Journal of Logic, Language and Information **13:** 159–171, 2004.

Semantic Opposition and WORDNET

SANDIWAY FONG
NEC Laboratories America, 4 Independence Way, Princeton, NJ 08540, U.S.A.;
Departments of Linguistics and Computer Science, University of Arizona, P.O. Box 210028, Tucson, AZ 85721, U.S.A.★
E-mail: sandiway@email.arizona.edu

(Received in final form 14 January 2004)

Abstract. We consider the problem of semantic opposition; in particular, the problem of determining adjective-verb opposition for transitive change of state verbs and adjectivally modified grammatical objects. Semantic opposition problems of this type are a sub-case of the classic Frame Problem; the well-known problem of knowing what is preserved or changed in the world as a result of some action or event. By definition, grammatical objects of change of state verbs undergo modification. In cases where the object is adjectivally modified, the problem reduces to determining whether the property denoted by the adjective still holds true after the event denoted by the verb. In this paper, we evaluate the efficacy of WORDNET, a network of concepts organized around linguistically relevant semantic relations including antonymy, for this task. Test examples are drawn from the linguistic literature. Results are analyzed in detail with a view towards providing feedback on the concept of a network as an appropriate model of semantic relations for problems in semantic inference.

Key words: Adjectives, change of state verbs, semantic inference, semantic networks, semantic opposition

1. Introduction

In the area of lexical semantics, there has been much interest in the class-based representation of verbs (Pustejovsky, 1995; Rappaport Hovav and Levin, 1998; Fillmore et al., 2001). A lexical semantic representation of verbs that predicts syntactic behavior has important implications not only for theories of human lexical knowledge but also for the large-scale construction of lexicons for natural language applications (Fong et al., 2001).

Lexical semantic knowledge about verbs can also be extended to apply to problems involving logical inference. The *Frame Problem*, defined in (1), is an important example of such a problem. The *Frame Axiom* in (2), from Kowalski (1979), represents the corresponding logical formulation. (2) states that a statement s holds in a state resulting from action a in state y if it already holds in state y to begin with and action a preserves the truth of statement s.

(1) *Frame Problem*: the problem of whether statements that hold true of a given state of the world continue to hold after some action has been performed.

★ Address for correspondence.

(2) *Frame Axiom*:
HOLDS(s,result(a,y)) $\leftarrow$ HOLDS(s,y), PRESERVES(a,s).

The problem of *persistence* for grammatical objects of the class of (transitive) change of state verbs in the sense of Pustejovsky (2000b) reduces to an instance of the Frame Problem. (Examples to follow immediately below.) Given lexical semantic and ontological knowledge about verbs and other grammatical elements such as adjectives, the central question to be explored in this paper is whether such knowledge can be used to evaluate if adjectival properties of a grammatical object still hold after an event denoted by a change of state verb applies to that object. More concretely, we will explore the issue of whether WORDNET (Fellbaum, 1998), a semantic network containing *inter alia* lexical and ontological relations about English verbs and adjectives, can be used to implement the Frame Axiom relation PRESERVES.

We consider minimal pairs of the form shown in (3); this example is adapted from Pustejovsky (2000a).

(3) a. Cathie mended the *torn* dress.

b. Cathie mended the *red* dress.

Change of state verbs like *mend* refer to some action resulting in a change of state for their direct objects. For example, in both (3a) and (3b), we can infer that as a result of the action of *mending*, the object *dress* becomes or has property *mended*. Using Rappaport Hovav and Levin 1998's notation, the event template for *mend* can be expressed as (4):

(4) [x CAUSE [BECOME [y <mended>]]],

where *x* and *y* denote the subject *Cathie* and the affected object *dress*, respectively. The sub-template BECOME [y <mended>] encodes the fact that *y* undergoes a change of state. An adjective modifying the object *y* expresses a property that is true at the onset of the event. The question is whether that property still holds true at the completion or culmination of the event.

Returning to (3a), to use Pustejovsky's terminology, *torn* and *mended* are in semantic opposition with respect to each other. That is, they cannot be simultaneously true for a given entity. From this, we can conclude *dress* is no longer in state *torn* as a result of the action of "mending." By contrast, if no semantic opposition obtains, there is no reason to assume any change in property status. In other words, there is no reason to conclude that *dress* changes color (from *red*) in (3b). Put another way, relative to the particular verb *mend*, *red* is a permanent property of *dress*, whereas *torn* is a defeasible one.

Pustejovsky (2000b) also discusses more general cases of semantic opposition not involving adjectival modification, such as that shown in (5a).

(5) a. The woman on the boat jumped into the water.

b. [BECOME$_{<\text{jumping}>}$ [IN y <the water>]]].

Jump here is a change of location verb. The corresponding event template is given in (5b). Here, *y*, the entity undergoing the change of location, represents the prepositional phrase (PP) modified noun phrase *the woman on the boat*. At this point, we draw an important distinction between the kind of semantic opposition found here, i.e., on the one hand, between PP modification (*on the boat*) and a change in location (*the water*), and, on the other hand, the adjective/change of state verb opposition introduced earlier in (3a). Unlike the corresponding situation in (3a), the PP modification in (5a) is not necessarily negated or opposed by the change of location implied by the verb. For example, the "boat" in question could be a cruise ship, and the "water" could be its on-board swimming pool. Deciding such cases involves real-world or situation-specific knowledge. In this paper, we restrict our attention to cases of inference involving only grammatical or lexical knowledge. In particular, we will make precise the nature of semantic opposition with respect to the network of synonym/antonym relations in WORDNET, and describe how WORDNET can be employed by an event semantics-aware parser to evaluate cases of adjective-verb opposition.

1.1. FURTHER EXAMPLES

Further examples of the paradigm from (Pustejovsky, 2000a) are given in (6) through (13).

(6) a. The plumber fixed every leaky faucet.

b. The plumber fixed every blue faucet.

(7) Mary fixed the flat tire.

(8) John mixed the powdered milk into the water.

(9) The father comforted the crying child.

(10) John painted the white house blue.

(11) Mary rescued the drowning man.

(12) Mary cleaned the dirty table.

(13) The waiter filled every empty glass with water.

An accurate algorithm must be able to discriminate between *leaky* and *blue* in (6a) and (6b), respectively, with respect to *fix*. Also, *flat* in (7) must be treated like *leaky* in the sense that both *leaky* and *flat* must be in semantic opposition to *fix*.

Semantic opposition is not limited to change of state verbs. Activity verbs such as *sweep*, *wipe*, *broom*, *paint* in (10) and *brush* have a change of state interpretation when they are modified by a resultative. For example, in (14), the cancellation of the state *dirty* is only implied in the latter case, i.e., the semantic opposition is between the resultative encoding the end state *clean* and *dirty*.

Relation	Description	Example
x HYP y	y is a hypernym of x	x:*repair*, y:*improve*
x ENT y	x entails y	x:*breathe*, y:*inhale*
x SIM y	(Adjectives) y similar to x	x:*achromatic*, y:*white*
x CS y	y is a cause of x	x:*anesthetize* , y:*sleep*
x VGP y	(Verbs) y similar to x	x:*behave* , y:*pretend*
x ANT y	x and y are antonyms x	x:*present* , y:*absent*
x SA y	x, see also y x	x:*breathe* , y:*breathe out*
x PPL y	x participle of y x	x:*applied* , y:*apply*
x PER y	x pertains to y x	x:*abaxial* , y:*axial*

Figure 1. WORDNET semantic relations.

(14) a. John brushed the dirty carpet.

b. John brushed the dirty carpet *clean*.

However, the same analysis does not extend to all verb classes. Verbs of creation and destruction are exceptions, as the examples in (15a) and (16a) show.

(15) a. Nero built the (gleaming) temple.

b. [x CAUSE [BECOME$_{<\text{built}>}$ [y EXIST(+)]]].

(16) a. Nero ruined/(destroyed) the (splendid) temple.

b. [x CAUSE [BECOME$_{<\text{ruined}>}$ [y EXIST(-)]]].

As the event templates in (15b) and (16b) indicate, by definition the "effected" object does not exist prior to and subsequent to the event for creation and destruction verbs, respectively. Hence, the solution to the Frame Problem for these cases is technically trivial in the sense it does not involve semantic opposition with respect to an individual verb.

2. WORDNET

WORDNET (version 1.6) organizes verb, nouns and adjectives into fairly distinct networks consisting of synonym set nodes called synsets.⋆ Synsets are linked via semantic relations such as hypernymy ("instance of") and antonymy. A list of

⋆ For example, the only direct link between the adjective and verb networks is the adjectival participle relation PPL which consists of a total of just 90 links, linking examples like *cooing* and *coo*. For this experiment, we use past and present participles derived from COMLEX to supplement these links (see Section 3).

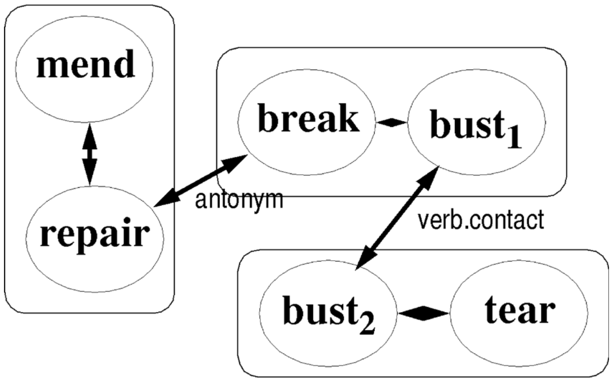

Figure 2. Shortest path between *mend* and *tear*.

semantic relations used in this experiment is given in Figure 1. Adjectives are organized using a flat structure consisting of (two) directly opposing antonyms, each with an accompanying cluster of similar adjectives (encoded by the SIM relation). Hence, these similar adjectives are indirect antonyms with respect to a central adjective (and the adjectives in its cluster). By exploiting these and other relations via transitivity, we can establish whether an appropriate semantic chain exists between a given adjective and verb.

Figure 2 illustrates the shortest path between *mend* and *tear*. *Mend* and *repair* belong to the same synset. *Repair* and *break* are antonyms. *Break* and one sense of *bust* are in the same synset. A second sense of *bust* and *tear* belong to another synset. The link between the two very similar senses of *bust* is made through a common lexicographers' source file, i.e., they are both classified as being verbs of contact. (The two senses being linked here are *bust* as in "ruin completely" and "separate or cause to separate abruptly.")

There are another five ways (all involving a longer chain) to get from *mend* to *tear* using WORDNET. Four of these also involve a single antonym relation consistent with the chain from Figure 2. However in general, it is reasonable to assume that the longer the chain, the less certain or reliable the information imparted by that chain. For example, the implausible-looking chain for *mend* to *tear* shown in Figure 3 has one dozen links. The path identified in Figure 3 is implausible due to the inclusion of intuitively irrelevant concepts (in relation to *mend* and *tear*) such as *touch*, *land* and *shoot down*. This example also illustrates an interaction between the (generally) polysemous nature of verbs and the (over)use of the lexicographers' file numbers proposed for Figure 2. For instance, *fix* (in line 2) as a verb of change has two distinct senses, namely (200177962,3) and (200339066,2), corresponding to dictionary senses *to restore* and *to make fixed, stable or stationary*, respectively.⋆ Nevertheless, these two senses of *fix* should not be conflated despite the fact that

⋆ An identifier like (200177962,3) in the case of *fix* is a WORDNET-internal (unique) identifier giving the synset number (200177962) and offset (3) within the referenced synset for that particular sense of *fix*.

1. *mend* and *fix* in same synset
2. *fix* (200177962,3) and *fix* (200339066,2) in synsets related by verb.change
3. *fix* instance of *attach*
4. *attach* (200885494,1) and *attach* (200881541,1) in synsets related by verb.contact
5. *attach* instance of *touch*
6. *touch* (200820504,1) and *touch* (200820743,1) in synsets related by verb.contact
7. *touch* see also *touch down*
8. *touch down* instance of *land*
9. *land* (201348563,1) and *land* (201349748,3) in synsets related by verb.motion
10. *land* and *shoot down* in same synset
11. *shoot down* (201349748,2) and *shoot down* (201405541,3) in synsets related by verb.motion
12. *shoot down* and *tear* in same synset

Figure 3. Antonym-less path from *mend* to *tear*.

they are both verbs of change. Finally, the chain of inference does not count as an instance of semantic oppostion since no antonym relation is involved. Obviously, this chain should be rejected in favor of the much shorter path shown in Figure 2. In this paper, we consider shortest paths only as candidates for plausible chains of inferences.

3. Implementation

A PROLOG-based sentence parser was used to analyze the Pustejovsky examples given in (3) through (13). The parser computes both a parse tree based on event structure and the shortest WORDNET path for relevant configurations. An example of the output for *Mary rescued the drowning man*, example (11), is given in Figure 4. As explained in Section 1, an appropriate configuration will contain an NP object modified by an adjective in the context of a change of state verb (or an activity verb modified by a resultative). If a single antonym link is present in the shortest path, the modifying adjective is marked with the feature `cancelled` to signify semantic opposition, i.e., the property no longer holds true on event completion. In the case where no antonym link is found or the verb does not allow a change of state reading, the property represented by the adjective remains uncancelled. We briefly describe the relevant components of the parser below:

– The PROLOG version of the WORDNET verb/adjective system was employed. This is a network of approximately 174 K nodes and 600 K links. Breadth-first search, using a hybrid PROLOG/C program, was used to compute shortest chains.*

* The WORDNET databases are stored in PROLOG for flexibility and accessed via (efficient) first-argument indexing during searches. The breadth-first tree is implemented in C for efficiency and to

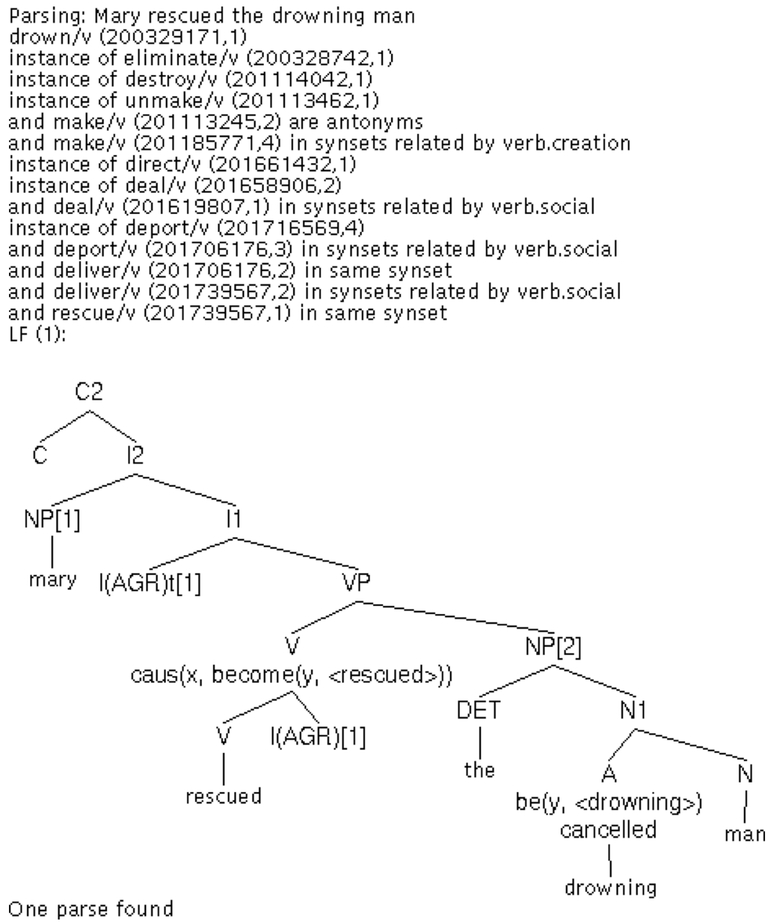

Figure 4. Mary rescued the drowning man.

– The broad-coverage lexicon used for the experiments was formed by combining parts of COMLEX(Grishman et al., 1994), with an event semantics verb lexicon derived from Levin (1993). The adjective subsystem is composed of 6.2 K basic adjectives from COMLEX (WORDNET has about 20 K entries) plus approximately 9.5 K past and present participle verb forms.

avoid triggering PROLOG garbage collection for reclaiming search tree storage. On a 200 MHz Sun Sparc, the search rate is approximately 26 K nodes/sec. As Figure 5 attests, finding the shortest chain takes no more than about a second even for the largest search.

Candidate Pair	Shortest Chain (No. of links)	Semantic Opposition	Size of Search (No. of Nodes)
mend-torn	5	Yes	1261
mend-red	–	No	11974
fix-leaky	5	Yes	12167
fix-blue	11	No	14553
fix-flat	–	No*	12286
mix-powdered	6	Yes	11931
comfort-crying	9	Yes	11359
blue-white	–	No*	24431
rescue-drowning	13	Yes	9142
clean-dirty	1	Yes	61
fill-empty	1	Yes	48

Figure 5. Shortest path results for WORDNET's verb/adjective system.

4. Results and Conclusions

The larger notion being tested in this paper is the general transitivity of semantic relations. A pertinent question to ask then is: What are the limits to transitivity, especically when applied to heterogeneous relations? Empirically, how many relations can be chained together before reliability is compromised and we end up with degraded or unwanted inferences? Can some sort of thresholding be applied to this problem? Also, are there peculiarities in the organization of WORDNET concepts that affect the results? In the following sections, we report experimental results on cases of adjective-verb opposition, and probe the underlying organization of WORDNET in detail.

4.1. RESULTS

The table in Figure 5 reports the results for WORDNET's adjective/verb system on examples (3), and (6) through (13). The system operates as a simple decider for semantic opposition, reporting "Yes" if the shortest chain contains a single antonym link, and "No" otherwise. Incorrect outcomes are marked in the table with an asterisk (*). An example of such an outcome is in the case of *fix* and *flat*, where the system fails to report a chain between these two concepts. Note that in the case of *mend* and *red*, the system also fails to detect a connection, but since this is the expected outcome, the result is not marked. Clearly, the shortest path algorithm appears to perform well, producing the correct value for the semantic opposition feature in 9 out of the 11 cases.

In the current implementation, the breadth-first search algorithm employed is a simple unidirectional one. Since both the source and goal nodes are known,

1. *blue* instance of *discolor*
2. *discolor* instance of *change*
3. *change* and *leave* are antonyms
4. *leave* see also *leave_out* (200416229,5)
5. *leave_out* (200416229,5) and *leave_out* (200416622,3) in synsets related by verb.cognition
6. *leave_out* (200416622,3) instance of *eliminate*
7. *eliminate* instance of *destroy*
8. *destroy* instance of *unmake*
9. *unmake* and *make* (201113245,2) are antonyms
10. *make* (201113245,2) and *make* (201142893,4) in synsets related by verb.creation
11. *make* (201142893,4) and fix in same synset

Figure 6. Shortest path from *blue* to *fix*.

substituting a bi-directional variant will probably halve the depth of the search space on average and provide for a concomitant gain in efficiency.

4.2. ANALYSIS

The results shown in Figure 5 are analyzed in detail in the following subsections. In particular, we discuss the shortest path heuristic as well as explaining why the program produces the wrong values for the two marked (*) examples, namely *fix-flat* and *blue-white*.

4.2.1. *The Value of Thresholding*

Currently, no upper limit is imposed on the length of the shortest chain. The reason is that a fixed length threshold is difficult to establish. For example, *fix* and *blue* are connected by an 11-link chain, as shown in Figure 6. Note that this long chain contains two antonym relations and thus does not count towards establishing semantic opposition between *fix* and *blue*.⋆ This spurious 11-link chain could be eliminated by a sub-11-link threshold. However, this would also eliminate positive data. For example, the program crucially relies on a 13 link antonym chain connecting *rescue* and *drowning* to establish semantic opposition.

4.2.2. *The Shortest Path Heuristic*

The shortest path heuristic may fail to find an appropriate antonym chain if a shorter non-antonym chain exists. For example, this situation obtains with *fix-flat*, example (7).

⋆ In other words, we make the assumption that two antonym relations in series cancel each other out. Theoretically, we can also generalize this to the assumption that a (non-trivial) odd number of antonym relations should be equivalent to a single antonym. However, in our experiments, none of the shortest paths require three or more antonyms.

1. *deflate* instance of *collapse*
2. *collapse* instance of *fold*
3. *fold* (200872449,1) and *fold* (201042123,1) in synsets related by verb.contact
4. *fold* instance of *lace*
5. *lace* (201042567,6) and *lace* (201045661,1) in synsets related by verb.contact
6. *lace* instance of *tie*
7. *tie* instance of *fix*

Figure 7. Shortest path from *deflate* to *fix*.

1. *deflate* instance of *cut* (200298808,7)
2. *cut* (200298808,7) and *cut* (200329545,1) in synsets related by verb.change
3. *cut* (200329545,1) instance of *eliminate*
4. *eliminate* instance of *destroy*
5. *destroy* instance of *unmake*
6. *unmake* and *make* (201113245,2) are antonyms
7. *make* (201113245,2) and *make* (201142893,4) in synsets related by verb.creation
8. *make* (201142893,4) and *fix* in same synset

Figure 8. 8-link path from *deflate* to *fix*.

There are 18 senses of *flat* as an adjective, one of them directly referring to a flat tire. *Flat* and *deflated* are clustered adjectives. *Fix* and *deflate* can be connected by the 7-link shortest chain shown in Figure 7. This chain lacks an antonym link. Furthermore, WORDNET is missing a link here between *deflated* as an adjective and *deflate* as a verb due to the incompleteness of the *adjective participle of verb* relation (PPL). However, even assuming the existence of this link, our algorithm will still fail to find any evidence of semantic opposition.

Note that there exists a chain, shown in Figure 8, containing the required single antonym relation between *deflate* and *fix*; the only problem being that it contains 8 links, and therefore it is blocked by the 7-link (non-antonym) shortest chain reported in Figure 7.

4.2.3. *The Color System in* WORDNET

Finally, the algorithm also fails to discover an antonym chain between *blue* and *white*, example (10), revealing pecularities about the organization of color terms. In WORDNET, there is no direct or indirect antonym relation between colors.

Colors such as *white* and *blue* are not considered to be traditional binary opposites. Rather, these two colors are contraries.* However, WORDNET does not make a principled distinction between contradictions and contraries. For example, it contains direct antonym links between scalar antonyms such as *rich* and *poor*, and and *early* and *late*, directional opposites such as *left* and *right*, as well for polar

* In traditional logic, a distinction is made between contradictions where one of the two items must hold true, and contraries where both items may be simultaneously false.

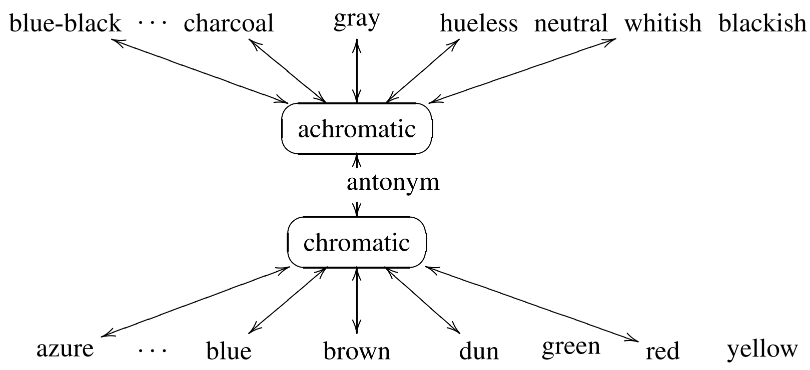

Figure 9. The color system in WORDNET.

1. *white* and *achromatic* (300367747,2) in same synset
2. *achromatic* (300364634,1) and *chromatic* (300355823,1) are antonyms
3. *chromatic* (300355823,1) and *blue* are similar

Figure 10. White and *blue*.

opposites such as *present* and *absent*. One might conclude that, in the case of the color system, WORDNET is simply incomplete.

However, a closer look at the organization of the color system reveals that an antonym chain for *blue* and *white* should have been possible in principle. WORDNET divides color adjectives into two clusters headed by the direct antonyms *achromatic* and *chromatic*, as shown in Figure 9.

White and *achromatic* belong to the same synset, and *blue* is similar or a satellite of *chromatic*. Hence, *white* and *blue* are also indirect antonyms. However, as Figure 10 shows, the algorithm fails to find a connection between these two colors because the two colors refer to two unconnected senses of *achromatic*.⋆ Of course, even if there is a link between the two senses of *achromatic*, the color model is still flawed with respect to semantic opposition. For instance, *red* and *blue* are both chromatic, and *white* and *grey* are both achromatic. This implies (incorrectly) that semantic opposition will not obtain in the case of examples like *John painted the red door blue* and *Mary painted the white tiles grey*.

A possible solution would be ignore the encoded chromatic/achromatic distinction, i.e., the antonymn link, altogether, and simply state that, to a first approximation, that all colors are contraries and therefore distinct. (Difficulties remain, for example, with *hueless* and *neutral*.) However, this solution is domain-specific and resists generalization as we shall see presently.

More abstractly, we can take advantage of the clustering given by WORDNET, and state that adjectives belonging to the same "semantic dimension" or common

⋆ *Achromatic* sense (300364634,1) refers to "having no hue" whereas sense (300367747,2) refers to "being of the achromatic color of maximum lightness; having little or not hue owing to reflection of almost all incident light."

Candidate Pair	Shortest Chain (No. of links)	Semantic Opposition	
		WORDNET Antonyms	Clustering
lengthen-short	1	Yes	Yes
shorten-short	0	No	No
warm-tepid	1	No	Yes
cool-tepid	2	Yes	Yes

Figure 11. Antonyms versus clustering.

"archisememe/archilexeme" in the sense of Mettinger (1994) (and cites therein) lie in systematic semantic opposition.⋆ Call this the *Clustering Model*. For our purposes, it suffices to state that the archilexeme should be COLOR and the dimension or scale WAVELENGTH. Extending this line of reasoning to adjective-verb combinations along other dimensions such as LENGTH or TEMPERATURE makes different predictions about semantic opposition as compared to the antonymy-based algorithm employed so far. Consider the examples in (17) and (18):

(17) a. John lengthened the short rope. (opposition)

b. John shortened the short rope.

(18) a. Mary warmed the tepid water. (opposition)

b. Mary cooled the tepid water.

The opposition judgements in (17) and (18) are based on the assumption that verbs like *lengthen/shorten* and *warm/cool* (may) have natural interpretions as change of state verbs. The table in Figure 11 compares the antonym-based approach to the clustering model. Summarizing, semantic opposition obtains for the antonym-based approach in two cases, *lengthen/short* and *cool/tepid* because *long/short* and *warm/cool* are direct antonymns in WORDNET. Adjectives *tepid* and *warm* are similar to one another with respect to temperature in much the same way as *neutral* and *achromatic* are with respect to color. By contrast, the clustering model predicts opposition for all distinct terms in each domain. However, neither model matches the judgements given previously in (17) and (18) because they do not take into account other factors such as gradability and direction. In other words, WORDNET's bipolar model of adjectives needs to be revised before such examples can be properly handled.

⋆ As Mettinger observes, in the case of nouns (unlike adjectives), these two concepts are quite distinct. For example, in the case of *boy* and *girl*, an appropriate dimension and archilexeme might be GENDER and CHILD, respectively.

4.3. CONCLUSIONS

To summarize, we have described how WORDNET can be used as part of a decider for the semantic opposition task. An experimental system was constructed and successfully tested on the Pustejovsky examples. All three of the difficulties discussed above reveal potential shortcomings in WORDNET's design for semantic inference. Thresholding and the shortest path heuristic can only work if the length of a chain can be inversely correlated with reliability. Furthermore, the case of the *blue*/*white* example also illustrates how WORDNET's achromatic/chromatic divide can affect semantic reasoning.

Well-defined linguistic tests such as semantic opposition can be useful empirical devices in the evaluation of the quality of WORDNET's synset links. By chaining together semantic relations, we also probe the limits on general transitivity of such relations, and, indeed, on the applicability of a network-style organization of semantic relations for semantic inference tasks.

Acknowledgements

The author is indebted to Christiane Fellbaum and an ICoS-3 reviewer for detailed comments.

References

Fellbaum, C., ed., 1998, *WordNet*, Cambridge, MA: MIT Press.

Fillmore, C.J., Wooters, C., and Baker, C.F., 2001, "Building a large lexical database which provides deep semantics," pp. 3–25 in *Proceedings of the 15th Pacific Asia Conference on Language, Information and Computation (PACLIC-15)*, B. Tsou and O. Kwong, eds.

Fong, S., Fellbaum, C., and Lebeaux, D., 2001, "Ghosts, shadows and resultatives: The lexical representation of verbs," *TAL* **42**, 755–789.

Grishman, R., Macleod, C., and Meyers, A., 1994, "Comlex syntax: Building a computational lexicon," pp. 268–272 in *COLING '94*, Kyoto, Japan.

Kowalski, R., 1979, *Logic for Problem Solving*, Amsterdam: North-Holland.

Levin, B., 1993, *English Verb Classes and Alternations: A Preliminary Investigation*, Chicago, IL: University of Chicago Press.

Mettinger, A., 1994, *Aspects of Semantic Opposition in English*, Oxford: Oxford University Press.

Pustejovsky, J., 1995, *The Generative Lexicon*, Cambridge, MA: MIT Press.

Pustejovsky, J., 2000a, "Event-based models of change and persistence in language," IRCS Colloquium Series handout, University of Pennsylvania.

Pustejovsky, J., 2000b, "Events and the semantics of opposition," pp. 445–482 in *Events as Grammatical Objects*, C. Tenny and J. Pustejovsky, eds., Stanford, CA: CSLI Publications.

Rappaport Hovav, M. and Levin, B., 1998, "Building verb meanings," pp. 97–134 in *The Projection of Arguments: Lexical and Compositional Factors*, Butt and Geuder, eds., CSLI Lecture Notes No. 83, Stanford, CA: CSLI Publications.