Invented Antonyms: Esperanto as a Semantic Lab[*]

Andreas van Cranenburgh, **Galit W. Sassoon**, and **Raquel Fernández**

Institute for Logic, Language & Computation
University of Amsterdam

Abstract

This paper uses Esperanto—a constructed language with transparent morphology but rich semantic-pragmatic components—to study antonymy and polarity. We investigate the distribution of the Esperanto antonymy morpheme '*mal-*' (as in, for instance, '*mal-alta*': *antonym-tall*, *short*) in a 4.3 million-word corpus, Tekstaro, and use it as an empirical basis to assess different theories of negative antonyms. Our methodology consists in investigating the extent to which the antonymy morpheme '*mal-*', which we take to denote negative polarity, bears the linguistic features predicted by traditional linguistic tests (such as incompatibility with measure and ratio phrases and low likelihood of nominalisation).

1   Introduction

It is widely accepted that antonymous gradable adjectives, such as '*tall*'-'*short*', '*strong*'-'*weak*', and '*big*'-'*small*', have different semantic properties. The common assumption is that one of the adjectives within an antonym pair is unmarked or *positive* (e.g. '*tall*', '*strong*', '*big*') while its antonym is taken to be marked or *negative* (e.g. '*short*', '*weak*', '*small*'). Although defining the set of negative as opposed to positive adjectives is notoriously difficult, it is intuitively clear that such a distinction exists. Evidence in this direction comes from different sources. For instance, psycholinguistic studies have shown that negative antonyms are acquired later than their positive counterparts, arguably due to their semantic complexity (Klatzky et al. 1973, Giora 2006). Several linguistic studies (Lehrer 1985, Horn 1989, Bierwisch 1989, Kennedy 2001, Sassoon 2010) have pointed to different properties that distinguish positive versus negative antonyms: negative antonyms are considered less frequent than positive ones, incompatible with measure (1a) and ratio (1b) phrases, and much less felicitous in nominalisations (1c).

(1)   a. Two meters tall / [#]Two meters short
      b. Twice as tall / [#]Twice as short
      c. Height / [#]Shortness

However, these distributional tendencies admit a variety of "exceptions" or non-paradigmatic cases, including many positive adjectives that do not license measure phrases ([#]*two degrees warm/cold*) and rarely also negative adjectives that do ('*two hours late/early*'). In addition,

many positive adjectives resemble negative ones in rarely licensing '*twice*' (e.g. '*glad*', '*wise*') and a few notorious exceptional negative adjectives ('*bad*', '*slow*') seem to license '*twice*' more often than their positive bases (Sassoon 2010).

A key problem underlying the investigation of the linguistic factors that characterise antonyms with different polarity is the fact that in English there are usually no overt morphological traces that allow us to make a clear-cut distinction between negative and positive gradable antonyms. In this paper, we use Esperanto—a constructed language with transparent morphology that features an antonymy morpheme ('*mal-*')—to examine the extent to which some of the linguistic tests put forward in the literature are actually indicative of semantic polarity. In particular, we concentrate on tests that link polarity with measure phrases, ratio modification, and nominalisations.

Section 2 introduces the relevant grammatical features of Esperanto and discusses the properties of the antonymy morpheme '*mal-*'. In Section 3, we briefly describe two proposals that aim at accounting for the different distributional patterns of antonymous gradable adjectives. After that, in Section 4, we present our study, carried out using a corpus of Esperanto, *Tekstaro* (Wennergren 2003-). We examine to what extent the linguistic tests allow us to distinguish negative from positive antonyms and to assess the validity of the predictions made by the two proposals introduced in Section 3.

2   Esperanto and Antonymy

Esperanto is a constructed language developed by Ludovic Lazarus Zamenhof in the 1880s (Zamenhof 1887). Although it was designed as an easy-to-learn language, with regular and transparent syntax and morphology, its semantic and pragmatic components have evolved naturally: Since its origin, Esperanto has enjoyed a continuous history of use, with an average of 200 new Esperanto books published per year and a lively speech community with an estimated 2 million speakers (Lewis 2009), including up to 2000 native speakers. Esperanto, thus, provides a special opportunity to study study the interface between morphology and interpretation.

In Esperanto, the lexicon consists of a stock of roots, derived from Romance and Germanic languages, that combine with affixes. Parts of speech are morphologically marked, with all adjectives ending with the suffix '*-a*'. Of special importance for our purposes is the fact that in Esperanto the vast majority of antonyms are formed using the very productive prefix '*mal-*'. For instance:

(2) a. sana (*healthy*) / malsana (*sick*)
     b. alta (*tall*) / malalta (*short*)

With very few exceptions, such as '*diferenca*' (*different*) and '*basa*' (*low*), antonyms are derived by the regular morphological process exemplified in (2), which is extremely productive. We can determine the productivity of an affix by looking at the frequency spectrum of a class of words (Baayen and Lieber 1996). Highly productive affixes combine with a large variety of words thus generating new words that have low frequency, while unproductive affixes combine only with a fixed set of words that are frequent enough to be memorised as correct by speakers. According to Baayen and colleagues, a *productivity index* corresponding to the rate at which new word types are expected when further tokens are sampled can be calculated by dividing the number of "hapax legomena" (word types which occur only once) of a given word formation process
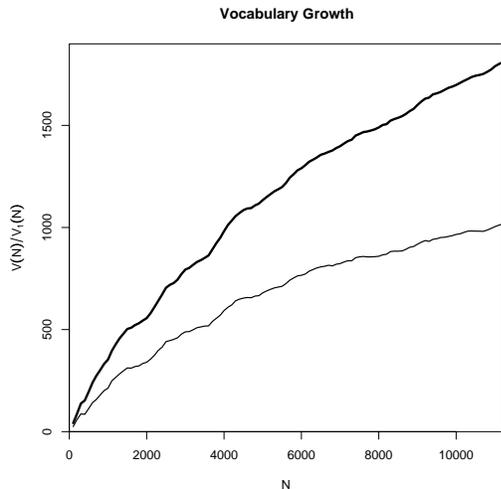
Vocabulary Growth

Figure 1: Empirical growth curve of '*mal-*'words. The thin line is the number of hapax legomena. X-axis: number of '*mal-*'word tokens considered. Y-axis: number of types encountered so far.

by the total number of its tokens. Hay and Baayen (2002) report a 0.005 productivity index for the English prefix '*un-*'. In contrast, the productivity index of the '*mal-*' prefix is 0.092 (1017/11025). This is exceedingly high, predicting almost 1 new '*mal-*'word type for every 10 '*mal-*'word tokens that are sampled. The graph in Figure 1 shows the empirical vocabulary growth curve for *mal*-words (thick line) together with the number of hapax legomena (thin line). As can be seen, the growth curve is still very steep in the end, which predicts that the number of types will keep rising steadily as more tokens are considered.

The regular morphological construction of antonyms in Esperanto seems in line with Heim's syntactic negation theory of antonymy (Heim 2008) according to which antonyms are not specified as independent entries in the lexicon but generated by a predicate negation operator, *little*, hidden away in the logical form of negative adjectives such as '*short*'. In Esperanto, we instead find the explicit antonym prefix '*mal-*' at the surface level.

We shall thus assume that when '*mal-*' combines with a gradable adjective it acts as a surface-level indicator of negative polarity, i.e. it reverses the scale of the positive base it combines with. With this assumption in place, we can straightforwardly investigate to what extend the distributional tests found in the literature give results that are indicative of negative antonymy.

3   Two Views on Antonymy

Before moving on to describe our study, we consider two competing proposals that explain some of the distributional patterns of antonymous gradable adjectives by appealing to different underlying mechanisms.

The Polarity Hypothesis.   According to Kennedy (2001), the distributional differences of gradable adjectives are due to their semantic *polarity*: Positive adjectives map their arguments onto positive degrees corresponding to closed (or bound) intervals, while negative adjectives map their arguments onto negative degrees corresponding to open (or unbound) intervals. Since measure and ration phrases correspond to closed intervals, this hypothesis predicts that such phrases will be incompatible with negative adjectives.

The Additivity Hypothesis. According to this hypothesis, put forward by Sassoon (2010), the distributional differences of gradable adjectives are due to the *additivity* of the kind of mapping they denote, which is orthogonal to their semantic polarity. Additive and non-additive mappings have the following characteristics:

(3) a. Additive mapping:
$$f_{\text{tall},t}(x) = f_{\text{tall},t}(x \oplus x)/2$$
b. Non-additive mapping:
$$f_{\text{tall},t}(x) - 1 \neq (f_{\text{tall},t}(x \oplus x) - 1)/2$$

From this characterisation it follows that non-additive adjectives are incompatible with measure and ratio phrases. Negative polarity antonyms are predominantly non-additive while positive polarity ones can be either. This hypothesis thus predicts that almost all negative polarity antonyms but also many positive ones will be infelicitous with measure and ratio phrases.

4 Our Study: Investigating Antonym Morphology

For our study, we use Tekstaro (Wennergren 2003-),[1] a 4.3 million-word corpus of Esperanto that includes translated and original literature as well as magazine articles. We examine the extent to which linguistic tests that link polarity with measure phrases, ratio modification, and nominalisations are in line with the distributional properties of the antonymy morpheme '*mal-*' and hence are actually indicative of semantic polarity.

4.1 Antonym Frequency

The total number of adjectives in the corpus is 343,120. Of these, 21,363 are '*mal-*'antonyms. The ratio of '*mal-*'antonyms to all adjectives is therefore **6.23%**. This will be the expected ratio against which we shall evaluate the results obtained with the linguistic tests that we describe in the following sections.

Given the 6.23% figure, it is clear that '*mal-*'antonyms are substantially less frequent than unmarked adjectives. The lower frequency of '*mal-*'antonyms supports our assumption that '*mal-*' is an indicator of negative polarity since low frequency is considered a side effect of negativity (Lehrer 1985).

4.2 Measure and Ratio Phrases

According to the polarity hypothesis, incompatibility with measure and ratio phrases identifies negative antonyms. Assuming that '*mal-*' is an indicator of negative polarity, this predicts that the number of '*mal-*'antonyms that appear with measure and ratio phrases will be significantly lower than the number of non-'*mal-*' adjectives appearing with these phrases. In contrast, the additivity hypothesis states that the measure and ration phrase test identifies non-additive adjectives (which are mostly negative but can also be positive). This predicts that there will not be a statistically significant difference between the number of non-'*mal-*' and '*mal-*'antonyms appearing with measure and ratio phrases.

---

[1] www.tekstaro.com

4.2.1   Test 1: Incompatibility with measure phrases

Measure phrases in Esperanto can occur with adjectives, nouns, adverbs, and verbs. The corpus contains 276 occurrences of '*metro(j)*' (*meter(s)*) preceded by a number or a quantity. Of these measure phrases, 41 are combined with an adjective, 23 with a noun, 11 with an adverb, and 5 with a verb. One example of each type is shown in (4):

(4) a. **adjective:** ... proksimume dek centimetrojn <u>longa</u>
         approximately 10 centimeters long
     b. **noun:** ... havas ordinare la <u>altecon</u> de 3-4 metroj
         has ordinarily the height of 3-4 meters
     c. **adverb:** Mi min klinis, kaj vidis lin, du metrojn <u>sube</u> de mi
         I bowed down, and saw him, two meters below of me
     d. **verb:** mi eltrovis, ke la tegmento <u>altiĝas</u> 12 metrojn au plu
         I found out that the roof becomes-the-height of 12 meters or more

The overwhelming majority of measure phrases appear with positive (non'*mal-*') words. Overall, we find only 4 instances of measure phrases with '*mal-*'words, but none of them are adjectives: '*malleviĝas*' in (5a) is a verb, and '*malsupren*' in (5b) and '*malproksime*' in (5b) and (6b) are adverbs.

(5) a. Kiam estas refluo kaj la akvo malleviĝas per du metroj a pli
         When there is low tide and the water falls with two meters or more
     b. Ni ne povis vidi pli malproksime ol unu metron
         We couldn't see farther than one meter

(6) a. mi falis tre rapide tridek metrojn malsupren
         I fell very rapidly thirty meters downwards
     b. ... milojn da kilometroj malproksime
         thousands of kilometers away

From the 4 examples in (5) and (6) only those in (6) appear to be actual instances of a measure phrase being applied to a '*mal-*'word. The '*mal-*'words in (5), in contrast, do not directly combine with the measure phrase, which appears within a PP argument of a verb (5a) or a comparative clause (5b).[2] Thus, only the instances in (6) can be considered true exceptions[3] of what generally seems to be a strong incompatibility of measure phrases with '*mal-*'words. This is extreme in the case of '*mal-*'adjectives, of which we do not find any instance at all. Given the expected frequency of '*mal-*'antonyms (6.23%), we would expect to find at least 2 instances out of 41 adjectival occurrences.

Thus, test 1 supports the polarity hypothesis, confirming that the incompatibility of adjectives with measure phrases is indicative of negative polarity.

---

[2]Note that this is also possible in English.

[3]It is interesting to note that the adverb '*proksime*' (*far* or *away*) is atypical: there are almost the same number of instances of '*proksime*' and '*malproksime*', while acriss the board '*mal-*'words are much less frequent than their unmarked counterparts.

Figure 2: Left: frequency of '-*oble pli*' with positive vs. negative antonyms. Right: frequency of '-*oble pli*' with comparatives of positive vs. negative antonyms

4.2.2 Test 2: Incompatibility with ratio phrases

In Esperanto, "x times as ADJ as" is expressed with a numeral suffixed with '-*oble*' (*times*) followed by '*pli*' (*more*), as in (7):

(7) Dudekoble pli granda ol
     twenty times as big as

The corpus contains 88 matches for the construction "NUM-*oble pli* ADJ". Of these, 5 occur with a '*mal-*'antonym. For each adjective A occurring in this construction, we compared the following two figures:

(8) a. the frequency of '-*oble pli* A' given the frequency of A
     b. the frequency of '-*oble pli mal-A*' given the frequency of '*mal-A*'

A paired t-test showed that the difference between (8a) and (8b) is not statistically significant ($p = 0.1$). Thus, ratio phrases do not occur significantly less often with '*mal-*'antonyms than with their unmarked positive antonyms, which goes against the polarity hypothesis.[4] These results are consistent with the additivity hypothesis. The graphs in Figure 2 show that the unmarked positive adjectives of the 5 occurrences of '*mal*'-antonyms that appear with ratio phrases (circled in red) are "well-behaved" negative adjectives (with very low frequency of ratio modification). It is their positive antonyms that resemble negative ones in not licensing ratio phrases, as expected by the additivity hypothesis but not by the polarity hypothesis.

We can thus conclude that, contra the polarity hypothesis, the low frequency of ratio modification of negative adjectives is not indicative of polarity *per se*. Rather, our results are

---

[4]In order to assess the validity of the additivity hypothesis, we also compared the following:

(i)   a. the frequency of '-*oble pli* A' given the frequency of '*pli* A'
      b. the frequency of '-*oble pli mal-A*' given the frequency of '*pli mal-A*'

(ii) a. the frequency of '-*pli* A' given the frequency of 'A'
      b. the frequency of '-*pli mal-A*' given the frequency of '*mal-A*'

Both the difference between (ia) and (ib) and the difference between (iia) and (iib) are statistically significant ($p = 0.03$ and $p = 0.04$, respectively). This indicates that ratio phrases modify significantly more often comparatives of bare adjectives than comparatives of their '*mal-*'antonyms, while comparatives modify more often '*mal-*'antonyms than their positive counterparts.

6

consistent with the additivity hypothesis, according to which negative adjectives are predominantly non-additive (i.e. less likely to label an additive scale than a non-additive one), and positive adjectives, though not reversed, often resemble negative ones in being predominantly non-additive.

4.3   Test 3: Nominalisations

In Esperanto, two suffixes can be used to generate nouns: '-o' and '-eco', the latter being compatible with more abstract qualities (similarly to the English suffixes '-ness' and '-ity'). For instance:

(9) a. grand-a (*big*) → grand-o (*size*), grand-eco (*size, greatness*)
     b. long-a (*long*) → long-o (*length*), long-eco (*length, longness*)
     c. hom-a (*human*) → hom-o (*human being*), hom-eco (*humanity*)

We found that 107 adjectival roots are nominalised with '-o' and '-eco'. In both cases, the frequency of the nominalisation with the unmarked positive adjective form is not significantly different from the frequency of the nominalisation with the 'mal-'antonym ($p = 0.23$ for '-o' and $p = 0.18$ for '-eco').

Thus, in contrast to what is typically assumed, nominalisations are not significantly less frequent with negative adjectives. However, interesting results are obtained when we consider the likelihood of using the abstract suffix '-eco' when a nominalisation occurs. In particular, for each adjective A, we compared the following:

(10) a. the frequency of the abstractness marker 'A-eco' given the total number of nominalisations ('A-o' + 'A-eco')
      b. the frequency of the abstractness marker with the corresponding 'mal-'antonym 'mal-A-eco' given the total number of nominalisations with that 'mal-'antonym ('mal-A-o' + 'mal-A-eco')

A paired t-test showed that the difference between (10a) and (10b) is statistically significant ($p = 0.01$), indicating that abstract nominalisations with '-eco' are significantly less frequent with 'mal-'antonyms than with their unmarked positive counterparts. This opens the door to using abstract nominalisation as a new linguistic test for identifying semantic polarity.

5   Conclusions

In this paper we have used Esperanto to study antonymy and polarity. Since Esperanto has a regular and transparent morphology while featuring rich semantics and pragmatics, it offers the possibility of investigating the interface between morphology and interpretation. Our study of the antonymy morpheme 'mal-' shows that the restricted distribution of negative antonyms cannot be due to their polarity *per se* (that is, to their reversed scale), since a similar ditribution characterises also many positive antonyms. Rather, non-additivity (which prevents representation of degree ratios) decreases the frequency of use of measure phrases and ratio modification. We have also seen that the assumption that nominalisations of positive polarity adjectives are more frequent than nominalisations of negative antonyms appears not to be empirically

grounded. We have, however, identified a new "abstractness" test that does appear to account for the positive vs. negative distinction, indicating that abstract nominalisation morphemes have a preference for positive poilarity items.

References

R. Baayen and R. Lieber. Word frequency distributions and lexical semantics. *Computers and the Humanities*, 30(4):281–291, 1996.

M. Bierwisch. The semantics of gradation. In *Dimensional Adjectives: Grammatical Structure and Conceptual Interpretation*. Springer-Verlag, 1989.

R. Giora. Is negation unique? On the processes and products of phrasal negation. *Journal of Pragmatics*, 38(7):979–980, 2006.

J. Hay and R. Baayen. Parsing and Productivity. In *Yearbook of Morphology*, volume 2001, pages 203–235. Kluwer Academic Publishers, 2002.

I. Heim. Decomposing antonyms? In *Proceedings of Sinn und Bedeutung*, volume 12, pages 212–225, 2008.

L. Horn. *A natural history of negation*. University of Chicago Press, 1989.

C. Kennedy. Polar opposition and the ontology of degrees. *Linguistics and philosophy*, 24(1): 33–70, 2001.

R. Klatzky, E. Clark, and M. Macken. Asymmetries in acquisition of polar adjectives: Linguistic or conceptual? *Journal of Experimental Child Psychology*, 16:32–46, 1973.

A. Lehrer. Markedness and antonymy. *Journal of Linguistics*, 21(02):397–429, 1985.

M. P. Lewis, editor. *Ethnologue: Languages of the World*. SIL International, Dallas, Texas, 16th edition, 2009.

G. Sassoon. The degree functions of negative adjectives. *Natural Language Semantics*, 2010.

B. Wennergren. Tekstaro de Esperanto. Available online at `www.tekstaro.com`, 2003-.

L. Zamenhof. *Unua Libro*. Warsow, 1887.