

# Discovering Polarity for Ambiguous and Objective Adjectives through Adverbial Modification

Irene Russo

Istituto di Linguistica Computazionale, C.N.R.  
Area della Ricerca Via Moruzzi 1, 56124 Pisa, Italy  
E-mail: irene.russo@ilc.cnr.it

## Abstract

The field of opinion mining has emerged in recent years as an exciting challenge for computational linguistics: investigating how humans express subjective judgments through linguistic means paves the way for automatic recognition and summarization of opinionated texts, with the possibility of determining the polarities and strengths of opinions asserted. Sentiment lexicons are basic resources for investigating the orientation of a text that can be performed considering polarized words included in it but they encode the polarity of word types instead that the polarity of word tokens.

The expression of an opinion through the choice of lexical items is context-sensitive and sentiment lexicons can be integrated with syntagmatic patterns that emerge as significant with statistical analyses. In this paper it will be proposed a corpus analysis of adverbially modified ambiguous (e.g. *fast*, *rich*) and objective adjectives (e.g. *chemical*, *political*) - that can be occasionally exploited to express a subjective judgments -. Comparing polarity encoded in sentiment lexicons and the results of a logistic regression analysis, the role of adverbial cues for polarity detection will be evaluated on the basis of a small sample of sentences manually annotated.

## 1. Introduction

The field of opinion mining has emerged in recent years as an exciting challenge for computational linguistics: investigating how humans express subjective judgments through linguistic means paves the way for automatic recognition and summarization of opinionated texts, with the possibility of determining the polarities and strengths of opinions asserted.

Studies for opinion mining belong generally to either the data-driven approach, where an annotated corpus is used to train a machine learning classifier, or to the lexicon-based approach, where a pre-compiled list of manually selected sentiment terms is used to build a polarity score function. However, sentiment lexicons are often basic resources for the first task because the general evaluation of the orientation of a text can be performed considering polarized words included in it. Some empirical methods are developed to automatically identify adjectives, verbs, and N-grams that are statistically associated with subjective language (e.g., Turney, 2002; Hatzivassiloglou and McKeown, 1997).

In sentiment lexicons, each lexical item could be tagged on the basis of its prior polarity. The starting seed list can be increased (automatically or semi-automatically) with other resources (thesaurus, WordNet, General Inquirer) or with corpus based techniques such as co-occurrence with words of known polarity (Turney & Littman, 2003) and statistical measure of word associations.

One serious limitation of lexical resources for opinion mining is that they include the polarity of word types instead that the polarity of word tokens -prior polarity and not contextual polarity-. Analyzing mainly the negative or positive polarity of lexical items out of context is not

sufficient. The basic polarity of a lexical item can be modified at the lexical and discourse level (Polanyi & Zaenen, 2004).

Several attempts to encode polarity beyond word units have been made. For example (Wiebe et al., 2001) aims is to identify collocational clues of subjectivity as fixed sequences of words which, when they appear together, tend to be subjective. (Riloff & Wiebe, 2003) suggest an extraction pattern learning technique that can learn subjective expressions linguistically richer and more flexible than single words or N-grams, characterizable in terms of subtle connotations that are more expressive than single words.

However, beyond the encoding of specific collocations, it could be interesting to find generalizations, i.e. specific linguistic patterns for the expression of opinions.

The expression of an opinion through the choice of lexical items is quite context-sensitive and sentiment lexicons could be integrated with syntagmatic patterns that emerge as significant with statistical analyses and, when possible, with heuristics about how to manage contextual polarity.

In this paper a corpus based analysis of two types of adjectives with ambiguous polarity is proposed. The focus is on adjectives that are labelled as positive and negative at the same time in SentiWordNet (Esuli & Sebastiani, 2006) (e.g. *fast*, *rich*) and on objective adjectives that are not included in this resource but that can be occasionally exploited to express a subjective judgment (e.g. *chemical*, *political*).

Comparing polarity encoded in sentiment lexicons and the results of a logistic regression analysis, the role of adverbial cues for polarity detection will be evaluated on

the basis of a small sample of sentences manually annotated. The discovery of regularities can help to improve lexical resources with strategies to manage polarity emerging in context.

## 2. Occasionally Polarized Adjectives

In terms of polarity an adjective can be positive, such as *beautiful*, or negative, such as *horrid*, but apart from these clear cases, several difficulties emerge. First of all, a single lexical item could be negative or positive depending on the context:

1a. It is a wine of impressive concentration, with an *intense fragrance* of black fruit. (+)

1b. Lightweights who find other beers of the *style too intense* would likely enjoy this brew. (-)

Moreover, objective adjectives that could be occasionally exploited as positively or negatively marked are interesting but quite hard to identify. Relational adjectives (e.g. *domestic*, *medical*) are supposed to have no orientation (Hatzivassiloglou & Wiebe, 2000) but it's possible to find counter-examples in review aggregation sites such as Epinions.com, especially when they are modified by degree adverbs:

2a. It has a *very chemical taste*, but with a hint of raspberry.

2b. Vianne and her child arrive in a small, *quite French town* in the winter.

The occasional polarity of objective adjectives should not be necessarily included in sentiment lexicons because it's not a stable property; however, it's a kind of information that is salient for the analysis of subjective judgments. Even if it's not really frequent, it's highly relevant, especially in product reviews and it has been demonstrated how rare terms and *hapax legomena* could be high-precision indicators of subjectivity (Yang et al., 2006).

In subjective contexts it is possible to find low-frequency words because people tend to be creative when they are being opinionated. Low-frequency words in specific syntagmatic slot can be very informative for recognizing subjectivity (Wiebe et al. 2001):

3. 'What a' NP: *What a divine dress/ heavenly sunset*.

Because of the high informativeness of these examples, a sentiment lexicon should be improved with heuristics that take into consideration the contextual influence of the surrounding lexical items, through the identification of syntagmatic contexts that potentially trigger an evaluative meaning. In theoretical terms, it is a matter of semantic prosody as "the spreading of connotational colouring

beyond single word boundaries"(Partington 1998: 68), that involve words that seem to be neutral out of context.

As Hoey (2005: 8) points out, our knowledge of language is deeply influenced by our massive exposure to co-occurrences that model our mental lexicon beyond our awareness: "Every word is mentally primed for collocational use. As a word is acquired through encounters with it in speech and writing, it becomes cumulatively loaded with the context and co-texts in which it is encountered, and our knowledge of it includes the fact that it co-occurs with certain other words in certain kinds of context."

For this reason, the evidence that in context the polarity of a word could be different or changing with respect to its prior encoded polarity should be included - at least partially - in resources created for opinion mining.

## 3. Adjectives and Adverbs in Sentiment Lexicons

In terms of class of words, adjectives are good features according to (Bruce & Wiebe, 2000) because there is a high correlation between the presence of adjectives and the subjectivity of a sentence: the probability of a sentence to be subjective with just one adjective is 56%. Also for SentiWordNet, adjectives and adverbs are very often subjective (Esuli & Sebastiani 2006). As a consequence, a fine grained representation for adjectives - through a detailed analysis of their changing polarity in context - will improve lexicographic resources created for opinion mining tasks.

The focus of this work is on adverbial modification of adjectives because it's a relevant linguistic context for the identification of occasional subjective exploitations of objective adjectives (Russo 2009) and it can be used to discover the polarity of ambiguous adjectives. Even in opinion mining gradability is considered an essential feature for subjectivity (Hatzivassiloglou & Wiebe 2000). For example, a detailed analysis of the core properties of products modified by degree adverbs has proven more effective than a simpler co-occurrences approach (Chklowski, 2006).

In the present work two sentiment lexicons - OpinionFinder lexicon and SentiWordNet - will be compared and evaluated with respect to the way they encode adjectives and adverbs. The relevance of these POSs in the two resources emerges from Table 1.

	Adj	Adv	Verbs	Nouns
OF	44.5%	10.9%	16.2%	27.9%
SWN	16%	3,2%	11,8%	69%

Table 1 Percentages of polarized POSs in the OpinionFinder lexicon and SWN

The subjectivity lexicon of OpinionFinder (Wiebe and Riloff, 2005) includes words and phrases manually annotated in texts that may be used to express private states; words that are subjective in most contexts were marked strongly subjective (strongsubj), and those that may only have certain subjective usages were marked weakly subjective (weaksbj). The list was expanded using a dictionary and a thesaurus, and also adding words from the General Inquirer positive and negative word lists which they judged to be potentially subjective.

This list of subjectivity cues is part of a system that processes documents and automatically identifies subjective sentences, identifying when opinions, sentiments, speculations and other private states are present in text.

SentiWordNet (Esuli & Sebastiani, 2006) is an automatically generated sentiment lexicon using a semi-supervised method. It is based on the quantitative analysis of the glosses associate to synsets of Word-Net, where each synset is assigned three probability scores (positive, negative, and objective) that add up to 1.

It is better than Opinion Finder's subjectivity lexicon (Wilson et al. 2005) to refine semantic analysis because it allows fuzzy values, while OpinionFinder's lexicon simply distinguishes between weak and strong subjectivity, but it doesn't provide rules to disambiguate locally the polarity of polysemous items. In this work, as an approximation of lexical items' polarity, the sums of positive and negative values will be considered for each term. For example, the adverb *so* has nine senses, some of them mixed in terms of polarity (positive and objective, negative and objective etc.) but it will be consider positive because the sum of positive value is higher than that of negative values.

Because of their different structures, several mismatches can be observed between these two resources and they will be relevant for the case study in par. 5. The focus of the present work is on adverbial modification and table 2 shows how many adverbs are encoded as positive, negative, neutral or ambiguous for the two resources, after the processing of SWN data based on the sums of positive and negative values.

	Pos	Neg	Neu	Amb	Tot
OF	314	506	76	4	900
SWN	2123	604	939	0	3666

Table 2 Adverbs in OpinionFinder lexicon and SWN grouped according to polarity.

#### 4. Data Analysis

To discover the adverbs that show a preference for positive or negative adjectives and potentially determine the polarity of objective and ambiguous adjectives both sentiment lexicons and corpus data are useful.

In this work, co-occurrences frequencies with modifying adverbs for a dataset of 50 strongly positive adjectives (e.g. *perfect*, *encouraging* etc.) and 50 strongly negative adjectives (e.g. *deficient*, *horrid* etc.) are considered. Co-occurrences frequencies where extracted from the Google Web 1T 5-Gram Database, a collection of frequent 5-grams extracted from approximately 1 trillion words of Web text collected by Google Research and consulted through the web-interface developed by Stefan Evert.

The adjectives are highly polarized both for Opinion Finder lexicon and SentiWordNet and are more frequent than 200 in the British National Corpus.

A very interesting finding is that OpinionFinder's list includes adverbs that frequently co-occur with negative and positive adjectives as weak neutral adverbs (*very*, *so*, *rather* etc.). Moreover, adverbs that frequently co-occur with negative and positive adjectives are generally ambiguous (e.g. *too*, *quite*) or without polarization (e.g. *most*, *almost*) in SWN, with several exceptions (*more* is slightly negative, *very* is positive).

If we consider the adverbs most frequently modifying the two sets (adverbs more frequent than 30), it's clear that positive adverbs tend to prefer positive adjectives but the same is not true for negative adverbs (Tables 3.1 and 3.2):

	PosAdv	NegAdv	NeuAdv
posAdj	10	4	11
negAdj	1	4	11

Table 3.1 OpinionFinder lexicon

	PosAdv	NegAdv	NeuAdv
PosAdj	19	4	5
NegAdj	15	4	4

Table 3.2 SentiWordNet lexicon.

In general terms, it's clear that these resources cannot be very useful for a detection of ambiguous and objective adjectives polarity based on adverbial modification. Concerning modified objective adjectives, just few cases will be recognized as opinion/sentiment oriented. For example, for the adjective *political*, 4a. will not be recognized as positively or negatively polarized by OpinionFinder lexicon and SWN while 4b. will be not detected by OpinionFinder, even if both sentences convey an evaluative meaning:

4a. I want to be clear, I didn't find the *book too political*.

4b. It's not *so political* and personal for him.

##### 4.1 Logistic Regression Analysis

A logistic regression analysis has been performed on the two datasets of 50 strongly positive (e.g. *perfect*,

encouraging etc.) and 50 strongly negative (e.g. *deficient, horrid* etc.) adjectives modified by adverbs.

The aim of this logistic regression analysis is to find which adverbs could help to label automatically polarity at phrasal level when they modify ambiguous or objective adjectives in context. Logistic regression analysis of data highlights two sets of 33 adverbs that prefer negative or positive adjectives, reported in Table 4:

Negative cues	<i>almost, clearly, enough, exceptionally, increasingly, incredibly, less, little, most, partly, pretty, quite, rather, seriously, slightly, thoroughly, too, totally, well</i>
Positive cues	<i>equally, extremely, fairly, fully, highly, largely, more, partially, perfectly, really, simply, so, very, wholly</i>

Table 4 Adverbial cues for ambiguous and objective adjectives.

## 5. Evaluation: a Case Study

To evaluate the effectiveness of sentiment lexicons and cues emerging from logistic regression analysis to determine polarity of ambiguous and occasional subjective adjectives, a dataset of 100 sentences has been created. They are randomly selected from Epinions.com and result from the agreement of two native speakers of English on a larger set. Sentences have been classified as positive, negative or undetermined in terms of polarity.

In 50 of them there are adverbially modified ambiguous adjectives that are highly polarized according to SentiWordNet, in 40 of them there are high frequency relational adjectives (ending with the suffixes *-ic /-ical*) that are neutral according to the same resource (e.g. *systematic, hierarchical*) while in 10 there are nationality adjectives (e.g. *English, French*). Judgments has been compared with results obtained on the basis of the polarity of adverbs encoded in the OpinionFinder lexicon, SentiWordNet and emerged through the logistic regression analysis (par. 4.1), using adverbs that prefer positive or negative adjectives as cues of opinionated content in terms of semantic prosody. With respect to human judgments, syntagmatic cues emerging from logistic regression analysis (Table 5) help to find the majority of positive and negative sentences with respect to OpinionFinder and SWN and this generalization is true both for the subset of sentences with ambiguous and objective adjectives, even if the data set is too small to warrant a well-founded comparison.

	Neg	Pos	Undetermined
Human Judgments	32	39	29
OF	11	7	15
SWN	8	20	6
LR	19	19	0

Table 5 – Comparison on 100 sentences between human judgments, lexical resources and syntagmatic cues from logistic regression analysis.

However, the polarized cues emerged through logistic regression analysis cannot find sentences labelled as undetermined by human judgments:

5. It is a *very German movie*, I'm sorry but I can't describe why.

Moreover, the incidence of other syntagmatic cues in the sentences is not considered: the presence of negations or polarized nouns could produce a shift in terms of polarity.

## 6. Conclusion

Procedures for opinion detection and classification can be ameliorated through automatic classifications that could be refined with heuristics, going beyond a bag-of-words approach, and through lexicographic resources that can be enriched, including semantic orientation at the phrasal level (Wilson et al., 2005) and fuzzy values (Andreevskaia & Bergler, 2006). Two strategies are possible: to enrich sentiment lexicons with values relative to collocational units or to encode specific rules that help to establish the polarity in context – in specific syntagmatic patterns.

In this paper the role of adverbial modification has been assessed and it has been evaluated the relevance of semantic preferences of adverbs in predicting polarity of ambiguous adjectives and occasional subjective uses of objective adjectives as items of a simple syntagmatic pattern. Considering the dataset manually annotated the results show how syntagmatic cues classified as positive or negative with the logistic regression analysis perform better than polarized adverbs encoded in OpinionFinder lexicon and SentiWordNet but both ambiguous and objective adjectives adverbially modified can be further investigated taking into account the influence of surrounding lexical items.

## 7. References

- Andreevskaia, A., Bergler, S. (2006). Mining WordNet for Fuzzy Sentiment: Sentiment Tag Extraction from WordNet Glosses. In *Proceedings of the 11th Conference of the European Chapter of the Association for the Computational Linguistics*, EACL-2006.
- Bruce, R., Wiebe, J. (2000). Recognizing subjectivity: A case study of manual tagging. *Natural Language Engineering*, v.5 n.2, p.187-205.
- Chklovski, T. (2006). Deriving Quantitative Overviews of Free Text Assessments on the Web. In *Proceedings of the 11th international conference on Intelligent user interfaces*.
- Esuli, A., Sebastiani, F. (2006). SentiWordNet: A Publicly Available Lexical Resource for Opinion Mining. In *Proceedings of LREC-06, 5th Conference on Language Resources and Evaluation*, Genova, IT, pp. 417-422.



- Hatzivassiloglou, V., McKeown, K. (1997). Predicting the Semantic Orientation of Adjectives. In *Proceedings of the 35th Annual Meeting of the ACL and the 8th Conference of the European Chapter of the ACL*, pp. 174–181.
- Hatzivassiloglou, V. , Wiebe, J. (2000). Effects of Adjective Orientation and Gradability on Sentence Subjectivity. In *Proceedings of COLING 2000*, pp. 299-305.
- Hoey, M. (2005). *Lexical Priming: A New Theory of Words and Language*. London: Routledge.
- Wiebe, J., Wilson, T., Bell, M. (2001). Identifying Collocations for Recognizing Opinions. In *Proc. ACL 01 Workshop on Collocation*.
- Wiebe, J., Riloff, E. (2005). Creating Subjective and Objective Sentence Classifiers from Unannotated Texts. In *Proceedings of the 6th International Conference on Computational Linguistics and Intelligent Text Processing (CICLing-05)*.
- Wilson, T., Wiebe, J., Hoffmann, P. (2009). Recognizing Contextual Polarity: An Exploration of Features for Phrase-Level Sentiment Analysis. *Computational Linguistics* 35(3): pp. 399-433.
- Partington, A. (1998). *Patterns and meanings*. Amsterdam and Philadelphia: Benjamins.
- Riloff, E., Wiebe, J. (2003). Learning Extraction Patterns for Subjective Expressions. In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing (EMNLP-03)*.
- Russo, I. (2009). *Usi qualificativi degli aggettivi relazionali in italiano e in inglese*. Unpublished PhD thesis.
- Turney, P. D. (2002). Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pp. 417–424.
- Turney, P.D., Littman, M.L. (2003). Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems (TOIS)*, 21(4), pp. 315-346.
- Yang, K., Valerio, N., Zhang, H. (2006), WIDIT in TREC-2006 blog track. In *Proceedings of TREC, 2006*.