The Ontological Level

Nicola Guarino

1. Introduction

In 1979, Ron Brachman discussed a classification of the various primitives used by KR systems at that time[1]. He argued that they could be grouped in four levels, ranging from the *implementational* to the *linguistic level* (Fig. 1). Each level corresponds to an explicit set of primitives offered to the knowledge engineer. At the *implementational* level, primitives are merely pointers and memory cells, which allow us to construct data structures with no *a priori* semantics. At the *logical* level, primitives are propositions, predicates, logical functions and operators, which are given a formal semantics in terms of relations among objects in the real world. No particular assumption is made however as to the nature of such relations: classical predicate logic is a general, uniform, neutral formalism, and the user is free to adapt it to its own representation purposes. At the *conceptual* level, primitives have a definite cognitive interpretation, corresponding to language-independent concepts like elementary actions or thematic roles. Finally, primitives at the *linguistic* level are associated directly to nouns and verbs.

Brachman noticed an evident gap in this classification: while primitives at the logical level are extremely general and content-independent, at the conceptual level they acquire a specific intended meaning that must be taken as a whole, without any account of its internal structure. He proposed the introduction of an intermediate *epistemological level,* where the primitives allow us to specify "the formal structure of conceptual units and their interrelationships *as conceptual units* (independent of any knowledge expressed therein)"[2]. In other words, while the logical level deals with abstract predicates and the conceptual level with *specific* concepts, at the epistemological level the *generic* notion of a concept is introduced as a *knowledge structuring primitive.*

Level	Primitives
Implementational	Memory cells, pointers
Logical	Propositions, predicates, functions, logical operators
Epistemological	Concept types, structuring relations
Conceptual	Conceptual relations, primitive objects and actions
Linguistic	Linguistic terms

Fig. 1. Classification of primitives used in KR formalisms (adapted from Brachman 79). Epistemological level was "the missing level".

Brachman's KL-ONE[3] is an example of a formalism built around these notions. Its main contribution was to give an epistemological foundation to cognitive structures like frames and semantic networks, whose formal contradictions had been revealed in the famous "What's in a link?" paper by Bill Woods[4]. Brachman's answer to Woods' question was that *conceptual links* should be accounted for by *epistemological links*, which represent the structural connections in our knowledge needed to justify conceptual inferences. KL-ONE focused in

Invited paper presented at the 16th Wittgenstein Symposium, Kirchberg, Austria, August 1993. To appear in R. Casati, B. Smith and G. White (eds.), *Philosophy and the Cognitive Sciences*, Vienna: Hölder-Pichler-Tempsky, 1994.

particular on the inferences related to the so-called IS-A relationship, offering primitives to describe the (minimal) formal structure of a concept needed to guarantee “formal inferences about the relationship (subsumption) between a concept and another”. This formal structure consists of the sum of the constituents of a concept (primitive concepts and role expressions) and the constraints among them, independently of any commitment as to: (i) the meaning of primitive concepts; (ii) the meaning of roles themselves; (iii) the nature of each role’s contribution to the meaning of the concept. The intended *meaning* of concepts remains therefore totally arbitrary: indeed, the semantics of current descendants of KL-ONE is such that – at the logical level – concepts correspond to *arbitrary* monadic predicates, while roles are *arbitrary* binary relations. In other words, at the epistemological level, emphasis is more on *formal reasoning* than on (formal) *representation:* the very task of representation, i.e. the structuring of a domain, is left to the user.

Current frame-based or object-oriented formalisms suffer from the same problem. For example, the advantage of a frame-based language over pure first-order logic is that some logical relations, such as those corresponding to classes and slots, have a peculiar, *structuring* meaning. This meaning is the result of a number of ontological commitments, which accumulate in layers starting from the very beginning of the process of developing a knowledge base[5]. For a particular knowledge base, its ontological commitments are however implicit and strongly dependent on the particular task being considered, since the formalism itself is in general *neutral* as concerns ontological choices[6].

In this paper I argue *against* this neutrality, claiming that a rigorous ontological foundation for knowledge representation can improve the quality of the knowledge engineering process, making it easier to build at least understandable (if not reusable) knowledge bases. We contrast the notion of formal ontology, intended as a theory of the *a priori* forms and natures of objects, to that of (formal) epistemology, intended as a theory of meaning connections[7]. We show in the following how theories defined at the epistemological level, based on structured representation languages like KL-ONE, cannot be distinguished from their "flat" first-order logic equivalents unless we make clear their implicit ontological assumptions by stating formally what it means to interpret a unary predicate as a concept (class) and a binary predicate as a "role" (slot). We need therefore to introduce the notion of *ontological level*, as an intermediate level between the epistemological and the conceptual one (Fig. 6)[8]. While the epistemological level is the level of structure, the ontological level is the level of *meaning*. At the ontological level, knowledge primitives satisfy formal meaning postulates, which restrict the interpretation of a logical theory on the basis of formal ontology, intended as a theory of *a priori distinctions:* [9]

- among the entities of the world (physical objects, events, processes...);
- among the meta-level categories used to model the world (concepts, properties, states, roles, attributes, various kinds of part-of relations...).

We focus here on the latter kind of distinctions, showing how the basic dichotomy existing in KR systems between *concepts* like *apple* and *assertional properties* like *red* can be understood in terms of the philosophical distinction between *sortal* and *characterising* universals[10]. In section 2 I present examples which show the necessity of making such a distinction explicit. In section 3 I introduce the notion of ontological commitment as a constrained interpretation of a logical theory, and I sketch a basic ontology of meta-level categories of unary predicates. In section 4 I discuss the role of the ontological level in current knowledge engineering practice.

2. Reds and apples.

Suppose we want to state that a red apple exists. In standard first order logic, it is straightforward to write down something like $\exists x.(Ax \wedge Rx)$. If we want however to impose some *structure* on our domain, the simplest formalism we may resort to is many-sorted logic. Yet we have to decide which predicates correspond to sorts: we may write $\exists x{:}A.Rx$ as well as $\exists x{:}R.Ax$ (or maybe $\exists(x{:}A,y{:}R).x=y$). All these structured formalisations are equivalent to the previous one-sorted axiom, but each contains an implicit structuring choice. At the epistemological level, this choice is up to the user, since the semantics of the primitive "sort" is the same as its corresponding first-order predicate. At the *ontological level*, what we want is a formal, restricted semantic account that reflects the ontological commitment intrinsic in the use of a given predicate as a sort. This means that the choice of a particular axiomatisation is still up to the user, but its *consequences* are formalised in such a way that another user can understand the meaning of the choice itself, and possibly *agree* on it on the basis of its semantics.

In our case, a statement like $\exists x{:}R.Ax$ sounds intuitively odd: what are we quantifying over? Do we assume the existence of "instances of redness" that can have the property of being apples? According to Strawson, the difference between the two predicates lies in the fact that *apple* "supplies a principle for distinguishing and counting individual particulars which it collects", while *red* "supplies such principle only for particulars already distinguished, or distinguishable, in accordance with some antecedent principle or method"[11]. This distinction is known in the philosophical literature as the distinction between *sortal* and non-sortal (*characterising*) universals, and is (roughly) reflected in natural language by the fact that the former are common nouns, while the latter are adjectives. The issue is also related to the difference between count and mass terms, and has been a matter of lively debate among linguists and philosophers[12]. The distinction is implicitly present in the KR literature, where sortal universals are usually called "concepts", while characterising universals are called "properties". The difference between the two is however the result of heuristic considerations, and nothing in the semantics of a concept forbids any arbitrary unary predicate from acquiring this status.

Our position is that, within a KR formalism, the meaning of structuring primitives as sorts (or concepts, in KR terminology) should be at least specified with formal, necessary conditions at the meta-level, which force the user to accept their consequences when he/she decides to use a given predicate as a sort. According to our previous discussion, a predicate like *red* – under its ordinary meaning – will not satisfy such conditions, and should be excluded therefore from being used as a sort. Notice however that this may be simply a matter of point of view: at the ontological level, it is still the user who decides which conditions reflect the *intended* use of the predicate *red*: a more rigid choice would be distinctive of higher levels, like the conceptual or the linguistic level. For example, compare the statement mentioned above with others where the same unary predicate *red* appears in different contexts (Fig. 2):

In case (2), *red* is still a unary predicate whose argument refers to a particular colour instead of a particular fruit; in (3) the argument refers to a particular colour gradation belonging to the set of "reds", while in (4) the argument refers to a human-being, meaning for instance that he/she is a communist.

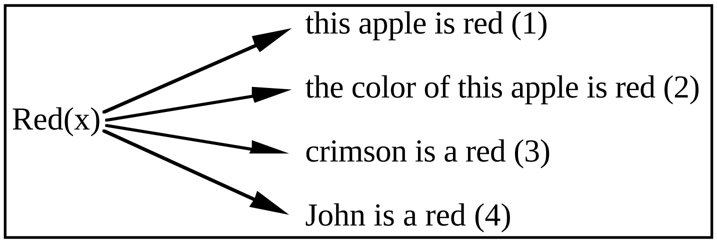

Fig. 2. Varieties of predication.

We face here the difference of positions between the stereotype Linguist and Philosopher discussed by Bill Woods in one of the historical papers on knowledge representation[13]: while the Linguist "is interested in characterising the fact that the same sentence can sometimes mean different things", the Philosopher "is concerned with specifying the meaning of a formal notation rather than a natural language". Woods goes on by stating that

> philosophers have generally stopped short of trying to actually specify the truth conditions of the basic atomic propositions, dealing mainly with the specification of the meaning of complex expressions in terms of the meanings of elementary ones. Researchers in artificial intelligence are faced with the need to specify the semantics of elementary propositions as well as complex ones.

In a knowledge representation formalism, we are constantly using natural language expressions within our formulas, relying on them to make our statements readable and to convey meanings we have not explicitly stated: however, since words are ambiguous in natural language, it may be important to "tag" these words with a semantic category, endowed with a suitable axiomatisation, in order to guarantee a consistent interpretation. This is unavoidable, in our opinion, if we want to share theories across different domains[14].

How can we account for the semantic differences in the use of *red* in the formulas given above? In our opinion, they are not simply related to the fact that the argument belongs to different domains: they are mainly due to different *ways of predication*, i.e. different subject-predicate relationships. Studying the formal properties of these relationships is a matter of formal ontology.

3. A basic ontology of unary predicate types

A basic ontology, which – according to Strawson's intuitions – classifies unary predicates on the basis of their ability to supply an identification principle for their arguments, is presented in Fig. 3.

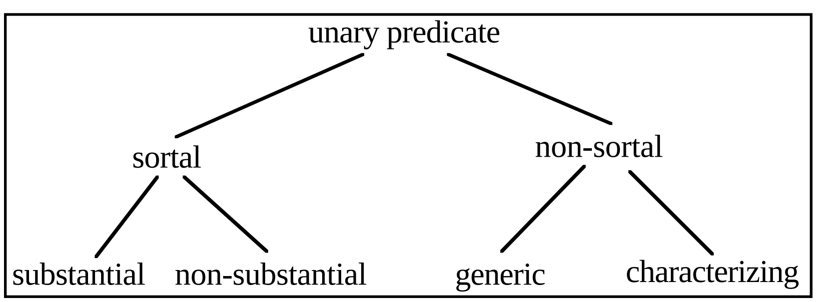

Fig. 3. A basic ontology of predicate types.

Following Wiggins[15], sortal predicates are here divided into substantial (like *apple* or *human-being*) and non-substantial (like *food* or *student*), while non-sortal predicates include *generic predicates* like *thing* and *characterising predicates* like *red*. A preliminary formalisation of such distinctions will be presented in the sequel.

But of trying to give a “universal” formal definition of the above categories, we shall pursue here a more modest account: our definitions will be related to a specific knowledge base described by a standard first order theory, which we are interested in “adding structure” to. This means that the basic knowledge-building blocks are taken as having been already fixed, being the predicates of the theory itself; our work will be to offer a formal instrument for clarifying their ontological implications for the specific purposes of knowledge bases understanding and reuse among users belonging to the same culture. We assume therefore that interpretations of our specific theory, rather than describing a real or hypothetical situation in a world that has the same laws of nature of ours[16], are states of affairs having an “idealised rational acceptability”[17]. This choice excludes unwanted metaphysical implications.

Within this framework, let us concentrate on a minimal problem. Suppose we have a first order, non-functional language L with signature $\Sigma=<C, R>$, where C is a set of constant symbols, R is a finite set of predicate symbols and $P \subseteq R$ is a set of monadic predicate symbols. We are interested in some formal criteria for translating L into an order-sorted language L_s with signature $\Sigma_s=<C, S, Q>$, where $S \subseteq P$ is a set of sort symbols called *sortals*, and $Q=R \setminus S$ is a set of ordinary predicates.

Def. 1. Let L_m be the *modal extension* of a first order language L obtained by adding to L the usual modal operators ■ and ◆, and *D* be a set. A *rigid model* [18] for L_m based on D is a structure $M=<W, \mathfrak{R}, D, \mathfrak{F}_C, \mathfrak{F}_P>$, where W is a set of possible worlds sharing the same interpretation $\mathfrak{F}_C$ for constant symbols of L_m, $\mathfrak{R}$ is a binary relation on W, D is a domain common to all possible worlds, and $\mathfrak{F}_P$ is a mapping that assigns to each predicate symbol P of L and world $w \in W$ a unary relation on D.

For a given rigid model M, $\mathfrak{R}$ is a relation between worlds (i.e., interpretations of L) that may differ in the interpretation of predicates while sharing the same interpretation for constants. We want to give $\mathfrak{R}$ the meaning of an *ontological compatibility* relation: two worlds are ontologically compatible if they describe plausible alternative states of affairs involving the same elements of the domain. $\mathfrak{R}$ will be in this case reflexive, transitive and symmetric (i.e., an equivalence relation), and the corresponding modal theory will be S5.

Def. 2. Let L be a first order language and D a domain. An *ontological commitment* for L based on D is a set **C** of rigid models for L_m based on D, where the relation $\mathfrak{R}$ is an equivalence relation. Such commitment can be specified by an S5 modal theory of L_m, being in this case the set of all its rigid models based on D. A formula Φ of L_m is valid under **C** (**C** $\models \Phi$) if it is valid in each model $M \in$ **C**.

Within this modal framework, preliminary reflection on the distinction between sortal and non-sortal predicates reveals that the former cannot be necessarily false for each element of the domain: they must be *natural predicates*, in the sense of the following definition[19]:

Def. 3. Let L be a first order language, P a monadic predicate of L, and **C** an ontological commitment for L. P is called *natural* under **C** iff **C** $\models \exists x.\blacklozenge Px$.

A more substantial observation that comes to mind when trying to formalise the nature of the subject-predicate relationship in the examples above, is that the “force” of this relationship is much higher in “x is an apple” than in “x is red”. If x has the property of being an apple, it cannot lose this property without losing its identity, while this does not seem to be the case in the second example. This observation goes back to Aristotelian essentialism, and can be easily formalised as follows[20]:

Def. 4 A predicate P is *ontologically rigid* under **C** iff it is natural under **C** and **C** $\models \forall x.(Px \supset \blacksquare Px)$.

Ontological rigidity seems a useful property for characterising sortals: stating that *apple* is rigid and *red* is not will clarify the intended meaning of these two predicates in the statement (1) of Fig. 2. In this case, if $a \in C$, the worlds satisfying $(A(a) \wedge R(a))$ or $(A(a) \wedge \neg R(a))$ will turn out to be mutually compatible, while those satisfying $(R(a) \wedge A(a))$ or $(R(a) \wedge \neg A(a))$ are not (due to the constraints imposed on $\mathfrak{R}$ by the rigidity of A). Assuming that rigidity is a necessary property for sortals, we can then exclude both $\exists x{:}R.Ax$ and $\exists(x{:}A,y{:}R).x=y$ from our axiomatisation choices for (1).

Notice that the naturalness condition in the above definition excludes cases where rigidity would be trivially true due to the impossibility of P. On the other hand, ontological rigidity will be trivially satisfied by predicates being necessarily true for each element of the domain, like *thing* or *entity*. Yet, according to traditional wisdom they are excluded from being sortals, since no clear distinction criteria are associated with them. Rigidity cannot be therefore be considered as a necessary condition for sortals. We call these “top level” predicates *generic predicates* [21]. In the same category other rigid predicates should be included, that, although being not trivially rigid, are still too general to supply a distinction criterion: *object*, *individual*, *event*... However, a distinctive characteristic of generic predicates is that they are rigid but *divisive*, in the sense that they can hold for parts of their arguments. Various divisivity criteria have been proposed in the literature in order to account for the distinction between countable and uncountable predicates[22]; to the purposes of the present paper, the following definition will be good enough:

Def. 5. Let P be a natural predicate under **C**, and < be a “proper part” relation assumed as primitive, satisfying the axioms of classical mereology[23]. P is *divisive* under **C** iff **C** $\models \exists x.\blacklozenge(Px \wedge \exists y.\ y<x) \wedge \forall x.\blacksquare(Px \supset (\exists y.(y<x \supset Py)))$.

In other words, a predicate P is divisive if its arguments can have proper parts, and, necessarily, if its instances have proper parts then P holds for one of these. We are now in a position to give a definition of substantial sortals:

Def. 6. Let P be ontologically rigid under **C**. It is a *generic predicate* in **C** if it is divisive in **C**, and a *substantial sortal* in **C** otherwise.

Within our KR framework, the above definition gives a formal characterisation to the notion of substantial sortals originally introduced by Wiggins, delimiting those rigid predicates that are sortals. We need now a distinction criterion between non-rigid predicates:

some of them (like *student*) will presumably be *non-substantial sortals*, while the others (like *red*) will be *characterising predicates*. The intuition behind the distinction between substantial and non-substantial sortals is that in the first case the identity criterion is given by the predicate itself, while in the second case it is provided by some superordinate sortal. We formalise this idea as follows:

Def. 7. Let P be a natural predicate which is not ontologically rigid under **C**. It is a *non-substantial sortal* in **C** iff there exists a substantial sortal S in **C** such that **C** $\models \forall x.\blacksquare(Px \supset Sx)$, and a *characterising predicate* otherwise.

Since the set of predicate symbols of L is finite and fixed for all possible ontological commitments of L, this definition does not imply any “real” second-order quantification in L_m, nor does it have any metaphysical implication. For example, suppose that *student* is a non-rigid predicate under some commitment **C** for L. If *human-being* also belongs to L and is a superordinate substantial sortal under **C**, then *student* will be a non-substantial sortal, while otherwise it may simply be a characterising predicate. This means that *a priori* considerations about the real world do not affect our definitions unless they force the user to revise the original first-order axiomatisation. However, one of the advantages of the ontological level is that an unwanted formal property for a predicate may trigger a knowledge elicitation process: in our case, if *student* sounds strange when used as a characterising predicate, the reason may be that we have forgotten to include *human-being* within our axiomatisation.

Adapting some definitions from Cocchiarella[24], we believe it is important, for knowledge representation purposes, to make some further assumptions on sortals, which characterise what we call a *well-founded* ontological commitment:

Def. 8. An ontological commitment **C** based on D is *well-founded* iff:

- each element of D belongs to a substantial sortal;
- if two substantial sortals are not in the subsumption relationship, then they are mutually disjoint.

From Def. 7 and 8 it follows that:

Theorem 1. Two overlapping non-substantial sortals are subordinate to the same substantial sortal.

Def. 9. Let **C** be a well-founded ontological commitment. If a substantial sortal S is subordinate to another substantial sortal T under **C**, then S is called a *kind* of T.

Def. 10. A predicate is called a *sortal* under a commitment **C** if it is either a substantial or a non-substantial sortal under **C**.

As a final comment concerning the taxonomy of unary predicates we have discussed in this section, we would like to make the following proposal regarding the relationship between the terminology currently used in KR formalisms and the philosophical terms we have defined here (Fig. 4):

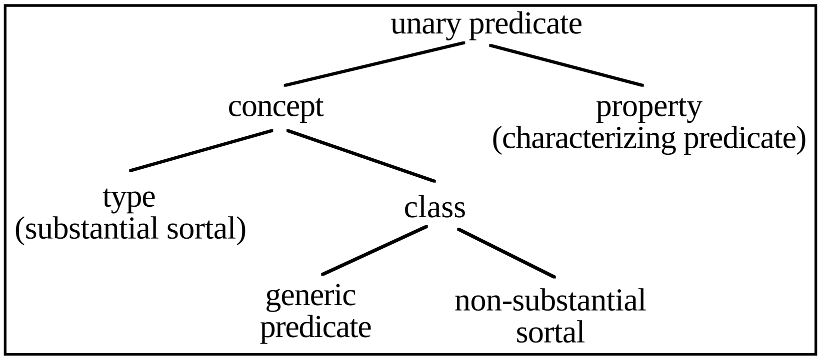

Fig. 4. A terminology proposal for KR formalisms

4. The Ontological Level

Just as the logical and epistemological levels are characterised by a (standard) formal semantics, so the ontological level is characterised by a formal ontological account, like the one introduced in the previous section. Although we have limited ourselves to a few very basic ontological distinctions, it should be clear that other important distinctions could be employed within a similar framework, like for instance those between attributes and arbitrary binary relations. However, the definitions we have given are enough to capture the different uses of the *red* predicate shown in Fig. 2, which correspond to distinct ontological commitments (Fig. 5).

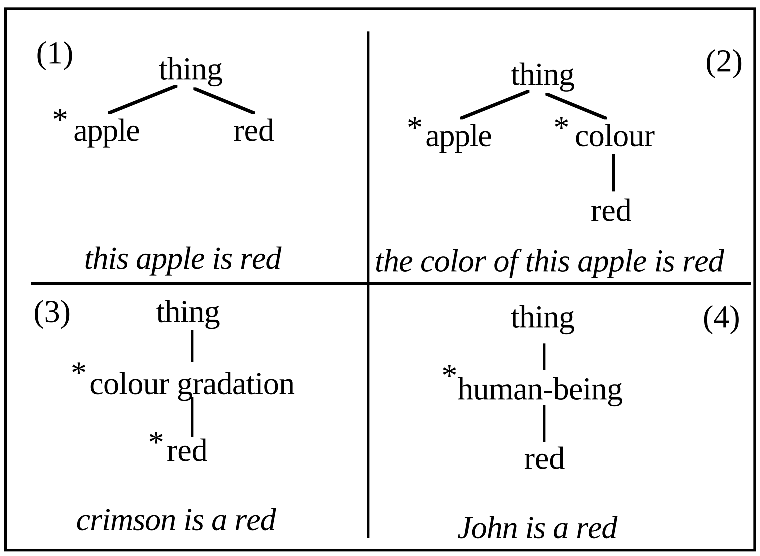

Fig. 5. Different ontological commitments capturing varieties of predication.
Arcs represent subsumption relationships, and asterisks mark substantial sortals.

In case (1), *red* is not rigid, and it has no superordinate substantial sortal: it is a characterising predicate, having as argument a physical object. In case (2), *red* is still not rigid, but it is subordinate to colour, which is assumed to be rigid and not divisive (a colour has no parts): it is a non-substantial sortal, having as argument the colour of a physical object. In case (3), *red* is rigid, since its argument is a colour gradation (crimson *has* to be a red): it is a substantial sortal, and also a kind of colour gradation. Finally, in case (4), *red* is used as a contingent property of human-beings and hence is not rigid, but it is not a characterising predicate since it is assumed that being a red implies being a human-being: *red* is therefore a non-substantial sortal like in (2), under a different ontological commitment.

It is important to stress that, although the notion of ontological commitment we have defined is bound to a quantified modal logic, the computational (bad) properties of such a theory have nothing to do with those of the first order language we started with. Even with a language of very limited expressiveness like a description logic[25], we can embed it in a full quantified modal logic, and use this to define the ontological commitment of the original language. This means that we give up the task of performing any automatic deduction on the modal theory, since we are only interested in its semantic properties. However, given a KR formalism at the epistemological level, we may be interested in somehow expressing its ontological commitment *within the formalism itself*. In other words, this is a matter of *ontological adequacy* of a KR language. We can get this ontological adequacy by suitably restricting the semantics of the epistemological level primitives (assuming for instance that “concepts” used in description logics have the semantics of sortals), or otherwise by having a syntactic way to “tag” a predicate symbol with an ontological category (stating for instance that *human-being* is a 'substantial sortal', where the latter is a primitive symbol). Some meta-level capability is necessary in the second case. In conclusion, while the ontological level is neutral with respect to the underlying epistemological level, not any epistemological level formalism may be considered as ontologically adequate.

Level	Primitive concepts...	Main feature	Interpretation
Logical	are predicates	Formalisation	Arbitrary
Epistemological	are structuring primitives	Structure	Arbitrary
Ontological	*satisfy meaning postulates*	*Meaning*	*Constrained*
Conceptual	are cognitive primitives	Conceptualisation	Subjective
Linguistic	are linguistic primitives	Language	Subjective

Fig. 6. Main features of the ontological level.

The main features of the ontological level are compared in Fig. 6 to that of other levels. The ontological level is the only level where the intended meaning of a KR language is constrained in a formal way. Lower levels have an arbitrary interpretation, since a logical theory admits a number of models much higher than the intended ones; higher levels (which can still be “implemented” at the ontological level) have a subjective interpretation, which can however be the refinement of a formal interpretation already constrained at the ontological level.

5. Conclusions

In Brachman, Fikes et al. 1983, the authors discussed the example reported in Fig. 7 below. They argued that a question like “How many kinds of rocks are there?” cannot be answered by simply looking at the nodes subsumed by ‘rock’ in the network, since the language allows them to proliferate easily. Hence they give up answering such dangerous questions within a KR formalism, by specifying a functional interface designed to answer “safe” queries about analytical relationships between terms independently of the structure of the knowledge base, like “a large grey igneous rock is a grey rock”. On the other hand, the same authors, in an earlier paper[26], stressed the importance of *terminological competence* in knowledge representation, stating for instance that an *"enhancement mode transistor"* (which is "a *kind* of transistor") should be understood as different from a *"pass transistor"* (which is "a *role* a transistor plays in a larger circuit").

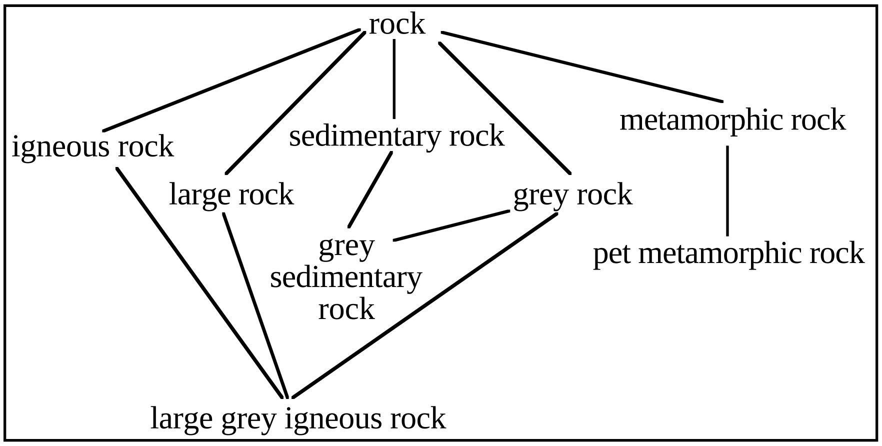

Fig. 7. Kinds of rocks (From Brachman, Fikes et al. 1983)

We hope to have shown in this paper that – in the spirit of Woods' statement cited in section 3 – terminological competence can be gained by formally expressing the ontological commitment of a knowledge base. If, in the example above, predicates corresponding to *rock*, *igneous-rock*, *sedimentary-rock* and *metamorphic-rock* are marked as substantial sortals (as they should be according to their ordinary meaning), while all the others are marked as non-substantial sorts (since they are not rigid), then a safe answer to the query "how many kinds of rocks are there?" would be "at least 3".

It is important to make clear however that the complete formal characterisation of the taxonomy described in section 3 – and the taxonomy itself – are still a matter of discussion. In particular, the notion of divisiveness is still problematic[27], since we may assume for instance that there are parts of an igneous rock which are still rocks, invalidating the example above (but not its spirit). Our answer to this objection is that in the notion of a rock, and of a physical object in general, there is implicit a notion of external boundary[28], such that an *undetached* part of a rock is not a rock, but just a part of it: this is why we can answer a question like "how many rocks are there?", while it is difficult to answer "how many parts of rock are there?". A more thoroughly account of the basic ontology sketched in section 3 will be the subject of a forthcoming paper; the preliminary distinction criteria introduced here have in our opinion the advantage of simplicity, avoiding the need to make use of subtle notions like ontological foundation, discussed in previous works[29].

The title chosen for this paper should however suggest to the reader that the particular ontology of unary predicates we have proposed is not the main issue here. Rather, we believe that the main contribution of section 3 is the notion of ontological commitment expressed in terms of a modal framework: the use of a modal logic, as a tool to constrain the intended semantics of the underlying non-modal theory, seems to be unavoidable if we wish to express ontological constraints unambiguously. In the perspective of formal ontology mentioned in the introduction, these constraints should also be related to *a priori* distinctions among the entities of the world, while we have limited ourselves to meta-level categories. Far from claiming to have said the last word on the latter issue, we tried to show here that (i) *some* formal properties which account for distinctions among predicate types can indeed be worked out, although complete, unproblematic definitions may not be given; (ii) when the semantics of structuring primitives used in KR languages is restricted in order to take into account such formal distinctions at the ontological level, the potential misunderstandings and inconsistencies due to conflicting intended models are reduced; (iii) further research in this area is needed, and it should be encouraged within the KR community, in co-operation with the philosophical and linguistic communities.

Acknowledgements

I am especially indebted to Massimiliano Carrara and Pierdaniele Giaretta for having introduced me to the philosophical distinctions concerning universals. I also thank Luca Boldrin, Matteo Cristani, Barry Smith and Claudio Sossai for their valuable comments on earlier versions of this paper.

Notes

1 Brachman 1979.
2 Brachman 1979, p. 30.
3 Brachman 1979, Brachman and Schmolze 1985.
4 Woods 1975.
5 See for instance Davis, Shrobe et al. 1993.
6 See Genesereth and Nilsson 1987.
7 The use of the term “epistemology” sounds somehow reductive here, but I believe this reflects its common understanding in the KR literature.
8 A preliminary ontological analysis of the primitives used in KL-ONE-like formalisms appeared in Guarino 1992, while the notion of ontological level has been first introduced in Guarino and Boldrin 1993a and Guarino and Boldrin 1993b.
9 Formal ontology has been recently defined as "the systematic, formal, axiomatic development of the logic of all forms and modes of being" Cocchiarella 1991.
10 See Strawson 1959, Wiggins 1980.
11 Strawson 1959 p. 168.
12 See Pelletier and Schubert 1989 for an overview.
13 Woods 1975. For a philosophical position different from the one criticized by Woods, see Mulligan *et al.* 1984.
14 Gruber 1993.
15 Wiggins 1980; Carrara 1992.
16 Cocchiarella 1993.
17 Putnam 1981, cited in Aune 1991, p. 543.
18 We use here the terminology introduced in Fitting 1993.
19 Cocchiarella 1993.
20 Barcan Marcus 1968.
21 In Pelletier and Schubert 1989 they are called “super sortals”.
22 Regarding the criticisms made for instance in Pelletier and Schubert 1989, see the comments in the conclusions. See also Guarino *et al.* 1993 for a finer account of this and other ontological distinctions.
23 See for instance Simons 1987. We assume here that L_m and C are suitably extended to nclude $<$.
24 Cocchiarella 1993, Cocchiarella 1977.
25 See a brief review in Woods and Schmolze 1992.
26 Brachman and Levesque 1982.
27 Pelletier and Schubert 1989.
27 Smith 1992.
27 Guarino and Boldrin 1993a.

References

Aune, B. 9. 1991. “Metaphysics of Analytic Philosophy”. In H. Burkhardt and B. Smith (Ed.), *Handbook of Metaphysics and Ontology*. Munich: Philosophia, 1991.

Barcan Marcus, R. 1968. “Essential Attribution”, *The Journal of Philosophy* 7.

Brachman, R., R. Fikes, et al. 1983. “Krypton: A Functional Approach to Knowledge Representation”, *IEEE Computer* (October), 67-73.

Brachman, R. and H. Levesque 1982. “Competence in Knowledge Representation”, *AAAI 82*.

Brachman, R. J. 1979. “On the Epistemological Status of Semantic Networks”. In N. V. Findler (Ed.), *Associative Networks: Representation and Use of Knowledge by Computers*. Academic Press, 1979.

Brachman, R. J. and J. G. Schmolze 1985. “An Overview of the KL-ONE Knowledge Representation System”, *Cognitive Science* 9, 171-216.

Carrara, M. 1992. Identità e persona nella riflessione filosofica di David Wiggins. Graduation thesis, Faculty of Philosophy, University of Padova.

Cocchiarella, N. 1977. “Sortals, Natural Kinds and Re-identification”, *Logique et Analyse* 80.

Cocchiarella, N. B. 1991. “Formal Ontology”. In H. Burkhardt and B. Smith (Ed.), *Handbook of Metaphysics and Ontology*. Munich: Philosophia Verlag, 1991.

Cocchiarella, N. B. 1993. “Knowledge Representation in Conceptual Realism”, *International Workshop on Formal Ontology in Conceptual Analysis and Knowledge Representation*, Padova, Italy, LADSEB-CNR Int. Rep. 01/93.

Davis, R., H. Shrobe, et al. 1993. “What is in a Knowledge Representation?”, *AI Magazine* (Spring 1993).

Fitting, M. 1993. “Basic Modal Logic”. In D. M. Gabbay, C. J. Hogger and J. A. Robinson (Ed.), *Handbook of Logic in Artificial Intelligence and Logic Programming*. Oxford: Clarendon Press, 1993.

Genesereth, M. R. and N. J. Nilsson 1987. *Logical Foundation of Artificial Intelligence*. Los Altos, California: Morgan Kaufmann.

Gruber, T. 1993. “Toward Principles for the Design of Ontologies Used for Knowledge Sharing”, *International Workshop on Formal Ontology in Conceptual Analysis and Knowledge Representation*, Padova, LADSEB-CNR Int. Rep. 01/93.

Guarino, N. 1992. “Concepts, Attributes and Arbitrary Relations: Some Linguistic and Ontological Criteria for Structuring Knowledge Bases”, *Data & Knowledge Engineering* 8, 249-261.

Guarino, N. and L. Boldrin 1993. “Concepts and Relations”, *International Workshop on*

Formal Ontology in Conceptual Analysis and Knowledge Representation, Padova, LADSEB-CNR Int. Rep. 01/93.

Guarino, N. and L. Boldrin 1993. “Ontological Requirements for Knowledge Sharing”, *IJCAI workshop on Knowledge Sharing and Information Interchange*, Chambery, France.

Guarino, N., Carrara, M., and Giaretta, P. "An Ontology of Meta-Level Categories", LADSEB-CNR Int. Rep. 06/93.

Mulligan, K., P. M. Simons, B. Smith 1984. “Truth Makers”, *Philosophy and Phenomenological Research* 44, 287-321.

Pelletier, F. J. and L. K. Schubert 1989. “Mass Expressions”. In D. Gabbay and F. Günthner (Ed.), *Handbook of Philosophical Logic*. Reidel, 1989.

Putnam, H. 1981. *Reason, Truth, and History*. Cambridge: Cambridge University Press.

Simons, P. 1987. *Parts: a Study in Ontology*. Oxford: Clarendon Press.

Smith, B. 1992. “Characteristica Universalis”. In K. Mulligan (Ed.), *Language, Truth and Ontology*. Dordrecht: Kluwer, 1992.

Strawson, P. F. 1959. *Individuals. An Essay in Descriptive Metaphysics*. London and New York: Routledge.

Wiggins, D. 1980. *Sanemess and Substance*. Oxford: Blackwell, 1980.

Woods, W. A. 1975. “What’s in a Link: Foundations for Semantic Networks.”. In D. G. Bobrow and A. M. Collins (Ed.), *Representation and Understanding: Studies in Cognitive Science*. Academic Press, 1975.

Woods, W. A. and J. G. Schmolze 1992. “The KL-ONE family.”. In F. W. Lehmann (Ed.), *Semantic Networks in Artificial Intelligence.* Oxford: Pergamon Press, 1992.

Author address:

Nicola Guarino, LADSEB-CNR, National Research Council, I-35020 Padova, Italy
tel. +39 49 8295751, fax +39 49 8295763, email guarino@ladseb.pd.cnr.it