# On the use of WordNet for semantic interoperability: towards "cognitive computing"

Csaba Veres

Dept. of Computer and Information Science (IDI) Norwegian University of Science and Technology (NTNU) N-7491 Trondheim-NTNU. Csaba.Veres@idi.ntnu.no

**Abstract.** WordNet is an electronic lexical database structured around psychological and linguistic principles. As such, it should play a part in any effort to integrate cognitive factors into knowledge based systems. Yet, some of its basic assumptions have been attacked, and the suggestion made that a major restructuring would make it more cognitively transparent. We investigate these allegations from a psycholinguistic perspective and conclude that WordNet is in fact rigorous in terms of the cognitive principles it embodies. What is lacking is a methodology for translating the explicit and implicit knowledge in WordNet into a usable, formal ontologies. We show some ways in which WordNet should be extended to facilitate this process. We agree that WordNet is not in itself ready for use as a formal ontology, but we argue that it is an invaluable tool for describing the conceptualized structure of our world, and should be used as a fundamental resource.

## 1 Introduction

In this paper we investigate some important issues concerning the use of WordNet in the construction of ontologies. Along the way we tackle questions regarding the way in which cognitive science can inform the design of computer systems. It is clear that there is a stated commitment to cognitive factors in a number of research areas. For instance DARPA's project on Cognitive Information Processing Technology [1], headed by Ronald Brachman, aims at completely novel, cognitively inspired approaches to computation. In the interoperability arena Oltramari et. al.(2002), and Gangemi. et. al. (2003a) (to be referred to together as O&G henceforth, since both make the same essential argument) have also argued that their principles have a strong "cognitive bias".

WordNet is a lexical database constructed on lexicographic and psycholinguistic principles, under active development for the past 20 years at the Cognitive Science Laboratory at Princeton University. It has been widely used in computational linguistics partly because of its vast coverage, containing 138,838 English words (Miller, (2002)), compared with 35,958 in the Longman Dictionary of Contemporary English, and 43,943 in Roget's Thesaurus (Stevenson, (2003)). In addition it contains hundreds of thousands of links that represent psycholinguistically important relationships between words (more precisely; synsets, which

[1] http://www.eps.gov/spg/ODA/DARPA/CMO/BAA02-21/SynopsisP.html

represent specific meanings of a word), which are in turn grouped according to syntactic category to reflect the different sorts of relationships observed in the different categories.

On the face of it, WordNet would appear to be an invaluable resource for documenting the important terms in a domain and the relationships between them, acting as a sort of "ontological dictionary". Yet, O&G's claim that "WordNet is used more and more today as an ontology" notwithstanding, the penetration of WordNet into the ontology community has been relatively shallow. Much of the published work tries to use it as a means for automated disambiguation of ontology terms (e.g. Missikoff, et. al., (2002)), or supplements it in some non trivial way for very specific purposes (e.g. Gangemi et. al., (2003b)). But very few attempt to use it as an ontology. There are many good reasons for this, not the least of which is that it is not formalized in terms of a readily usable ontology.

But a more serious set of criticisms by O&G challenge fundamental aspects of WordNet, suggesting major revisions to its ontology. They suggest that their 'cognitively biased' approach can be used to "clean up" WordNet. But the approach actually has its foundations in philosophy and metaphysics. On the face of it, there seems to be some oddness about a metaphysical approach claiming a more intuitive 'cognitive bias' than one that is built around psychological and linguistic principles! It therefore becomes doubly interesting to investigate these critiques to see if indeed WordNet has problems, and as an opportunity to discuss issues concerning the relationship between cognitive science, philosophy, and the information sciences.

## 2 Challenges for WordNet

DOLCE (Descriptive Ontology for Linguistic and Cognitive Engineering) is an attempt to define an upper level ontology, one module in the Foundational Ontology Library being developed in the WonderWeb project (O&G). In turn, the Library of Foundational Ontologies is a vision for assembling a collection of highly articulated formal ontologies with clearly defined commitments, to allow Semantic Web applications to chose common frameworks for interoperability. While DOLCE itself is not meant as a "universal standard ontology", it is intended as a "starting point for comparing and elucidating the relationships with other future modules of the library, and also for clarifying the hidden assumptions underlying existing ontologies or linguistic resources such as WordNet". Clearly, the authors think the principles contained in DOLCE have a certain primacy in questions of ontological "correctness". Ironically as already noted, in spite of the strong cognitive commitment claimed for the approach, the majority of their principles derive from philosophy and metaphysics. In order to see if these do in fact challenge the assumptions in WordNet, we take two poignant criticisms, and investigate their implications. To succeed in this task, we need some criterion by which we can judge if indeed O&G's criticisms are successfully refuted. We will adopt a two level criterion. First, O&G argue that the problems in WordNet prevent it from forming a useful foundation for the construction of a precise ontology. If we can show this to be false, we will have succeeded at level one. But if we can also show that WordNet presents new and constructive approaches to further develop the technology of ontologies, then we will have succeeded at the second level also.

## 3 Brief summary of the argument against O&G

In the remainder of this paper we show that, at least some of the major critiques posed by O&G are not correct because the information that is claimed to be absent or obscured is in fact available through the cognitive principles embodied in WordNet. We consider these psycholinguistic facts underlying the construction of WordNet, and show how this analysis helps us uncover information that is needed for constructing an ontology. Then we argue that the psycholinguistic facts suggest a number of extensions to WordNet as well as current web technologies, to significantly increase the power of ontologies. What is needed is a formal clarification of the facts behind WordNet, not its re structuring.

## 4 Two "critical problems" investigated

### 4.1 Confusion between concepts and individuals.

The essence of the first and "critical problem" (sic.) is that WordNet confuses concepts and individuals, freely mixing them in the taxonomy of nouns.[2] For example under the concept `composer` we find `Bach` and `Beethoven` together with `songwriter` and `contrapuntist`, where the first two are, putatively, instances of `composer` whereas the latter two are subclasses (i.e. concepts) denoting particular kinds of `composer`. This problem is supposedly due to an "expressivity lack" since the problem could be overcome with the inclusion of an INSTANCE-OF relation. So why is this relation missing in WordNet? We will argue that the lack of differentiation in WordNet is a direct consequence of the way human cognition treats classes and instances.

The taxonomic link for nouns in WordNet is described by the two terms *hyponymy* (subordinate) and *hypernym* (superordinate), with the intended interpretation as a "transitive, asymmetric relation that can be read 'IS-A' or 'IS-A-KIND-OF'" (Miller, (1998), p.25). Various authors have struggled with the formal interpretation of this relationship with many opting to simply equate it to subsumption as found in description logics (e.g. Alvarez, (2000)). This is the interpretation that leads to the problem currently being discussed, since subsumption holds only between concepts. In fact, O&G conclude that the assumption is a mistake, but their solution is to split the hypernyms into ones that are equivalent to subsumption, and those which are not. We will argue that indeed the 'IS-A'

[2] Only nouns and verbs are arranged in a taxonomy, since the other syntactic categories have different organizing principles.

relation is not the cognitive/linguistic counterpart of subsumption, but at the same time cannot be split into a number of other relations like 'IS-A-KIND-OF'.

Subsumption itself can be defined as follows

> A concept $C$ is subsumed by a concept $D$ if in every model of $\mathcal{T}$ the set denoted by $C$ is a subset of the set denoted by $D$. (Where $\mathcal{T}$ is the terminology constructed by a knowledge engineer). (Baader and Nutt, (2003))

As an initial demonstration that 'IS-A' cannot be equated with subsumption, Wierzbicka (1984) points out in what she calls the "fallacy of set inclusion", "every policeman is somebody's son, and not vice versa, but this does not mean that a policeman is conceptualized in English as a kind of son" (p.314). All conceivable instances of *policeman* must also be instances of *son*. But neither human intuition, nor WordNet would presume that *a policeman is a (kind of) son (who works for the police)*. This example shows that not every possible subsumption relation should appear as hypernymy. But is it also true that existing hypernym relations in WordNet are not in fact subsumptions?

Miller (1998) notes that the nature of the IS-A relation in language is semantically quite complex, possibly involving five distinct kinds of relation as identified in Wierzbicka (1984). In particular, the distinction between "taxonomic" (IS-A-KIND-OF) and "functional" (IS-USED-AS-A-KIND-OF) hyponimies is often confounded. For instance `written_agreement` has two hypernyms in WordNet: `legal_document` and `agreement`. The first is a taxonomic role that can be paraphrased as 'a written agreement is a (kind of) legal document summarizing agreement between parties, and as such, can be used in a court of law', while the second represents a functional role as in 'a written agreement is used as an agreement to make sure everyone does what they are supposed to'. One way to deal with this problem in some cases is to claim that the words themselves are polysemous, and treat each individual meaning separately in the lexical network. For example, consider the word *chicken* as in: *chicken is a bird that is used for food*. WordNet contains two separate entries, one with hypernym `bird`, the other with hypernym `food`, in which case the sense of both relationships can be captured without need for additional types of links. However, this solution presents two intimately related problems. First, it still lacks a proper IS-USED-AS-KIND-OF relation. Is chicken REALLY a kind of food, or is it a kind of bird that is used for food, after all? The second and related problem is whether or not the word *chicken* can be thought of as having two separate meanings at all? Maybe there is only one kind of chicken, the one that is a bird, and it just happens to be used for food.

But now we should be puzzled. Why is there so much confusion over the IS-A relation? If IS-A is 'really' a confound of so many distinct and important relations, why are they not sufficiently psychologically salient to have been included in WordNet? In similar vein, why were similar distinctions confounded in the computer science literature for years, where the semantics of IS-A has been highly fragmented (Brachman, 1983)? An answer may be found in the psychology behind the lexical item. Consider the following linguistic intuitions that can be employed for checking hypernymy (see Miller, (2002) for a discussion and further references). First, it is sensible to compare two nouns only if they are not in a hypernymy relation: you might in fact *prefer apples to oranges*, but it makes little sense to *prefer apples to fruit*! Similarly, it seems wrong to *prefer fruit to apples*. Another intuition involves anaphoric coreference. *I gave him a good novel, but the book bored him* seems like a perfectly natural sentence, but *I gave him a good novel, but the catsup bored him* is distinctly odd. The problem is that anaphoric nouns appear to be acceptable only if they are hypernyms of the antecedent. Finally, there appear to be selectional constraints for some verbs that involve hypernymic relations. For example one can *drink chamomile tea, tea, cafe royale, a cappuccino, coffee, a mixed drink, a Long Island Iced Tea*, or *a nice cold beverage*. In fact, any hyponym of beverage will do. But one cannot *drink a chair*, or *the square root of negative two*.

Such intuitions make it clear that language speakers have access to information about some sort of hypernymy relations, since these form the basis for general rules of sentence formation. But do these intuitions distinguish between the troublesome relationships we have been discussing? Consider the following examples (where the asterisk indicates ill formed expressions and the question mark, where certain uncertainties exist):

1. I prefer chicken to beef.
2. *I prefer chicken to food.
3. *?I prefer chicken to living things.
4. I prefer veal to living things.
5. I thought the chicken was very tasty, but she didn't like the food at all.
6. *I thought the chicken was very tasty, but she didn't like the telephone at all.
7. *?I thought the chicken was very tasty, but she didn't like the animal at all.
8. I thought the chicken was very pretty, but she didn't like the animal at all.
9. I love to eat chicken/poultry/fowl/meat/food/*chairs.
10. I love to eat chicken/?animal/??living thing.
11. I love to feed chickens/animals/living things.

The felicitous use of the comparative in 1. shows that *chicken* and *beef* are not in a hypernymy relation, but 2. and 3. show that *chicken* is in such a relation with both *food*, and *living thing*. Example 3. is also interesting because the apparent uncertainty about its status supports the claim that *chicken* is in fact polysemous as previously suggested (it only sounds bad under one interpretation). Thus, if we take the "food" interpretation of *chicken* in 3., it sounds fine, as further illustrated in example 4. where *veal* is a word used only for the food interpretation. Examples 5. - 8. make a similar point. The coreference in 5. is established between *chicken* and its hypernym, but this is not possible in 6. Intuitions on 7. might vary depending on its interpretation. But the "food" sense shows that *animal* is equally inappropriate as a hypernym for *chicken* as *telephone* is. The dominant "food" interpretation in this example is determined by the immediate context as is illustrated in 8. which sounds perfectly well formed because the context biases the reader towards the "animal" reading. Finally, 9.

shows the selectional constraints for *chicken* as a member of the *food* hierarchy, and 10. shows the problems with *chicken* if construed as a member of the *living_thing* hierarchy (once again determined by local context, as shown in 11).

Based on these observations one could argue that the two relations do not necessarily need to be distinguished. But this is not to say that there is no validity to Wierzbicka's (1984) observations. Instead, one suggestion is to keep the undifferentiated IS-A link and capture the different intuitions about the roles as properties of the lexical items in the relation. An interesting research question is to determine the generality of this solution with respect to the set of distinctions identified by Wierzbicka (1984). But of more immediate concern is to see the significance of these tests to the examples that began this section, to do with instantiation. Is there really an "expressivity lack" in WordNet, with respect to the "missing" INSTANCE-OF relation, or could we incorporate that into IS-A as well? Consider the following sentences, using the O&G examples

12. Beethoven was much more talented than Bach.
13. *Beethoven was much more talented than a/the composer. (*Composers were much more talented than Beethoven).
14. Beethoven was much more talented than a/the songwriter. (Songwriters were much more talented than Beethoven).
15. I thought Beethoven was a genius, but she thought the composer was arrogant.
16. *I thought Beethoven was a genius, but she thought Bach was arrogant.
17. *I thought Beethoven was a genius, but she thought the president was arrogant.
18. *I thought Beethoven was a genius, but she thought George Bush was arrogant.
19. *I thought Beethoven was a genius, but she thought the songwriter was arrogant.

It seems that individuals exhibit similar patterns of behavior as concepts do. It is possible to use the comparative on the instances in 12., which contains two individuals that are not in a hypernymy relation. But the comparative is infelicitous in 13. because *composer* is a hypernym of *Beethoven*. On the other hand 14. is a fine, as *songwriter* and *Beethoven* are both hyponyms of *composer*, and therefore one is not a hypernym of the other. The coreference in 15. - 19. tells a similar story: coreference obeys the hypernymy constraint with individuals and classes equally.

Clearly this pattern suggests that the concept/individual distinction is not salient at some level of cognitive structure since they are freely mixed in the example sentences, and appear to participate in the same taxonomic hierarchies. This somewhat surprising conclusion is independently supported by the lexicographic roots of the IS-A relation in WordNet. This relation between nouns reflects a common 'definitional formula' for the meaning of nouns in dictionaries, where a hypernym is combined with various modifiers to define a more specific case of an already known word. For instance the word *robin* might be defined as 'a migratory bird that has a clear and melodious song and a reddish breast with gray or black upper plumage'. Thus, *robin* IS-A *bird* that is specialized in various ways. In the current example WordNet defines *Beethoven* as a 'German composer of instrumental music (especially symphonic and chamber music)', and a *songwriter* as 'a composer of words or music for popular songs'. The same definitional pattern applies to both, demonstrating once again the ubiquity of the IS-A relation.

At this point of the discussion it becomes important to emphasize that we are not denying the sorts of formal and intuitive distinctions that have been identified by Wierzbicka (1984), Brachman (1983), or O&G. What we are trying to do is question the way in which these distinctions ought to be captured in a formal system, especially one that is cognitively inspired. O&G clearly wish to separate IS-A and INSTANCE-OF. This of course would appear to simplify the implementation of the ontology in terms of currently available technologies such as OWL and description logics, and agrees with conventional thinking in the clear separation of "classes" and "instances".

For example in description logics, the formal foundation of OWL, there is a fundamental distinction between concepts and named individuals. Concepts define the terminology of the domain whereas named individuals define an actual state of affairs in a given world. Individuals are, essentially, the data that is stored in a given application. A modeler has to make choices about the nature of the represented entities, based on the needs of the application data. It is often a matter of choice that a given entity becomes an individual rather than a concept in a given ontology. For example, Borgida and Brachman (2003) (p. 353) ask us to consider a modeling task involving books. If the domain concerns literature courses, a named individual might be something like "Dicken's HARD-TIMES". An Internet bookstore, on the other hand, would require representation of more concrete levels - BOOK-EDITIONS for instance, since different editions might have different prices. Finally, a library application would be more concrete still, keeping track of each individual BOOK-COPY. The decision is crucial since it determines if an entity should be represented as a concept or individual in OWL, or an atomic concept or a concept assertion in a description logic.

The linguist Ray Jackendoff argues strongly against the logic based distinction of classes and individuals, and proposes a different, psychological distinction between TYPES and TOKENS which then turns the IS-A judgment into an act of categorization (Jackendoff, (1983)). A TOKEN, then, is anything (concrete or abstract) that requires categorization and a TYPE is a conceptual category. Importantly, TYPES and TOKENS are claimed to have essentially identical internal conceptual structures, which is reflected in natural language by the fact that both are represented by the same syntactic category. Thus in the sentence *Clark Kent is a reporter*, the TOKEN individual *Clark Kent* as well as the TYPE *reporter* are Noun Phrases that are connected by the verb *be*. This contrasts with a description logic representation, for instance, where *reporter* would be an atomic concept (REPORTER) and *Clark Kent* a name in a concept assertion (REPORTER(CLARK_KENT)). On the other hand the cognitively important difference between the two elements is that TOKENS have projections onto real world entities while TYPES do not. To see why the cognitive interpretation is important for the current issue, and why it differs from the classical logic based one, consider the following sentences:

20. Clark Kent is a reporter. (TOKEN - TYPE: ordinary categorization)
21. Clark Kent is Superman. (TOKEN - TOKEN: token identity)
22. A reporter is a person. (TYPE - TYPE: generic categorization)
23. Clark Kent looks like a reporter. (TOKEN - TYPE)
24. Clark Kent looks like Superman. (TOKEN -TOKEN)
25. Reporters look like frogs. (TYPE - TYPE)

Examples 20., 21., and 22. show an BE relation between a TOKEN and a TYPE, two TOKENS, and two TYPES, respectively. Once again note that classes and individuals are freely mixed in this kind of expression. But notice that now there are conceptual differences depending on the nature of the NPs: while 20. reflects an ordinary categorization (instantiation), 21. shows an identity relation, and 22. is a categorization between two concepts, which corresponds to the hypernymy relation in WordNet. We have previously suggested that the IS-A relation ought to be undifferentiated, and Jackendoff (1983) argues toward a similar conclusion from the current observations (among others). He argues that 20. - 22. are all examples of the very same process of categorization, and the differences in the kind of categorization performed is due to the nature of the elements being compared, not the verb that connects them. Two TYPES yield a generic categorization with the expected properties like irreflexivity, and so on. Comparing a TOKEN with a TYPE yields an ordinary categorization. The differences in the nature of the categorization are not given by differences in the relations, but by differences in the arguments of a single relation. In other words the fact that a single verb *be* expresses all categorization judgments is not accidental but reflects a deep unity of process: it is not the case that *be* is polysemous, and it just happens that several relations are expressed by the same verb. Examples 23. - 25. support this argument by showing that *be* is not the only verb that behaves according to this pattern of arguments. The comparisons involved in evaluating the *looks like* sentences parallel those needed for the *be* sentences. Thus if one wants to argue that different relationships are involved in *be* sentences with various TYPE and TOKEN arguments then one would also have to argue for a similar distinction in *looks like*.

Interestingly, this psycholinguistically motivated hypothesis is surprisingly like the conclusions of Brachman (1983) who also notes that the precise interpretation of the IS-A link found in various implemented semantic networks depends largely on the nodes they connect. He concludes that

> "... things might be a lot clearer if IS-A were broken down into its semantic subcomponents and those subcomponents then used as the primitives of a representation system." (p. 36)

The most central of these subcomponents are the interpretation of the nodes in terms of whether they are supposed to be GENERIC/INDIVIDUAL, as in

Jackendoff's analysis. But the force of Jackendoff's thesis is that we need not pre classify entities as concepts or individuals. These can be dynamically assigned according to the particular situation, and the possible inferences will then be determined not by the nature of the relation but the nature of the connected concepts. This is quite a novel way of looking at the class/instance distinction, and we are currently considering ways in which such ideas may be incorporated in knowledge based systems.

Let us summarize the argument so far. We saw various sorts of possible problems with WordNet, but then argued that psycholinguistic evidence suggests that the problems may not be so severe, especially in light of Brachman's independently motivated conclusions that the formal properties of the roles connecting two concepts are predictable from the properties of the concepts. Thus as long as the relevant information is available in WordNet, the role properties should be easily recovered. But is such information available in WordNet?

In the majority of O&G's examples, the problematic "individuals" appear as proper names (with an initial capital) in WordNet, that are interspersed among the "concepts". But in this case a putative distinction between TYPES and TOKENS, or classes and instances is in fact available since proper names denote particular individuals (Bloom (2002)). This linguistic distinction can be easily extracted in a specific implementation using a particular knowledge representation tool such as OWL, where proper names can become individuals and common nouns can be classes. But unfortunately as Miller (1998) points out, this heuristic is not generally applicable. The coding of nouns along the proper/common dimension is far too intricate and context dependent to be formally included in WordNet. Thus we have inconsistent examples like `martial art` which has hyponyms like `karate`, `aikido`, `judo`, and `Kung Fu`. Here, all entries except `Kung Fu` are common nouns, yet seem to express the same sorts of concept, whatever that may be. Fortunately, syntax helps decide what sorts of concepts these might be. In English, common nouns can appear with a determiner or demonstrative pronouns, but proper names can not. The fact that we say *I study judo* instead of *I study a judo* shows that `judo` is a proper name.

There is, then, some information already available for coding TOKENS in WordNet. But a great deal remains to be done to supplement WordNet in this regard. We need to identify linguistic factors that can be found in written texts, that can be used to determine if a word is used as a TYPE or TOKEN, and to incorporate their use in tools for automatically classifying relevant entities. We have already seen a simple syntactic cue that is useful in English, but unfortunately has limitations for other languages Bloom (2002). A full analysis would yield ways to distinguish reference to types and tokens, which could then be used to establish and populate an ontology in a particular application.

The important result for the moment is that the single IS-A relation in WordNet is not necessarily limiting, if we can clearly determine the nature of the entities participating in the relation. This might be restricted simply to the TYPE/TOKEN variety, which can be used to deal with the class/instance problem. A more interesting possibility is that by linking WordNet synsets to highly articulated descriptions of semantic structure one could derive intricate patterns of inference through the simple IS-A link. For instance, Pustejovsky (1991) defines a TELIC component for lexical semantics which specifies the intended use, or purpose of an entity. In the *chicken/food* example above, *food* could have a TELIC role specified in terms of its use for nutrition, which would then be inherited by *chicken*, which in turn would then acquire the IS-USED-AS-KIND-OF interpretation.

### 4.2 Confusion between object-level and meta-level

O&G feel that in some cases WordNet hierarchies include both object-level and meta-level concepts. Examples of the former are `set`, `time`, and `space`, and of the latter `attribute` and `relation,` all of which are hyponyms of `abstraction` (defined as "a general concept formed by extracting common features from specific examples") in WordNet.

This criticism once again stems from a specific logical/formal view of knowledge, which admits several distinct levels of description. This is a very common assumption, even with semi formal languages like UML (the Unified Modeling Language) which has several meta-levels, with each meta-level providing an abstract description of the kinds of concepts used at the lower level. But it is not clear that meta-levels of this sort have any basis in cognition at all, and any approach that purports to display a "cognitive bias" should not be confused over this point. Once again this is not to say that we should not have meta-levels in a specific ontology application, only that the resource on which the ontology is based should not necessarily define all the meta level concepts for us; it need only to identify concepts that might be useful for a meta level. This is precisely the function that `abstraction` in WordNet can play.

In considering this example it is important to remember that WordNet attempts to define all meanings of words, otherwise confusion can arise if we assume the wrong interpretation of a word in an example. O&G consider the sense of *set* that appears under the root node `abstraction`. Actually there are two senses that O&G do not differentiate:

26. set – ((mathematics) an abstract collection of numbers or symbols; "the set of prime numbers is infinite")
27. set – (a unit of play in tennis or squash; "they played two sets of tennis after dinner")

A third sense, not considered by O&G in their example at all, appears under the root node `group, grouping`:

28. set – (a group of things of the same kind that belong together and are so used; "a set of books"; "a set of golf clubs"; "a set of teeth")

There is a distinction between *set* as applied to things, and the mathematical notion of set (which is presumably the sense that O&G had in mind – not a set of tennis!). The more common definition seen in 28. seems to apply to a bounded collection of specific things, which could “naturally” be construed as a concept belonging at the object level. The mathematical notion, on the other hand refers to abstract collections that might have infinite size. What defines the set is an abstract definition that makes the members all alike in some way. The notion of abstraction therefore fits the concept perfectly well, in spite of O&G’s claim to the contrary. Perhaps their intuition about the object-level status of *set* is based on the sense expressed in 28., not the one they were citing?

On the other hand, there is nothing inherently privileged about `attribute` that makes it “natural” to consider it a meta-level concept in preference to the other terms in the example. The gloss for the relevant sense of `space`, for example, reads as: *the unlimited expanse in which everything is located; "they tested his ability to locate objects in space"*. Space, then, is an abstraction that relates to all objects. True, it might not be a particularly useful abstraction for use in meta-level modeling, but it is an abstraction nonetheless.

We do not deny that it can be useful to introduce meta-level concepts to capture useful generalizations for a given application. These meta-level concepts are abstractions since they relate to a class of concepts at the object level. But meta-level is not a cognitive construct and as such not all abstractions become meta-level concepts. WordNet provides a candidate set of cognitive abstractions, of which only some are useful as meta-level concepts in a given application. But again, the cognitively inspired structure of WordNet should not be required to encode a distinction that is not cognitive in nature.

## 5 Conclusion

Implementing a knowledge based system involves making a large number of simplifying assumptions. The distinctions and classifications reflect particular views of the world enforced by our logical formalisms and cognitive apparatus. Cognitive science has always played an inspirational role in this enterprise because the mind seems to have solved many problems we struggle with. The problems facing researchers of interoperability and the semantic web pose such challenges. Should we try and use theories of cognition to inform our choices? O&G clearly think we should.

Yet in spite of the enthusiasm and pledges of allegiance, O&G propose to restructure the WordNet ontology according to their principles derived, as far as we can tell, primarily from philosophical and metaphysical conjecture. They are of course entitled to do this but, as they themselves point out, our choices must be made explicit. If we genuinely chose to pursue the cognitive track then we must face the complications this brings with it, and perhaps be prepared to re evaluate our long cherished views of the world. The real challenge is to formalize the cognitive facts and theories in useful ways, and not to change them. To this end, we are pursuing two primary research goals. First we attempt to make explicit the hidden cognitive assumptions in WordNet, as we have started in this paper. From this we can derive ways to map WordNet knowledge to existing formal methods and use them in constructing formal ontologies. The more difficult but ultimately rewarding goal is to re evaluate the existing formal methods with an eye towards improving them to the point of displaying "more human" like behaviors.

## References


Alvarez, J. Integrating the WordNet Ontology into a Description Logic System Research Report LSI-00-40-R, Universitat Politcnica de Catalunya (UPC) (2002)

Bloom, P. How Children Learn the Meanings of Words Bradford Books, MIT Press, Cambridge, MA. (2002)

Baader, F. and Nutt, W Basic description logics In F. Baader, D. Calvanese, D. McGuinness, D. Nardi, P. Patel-Schneider (Eds.) The Description Logic Handbook: Theory, Implementation and Applications Cambridge University Press (2003)

Borgida, A. and Brachman, R. Conceptual Modeling with Description Logics In F. Baader, D. Calvanese, D. McGuinness, D. Nardi, P. Patel-Schneider (Eds.) The Description Logic Handbook: Theory, Implementation and Applications Cambridge University Press (2003)

Brachman, R. What IS-A Is and Isn't: An Analysis of Taxonomic Links in Semantic Networks. IEEE Computer, 16(10): 30–36 (1983)

Gangemi A., Guarino N., Masolo C., Oltramari, A. AI Magazine 24(3): Fall 2003, (2003) 13-24.

Gangemi A., Navigli R., Velardi P. Axiomatizing WordNet Glosses in the OntoWordNet Project. The Workshop on Human Language Technology for the Semantic Web and Web Services, 2nd International Semantic Web Conference ( ISWC2003). Sanibel Island, Florida, 20-23 October (2003).

Jackendoff, R. Semantics and Cognition Cambridge, Mass. MIT Press (1983)

Miller, G. Wizard of the New Wordsmiths: His idea to link words rewrote the dictionary The Star Ledger, January 22, (2002) http://www.nj.com/news/ledger/index.ssf?/news/stories/0122wordnet.html (last accessed 10/3/2004).

Miller, G. Nouns in WordNet C. Fellbaum (Ed.) WordNet: An Electronic Lexical Database MIT Press, Cambridge, MA. (1998)

Missikoff, M., Navigli, R., and Velardi, P. I. Horrocks and J. Hendler (Eds.): ISWC 2002, LNCS 2342, (2002) pp. 39–53.

Oltramari, A., Gangemi A., Guarino N., Masolo C. Proceedings of LREC2002 (OntoLex workshop). Las Palmas, Spain. (2002)

Pustejovsky The Generative Lexicon. Computational Linguistics, 17 (1991) 409–441.

Stevenson, M. Word Sense Disambiguation: The Case for Combinations of Knowledge Sources, CSLI Publications, Stanford, California (2003)

Welty, C. A., and Ferrucci, D. A. What's in an Instance? RPI Computer Science Technical Report #94-18 (1994)

Wierzbicka, A Apples are not a 'kind of fruit': the semantics of human categorization. American Ethnologist (1984) 313–328