

Modeling Attitude, Polarity and Subjectivity in Wordnet

Isa Maks

Faculteit der Letteren
De Boelelaan 1105,
1081 HV Amsterdam
The Netherlands
e.maks@let.vu.nl

Piek Vossen

Faculteit der Letteren
De Boelelaan 1105,
1081 HV Amsterdam
The Netherlands
p.vossen@let.vu.nl

Abstract

In recent years techniques have been developed to mine wordnets for sentiment-bearing words. They annotate synsets with labels for subjectivity and polarity. These techniques assume that all members of a synset are similar with respect to these annotation labels. In this paper we show that this is often not true, especially not when fine-grained polarity and subjectivity is taken into account. Therefore, we propose an extension to the wordnet model which will be better tailored to the description of fine-grained subjectivity and polarity and yet remains useful for the above mentioned mining techniques.

1 Introduction

In recent years much attention has been paid to the automatic detection of opinions, sentiments, beliefs and emotions (*subjectivity*) in text. Most of the techniques use some kind of word list annotated with polarity and subjectivity features. Initially, simple word lists were compiled and automatically annotated for negative or positive polarity only (Kamps et al. (2004), Hatzivassiloglou (1997)). The word lists were limited in size, often consisting of adjectives only. In these lists *friendly* would be tagged as positive, *sad* as negative and *chemical* as neutral. One of the main problems with these lists is that the annotation is at the word level, ignoring the possibility that a word may

have both objective and subjective senses, like the Dutch word *burgerlijk*, meaning *civil* and *narrow-minded*, or may have both positive and negative senses, like the Dutch word *wreed* meaning *cruel* and *fantastic* (cf. examples 2 and 3 below).

Such subjectivity-ambiguous and polarity-ambiguous words may cause major errors in the applications they are used in (Andreevskaia et al. 2006). To overcome this problem, Wiebe (2003, 2006) introduced annotation at the word-sense level. The most recent and comprehensive annotation scheme (Su and Markert (2008)) combines labels for subjectivity and polarity and applies them at the word-sense level.

What most sense-based approaches have in common is that they use wordnet as a lexical resource. Automatic methods are developed to annotate synsets with subjectivity and polarity labels (Esuli et al. (2006), Andreeskva et al. (2006)) For example, the synset below would be labelled as ‘objective’ and having ‘no polarity’:

Ex. (1) (Objective: noPolarity)
lion, king_of_beasts, Panthera_leo
[large gregarious predatory feline
of Africa and India having a tawny
coat with a shaggy mane in the
male]

As the annotation refers to the whole synset, only one set of labels for all members of the synset is provided. The methods obviously assume subjectivity-unambiguous and polarity-unambiguous synsets. The question is whether this is correct? Do these different synonyms

for *lion* all have the same connotation and polarity? One might argue that the *king_of_beast* has a positive connotation while the other synonyms do not. In that case, subjectivity and polarity labels should be stored for each synonym separately. But should that be done by splitting the synset in two synsets or by further defining the synonyms?

In this paper we further elaborate these questions and describe the consequences of these choices for the architecture of the wordnet structure. In the next section we will briefly describe the concepts, models and resources which will be used and referred to in this paper. Section 3 discusses the nature of synsets and in section 4 this is related to subjective language. Section 5 concludes with answering the core question of this paper: is the synset the appropriate unit for the annotation of subjectivity and polarity and if not, what are the alternatives.

2 Method and Resources

2.1 Attitudinal Language

Various definitions have been given for the notion of subjective language. Within the domain of sentiment analysis and opinion mining the most widely used definition is by Wiebe et al. (2005, 2008). They define subjective language as language that is used to express private states such as emotions, beliefs, opinions, doubts, etc.

A detailed and comprehensive definition of subjective language is given by Martin and White's Appraisal Model (2005). They describe subjective language or appraisal language as part of a more comprehensive language model. It distinguishes the following interacting components: (1) Attitude, for the expression of feelings, including emotions, emotional reactions, evaluations of people, things and places; (2) Engagement, for the ways people express their commitments towards what they evaluate and (3) Graduation, for the amplification or weakening the strength of these evaluations and commitments. Attitude, is considered to be the core of the model and is described in three dimensions: Affect, Judgement and Appreciation. Affect words refer to emotions and emotional reactions (e.g. *sorrow, happy, fear, loath, cry, upset*), to desires (e.g. *want, abhor*), etc. Judgement refers to words used for giving moral evaluations of people and the way they behave (e.g. *decep-*

tive, intelligent, neurotic, unreliable, nag). Appreciation refers to the value and evaluations of 'things', performances and natural phenomena. It typically includes aesthetic evaluations (e.g. *beautiful, ugly, shapeless*).

Examples (2a)-(3b) illustrate the annotation of polarity and subjectivity in combination with Attitude Labels based upon the Appraisal Model.

Ex. (2a) (Objective: noPolarity) *burgerlijk* (*civil, civilian*) ex.: *burgerlijke recht-civil law disobedience* [objective - no Attitude]

Ex. (2b) (Subjective: negative) *burgerlijk, bekrompen* (*narrow-minded*) ex.: *zijn bureen zijn vreselijk burgerlijk his neighbours are terribly narrow-minded* [Judgement of character/behaviour]

Ex. (3a) (Subjective: negative) *wreed, hardvochtig, etc.* (*cruel*) ex.: *een wrede despoot a cruel tyrant* [Judgment (of character/behaviour)]

Ex. (3b) (Subjective: positive) *wreed, fantastisch, geweldig, etc.* *fantastic, cool* ex.: *ze rijden daar in vet wrede auto's rond They drive around in really cool cars/* [Appreciation]

In the next sections we will focus on the Attitude component of the Appraisal model. Subjective or attitudinal words are considered – within the scope of this paper – as words belonging to one of the Attitudinal categories.

2.2 Lexical Resources

In this paper we will refer to three different lexical resources: the Dutch Wordnet –(part of the Cornetto database); the Dutch Reference Lexicon (also part of the Cornetto database) and the Princeton Wordnet (versions 3.0 or other if mentioned).

The Cornetto database (Vossen et al. 2008a) combines two resources with different semantic organisations: the Dutch Wordnet with its synset organisation and the Dutch Reference Lexicon with its form-meaning composites or

lexical units. The description of the lexical units includes definitions, usage constraints, selectional restrictions, syntactic behaviours, illustrative contexts, etc. Within the Cornetto Database, each synonym in a synset is linked to the corresponding lexical unit of the Dutch Reference Lexicon.

3 What makes a synset?

To give an answer to the question whether wordnet synsets are appropriate for sentiment encoding, we must first know in more detail what synsets are. What are the criteria for composing a synset and for words to be members of the same synset? Synsets consist of interchangeable words or synonyms. However, according to Miller (1998) synsets do “not entail interchangeability in *all* contexts; by that criterion natural languages have few synonyms.” Therefore, the notion of synonymy in WordNet is less strict: synset synonyms can be interchanged in *some* contexts.

This notion, however, may lead to different criteria for forming synsets, as illustrated by the comparison of the synset of *lion* (cf. example (1)) and the synsets *dog* and *doggie* (cf. examples (5a and b)). There are three variant terms denoting the concept *lion*: a neutral one, a scientific one and an affective one which expresses admiration for the strength and beauty of the animal. All three terms denote the same concept, but one might wonder if there is any context in which *Panthera_Leo* and *King of Beasts* are indeed interchangeable. The notion of synonyms is not only vague with respect to what terms should be included in a synset, it is, of course, equally vague in setting criteria for terms that are *not* synonyms.

Ex. (5a) dog, domestic dog,
Canis familiaris (a member of
the genus Canis..) “the dog
barked all night”

Ex. (5b) pooch, doggie,
doggy, barker, bow-wow (infor-
mal terms for dogs)

Although the terms of synsets (5a) and (5b) refer to the same animal, this time they are not considered synonyms. Style labels (cf. the definition of (5b) *informal terms for dogs*) cause the distribution of the synonyms over different synsets. Usage constraints can be good criteria for splitting up synsets, but - as

can be seen in synset (1) - they are not applied consistently. Furthermore, they seem to contradict the general rule of synonymy. It is probably much easier to find contexts in which *dog* and *doggy* are interchangeable than contexts in which the synonyms of synset (1) are interchangeable.

Other criteria for splitting up synsets are selectional restrictions and syntactic behaviour, which we will not discuss further here.

Thus we see, that the lack of a precise definition of synonyms within the framework of Wordnet gives rise to different sets of criteria ranging from rather loose (example 1) to rather strict ones (example 5). This leads not only to inconsistencies within wordnets but also to differences among wordnets. In the case of EuroWordNet, the number of synonyms per synset, varies from 1.35 to 2.16. In the case of more recent wordnets used in the Kyoto project (Vossen et al. 2008b), the rate varies between 1 for the Chinese wordnet to 1.76 for the English WordNet (cf. Table 1).

Wordnet	Synonyms/synset
Chinese Wordnet 1.6	1.00
Italian Wordnet 2.0	1.36
Japanese Wordnet3.0	1.39
English WordNet 3.0	1.76
Dutch Wordnet 2.0	1.69
Basque Wordnet 1.6	1.66
Spanish Wordnet 1.6	1.49

Table 1: synonyms per synset in different wordnets

Clearly, the English WordNet has most synonyms per synset, even though it is the largest wordnet in size. As the difference between the English and the Dutch Wordnet is relatively small (1.76 versus 1.69), we assume that their criteria for synset splitting are rather similar.

4 Subjective language in Wordnet

In this section we explore how attitudinal language is described in the English and Dutch Wordnets.

4.1 Large synsets

Subjective words seem to cluster together in large synsets more than other words do. The top 50 of the largest synsets of the Dutch Wordnet – ranging from 62 to 16 members - consists for the major part of Judgement and

Appreciation words: mostly evaluations of persons and of the way they behave.

Ex. (6a) judgement of persons:
etterbuil, lelijkerd, schoft,
smeerlap, pokkenlijder, etc.
(*skunk, bastard, son-of-a-bitch*)

Ex. (6b) judgement of persons:
leeghoofd, sufkop, uilskuiken,
uilebal, etc. (*featherbrain,
nitwit, rattlebrain*)

Ex. (6c) judgement of behaviour:
apekool, geouwehoer, wi-
jvenpraat, leuterpraat, quatsch,
lulkoek (*drivel, piffle, rub-
bish, poppycock*)

Some of these synsets could easily be put together in even larger ones. At the same time, they might be split up according to certain criteria, like usage constraints: differences in frequency, differences in style (ranging from informal to vulgar), differences in social group, etc. They have two things in common: firstly, they all have a rather strong connotation and seem to have little denotational meaning. This common feature makes it possible to regard them as synonyms. Secondly, all of them have a clear polarity in most cases, but the intensity of the polarity may differ considerably within a synset. A synset containing *lelijkerd* (*rascal*) and *etterbuil* (*son-of-a-bitch*) might be considered ‘intensity-ambiguous’ as the first synonym is much less negative than the second one. It will be quite difficult to split up synsets using criteria like weak or strong polarity since polarity values are scalar rather than categorical.

4.2 Smaller synsets

Not all attitudinal words are members of large synsets. On the contrary, many of them have more explicit denotational meanings and are not as one-dimensional as the examples given in the previous section. Most Affect words, most Appreciations and many Judgement words as well, refer to specific and complex concepts, which makes them members of different synsets. These synsets, however, are often subject to another type of ambiguity. They are subjectivity-ambiguous as their members all have the same, clear denotation but some of them also have a positive or negative connotation, as in the following example:

Ex. (7)
Nederlander:1, Hollander:1,
kaaskop: 1 (inhabitant of the
Netherlands)

Nederlander and *Hollander* are quite synonymous and interchangeable in many contexts. The third synonym *kaaskop* also refers to a Dutch person. Its literal meaning is ‘cheese head’ and it has a negative, affective meaning which cannot be inferred from the way it is presented in this synset.

5 Pilot

In the previous section we saw some examples of inconsistent synsets. We carried out a small experiment to manually mark the subjectivity of synset members with the aim to determine (1) how consistent or inconsistent synsets are with respect to attitudinal language and (2) if there are any differences between different categories of attitudinal language.

We selected about 75 words including Affect, Engagement and Appreciation words covering nouns, verbs and adjectives. We manually linked the words to the correct synsets and collected some synonyms of the involved synsets. Synsets with only one synonym were excluded. The synset members (approx. 250) were presented in alphabetical order to avoid synset membership associations. Each synonym was given its own definition taken from the lexical unit in the Cornetto database (cf. section 2.2).

The synonyms were individually annotated by two people and were given labels for subjectivity (subjective, objective and both). ‘Subjective’ is defined as belonging to one of the three attitudinal categories (Affect, Engagement or Appraisal); ‘objective’ as not belonging to any of them; and ‘both’ for lexical units permitting both readings. An example of the category ‘both’ is *zwaar* (*heavy*) which can be used in contexts as referring to an objective and precise measure of weight or as merely expressing ‘not light’ often having a negative connotation.

The annotators disagreed in 45 cases. These were excluded from the analysis. For this evaluation, we are only interested in clear cases. The results (Table 2) consist of 205 lexical units linked to 75 synsets. A synset is considered as ambiguous when at least one synonym has a different annotation than the any of the others.

part-of-speech	number of synsets	subjectivity-ambiguous
Adjective	15	3
Noun	35	10
Verb	25	6
<i>total</i>	75	19

Table 2: subjectivity-ambiguous synsets

We see that 25% of the synsets in total (19) are inconsistent in terms of subjectivity. For nouns this is 28%, for verbs 24% and for adjectives 20%. Even though the total set is small compared to the full wordnet, we can still clearly see a pattern.

The following points are noticeable:

- Most ambiguities involve judgements of behaviors and persons, where words denoting the concept of that particular behaviour or trait are mixed with words which denote the same concept but have also a judgement connotation. Synsets (8) – (9) exemplify some of these cases, which are comparable to the earlier examples of *Dutch man/ kaaskop*:

Ex.(8)

[s:neg]scharrel (*flirt*)
[o:neutral]affaire (*love affair*)

Ex.(9)

[o:neutral] zijn (*be, stay*)
[s:pos] vertoeven (*stay, dwell*)
[s:neg] uithangen (*stay*)

- We found no synsets which refer to emotions and expressions of emotions (Affect) that are ambiguous.

To summarize, we have seen that:

(1) The vague criteria for synset membership lead to inconsistencies in a wordnet, also with regard to Attitudinal Language. (2) Attitudinal language is not homogeneous with respect to polarity- and subjectivity-ambiguity. Both Judgement and Appreciation words fall into different categories: they are either words with a strong connotational and weak denotational meaning or, vice versa, words with a strong denotational and a weak connotational meaning. In the first case, the words cluster in large synsets, which in most cases are ambiguous with respect to the intensity of the polarity. In the latter case, the words with connotation often cluster with their objective counterpart into subjectivity-ambiguous synsets.

6 Discussion: splitting or lumping?

How can we solve these synset internal ambiguities? As explained above (section 3), the vague concept of synset permits a strict and a loose approach to synset forming. These approaches lead to two systematic rules that we can apply:

- “One denotation and one connotation per synset”. This rule dictates that ambiguous synsets should be split as they cannot represent different connotations (including positive, negative and neutral). As a consequence, different synsets may have the same denotation (cf. examples 5a and 5b). or

- “One denotation and several connotations per synset”. This rule dictates that all words that denote the same concept are stored together in one synset regardless of their connotations (cf. examples (8) and (9)).

A strict application of the first rule results in a wordnet with a large amount of very small synsets and many hypernym relations to links these synsets to higher concepts. Basically this technique is being applied in many cases in the English WordNet as shown in the following examples:

Ex. (11a) Kraut, Krauthead, Boche, Jerry, Hun (*offensive term for a person of German descent*)

is a direct hyponym of

Ex. (11b) German (*a person of German nationality*)

There are however, some major disadvantages to this approach:

- It leads to an ‘unintuitive’ wordnet structure with hardly any synonyms in the synsets.

- One might ask if the result is correct. *Kraut, Boche, etc.* will then be presented as a direct hyponym of *German* just as *Berliner, Prussian* etc. But *Krautheads* and *Jerries* are not a kind of Germans; they are Germans. The use of these words tells us something about the attitude of the speaker. Splitting synsets will thus lead to ISA-overloading (Guarino 1998), where genuine subtypes of Germans are mixed with synsets that refer to roles or subjective labels. A solution could be to differentiate many different subtypes of hyponymy relations, one for each type of connotation.

An important question is which criteria should apply? Wiebe et al. (2006) propose that “subjectivity labels should provide principled criteria for refining word senses in Wordnet, as well as for clustering similar meanings to create more course-grained sense inventories”. If we return to example (1) with the three variant terms for *lion*, this would imply that this synset will be split into two synsets: an objective one for *lion* and *Pantera-leo* and a subjective one with positive polarity for *King-of-beasts*. This raises the question whether these newly formed synsets are interesting and useful for other applications as well? For example, ‘an informal vs. formal-language’ classifier, or ‘a scientific vs. non-scientific language’ classifier, would prefer yet another type of categorisation of these synonyms. In the end - when all splitting criteria have been applied - all synsets may have been reduced to one member, eliminating the basis of the wordnet structure. This extreme differentiation of information would reduce a wordnet database to a database with lexical units.

A strict application of the second rule results in a wordnet with extremely large synsets. This approach also has some disadvantages:

- For example, the lumping of the following terms into one synset denoting 'man', connotes all kinds of mostly negative and positive qualities and would lead to highly unconventional structures:

Ex.(12)

man, buster, fellow, old boy,
old man, stiff, hunk, rotter,
dirty dog, rat, skunk, stinker,
stinkpot, bum, puke,
crumb, lowlife, scum bag, so-
and-so, git, hunk, ironside,
studd, machoman etc., etc.,
etc.

- Such synsets permit very low interchangeability of synonyms as their connotational meanings differ too much.
- It leads to low interoperability between Wordnets: on the one hand it is relatively easy to link complete synsets of different languages to each other, but on the other hand precise translation equivalence is impossible.
- There is a major loss of information since all information would be specified at the lexical-

unit level and applications need to access both the synsets and each lexical unit individually.

- Just as with the previous solution, it leads to very ‘unintuitive’ synsets.

Therefore, we propose a two-layered and hybrid approach. Hybrid because it will be a combination of applying loose and strict criteria for synset forming depending on the explicitness of subjectivity and polarity characteristics. Thus, synsets are formed by basic connotation groups of positive, negative and neutral connotation. The solution is two-layered because subjectivity will be described at both the synset and the lexical unit level. Subjectivity and polarity are often characteristics which depend on context; lexical units are the typical units where illustrative contexts and usages of words (implicitly or explicitly) are given. At the level of the lexical unit this information can be associated with fine-grained subjectivity and polarity.

Figure (1) illustrates the proposed model with regard to nouns. Following Guarino (1998) a distinction is made between *Roles* and disjoint *Types* to avoid ISA-overloading. Regular *Roles* and *Types* are considered to be neutral with respect to connotation. Words with connotations are considered as “subjective near-synonyms” of their neutral denotational counterparts. They can be near-synonyms of either *Types* or *Roles*. They are themselves always considered to be attitudinal *Roles* as they are only true in a personal worldview (i.e. that of the speaker or writer or of a participant). In our example (cf. figure 1), *woman* is a neutral *Type* which may have neutral *Roles* like *nanny*, *nursemaid*, *prostitute*, and *girlfriend*. A *woman* can also ‘play the role’ of a *cow*, *bitch*, *stunner*, *super woman*, etc., but only according to the subjective opinion of the speaker or writer. The relations are therefore labelled as a subjective negative *Roles*.

By linking synonyms to their respective lexical units, it will be possible to assign a scalar value for fine-grained polarity to each synonym separately. The subtle, internal ambiguity with regard to the intensity of the polarity can thus be handled without splitting synsets any further. Figure (1) shows these synonym-to-lexicalunit-links of one synset only; they express the differences with regard to polarity between *bag*, *cow*, and *bitch* ranging from relatively weak to relatively strong.

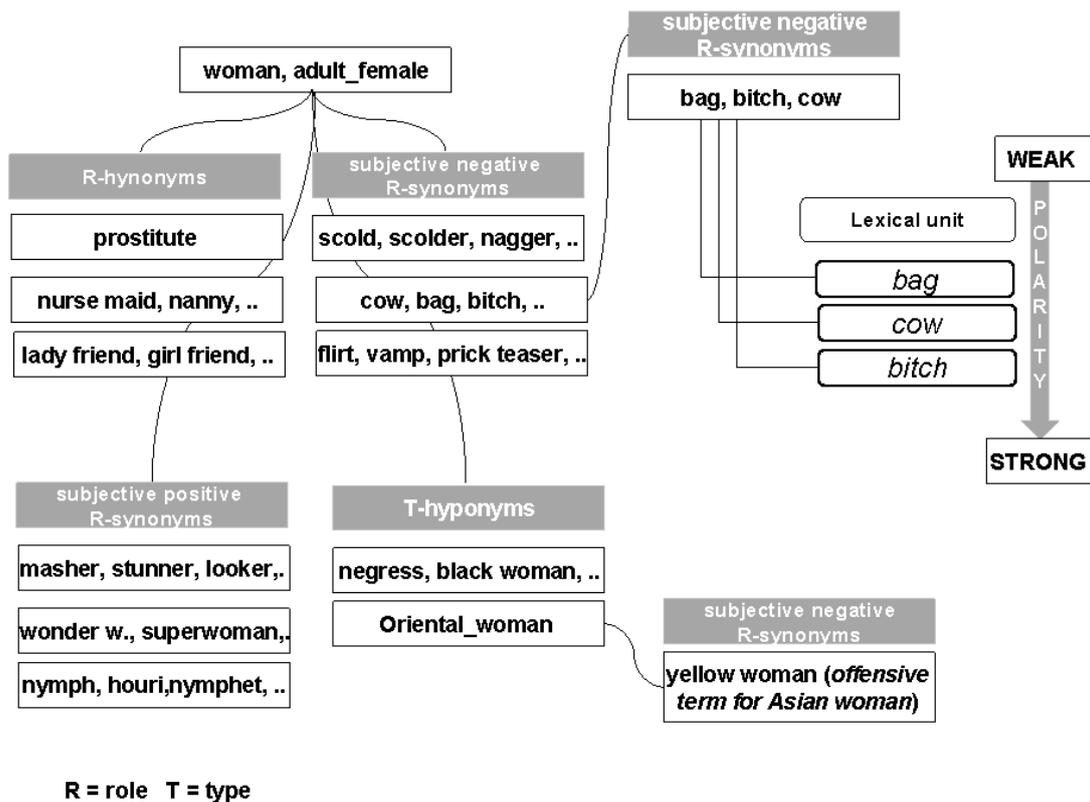


Figure 1: Attitudinal nouns in Wordnet

Likewise, with regard to verbs, the model allows us to make a distinction between words with a connotation and their neutral counterparts. Real troponyms, indicating an action performed in a particular manner or for a particular purpose, are distinguished from ‘subjective near-synonyms’. In our example (cf. figure 2), *gorge*, *stuff*, *devour*, *fress* etc. are not considered as manners of eating but they are interpreted as expressing the negative opinion of the speaker or writer about someone’s eating behavior.

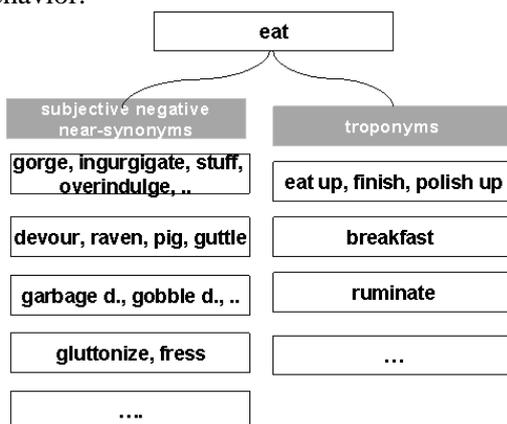


Figure 2: Attitudinal verbs in Wordnet

For adjectives the situation is different: it is not possible to form hierarchical networks of neutral adjectives and relate them to subjective counterparts as almost all adjectives do have some kind of negative or positive connotation. Adjective synsets, however, are often internally ambiguous with regard to polarity and would benefit greatly from the possibility to assign a scalar value for fine-grained polarity to each of its synonyms.

The combination of two resources offers the advantages of both. Applications can directly use the synset structure to determine the coarse-grained affective value of complete synsets. To use more fine-grained connotations they can access the lexical units individually. Moreover, the synset structures remain ‘intuitive’ with all neutral synonyms grouped together and with negative and positive synonyms grouped in distinct synsets. Furthermore, clearer criteria can be used to define synset-membership and less well-defined or agreed criteria can be used to encode fine-grained distinctions at the lexical unit level. This makes it

easier to match synsets across languages and to establish equivalence relations across wordnets. Finally, software for deriving a sentiment wordnet using page-rank and sentiment propagation will benefit from the homogeneous synsets and perform better.

7 Conclusions

We have shown that wordnet is not the appropriate structure for the description of Attitudinal Language. Especially if we want to describe fine-grained subjectivity and polarity, there are serious limitations caused by the possible internal ambiguity of synsets with regard to these features. To overcome these limitations we propose a modification and extension of the current wordnet architecture.

In future we will refine the guidelines for the new categories and perform annotation exercises to test and further improve the model.

8 Acknowledgments

This research has been carried out within the project From Text To Political Positions (<http://www2.let.vu.nl/oz/cltl/t2pp/>). It is funded by the VUA Interfaculty Research Institute CAMeRA.

References

- Andreevskaia, A. and S. Bergler (2006) Mining WordNet for Fuzzy Sentiment: Sentiment Tag Extraction from WordNet Glosses. In: EACL-2006, Trento, Italy.
- Argamon, S, K. Bloom, A. Esuli, F.o Sebastiani (2007) *Automatically Determining Attitude Type and Force for Sentiment Analysis*. In: Proceedings of the 3rd Language and Technology Conference, Poznan, Poland.
- Esuli, Andrea and Fabrizio Sebastiani. (2006). *SentiWordNet: A Publicly Available Lexical Resource for Opinion Mining*. In: Proceedings of LREC-2006, Genova, Italy.
- Guarino, N. (1998) *Some ontological principles for designing upper level lexical resources*. In Proceedings of LREC-1998, Granada, Spain.
- Hatzivassiloglou, Vasileios and Kathleen McKeown.1997. *Predicting the Semantic Orientation of Adjectives*.In: Proceedings of ACL'97, Madrid, Spain.
- Kamps, J., R. J. Mokken, M. Marx, and M. de Rijke (2004). *Using WordNet to measure semantic orientation of adjectives*. In: Proceedings LREC-2004, Paris.
- Martin, J. R. and P. R. R. White. 2005. *The Language of Evaluation: Appraisal in English*. Palgrave, London, UK. (<http://grammatics.com/appraisal/>).
- Miller, G. (1998) *Introduction*. In: WordNet: An Electronic Lexical Database (C. Fellbaum (ed.)). The MIT Press, Cambridge, Mass.
- Su, F.and K. Markert *Eliciting Subjectivity and Polarity Judgements on Word Senses*. In: Proceedings of Coling-2008, Manchester, UK.
- Wiebe, Janyce and Rada Micalcea.(2006) . *Word Sense and Subjectivity*. Proceedings of ACL'06, Sydney, Australia.
- Wiebe J., Theresa Wilson , Rebecca Bruce , Matthew Bell , and Melanie Martin (2004). *Learning subjective language*. Computational Linguistics 30 (3).
- Vossen, P., I.Maks, R. Segers and H. van der Vliet (2008). *Integrating Lexical Units, Synsets, and Ontology in the Cornetto Database*. In Proceedings of LREC-2008, Marrakech, Morocco, 2008
- Vossen, P., E. Agirre, N. Calzolari, C. Fellbaum, S. Hsieh, C. Huang, H. Isahara, K. Kanzaki, A. Marchetti, M. Monachini, F. Neri, R. Raffaelli, G. Rigau, M. Tescon, J. van Gent (2008). *KYOTO: A System for Mining, Structuring, and Distributing Knowledge Across Languages and Cultures* , in: Proceedings of GWC 2008, Szeged, Hungary.