

The Role of Identity Conditions in Ontology Design¹

Nicola Guarino

National Research Council
LADSEB-CNR, Padova, Italy
guarino@ladseb.pd.cnr.it

<http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html>

From: AAAI Technical Report WS-99-13. Compilation copyright © 1999, AAAI (www.aaai.org). All rights reserved.

Abstract

Current ontologies' upper-level taxonomic structure is often quite complicated and hard to understand. In this paper I show how the theoretical tools of so-called Formal Ontology, and especially the theory of identity, can help to formulate cleaner, more general, more rigorous, and more understandable upper-level ontologies. I focus in particular on some examples of multiple generalization, proposing a way of simplifying the domain structure by splitting some concepts according to different identity conditions, or by excluding them because of their limited ontological role.

1. Introduction

Currently, a number of efforts in the area of knowledge and language engineering are aimed to the development of systems of basic semantic categories (often called "upper-level ontologies"), to be used for various applications such as natural-language processing, multilingual information retrieval, information integration, intelligent systems design. Examples of such systems are CYC's upper level, Penman's Upper Model (Bateman et al., 1990), recently evolved into the Pangloss ontology (Knight & Luk, 1994) and the Revised Upper Model (Bateman, 1995), the Mikrokosmos ontology (Mahesh, 1996), and WordNet's upper structure (Miller, 1995). Unfortunately, the upper-level taxonomic structure of these ontologies is often quite complicated and hard to understand.

This paper intends to show how the theoretical tools of so-called Formal Ontology (Guarino, 1995)—and especially the Theory of Identity—can help to formulate some ontological distinctions and design principles able to produce cleaner, more general, more rigorous and at the same time more understandable upper-level ontologies.

I will focus here on a single design practice, which is the main responsible of most semantic difficulties: IS-A overloading. Notice that I do not refer with this expression to the old debate about the semantics of IS-A (Brachman, 1983), since I assume a standard set-inclusion semantics for it. Rather, I will discuss the cases where this standard semantics turns out to be violated, if we carefully analyze the ontological nature of the arguments. I will consider a

number of "bad practice" cases, proposing a way to simplify the domain's hierarchical structure by splitting some concepts according to different identity conditions, or by excluding them from an explicit representation because of their limited taxonomic role.




Figure 1. An example of multiple inheritance in Pangloss.

2. The Main Problem: IS-A Overloading

All ontologies are centered on a taxonomy, based on a partial ordering relation named in various ways. Such taxonomy is the ontology's main *backbone*, which can be "fleshed" with the addition of attributes and other relations among nodes. As usual, we shall generically call IS-A such taxonomic relation (not to be confused with InstanceOf, which links a node to the class it belongs to and is not a partial order). IS-A's semantics will be the standard one: if P and Q are unary predicate symbols, P IS-A Q iff $I(P) \subseteq I(Q)$, where I is an interpretation function mapping unary predicates into subsets of the domain.

One of the problems with IS-A when considering linguistic ontologies like WordNet is that it often reflects a *lexical* relation between words, rather than an *ontological* relation between classes of domain entities. Although this fact is well known, the tendency to confuse the two aspects is quite high, especially when linguistic ontologies are used for non-linguistic applications. For instance, a common habit in linguistic ontologies is to rely on multiple inheritance for representing polysemy, as in the case of Figure 1. This results in an *overloading* of IS-A's role, which, as we shall see, may cause serious semantic problems.

To overcome these problems, the solution I propose is to pay more attention to the ontological implications of taxonomic choices, limiting IS-A links to connect nodes sharing a common *identity criterion*. In this way, the tax-

¹ A longer version of this paper appears in the Proceedings of IJCAI99 workshop on Ontologies and Problem Solving Methods

onomy reflects a basic ontological structure with a clear semantics, while the extra-information currently encoded by IS-A links can be represented by means of specialized links and attributes.

I will consider the following examples of IS-A links:

Reduction of sense

1. A physical object is an amount of matter (Pangloss)
2. An association is a group (WordNet)

Overgeneralization

3. An amount of matter is a physical object (WordNet)
4. A place is a physical object (Mikrokosmos, WordNet)

Confusion of senses

5. A window is both an artifact and a place (Mikrokosmos)

Clash of senses

6. A person is both a physical object and a living being (Pangloss)
7. An animal is both a solid tangible thing and a perceptual agent (CYC)
8. A car is both a solid tangible thing and a physical device (CYC)
9. An organization is both a social-being and a group (CYC)
10. A communicative event is a physical, a mental, and a social event (Mikrokosmos, Penman)

Suspect type-to-role link

11. A person is both a living thing and a causal agent (WordNet)
12. An apple is both fruit and food (WordNet, CYC)

These five classes of examples represent different kinds of IS-A overloading due to ontological misconceptions, which are discussed in turn.

Reduction of sense

In Examples 1-2, the ontological nature of the parent concept does not fully account for the nature of its child. A physical object is more than *just* an amount of matter, and an association is more than *just* a group of people: in fact, the same group of people could constitute different organizations. An analysis of the different identity conditions associated to these concepts leads to a different conceptualization, where an explicit *constitution* relation is taken into account.

Overgeneralization

This case is opposite to the previous one. Here is the parent concept's ontological nature that is weakened by a

child whose nature is very different. Consider for instance WordNet's gloss for *physical-object*:

object, physical object: a physical (tangible and visible) entity; "it was full of rackets, balls and other objects"

[Wordnet 1.6]

The sense of "physical object" implicit in these words seems to contradict the two examples reported, since: ii) "amount of matter" is an *uncountable* term, while the statement reported in the gloss refers to a *countable* notion of physical objects; i) a place is not tangible. We see therefore how these questionable IS-A links force a weaker interpretation of the subsuming concept.

Confusion of senses

In Example 5, different senses of a word are collapsed into a single concept, inheriting from two parents. In WordNet, two different "synsets" are associated to the term *window*: one is subsumed by *opening*, and the other by *panel*. Here such distinction is not made, and polysemy is handled by multiple inheritance. While this choice may appear economical at a first sight, it presents in this case evident ontological difficulties: is there an entity that is *both* a panel and an opening? A cleaner solution is to represent the two senses as disjoint concepts, expressing their intrinsic connection by means of extra relations (for instance, the *purpose* of a window-panel is to cover a window-opening).

Clash of senses

In these examples of multiple inheritance, the parent concepts have incompatible meanings. Consider Example 6: in fact, one may think that a person has both the properties of physical objects and living beings. However, as we shall see, *physical-object* and *living-being* appear at a closer inspection as disjoint concepts, since they have no common identity criteria. So the situation is the same as in Example 5, with the difference that the link between the two senses is stronger: here a living being *depends* on a physical object (namely its body) for its existence (Figure 2). It is also *constituted* by it, since the body and the living being have the same parts. In WordNet, these two senses for "person" are recognized, but no explicit relationships hold between them.

Examples 7-8 are very similar to 6, since in both cases perceptual agents and physical devices have no identity criteria in common with solid tangible things. Example 9 is even more evident, since the same group of people can form different organizations. In Example 10 (Figure 5), three different senses are collapsed together, as if a communicative event was just a physical event with some extra properties. Indeed, a communication event *involves* some mental event as well as some physical event, but the identity criteria of these entities are different one each other, so they belong to disjoint concepts.




Figure 2. The multiple links of Fig. 1 have been eliminated by introducing disjoint concepts for the different senses of person. Dashed arrows represent dependence links, accounting for the relationship between these senses.

Suspect type-to-role link.

This last class of examples is different from the previous ones. In this case the two links involved in the multiple inheritance have a different “strength”: a person is *necessarily* a living thing, while her/he *plays the role* of causal agent only when involved in certain events. Analogously for the apple, which is necessarily a fruit while it *can* be a food. person and apple are *types*, while causal-agent and food are *roles*. As we shall see, person and apple are *types*, while causal-agent and food are *roles*. On the basis of ontological analysis, I shall bring a reason why forbidding a role to subsume a type.

3. The Role of Identity Criteria

Let us focus first on the notion of *identity criterion*, which plays a fundamental role in our discussion. Briefly, we can say that an identity criterion (IC) for a property P is a binary relation I_P such that (Noonan, 1993)

$$Px \wedge Py \wedge I_P xy \rightarrow x=y$$

If, for a given property P, we are able to define such an I_P , then we say that P *carries* an IC for its instances.

We see therefore that an IC determines a *sufficient* condition for identity. In practice, ICs for classes corresponding to natural language words are often difficult or impossible to express. However, it is relatively easy (and quite useful) to identify some *necessary conditions* for identity, which are required to:

- *classify* an entity as an instance of a class C
- *individuate* an entity as a *countably distinct* instance of C
- *identify* two entities at a given time (*synchronic identity*)
- *re-identify* an instance of C across time (persistence, or *diachronic identity*)

These conditions may involve the nature of the parts of a certain entity (e.g., the chemical constitution of water), the topological relationships among parts (e.g., the self-connectedness of a stone), the relationships with regions of time/space (e.g. the occupancy of the same region of space), or the persistence of a certain global property across time (e.g., the permanence of a certain shape). It is important to take in mind that the decision of ascribing any of these conditions to a certain class is the result of our *conceptualization* of the world, i.e. of our *ontology*.

After these clarifications, let us see how ICs may affect a taxonomic organization. Consider to this purpose the example reported in (Lowe, 1989): should *ordered-set* be subsumed by *set*? Despite the apparent answer to this question is “yes”, after some thoughts we recognize it cannot be the case, because of cardinality reasons (multiple different ordered-sets may correspond to the same ordinary set). This is because the two concepts do not have an IC in common: having the same members is a sufficient condition for the identity of sets, but not for the identity of ordered sets². For analogous reasons, organization should not be subsumed by group, and person should not be subsumed by physical object (as a physical object, a body has persistence conditions different from a living being). Jonathan Lowe (1989) has discussed these problems at length, defending the following principle:

“No individual can instantiate both of two sorts if they have different criteria of identity associated with them”.

(Lowe, 1989)

I assume this as the basic principle to be adopted for well-founded ontologies.

4. A Minimal Ontology of Particulars

In order to address the problems discussed in the introduction, we have to clarify first our minimal ontological commitments. I assume a distinction between

- the entities of the world to be included in our domain of discourse, or *particulars*
- the properties and relations used to talk about such entities, or *universals*

I will quickly present below what I consider the “basic backbones” of these ontologies, with the only purpose of

² We may observe that, according to Kuratowski's definition, the ordered set $\langle a, b \rangle$ can be written as $\{\{a\}, \{a, b\}\}$. In this case, denying that an ordered set is a set would be difficult. Lowe however underlines that this “definition” establishes a useful mathematical correspondence between ordinary sets and ordered sets, but it can hardly be considered as an ontological definition, since it does not give an account of the notion of “member” of an ordered set. This example however does not influence the main argument.

making clear the methodological discussion which follows.

Particular	
Location	
Space	(a spatial region)
Time	(a temporal region)
Object	
Concrete object	
Continuant	(an apple)
Occurrent	(the fall of an apple)
Abstract object	(Pythagoras' theorem)

Figure 3. The basic backbone of the ontology of particulars.

The first distinction among particulars is between objects and locations. A *location* is either a region of (absolute) space or an interval of (absolute) time. An *object* is a particular which is not a location. Objects are *concrete* if they have a location in time and/or in space, *abstract* otherwise.

Within concrete objects, I assume here for granted the classical distinction between *continuants* and *occurrents*.

Continuants have a location in space, but this location can vary with time. They have spatial parts, but they do not have a temporal location, nor temporal parts. They always have other continuants as parts.

Ocurrents are "generated" by continuants, according to the ways they behave in time. In order for an occurrent to exist, a specific continuant must *take part* to it. If the continuant changes its identity, the occurrent also changes its identity, so that continuants are *rigidly dependent* on continuants. Examples of occurrents are ordinary events like the change of location of a body, but also the permanence of a body in a given location (a state occurrence). Occurrents always have other occurrents as parts. They have a unique temporal location, while their exact spatial location is often not clear.

Ontological Levels

Let us see now how, by applying Lowe's principle, we can introduce systematic distinctions among ontological categories that overcome the IS-A overloading problems mentioned earlier.

Take for instance a person, which can be conceptualized as an intentional agent, a living being or just a physical object (see Example 6). I argue that, since these concepts imply different ICs, they correspond to *disjoint* categories: three distinct individuals, instances of the three concepts above, share a common spatial location. Since the person depends on the underlying biological organism, as well as the latter depends on the underlying amount of matter, there is an intrinsic order within these categories, which belong therefore to different *ontological levels*.

A *dependence relation* links higher levels to lower levels: an animal depends on its body, which depends on body parts having a certain functionality, which depend on pieces of matter having specific properties, and so on.

Atomic	(a minimal grain of matter)
Static	(a configuration)
Mereological	(an amount of matter)
Physical	
Topological	(a piece of matter)
Morphological	(a cubic block)
Functional	(an artifact)
Biological	(a human body)
Intentional	(a person or a robot)
Social	(a company)

Figure 4. Ontological levels correspond to *disjoint sets* of particulars, according to the different ICs adopted to conceptualize them.

We have therefore a classification of ontological levels based on different *kinds* of IC (Figure 4), corresponding to different sets of *individuation* and *persistence* conditions. Each level corresponds to a class of ICs. These classes are assumed to describe *disjoint* sets of entities, in accordance with Lowe's principle. They are *orthogonal* to the continuants/occurrents distinction introduced above.

At the *atomic level*, we consider entities having minimal spatial or temporal dimensions, according to a certain granularity dependent on our conceptualization. We assume spatio-temporal continuity as a necessary condition for their identity.

At the *static level*, all the non-temporal properties of a particular contribute to its identity: if one of these changes, identity is lost. In this level only very peculiar objects are defined, namely *configurations* of atoms and *situations* (occurrences of configurations). The former are continuants, the latter are occurrents.

At the *mereological level*, the IC is *extensional*: two entities are the same if and only if they have the same parts (*mereological essentialism*). Regions of space, temporal intervals and amounts of matter belong to this level. Subsequent levels are characterized by an *intensional* criterion of identity, in the sense that mereological identity is neither sufficient nor necessary for identity.

The *physical level* corresponds to ICs bound to spatial configuration of matter (i.e., to topo-morphological properties). It can be split into two separate layers.

At the *topological layer*, the IC is bound to topological properties: for instance, topological self-connection can be considered as a necessary property to maintain identity: a *piece* of matter belongs to this layer, while a (possibly disconnected) *amount* of matter belongs to the mereological level. The two things are *distinct entities*, since a piece of matter can cease to exist (generating new pieces) while the same amount of matter is still there.

At the *morphological layer*, the IC is bound to morphological properties (or, in general, *gestaltic* properties related to spatial proximity), like spatial shapes or temporal patterns. A change of these properties can influence identity. A cube-shaped block is an example of an instance of this level: if its shape changes (above a certain limit) it is not *the same cube* any more, while still being the same piece of matter.

The levels above the physical level are related to ICs bound to the way objects interact with the external world. At the *functional level*, the IC is bound to functional and pragmatic properties: identity is destroyed when functionality is destroyed. At the *biological level*, the IC is bound to properties related to life: identity is destroyed when biological activity ceases. At the *intentional level*, the IC is bound to capability of intentional behavior: identity is destroyed when such capability ceases. At the *social level*, the IC is bound to social rules and conventions involving the interaction of intentional objects. Identity is destroyed when some of these rules change.

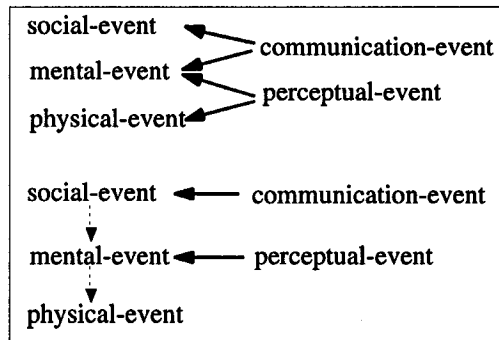


Figure 5. Above: communication and perceptual events in Mikrokosmos. Below: the simplification resulting from the assumption of ontological levels. Dashed arrows denote the dependence relation.

In conclusion, we can see how the adoption of Lowe's principle and the introduction of ontological levels solves the IS-A overloading problems discussed earlier (at least for examples 1-10). In all these cases, the proposed solution is the introduction of disjoint concepts (belonging to different ontological levels) accounting for the different senses involved (Figure 5). The costs of this choice are: i) a moderate proliferation (by a constant factor corresponding to the number of levels) of the number of entities in the domain; ii) the necessity to take into account different relations besides IS-A, such as *dependence*, spatio-temporal *colocalization* and *constitution* relations of various kinds. A formal characterization of these relations is outside the scope of this paper. See (Simons, 1987; Fine, 1995; Casati & Varzi, 1999) for a thorough technical discussion on these issues.

5. A Minimal Ontology of Universals

A minimal ontology of universals, based on a revision of (Guarino, Carrara & Giarretta, 1994), is reported in Figure 6. The first distinction is the usual one between *properties* and *relations*, according to the number of arguments. We only focus on *primitive* properties, which are not obtainable by Boolean combination of other properties.

Universal		
Property		
Type	(<i>person</i>)	(+I +R)
Category	(<i>location, object</i>)	(-I +R)
Role		(~R +D)
Material role	(<i>student</i>)	(+I)
Formal role	(<i>patient, part</i>)	(-I)
Attribution	(<i>red, decomposable</i>)	(-I -R -D)
Relation	(<i>part-of</i>)	

Figure 6. The basic *backbone* of the ontology of universals. I = identity, R = rigidity, D = dependence.

The purpose of studying the distinctions among properties is twofold. On one hand, we are interested in assessing their *organizational role* in a taxonomy, that is their practical relevance as *taxons*, i.e. nodes of a taxonomy; on the other hand, we want to study their attitude to generate clean and understandable hierarchies, with a minimum degree of "tangleness". With the help of formal ontology, we can characterize such distinctions on the basis of the following meta-properties:

1. *Identity* (+I). The property of carrying an IC.
2. *Rigidity* (+R). A property P is rigid if, for each x, if P(x) is true in one possible world, then it is also true in all possible worlds. *Person* and *location* are rigid, while *student* and *tall* are not.
3. *Anti-rigidity* (~R)³. A property P is anti-rigid if, for each x, P(x) is true in one possible world, and false in a different possible world. *Student* and *tall* are both nonrigid (~R) and anti-rigid (~R).
4. *Dependence* (+D). A property P is dependent if, necessarily, whenever P(x) holds, then Q(y) holds, with $x \neq y$ and $P \neq Q$ (see the *class dependence* mentioned before). *Father* is dependent, *person* is not.

A *type* is a property that is rigid and carries an IC. Types play the most important organizational role in a taxonomy. Assuming that each type has a distinct set of ICs, we have that, according to Lowe's principle, a taxonomy of types is always a *tree*. When a type specializes another type, it adds further ICs to those carried by the subsuming type. For instance, when the type *triangle* specializes *polygon*, it adds the ICs based on the equivalence of two sides and one angle (or two angles and one side) to those proper of polygons (same sides and same angles).

A *category* is a property that is also rigid but does not carry a specific IC. Since they cannot be subsumed by types (otherwise they would have an IC), categories only

³ See (Guarino, 1992; Guarino, Carrara & Giarretta, 1994) for a technical account of ontological rigidity and for a characterization of roles as non-rigid entities. The notion of anti-rigidity introduced here seems to better account for the ontological nature of roles, and explains the issue discussed in Figure 7.

appear in the uppermost levels of a taxonomy. Their role is to make clear the most general distinctions.




Figure 7. Types and roles in WordNet (roles marked with *). While it is OK for a type to subsume a role, the vice versa is forbidden according to the semantics we have given. Notice that, being types, *African* and *European* are disjoint.

Types and categories are both rigid, and can be either dependent or independent (*person* is independent, *event* is dependent). A *role* is a property that is *anti-rigid* and is always dependent. *Material* roles like *student* do have an IC, while *formal* roles like *part* do not. However, the IC of material roles is only indirect, since they do not introduce any specific IC, but rather they inherit it from a subsuming type. No explicit mutual disjointness assumption is made for roles, as they tend to generate tangled hierarchies. They have for this reason a limited organizational relevance. It seems therefore advisable to explicitly distinguish roles from types in order to easily isolate the main backbone of a taxonomy, and to perform inferences related to mutual disjointness (Figure 7). Notice that a role cannot subsume a type, since the former is anti-rigid and the latter is rigid.

Finally, an *attribution* is a property that is not rigid, is not dependent, and does not carry any IC. Attributions do not seem to play any useful organizational role in the ontology of particulars, as they may hold for disparate kinds of entity. Hence, they should not appear as taxons there, while the related information can be contained within the definitions of taxons whose instances exhibit such attribution. We see therefore how with this choice we solve the problems of confusion of organizational roles mentioned at the beginning of this paper.

Acknowledgements

I wish to thank Stefano Borgo, Massimiliano Carrara, Pierdaniele Giarretta, Claudio Masolo, Barry Smith, Achille Varzi.

References

- Bateman, J. A. 1995. On the Relationship Between Ontology Construction and Natural Language: a Socio-Semiotic View. *International Journal of Human-Computer Studies*, 43: 929-944.
- Bateman, J. A., Kasper, R. T., Moore, J. D., and Whitney, R. A. 1990. A General Organization of Knowledge for Natural Language Processing: the PENMAN upper model. USC/Information Sciences Institute, Marina del Rey, CA.
- Brachman, R. 1983. What IS-A Is and Isn't: An Analysis of Taxonomic Links in Semantic Networks. *IEEE Computer*, 16(10): 30-36.
- Casati, R. and Varzi, A. 1999. *Parts and Places. The Structure of Spatial Representation*. MIT Press, Cambridge, MA.
- Fine, K. 1995. Ontological Dependence. *Proceedings of the Aristotelian Society*, 95: 269-90.
- Guarino, N. 1992. Concepts, Attributes and Arbitrary Relations: Some Linguistic and Ontological Criteria for Structuring Knowledge Bases. *Data & Knowledge Engineering*, 8(2): 249-261.
- Guarino, N. 1995. Formal Ontology, Conceptual Analysis and Knowledge Representation. *International Journal of Human and Computer Studies*, 43(5/6): 625-640.
- Guarino, N., Carrara, M., and Giarretta, 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.
- Knight, K. and Luk, S. 1994. Building a Large Knowledge Base for Machine Translation. In *Proceedings of American Association of Artificial Intelligence Conference (AAAI-94)*. Seattle, WA: 773-778.
- Lowe, E. J. 1989. *Kinds of Being. A Study of Individuation, Identity and the Logic of Sortal Terms*. Basil Blackwell, Oxford.
- Mahesh, K. 1996. *Ontology Development for Machine Translation: Ideology and Methodology*. New Mexico State University, Computing Research Laboratory MCCS-96-292.
- Miller, G. A. 1995. WORDNET: A Lexical Database for English. *Communications of ACM*, 2(11): 39-41.
- Noonan, H. 1993. Introduction. In H. Noonan (ed.) *Identity*. Dartmouth, Aldershot: xi-xx.
- Simons, P. 1987. *Parts: a Study in Ontology*. Clarendon Press, Oxford.