Dealing with Concepts:
From Cognitive Psychology to Knowledge Representation

Marcello Frixione[*1], Antonio Lieto[2]

[1]DAFIST, University of Genova, via Balbi, 4 - 16126 Genova, Italy

[2]Dept. of Computer Science, University of Torino, C.So Svizzera, 185 - 10149 Torino, Italy

[*1]marcello.frixione@gmail.com; [2]lieto@di.unito.it

Abstract- **Concept representation is still an open problem in the field of ontology engineering and, more generally, of knowledge representation. In particular, the issue of representing "non classical" concepts, i.e. concepts that cannot be defined in terms of necessary and sufficient conditions, remains unresolved. In this paper we review empirical evidence from cognitive psychology, according to which concept representation is not a unitary phenomenon. On this basis, we sketch some proposals for concept representation, taking into account suggestions from psychological research. In particular, it seems that human beings employ both prototype-based and exemplar-based representations in order to represent non classical concepts. We suggest that a similar, hybrid prototype-exemplar based approach could also prove useful in the field of knowledge representation technology. Finally, we propose conceptual spaces as a suitable framework for developing some aspects of this proposal.**

Keywords- Concept Representation; Psychological Theories of Concepts; Formal Ontologies; Conceptual Spaces

I. INTRODUCTION

This article deals with the problem of representing non- classical concepts, with particular attention to artificial systems, such as formal ontologies. According to our approach, concept representation in artificial systems can draw considerable advantage from the empirical results of cognitive psychology. By non-classical concepts we mean concepts that cannot be represented in terms of sets of necessary and/or sufficient conditions. After introducing the problem (Sect. II), we review some empirical evidence from cognitive psychology (Sect. III). In Sect. IV we point out 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 prototypical effects separate from compositional representations of concepts; the possibility of developing hybrid prototype and exemplar-based representations of concepts. We concentrate in particular on the latter point. Prototype and exemplar-based models of non classical concepts are both plausible and can account for different aspects of human abilities (Sect. V). We argue that these results could suggest the adoption of a hybrid prototype and exemplar-based approach also in the field of knowledge representation and of formal ontologies. In Sect. VI we sketch the proposal of an architecture for concept representation based on both prototypes and exemplars. Then we introduce conceptual spaces (Sect. VII) as a suitable framework for the development of some aspects of our proposal (Sect. VIII). Some conclusions follow.

II. REPRESENTING NON CLASSICAL CONCEPTS

The representation of common sense concepts is still an open problem in ontology engineering and, more generally, in Knowledge Representation (KR) [1]. Cognitive Science has shown 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, whereas Eleanor Rosch's experiments [2] – historically preceded by the philosophical analyses of Ludwig Wittgenstein [3] – showed that ordinary concepts can be characterized in terms of prototypical information.

These results influenced early research into knowledge representation: KR practitioners initially tried to take the suggestions from cognitive psychology into account and designed artificial systems – such as frames [4] and early semantic networks [5] – able to represent concepts in "non classical" (prototypical) terms (for early KR developments, see also the papers collected in Brachman and Levesque, [6]).

However, these early systems lacked clear formal semantics and a satisfactory meta-theoretic account, and were later sacrificed in favor of a class of formalisms stemming from the so-called structured inheritance semantic networks and the KL-ONE system [7]. These formalisms are known today as *description logics* (DLs) [8]. DLs are logical formalisms, which can be studied by means of the traditional, rigorous meta-theoretic techniques developed by logicians (for example, the definition of semantics in formal terms, in such a way as to obtain rigorous correctness and completeness proofs; the study of the formalism's computational complexity properties, and so on). Technically, description logics are subsets of the first order predicate calculus (FOL) that, if compared to the full FOL, is computationally more efficient.

Hence, DLs are concept-oriented formalisms that provide a number of constructs for concept description. However, they do not allow for exceptions to inheritance, or for the possibility to represent concepts in prototypical terms. From this point of

view, therefore, such formalisms can be seen as a revival of the classical theory of concepts. As far as prototypical information is concerned, such formalisms offer only two possibilities: representing it by resorting to tricks or ad hoc solutions or, alternatively, ignoring it. The first solution is obviously unsuitable as it could have disastrous consequences for the soundness of the knowledge base and for overall system performance. The second drastically reduces the expressive power of the representation. For example, in information retrieval terms, this could severely affect the system's recall. Let us suppose that you are interested in documents about flying animals. A document about birds is likely to interest you, because most birds are able to fly. However, flying is not a necessary condition to being a bird (there are many birds that are unable to fly). So, the fact that birds usually fly cannot be represented in a formalism that allows only the representation of concepts in classical terms, and the documents about birds will be ignored by your query.

Nowadays, DLs are widely adopted within many fields of application, in particular within the area of ontology representation. For example, OWL (Web Ontology Language, see http://www.w3.org/TR/owl-features/) 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 fail to solve the problems of representing concepts in prototypical terms.

Within the field of logic oriented knowledge representation, rigorous approaches have been designed to make possible the representation of exceptions, and that are therefore, at least in principle, suitable for dealing with (some aspects of) "non-classical" concepts. Examples are fuzzy and non-monotonic logics. Thus, the adoption of logic oriented semantics is not necessarily incompatible with the representation of prototypical effects. Various fuzzy and non-monotonic extensions of DL formalisms have been proposed. Nevertheless, such approaches pose various theoretical and practical problems, which in part remain unsolved (for a more detailed account, see [9]).

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.

III. THE ADVANTAGE OF A COGNITIVE APPROACH TO TYPICALITY IN ARTIFICIAL SYSTEMS

As mentioned above, prototypical effects in categorisation and category representation generally are not only crucial for the empirical study of human cognition, but are also of the utmost 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 words, that they possess 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 information to be represented in typical/prototypical 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[1] 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 may happen that we know the classical definition of a concept, but typical/prototypical knowledge still plays a central role in many cognitive tasks. Consider the following example: nowadays most people know the necessary and sufficient conditions for being WATER: water is exactly the chemical substance whose formula is H_2O, i.e. the substance whose molecules are formed by one atom of oxygen and two atoms of hydrogen. However, in most cases, when we categorise a sample of stuff as WATER in everyday life, we do not take advantage of this piece of knowledge. We use such prototypical traits such as the fact that (liquid) water is usually a colourless, odourless and tasteless fluid. As a further example, consider the concept GRANDMOTHER. Everybody knows a classical definition for it: x is the grandmother of y if and only if x is the mother of a parent of y. However, in many cases we do not use this definition to categorise somebody as a grandmother. We resort to typical traits: grandmothers are old women who take care of children, who are tender and polite with them, and so on. Once more, the problem is not different in the case of artificial systems: generally a system that has to categorise some stuff as WATER cannot perform chemical analyses, and it must trust prototypical evidence.

Therefore, the use of prototypical knowledge in cognitive tasks such as categorisation 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

---

[1] A typical example of such a misuse of DLs is the inclusion in the description of concepts of information that is not necessary for the corresponding category – e.g., the information that DOGs have exactly four legs, that LEMONs are yellow, or that BIRDs are able to fly. But the possibilities in this sense are limited only by the creativity of the designer of the knowledge base.

recurrent errors in reasoning could have a deeper "rationality" within the general economy of cognition). It has to do with the constraints that concern every finite agent that has limited access to the relevant knowledge for a given task. This is the case of both natural and artificial cognitive systems.

IV. SOME SUGGESTIONS FROM COGNITIVE SCIENCE

Some recent trends in 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 point out 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, developed within the context of the so-called "dual process" accounts of reasoning (Sect. A. below); (ii) the proposal to keep prototypical effects separate from the compositional representation of concepts (Sect. B.); and (iii) the possibility to develop hybrid prototype and exemplar-based representations of concepts (Sect. C.).

A. A Dual Process Approach

Cognitive research on concepts [10, 11] 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, and in particular by the so-called dual process theories of reasoning and rationality [12, 13]. 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 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 and sequential). On the contrary, exceptions play an important role in processes such as categorization and inheritance, which are more likely to be Type 1 tasks: 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 benefits from the management of exceptions. The latter 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).

B. A Pseudo-Fodorian Proposal

According to Fodor, concepts cannot be prototypical representations, since concepts must be compositional, and prototypes do not compose. On the other hand, by virtue of the criticisms of "classical" theory, concepts cannot be definitions. Therefore, Fodor argues that (most) concepts are atoms, i.e. they 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 [14, 15]. Of course, Fodor acknowledges the existence of prototypical effects.

However, he claims that prototypical representations are not part of concepts. Prototypical representations make it possible to individuate the reference of concepts, but they must not be identified with concepts. Consider for example the concept DOG. Of course, in our minds there is some prototypical representation associated to DOG (e.g., that dogs usually have fur, that they typically bark, and so on). But this representation does not coincide with the concept DOG: DOG is an atomic, unstructured symbol.

We borrow Fodor's hypothesis that compositional representations and prototypical effects are assigned 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 prototypical representation of categories is responsible for such processes as categorisation, but it does not affect the inferential behaviour of the compositional component.

(It must be noted that our proposal is not entirely "Fodorian", at least in the following three senses: i. We leave out the problem of the nature of the semantic content of conceptual representations. Fodor endorses a causal, informational theory of meaning, according to which the content of concepts is constituted by a nomic mind-world relation. We are in no way committed to such an account of semantic content - in general, the philosophical problem of the nature of the intentional content of representations is largely irrelevant to our present purposes). ii. Fodor claims that concepts are compositional and that as prototypical representations are not compositional, they cannot be concepts. We do not take any position on which part of the system we propose must be considered as truly "conceptual". In our opinion, both the compositional and the prototypical components contribute to those abilities that we usually describe in terms of the possession of concepts. iii. According to Fodor, almost all lexical concepts are atomic. We maintain that many lexical concepts, even though indefinable in classical theory terms, should exhibit some form of structure, and that such a structure can be represented, for example, by means of a DL taxonomy).

C. Prototypes and Exemplars

Different positions and theories on the nature of concepts are available within the field of cognitive psychology. Usually, they are grouped in three main classes, namely: prototype views, exemplar views and theory-theories (see e.g. [10] and [11]), all of which are assumed to account for (some aspects of) prototypical effects in conceptualization.

According to the prototype view, 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, which 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 as *having fur*, *having retractile claws*, *meows*. The weights are numeric values that express the relevance of each feature. For example, having feathers could be considered more significant than having two legs in categorizing something as a bird, since birds are the only (living) animals with feathers, while there are many other animals with two legs.

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 have encountered during our lifetime.

Theory-theory approaches adopt some form of holistic point of view regarding concepts. According to some versions of 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 turn out to be 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 in order to explain psychological data (see again [10] and [11]). In this perspective, we propose to integrate some of them in computational representations of concepts. More precisely, we 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 prototypical effects (for a similar, hybrid prototypical and exemplar based proposal in a different field, see [16]). Here we do not take into consideration the theory-theory approach since, in one sense, it is more vaguely defined than the other two points of view. As a consequence, its computational treatment seems at present to be less feasible.

Prototype and exemplar-based approaches to concept representation are, as mentioned above, not mutually exclusive, and they succeed in explaining different phenomena. Exemplar-based representations can be useful in many situations. Various experiments have shown that instances of a concept that are somewhat dissimilar from the prototype, but are very close to a known exemplar, may be categorized quickly and with high confidence. For example, a penguin is quite dissimilar from the prototype of BIRD. However, if we already know an exemplar of penguin, and if we know that it is an instance of BIRD, it is easier for us to classify a new penguin as a BIRD. This is particularly relevant for concepts (such as FURNITURE, or VEHICLE) whose members differ significantly from one another.

Exemplar-based representations are easier and faster to acquire than prototypes. In some situations, there may be not enough time to extract a prototype from the available information. Moreover, the exemplar based approach makes it easier to acquire concepts that are not linearly separable (see [17]). In the following section we shall review some of the available empirical evidence concerning prototype and exemplar-based approaches to concept representation in psychology.

V. PROTOTYPES, EXEMPLARS AND KNOWLEDGE REPRESENTATION

As mentioned in the previous section, the available experimental evidence suggests that exemplar models are in many cases more successful than prototypes (for a more detailed review of these results, see [18]). In certain cases, for example, a less typical item of a given category is categorized more quickly and more accurately than a more typical one if it is similar to previously encountered exemplars of the same category [19]. Consider the penguin example from the previous section: 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, I may well classify new penguins as birds more quickly and more confidently than less atypical birds (such as, say, toucans or hummingbirds) that I have never encountered before.

Another important source of evidence for the exemplar model stems from the study of linearly separable categories (see, again, [17]). Two categories are linearly separable if, given an item *i*, it is possible to determine to which of them *i* belongs to by summing the evidence concerning each attribute of *i*. For example, let us suppose that two categories are characterized by two attributes, or dimensions, corresponding to the axes in Fig. 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 straight line separating the members of the categories can be drawn.
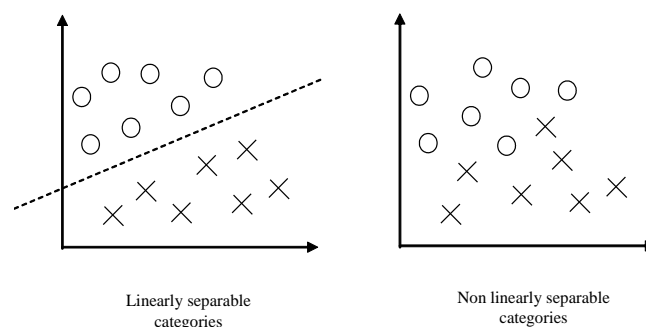
Fig. 1 Linearly separable and non separable categories

According to the prototype approach, it should be more difficult to form a concept of a non-linearly separable category and people should be faster at learning two categories that are linearly separable. However, Medin and Schwanenflugel [17] 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 do not confirm this conclusion. An empirical research supporting the hypothesis of a multiple mental representation of categories can be found in Malt [20]. This study aimed to establish whether people categorize and learn categories using exemplars or prototypes. The empirical data, consisting of behavioral 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. In our opinion, therefore, it is likely 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 better suited for representation in terms of exemplars, and concepts that seem better suited for representation in terms of prototypes. For example, in the case of concepts with a small number of instances, which differ greatly from one another, a representation in terms of exemplars should be more convenient. An exemplar-based representation may also be more suitable 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 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 a definition based on necessary/sufficient conditions is unlikely to be available or adequate for the concept APPLE, the system must incorporate some form of representation that exhibits typicality effects. However, an exemplar based representation is probably not convenient in this case: the system has to deal with thousands of apples, which are all very similar to 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. Let us consider again the concept BIRD. And let us suppose that a certain number of individuals $b_1, ..., b_n$ are known by the systems to be instances of BIRD (i.e. the system knows *for sure* that $b_1, ..., b_n$ are birds). Let us suppose also that one of these b_i's (say, b_k) is a penguin. Then, a prototype PBIRD is extracted from exemplars $b_1, ..., b_n$, and it is associated with the concept BIRD. Exemplar b_k concurs to the extraction of the prototype but, since penguins are rather atypical birds, it will turn out to be rather dissimilar from PBIRD. Let us suppose now that a new exemplar b_h of penguin must be categorized. If the categorization process was based only on the comparison between the target and the prototype, then b_h (which in turn is rather dissimilar from PBIRD) would be categorized as a bird only with a low degree of confidence, in spite of the fact that penguins are birds in all respects. On the other hand, let us suppose that the categorization process also takes advantage of a comparison with known exemplars. In this case, due to its high degree of similarity to b_k, b_h will be categorized as a bird with full confidence. Therefore, even if a prototype for a given concept is available, knowledge of specific exemplars should be valuable in many tasks involving conceptual knowledge. On the other hand, the prototype should be useful in many other situations.

VI. A HYBRID PROTOTYPE – EXEMPLAR BASED ARCHITECTURE

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 prototypical representations.

Concepts in the DL component are represented as in Fig. 2. As usual, every concept can be subsumed by a certain number of superconcepts, and 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.


Fig. 2 A concept in the DL component

Since only necessary/sufficient conditions can be expressed in this component, concepts can here 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.

As an example, consider the network fragment shown in Fig. 3 below:
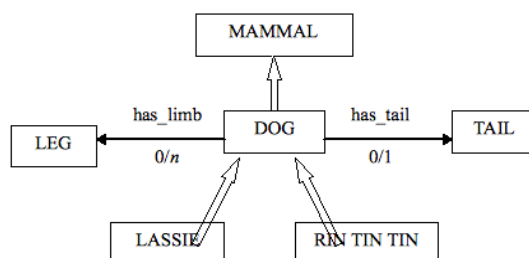

Fig. 3 An example of concept

Here 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, according to Fig. 3, a DOG may or may not have a tail (this is expressed by the number restriction 0/1 imposed on the attribute *has_tail*), and have 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 are 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 are 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, whose value 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 (It may also happen that some information concerning exemplars is represented outside the DL component, in the form of Linked Data [21] – for this notion see below in this section. Typically, this could be the case of "non symbolic" information, such as images, sounds, etc.).

It must be noted that prototypical 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.

The possibility of performing forms of non-monotonic reasoning (namely, non-monotonic categorization of instances) only outside the compositional component of the representation system is 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 and 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.

The possibility of developing hybrid non-monotonic categorization, based on both prototypes and exemplars, has also been

explored in the field of machine learning, where the prototype-exemplar dichotomy in concept representation has been investigated. Consider for example the *PEL-C algorithm*, where *PEL-C* stands for *Prototype-Exemplar Learning Classifier* [22]. The *PEL-C* is a hybrid instance-based algorithm used for machine learning tasks, which accounts for typicality effects in categorization using both prototypes and exemplars. It is based on a learning phase as well as a test phase, and it can also be adopted for a semi-automatic ontology population as well as for updating processes.

VII.   CONCEPTUAL SPACES: A PROPOSAL FOR GEOMETRIC REPRESENTATION OF NON CLASSICAL CONCEPTS

In the rest of this paper, we shall consider conceptual spaces [23] as a possible framework for developing 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 and 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 of 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. Concepts correspond to regions, and regions with different geometrical properties correspond to different kinds of concepts.

Let us briefly consider two examples that are strictly related to the representation of sensory data, namely a conceptual space for color and one for taste ([23], Sect. 1.5). One possibility to describe colors entails 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 and is, therefore, 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 each other. As a result, a possible conceptual space for colors is a tridimensional space with the structure of the familiar color spindle (Fig. 4).


Fig. 4 A conceptual space for colors

As far as taste is concerned, human taste perception seems to depend on four different types of receptors, namely, salt, sour, sweet, and bitter. Thus, a possible model to represent taste is a 4-dimensional space with the structure of a tetrahedron (as shown in Fig. 5).


Fig. 5 A conceptual space for representing taste

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 terms of both prototypes and exemplars.

Conceptual spaces are suitable for representing concepts in non classical terms. The regions representing concepts can have soft boundaries and, moreover, typicality effects can in many cases 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" ([23], p. 9). Indeed, given a convex region, each point can be associated with 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 also well suited for modeling concepts in terms of exemplars [24] 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 also easy to keep the information concerning (some) known exemplars of C within the same representation. This can facilitate many forms of conceptual reasoning, which are psychologically plausible, and which may prove useful in many application contexts.

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, under certain conditions, prototypes emerge as a consequence of the global geometric properties of the model. In this respect, conceptual spaces are profoundly different from traditional approaches in which prototypes are explicitly represented as local data structures - 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 ([17] - see Sect. IV. B. above) are exactly "non convex" concepts of this sort. Exemplar based representation are probably an adequate choice for representing 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 [25] proposed conceptual spaces as a tool for representing knowledge in the semantic web. His claim is that that description logics and traditional semantic web languages derived from them (e.g. OWL) do not allow two important aspects of semantics to be accounted for, 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 [26] to provide a foundation for datatypes, i.e. abstract structures delimiting the value space for data attributes. According to these authors, conceptual spaces introduce some additional benefit if compared to other conceptual modeling approaches (e.g. formal relations).

Following Gärdenfors' suggestions [25], Adams and Raubal [27] proposed 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 [28].

A specific feature of CSML is the possibility of encoding contrast classes within 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 [29] suggest that in this way it is possible to (non-monotonically) infer that warm weather in Sweden is not a particular case of what can generally be considered warm weather in Europe, even if Sweden is represented as an instance of the class EUROPEAN COUNTRY.

VIII.   INTEGRATING CONCEPTUAL SPACES IN DL REPRESENTATION

In this section we propose a way to integrate conceptual spaces in usual, DL based representations, in order to keep the DL representation fully monotonic and to assign the representation of prototypical effects to conceptual spaces (in the context of a different field of application, a similar solution has been proposed in [30] but in other respects, our approach is akin to that in [26]).

In many cases, much of the information associated to a concept can be represented in the terms of some conceptual space, in order to take advantage of its geometrical structure. However, given a certain ontology, it is generally not plausible to associate it with 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. What, for example, might be in an ontology the conceptual space associated to the class THING? Therefore, our approach consists of associating (possibly different) conceptual spaces to different parts of a taxonomy.

Let us suppose for example that a given taxonomy includes the concepts COLOR and TASTE. We can imagine associating

a "local" conceptual space to each of them, characterised in the terms described above in Sect. VII.

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 $d_1, \dots d_n$. The $d_1, \dots d_n$ can be represented within the DL ontology as attributes of the concept C, as shown in Fig. 6.


Fig. 6 Associating a conceptual space to a concept

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

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

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

Let us consider for example the color space described in Sect. VII, and suppose that a DL taxonomy includes a class COLOR. If we want to associate it with a conceptual space having the structure described above, then three attributes *brightness*, *saturation* and *hue* must be added to COLOR (Fig. 7). 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 assumed to be subclasses of some class like CONTINUOUS INTERVAL, and so on.


Fig. 7 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 need not necessarily specify which particular ranges of values correspond to a certain color class (say RED). Rather, each color class can be associated with 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 correspond to the known instances of that color class). A new instance can be categorised by calculating its degree of similarity with the prototypes or with the known exemplars of the various chromatic classes.

IX. CONCLUSIONS AND FURTHER DEVELOPMENTS

In conclusion, it is our opinion that a hybrid prototype/exemplar based representation of non classical concepts could make the representation of common-sense concepts more flexible and realistic, thus also avoiding some frequent misuses of DL formalisms. In this article we have proposed the adoption of conceptual spaces as a common computational framework to develop some aspects of our proposal.

The main novelty of our proposal compared to Gärdenfors's approach probably lies in the adoption of both prototype and exemplar based representation of non classical concepts. Gärdenfors concentrates almost exclusively on prototypes (this is probably due partially to the fact that Gärdenfors has a philosophical background and prototypes have been much more widely discussed within the philosophical tradition than the exemplar based approach). However, as stated in Sect. VII, in our opinion, conceptual spaces are also well suited for representing concepts in terms of exemplars, and offer the advantage of a unique framework in which to integrate both these approaches. Another difference lies in the fact that we place greater emphasis on the forms of reasoning that can be performed by the DL component. In particular, according to a dual process perspective (see the subsection A of Sect. IV), we assume that conceptual spaces are responsible for Type 1 processes, and assign "difficult", Type 2 deductive tasks to the DL formalism. Another innovative aspect is the possibility to associate different CSs to different parts of a logic oriented knowledge base (as outlined in Sect. VIII).

A possible technological solution for enriching the representational level of semantic web technologies could be the encoding of conceptual spaces as Linked Data [21]. 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 of logic based (e.g. OWL based) representations and, then, be integrated with the results coming from the latter[2]. This makes it possible to extend the reasoning capabilities of existing semantic web representations, thus enabling the semantic technologies to respond properly to complex queries based on typical information. More specifically, there are at least two tasks that current semantic web technologies are not yet able to perform and for which a hybrid proposal, based on DLs and conceptual spaces, could be fruitful. Namely: (i) question answering based on prototypical 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[3]. As an example of the second kind, consider the WARM vacation example of Section VII: 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.

In order to implement our proposal, a promising research direction could be the extension of the above mentioned Conceptual Space Markup Language (CSML) formalism. Since extensibility is one the fundamental design principles of CSML, then this proposal seems feasible and, in our opinion, deserves further investigation.

REFERENCES

[1] M. Frixione and A. Lieto, "Representing Concepts in Formal Ontologies: Compositionality vs. Typicality Effects," *Logic and Logical Philosophy*, 21(4), December 2012, 391-414 .

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

[3] L. Wittgenstein, *Philosophische Untersuchungen*, Blackwell, Oxford, 1953.

[4] M. Minsky, "A framework for representing knowledge," in Patrick Winston (ed.), *The Psychology of Computer Vision*, McGraw-Hill, New York, 1975. Also in [6].

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

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

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

[8] F. Baader F., D. Calvanese, D. McGuinness, D. Nardi and P. Patel-Schneider, *The Description Logic Handbook: Theory, Implementations and Applications*, 2nd edition, Cambridge University Press, Cambridge, 2010.

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

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

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

[12] K.E. Stanovich and R. West, "Individual differences in reasoning: Implications for the rationality debate?," *The Behavioural and Brain Sciences* 23 (5), 2000. 645- 665.

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

[14] J. Fodor, *Psychosemantics*, Cambridge, MA: The MIT Press/A Bradford Book, 1987.

[15] J. Fodor, *Concepts: Where Cognitive Science Went Wrong*, Oxford, UK: Oxford University Press, 1998.

[16] F. Gagliardi, "A Prototype-Exemplars Hybrid Cognitive Model of "Phenomenon of Typicality," *Proc. 30th Annual Conf. of the Cognitive Science Society*, Austin, TX, 1176-1181, 2008.

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

[18] M. Frixione and A. Lieto, "Prototypes Vs Exemplars in Concept Representation," *Proc. KEOD 2012*, Barcelona, Spain, 2012.

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

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

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

[22] F. Gagliardi, "Cognitive models of typicality in categorization with instance-based machine learning," *Practices of Cognition. Recent Researches in Cognitive Science*, University of Trento Press, Trento, 2010, 115-130.

---

[2] Of course the integration of different types of reasoning processes requires some conciliation strategies. For this aspect we refer the reader to [1].

[3] As yellow is not a necessary condition for being a lemon, this property cannot be associated to the class LEMON in the 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.

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

[24] M. Frixione and A. Lieto, "Exemplars, prototypes and conceptual spaces," *Proc. BICA 2012 Conference*, Palermo, Italy, 2012.

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

[26] G. Guizzardi, G. Wagner and H. Herre, "On the foundations of UML as an ontology representation language," *Proc. 14th Int. Conf. on Knowledge Engineering and Knowledge Management (EKAW)*, 2004, Springer Verlag, Berlin, LNCS.

[27] B. Adams and M. Raubal, "A metric conceptual space algebra," *Spatial Information Theory - 9th Int. Conf., COSIT 2009*, Aber Wrac'h, France, September 2009 (Vol. 5756, pp. 51-68). Berlin: Springer.

[28] B. Adams and M. Raubal, "The conceptual space markup language (CSML): Towards the cognitive semantic web," *3rd IEEE Int. Conf. on Semantic Computing (ICSC 2009)*, September 14-16, Berkeley, CA, pp. 253-260.

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

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

Marcello Frixione (Genova, 1960) is full professor of Philosophy of Language at the University of Genova.

His research interests include themes of knowledge representation in artificial intelligence and of philosophy of cognitive science. He is author of several essays on the above mentioned areas published on various scientific journals.

Antonio Lieto (Garbagnate Milanese, 1983) is a post-doc research fellow at the Department of Computer Science of the University of Turin, Italy.

His research interests are in the areas of knowledge engineering (ontology engineering) and cognitive modeling. He is author of international and national referred publications on the following topics: formal ontologies, non classical information representation, semantic technologies.