explaining different classes of cognitive phenomena, and many researchers hold that

all of them are needed to explain psychological data (see again [Murphy, 2002] and [Machery, 2009]). In this perspective, we propose to integrate some of them in computational representations of concepts. More precisely, we mainly focus on prototypical and exemplar based approaches, and propose to combine them in a hybrid representation architecture in order to account for category representation and typicality effects. Furthermore we argue that the theory-theory approaches can be seen as a possible general framework in which different types of conceptual representations, including prototypes and exemplars, are connected via rules.

5 Prototypes and Exemplars

According to the available experimental evidence, exemplar models are in many cases more successful than prototypes (for a more detailed review of these results, see [Frixione and Lieto 2012b]). It can happen for example that a less typical item is categorized more quickly and more accurately than a more typical category member if it is similar to previously encountered exemplars of the category [Medin and Schaffer, 1978]. Let us consider the following example: a penguin is a rather atypical bird. However, let us suppose that some exemplar of penguin is already stored in my memory as an instance of the concept BIRD. In this case, it can happen that I classify new penguins as birds more quickly and more confidently than less atypical birds (such as, say, toucans or hummingbirds) that I never encountered before.

Another important source of evidence for the exemplar model stems from the study of linear separable categories (see, again, [Medin and Schwanenflugel, 1981]). Two categories are linearly separable if and only if it is possible to determine to which of them an item belongs by summing the evidence concerning each attribute of this item. For example, let us suppose that two categories are characterized by two attributes, or dimensions, corresponding to the axes in figure 1. These categories are linearly separable if and only if the category membership of each item can be determined by summing its value along the *x* and *y* axes, or, in other terms, if a line can be drawn, which separates the members of the categories.

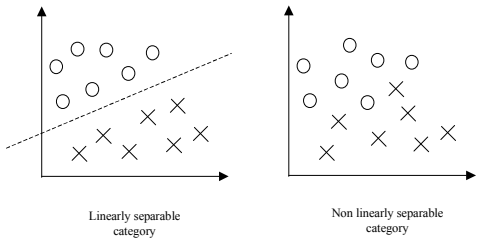

Figure 1: Linearly separable and non separable categories

According to the prototype approach, people should find it more difficult to form a concept of a non-linearly separable category. Subjects should be faster at learning two categories that are linearly separable. However, Medin and Schwanenflugel [Medin and Schwanenflugel, 1981] experimentally proved that categories that are not

linearly separable are not necessarily harder to learn. This is not a problem for exemplar based theories, which do not predict that subjects would be better at learning linearly separable categories. In the psychological literature, this result has been considered as a strong piece of evidence in favor of the exemplar models of concept learning.

The above mentioned results seem to favor exemplars against prototypes. However, other data does not confirm this conclusion. An empirical research supporting the hypothesis of a multiple mental representation of categories is in Malt [Malt, 1989]. This study was aimed to establish if people categorize and learn categories using exemplars or prototypes. The empirical data, consisting in behavioural measures such as categorization probability and reaction time, suggest that subjects use different strategies to categorize. Some use exemplars, a few rely on prototypes, and others appeal to both exemplars and prototypes.

Summing up, prototype and exemplar approaches present significant differences, and have different merits. Therefore, it is likely, in our opinion, that a dual, 'prototype- and exemplar-based', representation of concepts could turn out to be useful also from a technological point of view, for the representation of non-classical concepts in ontological knowledge bases.

In the first place, there are kinds of concepts that seem to be more suited to be represented in terms of exemplars, and concepts that seem to be more suited to be represented in terms of prototypes. For example, in the case of concepts with a small number of instances, which are very different from one another, a representation in terms of exemplars should be more convenient. An exemplar based representation could be more suitable also for non linearly separable concepts (see above).

On the other hand, for concepts with a large number of very similar instances, a representation based on prototypes seems to be more appropriate. Consider for example an artificial system that deals with apples (for example a fruit picking robot, or a system for the management of a fruit and vegetable market). Since it is unlikely that a definition based on necessary/sufficient conditions is available or adequate for the concept APPLE, then the system must incorporate some form of representation that exhibits typicality effects. But probably an exemplar based representation is not convenient in this case: the system has to deal with thousands of apples that are all very similar one another. A prototype would be a much more natural solution.

In many cases, the presence of both a prototype and an exemplar based representation seems to be appropriate.

6 Theory-theory

Theory-theory approaches [Murphy and Medin, 1985] assume that concepts consists of more or less complex mental structures representing (among other things) causal and explanatory relations. During the 80's, these approaches stem from a critique to the formerly dominant theory of concepts as prototypes.

In our view, the theory-theory approach to concepts is, at the current state of the art, partially covered within the field of knowledge representation and formal ontologies. Classical ontological languages (e.g. [OWL]), in combination with some ontology-compliant rule language allow the expression of ontological concepts as part of a theory. The SWRL (Semantic Web Rule Language, [SWRL]) is one of the most

known ontology languages. It integrates OWL with a rule layer built on top of it [Eiter et al. 2008]. SWRL's goal of enhancing description logics with rules is aimed at overcoming some well known expressive limitations in ontology languages, which can be easily fixed by adding rules to an ontological knowledge base. More specifically, SWRL adds the possibility to declare arbitrary Horn clauses expressed as IF … THEN … rules. A SWRL ontology is therefore composed of ordinary OWL axioms plus SWRL rules. The antecedents and consequents of the rules consist of lists of atoms, which may be OWL class expressions, property definitions, or built-ins.

Most current DL reasoners support inferences based on SWRL. However, SWRL is restricted to monotonic inference and only certain types of theories (namely causal theories) seem to be covered by the integration of the current state of the art ontology languages and rules. Furthermore, as already pointed out before, common sense knowledge is mostly characterized in terms of "theories" which are based on arbitrary, i.e. experience-based, rules. Therefore, in order to represent, within ontology based systems, more realistic "theories", as intended in the theory-theory approaches, there is the need of going beyond classical logic rules. A possible solution, in our opinion, could be achieved by partially reconsidering the "conceptual environment" to which traditionally these rules have been applied. More specifically we claim that there is the need of a renewed, hybrid, conceptual structure to which rule based mechanism can be applied. Within this new environment, the role of the rules, expressed through different formalisms, could be that of providing both causal and associative connections between concepts and/or between different conceptual components referred to the same concept. In the next section a proposal for the realization of such a hybrid structure is proposed.

7 An Extended, Typicality-Based, Model for Concept Representation

In this section we outline the proposal of a possible architecture for concept representation, which takes advantage of the suggestions presented in the sections above. It is based on a hybrid approach, and combines a component based on a Description Logic (DL) with a further component that implements typical representations including both prototypes and exemplars models of non-classical concepts. In this sense, our approach can be seen as "hybrid" at different levels: the first level of hybridization is based on the above mentioned distinction between classical (or compositional) and non-classical (or typical) conceptual components. The second level of hybridization is, on the other hand, internal to the typical component and is aimed at representing the different types of typicality-based inferential processes coming from the representation of non-classical concepts in terms of both prototypes and exemplars.

Our solution has some analogies with the approach considering concepts as "semantic pointers" (Eliasmith et al. 2012, Thagard 2012) proposed in the field of the computational modelling of brain. In such approach, different informational components are supposed to be attached to a unifying concept identifier. The similarity with this approach is limited to the idea that concepts consist of different types of information which, combined together, provide different ways of accessing

the conceptual knowledge. However, while Eliasmith and colleagues and Thagard concentrate specifically on the different modalities of the stimuli contributing to conceptual knowledge, and, therefore, identify the different components of the concepts according to the different information carriers through which the content of the information is provided (e.g. 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) our focus is on the content of conceptual information itself (e.g. classical vs. typical information). More precisely we do not take into account from *where* the different types of information come from (e.g. visual or verbal stimuli). Rather, our focus is on *what* type of information is combined in a hybrid conceptual architecture.

Concepts in the DL component are represented as in figure 2. As usual, every concept can be subsumed by a certain number of superconcepts, and it can be characterized by means of a number of attributes, which relate it to other concepts in the knowledge base. Restrictions on the number of possible fillers can be associated to each attribute. Given a concept, its attributes and its concept/superconcept relations express necessary conditions for it. DL formalisms make it possible to specify which of these necessary conditions also count as sufficient conditions.

Since in this component only necessary/sufficient condition can be expressed, here concepts can be represented only in classical terms: no exceptions and no prototypical effects are allowed. Concepts can have any number of individual instances, that are represented as individual concepts in the taxonomy.

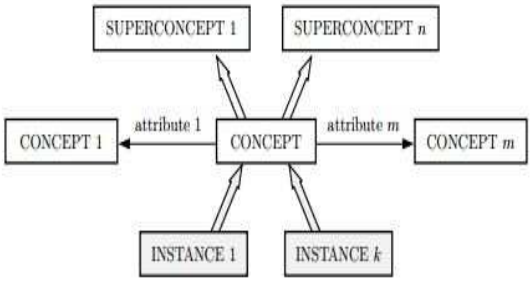

Figure 2: A concept in the DL

For example: the concept DOG is represented as a subconcept of MAMMAL. Since DL networks can express only necessary and/or sufficient conditions, some details of the representation are very loose. For example, a DOG may or may not have a tail (this is the expressed by the number restriction 0/1 imposed on the attribute *has_tail*), and has an unspecified number of limbs (since some dogs could have lost limbs, and teratological dogs could have more than four legs). LASSIE and RIN TIN TIN can be represented as individual instances of DOG (of course, concepts describing individual instances can be further detailed, fully specifying for example the values of the attributes inherited from parent concepts).

Prototypes describing typical instances of concepts can be represented as data structures that are external to the DL knowledge base. Such structures could, for example, be lists of (possibly weighted) attribute/value pairs that are linked to the

corresponding concept. Some attributes of the list should correspond to attributes of the DL concept, the values of which can be further specified at this level. For example, the prototypical dog is described as having a tail and exactly four legs. Other attributes of the prototype could have no counterpart in the corresponding DL concept.

As far as the exemplar-based component of the representations is concerned, exemplars are directly represented in the DL knowledge base as instances of concepts, while some typical information concerning an exemplar can be demanded, as for the prototypes, to an external component.

In the field of web ontology languages, the development of the architecture sketched above is nowadays, technologically easier to implement. Within the Semantic Web research community, in fact, the Linked Data perspective is assuming a prominent position [Bizer et al. 2009]. According to this view the main goals of the Semantic Web community is the integration of different data representations (often stored in different data sources) within a unique, semantically linked, representational framework. The main technical result coming from this integration is represented by the possibility of enlarging the answer-space of a query through the realization of "semantic bridges" between different pieces of data (and, often, data sources). Such integration is made possible through constructs provided by Semantic Web languages, such as OWL (e.g. the *owl:sameAs* construct), or schemas such as SKOS[3]. It must be noted that, in this setting, typical information about concepts (either stored in the form of prototypes or extracted from the representation of exemplars) extends the information coded within the DL formalism: the semantic network provides necessary and/or sufficient conditions for the application of concepts. As a consequence, such conditions hold for every instance of concepts, and cannot be violated by any specific exemplar. Therefore, what can be inferred on the basis of prototypical knowledge can extend, but can in no way conflict with what can be deduced from the DL based component (this aspect is synthesized by the direction of the arrow in the next figure).

The figure 3 below shows the general structure of our proposal for a hybrid conceptual representation. A concept X is assumed to be composed by different bodies of conceptual knowledge (the meaning of the arrows connecting the two components has to be intended as a predicate modeling a mereological relation). The final arrow connecting the two components, instead, indicates the conditions, expressed by logical rules, regulating the access to the different body of knowledge.

[3] As an example of information integration through the linked data approach is he following:
SUBJECT: http://dbpedia.org/page/Venus;
PREDICATE: http://www.w3.org/2002/07/owl#sameAs;
OBJECT: http://dbpedia.org/page/Phosphorus_(morning_star).

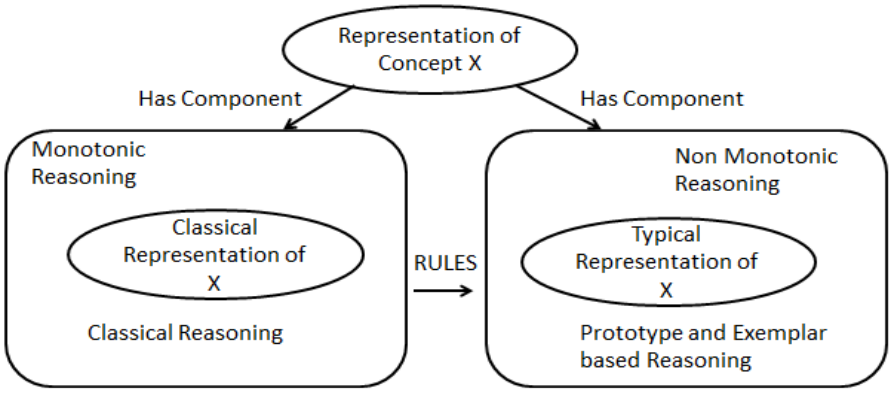

Figure 3: A general picture of the hybrid conceptual representation proposal

From a conceptual modelling point of view, such an architecture extends the modeling possibilities of the standard conceptual representations in formal ontologies. Let us suppose, in fact, that we want to model the assertion "A typical rose is red". If we consider, for example, the foundational approach proposed in DOLCE [Masolo et al. 2009], based on the distinction between the categories of "Quality" and "Quality Regions" inspired by Gärdenfors [2000], it is possible to model the information that a certain rose (rose#1, in the figure 4) has a certain Color (this is expressed via the *inherence relation*, qtc. Ex: qtc(#rose1)) and that the particular color of the rose#1 has a particular *redness* at a certain time *t* (this is expressed via the *quale, ql,* relations: ql(qtc(rose#1, *t*))). The figure 4 below (from Masolo et al. 2009) illustrates this solution.

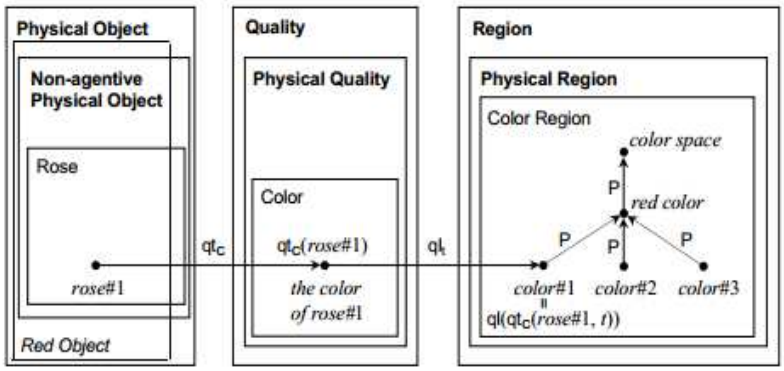

Figure 4: Connecting Concepts to Qualities and Quality Regions in an Foundational Ontology (from Masolo et al. 2009)

However, this approach, despite its powerful, does not cope with the problem of representing a prototypical piece of information such as "A typical rose is red". In fact, the standard - DL based - ontological languages, do not allow to represent and

identify an instance of rose (let us suppose #roseP) as prototypical, nor is it possible to represent a prototypical class of roses ("to be red" is neither a necessary nor a sufficient condition for being a rose).

Our proposal allows to represent typical information (such as that prototypical roses are red) because the component representing the typical information is assumed to be external, but connected, to the ontological knowledge base. We assume that this component is expressed in a separate formalism. An example of how the information that "typical roses are red" can be modeled in our framework is given in figure 5.

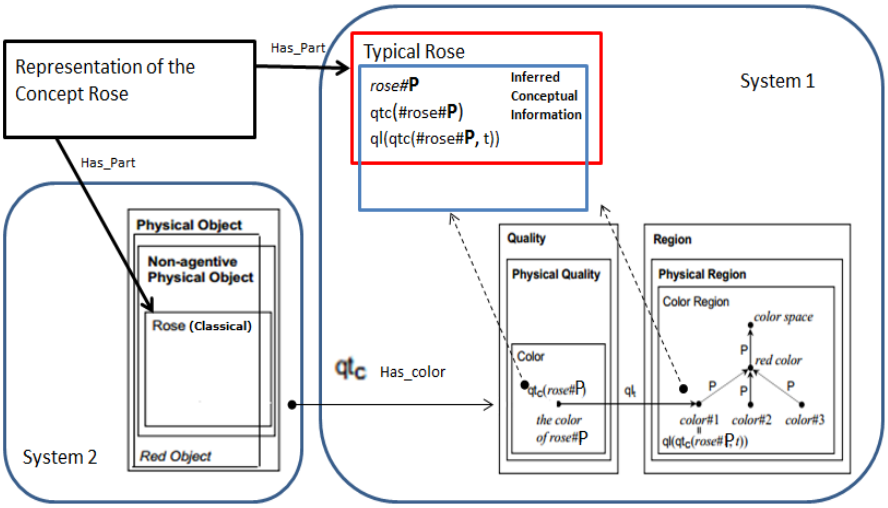

Figure 5: Modeling Typical Information in a Hybrid Conceptual System

The picture shows that the typical component is modelled externally with respect to the classical knowledge base. As we will show in sections 8 and 9, the framework that we propose for representing this component makes it possible to model and retrieve typical information associated to a certain concept. Such external solution allows to perform, in principle, forms of typicality based reasoning only outside the compositional component of the representational system. This aspect represents one of the main features of our proposal. Among other things, this solution makes it possible to avoid consistency problems in the compositional part (this was one of the main problems both in frame based systems as well as in hybrid knowledge representation approaches) introducing at the same time within the ontology (intended in a broad sense) the possibility to expand the allowed types of reasoning.

8 Conceptual Spaces: A Proposal for Geometric Representation of Non Classical Concepts

In the rest of this paper, we shall consider conceptual spaces [Gärdenfors, 2000] as a possible framework to develop some aspects of the ideas presented in the above sections. Conceptual spaces are geometrical representations of knowledge that consist

of a number of quality dimensions. In some cases, such dimensions can be directly related to perceptual data; examples of this kind are temperature, weight, brightness, pitch. In other cases, dimensions can be more abstract in nature. To each quality dimension is associated a geometrical (topological or metrical) structure. The central idea beyond this approach is that the representation of knowledge can take advantage from the geometrical structure of conceptual spaces. For example, 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 (often Euclidean and Manhattan distances) based on the geometry and metric of the space.

In this framework, concepts correspond to regions and regions with different geometrical properties correspond to different kinds of concepts.

Let us briefly consider some example that are strictly related to the representation of sensory data, namely a conceptual space for color ([Gärdenfors, 2000], sect. 1.5). One possibility to describe colors consists in choosing three parameters: *brightness*, *saturation* and *hue*. Such parameters can be viewed as the dimensions of a chromatic conceptual space: brightness varies from white to black, so it can be represented as a linear dimension with two end points; saturation (i.e., color intensity) ranges from grey to full intensity, therefore, it is isomorphic to an interval of the real line; hues can be arranged in a circle, on which complementary colors (e.g. red and green) lie opposite to each other. As a result, a possible conceptual space for colors is a tridimensional space with the structure of the familiar color spindle (figure. 6).

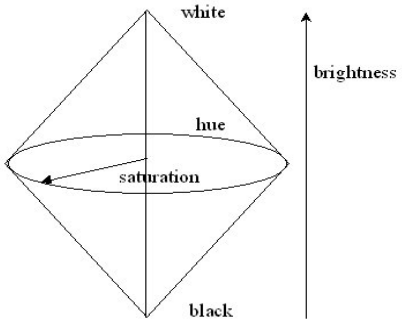

Figure 6: A conceptual space for colors

From our point of view, conceptual spaces could offer a computational and representational framework to develop some aspects of our proposal of representing concepts in the terms of both prototypes and exemplars.

Conceptual spaces are suitable to represent concepts in non-classical, "typical", terms. The regions representing concepts can have soft boundaries. Moreover, 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 (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" ([Gärdenfors, 2000], p. 9). Given a convex region, in fact, to each point can be associated a certain degree of centrality, which can be interpreted as a measure of its typicality. Moreover, single exemplars correspond to single points of the space.

Gärdenfors concentrates almost exclusively on representations based on prototypes. However, in our opinion, conceptual spaces are well suited also for modeling of concepts in terms of exemplars [Frixione and Lieto, 2013a and 2013b], and, therefore, for developing hybrid, 'prototype- and exemplar-based', solutions. As said before, in conceptual spaces exemplars are represented as points. Therefore, if the prototypical representation of some concept C corresponds to a convex region (with the prototype of C corresponding to the center of the region), then it is easy to keep within the same representation also the information concerning (some) known exemplars of C. This can facilitate many forms of conceptual reasoning, which are psychologically plausible, and which can turn useful in many application contexts related to the semantic search such as, for example, that one regarding the area of question answering or concept retrieval in big data[4].

This is due to the fact that conceptual spaces are not specifically designed to represent prototypes; rather, they are a general framework for knowledge representation in which, at certain conditions, prototypes emerge as a consequence of the global geometric properties of the model. In this respect, conceptual spaces deeply differ from traditional approaches in which prototypes are explicitly represented as local data structure - for example, as frames, or as (possibly weighted) lists of features.

As a consequence, the theory of conceptual spaces is compatible with the possibility of representing concepts that do not correspond to properties, i.e. to convex regions in the space. Non-linearly separable categories ([Medin and P.J. Schwanenflugel, 1981] - see sect. 4.3 above) are exactly "non convex" concepts of this sort. Exemplar based representation are probably an adequate choice for representing of a non linearly separable categories, and this can be achieved in the framework of conceptual spaces. Also in this case, however, it is likely that the geometrical structure of the space allows many relevant forms of reasoning (based, for example, on the metric associated to the space itself).

Gärdenfors [Gärdenfors, 2004] proposed conceptual spaces as a tool for representing knowledge in Web ontologies. His claim is that that description logics and traditional Semantic Web languages derived from them (e.g. OWL) do not allow to account for two important aspects of semantics. Namely: the representation of semantic similarity and the combination of concepts that cannot be expressed by conjunction of properties (consider, for example, the fact that a red face is not exactly red).

Within the field of conceptual modeling, conceptual spaces have been proposed [Guizzardi et al. 2004] to provide a foundation for datatypes i.e., abstract structures delimiting the value space for data attributes. According to these authors, conceptual spaces bring in some additional benefit if compared to other conceptual modeling approaches (e.g. formal relations).

---

[4] Let us consider, for example, the biomedical domain, characterized by an enormous amount of distributed information and by the fact that the conceptual characterization of the symptoms related to a particular disease is often based on typical patterns. In such case, the use of such hybrid architecture - and the exploitation of the corresponding double level of reasoning (compositional and typicality based) - could be useful to extend the knowledge of the classical biomedical ontologies by allowing to retrieve, given a set of symptoms, the corresponding typical diseases (and treatments) associated to them.

Following Gärdenfors' suggestions [Gärdenfors, 2004], Adams and Raubal [Adams and Raubal, 2009] proposed a conceptual space algebra as the basis for what they call the Conceptual Space Markup Language (CSML), an XML based interchange format for conceptual spaces, which facilitates the creation and sharing of conceptual structures using geometric information.

A specific feature of CSML is the possibility of encoding contrast classes within the conceptual space representations. Through the introduction of contrast classes it is possible to take into account the different nuances of meaning that a concept assumes when combined with another contrast class. For example: the term *warm Swedish vacation* involves different semantics for the term "warm" than does *warm California vacation*. In order to account for this difference, WARM can be represented as a contrast class in a conceptual space, i.e. as a sub-region of the entire climate domain class. The combination of the contrast class WARM with another class is not compositional (i.e., it is not obtained as the intersection of the two classes). Rather, it is the result of a geometric projection of the WARM region onto the other class climate property. Adams and Raubal [Adams and Raubal, 2010] suggest that in this way it is possible to (non-monotonically) infer that warm weather in Sweden is not a particular case of what, in general, can be considered warm weather in Europe, even if Sweden is represented as an instance of the class EUROPEAN COUNTRY.

9 Integrating Conceptual Spaces in DL Representations

In this section we propose a way to integrate conceptual spaces in usual, DL ontological representations, in order to keep the DL representation fully monotonic, and to demand the representation of prototypical effects to conceptual spaces.

In the context of a different field of application, a solution in some sense similar has been proposed in [Chella, Frixione and Gaglio, 1997]. Also in their proposal, in fact, conceptual spaces are coupled with a DL knowledge base (built in KL-ONE), in such a way to keep the DL representation fully monotonic, and to charge the conceptual space with the representation of typicality effects. However, deep differences exist with the present proposal. In fact in [Chella, Frixione and Gaglio 1997] the authors were not aimed to the development of a general approach to ontological representation; rather, the aim was the representation of a particular domain of knowledge (complex three dimensional shapes) with a specific task in mind (the reconstruction of perceived visual scenes in an autonomous agent). So, the DL representation were coupled just with a homogeneous, specific type of conceptual space. The needs of ontological representations are different: here we are not interested in a specific type of conceptual space, suited for the representation of a particular kind of knowledge; rather, in principle we keep open the possibility of associating to a DL ontology any type of conceptual space. Moreover, we also admit the possibility of associating to different portion of an ontology different "local" conceptual spaces. Finally, in [Chella, Frixione and Gaglio, 1997] the possibility of representing typicality effects in terms of exemplar within conceptual spaces is not taken into account. In other respects our modelling approach is also akin to that in [Guizzardi et al., 2004]. However while these authors, following the modeling approach proposed in DOLCE [Masolo et al.2009], simply use the conceptual spaces notions of "Quality dimension" and "Quale" in order to propose well founded representations of datatypes in the ontologies, in our case we maintain the use the full theoretical power of the conceptual spaces in order to model both properties

corresponding to datatypes as well as standard logical predicates known as *object properties* (i.e. relational properties asserting information between members of two classes). Furthermore, another relevant difference of their approach w.r.t. our proposal is represented by the fact that, in [Guizzardi et al., 2004], the representation of the information regarding the "Qualia", is supposed to be completely internal to the ontological component. Otherwise, we propose to demand such typical aspects to an external conceptual space and to combine, via rules, such component with a classical ontological one in order to allow the possibility of performing, externally, non-monotonic inferences. The possibility of using the conceptual spaces representations for performing typicality based reasoning is not considered in [Guizzardi et al., 2004].

In many cases, a consistent part of the information that is associated to a concept can be represented in the terms of some conceptual space, in order to take advantage of its geometrical structure. However, in general, given a certain ontology, it is not plausible to associate to it a unique "global" conceptual space in which every concept of the ontology can be represented. And it is not plausible that *every* concept can be represented in the terms of some conceptual space, or that *every* characteristic of a given concept can be represented in the terms of some conceptual space. Which could be, for example, the conceptual space associated in an ontology to the class THING? Therefore, our approach consists in associating (possibly different) conceptual spaces to different parts of a taxonomy.

Let us suppose for example that a given taxonomy includes the concepts COLOR. We can imagine to associate a "local" conceptual space to each of them, characterised in the terms described above in section 8.

Let us suppose that we want to associate a certain conceptual space CS_C to a given concept C in a taxonomy, and that space CS_C is characterized by the dimensions $d_1, \ldots d_n$. The $d_1, \ldots d_n$ can be represented within the DL ontology as attributes of the concept C, as showed in the figure below.

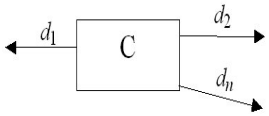

Figure 7: Associating a conceptual space to a concept

The value restrictions of the attributes $d_1, \ldots d_n$ can be used to characterise the various dimensions d_i (by specifying e.g. their range of values, their topology, whether they are continuous or discrete, and so on).

In this way, we can associate to the various sub-concepts $C_1, \ldots C_m$ of C different regions (or sets of points) in the space CS_C. We can admit also the possibility that, in the case of some specific sub-concepts of C, the conceptual space is enriched by adding further quality dimensions[5].

It must be noted that, in general, it is not necessary to specify within the DL formalism which region of the conceptual space exactly corresponds to each class C_i. It can be calculated on the basis of the geometrical structure of the conceptual space itself.

[5] As previously shown this solution is compliant with the representation of the classes "Quality" and "Qualia" in the foundational ontology DOLCE (see Masolo et al. 2009).

Let us consider for example the color space described in sect. 8, and suppose that a DL taxonomy includes a class COLOR. If we want to associate to it a conceptual space with the structure described above, then three attributes *brightness*, *saturation* and *hue* must be added to COLOR (figure 8). The values of such attributes must be restricted to suitable classes: the black-white interval in the case of brightness; the grey-full intensity interval in the case of saturation; and the set of polar co-ordinates in the case of hue (since the hue dimension is circular, the hue of a color can be expressed in terms of polar co-ordinates). Classes black-white interval and grey-full intensity are supposed to be subclasses of some class like CONTINUOS INTERVAL, and so on.

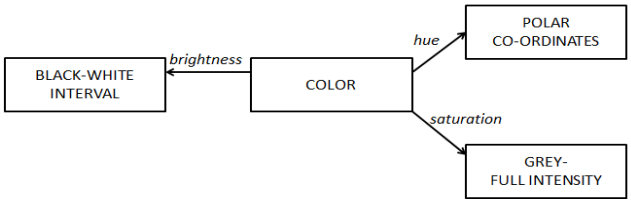

Figure 8: The COLOR concept with its quality dimensions

The various subclasses of COLOR (i.e., the various types of color) correspond to particular regions in the conceptual space (i.e., to particular ranges of values for the three attributes brightness, saturation and hue). However, the DL taxonomy must not necessarily specify which particular ranges of values correspond to a certain color class (say, to RED). Rather, to each color class can be associated a prototype (i.e., a point in the conceptual space which represents the best sample of that class), or a set of exemplars (a set of points that the correspond the known instances of that color class). Given a new instance, it can be categorised by calculating its degree of similarity with the prototypes or with the known exemplars of the various chromatic classes. For example the distance between two instances p1, p2 can be calculated combining the Euclidean distance and the angular distance intervening between the points.

With respect to Gärdenfors's approach, we adopt both prototypes and exemplar based representation of non-classical concepts while Gärdenfors concentrates almost exclusively on prototypes. In fact, conceptual spaces are well suited for representing concepts in terms of exemplars and offer the advantage of a unique framework in which to integrate both approaches (sect. 7). Another difference is that we put greater emphasis on the forms of reasoning that can be performed by the symbolic (DL) component. In particular, according to a dual process perspective (sect. 4.1), we assume that conceptual spaces are responsible for type 1 processes, while type 2, "difficult" deductive tasks are demanded to the DL formalism. Another innovative aspect is the possibility to associate different CSs to different parts of a logic oriented knowledge base.

As mentioned above, a possible technological solution for enriching the representational level of ontology based technologies could be the encoding of conceptual spaces as Linked Data [Bizer et al. 2009]. In this way, reasoning based on geometric representations (e.g. similarity calculation, prototypical effects and non-classical concept combination) could be performed on conceptual spaces, independently from logic based (e.g. OWL) representations and, then, be integrated

with the results coming from the latter[6]. This allows to extend the reasoning capabilities of existing ontological representations, thus enabling the semantic technologies to answer at complex queries based on typical information. More specifically, there are at least two tasks that current ontology based technologies are not yet able to perform, for which a hybrid proposal, based on DLs and conceptual spaces, could be fruitful. Namely: (i) question answering based on typical information, and (ii) question answering based on contextual information.

As an example of the first kind, consider a query like "which kind of citrus fruit is yellow?"; the relevant answer would be LEMON[7]. As an example of the second kind, consider the WARM vacation example of section 8: in the conceptual space component it would be possible to model certain concepts in order to obtain a non compositional concept combination and, therefore, context sensitive inferences and retrieval.

A first implementation of the proposed architecture (the whole details are provided in [Ghignone, Lieto and Radicioni, 2013]), has been realized and preliminarly tested in a system involved in a question answering task for typicality based queries (i.e. queries in natural language where the questions are provided through imprecise, common sense, descriptions of a given concept/target to be retrieved). The task designed for the system evaluation consisted in the individuation of the appropriate concept, to which a given description was referred to, by exploiting the inferential capability of the proposed hybrid conceptual architecture. Examples of these common-sense descriptions are: "the big carnivore with black and yellow stripes" denoting the concept of tiger, or "the sweet water fish that goes upstream" denoting the concept of salmon, and so on. The obtained preliminary results are encouraging and show that the identification and retrieval of concepts described with typical features is considerably improved by such hybrid architecture w.r.t. the classical case, based symply on the use of ontological knowledge. A deeper evaluation on a larger set of coupled local conceptual spaces and ontologies containing a large amount of structured knowledge (such as OpenCYC and SNOWMED) represents an ongoing work, as well as the publication of the local conceptual spaces in a linked open data format and the corresponding alignment with ontologies, aimed at the realization of a distributed, web based, service-oriented conceptual system.

---

[6] Of course the integration of different types of reasoning processes requires some conciliation strategies. We have considered different strategies in some previous works. One, similar to what is proposed here, is more conservative and safe [Frixione and Lieto 2012a] since the typicality based reasoning is considered as an extension of the classical one. Another one, more cognitively grounded on the tenets of the Dual Process Theory, assumes that the output of the typical reasoning is the input of a second level of processing based on classical reasoning [Frixione and Lieto 2013b].

[7] Being yellow is not a necessary condition for being a lemon and, therefore, this property cannot be associated to the class LEMON in a DL ontology. However, from a cognitive point of view, the property of being yellow is relevant to characterize the concept LEMON. According to our proposal, this piece of knowledge can be represented, and therefore retrieved, in the prototypical component associated to the concept LEMON.

10 Conclusions and Future Developments

In this paper we proposed a hybrid 'prototype and exemplar-based' approach to the representation of non-classical concepts in formal ontologies. The main elements of our proposal are the following: i) splitting of the conceptual representation in compositional and typical components in order to extend the classical conceptual model hypothesized by the standard ontology based systems and formalisms (ii) possibility to integrate these representations using the Linked Data approach and (iii) division, and successive integration, of the type of reasoning processes operating on the interconnected components of the knowledge base. In order to implement our proposal, a promising research direction could be the extension, to RDF semantics and beyond, of the above mentioned Conceptual Space Markup Language (CSML) formalism. This would allow the creation , in the so called Linked Data Cloud, of an interrogable Conceptual Space bubble through which to extend the classical conceptual representations. Since extensibility is one of the fundamental design principles of CSML, this proposal seems feasible and, in our opinion, deserves to be further investigated.

References

[Adams and Raubal, 2010], Adams B. and Raubal M., "The semantic web needs more cognition", *Journal of Semantic Web*, **1**(1-2), 2010, pp. 69-74.

[Adams and Raubal, 2009], Adams B. and Raubal M., "The conceptual space markup language (CSML): Towards the cognitive semantic web", *3rd IEEE ICSC 2009*, Berkeley, pp. 253-260.

[Baader et al. 2010], Baader F., Calvanese D, D. McGuinness, D. Nardi and P. Patel-Schneider, *The Description Logic Handbook*, 2nd edition, Cambridge University Press, Cambridge, 2010.

[Bizer et al., 2009], Bizer C., Heath T. and Berners-Lee T., "Linked data - The story so far", *International Journal on Semantic Web and Information System* **5**(3), 2009, 1-22.

[Brachman and Levesque, 1985] Brachman R. and H. Levesque (eds.), *Readings in Knowledge Representation*, Morgan Kaufmann, Los Altos, CA, 1985.

[Brachman and Schmoltze, 1985] Brachman R. and J.G. Schmolze J.G., An overview of the KL-ONE knowledge representation system, *Cognitive Science* **9**, 1985. 171-216.

[Chella, Frixione and Gaglio, 1997] Chella A., Frixione M. and Gaglio S., "A cognitive architecture for artificial vision", *Artificial Intelligence*, 89 (1–2), 1997, 73-111.

[Evans and Frankish, 2008] Evans J.S.B.T: and Frankish K. (eds.), *In Two Minds: Dual Processes and Beyond*. New York, NY: Oxford UP, 2008.

[Eiter et al. 2008], Eiter T., Ianni G., Krennwallner and Polleres. Rules and Ontologies for the Semantic Web, Springer, 2008.

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

[Fodor, 1987] Fodor J., *Psychosemantics*, Cambridge, MA: The MIT Press, 1987.

[Frixione and Lieto, 2012a] Frixione M. and Lieto A., "Representing Concepts in Formal Ontologies: Compositionality vs. Typicality Effects" in *Logic and Logical Philosophy*, **21**(4).

[Frixione and Lieto, 2012b] Frixione M and Lieto A., "Representing Non Classical Concepts in Formal Ontologies: Prototypes and Exemplars", in Vol 439, LNCS, Springer, 2013, pp.171-182.

[Frixione and Lieto, 2013a] Frixione M. and Lieto A., "Exemplars, prototypes and conceptual spaces", Advances in Intelligent Systems and Computing, Volume 196, 131-136, Springer.

[Frixione and Lieto, 2013b] Frixione M. and Lieto A., "Dealing with Concepts: From Cognitive Psychology to Knowledge Representation", Frontiers in Psych. and Behav. Sci., 2013, 96-106.

[Frixione and Lieto, 2010] Frixione M. and Lieto A., "The computational representation of concepts in formal ontologies: Some general consideration", *Proc. KEOD 2010*, pp. 396-403.

[Gärdenfors, 2004] Gärdenfors P., "How to make the semantic web more semantic", *Proc. 3rd Int. Conf. on Formal Ontology in Information Systems* (*FOIS*), Torino, Italy, 2004.

[Gärdenfors, 2000] Gärdenfors P., *Conceptual Spaces: The Geometry of Thought*. The MIT Press/Bradford Books, Cambridge, MA, 2000.

[Ghignone, Lieto and Radicioni, 2013], Typicality based Inference by Plugging Conceptual Spaces into Ontologies*, Proceedings of AIC 2013, Int. Workhop on AI and Cognition*, pp.68-79.

[Gruber, 1993] T. Gruber. A translation approach to portable ontologies. *Knowledge Acquisition*, 5(2):199-220, 1993.

[Guizzardi et al., 2004] Guizzardi G., G. Wagner and H. Herre, "On the foundations of UML as an ontology representation language", *Proc. of EKAW*, 2004, Springer Verlag, Berlin, LNCS.

[Machery, 2009] Machery E., *Doing without Concepts*. Oxford University Press, Oxford, 2009.

[Malt, 1989] Malt B.C., "An on-line investigation of prototype and exemplar strategies in classification", *Journal of Experimental Psychology* **15**(4), 1989, 539-555.

[Masolo et al. 2009], Masolo, C., Borgo S., Gangemi A., Guarino N., Oltramari Al, "WonderWeb DeliverableD18", Technical Report, CNR.

[Medin and Schaffer, 1985] Medin D.L. and M.M. Schaffer, "Context theory of classification learning", *Psychological Review* **85**(3), 1978, 207-238.

[Medin and Schwanenflugel, 1981] Medin D.L.and Schwanenflugel P.J., "Linear separability in classification learning", *J. of Exp. Psyc.: Human Learning and Memory* 7, 1981, 355-368.

[Minsky M., 1975] Minsky M., "A framework for representing knowledge", in Patrick Winston (ed.), *The Psychology of Computer Vision*, McGraw-Hill, New York, 1975.

[Murphy and Medin, 1985] Murphy G. and Medin D., "*The role of theories in conceptual coherence*", Psychological Review, 1985.

[Murphy, 2002] Murphy G.L., *The Big Book of Concepts*. The MIT Press, MA, 2002.

[OWL] http://www.w3.org/TR/owl-features

[Quillian, 1968] Quillian M.R., Semantic memory. In M. Minsky (ed.), *Semantic Information Processing*, The MIT Press, Cambridge, MA, 1968.

[Rosch, 1975] Rosch E., "Cognitive representation of semantic categories", *Journal of Experimental Psychology*, **104,** 1975, 573-605.

[SWRL] http://www.w3.org/Submission/SWRL

[Thagard, 2012], Thagard, P. (2012). The cognitive science of science: Explanation, discovery, and conceptual change. Cambridge, MA: MIT Press.