of Knowledge Representation and Reasoning: Proceedings of the Fourth International Conference (KR94), Morgan Kaufmann, San Mateo, CA.

An Ontology of Meta-Level Categories

Nicola Guarino

LADSEB-CNR,
National Research Council,
Corso Stati Uniti, 4
I-35129 PADOVA [Italy]

Massimiliano Carrara

LADSEB-CNR
National Research Council,
Corso Stati Uniti, 4
I-35129 PADOVA [Italy]

Pierdaniele Giaretta

Institute of History of Philosophy,
University of Padova,
P.zza Capitaniato, 3
I-35100 PADOVA [Italy]

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.

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 [11]. 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 ([13], p. 13) explicitly state the "essential ontological promiscuity of AI". We have argued elsewhere *against* this neutrality [18,20,21], 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 or order-sorted logics, cannot be distinguished from their "flat" first-order logic equivalents unless we make clear their implicit ontological assumptions. Referring to the classification proposed in [5], we have introduced therefore the notion of *ontological level*, intermediate between the epistemological and the conceptual levels [19]. 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 [23,16]. A second important advantage of clear ontological distinctions is the possibility of a *method-*

[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.

ological 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 [31]; 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[1].

We focus in this paper on some fundamental ontological distinctions among unary predicates, refining and extending some previous work [19]. 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 [32] and Wiggins [34]. 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" [32]. 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. The issue is also related to the semantic difference between count and non-count (or mass) terms. Philosophers have characterised count terms as denoting integral wholes, whereas entities denoted by mass terms are cumulative and divisive. This criterion has been a matter of lively debate [25], since such semantic-pragmatic distinctions not always correspond to the syntactical "count/mass" distinction, according to which, while mass-terms admit quantifiers like *much*, or *a little* and the indefinite article *some*, count-terms use the quantifiers *each*, *every*, *some*, *few*... and the indefinite article *a*.

Distinctions among unary predicates are also present in the KR literature, where sortal predicates are usually called "concepts", while characterising predicates are called "properties", or sometimes "qualities". The necessity of a distinction between the two kinds of predicates has always been acknowledged by advocates of the logicist approach in KR, as emerges clearly from the following quotation from David Israel [22]:

> "There is to be one tree for kinds of things and another for qualities of things. Kinds must be distinguished from qualities: being a cat must be distinguished (in kind, no doubt) from being red"

Within current KR formalisms, however, the difference between the two kinds of predicates is only based on heuristic considerations, and nothing in the semantics of a concept forbids it from being treated like any other unary predicate. Our task here is to formalize such a difference: our job is simpler than that of a linguist, since we do not try to classify a linguistic item as belonging to a particular category, but simply to make explicit its intended meaning when it is used as a predicate symbol with a specific representation purpose.

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 the philosophical literature, such a term was first used by Quine [27]. According to him, a logical theory is ontologically commited to the entities which it quantifies over. Quine expressed his criterion for ontological commitment with the slogan: "to be is to be the value of a variable". Such criterion was further refined by Church [6] and Alston [1], and finally modified by Searle in order to defend his argument that the ontological commitment of a theory simply coincides with what it asserts. We reject such a position, holding that different theories can share the same ontological commitment. In the AI community, this claim is at the basis of current projects for knowledge sharing and reuse [23]. In the knowledge acquisition literature, the notion of ontological commitment has been introduced by Gruber [16,17] as an agreement to use a shared vocabulary. We focus in this paper on the formal semantic interpretation of such a vocabulary: specifying 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 domain elements and relations and their *a priori* relationships. In order to capture such *a priori* knowledge we believe it is necessary to use a modal semantics, in contrast with Quine's view.

The notion of ontological commitment is exploited in section 4 to introduce some meta-level properties of unary predicates such as *countability*, *temporal stability* and *rigidity*. These properties allow us to establish an ontology meta-level categories of predicates, where the basic sortal/non-sortal distinction is further explored and refined. The impact of these distinctions on the current practice of knowledge engineering is discussed in section 5.

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

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.

As discussed in the previous section, the introduction of formal, necessary conditions for being a sort has a general ontological motivation. Besides that, ontological distinctions among predicates can be useful to make explicit a particular meaning of a lexical item. For example, compare the statement "a red apple exists" with others where the same term *red* appears in different contexts (Fig. 1):

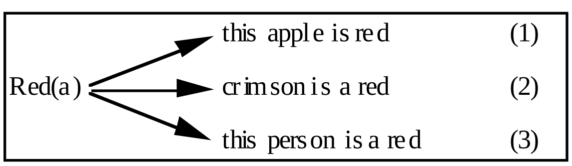

Fig. 1. Varieties of predication.

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. Clearly, *red* is a case of lexical ambiguity. The use of a lexically ambiguous predicate can be specified by stating, for each context, the *intended* meaning. It is interesting, however, that at least for some predicates the possible intended meanings are not simply related to the fact that the arguments belong to different domains: they correspond to different *ways of predication*, i.e. different types of subject-predicate relationships, corresponding to meta-level kinds 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" [9]. In practice, formal ontology can be intended as a theory of *a priori distinctions*:

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

The latter kind of distinctions are the subject of the present paper.

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 [10], are states of affairs having an "idealised rational acceptability" [26].

Notation. In the following, we shall use bold capital letters for sets, plain capital letters for predicate symbols and handwritten-style capital letters for relations.

Suppose we have a first-order language $\mathbf{L}$ with signature $\Sigma=<\mathbf{K}, \mathbf{R}>$, where $\mathbf{K}$ is a set of constant symbols, $\mathbf{R}$ is a finite set of n-ary predicate symbols and $\mathbf{P} \subseteq \mathbf{R}$ is the set of monadic predicate symbols. Let $\mathbf{T}$ be a theory of $\mathbf{L}$, $\mathbf{D}$ its intended domain and $\mathbf{M}$ the set of its models $M=<\mathbf{D}, \mathfrak{I}>$, 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 $\mathbf{P}$ which are considered as relevant to the purposes of $\mathbf{T}$ as applied to $\mathbf{D}$. 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 $\mathbf{T}$ into an order-sorted theory $\mathbf{T}_S$ with signature $\Sigma_S=<\mathbf{K}, \mathbf{S}, \mathbf{Q}>$, where $\mathbf{S} \subseteq \mathbf{P}$ is a set of sortal predicates and $\mathbf{Q} = \mathbf{R} \backslash \mathbf{S}$ a set of ordinary predicates. We shall see how this and other distinctions will be expressed in terms of constraints on the set of models.

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

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[1] with signature Σ. The *tense-modal extension* of **L** is the language $\mathbf{L_m}$ obtained by adding to the logical symbols of **L** the usual modal operators $\blacksquare$ and $\blacklozenge$ and the tense operators $\mathbb{F}$ and $\mathbb{P}$, respectively standing for "sometimes in the future" and "sometimes in the past".

Def. 2 Let **L** be a first-order language with signature Σ=<**K**, **R**>, $\mathbf{L_m}$ its tense-modal extension and **D** a domain. A *constant-domain rigid model* for $\mathbf{L_m}$ based on **D** is a structure M = <**W**, $\mathfrak{R}$, $\mathfrak{B}$, **D**, $\mathfrak{F}_K$, $\mathfrak{F}_R$>, where:

- **W** is a set of possible worlds;
- $\mathfrak{R}$ and $\mathfrak{B}$ are binary relations on **W** such that $\mathfrak{B}$ is a union of linear orders and for each $w_i,w_j \in \mathbf{W}$ if $<w_i,w_j> \in \mathfrak{B}$ then $<w_i,w_j> \notin \mathfrak{R}$.
- $\mathfrak{F}_K$ is a function that assigns to each $c \in \mathbf{K}$ an element $\mathfrak{F}_K(c)$ of **D.**
- $\mathfrak{F}_R$ 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 $\mathfrak{F}_R(w,r_n)$ on **D**.[2]

We want to give $\mathfrak{R}$ 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* (intended as real world relations, not as predicate symbols). 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, $\mathfrak{R}$ must be reflexive, transitive and symmetric (i.e., an equivalence relation), and the corresponding modal theory will be therefore S5.

Def. 3 Let **L** be a first-order language, $\mathbf{L_m}$ its tense-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 $\mathfrak{R}$ is the ontological compatibility relation between worlds.

The notion of truth in a model at a world is pretty standard, and it will not defined here in detail because of space limitations. 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 |= $\blacksquare\Phi$) iff M,v |= Φ for every v such that $\mathfrak{R}$(w,v);
- M,w |= $\blacklozenge$ Φ) iff M,v |= Φ for some v such that $\mathfrak{R}$(w,v);
- M,w |= $\mathbb{F}\Phi$ iff M,v |= Φ for some v such that $\mathfrak{B}$(w,v);
- M,w |= $\mathbb{P}\Phi$ iff M,v |= Φ for some v such that $\mathfrak{B}$(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 $\mathbf{L_m}$ 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 constraints by restricting the set of all possible compatibility models of $\mathbf{L_m}$:

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** (**C** |= Φ) 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 **R**'=**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 tense-modal extension of **L**' is called the *ontological extension* $\mathbf{L_o}$ of **L**.

[1] We assume L as non functional just for the sake of simplicity.

[2] This definition is taken from [12], extended with a relation $\mathfrak{B}$ intended to express the temporal precedence relationship between worlds. The latter is a union of linear orders, each of whom represents a possible history. Notice that, due to the fact that $\mathbb{B}$ and $\mathfrak{R}$ are disjoint, modal necessity does not imply temporal necessity.

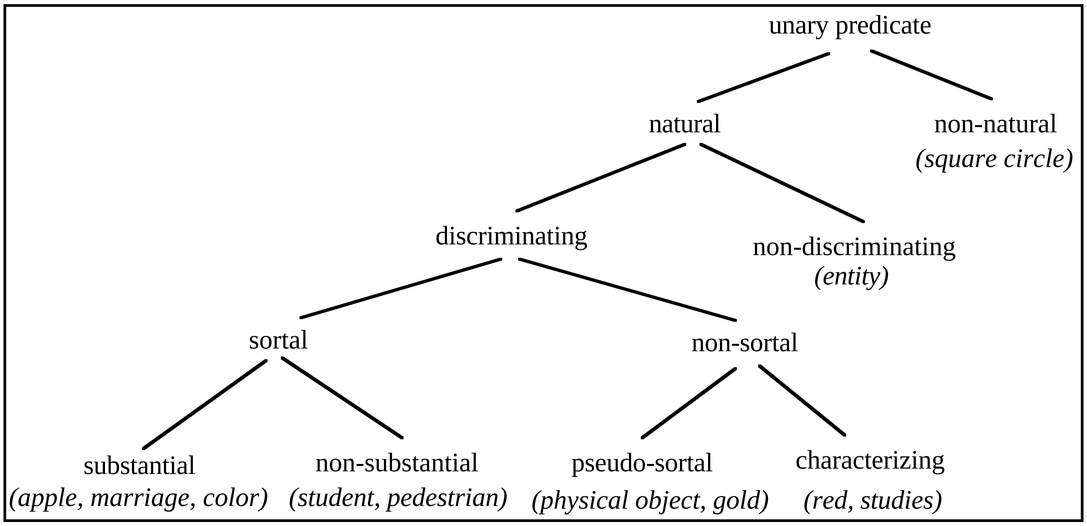

Fig. 2. Preliminary distinctions among unary predicates.

The properties of the part-of relation have been extensively studied in [28]. Connection has been used as a topological primitive in [7] and more recently in [28]. 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 [28]. We assume here that two entities are spatially connected if *their spatio-temporal extensions* are connected in the sense defined in [28] (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 $\mathbf{O}$[2].

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)
D1	$x \leq y =_{def} x < y \vee x = y$	(part)
D2	$Oxy =_{def} \exists z.\ z \leq x \wedge z \leq y$	(overlap)

[1] See [33] for a discussion of the relationsbips between mereology and topology.

[2] Axioms A1-A3 are taken from [29], while A4-A5 from [28].

4 A BASIC ONTOLOGY OF UNARY PREDICATE TYPES

Let us now stipulate some preliminary distinctions among unary predicates (Fig. 2). Notice that we are interested in very general, purely formal distinctions at the meta-level, completely independent on 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, adjectives, or verbs.

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 [8]. Moreover, we are interested in predicates that tell us something non-trivial about the domain, excluding therefore those which are always necessarily true.

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 *natural* in **O** iff **O** $\models \blacklozenge\exists x.Px$. A natural predicate is *discriminating* in **O** iff **O** $\models \blacklozenge\exists x.\neg Px$.

4.1 COUNTABILITY AND REIDENTIFIABILITY

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* [15] and temporal *reidentifiability* [34]. 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 [15], 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 within the ontological extension $\mathbf{L_O}$ of a language **L**:

D3 $\sigma x\phi x =_{def} \iota x\forall y(Oyx \equiv \exists z(\phi z \wedge Ozy))$[1]
(sum of all ϕers)

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)

D7 $\blacksquare_t \phi =_{def} \neg \mathbb{F} \neg\phi \wedge \phi \wedge \neg \mathbb{P} \neg\phi$
(temporal necessity)

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 relation of part-of 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. The following theorem follows immediately from the definition:

Theorem 1 A predicate is countable if it only applies to atomic entities, i.e. entities having no parts.

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 according to theorem 1, assuming that a color has no parts. On the other hand, according to its ordinary sense a predicate like *red* is not countable[2], 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 [30]. 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[3]. 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.

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] 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] 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).

[3] In contrast with [30], we do not assume that detaching an arm is an event such that the arm befor it is not the same arm as the arm after it.

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 [14] 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[1].

Def. 10 A discriminating predicate P is called *temporally stable* under **O** iff **O** $\models \forall x.(Px \supset \mathbb{F}Px \vee \mathbb{P}Px)$.

In conclusion, both mereo-topological and temporal modality are needed to characterize sortal predicates within an ontological commitment:

Def. 11 A discriminating predicate P is called *sortal* in **O** iff it is both countable and temporally stable 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 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 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” 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 [2]:

Def. 12 A discriminating predicate P is *ontologically rigid* in **O** iff **O** $\models \forall x(Px \supset (\blacksquare Px \wedge \blacksquare_t Px))$.

An immediate theorem following from Definition 10 is the following:

Theorem 2 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*[2]. They are all rigid (and therefore stable) but not countable.

Def. 13 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 [34], we call such predicates *non-substantial sortals*[3].

Def. 14 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 [24], p. 10). Finally, in case (3), *red* is used as a contingent property of human-beings and hence is not rigid, while it is countable and temporally stable: *red* is therefore a non substantial sortal.

---

[1] This definition is not completely satisfactory, since, according to the intuition, a temporally stable predicate should hold in a *neighbour* of the time where it is true, but this fact cannot be expressed in terms of $\mathbb{F}$ and $\mathbb{P}$. A more accurate definition would require the use of non-standard modal operators.

[2] They are called "super sortals" in [25]. Notice that *physical object* is not intended here in the sense of spatially isolated *thing.*

[3] 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 [31]), while the terms *class* or *concept* should be reserved to the union of sortal and pseudo-sortal predicates.

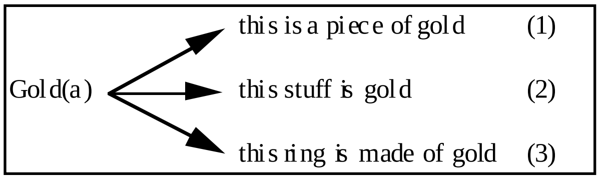

Fig. 3: Different interpretations of mass nouns.

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, 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 can be assumed as non-rigid, and *gold* becomes a characterizing predicate.

5 ONTOLOGICAL ENGINEERING

We would like to show in this section how the ontological distinctions introduced above can be of concrete utility in the current practice of knowledge engineering. The first result of the formal framework presented above is the possibility to draw a clear distinction between concepts[1] 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[2], 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 to our ontological commitment, which lead to the notion of *well-founded ontological commitment*.

Def. 15 Let P and Q be two natural predicates in **O**. P is *subordinate* to Q in **O** iff **O** $\models \forall x(Px \supset Qx) \wedge \neg\forall x(Qx \supset Px)$. P and Q are *disjoint* in **O** iff **O** $\models \neg\exists x.Px \wedge Qx$. A set $\mathbf{P}=\{P_1, ..., P_n\}$ of mutually disjoint natural predicates in **O** is a *domain partition* in **O** iff **O** $\models \forall x.(P_1x \vee ... \vee P_nx)$.

[1]The term *concept* is often used interchangeably with *type*, but we deserve to the latter a more specific meaning (see below).

[2]A similar proposal has been made by Sowa [31], which however refers to an unspecified notion of "natural type".

Def. 16 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[3].
- For each categorial predicate $C \in \mathbf{C}$, there is a set $\mathbf{S}_C$ of disjoint substantial sortals such that, for each $S \in \mathbf{S}_C$, S is subordinate to C and there is no substantial sortal S' such that S is subordinate to S'.
- 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*[4]. 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[5], except that top-level substantial sortals should specify natural kinds within general categories: therefore, they must be disjoint and cannot overlap general categories. A useful definition related to substantial sortals is the following one:

Def. 17. Let **O** be a well-founded ontological commitment. If a substantial sortal S is subordinate to another substantial sortal T under **O**, then S is called a *kind* of T.

[3]A possible further constraint for categorial predicates could be **O** $\models \forall xy((Cx \wedge y<x) \supset Cy)$.

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

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

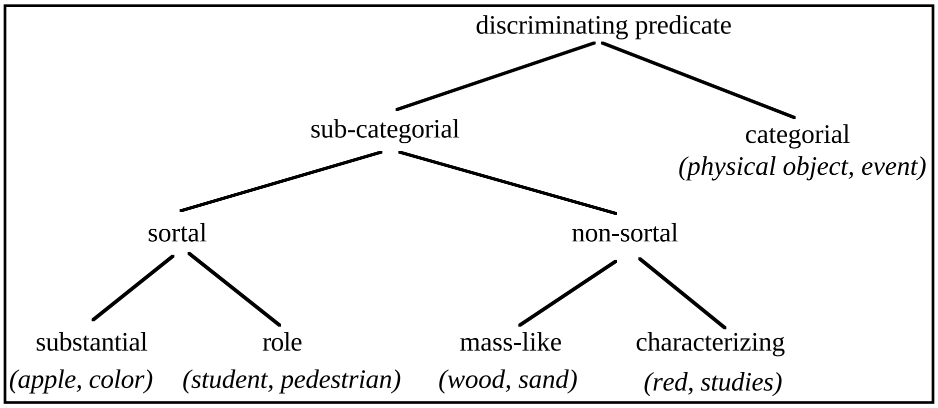

Fig. 4. Basic distinctions among discriminating predicates within a well-founded ontological commitment

Finally, the intuition behind the third constraint in Def. 16 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" [31]. We suggest to adopt the term "role" for non-substantial sortals within the KR community, avoiding to use it at a synonym for an (arbitrary) binary relation as common practice in the KL-ONE circles.[1] A useful theorem following from Def. 16 is the following one:

Theorem 3. Within a well-founded ontological commitment, any two overlapping non-substantial sortals are subordinate to the same substantial sortal.

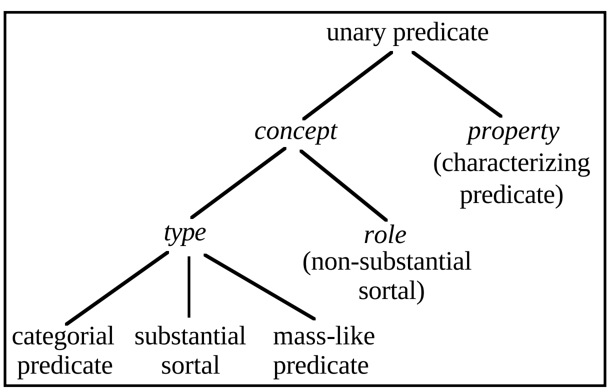

Fig. 5. A terminological proposal for KR formalisms. Commonly used KR terms are shown in italics.

[1]See [18] 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.

On the basis of the above considerations related to the practice of knowledge engineering, we are now in the position to formulate a terminological proposal regarding the relationship between the terminology currently used in KR formalisms and the philosophical terms we have defined here (Fig. 5 above). Rigid (both countable and uncountable) unary predicates are called *types*, while as noticed before non-substantial sortals correspond to *roles*. Types and roles are collectively called *concepts*, and are distinguished from *properties* since the latter are characterizing predicates, i.e. they are uncountable and non-rigid. Notice that we prefer to speak of *properties* rather than of *qualities*, since it seems more appropriate to adopt the latter term for substantial sortals having universals as arguments, as for instance *color*.

6 CONCLUSIONS

In [3], the authors discussed the example reported in Fig. 6 above. 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

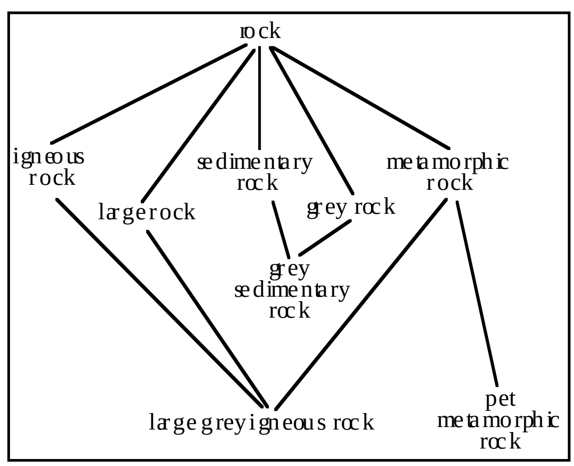

Fig. 6. Kinds of rocks (From [3])

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 [4], 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”).

We hope to have shown in this paper that 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”.

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 some of these laws are suitable to formal characterization: independently of the particular formalization we have adopted, which can be of course changed or revised, we would like to stress that the ontological distinctions we have introduced can have a profound impact on the current practice of knowledge engineering.

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 cooperation 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] W. P. Alston 1958. Ontological Commitments. *Philosophical Sttudies*, **9**: 8-16.
[2] R. Barcan Marcus 1968. Essential Attribution. *The Journal of Philosophy*, **7**.
[3] R. Brachman, R. Fikes, and H. Levesque 1983. Krypton: A Functional Approach to Knowledge Representation. *IEEE Computer*, October: 67-73.
[4] R. Brachman and H. Levesque 1982. Competence in Knowledge Representation. In *Proceedings of AAAI 82*.
[5] 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.
[6] A. Church 1958. Ontological Commitment. *The Journal of Philosophy*, **55**: 1008-14.
[7] B. L. Clarke 1981. *A Calculus of Individuals Based on "Connection"* **22**.
[8] N. Cocchiarella 1989. Philosophical Perspectives on Formal Theories of Predication. In D. Gabbay and F. Günthner (ed.), *Handbook of Philosophical Logic*. Reidel: 253-326.
[9] N. B. Cocchiarella 1991. Formal Ontology. In H. Burkhardt and B. Smith (ed.), *Handbook of Metaphysics and Ontology*. Philosophia Verlag, Munich.
[10] N. B. Cocchiarella 1993. Knowledge Representation in Conceptual Realism. In N. Guarino and R. Poli (ed.), *Formal Ontology in Conceptual Analysis and Knowledge Representation.* Kluwer (in preparation).
[11] R. Davis, H. Shrobe, and P. Szolovits 1993. What is in a Knowledge Representation? *AI Magazine*, Spring 1993.
[12] 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.
[13] M. R. Genesereth and N. J. Nilsson 1987. *Logical Foundation of Artificial Intelligence.* Morgan Kaufmann, Los Altos, California.
[14] T. Givón 1979. *On Understanding Grammar*. Academic Press, New-York.
[15] N. Griffin 1977. *Relative Identity*. Oxford University Press, Oxford.

[16] T. Gruber 1993. 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*. Kluwer (in preparation).

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

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

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

[20] N. Guarino and L. Boldrin Formal Ontology and Knowledge Representation. In N. Guarino and R. Poli (ed.), *Formal Ontology in Conceptual Analysis and Knowledge Representation*. Kluwer (in preparation).

[21] 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.

[22] D. Israel 1983. Interpreting Network Formalisms. *International Journal of Computers and Mathematics*, **9** 1.

[23] 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.

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

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

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

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

[28] 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.

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

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

[31] 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.

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

[33] 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.

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