An Ontology of Meta-Level Categories

Nicola Guarino, Massimiliano Carrara,
and Pierdaniele Giaretta

LADSEB-CNR Int. Rep. 6/93
Substantially revised version - January 16, 1994

Abstract

We focus in this paper on some meta-level ontological distinctions among unary predicates, like those between concepts and assertional properties. Three are the main contributions of this work, mostly based on a revisitation of philosophical (and linguistic) literature in the perspective of knowledge representation. The first is a formal notion of ontological commitment, based on a modal logic endowed with mereological and topological primitives. The second is a formal account of Strawson's distinction between *sortal* and *non-sortal* predicates. Assertional properties like *red* belong to the latter category, while the former category is further refined by distinguishing *substantial* predicates (corresponding to *types* like *person*) from *non-substantial* predicates (corresponding to *roles* like *student*). The third technical contribution is definition of countability which exploits the topological notion of connection to capture the intended semantics of unary predicates.

An Ontology of Meta-Level Categories

Nicola Guarino, Massimiliano Carrara, and Pierdaniele Giaretta

1. Introduction

Most KR formalisms differ from pure first-order logic in their *structuring power*, i.e. their ability to make evident the "structure" of a domain. For example, the advantage of frame-based languages 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 from the very beginning of a knowledge base development process [6]. For a particular knowledge base, such ontological commitments are however implicit and strongly dependent on the particular task being considered, since the formalism itself is in general deliberately *neutral* as concerns ontological choices: in their well-known textbook on AI, Genesereth and Nilsson ([8], p. 13) explicitly state the “essential ontological promiscuity of AI”. We have argued elsewhere *against* this neutrality [12,14,13], claiming that a rigorous ontological foundation for knowledge representation can result in better methodologies for conceptual design of data and knowledge bases, facilitating knowledge sharing and reuse. We have shown how theories defined at the (so-called) 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. Referring to the classification proposed in [2], we have introduced therefore the notion of *ontological level*, intermediate between the epistemological and the conceptual levels. At the ontological level, formal distinctions are made among logical predicates, distinguishing between (meta-level) categories such as *concepts, roles,* and *assertional properties*.

Such distinctions have three main purposes. First, they allow the knowledge engineer to make clear the *intended meaning* of a particular logical axiomatization, which is of course much more restricted than the set of all its Tarskian models. This is especially important since we are constantly using natural language words within our formulas, relying on them to make our statements readable and to convey meanings not explicitly stated. However, since words are ambiguous in natural language, it may be important to “tag” these words with a semantic category, in association with a suitable axiomatisation, in order to guarantee a consistent interpretation[1]. This is unavoidable, in our opinion, if we want to share theories across different domains [15,11]. A second important advantage of clear ontological distinctions is the possibility of a *methodological foundation* for deciding between the various representation choices offered by a KR formalism: for example, within a hybrid terminological framework, for deciding whether a predicate should go in the TBox or ABox, or how a KL-ONE role should be related to a correponding concept. Finally, these distinctions may impact the *reasoning services* offered by a KR formalism: for example, a terminological reasoner can forbid certain kinds of update on the basis of ontological considerations; it may take advantage of the fact that some kinds of concepts form a tree, while in general they do not [22]; it may maintain indices for instances of concepts but not for instances of properties; it may provide domain-checking facilities for properties but not for concepts[2].

We focus in this paper on some fundamental ontological distinctions among unary predicates, refining and extending some previous work [13]. Most of our results come from a revisitation, from the point of view of KR, of philosophical (and linguistic) work largely extraneous to the KR tradition. The main distinction we focus on is that between *sortal* and *non-sortal* predicates, originally introduced by Locke and discussed in more detail e.g. by Strawson [21] and Wiggins [25]. According to Strawson, a sortal predicate (like *apple*) “supplies a principle for distinguishing and counting individual particulars which it collects”, while a non-sortal predicate (like *red*) “supplies such a principle only for particulars already distinguished, or distinguishable, in accordance with some antecedent principle or method”[23]. This distinction is (roughly) reflected in natural language by the fact that the former terms are common nouns, while the latter are adjectives and verbs. It is implicitly present in the KR literature, where sortal predicates are usually called “concepts”, while characterising predicates are called “properties”. Within current KR formalisms, however, the difference between the two is only based on heuristic considerations, and nothing in the semantics of a concept forbids it from being treated like any other unary predicate.

---

[1] Notice that we do not mean that the user is forced to accept some *one* fixed interpretation of a given word: simply, we want to offer some instruments to help specifying the intended interpretation.

[2] The last two examples are due to Bob MacGregor.

After giving a simple example showing the necessity of the above distinction, we introduce in section 3 a formal notion of *ontological commitment,* based on a modal logic endowed with mereological and topological primitives. In section 4 we present an ontology of kinds of unary predicates (i.e., meta-level categories), where the basic sortal/non-sortal distinction is further explored and refined.

2. Reds and apples.

Suppose we want to state that a red apple exists. In standard first-order logic, it is a simple matter to write down something like $\exists x.(Ax \wedge Rx)$[1]. If we want to impose some *structure* on our domain, then we may resort to a many-sorted logic. Then, however, we have to decide which of our 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. How can such a choice be motivated, if the semantics of a primitive sort is the same as that of its corresponding first-order predicate?

A statement like $\exists x{:}R.Ax$ sounds intuitively odd. What are we quantifying over? Do we assume something like the existence of “instances of redness” that can have the property of being apples? Our position is that structured representation languages like many-sorted logics should be constructed in such a way that predicates can be taken as sorts (or concepts, in KR terminology) only when they satisfy formal, necessary conditions at the meta-level, grounded on common-sense intuitions. According to our previous discussion, a predicate like *red* should not satisfy such conditions, and thus it should be excluded from being used as a sort. Notice that this may still be a matter of point of view, since it is still the user who must decide which conditions reflect the *intended* use of a predicate like *red*. For example, compare the statement mentioned above with others where the same predicate appears in different contexts (Fig. 1):

---

[1] As usual, predicates are symbolized via the capitalized first letter of the word used in the text.

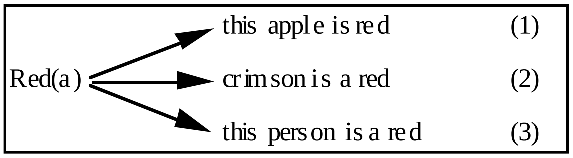

Fig. 1. Varieties of predication.

In case (2) the argument refers to a particular colour gradation belonging to the set of “reds”, while in (3) the argument refers to a human-being, meaning for instance that he is a communist.

How can we account for such semantic differences? 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 types of subject-predicate relationships, corresponding to meta-level categories of predicates. Studying the formal properties of such categories is a matter of *formal ontology*, recently defined by Cocchiarella as "the systematic, formal, axiomatic development of the logic of all forms and modes of being" [4].

3. The formal framework

Instead of trying to give a “universal” definition of the main predicate categories, we shall pursue here a more modest goal: our definitions will be related to a specific first-order theory whose intended meaning we are interested in specifying. This means that the basic building blocks of knowledge are already fixed, being the atomic predicates of the theory itself; our job will be to offer a formal instrument for clarifying their ontological implications, for the specific purposes of knowledge understanding and reuse among users belonging to a single culture. We assume therefore that the intended models of our theory, rather than describing a real or hypothetical situation in a world that has the same laws of nature of ours [5], are states of affairs having an “idealised rational acceptability” [18].

Suppose we have a non-functional first-order language L with signature $\Sigma=<K, R>$, where K is a set of constant symbols, R is a finite set of n-ary predicate symbols and $P \subseteq R$ is the set of monadic predicate symbols. Let T be a theory of L, D its intended domain and **M** the set of its models of the form <I, D>, where I is the usual interpretation function for constants and predicate symbols. We are interested in some formal criteria accounting for those ontological distinctions among the elements of P which are considered as relevant to the purposes of T as applied to D. For example, we are looking for a clear distinction between sortal and non-sortal predicates which can account for the structuring choices implicit in the translation of T into an order-sorted language T_s with signature $\Sigma_s = \langle K, S, Q\rangle$, where $S \subseteq P$ is a set of sortal predicates and $Q = R \backslash S$ a set of ordinary predicates. We shall see how this and other distinctions will be expressed in terms of constraints which restrict the set **M**.

Our main methodological assumptions here are that (i) we need some notion of tense and modality in order to account for the intended meaning of predicate symbols; (ii) we need mereology and topology in order to capture the *a priori* structure of a domain. In the following, we first extend our first order language by introducing a semantics of tense and modality which satisfies our purposes, then we further extend both the language and the domain on the basis of mereo-topological principles, in order to formalize the notion of *ontological commitment* for the original language applied to the original domain.

Def. 1. Let L be a first-order language with signature Σ. The *tense-modal extension* of L is a language L_m with signature Σ, which adds to the logical symbols of L the usual modal operators " and Æ and the tense operators **F** and **P**, respectively standing for "sometimes in the future" and "sometimes in the past".

Def. 2. Let L be a first-order language with signature $\Sigma=\langle K, R\rangle$, L_m its tense-modal extension and D a domain. A *constant-domain model* for L_m based on D is a structure $M = \langle W, r, b, D, f_K, f_R\rangle$, where W is a set of possible worlds, r is a binary relation on W, b is a linear order on W, f_K is a mapping that assigns to each $w \in W$ and each $c \in K$ an element $f_K(w,c)$ of D, and f_R is a mapping that assigns to each $w \in W$ and each n-ary predicate symbol $r_n \in R$ an n-ary relation $f_R(w,r_n)$ on D.[1]

Def. 3. Let L be a first-order language with signature $\Sigma=\langle K, R\rangle$, L_m its tense-modal extension and $M = \langle W, r, b, D, f_K, f_R\rangle$ a constant-domain model for L_m based on D. Two worlds $w_i, w_j \in W$ are *compatible* in M iff, for each $c \in K$, $f_K(w_i,c)=f_K(w_j,c)$.

In other words, two worlds are compatible if they describe alternative states of affairs *involving the same elements of D*. The notion of modality we are interested in accounts for rational possibility among such alternative states of affairs, i.e. among sets of mutually compatible worlds. We choose therefore the accessibility

---

[1] This definition is taken from [7], extended with a relation b intended to express the temporal precedence relationship between worlds.

relation r to coincide with the compatibility relation. Since the latter is an equivalence relation, our modal theory will be S5.

Def. 4. Let L be a first-order language, L_m its tense-modal extension and D a domain. A *compatibility model* for L_m based on D is a constant-domain model for L_m based on D, where r is the compatibility relation between worlds.

Notice that, although the mapping f_K is not rigid, the particular choice for r we have made within compatibility models allows us to ignore the distinction between *de dicto* and *de facto* interpretations of modal formulas typical of non-rigid models, since the worlds connected by the accessibility relation always share the same interpretation for constants. The notion of truth in a model at a world is therefore pretty standard, and it will not defined here in detail because of space limitiations. The only slight deviation from standard truth conditions regards formulas that involve tense operators.

In particular, a formula Φ is necessary in a compatibility model M at a world w (written M,w |= Æ Φ) iff M,v |= Φ for every v such that r(w,v); M,w |= " Φ) iff M,v |= Φ for some v such that r(w,v); M,w |= **FΦ** iff M,v |= Φ for some v such that r(w,v) and b(w,v); M,w |= **PΦ** iff M,v |= Φ for some v such that r(w,v) and b(v,w). Φ is valid in M (M |= Φ) iff M,w |= Φ for each world w of M.

Given a domain D, consider now the set of all compatibility models based on D of the tense-modal extension L_m of a language L. In order to account for our ontological assumptions about D, we should somehow restrict such a set, excluding those models allowing for non-intended worlds or too large sets of compatible worlds. Suppose for instance we want to clarify the assumptions underlying the formula $\exists x.(Ax \wedge Rx)$ discussed in section 2, whose intended interpretation is that "a red apple exists". In this case, if $a \in K$ is a constant symbol of our language L, the two worlds satisfying respectively $(Aa \wedge Ra)$ and $(Aa \wedge \neg Ra)$ should turn out to be compatible, while those satisfying $(Ra \wedge Aa)$ and $(Ra \wedge \neg Aa)$ should not: a world where something is an apple cannot be compatible with another where *the same thing* is not an apple, since being an apple affects the *identity* of an object. Within our framework, we can express such constraints by restricting the set of all possible compatibility models of L_m:

Def. 5. A *commitment* for L_m based on D is a set $\mathbf{C}$ of compatibility models for L_m based on D. Such a commitment can be specified by an S5 modal theory of L_m, being in this case the set of all its compatibility models based on D. A formula Φ of L_m is valid in $\mathbf{C}$ ($\mathbf{C} \models \Phi$) iff it is valid in each model $M \in \mathbf{C}$.

We shall see in the next section how we can express the constraints mentioned in the example above by choosing a commitment $\mathbf{C}$ such that $\mathbf{C} \models \forall x(Px \supset \text{Æ}\ Px)$. Before that, we need first to further extend both L_m and D in order to be able to express our ontological assumptions about D itself:

Def. 6. Let L be a first order language with signature Σ=<K, R>, and L_m its tense-modal extension. The *ontological extension* L_o of L is the tense-modal extension of a language L' with signature Σ_o=<K, R_o>, where $R_o = R \cup \{<, C\}$, while < and C are two binary predicate symbols used to represent the mereological relation of "proper part" and the topological relation of "spatial connection".

Since our domain is not restricted to topological entities only, the relation of spatial connection can have arguments which are physical bodies and not only regions as in [19]. We assume here that two physical bodies are spatially connected if *their spatial extensions* are connected in the sense defined in [19] (i.e. two *regions* are connected if their topological closures share a point). Notice that we do not share with Randell and colleagues the choice to define parthood in terms of connection[1].

Def. 7. The *mereological closure* of a domain D is the set D_o obtained by adding to D the set of all proper parts of the elements of D.

Def. 8. An *ontological commitment* $\mathbf{O}$ for L based on D is a commitment for L_o based on D_o, such that the following minimal mereo-topological theory is valid in $\mathbf{O}$[2].

A1	$x < y \supset \neg (y < x)$	(asymmetry of parthood)
A2	$x < y \wedge y < z \supset x < z$	(transitivity of parthood)
A3	$x < y \supset \exists z.(z < y \wedge \neg Ozx)$	(supplementation)
A4	$\forall x.Cxx$	(reflexivity of connection)

---

[1] See [24] for a discussion of the relationsbips between mereology and topology.

[2] Axioms A1-A3 are taken from [20], while axioms A4-A5 from [19].

A5 $\forall x \forall y. Cxy \supset Cyx$ (symmetry of connection)

D1 $x \leq y =_{def} x < y \vee x = y$ (part)

D2 $Oxy =_{def} \exists z.\ z \leq x \wedge z \leq y$ (overlap)

4. A basic ontology of unary predicate types

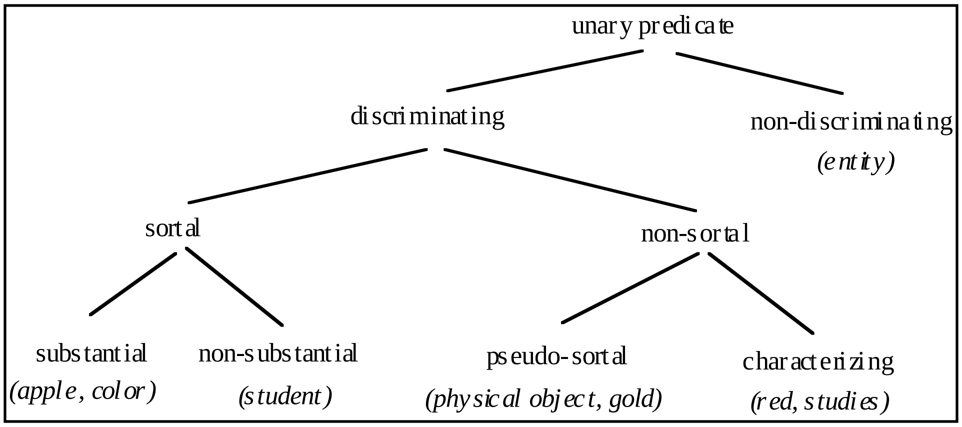

Fig. 2.
A basic ontology of unary predicate types. Examples shown in italic refers to usual interpretation.

The basic ontology we propose is sketched in Fig. 2 above. It looks like a tree, where each layer corresponds to a single meta-level property used in various ways to discriminate among different nodes.

Within our modal framework, the first fundamental distinction we make among unary predicates regards their “discriminating power”. If we want to use a predicate for knowledge-structuring purposes it cannot be necessarily false for each element of the domain, i.e. it must be *natural* in the sense of [3]. Moreover, we are interested in predicates that tell us something non-trivial about the domain, excluding therefore those which are always necessarily true.

Def. 9. Let L be a first order language, P a monadic predicate of L, and **O** an ontological commitment for L. P is called *natural* under **O** iff **O** $\models$ "$\exists x.Px$. A natural predicate is *discriminating* under **O** iff **O** $\models$ "$\exists x.\neg Px$.[1]

4.1 Sortals vs. Non-sortals

---

[1] In the following, we shall omit to specify the words "under **O**" where this is obvious from the context.

Among discriminating unary predicates, the relevant distinction is the classical one between sortals and non-sortals. To this end, we introduce two meta-level properties which give a minimal characterization of individuality, and are therefore distinctive of sortal predicates. They bear on two main notions proposed in the philosophical literature: *countability* [10] and temporal *reidentifiability* [25]. The former is bound to the capacity of a predicate to isolate a given object among others: "*this* is a P, this is *another* P, this *is not* a P". In other words, if P is a sortal predicate, then it is possible to answer: "how many Ps are there?" In the literature, various "divisivity" criteria have been proposed to account for the countable/non-countable distinction. Excluding those based on universal quantification on all parts of an object for reasons having to do with the problem of granularity, a quite satisfactory criterion is the one proposed by Griffin [10], which can be formulated in such a way that P is a countable predicate iff $\forall x.(Px \supset \neg\exists z.(z < x \wedge Pz))$. But such a criterion does not take into account a notion of topological connection which seems to be related to the notion of countability. In our opinion, the main feature of countable predicates is that they cannot be true of an object and of a non-isolated part of it. For example, we think it is natural to consider *piece of wood* as a countable predicate, but it turns out to be uncountable according to Griffin's definition. The point is that in its ordinary meaning such a predicate does not apply to any part of a single, integral, piece of wood. We think that this is a structural feature of countable predicates. In order to capture it within our formal framework, let us introduce the following definitions within the ontological extension L_O of a language L:

D3 $x – y =_{def} \sigma z.(z \leq x \wedge \neg Ozy)$ [1] (mereological difference)

D4 $x <_i y =_{def} x<y \wedge \neg Cx(y–x)$ (isolated part)

D5 $x <_c y =_{def} x<y \wedge Cx(y–x)$ (connected part)

Def. 10. A discriminating predicate P is called *countable* in **O** iff **O** $\models \forall x. (Px \supset \neg\exists z.(z <_c x \wedge Pz))$.

In the above definition, we have simply substituted the relation of connected part to the relation of part appearing in Griffin's definition. The following two theorems follow from Definition 10:

---

[1] $\sigma x. \phi$ is the mereological sum of all x such that ϕ is true.

Theorem 1: A countable predicate P only holds for maximally connected entities having the property of being parts of the same P.

Theorem 2: A predicate is countable if it only applies to atomic entities, i.e. entities having no parts.

According to theorem 1, the predicate *piece of wood* is countable if (as seems natural) it only applies to isolated pieces of wood, while the monadic predicate *color* turns out to be countable according to theorem 2, assuming that a color has no part. On the other hand, according to its ordinary sense a predicate like *red* is not countable[1], since while holding for a physical object it can also apply to non-isolated parts of it, such as its surface.

The above definition allows us to consider predicates denoting physical structures like *stack* (of blocks), *chain* or *lump* (of coal) as countable predicates only if it can be claimed, perhaps on the basis of Gestalt-theoretical considerations, that no connected part of a physically realized structure can be a structure of the same kind [21]. In this sense, a substack can be a stack only as an isolated whole. There are some intuitive and practical reasons in favour of this way of thinking. For example, a request to count the chains put in a box is not usually understood as a request to count also the subchains of such chains.

Notice that we do not require instances of countable predicates to be isolated entities: for example, we want *arm* to be countable and such that both detached and undetached arms are instances of it. However, it is reasonable to hold that *tube* is countable. It follows that no part of a tube is a tube, otherwise it would violate the assumption of countability. So while arms are instances of *arm* even before a possible detaching event, the same does not hold for halves of tubes. Lack of analogy between an undetached arm and an undetached half of a tube is motivated by the fact that in the former case an object is connected to something of a different kind.

We may be tempted to conclude now that countability is enough to decide about sortality. Things are not so easy, however. Think of a unary predicate expressed by a verb, like *studies*. It seems to be countable according to our definition, and in fact we can count those entities x such that the statement *x studies* is true, but still it seems odd to consider *studies* as a sortal predicate. The

---

[1] Notice that when we attach an ontological category to a linguistic term we do not imply any *a priori* meaning attribution: we simply assume, for simplicity reasons, the ontological commitment corresponding to the usual meaning of the term (in this case the meaning of case 1 in Fig. 1).

reason is that sortality implies a notion of *reidentifiability* across time, which is not implied by the semantics of a verb. Linguists such as Givòn [9] have pointed out that *temporal stability* can be a useful criterion to distinguish verbs from nouns. We say that a predicate is temporally stable when, if it holds for an object at a given time, then it must hold for the same object at another time.

Def. 11. A discriminating predicate P is called *temporally stable* under **O** iff $\mathbf{O} \models \forall x.(Px \supset \mathbf{F}Px \lor \mathbf{P}Px)$.

In conclusion, we characterize sortal predicates by the following definition:

Def. 12. A discriminating predicate P is called *sortal* under **O** iff it is both countable and temporally stable under **O**, and *non-sortal* otherwise.

According to this definition, we have a criterion to distinguish between the two predicates involved in the statement "a red apple exists". *apple* will be in this case a sortal predicate being countable and temporally stable, while *red* will be non-sortal being not countable under our intended intepretation. Both $\exists x{:}R.Ax$ and $\exists(x{:}A, y{:}R).x=y$ will be therefore excluded from a many-sorted axiomatisation.

4.2 Rigidity

Although useful for many purposes, the distinction between sortal and non-sortal predicates discussed above is not fine enough to account for the difference in the interpretation of *red* in cases (2) and (3) of Fig. 1, since in both of them *red* is used as a countable predicate. Let us therefore further explore the ontological distinctions we can draw among both sortal and non-sortal predicates. An observation that comes to mind, when trying to formalise the nature of the subject-predicate relationship, 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 latter example. This observation goes back to Aristotelian essentialism, and can be formalised as follows [1]:

Def. 13. A discriminating predicate P is *ontologically rigid* iff $\mathbf{O} \models \forall x(Px \supset \text{Æ}\, Px)$.

An immediate theorem following from Definition 11 is the following:

Theorem 3. Any ontologically rigid predicate is also temporally stable.

However, the example above notwithstanding, ontological rigidity is not a sufficient condition for sortality. In fact, there are a number of rigid predicates which should be excluded from being sortals, since no clear distinction criteria are associated with them. Predicates corresponding to certain mass nouns belong to this category (at least if their arguments denote an amount of stuff and not a particular object), as well as "high level" predicates like *physical object*, *individual*, *event*. We call these predicates *pseudo-sortals*[1]. They are all rigid (and therefore stable) but not countable.

Def. 14. Let P be a non-sortal predicate under **O**. It is a *pseudo-sortal* iff it is ontologically rigid under **O**, and a *characterising predicate* otherwise.

Rigidity cannot be considered as a necessary condition for sortality, either. According to our definition, sortals include predicates like *student*, which – although not rigid – are still countable and stable enough to guarantee distinguishability and reidentification. Following [25], we call such predicates *non-substantial sortals*[2].

Def. 15. Let P be a sortal predicate under **O**. It is a *substantial sortal* iff it is ontologically rigid under **O**, and a *non-substantial sortal* otherwise.

As noticed before, temporal stability plays here a crucial role for distinguishing *student* from *studies*: both are countable and not ontologically rigid, but the latter is not temporally stable and is therefore a characterizing predicate, while the former is a non-substantial sortal.

We are now in a position to exploit the above distinctions in order to specify the ontological commitment of a first order theory: for instance, stating that *red* is a characterizing predicate will clarify its intended meaning in the case (1) of Fig. 1. In case (2), *red* is rigid and countable, since its argument is a colour gradation: it will be therefore a substantial sortal (crimson *has* to be a red: see [16,17], p. 10). Finally, in case (3), *red* is used as a contingent property of human-beings and

---

[1] They are called "super sortals" in [17].

[2] According to the current terminology used in knowledge representation, substantial sortals should in our opinion correspond to *types* and non-substantial sortals to *roles* (in the sense of [22]), while the terms *class* or *concept* should be reserved to the union of sortal and pseudo-sortal predicates.

hence is not rigid, while it is countable and temporally stable: *red* is therefore a non substantial sortal.

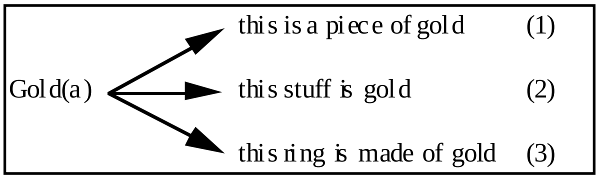

Fig. 3. Different interpretations of mass nouns.

Another interesting example regards the different interpretations of a mass noun like *gold*, reported in Fig. 3 above. In case (1), *gold* is countable, stable but not rigid (since that piece can have been taken from a rock, for instance), and it is used as a non-substantial sortal; in cases (2) and (3) the predicate is non-countable, but in the former case it is rigid (and *gold* is therefore a pseudo sortal), while in the latter it is non-rigid, and *gold* is a characterizing predicate.

5. Conclusions

In our opinion, three are the main contributions of this paper. The first one is the formal account of ontological commitment we have given within 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. The second one is our definition of countability, which seems to solve some of the puzzling cases reported in the literature. The third one is the formalization of Strawson's distinction between sortal and non-sortal predicates, which has been further refined by taking into account Wiggins' distinction between substantial and non-substantial predicates. Far from claiming to have said any definitive word on these issues, we would like to underline here that (i) *some* formal properties which account for distinctions among predicate types can indeed be worked out, even if complete, unproblematic definitions may never be given; (ii) when the semantics of structuring primitives used in KR languages is restricted in such a way as to take into account of such formal distinctions at the ontological level, then 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

We are indebted to Luca Boldrin, Dario Maguolo and Barry Smith for their valuable comments on earlier drafts of this paper.

References

[1] R. Barcan Marcus 1968. Essential Attribution. *The Journal of Philosophy*, **7**.

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

[3] N. Cocchiarella 1989. Philosophical Perspectives on Formal Theories of Predication. In D. Gabbay and F. Günthner (ed.), *Handbook of Philosophical Logic*. Reidel: 253-326.

[4] N. B. Cocchiarella 1991. Formal Ontology. In H. Burkhardt and B. Smith (ed.), *Handbook of Metaphysics and Ontology*. Philosophia Verlag, Munich.

[5] N. B. Cocchiarella 1993. Knowledge Representation in Conceptual Realism. In *Proceedings of International Workshop on Formal Ontology in Conceptual Analysis and Knowledge Representation*. Padova, Italy, LADSEB-CNR Int. Rep. 01/93.

[6] R. Davis, H. Shrobe, and P. Szolovits 1993. What is in a Knowledge Representation? *AI Magazine*, Spring 1993.

[7] M. Fitting 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*. Clarendon Press, Oxford.

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

[9] T. Givón 1979. *On Understanding Grammar.* Academic Press, New-York.

[10] N. Griffin 1977. *Relative Identity.* Oxford University Press, Oxford.

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

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

[13] N. Guarino 1994. The Ontological Level. In R. Casati, B. Smith and G. White (ed.), *Philosophy and the Cognitive Science*. Hölder-Pichler-Tempsky, Vienna.

[14] N. Guarino and L. Boldrin 1993. Concepts and Relations. In *Proceedings of International Workshop on Formal Ontology in Conceptual Analysis and Knowledge Representation*. Padova, LADSEB-CNR Int. Rep. 01/93.

[15] R. Neches, R. Fikes, T. Finin, T. Gruber, R. Patil, T. Senator, and W. R. Swartout 1991. Enabling Technology for Knowledge Sharing. *AI Magazine*, : fall.

[16] F. J. Pelletier 1979. Non-Singular References: Some Preliminaries. In F. J. Pelletier (ed.), *Mass Terms: Some Philosophical Problems*. Reidel, Dordrecht: 1-14.

[17] F. J. Pelletier and L. K. Schubert 1989. Mass Expressions. In D. Gabbay and F. Günthner (ed.), *Handbook of Philosophical Logic*. Reidel.

[18] H. Putnam 1981. *Reason, Truth, and History.* Cambridge University Press, Cambridge.

[19] D. Randell, Z. Cui, and A. Cohn 1992. A spatial logic based on regions and connection. In *Proceedings of KR '92.* San Mateo (CA), Morgan Kaufmann.

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

[21] B. Smith 1992. Characteristica Universalis. In K. Mulligan (ed.), *Language, Truth and Ontology*. Kluwer, Dordrecht: 48-77.

[22] J. F. Sowa 1988. Using a lexicon of canonical graphs in a semantic interpreter. In M. W. Evens (ed.), *Relational models of the lexicon*. Cambridege Univerity Press.

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

[24] A. Varzi 1994. On the Boundary Between Mereology and Topology. In R. Casati, B. Smith and G. White (ed.), *Philosophy and the Cognitive Science*. Hölder-Pichler-Tempsky, Vienna.

[25] D. Wiggins 1980. *Sameness and Substance.* Blakwell, Oxford.