Avoiding IS-A Overloading: The Role of Identity Conditions in Ontology Design

Nicola Guarino
National Research Council
LADSEB-CNR, Padova, Italy
guarino@ladseb.pd.cnr.it
http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html

extended abstract
(a longer version appears in the Proceedings of IJCAI99 workshop on Ontologies and Problem Solving Methods)

Abstract

The upper-level taxonomic structure of current ontologies is often quite complicated and hard to understand. In this paper I analyze the ontological nature of upper-level categories, considering the necessity of splitting them according to more subtle distinctions or the opportunity of excluding them because of their limited organizational 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 the CYC upper level, the Penman Upper Model [Bateman *et al.* 1990], recently evolved into the Pangloss ontology [Knight and Luk 1994] and the Revised Upper Model [Bateman 1995], the Mikrokosmos ontology [Mahesh 1996], and the Wordnet 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] 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 does not reflect 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 more subtle ontological distinctions, 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 x and y are unary predicate symbols, x IS-A y iff $I(x) \subseteq I(y)$, 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 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 Fig. 1. This results in an *overloading* of IS-A' 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 similar *identity conditions*. In this way, the taxonomy 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.

Let us consider the following examples of IS-A links:

Confusion of senses

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

Reduction of sense

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

Collapse of senses

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

Overgeneralization

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

The four classes of examples above represent different kinds of ontological misconceptions, which are discussed in turn.

Confusion of senses

In Example 1 above, different senses of a word are collapsed into a single class, inheriting from multiple superclasses. In WordNet, two different “synsets” are associated to the term `window`: one is subsumed by `opening`, while the other by `panel`. Here such distinction is not made, and the two senses collapse into one. While the choice of using multiple inheritance to handle polysemy may appear economical at a first sight, it presents in this case evident ontological inconveniences: 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 the intrinsic dependence by means of an extra relation. For instance, an opening needs to exist in order for a window to function properly.

[1] It is interesting to compare this with WordNet, where an amount of matter is a physical object (example 10).

Reduction of sense

In Examples 2-3, the IS-A link points to an *aspect* of the meaning of a given concept that does not fully account for its identity: 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. As we shall see, an analysis of the different identity conditions associated to these entities leads to a different conceptualization, where an explicit *constitution* relation is taken into account.

Collapse of senses

In these examples of multiple inheritance, one of the subsuming concepts has a very basic meaning, which is incompatible with those assigned to other subsumers: in this way, the subsumee's sense collapses to a very basic one. Consider Example 4: 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, due to their different identity conditions. So the situation is the same as in the window example, with the difference that the relation linking together the two senses is now a stronger one: not only a living being *depends* on a physical object (namely its body), but it is also *constituted* by it. Indeed, in WordNet two different senses for the word “person” (i.e., “person” as a body and “person” as a living being) are postulated, but no explicit relationship holds between them (Figure 2).

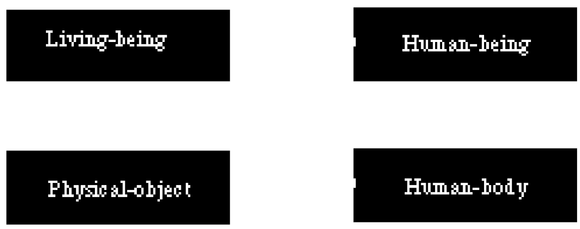

Figure 2. The multiple links appearing in Fig. 1 have been eliminated by introducing disjoint concepts for the different senses of `person`. Dashed arrows represent dependence links, accounting for the relationship between these senses.

Examples 5-6 are very similar to 4, since in both cases "being a perceptual agent" and "being a physical device" impose stringent identity conditions on solid tangible things. Example 7 is even more evident, since the same group of people can constitute different organizations. In

Example 8 (Figure 3), 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 these entities are different one each other, and belong to disjoint classes.
In all these cases, the proposed solution is to introduce separate concepts accounting for the different senses involved. Relationships between these concepts can be represented by means of extra relations, which in many cases turn to be *dependence* and *colocalization* relations of various kinds. For instance, a living being depends on the existence of a material body; a communicative event depends for its existence on some mental events and some physical events. In all these cases the depending entity is spatio-temporally *colocalized* with the entity it depends on.

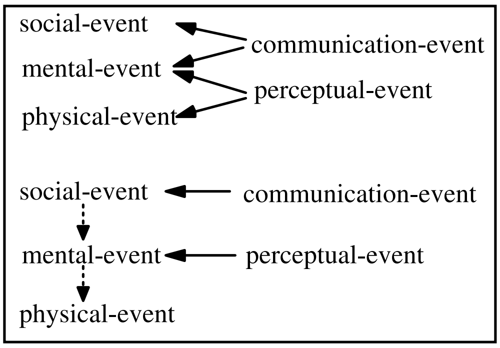

Figure 3. Above: communication and perceptual events in Mikrokosmos. Below: the simplification resulting from the assumption of ontological levels. Dashed arrows denote the dependence relation.

Overgeneralization
This case is opposite to the previous one. Here is the subsumer's ontological nature that is weakened by the existence of an IS-A link coming from a concept that is very different in nature. 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: i) a place is not tangible; ii) "amount of matter" is an *uncountable* term, while the statement reported in the gloss refers to a *countable* notion of physical objects. We can see therefore how these questionable IS-A links force a weaker interpretation of the subsuming concept.

The Role of Identity Criteria

Let us focus now 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 difficult or impossible to express. Rather, it is relatively easy (and quite useful) to identify some *necessary conditions* for identity, which allow one 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 space as a criterion of synchronic identity), 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 have different ICs. 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). These problems have been discussed at length by Jonathan Lowe [1989], who formulated 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.

A Minimal Ontology of Particulars

In order to address the problems discussed in the introduction, it is useful to clarify my minimal commitments concerning the upper-level structure. First of all, 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 talk therefore of two separate ontologies, each one with its "basic backbone". For reasons of space, I will focus here on the basic ontology of particulars, presenting it the only purpose of making clear the discussion on the design principles, which is the goal of this paper.

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 4. 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.

Occurrents 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 considered 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 5. Ontological levels correspond to *disjoint sets* of particulars, according to the different ICs adopted to conceptualize them. Examples are reported in parenthesis.

We have therefore a classification of ontological levels based on different *kinds* of IC (Figure 5), 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.

In conclusion, we can see how the introduction of ontological levels leads to simpler and (I believe) more understandable taxonomies, as shown by Fig. 4. 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 co-localization and dependence.

Some Basic Design Principles

In conclusion, let us summarize the ontology design principles emerging from this discussion, which can solve the *IS-A overloading* problems we have mentioned in precedence.

1. *Be clear about the domain*. Any formal theory is a theory about a domain. Such a domain must be clarified in advance[2].

2. *Take identity seriously*. We have seen how the notion of identity criterion (and especially Lowe's principle) can play a crucial role in clarifying ontological distinctions.

Acknowledgements

I am much indebted to many people. I wish to thank in particular Stefano Borgo, Massimiliano Carrara, Pierdaniele Giaretta, 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.

Guarino, N. 1995. Formal Ontology, Conceptual Analysis and Knowledge Representation. *International Journal of Human and Computer Studies*, **43**(5/6): 625-640.

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.

[2] A further design principle based on the distinction between different predicate kinds on the basis of their identity criteria is discussed in the extended version of this paper.

Miller, G. A. 1995. WORDNET: A Lexical Database for English. *Communications of ACM*, **2**(11): 39-41. htpp://clarity.princeton.edu/pub/wordnet.

Noonan, H. 1993. Introduction. In H. Noonan (ed.) *Idenity*. Dartmouth, Aldershot: xi-xx.