Identity and Subsumption

Nicola Guarino[1] and **Christopher Welty**[2]

LADSEB-CNR Internal Report 01/2001

Final version - August 7, 2001

Abstract. The intuitive simplicity of the so-called *is-a* (or *subsumption*) relationship has led to widespread ontological misuse. Where previous work has focused largely on the semantics of the relationship itself, we concentrate here on the ontological nature of its *arguments*, in order to tell whether a single *is-a* link is ontologically well-founded. For this purpose, we introduce some techniques based on the philosophical notions of *identity, unity,* and *essence*, which have been adapted to the needs of taxonomy design. We demonstrate the effectiveness of these techniques by taking real examples of poorly structured taxonomies, and revealing cases of invalid generalization.

1 Introduction

Taxonomies based on a partial-ordering relation commonly known as *is-a*, class inclusion, or subsumption have become an important conceptual modeling tool for database schemas, knowledge-based systems, and semantic lexicons. Properly structured taxonomies help bring substantial order to conceptual models, are particularly useful in presenting limited views for human interpretation, and play a critical role in reuse and integration tasks. Improperly structured taxonomies have the opposite effect, making models confusing and difficult to reuse or integrate.

Many previous efforts at providing some clarity in organizing taxonomies have focused on the semantics of the subsumption relationship [Brachman 1983], on various kinds of related relations (generalization, specialization, subset hierarchy) [Storey 1993], or on its role in the more general framework of data abstractions [Goldstein and Storey 1999]. Our approach differs in that we focus on the arguments (i.e. the properties) involved in the subsumption relationship, rather than on the semantics of the relationship itself. The latter is taken for granted, as we take the statement "ψ subsumes φ" to mean that, *necessarily*:

$$x\ \varphi(x) \rightarrow \psi(x) \tag{1}$$

The modal reading of the above formula is an important qualification: we take subsumption as an *ontological constraint*, therefore we assume that it must hold in all possible worlds. Indeed, we believe that modal necessity is what distinguishes – within a particular conceptualization – an ontological truth from a contingent assertion. So we focus here on *necessary subsumption.* Our task will be to verify its plausibility on the basis of the *ontological nature* of its arguments.

1. LADSEB-CNR, Corso Stati Uniti 4, I-35127 Padova, Italy, email: Nicola.Guarino@ladseb.pd.cnr.it

2. Vassar College, Poughkeepsie, NY, USA, email: welty@cs.vassar.edu.

This chapter is organized as follows. First, we clarify some major issues lying behind the generic notion of identity, and the related notions of unity and essence. Then we show how these notions impose ontological constraints on the subsumption relationship, and discuss some concrete examples of problems and misconceptions concerning subsumption taxonomies.

Most of the ideas discussed here have been introduced in previous papers [Guarino and Welty 2000a, Guarino and Welty 2000b, Guarino and Welty 2001]. The present work does however refine and simplify most of the core definitions presented in the past, and aims at offering a self-contained overview for what concerns the ontological analysis of the subsumption relationship.

2 Identity, Unity, and Essence

Identity is one of the most fundamental notions in ontology, yet the related issues are very subtle, and isolating the most relevant ones is not an easy task; see [Hirsch 1982] for an account of the identity problems of ordinary objects, and [Noonan 1993] for a collection of philosophical papers in this area. In particular, the relationship between *identity* and *unity* appears to be crucial for our interest in ontological analysis. These notions are different, albeit closely related and often confused under a generic notion of identity. Strictly speaking, identity is related to the problem of distinguishing a specific instance of a certain class from other instances by means of a *characteristic property*, which is unique for *it* (that *whole* instance). Unity, on the other hand, is related to the problem of distinguishing the *parts* of an instance from the rest of the world by means of a *unifying relation* that binds them together (not involving anything else).

For example, asking "Is that my dog?" would be a problem of identity, whereas asking "is the collar part of my dog?" would be a problem of unity. As we shall see, the two notions are complementary: when something can be both recognized as a whole and kept distinct from other wholes then we say that it is an *individual*, and can be counted as *one*.

The actual conditions we use to support our answers to these questions for a certain class of things vary from case to case, depending on the properties holding for these things. If we find a condition that consistently supports identity or unity judgements for *all* instances of a certain property, then we say that such property *carries* an identity or a unity condition.

These notions encounter problems when time is involved. The classical one is that of *identity through change*: in order to account for common sense (i.e. for the way we, as cognitive agents, interact with the world around us), we need to admit that an individual may keep its identity while exhibiting different properties at different times. But which properties can change, and which must not? And how can we reidentify an instance of a certain property after some time? The former issue leads to the notion of an *essential property*, on which we base the definition of *rigidity*; the latter is related to the distinction between *synchronic* and *diachronic* identity. Both issues will be discussed below.

Before going on, it is important to make clear that all the assumptions related to the notions above depend on our *conceptualization* of the world [Guarino 1998]. For example, the decision as to whether cats keep their identity after losing their tail, or whether statues are identical with the marble they are constituted of, are ultimately the result of our sensory system, our culture, and so on. The examples we shall use in this paper concerning the ontological nature of certain properties (e.g. STUDENT, SPHERICAL) are merely indicative of our own intuitions. The aim of the present analysis is not so much to discuss these assumptions, but rather to clarify the formal tools needed to make them explicit, and to explore their logical consequences. These formal tools form the core of a methodology for *ontology-driven conceptual modeling* called OntoClean, which is discussed in more detail elsewhere [Guarino and Welty 2001].

3 The Formal Framework

In this section we present a formal analysis of the basic notions discussed above, and we introduce a set of *meta-properties* that represent the behaviour of a property with respect to these notions. Our goal is to show how these meta-properties impose relevant ontological constraints on the subsumption relationship.

In the following, we shall denote meta-properties by bold letters preceded by the sign "+", "-" or "~", whose meaning will be described for each meta-property. We use the notation $\varphi^{\mathbf{M}}$ to indicate that the property φ has the meta-property **M**.

We shall adopt a simple temporal logic, where all predicates are temporally indexed by means of an extra argument. If the time argument is omitted for a certain predicate P, then the predicate is assumed to be time invariant, that is $\; tP(x,t) \rightarrow \; tP(x,t)$. Our domain of quantification will be that of *possibilia*: this means that we include all possible entities, independent of their actual existence [Lewis 83]. Therefore, we shall quantify over a constant domain in every possible world. Worlds will be considered "histories" (temporally ordered sequences of maximal states of affairs) rather than "snapshots", and we shall consider all of them as equally accessible. As a result, we shall adopt the simplest quantified modal logic, namely S5 plus the Barcan Formula [Hughes and Cresswell 1996].

For example, the prroperty *UNICORN* will not be empty in our world, although no instance has actual existence there. Actual existence is therefore different from existential quantification ("logical existence"), and will be represented by the temporally indexed predicate E*(x,t)*, meaning that *x* has actual existence at time *t* [].

Finally, in order to avoid trivial cases in our meta-property definitions, we shall implicitly assume the property variables as restricted to *discriminating properties*, properties φ such that $\diamond \; x\, \varphi(x) \quad \diamond \; x\, \neg\varphi(x)$. In other words, discriminating properties are properties for which there is possibly something which exhibits that property, and possibly something that does not exhibit that property; they are neither tautological nor vacuous.

3.1 Essential Properties and Rigidity

Before addressing the core issues of identity and unity, it is useful to clarify the notion of *essential property*, and to introduce a set of related meta-properties.

Definition 1 A property φ holding for a certain individual *a* in a certain state of affairs at time t is said to be *essential* to *a* iff it necessarily holds for *a* at every possible time in every possible world, i.e.

$$\square \; t\varphi(a, t) \qquad (2)$$

Examples of essential properties for a human being would be *PERSON* and *HAVING A BRAIN.*

Definition 2 A property φ is *rigid* iff, necessarily, it is essential to *all* its instances, i.e.

$$\Box \quad xt\,(\varphi(x,t) \rightarrow \Box \quad t'\varphi(x,t')) \tag{3}$$

For example, *HAVING A BRAIN* would be essential to human beings but not to all its instances, since it is not essential to -say- a dead corpse. On the other hand, *PERSON* can be safely be taken as rigid. Other examples of non-rigid properties could be *TIRED* or *STUDENT*.

Definition 3 A property φ is *non-rigid* iff it is not rigid, that is there is at least one instance such that φ is not essential to it:

$$\Diamond \quad x\,(\varphi(x,t) \quad \Diamond \quad t'\neg\varphi(x,t')) \tag{4}$$

Non-rigidity can be further restricted as follows:

Definition 4 A property φ is *anti-rigid* iff all its instances are such that φ is not essential to them:

$$\Box(\quad xt\;\varphi(x,t) \rightarrow \Diamond \quad t'\neg\varphi(x,t')) \tag{5}$$

Consider for example the properties *SPHERICAL* and *STUDENT*: the former is non-rigid but not anti-rigid, since it may be the case that something (like a lump of clay) is spherical by accident, but it is also possible that there are things that are essentially spherical (spheres, for instance). The *STUDENT* property, on the other hand, appears to be anti-rigid, since it is always possible for a student to become a non-student while being the same individual.

Rigid properties are marked with the meta-property +**R**, non-rigid properties are marked with -**R**, and anti-rigid properties with ~**R**. Note that rigidity as a meta-property is not "inherited" by sub-properties of properties that carry it, so the markings *PERSON*$^{+\mathbf{R}}$ and *STUDENT*$^{\sim\mathbf{R}}$ are perfectly consistent with the fact that *PERSON* subsumes *STUDENT*.

3.2 Identity and Identity Conditions

Before discussing the formal structure of identity conditions (ICs), some clarifications about their intuitive meaning may be useful. If we say, "Two persons are the same if they have the same SSN," we seem to create a puzzle: how can they be *two* if they are the same? The puzzle can be solved by recognizing that two (incomplete) descriptions of a person (like two records in different databases) can be different while referring to the same individual. The statement "two persons are the same" can be therefore rephrased as "two descriptions of a person refer to the same object". A description can be seen as a set of properties that apply to a certain object. Our intuition is that two incomplete descriptions denote the same object if they have an identifying property in common.

Depending on whether the two descriptions hold at the same time, we distinguish between *synchronic* and *diachronic* ICs. The former are needed to tell, e.g., whether the statue is identical with the marble it is made of, or whether a hole is identical with its filler [Casati and Varzi, 1998], while the latter allow us to re-identify things over time.

In the philosophical literature, an *identity criterion* is generally defined as a condition that is both necessary and sufficient for identity. According to [Strawson, 1959], a property φ *carries* an IC iff the following formula holds for a suitable ρ:

$$\varphi(x) \quad \varphi(y) \rightarrow (\rho(x,y) \leftrightarrow x{=}y) \tag{6}$$

Since identity is an equivalence relation, it follows that ρ restricted to φ must also be an equivalence relation. For example, the property *PERSON* can be seen as carrying an IC if relations like *having-the-same-brain* or *having-the-same-SSN* are assumed to satisfy (6).

Properties carrying an IC are called *sortals* [Strawson 1959]. In many cases, their linguistic counterparts are *nouns* (e.g., Apple), while non-sortals correspond to *adjectives* (e.g., Red). Distinguishing sortals from non-sortals is of high practical relevance for conceptual modeling, as we tend to naturally oranize knowledge around nouns.

When trying to use (6) for conceptual modeling purposes we encountered however a number of problems. First, the nature of the ρ relation remains mysterious: what makes it an IC, and how can we index it with respect to time to account for the difference between synchronic and diachronic identity? Second, it only accounts for identity *under a certain property*: in principle it may be that $x=y$ without being $\varphi(x,t) \quad \varphi(y,t')$ (with $t \neq t'$), but the formula (6) does not help us in this case. Third, deciding whether a property carries an IC or not may be difficult, since finding a ρ that is both necessary *and* sufficient for identity is often hard, especially for natural kinds and artifacts.

3.3 The "Sameness" Relation

Our intuition is that the nature of the ρ relation in (6) is based on the "sameness" of a certain *property*, which is unique to a specific instance. Suppose we stipulate, e.g., that two persons are the same iff they have the same brain: the reason why this *relation* can be used as an IC for persons lies in the fact that a property like "having this particular brain" is an *identifying property*, since it holds exactly for one person.

Identifying properties can be seen as relational properties, involving a *characteristic relation* between a class of individuals and their *identifying characteristics*[1]. In the above example, brains are taken as identifying characteristics of persons. Such characteristics can be internal to individuals themselves (parts or qualities) or external to them (other "reference" entities). So two things can be the same because they have some parts or qualities in common, or because they are related in the same way to something else (for instance, we may want to say that two material objects are the same if they occupy the same spatial region). Of course, an individual's characteristic cannot be identical to the individual itself, so a characteristic relation must be *irreflexive*.

This means that, if χ denotes a suitable characteristic relation for φ, we can assume:

$$\rho(x,y) = \quad z(\chi(x,z) \leftrightarrow \chi(y,z)) \tag{7}$$

The scheme (6) becomes therefore:

$$\varphi(x) \quad \varphi(y) \rightarrow (\quad z(\chi(x,z) \leftrightarrow \chi(y,z)) \leftrightarrow x=y) \tag{8}$$

For instance, if we take φ as the property of being a set, and χ as the relation "has-member", this scheme tells us that two sets are identical iff they have the same members.

An important advantage of (8) over (6) is that it is based on a characteristic relation χ holding separately for x and y, rather than on a relation ρ holding between them. This

1.Of course, this characteristic relation must be defined for each instance of the class. This means that fingerprints cannot be used as identifying characteristics for persons, since a person may have no fingerprints (while we can assume that each person must have a brain).

allows us to take time into account more easily, clarifying the distinction between *synchronic* and *diachronic* identity:

$$E(x,t) \wedge E(y,t') \wedge \varphi(x,t) \wedge \varphi(y,t') \rightarrow (\forall z(\chi(x,z,t) \leftrightarrow \chi(y,z,t')) \leftrightarrow x=y) \qquad (9)$$

Since we have now explicitly introduced time, it seemed safe to restrict our analysis to the case where the entities to be identified do actually exist, avoiding identity problems related to non-actually existing entities. So the above formula means that, if x and y do actually exist at times t and t' respectively, and they are both instances of the property φ, then they are identical if and only if they share the same characteristics. We shall have a synchronic criterion if $t=t'$, and a diachronic criterion otherwise.

Note that accounting for the difference between synchronic and diachronic identity would have been difficult with (6): we may have tried adding two temporal arguments to the ρ relation, but in this case its semantics would have become quite unclear, being a relation that binds together two entities at different times. Note also that synchronic identity criteria are weaker than diachronic ones. For instance, the sameness of spatial location is usually adopted as a synchronic identity criterion for material objects, but of course it does not work as a diachronic criterion.

A possible criticism of (9) is that it looks circular, since it defines the identity between x and y in terms of the identity between something else (in this case, the identifying characteristic(s) z common to x and y). However, as observed by Lowe ([Lowe 1998], p. 45), we must keep in mind that *ICs are not definition*s, as identity is a primitive.This means that the circularity of identity criteria with respect to the very notion of identity is just a fact of life: identity can't be defined. Rather, we may ask ICs to be *informative*, in the sense that identity conditions must be non-circular with respect to the properties involved in their definition. For instance, Lowe points out that Davidson's identity criterion for events, stating that two events are the same if they have the same causes and they are originated by the same causes [Davidson, 1980], is circular in this sense since it presupposes the identity of causes, which are themselves events. In many cases, however, even this requirement cannot easily be met, and we must regard ICs as simple constraints.

For brevity, the formula (9) above can be rewritten as:

$$E(x,t) \wedge E(y,t') \wedge \varphi(x,t) \wedge \varphi(y,t') \rightarrow (\Sigma\chi(x,y,t,t') \leftrightarrow x=y) \qquad (10)$$

where $\Sigma\chi$ is a *sameness formula* depending on χ, defined as

$$\Sigma\chi(x,y,t,t') =_{def} \forall z(\chi(x,z,t) \leftrightarrow \chi(y,z,t')) \qquad (11)$$

We may conclude therefore that an IC for φ is a sameness formula $\Sigma\chi$ that satisfies (10) and is based on a suitable characteristic relation χ. It is safe however to make sure that the IC really depends on φ, imposing a *non-triviality constraint* such as:

$$\neg \forall xy(\Sigma\chi(x,y,t,t') \leftrightarrow x=y) \qquad (12)$$

For instance, suppose that χ is the *proper-part* relation: in this case the statement $\forall xy(\Sigma\chi(x,y,t,t') \leftrightarrow x=y)$ would represent the extensionality principle, which says that two (non-atomic) entities are the same iff they have the same proper parts. Without the constraint (12), any property holding for non-atomic entities would trivially carry an identity criterion if the extensionality principle was assumed.[1]

3.4 Local and global Identity Conditions

The formula (10) above may hold for *rigid* properties like *PERSON*, as well as for *non-rigid* properties like *STUDENT* or *TIRED* . In the latter case, $\Sigma\chi$ may act only as *local* IC for φ, as we can't be sure it also accounts for identity among entities that are not both instances of φ. Consider e.g. the properties *CATERPILLAR* and *BUTTERFLY*: in this case, a formula $\Sigma\chi$ based (for instance) on the sameness of a certain wing pattern could count as an IC for butterflies, but it would not account for the identity relation holding, at different times, between butterflies and caterpillars.

It seems useful therefore to distinguish between *local* and *global* ICs. The following formal definitions are intended to account for such a distinction, as well as for the related issues concerning the way ICs are inherited along subsumption hierarchies. They refine some previous definitions reported in [Guarino and Welty 2000a, 2000b, 2001].

Definition 5 Let φ be a property, and $\chi(x,z,t)$ a non-trivial characteristic relation satisfying (12), and such that $\Box(\varphi(x, t) \rightarrow \quad z\chi(x, z, t))$. Then

- φ *carries a local identity condition* $\Sigma\chi$ iff, necessarily:

$$E(x,t) \quad E(y,t') \quad \varphi(x,t) \quad \varphi(y,t') \rightarrow (\Sigma\chi(x,y,t,t') \leftrightarrow x=y) \qquad (13)$$

- φ *carries a global identity condition* $\Sigma\chi$ iff, necessarily:

$$E(x,t) \quad E(y,t') \quad \varphi(x,t) \rightarrow (\Sigma\chi(x,y,t,t') \leftrightarrow x=y). \qquad (14)$$

Since ICs can be inherited along subsumption hierarchies, it is useful to distinguish between *supplying* an IC and just *carrying* it:

Definition 6 A property φ is a *sortal* iff it carries a (local or global) IC. Sortals are marked with the metaproperty **+I.**

Definition 7 A property $\varphi \quad$ **O** *supplies* a (local or global) IC $\Sigma\chi$ in **O** iff (i) **O** is the set of *explicit properties* introduced in a certain ontology (i.e., corresponding to predicate names);(ii) φ carries $\Sigma\chi$; and (iii) $\Sigma\chi$ is not carried by all the properties in **O** directly subsuming φ. This means that, if φ inherits different (but compatible) ICs from multiple properties, it still counts as supplying an IC (see section 3.6).

Properties supplying global identity are marked with the metaproperty **+G**[1]; those supplying local identity are marked with **+L**.

Definition 8 A property φ is a *type* iff it supplies a global IC.

Let us now introduce an important principle, adapted from [Lowe 89]:

Sortal Expandability Principle (SEP). If two entities are identical, they must be instances of a common sortal that *accounts* for their identity, i.e. it carries an IC they satisfy.

1.Note that other trivial identity conditions, such as those discussed in [Kaplan 2001] and [Carrara and Giaretta 2001], are excluded due to the irreflexivity constraint imposed on χ, which, taking time into account, corresponds to assuming $t\neg \quad x\chi(x,x,t)$. In particular, we believe that this constraint eliminates many trivial instances of $\Sigma\chi(x,y,t,t')$ that imply the indentity of *x* and *y*.

1.This corresponds to the **+O** (own identity) mark in previous papers

We also add a further principle, which seems to be a plausible constraint for a *well-founded ontology*:

***Unique Sortal Principle*:** if an IC holds between two entities, there must be a *unique* sortal that *supplies* such IC.

On the basis of these principles, we can prove the following:

Theorem 1 In a well-founded ontology, *types are rigid properties*.

To see this, let's suppose that a type φ, supplying a certain IC $\Sigma\chi$, is not rigid. This means that, for some x, y and t, with $x=y$, there are two worlds w_1 and w_2 such that $\varphi(x,t)$ holds in w_1 and $\neg\varphi(y,t')$ holds in w_2. Because of the Unique Sortal Principle, there must be a unique sortal ψ that supplies $\Sigma\chi$. But, by hypothesis, $\Sigma\chi$ is supplied by φ, so it must be $\psi=\varphi$. This means that φ must be rigid, contradicting the original hypothesis.

Recognizing types is of utmost importance in ontology design. Since they must be rigid in order to supply global IC, they represent *invariant properties* that characterize the nature of a domain element by supplying identity criteria to it. If we assume Quine's motto "No entity without identity," this implies that every element of our domain of discourse must be an instance of a type. Types and other property kinds defined on the basis of the metaproperties discussed here have been presented in more detail (although with a few formal differences) in [Guarino and Welty 2000a].

3.5 Heuristics for Identity

Unfortunately, recognizing that a property carries a *specific* IC is often difficult in practice. However, in many cases it suffices to recognize whether a property carries *some* IC (being therefore a sortal) or not , without telling exactly *which* IC. In these cases, we may want to check for some *minimal* ICs, which are (only)-necessary or (only)-sufficient for identity, but close enough to "true" ICs. If none of these weak conditions holds for a given property, we may safely conclude it is not a sortal. Otherwise, we may have some heuristic evidence that some "true" IC exists.

Only-necessary and only-sufficient conditions for global identity can be defined by considering separately the two senses of the double implication in (13). We need however to assume in advance that φ is a rigid property, and to make sure to exclude trivial cases (in the following, (16) is needed to guarantee that the second literal in (15) is relevant and not tautological).

Definition 9 A *necessary global identity condition* for a rigid property φ is a formula $\Sigma\chi$, satisfying (11) and (12), such that:

$$\varphi(x,t) \quad x=y \rightarrow \Sigma\chi(x,y,t,t') \tag{15}$$

$$\neg \quad xy(\varphi(x,t) \quad \varphi(y,t') \rightarrow \Sigma\chi(x,y,t,t')) \tag{16}$$

Definition 10 A *sufficient identity condition* for a rigid property φ is a formula $\Sigma\chi$, satisfying (11) and (12), such that:

$$\varphi(x,t) \quad \Sigma\chi(x,y,t,t') \rightarrow x=y \tag{17}$$

Besides being useful for recognizing sortals, minimal ICs have also a practical relevance in taxonomy design, since of course they also follow the Identity Disjointness Constraint below. As we shall see, this means that in practice we can assume a property carries identity on the basis of the evidence given by the minimal ICs, and use them in place of true ICs to constrain the taxonomy.

3.6 Inheriting Multiple Identity Conditions

As described above, ICs can be inherited along subsumption hierarchies. They can also *specialize* along hierarchies, in the sense that new identity criteria can be supplied by a given property in addition to those inherited from the subsuming properties. Consider for instance the domain of abstract geometrical figures, where the property *POLYGON* subsumes *TRIANGLE.* A necessary and sufficient IC for polygons is “Having the same edges and the same angles”. On the other hand, an *additional* necessary and sufficient IC for triangles is “Having two edges and their internal angle in common” (note that this condition is only-necessary for polygons). Again, the property *EQUILATERAL TRIANGLE* may inherit from *REGULAR POLYGON* the IC “Having the same edges”, while also inheriting the ICs carried by *TRIANGLE*.

As this example shows, nothing prevents sorts from using multiple inheritance to form tangled hierarchies, at least in principle. However, the presence of explicit ICs attached to them imposes an important (as well as natural) constraint on the inheritance of multiple ICs, which we shall call the *Identity Disjointness Constraint* (IDC):

Properties carrying incompatible ICs are necessarily disjoint (18)

We shall see in the following how this simple principle, whose philosophical implications have been discussed in the seminal work by Lowe [Lowe 1989, Lowe 1998], has a deep impact on apparently innocent taxonomic assumptions.

3.7 Unity

We have discussed and formalized the notion of unity in some detail in previous work [Guarino and Welty 2000b]. At the time the present paper was being finalized, a new formalization was proposed [Gangemi et al. 2001], which overcomes some difficulties of the previous one[1]. Since such new formalization does not affect the main point of this paper, which is about the constraints that identity and unity conditions impose on the subsumption relation, we stick here to our former formalization, which seems also easier to grasp.

The notion of unity is closely tied to that of parthood, so that we need to introduce some basic axioms and definitions. We adopt a time-indexed mereological relation P(x,y,t), meaning that x is a (proper or improper) part of y at time t, satisfying the minimal set of axioms and definitions (adapted from [Simons 1987], p. 362) shown in Table 2. Differently from Simons, this mereological relation will be taken as completely general, holding on a domain which includes individuals, collections, and amounts of matter.

Definition 11 At a given time t, an entity x is a *whole under* R if R is an equivalence

1.Consider the following counterexample: suppose you want to say that all the children a, b, c of a certain person form a whole. So all the parts of a+b+c must be linked together by the unifying relation "having the same parent". But two of them, namely a+b and b+c, are not linked by such relation, since they are not persons. Another problem is linked to the fact that the previous definition excludes the possibility of overlapping of entities that are wholes [Kaplan 2001].

Table 2. Axiomatization of the part relation.

$PP(x,y,t) =_{def} P(x,y,t) \quad \neg x{=}y$	(proper part)
$O(x,y,t) =_{def} \quad z(P(z,x,t) \quad P(z,y,t))$	(overlap)
$P(x,y,t) \quad P(y,x,t) \rightarrow x{=}y$	(antisymmetry)
$P(x,y,t) \quad P(y,z,t) \rightarrow P(x,z,t)$	(transitivity)
$PP(x,y,t) \rightarrow \quad z(PP(z,y,t) \quad \neg O(z,x,t))$	(weak supplementation)

relation (called *unifying relation*) such that:

$$y(P(y,x,t) \rightarrow \quad z(P(z,x,t) \leftrightarrow R(z,y,t))) \tag{19}$$

$$\neg(PP(y,x,t) \quad PP(z,x,t) \leftrightarrow R(z,y,t)) \tag{20}$$

We can read the above definition as follows: *at time t, each part of x must be bound by R to all other parts and to nothing else.* (19) expresses a condition of *maximal self-connectedness* according to a suitable relation of “generalized connection” *R*. (20) is a non-triviality condition on *R*, that avoids considering any mereological sum as a contingent whole.

Depending on the ontological nature of such relation, we may have different kinds of unity. For example, we may distinguish *topological unity* (a piece of coal, a lump of coal), *morphological unity* (a ball, a constellation), *functional unity* (a hammer, a bikini). As a further example, an atomic object (i.e., an object with no proper parts), is a whole under the identity relation. As these examples show, nothing prevents a whole from having parts that are themselves wholes (under different unifying relations). Indeed, a *plural whole* can be defined as a whole which is a mereological sum of wholes.

According to Definition 11, an entity may be a whole only in a particular possible world, at a certain time. Consider for instance an isolated piece of clay. This certainly has a certain topological unity, which is however lost as soon as we attach it to a much larger piece: the original piece of clay is not a whole any more, while the new piece is.

A stronger and more useful notion of whole can be introduced by assuming that the same conditions for unity must *necessarily* hold for an object, i.e. by assuming unity as an *essential* property:

Definition 12 An entity *x* is an *essential whole under R* if, necessarily, it is always a whole under *R*.

We are now in the position to state the following:

Definition 13 A property φ *carries a unity condition* (UC), or simply *carries unity*, if there is a common unifying relation *R* such that all its instances are essential wholes under *R*. Properties carrying unity are marked with the meta-property **+U** (**-U** otherwise).

Within properties that do not carry unity, we distinguish properties that do not carry a *common* UC for all their instances from properties all of whose instances are not wholes. An example of the former kind may be *LEGAL AGENT*, all of whose instances are wholes, although with different UCs (some legal agents may be people, some companies). *AMOUNT OF MATTER* is usually an example of the latter kind, since none of its instances can be wholes (assuming that a single molecule does not count as an *amount* of matter). Therefore we define:

Definition 14 A property φ carries *anti-unity* (marked with the meta-property **~U**) if no instance of it is an essential whole.

4 Constraints on Subsumption

Let us see now how identity and unity affect the subsumption relationship. Our point is to check the consistency and the ontological plausibility of a subsumption relationship between two properties on the basis of their behavior with respect to identity and unity.

A first important constraint has been introduced in Section 3.6. It follows that a constraint similar to the Identity Disjointness Constraint holds also for Unity, so that we can state:

Properties with incompatible ICs/UCs are necessarily disjoint. (21)

Indeed, the statement above is just the consequence of the fact that having a certain IC or UC is an *essential property,* and incompatible essential properties must be necessarily disjoint. In many cases, just considering the essential properties associated to a given property (independently of any considerations related to identity or unity) is enough to conclude that a certain subsumption link is invalid, just because the two arguments are associated to incompatible essential properties [Akiba 2000]. For instance, to see that a vase is not an amount of matter, we may just consider that the property of having a certain shape is essential for vases, while the same property is "anti-essential" for amounts of matter, in the sense that any amount of matter can possibly have a different shape.

Besides this, we have several constraints involving the meta-properties we have introduced. Let us represent with $\neg(\psi \rightarrow \varphi)$ a constraint stating that property φ can't subsume ψ. The following constraints are a simple consequence of our definitions:

$$\neg(\psi^{+R} \rightarrow \varphi^{-R}) \quad (22)$$

$$\neg(\psi^{-I} \rightarrow \varphi^{+I}) \quad (23)$$

$$\neg(\psi^{-U} \rightarrow \varphi^{+U}) \quad (24)$$

$$\neg(\psi^{+U} \rightarrow \varphi^{\sim U}) \quad (25)$$

5 Some Problematic Subsumption Relationships

Let us finally examine some examples of subsumption relationships in the light of the above discussion (Table 3). All these examples appear acceptable at a first sight, but immediately become problematic as soon as identity and unity are taken into account.

Table 3: Problematic subsumption relationships in some current ontologies.

1	A physical object is an amount of matter	Pangloss
2	An amount of matter is a physical object	WordNet
3	An organization is a group	WordNet
4	An organization is both a social being and a group	CYC
5	A place is a physical object	Mikrokosmos, WordNet
6	A window is both an artifact and a place	Mikrokosmos

Table 3: Problematic subsumption relationships in some current ontologies.

7	A person is both a physical object and a living being	Pangloss
8	An animal is both a solid tangible thing and a perceptual agent	CYC
9	A car is both a solid tangible thing and a physical device	CYC
10	A communicative event is a physical, a mental, and a social event	Mikrokosmos

Note that a complete ontological analysis of these examples would go much beyond the scope of this paper, so we shall keep the discussion below rather informal. A far more in-depth example is available in [Guarino and Welty, 2000c].

Examples 1 and 2 clearly represent incompatible ontological commitments, unless we assume that amounts of matter and physical objects collapse into the same concept. As usual, the problem is that the underlying commitment has not been made explicit enough by the authors of these ontologies, and we only have the taxonomy to judge what the intended meaning of the terms “physical object” and “amount of matter” is. The analysis we have presented can help us to solve the puzzle, at least with respect to our own understanding of these terms. According to commonsense, amounts of matter can be assumed to carry an extensional IC (two amounts of matter are the same iff they have the same parts) and anti-unity (since every amount of matter is not an essential whole). Physical objects, on the other hand, allow for two possible options concerning their identity: in one account, they seem to have a non-extensional IC, since a physical object may keep its identity after replacing or removing some of its parts (e.g., a car with new tires); in a different account, they may have an extensional IC, if we assume that two physical objects are different if they don’t have the same parts. In any case, it seems natural to assume that physical objects *do* have unity, since we normally *count* them. In conclusion, the property *AMOUNT OF MATTER* can be labelled +I, ~U, while *PHYSICAL OBJECT* can be labelled +I, +U. This means that example 1 violates constraint (25), and example 2 violates constraint (24). Moreover, constraint (21) would be violated in both cases under the assumption of non-extensionality for physical objects. Our conclusion is that physical objects are *constituted* by amounts of matter, but they are *disjoint* from amounts of matter.

Examples 3 and 4 are similar. To analyze them, we have to decide what the IC for group [of people] are. If we admit that a group of people looses its identity when a member is replaced or deleted (as we believe is plausible), then the IC of *ORGANIZATION* becomes incompatible, since we clearly admit that organizations can change members. In both examples, therefore, constraint (21) is violated. Examining why this is so reveals our assumption that an organization is more than just a group of people: in fact, the same group of people could *constitute* different organizations.

Examples 5 and 6 include the notion of a place. We have (at least) two possible ontological choices regarding this notion: (i) we think of a place as a region of space (either absolute or relative space, the issue doesn’t matter here), adopting therefore an extensional IC (two regions of space are the same iff they have the same parts); (ii) we think of a place as a *feature* of something else (for instance, a hole in a wall). In the latter case, it seems plausible to give up the assumption of extensionality (if we think that the same hole can change its size), and introducing an assumption of unity.

Let us now consider example 5. If we take physical objects as non extensional (as above), then option (i) violates constraint (21). Option (ii) is consistent with (21), but this is a case where a further check of the *essential properties* of places and physical objects would be useful. For instance, if we take physical objects as being (essentially) *material* objects then we have an obvious inconsistence. So to account for example 5 we need to allow for *immaterial* physical objects. A further issue concerns however the ontological assumptions about *dependence*. If we take physical objects as ordinary objects, like a table or a glass, then we usually assume that they are (essentially) independent of everything else, i.e. they can actually exist even if nothing else does actually exist. Such assumption would be incompatible with option (ii), since a feature (like a hole) is an essentially dependent object, in the sense that it cannot exist unless something else (its host) exists. In conclusion, example 5 is consistent only if we consider places as features, and take physical objects in a very general sense, with no commitments regarding their materiality and their independence.

As for example 6, this is a classical case where multiple subsumption risks to be improperly used to account for lexical polysemy. If a window is assumed to be a *material* physical object (e.g, a suitably framed glass pane), then it cannot be a place at the same time. So there is a multiple *lexical* link that links *WINDOW* with its hypernyms, but this can't correspond to a subsumption link. The solution is to introduce two separate nodes, *WINDOW-ARTIFACT* and *WINDOW-PLACE*, which account for the two meanings of the word[1].

Examples 7 and 8 resemble each other, in the sense that *PERSON* behaves similarly to *ANIMAL*, *PHYSICAL OBJECT* to *SOLID TANGIBLE THING*, and *LIVING BEING* to *PERCEPTUAL AGENT.* The problem here, again, comes from incompatible ICs: if aa person is a physical object, there is no reason to claim she ceases to exist when she dies, since her body is still there... Indeed, life is considered to be an essential property for a person, while it seems to be an *anti-essential* property for a physical object, in the sense that any physical object (namely, a body) can possibly be a non-living object. If these assumptions are valid, then we must conclude that living beings are not physical objects, but they are rather *constituted* by physical objects.

Example 9 is similar to the previous two, with the difference that here the peculiar IC exhibited by physical devices is a *functional* one rather than a biological one.

Finally, example 10 involves events. The identity conditions for events may be complicated, but it seems plausible to assume as a necessary condition that if two events are the same then they must have the same participants. Now, the participants involved in physical, mental, and social events are different: we have a physical object, a perceptual agent, and a society, respectively. So the three events are different, although temporally co-located.

6 Conclusion

We have presented a compact formalization of some basic issues underlying the notion of identity and unity, by assembling, clarifying, and adapting philosophical insights forr

1.See {Nirenburg and Raskin 2001] for an objection to this argument, based on a rejection of the role of formal semantics for linguistically-motivated ontologies.

the purposes of practical knowledge engineering. We believe that the formal meta-properties we have introduced help to make explicit the ontological nature of the concepts used to structure a certain domain, and the constraints they impose on the subsumption relationship force the design of simpler, cleaner, and ultimately more reusable taxonomies.

Unlike previous efforts to clarify the nature of the subsumption relationship, our approach differs in that:

- It focuses on the nature of the properties involved in a single subsumption relationship, not on the semantics of the subsumption relation itself
- It focuses on the validation of single subsumption relationships based on the *intended meaning* of their arguments in terms of the meta-properties defined here, as opposed to focusing on structural similarities between property descriptions.
- It is founded on formal notions drawn from Ontology, and augmented with practical conceptual design experience, as opposed to being founded solely on the former or latter.

7 Acknowledgments

We are indebted to Bill Andersen, Stefano Borgo, Massimiliano Carrara, Pierdaniele Giaretta, Dario Maguolo, Claudio Masolo, Chris Partridge, Milena Stefanova, Mike Uschold, and Silvio Valentini for their useful comments on earlier versions of this paper. This work was supported in part by the Eureka Project (E! 2235) IKF, the Italian National Project TICCA (Tecnologie cognitive per l'interazione e la cooperazione con agenti artificiali), and a Clarke Fund research grant from Vassar College.

8 Bibliography

1. Akiba, K. 2000. Identity is simple. American Philosophical Quarterly, 37: 389-404.
2. Brachman, R. 1983. What IS-A Is and Isn't: An Analysis of Taxonomic Links in Semantic Networks. *IEEE Computer*, **16**(10): 30-36.
3. Casati, R. and Varzi, A. C., Holes and Other Superficialities, MIT Press/Bradford Books, Cambridge (MA) and London (UK), 1994.
4. Gangemi, A., Guarino, N., Masolo, C., and Oltramari, A. 2001. Understanding Top-Level Ontological Distinctions. In A. Gomez Perez, M. Gruninger, H. Stuckenschmidt, M. Uschold (eds.), *Proceedings of IJCAI2001 Workshop on Ontologies and Information Sharing*, Seattle, USA, 4-5 August 2001.
5. Goldstein, R. C. and Storey, V. C. 1999. Data abstractions: Why and how? *Data and Knowledge Engineering*, **29**: 293-311.
6. Guarino, N. 1998. Formal Ontology in Information Systems. In N. Guarino (ed.) *Formal Ontology in Information Systems. Proceedings of FOIS'98, Trento, Italy, 6-8 June 1998*. IOS Press, Amsterdam: 3-15. See amended version at www.ladseb.pd.cnr.it/infor/ontology/OntologyPapers.html.
7. Guarino, N., Carrara, M., and Giaretta, P. 1994. An Ontology of Meta-Level Categories. In D. J., E. Sandewall and P. Torasso (eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourth International Conference (KR94)*. Morgan Kaufmann, San Mateo, CA: 270-280.

8. Guarino, N. and Welty, C. 2000a. A Formal Ontology of Properties. In *Proceedings of 12th International Conference on Knowledge Engineering and Knowledge Management.* France, Springer Verlag.
9. Guarino, N. and Welty, C. 2000b. Identity, Unity, and Individuality: Towards a Formal Toolkit for Ontological Analysis. In *Proceedings of ECAI-2000: The European Conference on Artificial Intelligence*. Berlin, Germany, IOS Press.
10. Guarino, N. and Welty, C. 2000c. *Ontology driven conceptual analysis*. AAAI-00 Tutorial Notes. Available at http://www.cs.vassar.edu/faculty/welty/aaai-2000/.
11. Guarino, N. and Welty, C. 2001. Supporting Ontological Analysis of Taxonomic Relationships.To appear on *Data and Knowledge Engineering*.
12. Hirsch, E. 1982. *The Concept of Identity.* Oxford University Press, New York, Oxford.
13. Hughes, G. E. and Cresswell, M. J. 1996. *A New Introduction to Modal Logic.* Routledge, London.
14. Lewis, D. 1983. New Work for a Theory of Universals. *Australasian Journal of Philosophy*, **61**(4).
15. Lowe, E. J. 1989. *Kinds of Being. A Study of Individuation, Identity and the Logic of Sortal Terms.* Basil Blackwell, Oxford.
16. Lowe, E. J. 1998. *The possibility of metaphysics.* Clarendon Press, Oxford.
17. Nirenburg, S. and Raskin, V. 2001. Ontological Semantics, Formal Ontology, and Ambiguity. In B. Smith and C. Welty (eds.), *Proceedings of 2nd International Conference on Formal Ontology in Information Systems (FOIS2001).* ACM Press.
18. Noonan, H. (ed.) 1993. *Identity*. Dartmouth, Aldershot, USA.
19. Simons, P. 1987. *Parts: a Study in Ontology.* Clarendon Press, Oxford.
20. Storey, V. C. 1993. Understanding Semantic Relationships. *Very Large Databases Journal*, **2**: 455-488.
21. Strawson, P. F. 1959. *Individuals. An Essay in Descriptive Metaphysics.* Routledge, London and New York.