Formal redundancy and consistency checking rules for the lexical database WordNet™ 1.5

Dietrich H. Fischer, GMD-IPSI

Abstract

In a manually built-up semantic net in which not the concept definitions automatically determine the position of the concepts in the net, but rather the links coded by the lexicographers, the formal properties of the encoded attributes and relations provide necessary but not sufficient conditions to support maintenance of internal consistency and avoidance of redundancy. According to our experience the potential of this methodology has not yet been fully exploited due to lack of understanding of applicable formal rules, or due to inflexibility of available software tools. Based on a more comprehensive inquiry performed on the lexical database WordNet™ 1.5, this paper presents a selection of pertinent checking rules and the results of their application to WordNet 1.5. Transferable insights are: 1. Semantic relations which are closely related but differing in a checkable property, should be differentiated. 2. Inferable relations - such as the transitive closure of a hierarchical relation or semantic relations induced by lexical ones - need to be taken into account when checking real relations, i.e. directly stored relations. 3. A semantic net needs proper representation of lexical gaps. A disjunctive hypernym, implemented as a set of hypernyms, is considered harmful.

1 Introduction

When building large-scale lexical/semantic resources, subsequent - or better, simultaneous - validation of content is essential. Basic validation includes formal redundancy and consistency checks. Taking WordNet 1.5 as an example, we illustrate the development and application of such rules. The computational environment for this enquiry is TerminologyFramework, an object-oriented generic tool developed to represent and consistently maintain concept-oriented dictionaries of different types, including WordNet [Fischer et al., 1996]. Into this system, the content of WordNet 1.5 was downloaded from the "dict/data.*"-files; some checks were performed during download as part of the operational semantics of our relation definitions, but most of the checks were simulated a-posteriori by database queries. The main aim of the enquiry is not to produce an error report on WordNet 1.5, but to develop a methodology of redundancy and consistency checks for re-use. Therefore, not only WordNet 1.5 has been checked, but our ideas have been developed and checked for validity and relevance using WordNet 1.5. Compared to the thesauri we had previously modeled and downloaded, WordNet 1.5 offers a richer set of semantic and lexical relations which give rise to new questions of redundancy or consistency (cf. [Fischer, 1993]). The more relations introduced into a manually built up net, the more dependencies are created which hold each other in check; once they are explicated, formulated as guidelines and implemented, they can be used to support internal consistency - a necessary, but not sufficient condition for the correctness of a semantic net.

In the following - after having characterized in general the status of formal checks in semantic nets of the WordNet-type - we present and comment a series of constellations which shall give exemplary insight into the topic.

How to read the pictures in this paper

All pictures in this paper are snapshots from the screen, and clippings from a window of TerminologyFramework's graphical browser [Möhr and Rostek, 1993]. The original, and hence uncorrected WordNet 1.5 data is shown as a graph where a node may represent a concept (i.e. synset, named by its first synset-element) or occasionally a term, i.e. synset element linked to its concept by the designation relation which is shown by its inverse from the concept side and is thus labeled *term*. Terms are represented in TerminologyFramework as unique objects having the synset element word/phrase as their possibly homographic name, and a system-generated and maintained homograph counter number which is stored separated from the name string as another term attribute. A term node label is distinguished from a concept label by a lowercase prefix indicating the part of speech (hence *N* for noun concepts / *n* for noun terms, *V* for verb concepts / *v* for verb terms). A concept's label body takes its content from its first synset element which is also transformed on download into a term with this label body. We do not present examples from the adjectives and adverbs which are also contained in the database.

2 Formal checks versus checks based on concept definitions

We suggest that Woods and Schmolze [1992] implicitly also critize semantic nets of the WordNet-type when they characterize "typical frame-based representation systems" by stating that in these systems the semantics of manually linking concepts by *is-a* or *a-kind-of* relations is "strictly operational - defined by how they work and what they cause to happen." Then they continue: "There is no external criterion of correctness to which these decisions should adhere" (p.135).

We wonder what would be an *external* criterion of correctness even in a system of terminological logic, or is there a criterion which checks the adequacy of a manually supplied logic concept definition? In any case, we would call the criteria which are based on the reasoning of the system *internal* ones. However, in the framework of a KL-ONE-type system they can be necessary *and* sufficient - relying on a set of undefined concepts, axioms, and a logic calculus - while in a manually built-up thesaurus or semantic net of the WordNet-type only necessary conditions of correctness can be checked.

Note that Woods and Schmolze [1992] assume that something "happens" when the lexicographer draws a semantic link; they do not criticize - as we do - that sometimes nothing happens or not enough happens which differentiates this type of link from some other, except that they have different labels. Then linking would be nothing but connecting nodes. Our aim is to dig out and apply rules which approximate the intended semantics of the links, or to populate the inventory of those checks which are based on formal properties of the used relations and attributes and their logical dependencies, thus constituting their operational semantics. They are not based - as in a terminological logic system - on an "understanding" of the individual concept definitions. Therefore we call them *formal*.

WordNet concept definitions are given by glosses in natural language - not referring explicitly to the net itself, thus not making use of its disambiguated polysems. These glosses are attached to the concepts, but sometimes they are not available and sometimes they are intermixed with or replaced by usage examples. Checks may be conceived which are based on an understanding of these definitions whether the reader of the glosses is a human reader or a program, based on linguistics. However, sometimes the glosses are faulty and the link structure is correct, and sometimes the contrary is true; thus they could hold each other in check, as long as they are independent of each other.

We continue these notes on formal versus content based checks at the end of this paper (see section 5) by presenting three examples, contrasting both types of checks when we already have explicated the formal ones we then can refer to.

3 Checks to avoid redundancy

From the results presented in the following, one can infer that redundancy-free data was the aim of the WordNet lexicographers, but apparently if by insertions redundant data was generated this occasionally was missed, if the redundant data was out of sight of the human checker or the checking program, i.e. distant by some hierarchical levels from the point of update.

3.1 Short cuts in the generic hierarchy

Figure 1 shows a short cut, spanning four levels of the hyponymy hierarchy from noun concept *lipid* to noun concept *triglyceride*. Because of the **transitivity of the generic hierarchy relation** this link is redundantly stored in WordNet 1.5. Note however, that the short cut, so obvious in this picture, will probably be hidden to the eye, even in a hierarchical line print-out, when all hyponyms of *lipid*, *oil [2]*, *animal oil*, and *glyceride* would be displayed. The numbers of hyponyms from *lipid* down to *glyceride* are 4, 32, 29 and 1 respectively.

Note that *'2'* in the label *N. oil [2]* is the homograph counter number, generated by TerminologyFramework. This number need not be identical with the sense number shown by WordNet's read-only browser.

The database contained only 11 short cuts with respect to the hyponym / hypernym hierarchy for noun concepts and 2 short cuts with respect to the troponym / troponymOf hierarchy for verb concepts.

Note however, that the result of the redundancy check is valid if and only if the premises are valid: If one of the non-redundant hyponym links is wrong, and has to be removed the link diagnosed as redundant may be correct and non-redundant. For the cases of redundancy mentioned here the premises apparently were correct.

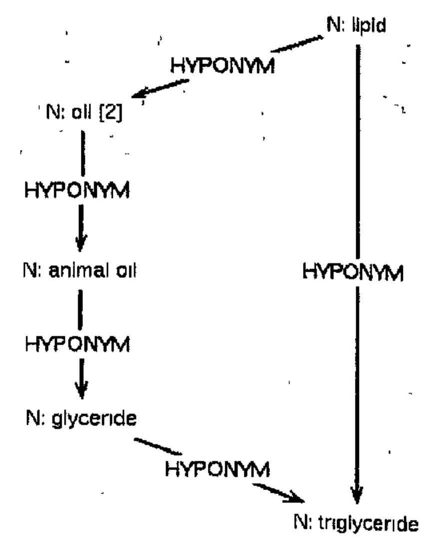

Figure 1 A short cut in the generic noun hierarchy

3.2 Are short cuts in the meronymic hierarchy redundancies?

Short cuts in the meronymic hierarchy - as represented in Wordnet 1.5 by the semantic relations *part/partOf*, *member/group*, and *ingredient/substance* - have also been supervised by the system on download. 73 short cuts were detected, all in the *part/partOf* hierarchy However these can be judged a-priori to be redundancies if and only if transitivity holds for the *part/partOf* relation, and this depends on definitions (not given by WordNet) and actual usage.

A preliminary result of our analysis is: 1. The lexicographers implicitly made heavy use of transitivity, otherwise the data would be highly incomplete; in other words, if transitivity of the partOf-relation had not been presupposed there should be many more short cuts. 2. By an occasional distorted use of the partOf-relation applied to individual concepts transitivity was invalidated; this pertains to cases of the type: "The Alps are part of Yugoslavia; the Alps are part of France, etc.... ". With respect to generic concepts (not individual ones, as the *Alps*) there is also an acceptable example of an implicit exclusive disjunctive partOf-value-set. These disjunctive partOf-value-sets invalidate transitivity and inheritance of meronymic relationships through generic relationships. These cases have to be singled out, separated into a special relationship or simply corrected by introducing e.g. the concept *French Alps* etc.; then for the rest transitivity holds and short cut checking and checking for redundancy by inheritance is meaningful.

Meronymy in WordNet needs the space of another paper, see also [Bloksma et al., 1996], and [Priss, 1996].

3.3 Redundancy by implication of different semantic relations

On the class of verb concepts two semantic relations are defined which are not logically independent: Troponymy (i.e. hypernymy on verbs) and entailment. Example: *limp* is a troponym of (or a special way to perform) *walk*, and *snore* is an entailment of *sleep*, if simulated snoring is not snoring. It is stated in Fellbaum's WordNet paper ([Miller et al., 1993]) that "Troponymy is a particular kind of entailment, in that every troponym V1 of a more general verb V2 also entails V2" (p. 47).

Accordingly, the database does not include a verb V1 which is a *direct* troponym of a verb V2 and directly entails V2, because the latter link would be redundant. However, because of the assumed transitivity of troponymy, the database stores only direct troponym links (2 exceptions, see above subsection 3.1), and taking this into account, the rule should be formulated more explicitly: If a verb concept V1 is a *direct or indirect* troponym of V2, then V1 entails V2. As a formula: (V1 troponymOf* V2) => (V1 entailment* V2). (Here the star operator designates the transitive closures of the relations it is applied to.) Hence the check must look for non-empty overlaps of these two virtual relations. Result: The database included 15 redundant *direct* entailments; one is shown in Figure 2.

Figure 2 also includes an overlap with indirect entailment (from *massage [2]* via *rub [1]* to *touch [7]*); this however vanishes if the redundant direct link (from *rub [1]* to *touch [7]*) is removed, but what is to do, if the troponym chain from *manipulate [2]* to *rub [1]* would not exist in the database? The constellation then would not be tractable by automatic update. The database includes another three *indirect* redundant entailments of this type.

With hindsight and foresight we are emphasizing the following: The logic of the generic subsumption demands that *every* instance of a subconcept is also an instance of its superconcepts, otherwise the logic, supposed to be started from, is changed. From this follows that poly-hypernymy has to be treated with care: Multiple superconcepts to be implicitly combined by *or* (disjunction) and not by *and* (conjunction) invalidate "that every troponym V1 of a more general verb V2 also entails V2" (see below Figure 12 and Figure 13).

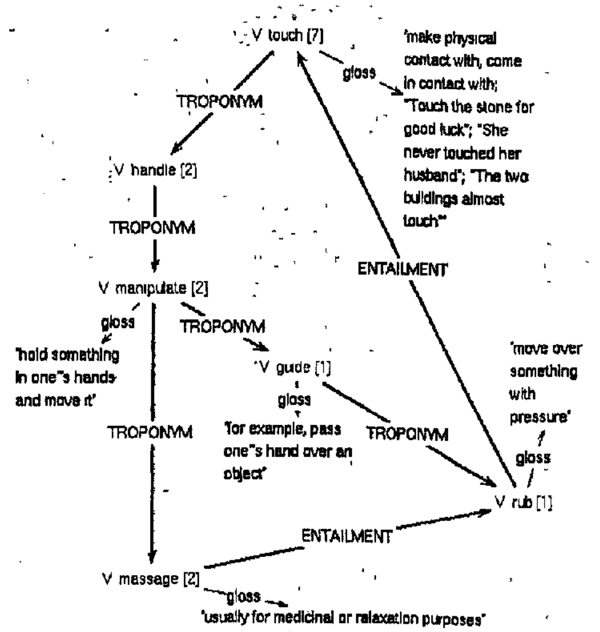

Figure 2 The entailment link from *rub [1]* to *touch [7]* is redundant However the entailment link from *massage [2]* to *rub [1]* is not redundant, - or could *massage [2]* be a troponym of *rub [1]*?

May a troponym of x also be an entailment of x?

It has been pointed out in this subsection that a troponym t of x also entails x. Therefore, if in addition x entails t, then x and t are equivalent, logically. It follows that - if such a constellation occurs, and the premises are valid - the concepts t and x should be merged.

The database contains exactly three constellations of the type in question. They are shown in Figure 3 and in Figure 4 which contains two cases (the pair *exhale [1] / breathe [1]* and the pair *inhale [1] / breathe [1]*).

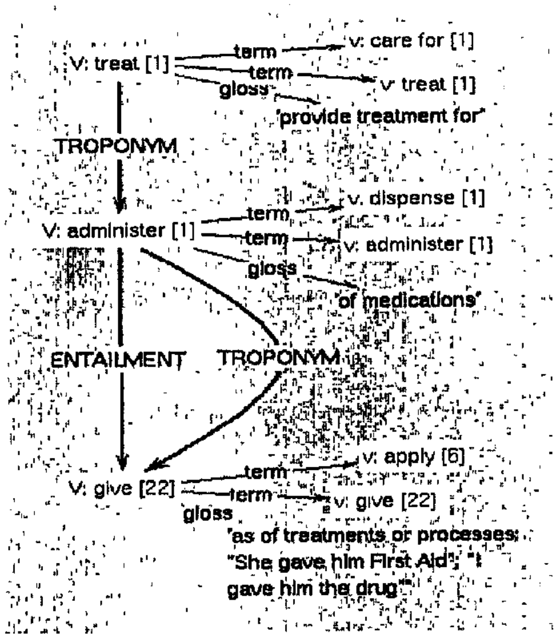

Figure 3 Are *administer [1]* and *give [22]* equivalent?

According to our view the premises in Figure 4 are not valid: *exhale [1]* and *inhale [1]* are not troponyms of *breathe [1]*, because breathing needs exhaling *and* inhaling; both troponym links should be removed. On the other hand, for the constellation in Figure 3 we would assume that the synsets {administer, dispense} and {give, apply} should be merged.

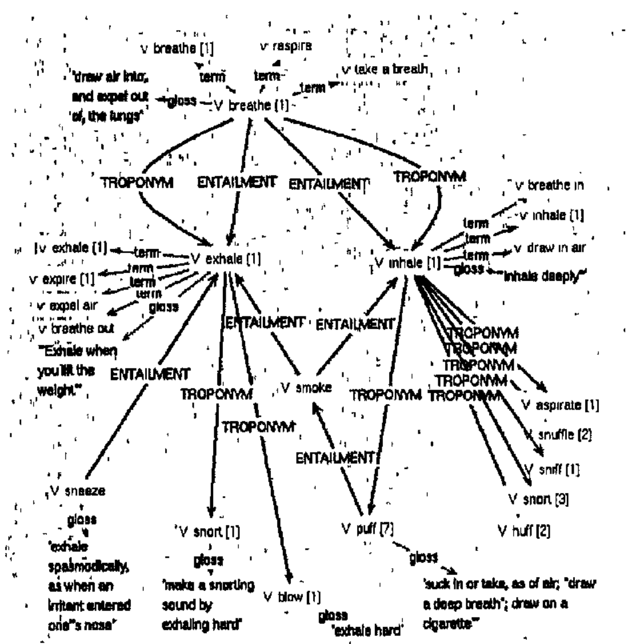

Figure 4 *exhale [1], inhale[1], breathe [1],* and their environment

3.4 Redundancy by inheritance via troponymy

We did not exemplify what is meant by redundancy caused by inheritance via the generic and meronymic relations for nouns (cf. subsection 3.2). Instead we show an example for inheritance via troponymy, see Figure 5.

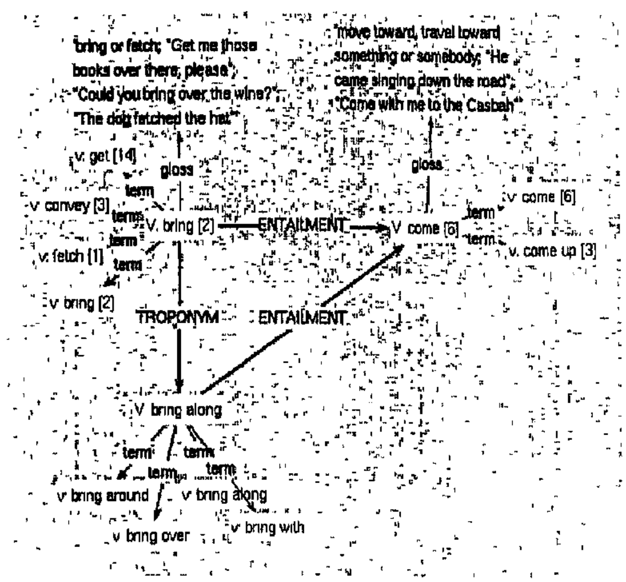

Figure 5 A redundant entailment link from *bring along* to *come [6]*

Nine concepts qualify with redundant inherited *entailment* links, however five of them also qualify as redundant by implication. (see above subsection 3.3). One example of the rest (4 concepts) is shown in Figure 5.
The database contains some further redundancies by inheritance via troponymy.

3.5 Redundancy by synecdoche versus auto-relationships

Our trainee students, in their search for synonyms of a concept *c*, tended to take designations of some superconcept of *c* -whether the superconcept already exists in the thesaurus or not - and added them as synonyms of *c* because these words are also used to denote the concept *c*. The examples given then may be related to the phenomenon of synecdoche: According to the OED synecdoche is a rhetorical "figure by which a more comprehensive term is used for a less comprehensive term and vice versa, a whole for a part or a part for a whole, genus for species or species for genus". While such a synecdochical use of designations, carefully applied in discourse does not lead to polysemy, it would lead to an absurd overload of polysems in a dictionary, if the principle would be transferred to it, even if reduced to *inheritance of designations*, i.e. top-down, not bottom up, and along the generic relation only.

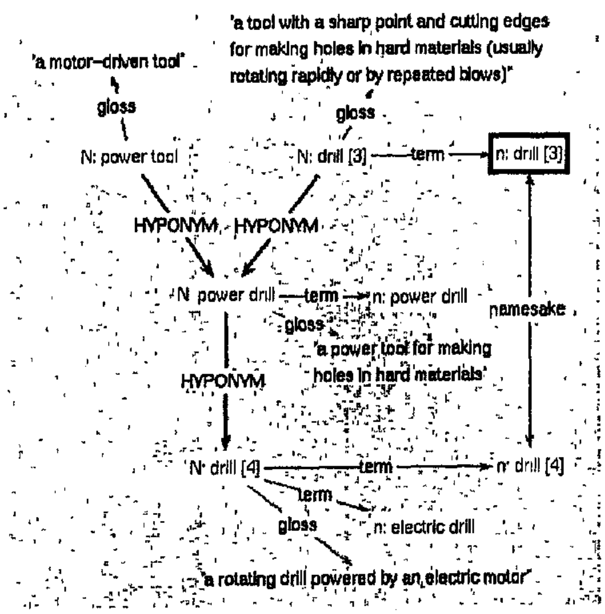

Figure 6 Redundancy by synecdoche or an auto-relationship?

An interactive or batch check may look for direct or indirect superconcept-subconcept pairs which have assigned terms with identical names, thus giving rise to homographs. However, the interactive check's action part should be a warning and not a roll back, otherwise the system would exclude that polysems can be arranged in a generic hierarchy as superconcept and subconcept. Fellbaum [1996] calls this an "Autorelation" (in German). A good example from WordNet is the noun *drink* in the sense of *beverage (drink[2])*,, and in the sense of *alcoholic beverage (drink [3])*. However, what about the example shown in Figure 6?

Is the assignment of term *drill* as a synonym of *electric drill* avoidable redundancy leading to avoidable homography, or not? If not, why was *drill* not also a synonym of *power drill*? If yes or no, and even if we confine ourselves to a synchronic view of language, what would be a working guideline which delineates the assumed good example *drink* from the alleged not so good one (Figure 6)?

In the bilingual *DUDEN OXFORD* we find as a translation of *drill* into German not only the general meaning *Bohrer* (WordNet: *drill [3]*), but also *Bohrmaschine* (WordNet: *power drill*) in the context *(Carpentry, Building)*. Given that the concept or meaning unit is context independent, this may bring to mind that not only translation needs context, but also synset membership.

The database contains only 23 noun homographs selected by the mentioned check. These homographs belong to 21 pairs of noun concepts. Five of them are singled out due to the type of mistake which is shown below in Figure 15: A hyponym link has been coded where a meronymic one would have been correct. The corresponding number of verb homographs is 283, and they belong to 336 pairs of verb concepts, among them five cases which are synecdochical or auto-related triplets.

A similar type of questionable synecdoche pertains to cases such as *{drumhead, head}*, a synset which is a hyponym of *membrane*, but also creates a new sense of *head* (total number of senses: 24).

4 Consistency checks obtained for semantic relations which are induced by lexical relations

When asking for possible checks not yet performed for WordNet's term-term relations (lexical relations) we became aware of their interference with the concept-concept relations (semantic relations). This enquiry is well supported in TerminologyFramework by its easy-to-handle facility to define virtual (i.e. inferable or computable) relations on top of real or virtual relations.

4.1 Exclusivity of hypernymy and antosemy

The WordNet antonymy relation is a lexical relation, i.e. a relation between synset elements of different synsets [Miller *et al.*, 1993]. Therefore it was modeled in TerminologyFramework as a term-term relation. With regards to concepts, two concepts can be defined to be ***opposed*** if at least two of their terms are antonyms. TerminologyFramework implements the *opposed* relation (synonymously: ***antosemy***) as a *virtual*, i.e. computable concept-concept relation (see Fischer et al. 1996). We say that the semantic relation antosemy *is induced by* the lexical relation antonymy.

Antosemy relationships and hypernymy or hyponymy relationships are exclusive to each other, i.e. both relationships cannot hold in conjunction between any pair of concepts. This rule may be based on the feature model of concepts: If we assume a concept representation by features, then hypernymy entails inclusion of all features of the superconcept, and this cannot be compatible with an antosemy between superconcept and subconcept, which may, for example, be based on a meta-antosemy relation between features.

For a check of this rule the overlap of two *virtual* relations has to be determined: antosemy and the transitive closure of hypernymy or troponymy. The number of antosemy relationships is about 800 for nouns, and about 500 for verbs; there was no overlap with direct or indirect hypernymy, and in only two cases there was a non-empty overlap with indirect troponymy. These two violations of the exclusion rule are shown in Figure 7.

Note that there is no overlap with direct troponymy, but with indirect or inferable troponymy. What is wrong with this counter-example? The troponym link from *prove [1]* to *negate [1]* is an error, and it may have its origin in a fallacy: *To prove by negation* is a troponym of *to prove*, but this is different from *to negate* in the sense of *to show to be false*. In other words, one may prove *A* by showing that the *negation of A* is false, but the point is, that the *negation of A* is another object than *A*, i.e. the object to be proved has changed, and indeed, it cannot reasonably be maintained that to negate *A* is a special way to prove *A*

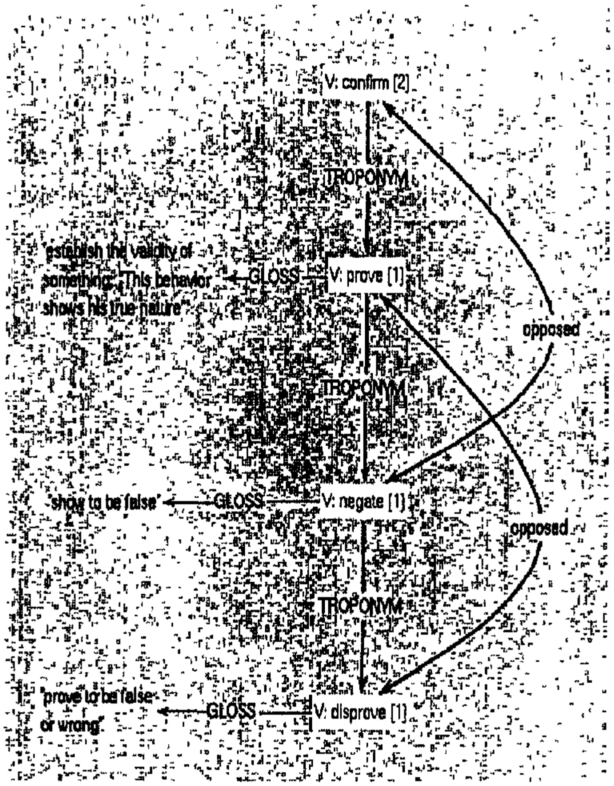

Figure 7 The two violations of the exclusion rule

4.2 About binary and n-ary (n > 2) antonymy and antosemy in WordNet

WordNet 1.5 permits cardinality > 1 for antonymy, however a cardinality check of antonymy is blind with respect to fundamental *semantic* differences of which the cardinality check for the induced relation, the antosemy relation, is more sensitive. We explain this by the following examples from WordNet 1.5:

The verb term *trust [1]* is an antonym of the verb term *mistrust* and of the verb term *distrust [1]*, and both - as variants - are synonyms of each other. Thus, the cardinality of the antonym set of *trust [1]* is 2. However, because these two antonyms are synonyms, the value set of the induced antosemy relation of the concept belonging to *trust [1]* has cardinality 1. Therefore we say that *trust [1]* has *no* genuine multi-value antonymy or, synonymously, it has binary antonymy - although its antonym set has cardinality 2.

On the other hand, there are cases (one is presented in Figure 8 below) where the induced relationship has cardinality 2 and the basic antonym relationship has cardinality 2: In Figure 8 the term *v. arise [3]* has genuine multi-value antonymy or (in this case) ternary antonymy. However, not all cases of cardinality 2 of antosemy are caused by deliberate ternary antonymy. Cardinality 2 may also be caused by an error of a type we try to describe in the next subsection. WordNet does not differentiate between binary and n-ary opposition (n > 2), or, basically, between binary and n-ary antonymy (n > 2), and due to this implicit merging of relations, the lexicographers intentions cannot be automatically checked by a simple cardinality test for antosemy; i.e. we do not know which of the cases of cardinality 2 of antosemy are inconsistent or not. The comparatively low number of these cases (about 200, but only 40 for nouns and verbs) allows for intellectual perusal.

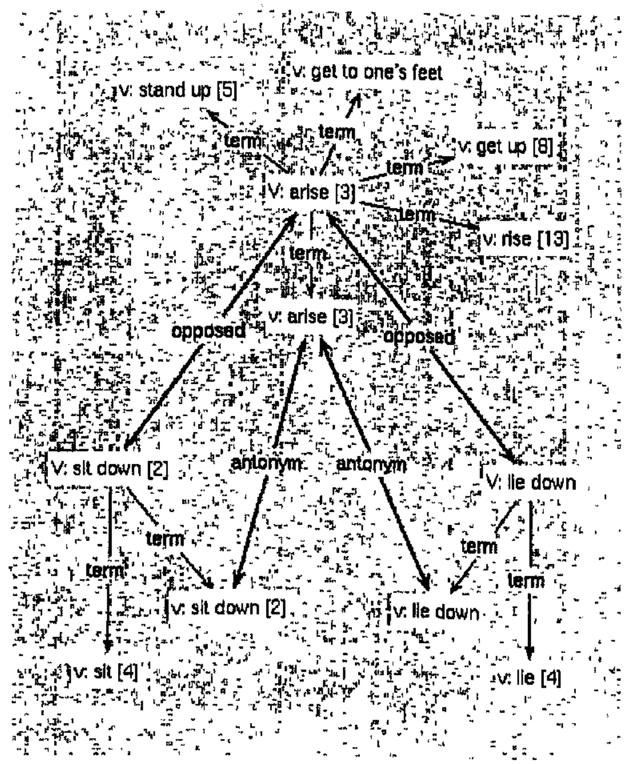

Figure 8 An interesting violation of transitivity

However, if n-ary antonymy is admitted the **law of transitivity** is at stake. We singled out all cases where the transitivity rule for antosemy was not fulfilled, resulting in 29 connection components or 85 involved concepts. All these constellations seem to need correction. The corrections we would suggest are: Cutting of an antonym link, or a merging or a splitting of concepts, and all these cases belong to the type characterized in the next subsection. There was just one case, shown in Figure 8, where we first thought that an antonym link was forgotten (between *v: sit down [2]* and *v: lie down*), and this was caused by the fact that the database contains the following transitive complete group: *lie [2]*, *sit [1]*, *stand [3]* are antonyms of each other (see Figure 13). Coming across other perhaps more harmful cases (see below Figure 12 and Figure 13) of a concept which is to be represented as a disjunction of concepts and itself a lexical gap, we now suggest that the constellation of Figure 8 was intended to say: *arise [3]* is an antosem of the one concept *sit down [2] or lie down*. However, if this single-valued binary relationship is split here into two binary ones linked to the components of the disjunction, then all antonym value sets must be interpreted as *disjunctions*, and transitivity need *not* hold for n-ary antosemy with n > 2. On the other hand, the antonym value set of *trust [1]* is to be interpreted as a conjunction. If the example of Figure 8 is interpreted in this way, we also loose simple symmetry of antosemy: From the mere fact that *the* antosem of *sit down [2]* is *arise [3]* one can no longer infer that *the* antosem of *arise [3]* is *sit down [2]*

4.3 Commutativity of synonymy and antonymy is equivalent to maximal cardinality 1 for antosemy

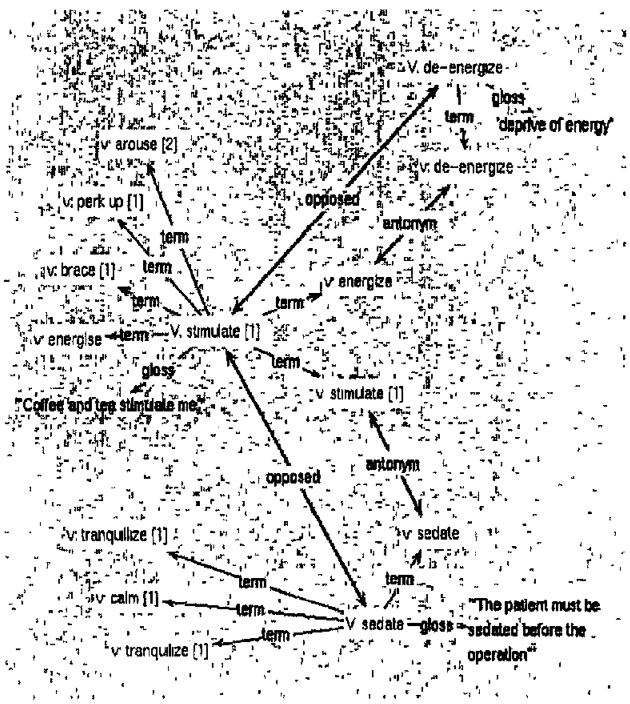

Figure 9 A case which qualifies as a violation of transitivity of antosemy.

Was the constellation shown in Figure 9 intended as ternary antosemy? Is the case equivalent to Figure 8? We do not think so, and this may be backed also by the fact (not shown in this figure) that *V. sedate* is a top concept while *V· de-energize* is not. The constellation shown in Figure 9 may be characterized by another law which is pertinent here if we may assume that binary antosemy was intended. It is an abstraction from one *possible* correction of the constellation: The antonyms of the synonyms *v: energize* and *v: stimulate [1]* should be synonyms, i.e. *v. de-energize* and *v: sedate* should be synonyms, and in order to achieve that, the concepts *V. de-energize* and *V sedate* need to be merged.

Abstracting, the rule postulated for binary antonymy is: For each set of synonyms *S* the set of antonyms of the elements of *S* must again be a set of synonyms or it is empty. By "a set of synonyms" we mean a set of elements which are defined as synonyms. If we would interpret "a set of synonyms" as a synset then the rule reads: For each synset *S* the set of the antonyms of the elements of *S* must be a subset of another synset, the subset may even be empty or be the whole synset. Transforming the rule into a formula, would help to better see that the rule is a kind of law of commutativity between antonymy and synonymy. Another paraphrase of this rule is that binary antonymy is a structure-preserving or homomorphic mapping with respect to synonymy.

There is a simple checking rule for this because the rule is equivalent to maximal cardinality 1 of antosemy. However, this would help for future checking only if WordNet is modified and updated so that binary and n-ary (n >2) antonymy would be different lexical relations. Note that the value sets of binary antonymy need not have maximal cardinality 1, although this is true for binary antosemy.

4.4 Commutativity of antosemy and hypernymy / troponymy

The commutativity rule for antosemy and hypernymy or troponymy - as a *heuristic* rule - has already been suggested by Fischer *et al* . It may be justified by feature inheritance. Formally the heuristic rule of commutativity states: For each concept *c*: If *antosem(c)* is not empty, then the equation *hypernym (antosem(c))* = *antosem (hypernym(c))* or set inclusion in one direction or the other should hold. Applying the rule on the WordNet data has resulted in 91 noun concept and 260 verb concepts fulfilling neither strong (equality) nor weak (set inclusion) commutativity. These cases have not yet been evaluated by a native speaker, but we suggest that they deserve revision. An example is shown in Figure 10.

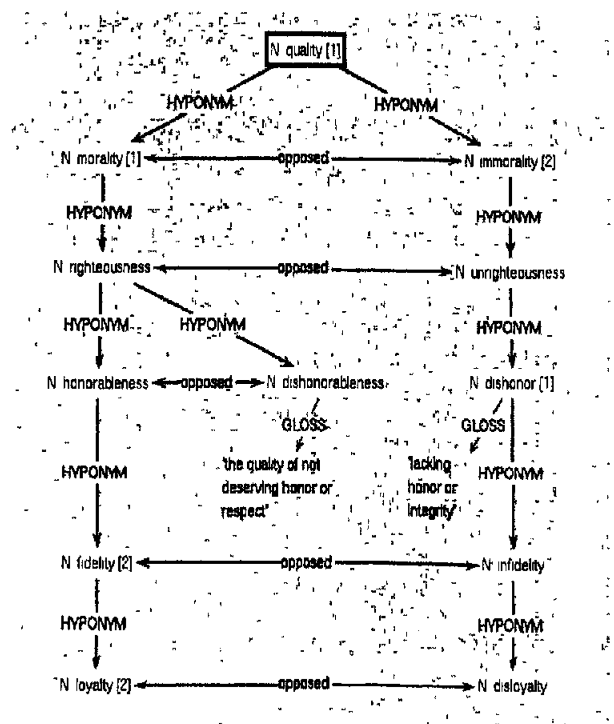

Figure 10 Four of 91 noun concepts fulfilling neither weak nor strong commutativity of antosemy and hypernymy· The four concepts are *honorableness / dishonorableness,* and *fidelity [2] / infidelity.* For *morality [1]* and *immorality [2]* weak commutativity holds

4.5 Disjunctive hypernyms - or may antosems share hyponyms?

If we rely on monotonic feature inheritance the above question needs a negative answer. All the more the empiricist may be interested in convincing counter-examples. These are the results of the respective database search:

For noun concepts (after we have corrected that *irresponsibility* was an antonym of itself) there is exactly one hit, shown in Figure 11.

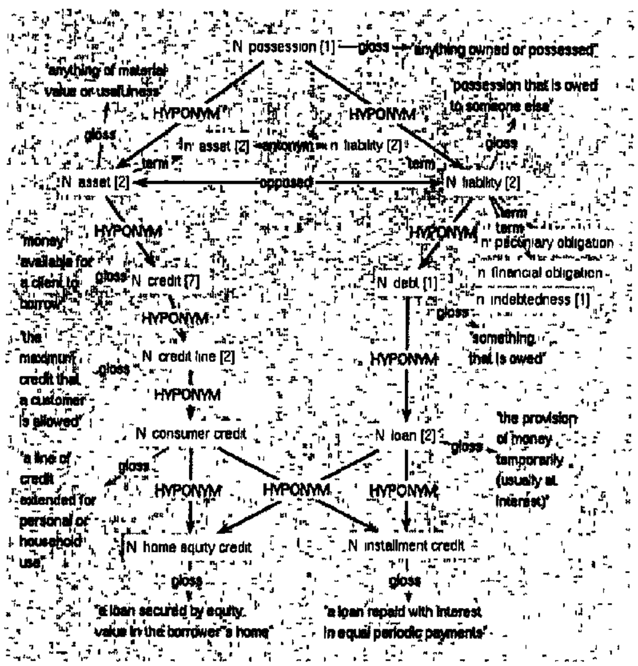

Figure 11 Special credits represented as oxymora (Webster· "a figure of speech in which opposite or contradictory ideas or terms are combined")? Or· Can a credit which has become a real debt be accepted as an asset entry? If that meaning of *credit* exists, where *credit [7]* (an offer) has been transformed into a debt by drawing on the offer, should it not be differentiated from *credit [7]*?

For verb concepts - if we subtract the two cases produced by the error treated in subsection 4.1 - the check detects exactly six pairs or two constellations shown in Figure 12 and Figure 13.

In Figure 12 the verb concept *smuggle* is a troponym, directly shared by *export* and by *import [2]*. Although the gloss says that smuggle is a troponym of *transport*, we may perhaps redefine: "export *or* import illegally". However, if someone is smuggling it cannot be inferred that he is exporting illegally, nor that he is importing illegally, only the disjunction can be inferred. The semantics of the generic subsumption is that *every* instance of a subconcept is also an instance of its superconcept, otherwise the subsumption is not justified. Therefore *smuggle* is a troponym of a concept *export or import* (or yet *transport*?). The stored troponym links have to be removed. Note also that otherwise transitivity of troponymy would be invalidated and short cuts could no longer be deemed a priori to be redundant.

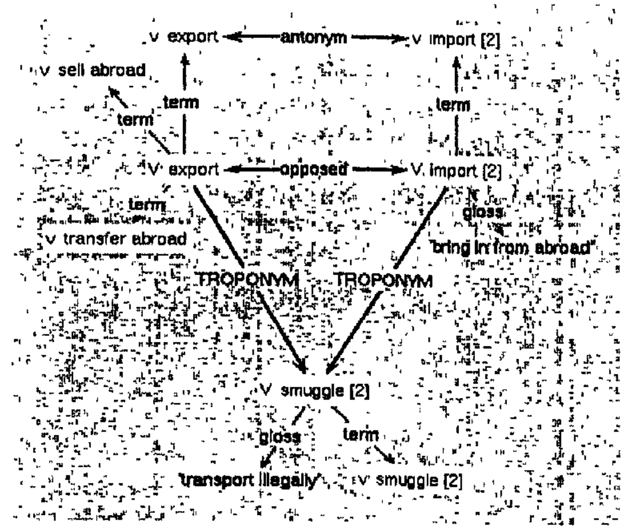

Figure 12 The "disjunctive hypernym", implemented in a way which is harmful: *smuggle [2]* should be a troponym of a concept *export or import [2]*

From this diagnosis we are led to the question what was the general practice in WordNet with respect to multiple direct superconcepts? The database contains 558 noun and only 25 verb concepts with more than one direct superconcept. Among them, of course, also those found by the short cut check (see subsection 3.1), and for them the value set of superconcepts implicitly is a conjunction. Apparently this was aimed at for all nouns, but we saw a case where the hypernym relation implicitly was substituted by *is-used-for* and this made up an *or*. For the verbs we see far more *ands* than *ors*, among the latter also undifferentiated verbs which are troponyms, e.g. of both *change [1]* ("undergo a change") and of *change [3]* ("cause to change"); cf. also below section 5.

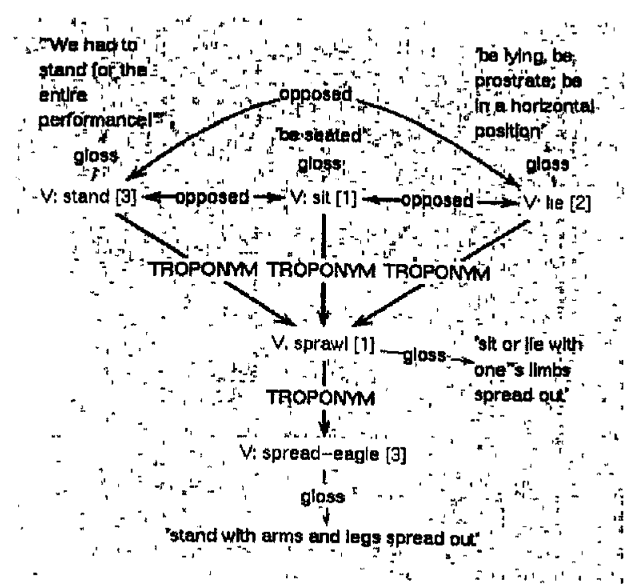

Figure 13 Beyond details to correct we want the reader to see that again subsumption is of the type of the harmfully implemented disjunctive hypernym.

5 Continued: Formal versus content based checking

At the end of section 2 we promised to present three examples which illustrate the difference and the limits of formal checks in contrast to checks based on concept definitions; here they are:

Figure 14 shows an example which was retrieved by the check described in subsection 4.4: Neither weak nor strong commutativity holds for the concept pair *N. employee* and *N: employer*, and the formal rule suggests that their direct superconcept should be opposed. However, reading the gloss of *N. employer* the human reader (or the program analysing the gloss) may infer that the superconcepts need revision: The gloss logically demands the existence of a concept *person or firm* which does not exist in the database, although Fellbaum [1996] argues to consider lexical gaps; however, if one would propose that an *additional hypernym* link to *N. firm* might help, that would be an error of the type discussed in subsection 4.5. Another unfortunate decision would be to create two polysems: 1. An employer who is person, 2. an employer who is a firm. Bloksma *et al* criticized that WordNet practice.

The second example, shown in Figure 15, was among the 21 noun pairs which were retrieved by the check for generic synecdoche or generic auto-relationship. This topic was treated above, see subsection 3.5. The concept pair we present here is *satinwood [1]* and *satinwood [2]*, linked by a hyponym link. Because superconcept and subconcept are both designated by the same word, thus creating homography, they were detected by the check which relates to generic synecdoche. However, this example is neither a case of *generic* auto-relationship nor is it a case of avoidable generic synecdoche: The *hyponym* link between *satinwood [1]* and *satinwood [2]* must be replaced by a *substance* link, also available in WordNet. No formal check could prevent the lexicographer to select the wrong link type unless the synecdoche checker would lie in wait and catch this special case because of name equality. Asking for a check which would be adequate in more generality, we draw the reader's attention to the two superconcept chains which - before they end up in a common ancestor - are headed by *life form* and *object [1]*, in the sense of the negation of life form (see glosses). Assuming that there exist entities which have aspects of a life form as well as aspects of not being a life form, we may be interested whether WordNet reflects this non-Boolean logic view. In fact, in Wordnet 1.5 we find 15,087 subconcepts of *object [1]*, and 13,806 subconcepts of *life form*, and only 6 concepts in the overlap, among them, of course *satinwood [2]*. In 4 of these 6 cases the same error occurred, and in the other 2 cases the *hyponym* link was mistaken for a *member* link, also available in WordNet. WordNet does not encode plain logical incompatibility, to express disjointness of hierarchies, but what about antosemy? Was it by chance that *life form* did not get the term *animate object* which then might associate the used term *inanimate object* as an antonym? Then we would be led back to the check of subsection 4.5: May antosems share hyponyms?

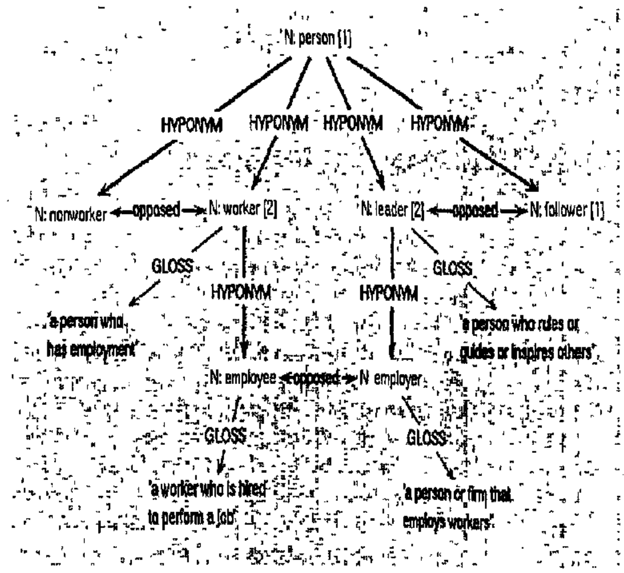

Figure 14 The position of *employer* in the net does not reflect its gloss.

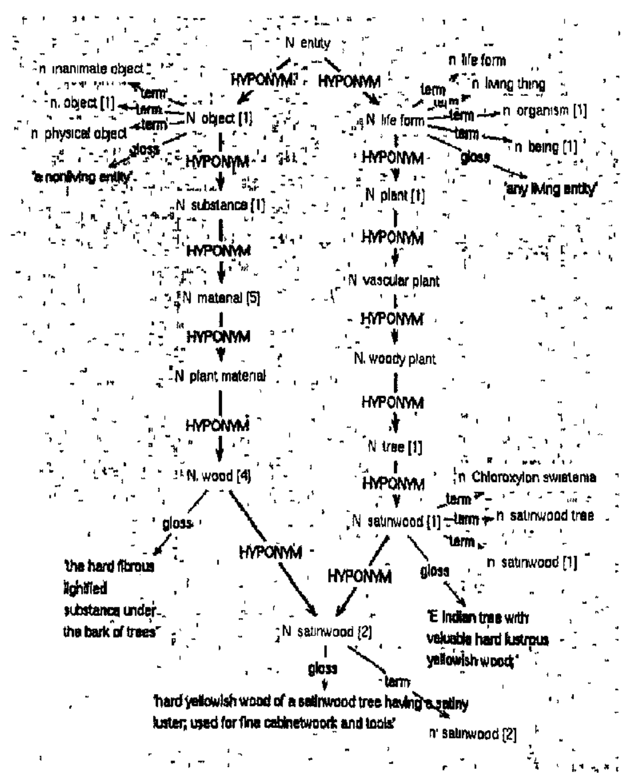

Figure 15 How could this type of false link selection between *satinwood [1]* and *satinwood [2]* be prevented?

We would term it a formal check when for a given generic relation of a concept class the system is asked to retrieve all top concepts (or the direct subconcepts of a possible unifying root), and especially select those which are isolated, i.e. which of these do not have subconcepts themselves. With respect to this check the troponymy hierarchy of WordNet's verb concepts has 573 top verbs, however it is surprising to find that 236 of them are isolated. Now, we already have reached the limits of formal checking because formal checking cannot tell us whether all this was intended or is acceptable, and if not, what to do. A human reader of the glosses may infer that this or that top concept should be subsumed under a more general one, or may have already existing or not existing subconcepts.

6 Conclusion

To our knowledge, the WordNet lexicographers were not supported by dynamic checking on update, or by an easy-to-use database query language for batch-checking, nor by a graphical browser/editor for visual feedback. WordNet's database set-up program, the "grinder", obviously controls consistency, however we are not informed about this. According to TerminologyFramework's error report on download, the "grinder" did not faultlessly perform **range-checking** because it allowed a few semantic (synset-to-synset) pointers where they should be lexical ones (synset-element to synset-element), and in very few cases it missed that a synset-element should not be an antonym of itself (**irreflexivity of antonymy**), and that a synset should not be a hyponym of itself (**irreflexivity, as an entailment of acyclycity of the generic relation**). Beyond these few slips we found more interesting examples of errors or of redundancies which were not detected by chance, but by triggers (created by TerminologyFramework from the specification of the operational semantics of the relations) or mostly by queries to the database, guided by our methodological interest.

A practical lesson is that the design of dictionary relations should be such that they are tractable by formal checking, and this is severely impeded if different relations are merged of which one has the checkable property p and another lacks it; the merged relation has then lost the checkable property p. Examples from WordNet 1.5 are meronymy (which is treated in this paper only by a short note) and antonymy. Another point we had to struggle with was WordNet's treatment of disjunctive hypernyms, especially when they are lexical gaps. The topic of implicit logical junctors in value sets of WordNet's generic and meronymic relations was also treated by Bloksma *et al.*, but the therapy they propose we could not get to like.

This paper does not contain a complete list of checking rules for WordNet 1.5: Whenever we tried to evaluate a rule we got hints for another rule, and we have not yet taken into account all WordNet relations and attributes. Of course, there are important and less important relations, but note, if one takes only the important ones, or the most important relation, the generic relation, then *formal* checking in this type of semantic net is very limited. In any case, formal checking is only a kind of syntax checking, the next step after spelling checking, but 100 new pennies will make up an Euro. Some of our concrete diagnoses may be wrong, or fall short, or become obsolete by a new release, but the questions to be posed for this type of semantic net remain.

Acknowledgments

I am indebted to Melina Alexa and John Bateman for encouraging this work, and to them both and Wiebke Möhr, Renato Reinau, Lothar Rostek, and Ingrid Schmidt for valuable help to improve this paper.

References

[Bloksma *et al.*,1996] Laura Bloksma, Pedro Luis Díez-Orzas, and Piek Vossen. User Requirements and Functional Specification of the EuroWordNet project. Deliverable D001, WP1, LE2-4003, Version 4, Draft, May 14, 1996.

[Fellbaum, 1996] Christiane Fellbaum. WordNet: Ein semantisches Netz als Bedeutungstheorie. In *Bedeutung, Konzepte, Bedeutungskonzepte. Theorie und Anwendung in Linguistik und Psychologie,* pages 211-230. Joachim Grabowski, Gisela Harras and Theo Herrmann, editors. Westdeutscher Verlag, Opladen, 1996.

[Fischer, 1993] Dietrich H. Fischer. Consistency Rules and Triggers for Multilingual Terminology. In *TKE'93 Terminology and Knowledge Engineering,* pages 333-342. Klaus-Dirk Schmitz, editor. INDEKS-Verlag, Frankfurt/M., 1993.

[Fischer *et al.*, 1996] Dietrich H. Fischer, Wiebke Möhr, and Lothar Rostek. A Modular, Object-oriented and Generic Approach for Building Terminology Maintenance Systems. In *TKE'96. Terminology and Knowledge Engineering,* pages 245-258. Christian Galinski and Klaus-Dirk Schmitz, editors. INDEKS-Verlag, Frankfurt/M., 1996.

[Miller *et al*, 1993] George A. Miller,. Richard Beckwith, Christiane Fellbaum, Derek Gross, Katherine Miller, and Randee Tengi. Five Papers on WordNet™. CSL Report 43, Cognitive Science Laboratory, Princeton University , July 1990, Revised August 1993.

[Möhr and Rostek, 1993] Wiebke Möhr and Lothar Rostek. TEDI: An Object-oriented Terminology Editor. In *TKE'93 Terminology and Knowledge Engineering,* pages 363-374. Klaus-Dirk Schmitz, editor. INDEKS-Verlag, Frankfurt/M., 1993.

[Priss, 1996] Uta Priss. Relational Concept Analysis: Semantic Structures in Dictionaries and Lexical Databases. Thesis, TH Darmstadt, October 1996 http://php.indiana.edu/~upriss/papers/cdissi.ps

[Woods and Schmolze, 1992] William A. Woods and James G. Schmolze. The KL-ONE family. In *Semantic Networks in Artificial Intelligence, Part 1* Fritz Lehman, editor. Computer&Mathematics, 23(2-5):133-177, 1992.