

Roles: One Dead Armadillo on WordNet's Speedway to Ontology

Martin Trautwein and Pierre Grenon

Institute for Formal Ontology and Medical Information Science, University of Leipzig
Email: martin.trautwein@ifomis.uni-leipzig.de,
pierre.grenon@ifomis.uni-leipzig.de

Abstract. We assume that the ontological structure of the common-sense world, and thus of human knowledge about this world, is organized in networks rather than in hierarchies. Thus, using the taxonomies that semantic relations generate in WordNet as the only source for the reconstruction of ontological information must fail at some point. Comparing the ontological structures underlying roles to WordNet representations, we demonstrate that the power of lexical semantics to abstract over contexts distorts the taxonomic order of a conceivable ontology. Approaches trying to adjust the semantics of WordNet relations, in order to reach a higher ontological adequacy, unintentionally produce artifacts deriving from differences between the frequency of contexts, and from metonymy-like reference to ontological relations.

1 Introduction

Although WordNet was designed as a semantic dictionary, many applications have put emphasis on the fact that the semantic content of the conceptual entries of WordNet depict common-sense world knowledge and thus reflect common-sense ontology. [4] notices however that “WordNet is only really serviceable as an ontology if some of its lexical links are interpreted according to a formal semantics that tells us something about (our conceptualization of) ‘the world’ and not (just) about the language.” In this light, the authors cited propose several improvements of the hierarchizing semantic relations in WordNet [3,4] in accordance with constraint-based Formal Ontology. In contrast, we want to put more emphasis on the question of which kinds of semantic links can be readily interpreted semantically as ontological relations and which can not. We show that in fact, some of WordNets semantic relations can locally be regarded as taxonomic, whereas the network as a whole can not or should not be converted one-to-one into a taxonomic ontological framework.

The semantic descriptions of WordNet concepts provides links to two distinct forms of knowledge: common-sense knowledge and semantic knowledge. Thus, a concept involves two aspects of information: (i) information about concrete world contexts and their spatiotemporal structure, and (ii) information about epistemology and grammaticalization, merging (our knowledge about) a set of concrete world contexts by means of abstraction. This second kind of information provides the criteria for associating an entity to a certain concept or lexical entry (classifying criteria).

Only the first aspect provides genuine ontological information. When referring and asserting (i.e. constructing concrete contexts of language usage), however, we merge both

the knowledge issued from concrete contexts and the abstracted classifying criteria at hand. Moreover, in order to identify referents or to establish truth, we have to match concrete contexts with the relevant classifying criteria carried by our word semantics. Accordingly, both modules of meaning have to interact and this in a manner prone to reveal their contribution to semantic meaning. WordNet's concepts are based on the synonymy of certain word meanings in a certain set of linguistic (discourse) contexts, but in the end the lexical data is detached from context. As a consequence, most parts of the ontological assumptions and intuitions standing behind WordNet are only implicitly represented by the set-up of synsets and by the semantic relations holding between the corresponding relations. These assumptions and intuitions, however, may be regained through the analysis of the classifying criteria encoded in the semantic aspects of a word meaning. In what follows, we apply our hypothesis to the category of ROLE. We compare the WordNet strategy of representing roles to the rich system of ontological structures of reality that roles crystallize.

2 Types vs. Roles: the WordNet Strategy from the Ontological and Linguistic Standpoints

Recent publications (cf. [4,1] discuss the type-role distinction and WordNet's representation of the two corresponding categories. One of the main arguments for a clear-cut distinction between the two categories is that they differ in the way their instances inherit their properties. This distinction is also crucial for applications which use WordNet as a source for common-sense reasoning. WordNet, however, does not distinguish between roles and types, but organizes both in the same taxonomy. In accordance with the ontological approach to properties of [3,2] proposes that roles should not subsume types in the hyponyme taxonomy. A more radical ontological account, however, could even demand that roles and types should not be represented at all within the same taxonomy.

Although an ontological characterization of properties in the spirit of [3] is indispensable for a proper treatment of classifications and of the qualitative information they provide, such a theory does not anticipate the way in which individual natural languages grammaticalize types and roles, nor how these grammaticalized forms are applied to language usage. With regard to the synonymy criterion that constitutes the conceptualization of WordNet, it is remarkable that the lexicalizations of types and roles generally do not differ in the way they are used in linguistic contexts since in most contexts, a role expression can replace a type expression and vice versa without changing the referential or truth-conditional value. Ranking immediately behind proper names, both type and role expressions (such as *man*, *teacher*, *speaker*, or *speaking*, *declaring*, *verbalizing*) have a strong identifying potential and thus sufficiently restrict a class of referents, e.g. in a nominal description (cf. (1)).

Claire had a dispute with James / her friend / a theologian. (1)

Verbs do not differ in this point: the sortal verb *swallow* can be replaced using verbs such as *eat* or *ingest* which express roles of the process referred to (cf. (2)).

Claire swallowed / ate / ingested a fly. (2)

Beyond this, co-referential descriptions often use hypo/hyponyms and role-expressions as sortals in order to bridge the gap between co-referential terms (cf. sample sentences (3))

and (4); co-referring terms are underlined).

- (3) *Hannah observed a hedgehog. She picked the animal up.* (via hypernymy)
 (4) *She turned quickly.Her vigilant reaction saved the armadillo's life.* (via a role)

3 Lexicalization of Roles and Context Frequency

The usage of nouns and verbs in nominal and verbal descriptions demonstrates that, from the perspective of linguistics, the need for differentiating between the treatments of type-, subtype-, supertype-, and role-denoting lexical items is not conspicuous. Nonetheless, since distinguishing types and roles is all the more desirable from the ontologist view, some authors are looking for ways out. [1] observes that some cases of troponymy are heterogeneous, i.e. they do not fit properly in the troponymy, e.g. *swim*^{v1} < *move*^{v1} (troponym) vs. *swim*^{v1} < / *exercise*^{v4} (non-troponymy).¹ In view of the expressiveness of WordNet used by NLP applications, however, [1] pleads for including links as those between *swim*^{v1} and *exercise*^{v4} in the lexical DB, and proposes to introduce additional, autonomous semantic relations into WordNet in order to capture the specific relationship between entities and the roles they carry. These relations, *para-hyponymy* for nouns and *para-troponymy* for verbs, omit the necessity condition which, in contrast, holds for the regular hyponymy and troponymy relations. The linguistic tests [1, p. 27f] are shown in (5) to (8).

- (5) *X's and other Y's & ¬(It's an X, but it's not a Y)* X is a hyponym of Y
 (6) *X's and other Y's & It's an X, but it's not a Y* X is a para-hyponym of Y
 (7) *X'ing and other manners of Y'ing* X is a troponym of Y
 & ¬(It's an X event, but it's not a Y event)
 (8) *X'ing and other manners of Y'ing* X is a para-troponym of Y
 & It's an X event, but it's not a Y event

[1] assumes that for some concepts the 'role' aspect is more important than the 'supertype' aspect, e.g. *jog*^{v1} tends to be interpreted as a para-troponym of *exercise*^{v4} rather than a troponym of *move*^{v1}. This point signifies that the weightiness of the semantic relations is gradual and vague. The links in Fellbaum's examples seem plausible since the sample word meanings occur frequently in contexts. But what about less frequent, less common, or less expectable contexts? Attending a course of survival training, you will probably soon discover that flies can perfectly adopt the role of food. A bottle can serve as a musical instrument. Singing (and, in particular, bad singing) can not only produce sound but may also amuse people. Grinning, finally, may be more than just a sign of amusement but may offend somebody (e.g. a bad singer). So do we also have to include links for these cases?

¹ Words given in bold italics and marked with an superscript symbolize word meanings taken from WordNet v1.7.1. The superscript indicates the entry number of the word sense of the noun (n) or the verb (v) database. Words indexed in such a way stand for the entire synset to which the corresponding word meanings belongs.

- (9) $fly^{n1} <_{para-hyponymy} food^{n2}$
 $sing^{v1,v2,v3} <_{para-troponymy} amuse^{v2}$

At least, they pass the linguistic test for para-hyponymes and para-troponymes.

- (10) *flies and other forms of food* & *It's a fly, but it's not food.*
 (11) *singing and other manners of amusing somebody*
 & *It's singing, but it's not amusing people.*

4 The Ontological Nature of Roles

The above examples indicate a general problem: some roles have (almost) no specific range of types of entities to which, in an adequate context, they may apply. Think of roles such as TOOL or INSTRUMENT which may be carried by almost every (natural or artificial) entity. The options available to the WordNet designers or users for escaping these problems are not really convenient. They (i) could drop all such links and loose the information. They (ii) could add more and more para-links to the DB. But where is the limit? Finally, they (iii) could link the more unspecific roles, not to possible candidates of a set of hyponymes of a certain concept, but to the hypernymic concept itself: *living_thing*ⁿ >_{meronymy} . . . *nutrient*ⁿ . . . <_{para-hyponymy} *food*ⁿ². Thereby the semantic relations would sketch approximately the encyclopedic information which is enclosed in the gloss of *food*ⁿ²: “any solid substance [. . .] that is used as a source of nourishment”. At least, this strategy would regain the information lost in the abstraction process, when the ontological richness of the possible contexts was reduced to underspecified classifying criteria. But it would also call for an ontological analysis of the world structure underlying every single concept. In order to avoid arbitrary or ad-hoc local decisions, such an ontological analysis has to satisfy certain constraints, e.g. it has to be systematic, consistent, has to optimize disambiguation, and so forth.

We will not give such an exhaustive analysis here, but restrict to some considerations which shed light on the origins of the problems mentioned above. We adopt the characterization of roles by [3] who distinguish between material and formal roles. We agree with the authors that roles generally are dependent entities, in the sense that they always depend on the existence of a further entity. A further interesting point with respect to NL encoding of roles is a different kind of dependency, however. Role expressions and their contextual interpretations always depend on a certain ‘domain of obtainment’, which is some kind of reference point, a certain respect, a particular perspective, or value of comparison. Examples for domains of obtainment are:

- a space of subjective mental or emotional states, e.g. *joy*ⁿ¹, *amuse*^{v1,v2};
- a quality space or scale: *relief*^{n7,n8,n9}, *fail*^{v2, v5};
- comparison to subjective expectations, e.g. *delay*ⁿ¹, *surprise*^{v1,v2}.

Corresponding to the distinction between common-sense world knowledge and linguistic knowledge, NL semantics always includes two aspects of meaning: on the one hand, it

denotes structural properties of possible referents; on the other hand, it restricts the set of possible contexts in which the item can be used (i.e. grammatical contexts, contexts of utterance, and world contexts in which referring expressions are linked to referents and the truth of propositions is evaluated). Therefore it is not enough to focus on the first, the denotational, aspect of word meaning in order to extract world knowledge. We also have to understand how the second, the abstracting, aspect interacts with the first aspect. Hence, an interesting factor for a characterization of roles is how these contextual dependencies and the domain of obtainment of a role-expression are realized or predetermined by lexical entries.

Material Roles in the lexicalization of WordNet are, for instance, *student*ⁿ¹ or *announce*^{v2}. Lexical items encoding material roles often specify the types of entities which may carry the role, and they also co-lexicalize the domain of obtainment (e.g. cultural status). This means that these items do not so much depend on the perspective of a given linguistic context, but rather predetermine the range of possible contexts which might enclose them.

Among **formal roles**, we distinguish between **thematic roles** and **schematic roles**. The ontological structure of thematic roles is the least complex. They correspond to a formal relation holding between two entities of the same or of different top-level categories. Examples are *causal_agent*ⁿ, *product*ⁿ³ or *perform*^{v1}. Lexicalizations of thematic roles do not specify an obtainment domain. Therefore this domain has to be specified by the context. The product of an orchestral performance can be seen as a sound (i.e. a particular), but also as the joy or as a headache of (parts of) the audience.

Examples of WordNet's lexicalizations of **schematic roles** are *speaker*ⁿ¹, *food*ⁿ², *exercise*^{v4}. Schematic roles are based on schemas, i.e. complex chains or networks of categories and relations. Schematic-role expressions may also provide more concrete information about the domains to which those concrete entities and relations belong that are given by the linguistic or world context and which fit into the schema.

So FOOD, for instance, suggests that the speaker thinks of groceries. But as we have seen in the fly case, also entities of an 'exotic' type may fall under a schema node, provided that it satisfies the ontological requirements. An illustration for the schema FOOD is given in Figure 1.

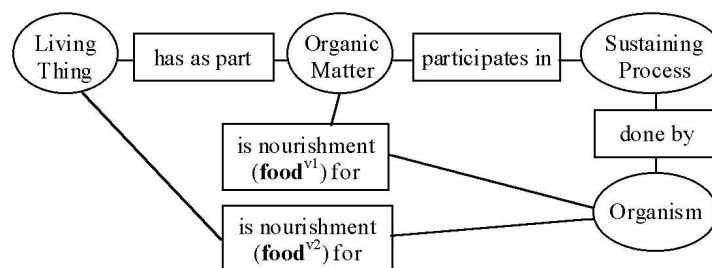


Fig. 1. A possible schema for the schematic role FOOD. Circles symbolize types, boxes stand for relations

Food may be both stuff that an organism ingests and processes (*food*ⁿ²) or the narrower notion of what actually is nourishing in *food*ⁿ², namely *food*ⁿ¹. It is not surprising that some

$food^{n2}$ will have $food^{n1}$ as parts (we abstract from the fact that food can be relative to a consumer). According to the schema in Figure 1, an instance of a living thing such as a fly is not itself $food^{n1}$. Rather, it is a source of $food^{n1}$ in that it contains parts which may play the role of a nutrient. Source of nourishment is then ambiguous and the different senses ought not to be conflated. Claire swallowing a fly brings about a role played by the fly, not a type. The fly plays the role of nourishment insofar as it contains nutritive parts. It is food insofar as it instantiates the schema of Figure 1. The nutritive claim applying to the fly is thus in essence metonymical. It is the most direct link, $food^{v2}$, and not the more specific and rigorous link, $food^{v1}$, which applies between a living thing (e.g. the fly) and an organism (e.g. Claire).

5 Conclusions

We have used the example of roles to motivate our claim that ontologizing WordNet means unveiling the implicit ontological structures which support lexicalization rather than merely turn the network into a taxonomy. The ontological variety of kinds of roles allows to differentiate between aspects of common-sense knowledge which purely linguistically motivated features race over. Driving WordNet towards ontological adequacy will in effect transform the lexical database into a knowledge base. The challenge is to preserve the richness of semantic information while operating this transformation. This would mean mobilizing other representational tools rather than merely altering the existing semantic relations.

References

1. Fellbaum, C.: Parallel Hierarchies in the Verb Lexicon. In: Proceedings of ‘The Ontologies and Lexical Knowledge Bases’ Workshop (OntoLex02), 27th May 2002, Las Palmas, Spain. Online source <http://www.bultreebank.org/OntoLex02Proceedings.html> (2002).
2. Gangemi, A., Guarino, N., Oltramari A.: Conceptual Analysis of Lexical Taxonomies – the Case of WordNet Top-Level. In C. Welty, B. Smith (eds.), Proceedings of FOIS 2001, ACM press (2001) 285–296.
3. Guarino, N. and Welty, C. A Formal Ontology of Properties. In R. Dieng and O. Corby (eds.), Knowledge Engineering and Knowledge Management: Methods, Models and Tools. 12th International Conference, EKAW2000. Springer Verlag (2000) 97–112.
4. Oltramari A., Gangemi A., Guarino N., Masolo C.: Restructuring WordNet’s Top-Level: The OntoClean approach. Proceedings of LREC2002 (OntoLex workshop). Las Palmas, Spain (2002).