THE COMPUTATIONAL REPRESENTATION OF CONCEPTS IN FORMAL ONTOLOGIES

Some General Considerations

Marcello Frixione and Antonio Lieto
Dipartimento di Scienze della Comunicazione, Università di Salerno
via Ponte Don Melillo, I-84084 Fisciano (Salerno), Italy

Keywords: Concept Representation, Description Logics, Non Monotonic Reasoning, “Dual Process” Theories.

Abstract: Within cognitive science, the “concept of concept” results to be highly disputed and problematic. In our opinion, this is due to the fact that the notion itself of concept is in some sense heterogeneous, and encompasses different cognitive phenomena. This results in a strain between conflicting requirements, such as, for example, compositionality on the one side and the need of representing prototypical information on the other. This has several consequences also for the practice of knowledge engineering and for the technology of formal ontologies. In this paper we propose an analysis of this state of affairs. As a possible way out, in the conclusions we suggest a framework for the representation of concepts, which is inspired by the so called *dual process theories* of reasoning and rationality.

1 INTRODUCTION

Computational representation of concepts is a central problem for the development of ontologies and for knowledge engineering. Concept representation is a multidisciplinary topic of research that involves such different disciplines as Artificial Intelligence (AI), Philosophy, Cognitive Psychology and, more in general, Cognitive Science. However, the notion of concept itself results to be highly disputed and problematic. In our opinion, one of the causes of this state of affairs is that the notion itself of concept is in some sense heterogeneous, and encompasses different cognitive phenomena. This results in a strain between conflicting requirements, such as, for example, compositionality on the one side and the need of representing prototypical information on the other. This has several consequences for the practice of knowledge engineering and for the technology of formal ontologies.

In this paper we propose an analysis of this situation. The paper is organised as follows. In sect. 2 we point out some differences between the way concepts are conceived in philosophy and in psychology. In sect. 3 we argue that AI research in some way shows traces of the contradictions individuated in sect. 2. In particular, the requirement of compositional, logical style semantics conflicts with the need of representing concepts in the terms of typical traits that allow for exceptions. In sect. 4 we review some attempts to resolve this conflict in the field of knowledge representation, with particular attention to description logics. In the conclusions (sect. 5) we sketch a possible way out, which is inspired by the so called *dual process theories* of human reasoning and rationality, according to which the existence of different types of reasoning systems is assumed. Indeed, it is our opinion that a mature methodology to approach knowledge representation and knowledge engineering should take advantage also from the empirical results of cognitive psychology that concern human abilities.

2 CONCEPTS IN PHILOSOPHY AND IN PSYCHOLOGY

Within the field of cognitive science, the notion of concept is highly disputed and problematic. Artificial intelligence (from now on AI) and, more in general, the computational approach to cognition reflect this state of affairs. Conceptual representation

Frixione M. and Lieto A..
THE COMPUTATIONAL REPRESENTATION OF CONCEPTS IN FORMAL ONTOLOGIES - Some General Considerations.
DOI: 10.5220/0003095903960403
In *Proceedings of the International Conference on Knowledge Engineering and Ontology Development* (KEOD-2010), pages 396-403
ISBN: 978-989-8425-29-4

seems to be constrained by conflicting requirements, such as, for example, compositionality on the one side and the need of representing prototypical information on the other.

A first problem (or, better, a first symptom that some problem exists) consists in the fact that the use of the term "concept" in the philosophical tradition is not homogeneous with the use of the same term in empirical psychology (see e.g. Dell'Anna and Frixione 2010). Briefly[1], we could say that in cognitive psychology a concept is essentially intended as the mental representations of a category, and the emphasis is on such processes as categorisation, induction and learning. According to philosophers, concepts are above all the components of thoughts. Even if we leave aside the problem of specifying what thoughts exactly are, this requires a more demanding notion of concept. In other words, some phenomena that are classified as "conceptual" by psychologists turn out to be "nonconceptual" for philosophers. There are, thus, mental representations of categories that philosophers would not consider genuine concepts. For example, according to many philosophers, concept possession involves the ability to make explicit, high level inferences, and sometimes also the ability to justify them (Peacocke 1992; Brandom 1994). This clearly exceeds the possession of the mere mental representation of categories. Moreover, according to some philosophers, concepts can be attributed only to agents who can use natural language (i.e., only adult human beings). On the other hand, a position that can be considered in some sense representative of an "extremist" version of the psychological attitude towards concepts is expressed by Lawrence Barsalou in an article symptomatically entitled "Continuity of the conceptual system across species" (Barsalou 2005). He refers to knowledge of scream situations in macaques, which involves different modality-specific systems (auditory, visual, affective systems, etc.). Barsalou interprets these data in favour of the thesis of a continuity of conceptual representations in different animal species, in particular between humans and non-human mammals: "this same basic architecture for representing knowledge is present in humans. [...] knowledge about a particular category is distributed across the modality-specific systems that process its properties" (p. 309). Therefore, according to Barsalou, a) we can speak of a "conceptual system" also in the case of non human animals; b) also low-level forms of categorisation, that depend on some specific perceptual modality pertain to the conceptual system. Elizabeth Spelke's experiments on infants (see e.g. Spelke 1994; Spelke and Kinzler 2007) are symptomatic of the difference in approach between psychologists and philosophers. Such experiments demonstrate that some extremely general categories are very precocious and presumably innate. According to the author, they show that newborn babies already possess certain *concepts* (e.g., the concept of physical object). But some philosophers interpreted these same data as a paradigmatic example of the existence of *nonconceptual* contents in agents (babies) that had not yet developed a conceptual system.

[1] Things are made more complex by the fact that also within the two fields considered separately this notion is used in a heterogeneous way, as we shall synthetically see in the following. As a consequence, the following characterisation of the philosophical and psychological points of view is highly schematic.

2.1 Compositionality

The fact that philosophers consider concepts mainly as the components of thoughts brought a great emphasis on *compositionality*, and on related features, such as productivity and systematicity, that are often ignored by psychological treatments of concepts (compositionality, productivity and systematicity of representations are defined in the following of this section). On the other hand, it is well known that compositionality is at odds with *prototypicality effects*, which are crucial in most psychological characterisations of concepts.

Let us consider first the *compositionality* requirement. In a compositional system of representations we can distinguish between a set of *primitive*, or *atomic symbols*, and a set of *complex symbols*. Complex symbols are generated starting from primitive symbols through the application of a set of suitable recursive syntactic rules (usually, starting from a finite set of primitive symbols, a potentially infinite set of complex symbols can be generated). Natural languages are the paradigmatic example of compositional systems: primitive symbols correspond to the elements of the lexicon (or, better, to morphemes), and complex symbols include the (potentially infinite) set of all sentences.

In compositional systems the meaning of a complex symbol *s* functionally depends on the syntactic structure of *s* and from the meaning of primitive symbols in it. In other words, the meaning of complex symbols can be determined by means of recursive semantic rules that work in parallel with syntactic composition rules. In this consists the so-called *principle of compositionality of meaning*, which Gottlob Frege identified as one of the main features of human natural languages.

In classical cognitive science it is often assumed that mental representations are compositional. One of the most clear and explicit formulation of this assumption is due to Jerry Fodor and Zenon Pylyshyn (1988). They claim that compositionality of mental representations is mandatory in order to explain some fundamental cognitive phenomena. In the first place, human cognition is *generative*: in spite of the fact that human mind is presumably finite, we can conceive and understand an unlimited number of thoughts that we never encountered before. Moreover, also *systematicity* of cognition seems to depend on compositionality: the ability of conceiving certain contents is related in a systematic way to the ability of conceiving other contents. For example, if somebody can understand the sentence *the cat chases a rat*, then she is presumably able to understand also *a rat chases the cat*, in virtue of the fact that the forms of the two sentences are syntactically related. We can conclude that the ability of understanding certain propositional contents systematically depends on the compositional structure of the contents themselves. This can be easily accounted for if we assume that mental representations have a structure similar to a compositional language.

2.2 Against "Classical" Concepts

Compositionality is less important for many psychologists. In the field of psychology, most research on concepts moves from the critiques to the so-called classical theory of concepts, i.e. the traditional point of view according to which concepts can be defined in terms of necessary and sufficient conditions. Rather, empirical evidence favours those approaches to concepts that accounts for prototypical effects. The central claim of the classical theory of concepts (i.e.) is that every concept *c* is defined in terms of a set of features (or conditions) $f_1, ..., f_n$ that are individually necessary and jointly sufficient for the application of *c*. In other words, everything that satisfies features $f_1, ..., f_n$ is a *c*, and if anything is a *c*, then it must satisfy $f_1, ..., f_n$. For example, the features that define the concept *bachelor* could be *human*, *male*, *adult* and *not married*; the conditions defining *square* could be *regular polygon* and *quadrilateral*. This point of view was unanimously and tacitly accepted by psychologists, philosophers and linguists until the middle of the 20th century.

The first critique to the classical theory is due to a philosopher: in a well known section from the *Philosophical Investigations*, Ludwig Wittgenstein observes that it is impossible to individuate a set of necessary and sufficient conditions to define a concept such as GAME (Wittgenstein, 1953, § 66). Therefore, concepts exist, which cannot be defined according to classical theory, i.e. in terms of necessary and sufficient conditions. Rather, concepts like GAME rest on a complex network of *family resemblances*. Wittgenstein introduces this notion in another passage in the *Investigations*: «I can think of no better expression to characterise these similarities than “family resemblances”; for the various resemblances between members of a family: build, features, colour of eyes, gait, temperament, etc. etc.» (*ibid*., § 67).

Wittgenstein's considerations were corroborated by empirical psychological research: starting from the seminal work by Eleanor Rosch, psychological experiments showed that common-sense concepts do not obey to the requirement of the classical theory[2]: usually common-sense concepts cannot be defined in terms of necessary and sufficient conditions (and even if for some concept such a definition is available, subjects do not use it in many cognitive tasks). Rather, concepts exhibit *prototypical effects*: some members of a category are considered better instances than others. For example, a robin is considered a better example of the category of birds than, say, a penguin or an ostrich. More central instances share certain typical features (e.g., the ability of flying for birds, having fur for mammals) that, in general, are not necessary neither sufficient conditions.

Prototypical effects are a well established empirical phenomenon. However, the characterisation of concepts in prototypical terms is difficult to reconcile with the requirement of compositionality. According to a well known argument by Jerry Fodor (1981), prototypes are not compositional (and, since concepts in Fodor's opinion must be compositional, concepts cannot be prototypes). In synthesis, Fodor's argument runs as follows: consider a concept like PET FISH. It results from the composition of the concept PET and of the concept FISH. But the prototype of PET FISH cannot result from the composition of the prototypes of PET and of FISH. For example, a typical PET is furry and warm, a typical FISH is greyish, but a

[2] On the empirical inadequacy of the classical theory and on the psychological theories of concepts see (Murphy 2002).

typical PET FISH is not furry and warm neither greyish.

Moreover, things are made more complex by the fact that, also within the two fields of philosophy and psychology considered separately, the situation is not very encouraging. In neither of the two disciplines does a clear, unambiguous and coherent notion of concept seem to emerge. Consider for example psychology. Different positions and theories on the nature of concepts are available (prototype view[3], exemplar view, theory theory), that can hardly be integrated. From this point of view the conclusions of Murphy (2002) are of great significance, since in many respects this book reflects the current status of empirical research on concepts. Murphy contrasts the approaches mentioned above in relation to different classes of problems, including learning, induction, lexical concepts and children's concepts. His conclusions are rather discouraging: the result of comparing the various approaches is that "there is no clear, dominant winner" (ibid., p. 488) and that "[i]n short, concepts are a mess" (p. 492). This situation persuaded some scholars to doubt whether concepts constitute a homogeneous phenomenon from the point of view of a science of the mind (see e.g. Machery 2005 and 2009; Frixione 2007).

3 CONCEPT REPRESENTATION IN AI

The situation sketched in the section above is in some sense reflected by the state of the art in AI and, more in general, in the field of computational modelling of cognition. This research area seems often to hesitate between different (and hardly compatible) points of view. In AI the representation of concepts is faced mainly within the field of knowledge representation (KR). Symbolic KR systems (KRs) are formalisms whose structure is, in a wide sense, language-like. This usually involves that KRs are assumed to be compositional.

In a first phase of their development (historically corresponding to the end of the 60s and to the 70s) many KRs oriented to conceptual representations tried to keep into account suggestions coming from psychological research. Examples are early semantic networks and frame systems. Frame and semantic networks were originally proposed as alternatives to the use of logic in KR. The notion of frame was developed by Marvin Minsky (1975) as a solution to the problem of representing structured knowledge in AI systems[4]. Both frames and most semantic networks allowed the possibility to characterise concepts in terms of prototypical information.

However, such early KRs where usually characterised in a rather rough and imprecise way. They lacked a clear formal definition, and the study of their meta-theoretical properties was almost impossible. When AI practitioners tried to provide a stronger formal foundation to concept oriented KRs, it turned out to be difficult to reconcile compositionality and prototypical representations. As a consequence, they often choose to sacrifice the latter.

In particular, this is the solution adopted in a class of concept-oriented KRs which had (and still have) wide diffusion within AI, namely the class of formalisms that stem from the so-called structured inheritance networks and from the KL-ONE system (Brachman and Schmolze 1985). Such systems were subsequently called terminological logics, and today are usually known as *description logics* (DLs) (Baader et al. 2002).

A standard inference mechanism for this kind of networks is *inheritance*. Representation of prototypical information in semantic networks usually takes the form of allowing exceptions to inheritance. Networks in this tradition do not admit exceptions to inheritance, and therefore do not allow the representation of prototypical information. Indeed, representations of exceptions can be hardly accommodated with other types of inference defined on these formalisms, concept classification in the first place (Brachman 1985). Since the representation of prototypical information is not allowed, inferential mechanisms defined on these networks (e.g. inheritance) can be traced back to classical logical inferences.

In more recent years, representation systems in this tradition have been directly formulated as logical formalisms (the above mentioned description logics, Baader et al., 2002), in which Tarskian, compositional semantics is straightly associated to the syntax of the language. Logical formalisms are paradigmatic examples of compositional representation systems. As a consequence, this kind

[3] Note that the so-called prototype view does not coincide with the acknowledgement of prototypical effects: as said before, prototypical effects are a well established phenomenon that all psychological theories of concepts are bound to explain; the prototype view is a particular attempt to explain empirical facts concerning concepts (including prototypical effects). On these aspects see again Murphy 2002.

[4] Many of the original articles describing these early KRs can be found in (Brachman & Levesque 1985), a collection of classical papers of the field.

of systems fully satisfy the requirement of compositionality. This has been achieved at the cost of not allowing exceptions to inheritance. By doing this we gave up the possibility of representing concepts in prototypical terms. From this point of view, such formalisms can be seen as a revival of the classical theory of concepts, in spite of its empirical inadequacy in dealing with most common-sense concepts.

Nowadays, DLs are widely adopted within many application fields, in particular within the field of the representation of ontologies. For example, the OWL (Web Ontology Language) system is a formalism in this tradition that has been endorsed by the World Wide Web Consortium for the development of the semantic web[5].

4 NON-CLASSICAL CONCEPTS IN COMPUTATIONAL ONTOLOGIES

Of course, within symbolic, logic oriented KR, rigorous approaches exist, that allow to represent exceptions, and that therefore would be, at least in principle, suitable for representing “non-classical” concepts. Examples are fuzzy logics and non monotonic formalisms. Therefore, the adoption of logic oriented semantics is not necessarily incompatible with prototypical effects. But such approaches pose various theoretical and practical difficulties, and many unsolved problems remain.

In this section we overview some recent proposal of extending concept-oriented KRs, and in particular DLs, in order to represent non-classical concepts.

Recently different methods and techniques have been adopted to represent non-classical concepts within computational ontologies. They are based on extensions of DLs and of standard ontology languages such as OWL. The different proposals that have been advanced can be grouped in three main classes: a) fuzzy approaches, b) probabilistic and Bayesan approaches, c) approaches based on non monotonic formalisms.

a) As far as the integration of *fuzzy logics* in DLs and in ontology oriented formalisms is concerned, see for example Gao and Liu 2005, and Calegari and Ciucci 2007. Stoilos et al. (2005) propose a fuzzy extension of OWL, f-OWL, able to capture imprecise and vague knowledge, and a fuzzy reasoning engine that lets f-OWL reason about such knowledge. Bobillo and Staccia (2009) propose a fuzzy extension of OWL 2 for representating vague information in semantic web languages. However, it is well known (Osherson and Smith 1981) that approaches to prototypical effects based on fuzzy logic encounter some difficulty with compositionality.

b) The literature offers also several *probabilistic generalizations* of web ontology languages. Many of these approaches, as pointed out in Lukasiewicz and Straccia (2008), focus on combining the OWL language with probabilistic formalisms based on Bayesian networks. In particular, Da Costa and Laskey (2006) suggest a probabilistic generalization of OWL, called PR-OWL, whose probabilistic semantics is based on multi-entity Bayesian networks (MEBNs); Ding et al. (2006) propose a probabilistic generalization of OWL, called Bayes-OWL, which is based on standard Bayesian networks. Bayes-OWL provides a set of rules and procedures for the direct translation of an OWL ontology into a Bayesian network. A general problem of these approaches could consist in avoiding arbitrariness in assigning weights in the translation from traditional to probabilistic formalisms.

c) The role of *non monotonic reasoning* in the context of formalisms for the ontologies is actually a debated problem. According to many KR researches, non monotonic logics are expected to play an important role for the improvement of the reasoning capabilities of ontologies and of the Semantic Web applications. In the field of non monotonic extensions of DLs, Baader and Hollunder (1995) propose an extension of ALCF system based on Reiter’s default logic[6]. The same authors, however, point out both the semantic and computational difficulties of this integration and, for this reason, propose a restricted semantics for open default theories, in which default rules are only applied to individuals explicitly represented in the knowledge base. Since Reiter’s default logic does not provide a direct way of modelling inheritance with exceptions in DLs, Straccia (1993) proposes an extension of H-logics (*Hybrid KL-ONE style logics*)

---

[5] The problem of representing of non-classical concepts in DLs is related with, but independent from, another theoretical and practical knowledge representation problem, namely the problem of representing instances in taxonomies. The notion of "instance" itself in taxonomic formalisms is problematic (for some aspects see e.g. Frixione et al. 1989). However, we do not face these problems here.

[6] The authors pointed out that “Reiter's default rule approach seems to fit well into the philosophy of terminological systems because most of them already provide their users with a form of ‘monotonic’ rules. These rules can be considered as special default rules where the justifications - which make the behavior of default rules nonmonotonic – are absent”.

able to perform *default inheritance reasoning* (a kind of default reasoning specifically oriented to reasoning on taxonomies). This proposal is based on the definition of a priority order between default rules. Donini et al. (1998, 2002), propose an extension of DL with two non monotonic epistemic operators. This extension allows one to encode Reiter's default logic as well as to express epistemic concepts and procedural rules. However, this extension presents a rather complicated semantics, so that the integration with the existing systems requires significant changes to the standard semantics of DLs. Bonatti et al. (2006) propose an extension of DLs with circumscription. One of motivating applications of circumscription is indeed to express prototypical properties with exceptions, and this is done by introducing "abnormality" predicates, whose extension is minimized. Giordano et al. (2007) propose an approach to defeasible inheritance based on the introduction in the *ALC* DL of a typicality operator **T**[7], which allows to reason about prototypical properties and inheritance with exceptions. This approach, given the non monotonic character of the **T** operator, encounters some problems in handling inheritance (an example is what the authors call the *problem of irrelevance*). Katz and Parsia argue that ALCK, a non monotonic DL extended with the epistemic operator K[8] (that can be applied to concepts or roles) could represent a model for a similar non monotonic extension of OWL. In fact, according to the authors, it would be possible to create "local" closed-world assumption conditions, in order the reap the benefits of non monotonicity without giving up OWL's open-world semantics in general.

A different approach, investigated by Klinov and Parsia (2008), is based on the use of the OWL 2 annotation properties (APs) in order to represent vague or prototypical, information. The limit of this approach is that APs are not taken into account by the reasoner, and therefore have no effect on the inferential behaviour of the system (Bobillo and Straccia 2009).

5 CONCLUSIONS: A "DUAL PROCESS" PROPOSAL

Though the presence of a relevant field of research, there is not, in the scientific community, a common view about the use of non monotonic and, more in general, non-classical logics in ontologies. For practical applications, systems that are based on classical Tarskian semantics and that do not allow for exceptions (as it is the case of "traditional" DLs), are usually still preferred. Some researchers, as, for example, Pat Hayes (2001), argue that the non monotonic logics (and, therefore, the non monotonic "machine" reasoning for Semantic Web) can be maybe adopted for local uses only or for specific applications because it is "unsafe on the web". Anyway, the question about which "logics" must be used in the Semantic Web (or, at least, until which degree, and in which cases, certain logics could be useful) is still open. At the same time, the empirical results from cognitive psychology show that most common-sense concepts cannot be characterised in terms of necessary/sufficient conditions. Classical, monotonic DLs seem to capture the compositional aspects of conceptual knowledge, but are inadequate to represent prototypical knowledge.

As seen before, cognitive research about concepts seems to suggest that concept representation does not constitute an 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 (Stanovich and West 2000, Evan 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, classification, a monotonic form of reasoning which is defined on semantic networks, and which is typical of DL systems, is a task of the type 2 (it is a difficult, slow, sequential task). 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 classification and in similar "difficult" tasks, and a non monotonic module involved in the *management of exceptions*. This last module should be a "weak" non monotonic system, able to perform

[7] For any concept *C*, **T**(*C*) are the instances of *C* that are considered as "typical" or "normal".

[8] The K operator could be encoded in RDF/XML syntax of OWL as property or as annotation property.

only some simple forms of non monotonic inferences (mainly related to categorization and to exceptions inheritance). This solution goes in the direction of a "dual" representation of concepts within the ontologies, and the realization of hybrid reasoning systems (monotonic and non monotonic) on semantic network knowledge bases. Some of the proposal reviewed in the sect. 4 above could be probably interpreted in this perspective. It could be objected that proposals based on non monotonic extensions of classical DLs are not suitable to model type 1 reasoning systems, since their computational properties are even worst of those of traditional, monotonic DLs. In this perspective, an alternative solution should be combining ontologies and logic programming rules, endowed with usual semantics for non monotonic logic programs (see e.g. Eiter et al. 2008).

REFERENCES

Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P., 2003. *The Description Logic Handbook: Theory, Implementations and Applications*. Cambridge University Press.

Baader, F., Hollunder, B., 1995. *Embedding defaults into terminological knowledge representation formalisms.* J. Autom. Reasoning 14(1), 149–180.

Barsalou, L. W., 1985. Continuity of the conceptual system across species. *Trends in Cognitive Science*, 9(7), 305-311.

Bobillo, F., Straccia, U., 2009. An OWL Ontology for Fuzzy OWL 2. Proceedings of the 18th *International Symposium on Methodologies for Intelligent Systems (ISMIS-09*). *Lecture Notes in Computer Science*. Springer Verlag.

Bonatti, P.A., Lutz, C., Wolter, F., 2006. *Description logics with circumscription*. In: Proc. of KR, pp. 400–410.

Brachman, R., 1985. I lied about the trees. *The AI Magazine*, 3(6), 80-95.

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

Brachman, R., Schmolze, J. G. (eds.), 1985. *Readings in Knowledge Representation*, Los Altos, CA: Morgan Kaufmann.

Brandom, R., 1994. *Making it Explicit*. Cambridge, MA: Harvard University Press.

Calegari, S., Ciucci, D., 2007. Fuzzy Ontology, Fuzzy Description Logics and Fuzzy-OWL. *Proc. WILF 2007*, Vol. 4578 of LNCS.

Da Costa P. C. G., Laskey, K. B., 2006. PR-OWL: A framework for probabilistic ontologies. *Proc. FOIS-2006*, 237–249.

Dell'Anna, A., Frixione, M., 2010. On the advantage (if any) and disadvantage of the conceptual/ nonconceptual distinction for cognitive science. *Minds & Machines*, 20, 29-45.

Ding, Z., Peng, Y., Pan, R.. 2006. BayesOWL: Uncertainty modeling in Semantic Web ontologies. In Z. Ma (ed.), *Soft Computing in Ontologies and Semantic Web*, vol. 204 of *Studies in Fuzziness and Soft Computing*, Springer.

Donini, F. M., Lenzerini, M., Nardi, D., Nutt, W., Schaerf, A., 1998. An epistemic operator for description logics. *Artificial Intelligence*, 100(1-2), 225–274.

Donini, F. M., Nardi, D., Rosati, R., 2002. Description logics of minimal knowledge and negation as failure. *ACM Trans. Comput. Log.*, 3(2), 177–225.

Eiter, T., Ianni, G., Luckasiewicz, T., schinklauer, R., Tompits, H. 2008. Combining answer set programming with description logics for the Semantic Web. *Artificial Intelligence*, 172, 1495-1539.

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

Fodor, J., 1981. The present status of the innateness controversy. In J. Fodor, *Representations*, Cambridge, MA: MIT Press.

Fodor, J., Pylyshyn, Z., 1988. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3-71.

Frixione, M., 2007. Do concepts exist? A naturalistic point of view. In C. Penco, M. Beaney, M. Vignolo (eds.). *Explaining the Mental*. Cambridge, UK: Cambridge Scholars Publishing.

Frixione, M., Gaglio, S., Spinelli, G. 1989. Are there individual concepts? Individual concepts and definite descriptions in SI-Nets. *International Journal of Man Machine Studies*, 30, 489-503.

Gao, M., Liu, C., 2005. Extending OWL by fuzzy Description Logic. Proc. 17th *IEEE Int. Conf. on Tools with Artificial Intelligence (ICTAI 2005)*, 562–567. IEEE Computer Society, Los Alamitos.

Giordano, L., Ghiozzi, V., Olivetti, N., Pozzato, G., 2007. Preferential Description Logics. Logic for Programming, In *Artificial Intelligence, and Reasoning*, *LNCS*, Springer Verlag.

Hayes, P., 2001. Dialogue on rdf-logic. Why must the web be monotonic?. World Wide Web Consortium (W3C). Link: http://lists.w3.org/Archives/public/www-rdf-logic/2001Jul/0067.html

Katz, Y., Parsia, B., 2005. Towards a non monotonic extension to OWL. *Proc. OWL Experiences and Directions*, Galway, November 11-12.

Klinov, P., Parsia, B., 2008. Optimization and evaluation of reasoning in probabilistic description logic: Towards a systematic approach. In: Sheth, A.P., Staab, S., Dean, M., Paolucci, M., Maynard, D., Finin, T., Thirunarayan, K. (eds.) *ISWC 2008*. LNCS, vol. 5318, 213–228. Springer, Heidelberg.

Lukasiewicz, L., Straccia, U., 2008. Managing uncertainty and vagueness in description logics for the Semantic Web**.** *Journal of Web Semantics*, 6, 291-308.

Machery, E., 2005. Concepts are not a natural kind. *Philosophy of Science*, 72, 444–467.
Machery, E., 2009. *Doing without Concepts*. Oxford, UK: Oxford University Press.
Minsky, M., 1975. A framework for representing knowledge, in Patrick Winston (a cura di), *The Psychology of Computer Vision*, New York, McGraw-Hill. Also in Brachman & Levesque (2005).
Murphy, G. L., 2002. *The Big Book of Concepts*. Cambridge, MA: The MIT Press.
Osherson, D. N., Smith, E. E., 1981. On the adequacy of prototype theory as a theory of concepts. *Cognition*, 11, 237-262.
Peacocke, C., 1992. *A Study of Concepts*. Cambridge, MA: The MIT Press.
Rosch, E., 1975. Cognitive representation of semantic categories. *Journal of Experimental Psychology*, 104, 573-605.
Spelke, E. S., 1994. Initial knowledge: six suggestions. *Cognition*, 50, 431-445.
Spelke, E. S., Kinzler, K.D., 2007. Core knowledge. *Developmental Science*, 10(1), 89–96.
Stanovich, K. E., West, R., 2000. Individual Differences in Reasoning: Implications for the Rationality Debate?. *The Behavioural and Brain Sciences* 23 (5), 645- 665.
Stoilos, G., Stamou, G., Tzouvaras, V., Pan, J.Z., Horrocks, I., 2005. Fuzzy OWL: Uncertainty and the Semantic Web. *Proc. Workshop on OWL: Experience and Directions (OWLED 2005)*. CEUR Workshop Proceedings, vol. 188.
Straccia, U., 1993. Default inheritance reasoning in hybrid kl-one-style logics, *Proc. IJCAI*, 676–681.
Wittgenstein, L., 1953. *Philosophische Untersuchungen*. Oxford, Blackwell.