Formalizing Ontological Commitments

Nicola Guarino

LADSEB-CNR, National Research Council,
Corso Stati Uniti, 4
I-35127 Padova, Italy
guarino@ladseb.pd.cnr.it

Massimiliano Carrara

Viale Ungheria, 43a
I-37046 Minerbe (VR)
Italy

Pierdaniele Giaretta

Institute of History of Philosophy,
University of Padova,
Piazza Capitaniato, 3
I-35100 Padova, Italy

Abstract

Formalizing the ontological commitment of a logical language means offering a way to specify the intended meaning of its vocabulary by constraining the set of its models, giving explicit information about the intended *nature* of the modelling primitives and their *a priori* relationships. We present here a formal definition of ontological commitment which aims to capture the very basic ontological assumptions about the intended domain, related to issues such as identity and internal structure. To tackle such issues, a modal framework endowed with mereo-topological primitives has been adopted. The paper is mostly based on a revisitation of philosophical (and linguistic) literature in the perspective of knowledge representation.

1 Introduction

First order logic is notoriously neutral with respect to ontological choices: when a logical language is used with the purpose of modelling a particular aspect of reality, the set **M** of all its models is usually much larger than the set $\mathbf{M_i}$ of the *intended* ones, which describe only those states of affairs which are compatible with some underlying *ontological commitment.* Such a commitment is usually implied by the *vocabulary* used, i.e. by the symbols chosen as constants and predicates: we sistematically use natural language words within our theories, relying on them to make our statements readable and to convey meanings not explicitly stated. However, since words are often vague and ambiguous in natural language, it may be important to constrain their semantics in order to guarantee a consistent interpretation. This is unavoidable, in our opinion, if we want to share theories across different domains (Neches et al. 1991, Gruber 1993).

In the philosophical literature, the notion of ontological commitment was first introduced by Quine (1961). According to him, a theory is ontologically commited to the entities which it quantifies over: "to be is to be the value of a variable". Such criterion was further refined by Church (1958) and Alston (1958), and finally modified by Searle (1969) in order to defend his argument that the ontological commitment of a theory simply coincides whith what it asserts. We reject the latter position, holding that non-equivalent theories can share the same commitment. On the other hand, Quine's proposal seems to be too weak for our purposes, since we want to include in the commitment some basic assumptions and distinctions presupposed by the theory.

In the AI community, the above position is at the basis of current projects for knowledge sharing and reuse (Neches et al. 1991). In the knowledge acquisition literature, the notion of ontological commitment has been introduced by Gruber (1993-1994) as an agreement to use a shared vocabulary specification: such a specification is a set of terminological axioms, and ontological commitment amounts to syntactical consistency with such axioms. This syntactical notion does not fit our intuitions, since it seems natural to allow two different vocabularies (using English or Italian words, for instance) to share the same ontology. In other words, the notion of ontological commitment should be a semantic one, not a syntactic one.

A semantic notion which gets closer to our purposes is that of *conceptualization*, defined in (Genesereth & Nilsson 1987) as a triple consisting of a domain, a set of functions (which we ignore for our purposes) and a set of relations on that domain. For instance (pp. 9-12), the triple <{*a, b, c*}, { }, {*on, above, table*}> is a conceptualization of a situation describing some block on a table. The authors note however that *names* of objects and relations refer to purely extensional entities. They describe therefore a *particular* state of affairs, without telling us anything about *other* possible states of affairs. On the other hand, the intended meaning implied by the names chosen for the relevant relations constraints *all* the possible states of affairs.

In conclusion, ontological commitment cannot be understood as "an explicit specification of a conceptualization" (Gruber 1993, p. 199), at least in the technical sense of the latter term. Rather, an ontological commitment should capture and constrain a *set* of conceptualizations. Formalizing the ontological commitment of a logical language means offering a way to specify the intended meaning of its vocabulary by constraining the set of its models, giving explicit information about the intended *nature* of the modelling primitives used and their *a priori* relationships. In this sense, an ontological commitment is a mapping between a language and something which can be called an *ontology*.

Consider a first-order language **L**, and a particular theory **T** of **L**. A possible way to formalize the ontological commitment of **L** is by specifying the set $\mathbf{M_i}$ of its intended models[1] by means of a suitable theory which uses the same language **L**. The only purpose of such a theory is to specify (or at least approximate) the meaning of the vocabulary used. Such a theory should be kept separated from theories which use the same vocabulary making assertions about particular states of affairs. Current approaches to the problem of *knowledge sharing* (Neches et al. 1991, Gruber 1993-1994) are along this line: a common (sub)theory called *terminology* describes the shared ontology, while task dependent knowledge is specified by separate theories.

The approach described above is not satisfactory, however, if our purpose is the formalization of the ontological commitment of an *arbitrary* language **L**. The reason is that nothing guarantees us that the vocabulary of **L** is adequate to express the ontological constraints we are interested in: if we want to capture the *a priori structure* of individuals we need enough *granularity* to be able to speak of their internal constitution, while to capture the *nature* of individuals and relations we also need suitable primitive categories. For instance, nothing is said in the example mentioned above about the nature of the domain: are *a, b,* and *c* physical objects or spatial regions?

A further limitation to the formalization of the ontological commitment of a language **L** by means of a first-order theory of **L** comes from the fact that ontology, being knowledge about *a priori* structure of reality, is intimately related to a notion of modality: choosing a particular intended model for a logical theory implies making implicit assumptions about *other* models compatible with the chosen one. In other words, there are *constraints* among possible models which reflect some important aspects of reality: for instance, models describing the temporal evolution of a situation should share the same interpretation for the individual constants used in the description of that situation.

We present in this paper a formal notion of the ontological commitment of a language **L**, expressed by means of a theory **T'** which uses a language **L'** *richer* than **L**. Such a language extends both the logical symbols and the vocabulary of **L** by adding modal operators, mereo-topological relations and basic domain categories. Since the only purpose of **T'** is to specify the intended use of **L**, it is not necessary to replace **T** with a larger theory **T"=T** $\cup$ T' of **L'**: for instance, a particular ontological property of a predicate, derivable in **T'**, does not need to be derived in **T**. Basically, deductions in **T'** are made by an external agent (e.g., a human being) which wants to understand or specify the ontological commitment of an agent holding the theory **T**; therefore, the computational properties of **L'** do not affect the behavior of the latter agent.

The main purpose of the present paper is to show how the intended interpretation of the primitive predicates used to model a particular domain can be formally specified in order to facilitate knowledge sharing and reuse. We shall base our work on a revisitation, from the point of view of KR, of philosophical (and linguistic) work largely extraneous to the KR tradition. In section 2 we give an example intended to motivate the kinds of distinctions we want to account for within an ontological commitment. After the presentation of our formal framework in section 3, we show in section 4 how some fundamental ontological properties of predicates can be expressed within that framework. A detailed analysis of meta-level ontological categories of unary predicates has been carried out in (Guarino, Carrara & Giaretta 1994). In this paper, we underline the necessity to adopt such distinctions as an uneludible part of any formal attempt to capture ontological commitment. We focus in particular on unary predicates, arguing that – in order to specify their intended meaning – a first, fundamental choice regards the distinction between so-called *sortal* and *non-sortal* predicates.

Such a distinction was originally introduced by Locke and discussed in (Strawson 1959) and (Wiggins 1980). 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". 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.

The notion of *well-founded* ontological commitment is introduced in section 5, with the purpose of offering some concrete methodological guidelines to the current practice of knowledge engineering.

2 A Preliminary Example

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 $\exists x.(Ax \wedge Rx)$[2]. What is the ontological commitment of such a simple theory? First of all, we must specify what we are quantifying over. Do we assume something like the existence of "instances of redness" that can have the property of being apples? How can we state that our commitment is exactly the opposite one, where *red* is considered as a property and *apple* as a concept? Sure, the solution cannot consist of an *a priori* classification of predicate symbols, since – being them words of a natural language – their intended meaning depends on the context. For example, compare the statement mentioned above with others where the same predicate *red* appears in different contexts (Fig. 1): in case (2) the argument refers to a par-

[1]We refer to the models of a *language*, not to the models of a particular *theory* expressed in that language. See for example (Chang & Keisler 1973), p. 20.

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

ticular 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.

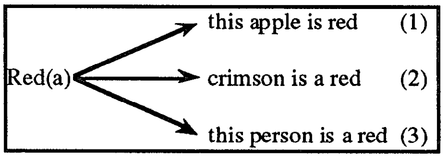

Fig. 1. Varieties of predication.

How can we account for such semantic differences? In this particular case, they are not simply related to the fact that the argument belongs to different domains: they are mainly due to 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 in (Cocchiarella 1991) as "the systematic, formal, axiomatic development of the logic of all forms and modes of being". In conclusion, although ontological distinctions not always can completely account for different semantic interpretation of linguistic terms, they may offer a significant help to the characterization of their intended meaning, as the present example shows.

3 The Formal Framework

Assuming as given the intended meanings of the predicates of a specific first order theory, we want to formally state their structural features, for the specific purposes of knowledge understanding and reuse among users belonging to a single culture. We assume here that the intended models of our theory, rather than describing merely hypothetical situations, are states of affairs having an "idealised rational acceptability" (Putnam 1981).

Suppose we have a first-order language **L** with signature Σ=<**K**, **R**>, where **K** is a set of constant symbols, **R** is a finite set of n-ary predicate symbols and $\mathbf{P} \subseteq \mathbf{R}$ is the set of monadic predicate symbols[1]. Let **T** be a theory of **L**, **D** its intended domain and **M** the set of its models M=<**D**, $\mathfrak{I}$>, where $\mathfrak{I}$ 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**.

Our main methodological assumptions here are that (i) we need some notion of modality in order to account for the intended meaning of predicate symbols, and (ii) we need mereology and topology in order to capture the *a priori* structure of a domain[2]. In the following, we first extend our first order language by introducing a semantics of 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 *modal extension* of **L** is the language $\mathbf{L_m}$ obtained by adding to the logical symbols of **L** the usual modal operators $\Diamond$ and $\Box$.

Def. 2 Let **L** be a first-order language with signature Σ=<**K**, **R**>, $\mathbf{L_m}$ its modal extension and **D** a domain. A *constant-domain rigid model* for $\mathbf{L_m}$ based on **D** is a structure M = <**W**, $\mathcal{R}$, **D**, $\Im_K$, $\Im_R$>, where **W** is a set of possible worlds, $\mathcal{R}$ is a binary relation on **W**, $\Im_K$ is a function that assigns to each $c \in \mathbf{K}$ an element $\Im_K(c)$ of **D**, and $\Im_R$ is a mapping that assigns to each $w \in \mathbf{W}$ and each n-ary predicate symbol $r_n \in \mathbf{R}$ an n-ary relation $\Im_R(w,r_n)$ on **D**.[3]

We want to give $\mathcal{R}$ the meaning of an *ontological compatibility* relation: intuitively, two worlds are ontologically compatible if they describe alternative states of affairs which do not disagree on the *a priori* nature of the domain. For instance, referring to the example discussed in the previous section, consider a world where a given individual is an instance of the two relations *apple* and *red* . Such a world will be compatible with another where such individual is still an apple but is not red, while it cannot be compatible with a world where *the same individual* is not an apple, since being an apple affects the *identity* of an object. To capture such intuitions, $\mathcal{R}$ must be reflexive, transitive and symmetric (i.e., an equivalence relation), and the corresponding modal logic will be therefore S5.

Def. 3 Let **L** be a first-order language, $\mathbf{L_m}$ its modal extension and **D** a domain. A *compatibility model* for $\mathbf{L_m}$ based on **D** is a constant-domain rigid model for $\mathbf{L_m}$ based on **D**, where $\mathcal{R}$ is the ontological compatibility relation between worlds.

The notion of truth in a compatibility model at a world is pretty standard, and it will not be defined here in detail because of space limitations. Given a domain **D**, consider now the set of all compatibility models based on **D** of the modal extension $\mathbf{L_m}$ of a language **L**. In order to account for our ontological assumptions about **D**, we should somehow restrict such a set, excluding those models that allow for non-intended worlds or too large sets of compatible worlds. Within our framework, we can express such con-

[1] We assume **L** as non functional just for the sake of simplicity. In the following, we shall use bold capital letters for sets, plain capital letters for predicate symbols and handwritten capital letters for relations.

[2] Also some notion of tense seems necessary (see section 4.1), but it will not appear here in our simplified formalization.

[3] This definition is taken from (Fitting 1993).

straints by restricting the set of all compatibility models of L_m:

Def. 4 A *commitment* for **L** based on **D** is a set **C** of compatibility models for $\mathbf{L_m}$ based on **D**. Such a commitment can be specified by an S5 modal theory of $\mathbf{L_m}$, being in this case the set of all its compatibility models based on **D**. A formula Φ of $\mathbf{L_m}$ is valid in **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 of the red apple by choosing a suitable commitment **C**. Before that, we need first to further extend both $\mathbf{L_m}$ and **D** in order to be able to express our ontological assumptions about **D** itself:

Def. 5 Let **L** be a first order language with signature Σ=<**K**, **R**>, and **L'** a language with signature Σ'=<**K**, **R'**>, where $\mathbf{R'}=\mathbf{R} \cup \{<, C\}$, while < and C are two binary predicate symbols used to represent the mereological relation of "proper part" and the topological relation of "connection". The modal extension of **L'** is called the *ontological extension* $\mathbf{L_o}$ of **L**.

The properties of the part-of relation have been extensively studied in (Simons 1987). Connection has been used as a topological primitive in (Clarke 1981) and more recently in (Randell, Cui & Cohn 1992). Since our domain is not restricted to topological entities only, the connection relation can have arguments which are physical bodies or events and not only regions as in (Randell, Cui & Cohn 1992). We assume here that two entities are connected if *their spatio-temporal extensions* are connected in the sense defined in (Randell, Cui & Cohn 1992) (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. 6 The *mereological closure* of a domain **D** is the set $\mathbf{D_o}$ obtained by adding to **D** the set of all proper parts of the elements of **D**.

Def. 7 An *ontological commitment* **O** for **L** based on **D** is a commitment for $\mathbf{L_o}$ based on $\mathbf{D_o}$, such that the following minimal mereo-topological theory is valid in **O**[2].

D1	$x \leq y =_{def} x < y \vee x = y$	(part)
D2	$Oxy =_{def} \exists z.\ z \leq x \wedge z \leq y$	(overlap)
A1	$x < y \supset \neg (y < x)$	(asymmetry)
A2	$x < y \wedge y < z \supset x < z$	(transitivity)
A3	$x < y \supset \exists z.(z < y \wedge \neg Ozx)$	(supplementation)
A4	$\forall x.Cxx$	(reflexivity)
A5	$\forall x \forall y.Cxy \supset Cyx$	(symmetry)

[1]See (Varzi 1994) for a discussion of the relationsbips between mereology and topology.
[2]Axioms A1-A3 are taken from (Simons 1987), while A4-A5 from (Randell, Cui & Cohn 1992).

4 Ontological Categories of Unary Predicates

In principle, any consistent set of formulas of $\mathbf{L_o}$ can be used to specify an ontological commitment for **L**; what is important of course is that the particular formulas chosen be suitable to really capture the underlying ontological intuitions. To this purpose, we must first define the relevant ontological properties of our predicates, and then explicitly declare the properties holding for each predicate of the language. A particular ontological commitment corresponds to a particular set of such declarations.

Let us see now how some important ontological properties of unary predicates can be easily formalized within the framework sketched in the previous section. The first fundamental distinction regards the "discriminating power"of unary predicates. If we want to use a predicate for knowledge-structuring purposes, it must tell us something non-trivial about the domain, and therefore it cannot be either necessarily true or necessarily false.

Def. 8 Let **L** be a first order language, P a monadic predicate of **L**, and **O** an ontological commitment for **L**. P is called *discriminating* in **O** iff $\mathbf{O} \models \Diamond\exists x.Px \wedge \Diamond\exists x.\neg Px$.

Some general distinctions among discriminating unary predicates are shown in Fig. 2. They are defined in the following as purely formal distinctions at the meta-level, completely independent of the nature of the domain. This means that our distinctions are intended to hold not only for standard examples related to the domain of physical objects, but also for predicates such as *color* or *marriage* whose arguments are universals like *red* or temporal entities like a particular marriage event. Analogously, no linguistic assumption is made on the names of predicates, which can be either nouns or adjectives.

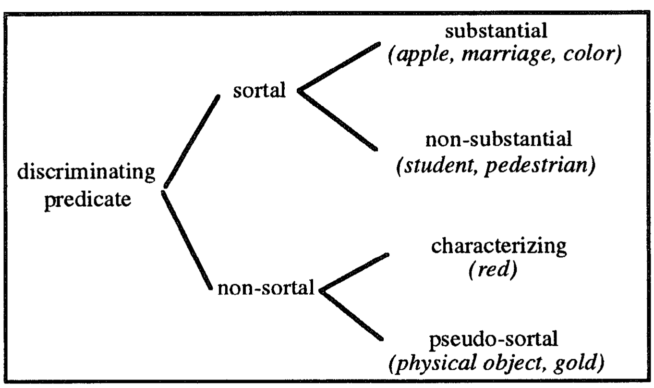

Fig. 2. Preliminary distinctions among unary predicates.

4.1 Countability

Among discriminating unary predicates, the distinction we focus on is the classical one between sortals and non-sortals. In the philosophical literature, such a distinction bears on two main notions: *countability* (Griffin 1977) and temporal *reidentifiability* (Wiggins 1980). 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 ask: "how many Ps are there?" The latter property is related to the fact that a predicate holds for the same individual through time, in the strong sense that it is possible to state "this is now the same P as before".

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 (1977), which can be formulated in such a way that P is a countable predicate iff $\forall x.(Px \supset \neg\exists z.(z < x \wedge Pz))$. Such a criterion, however, 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 cannot be excluded from being 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. In order to capture such a structural feature of countable predicates within our formal framework, let us introduce the following definitions for the ontological extension $\mathbf{L}_O$ of a language $\mathbf{L}$:

D3 $\sigma x\phi x =_{def} \iota x \forall y(Oyx \equiv \exists z(\phi z \wedge Ozy))$ (sum of all ϕers)[1]
D4 $x{-}y =_{def} \sigma z.(z \leq x \wedge \neg Ozy)$ (mereological difference)
D5 $x <_i y =_{def} x<y \wedge \neg Cx(y{-}x)$ (isolated part)
D6 $x <_c y =_{def} x<y \wedge Cx(y{-}x)$ (connected part)

Def. 9 A discriminating predicate P is called *countable* in **O** iff $\mathbf{O} \models \forall x.\ (Px \supset \neg\exists z.(z <_c x \wedge Pz))$.

In the above definition, we have simply substituted "connected part-of" to the part-of relation appearing in Griffin's definition. In other words, a countable predicate P only holds for entities which are "maximally connected" with respect to P, in the sense that they cannot have connected parts which are instances of P.

According to Def. 9, 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 if we assume that a color has no parts. On the other hand, according to its ordinary sense a predicate like *red* is not countable, since while holding for a physical object it can also apply to non-isolated parts of it, such as its surface.

The definition we have given allows us to consider predicates denoting physical structures like *stack* (of blocks)or *chain* 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 (Smith 1992). 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[2]. 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 the two cases is due to the fact that in the former case the argument of the predicate is connected to something of a different kind.

In (Guarino, Carrara & Giaretta 1994) we introduced a notion of *temporal stability* to take into account the idea of reidentifiability, and we defined sortality as the conjunction of countability and temporal stability. The former property is however much more important in practice, and so we omit in the following any reference to the latter because of space limitations.

Def. 10 A discriminating predicate P is called *sortal* in **O** iff it is countable in **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, while *red* will be non-sortal being not countable under our intended interpretation.

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 sortal 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"

[1] In order to avoid troubles with the satisfiability conditions for modal formulas involving the *iota* operator, we assume that terms built by means of such operator are contextually defined *a la* Russell. For instance, a formula like $P(\iota x.\phi x)$ is translated in $\exists x(Px \wedge \phi x \wedge (\forall y.\phi y \supset y{=}x))$.

[2] In contrast with (Smith 1992), we do not assume that detaching an arm is an event such that the arm before it is not the same arm as the arm after it.

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 (Barcan Marcus 1968):

Def. 11 A discriminating predicate P is *ontologically rigid* in **O** iff $\mathbf{O} \models \forall x(Px \supset \Box Px)$.

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 but not countable.

Def. 12 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. Following (Wiggins 1980), we call such predicates *non-substantial sortals*[2].

Def. 13 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.

We are now in a position to exploit the above distinctions in order to specify the ontological commitment of a first order language. Consider, for example, the statement (1) of Fig. (1). The formal language used to express such a statement includes the two predicate symbols R and A, standing respectively for *red* and *apple*. The ontological commitment of such a language which corresponds to the intended use of these two symbols can be specified by the following declarations (expressed either in a suitable metalanguage for $\mathbf{L_O}$ or via the corresponding axioms of $\mathbf{L_O}$):

A is a substantial sortal .

R is a characterizing predicate.

In statement (2), *red* is intended to be rigid and countable, since its argument is a colour gradation: it will be therefore declared as a substantial sortal (crimson *has* to be a red: see (Pelletier 1979), p. 10). Finally, in statement (3), *red* is used as a contingent property of human-beings and hence is not rigid, while it is still countable: it will be therefore declared as a non substantial sortal.

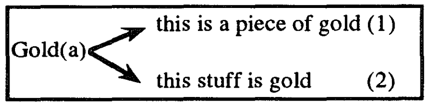

Fig. 3: Different interpretations of mass nouns.

Another interesting example regards the different interpretations of a mass noun like *gold*, reported in Fig. 3 above. In case (1), *gold* is intended as countable, 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 case (2) the predicate is non-countable and rigid, and *gold* is therefore a pseudo sortal.

5 Well-Founded Ontological Commitment

We would like to show in this section how the formal framework introduced above can be of concrete utility in the current practice of knowledge engineering. The first result of our approach is the possibility to draw a clear distinction between concepts and properties, in the sense usually ascribed to such terms within the KR community. Our proposal is that properties should coincide with what we called characterizing predicates, while all other kinds of unary predicates should be thought of as concepts.

Besides this first important distinction, our meta-level classification of unary predicates allows us to impose some further structure on the set of concepts, usually represented as an oriented graph where arcs denote subsumption relationships. As the size of this graph increases, it may be very useful to isolate a skeleton to be used for indexing and clustering purposes. Substantial sortals are a natural candidate to constitute such a skeleton[3], since their rigidity reduces the "tangleness" of the corresponding graph. However, to effectively use substantial sortals as a skeleton, we must introduce some further constraints, which lead to the notion of *well-founded ontological commitment*.

Def. 14 Let P and Q be two discriminating predicates in **O**. P is *subordinate* to Q in **O** iff $\mathbf{O} \models \forall x(Px \supset Qx) \wedge \neg\forall x(Qx \supset Px)$. P and Q are *disjoint* in **O** iff $\mathbf{O} \models \neg\exists x.Px \wedge Qx$. A set $\mathbf{P}=\{P_1, ..., P_n\}$ of mutually disjoint discriminating predicates in **O** is a *domain partition* in **O** iff $\mathbf{O} \models \forall x.(P_1x \vee ... \vee P_nx)$.

[1]They are called "super sortals" in (Pelletier & Schubert 1989). Notice that *physical object* is not intended here in the sense of spatially isolated *thing*. Therefore, it is assumed to be not countable.

[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 (Sowa 1988)), while the terms *class* or *concept* should be reserved to the union of sortal and pseudo-sortal predicates. Such terminological proposal is discussed in (Guarino, Carrara & Giaretta 1994).

[3]A similar proposal has been made by Sowa (1988), which however refers to an unspecified notion of "natural type".

Def. 15 An ontological commitment **O** based on **D** is *well-founded* iff:

- There is a set $\mathbf{C} \subseteq \mathbf{P}$ of mutually disjoint pseudo-sortal predicates called *categorial predicates,* such that (i) **C** is a domain partition in **O**, and (ii) no element $C \in \mathbf{C}$ is subordinate to a discriminating predicate.
- For each categorial predicate $C \in \mathbf{C}$, the substantial sortals subordinate to C and not subordinate to any other substantial sortal are mutually disjoint.
- Each non-substantial sortal is subordinate to a substantial sortal.

A well-founded ontological commitment introduces therefore a further subclass of discriminating predicates, i.e. categorial predicates, which belong to the class of pseudo-sortals according to the preliminary distinctions shown in Fig. 2. We call *mass-like predicates* those pseudo-sortals which are not categories; therefore, the final relevant distinctions within a well-founded commitment are those shown in Fig. 4.

Let us briefly motivate our definition of a well-founded ontological commitment. Categorial predicates are intended to represent what traditional ontology would call *summa genera.* A set of categorial predicates useful for a very broad domain is given by *physical object, event, spatial region, temporal interval, amount of matter*[1]. The fact that such predicates are assumed to be pseudo-sortals (and therefore uncountable) underlines their very general nature.

As for the second constraint mentioned in the definition, no particular structure is imposed on substantial sortals within a well-founded commitment[2], except that top-level substantial sortals should specify natural kinds within general categories: therefore, they must be disjoint and cannot overlap general categories.

Finally, the intuition behind the third constraint in Def. 15 is that in the case of substantial sortals the identity criterion is given by the predicate itself, while for non-substantial sortals it is provided by some superordinate sortal. Under this constraint, non-substantial sortals conform to the notion of "role type" proposed by Sowa, which fits well with the general meaning of the term "role": "Role types are subtypes of natural types in some particular patterns of relationships" (Sowa 1988). We suggest to adopt the term "role" for non-substantial sortals within the KR community, avoiding to use it as a synonym for an (arbitrary) binary relation as common practice in the KL-ONE circles.[3] An interesting consequence of Def. 15 is that, within a well-founded ontological commitment, any two overlapping non-substantial sortals are subordinate to the same substantial sortal.

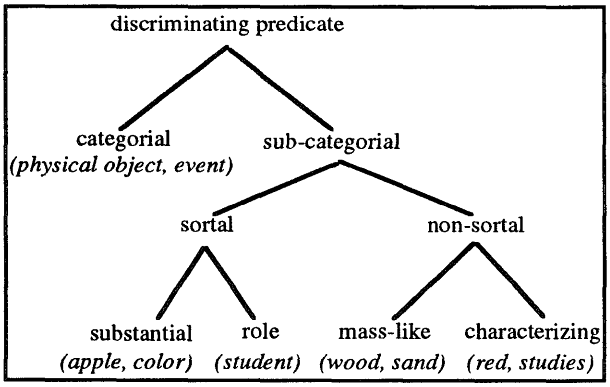

Fig. 4. Basic distinctions among discriminating predicates within a well-founded ontological commitment

6 Conclusions

We hope to have clarified in this paper the notion of ontological commitment, which has been understood in the past in various ways, both in the philosophical and AI tradition. We would like to stress that the notion we have defined is based on a fine-grained perspective of common-sense reality, where mereological and topological properties play an important role. We have defined ontological commitment as a map between a logical language and a *set* of semantic structures, and we have shown how modal logic, endowed with mereological and topological primitives, can be used to express ontological constraints among such structures. We have based our discussion on a simplified account of the ontological properties of unary predicates introduced in (Guarino, Carrara & Giaretta 1994), where a novel formalization of Strawson's distinction between sortal and non-sortal predicates has been presented in detail.

A number of extensions and refinements to the ontological distinctions described here are however necessary to obtain a satisfactory account of common-sense ontology. For instance, further distinctions among monadic predicate types must be defined in order to better characterize domain categories such as *physical objects, events, amounts of matter, spatial* or *temporal regions.* Moreover, ontological distinctions among binary relations must be introduced as well, possibly in a way similar to that described in (Guarino 1992).

We think we have still to learn a lot, to understand and represent the *a priori* laws that govern the structure of reality. Bearing on insights coming from the philosophical tradition of formal ontology, we have tried to show that

[1] These predicates should be characterized by suitable axioms, but such a task is beyond the scope of the present paper.

[2] It may be desirable, both for conceptual and computational reasons, to impose the condition that substantial sortals form a forest of trees; such a condition seems however not obtainable in many cases.

[3] See (Guarino 1992) for a general discussion on roles and attributes. Notice however that the distinctions among unary predicates discussed in that paper have been here drastically revised and simplified; in particular, no notion of ontological foundation is here advocated to distinguish between concepts and properties.

some of these laws are suitable to formal characterization, and we are convinced that they can have a profound impact on the current practice of knowledge engineering.

Acknowledgements

We are indebted to Luca Boldrin, Dario Maguolo and Barry Smith for their valuable comments on earlier drafts of this paper. Massimiliano Carrara's contribution has been made in the framework of a cooperation with LADSEB-CNR.

References

Alston, W. P. 1958. Ontological Commitments. *Philosophical Studies*, **9**: 8-16.

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

Chang, C. C. and Keisler, H. J. 1973. *Model Theory.* North-Holland-Elsevier Publishing Company, Amsterdam, London and New York.

Church, A. 1958. Ontological Commitment. *The Journal of Philosophy*, **55**: 1008-14.

Clarke, B. L. 1981. A Calculus of Individuals Based on "Connection". *Notre Dame Journal of Formal Logic*, **22**: 204-18.

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

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

Griffin, N. 1977. *Relative Identity.* Oxford University Press, Oxford.

Gruber, T. R. 1993. A translation approach to portable ontology specifications. *Knowledge Acquisition*, **5**: 199-220.

Gruber, T. 1994. Toward Principles for the Design of Ontologies Used for Knowledge Sharing. In N. Guarino and R. Poli (ed.), *Formal Ontology in Conceptual Analysis and Knowledge Representation*. (in preparation).

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., Carrara, M., and Giaretta, P. 1994. An Ontology of Meta-Level Categories. In J. Doyle, E. Sandewall and P. Torasso (eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourth International Conference (KR94)*. Morgan Kaufmann, San Mateo, CA.

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

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

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

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

Quine, W. O. 1961. *From a Logical Point of View, Nine Logico-Philosophical Essays.* Harvard University Press, Cambridge, MA.

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

Searle, J. 1969. *Speech Acts: An Essay on the Philosophy of Language.* Cambridge U. P., Cambridge.

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

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

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

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

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

Wiggins, D. 1980. *Sameness and Substance.* Blakwell, Oxford.