

The Role of Adverbs in Sentiment Analysis

Eduard C. Dragut

Computer and Information Sciences Dept.
Temple University
edragut@temple.edu

Christiane Fellbaum

Department of Computer Science
Princeton University
fellbaum@princeton.edu

Abstract

Sentiment Analysis, an important area of Natural Language Understanding, often relies on the assumption that lexemes carry inherent sentiment values, as reflected in specialized resources. We examine and measure the contribution that eight intensifying adverbs make to the sentiment value of sentences, as judged by human annotators. Our results show, first, that the intensifying adverbs are not themselves sentiment-laden but strengthen the sentiment conveyed by words in their contexts to different degrees. We consider the consequences for appropriate modifications of the representation of the adverbs in sentiment lexicons.

1 Introduction

It was probably Chuck who coined the term “arm-chair linguist” (Svartvik, 1991). Chuck Fillmore’s deep commitment to the study of language — in particular lexical semantics — on the basis of corpus data served as a model that kept many of us honest in our investigation of language. Today, we are lucky to be able to work from our office chairs while collecting data from a broad speaker group by means of crowdsourcing. And Chuck’s FrameNet taught us the importance of considering word meanings in their contexts. Our paper presents work that tries to take this legacy to heart.

2 Sentiment Analysis

Broadly speaking, sentiment analysis (SA) attempts to automatically derive a writer’s “sentiment” about the topic of a text. “Sentiment” is usually categorized into “positive,” “neutral” and

“negative,” where positive corresponds to satisfaction or happiness and “negative” to dissatisfaction or unhappiness. Some work in SA further distinguishes degrees of positive and negative sentiment. SA often refers to lexical resources where words are annotated with a sentiment value. SentiWordNet (SWN) (Esuli and Sebastiani, 2006) assigns one of three sentiment values to each synset in WordNet (Fellbaum, 1998). Opinion Finder (OF) (Wilson et al., 2005) identifies the sentiment of the writer. Other resources include Appraisal Lexicon (AL) (Taboada and Grieve, 2004) and Micro-WNOp (Cerini et al., 2007).

Much of this work relies on the assumption that specific lexemes (unique mappings of word forms and word meanings) carry an inherent sentiment value. This seems intuitively correct for words like *enjoy* (positive), *pencil* (neutral) and *pain* (negative).

Other words may not carry inherent sentiment value yet, in context, contribute to that of the words they co-occur with or modify. One such class of words comprises what we call *polarity intensifiers*. In this preliminary study, we analyze the contribution of *adverbial intensifiers* to the sentiment value of the sentences in which they occur.

Consider the adverb *absolutely* in two sample sentences from movie reviews:

S1 He and Leonora have absolutely no chemistry on screen whatsoever.

S2 I was absolutely delighted by the simple story and amazing animation.

The goal of this preliminary experimental study is to seek answers to the following questions

Adverbs	OF	AL	SWN
absolutely	Neu.	–	Neu.
awfully	Neg.	Neg.	Neu.
enormously	Neg.	–	Neu.
extremely	Neg.	–	Pos.
horribly	Neg.	Neg.	Neu.
incredibly	Pos.	Pos.	Neu.
pretty	Pos.	Pos.	Neu.
seriously	Neg.	–	Neu.

Table 1: Eight intensifying adverbs and their polarities in sentiment lexicons.

1. Do the adverbs we investigate carry inherent sentiment values, as postulated by some sentiment lexicons?
2. Which adverbs have the strongest sentiment intensifying effect?
3. Do some adverbs have a stronger effect on sentences with a negative polarity or on sentences with a positive polarity?
4. Does the presence or absence of each adverb affect the direction of the polarity of the sentence?

3 The Experiment

We analyze whether human judgments show an effect on the sentiment ratings of sentences in the presence or absence of selected adverbs, and how strong the effect of each adverb is.

Let $S1'$ be the sentence $S1$ from which an adverb like *absolutely* is removed. $S2'$ is defined similarly. Three main observations can be made: (1) the adverb appears in both positive and negative sentiment-bearing sentences ($S1$ is negative and $S2$ is positive); (2) its removal from either $S1$ or $S2$ does not change the overall polarity of the sentence; (3) intuitively, $S1$ has a stronger negative polarity value than $S1'$ and $S2$ has a stronger positive polarity value than $S2'$. We conduct a preliminary study of polarity intensifier words and show that they all have characteristics (1) - (3). We examine data with eight different adverbs (Table 1).

3.1 Data

We extracted sentences containing the target adverbs from a corpus of 50,000 movie reviews (Maas et al., 2011). Each sentence is extracted from a review that is labeled either “positive” or “negative” and correlated with a star rating. We

manually inspected the sentences and discarded those where the target adverb was used in a modal sense, as in *Seriously, there was not one respectable character in the entire script* while retaining sentences like *There is no doubt that Alfred Hitchcock was a seriously talented director*. For each adverb, we retained ten sentences from positive and negative reviews each, for a total 160 sentences. We copied the original sentences, removed the adverbs without making additional alterations. Our final dataset consisted of a total of 320 sentences with 160 sentence pairs whose members were identical except for the presence or absence of the target adverbs. Below is an example of a sentence pair, where the original sentence with the adverbs was pre-classified by (Pang and Lee, 2004) as carrying positive sentiment.

1. I was absolutely delighted by the simple story and amazing animation.
2. I was delighted by the simple story and amazing animation.

3.2 Collecting Judgments via Crowdsourcing

We submitted single sentences (not pairs) to be annotated with sentiment scores for crowdsourcing, using Amazon Mechanical Turk (AMT). To avoid any bias we shuffled the sentences and displayed them individually. We asked the Turkers to select, for each sentence, one of five sentiment scores: strong positive (2), positive (1), neutral (0), negative (-1), strong negative (-2). Each sentence was rated by five annotators. Altogether, twenty annotators completed the task within eight hours. Since the annotators did not all judge the same set of sentences, we computed the agreement between annotators as follows. For each annotator, his/her agreement with the others is given by the following formula:

$$\frac{1}{|S(i)|} \sum_{j \in S(i)} ps_{ji},$$

where $S(i)$ is the set of sentences annotated by the i^{th} Turker and ps_{ji} is the percentage of Turkers who have the same annotation with the i^{th} Turker for sentence j . $|S(i)|$ is the cardinality of set $S(i)$. The agreement ranges from 0.52 to 0.8. Although the annotation of some Turkers is close to that of flipping a coin, all judgments were retained and included in the results reported here.

3.3 Results

We report the main results. The polarity rating of a sentence j is the (un-weighted) average rating

Adverbs	Avg. Pol. Change	Pol. Reversal
absolutely	0.2	0/20
awfully	0.6	2/20
enormously	0.2	1/20
extremely	0.2	2/20
horribly	0.2	0/20
incredibly	0.2	4/20
pretty	0.2	1/20
seriously	0.4	3/20

Table 2: Effects of adverbs on sentiment ratings.

of the five annotators for the sentence, denoted α_j and $\alpha_j = \sum_i ps_{ji}$. We use uniform weighting. A sentence j is classified into one of the five polarity categories according to the following criteria:

strong positive	if $\alpha_j \in [1.5, 2]$
positive	if $\alpha_j \in (1.5, 0.5]$
neutral	if $\alpha_j \in (0.5, -0.5)$
negative	if $\alpha_j \in [-0.5, -1.5]$
strong negative	if $\alpha_j \in [-0.5, -2]$

3.3.1 Do Adverbs Change Sentiment Rating?

We first examine the polarity intensifying effects of the eight adverbs and determine their relative intensifying effects. For each adverb we compute the average polarity rating change between the members of the 20 sentence pairs with and without the target adverb. The second column of Table 2 shows the average polarity rating change for the adverbs. All adverbs have polarity intensifying effect, which ranges from 0.2 to 0.6. *Awfully* and *seriously* have the strongest effect.

3.3.2 Change of Sentiment Rating in Positive vs. Negative Contexts

Next we ask whether the adverbs have a stronger polarity intensifying effect on sentences with a negative, positive or neutral ratings. We partition the 20 sentences with/without each adverb into the three polarity categories according to their average polarity ratings. A sentence j is negative (positive) if $\alpha_j \leq -0.5$ ($\alpha_j \geq 0.5$). Figure 1 shows the results. For six out of the eight adverbs, the graph follows a V-shaped pattern, indicating that the adverbs have stronger polarity influence on sentences conveying opinionated, but not neutral, statements. *Pretty* shows the weakest effect across, which makes intuitive sense, as this




Figure 1: The polarity intensifying effects of adverbs over the sentiment categories.

adverb seems to have a “softening/weakening” effect: consider “pretty good,” which one could judge to be slightly less good than “good.” For example, the sentence

He has a pretty strident rant about how important it is.

received an average rating score of 0 with the adverb present and -0.2 without it. The results for *awfully* and *extremely* are surprising. A closer look at the annotations revealed some possible unreliable ratings. For example, the sentence

The part of the movie set in Vietnam was extremely inaccurate.

has average polarity score of 0 (i.e., neutral) with the adverb and -0.8 without. Intuitively, it seems that the first sentence conveys a strong negative sentiment. Such data indicate the need for further study. A more complex scheme for computing the average polarity scores, such as weighted by inter-annotator agreement, might produce better results.

3.3.3 Can Adverbs Reverse Sentiment Orientation?

We ask whether their presence can have the effect of reversing the polarity of a sentence. We again consider three sentiment categories: positive, negative and neutral. The third column in Table 2 shows for each adverb, how many sentences out of the total of 20 were judged to have a reversed polarity when the adverb was removed. Overall, the polarities of only 13 out of 160 sentences (i.e., about 8%) change.

3.3.4 Do Adverbs Have an Inherent Sentiment Value?

Our target adverbs have inherent polarity as claimed in some sentiment lexicons (see Table 1).

If the polarity of a sentence does not change when the adverbs is present or absent, we conclude that the adverb has no inherent polarity but may merely affect the intensity of the constituents that it modifies. These results, as displayed in Figure 1 indicate that our target adverbs do not carry inherent polarity. Instead, they modify the intensity of the sentiment connoted by the context.

4 Discussion

We examined the effect of eight intensifying adverbs on the sentiment ratings of the sentences in which they occur. Our study showed that, contrary to their representation in some widely used sentiment lexicons, these adverbs do not carry an inherent sentiment polarity, but merely alter the degree of the polarity of the constituents they modify; corrections of the corresponding entries in the sentiment resources seem warranted. Our results show further that all adverbs strengthen the polarity of the context to different degrees. If confirmed on a larger data set, this indicates that the intensifying force of different adverbs should be reflected in lexical resources, perhaps along an ordered scale.

5 Related Work

Two recent surveys give a detailed account of the SL acquisition techniques (Feldman, 2013; Liu, 2012). We give only an overview of the related work here. SLs are acquired by one of three methods. Manual tagging is performed by human annotators: e.g., OF, and AL. Dictionary-based acquisition relies on a set of seed words that is expanded by using external resources, such as WordNet: e.g., (Dragut et al., 2010; Hassan and Radev, 2010; Mohammad et al., 2009; Dragut et al., 2012; Takamura et al., 2005). In corpus-based acquisition a set of seed words is expanded by using a large corpus of documents (Feng et al., 2013; Lu et al., 2011; Yu et al., 2013; Wu and Wen, 2010).

To our knowledge, none of these works include the polarity intensifiers that we introduce in this paper.

References

S. Cerini, V. Compagnoni, A. Demontis, M. Formentelli, and G. Gandini, 2007. *Language resources and linguistic theory: Typology, second language acquisition, English linguistics*.

Eduard C. Dragut, Clement T. Yu, A. Prasad Sistla, and

Weiyi Meng. 2010. Construction of a sentimental word dictionary. In *CIKM*, pages 1761–1764.

Eduard C. Dragut, Hong Wang, Clement Yu, Prasad Sistla, and Weiyi Meng. 2012. Polarity consistency checking for sentiment dictionaries. In *ACL*.

A. Esuli and F. Sebastiani. 2006. Sentiwordnet: A publicly available lexical resource for opinion mining. In *LREC*.

Ronen Feldman. 2013. Techniques and applications for sentiment analysis. *CACM*, 56(4):82–89, April.

C. Fellbaum. 1998. *WordNet: An On-Line Lexical Database and Some of its Applications*. MIT Press.

Song Feng, Jun Sak Kang, Polina Kuznetsova, and Yejin Choi. 2013. Connotation lexicon: A dash of sentiment beneath the surface meaning. In *ACL*.

Ahmed Hassan and Dragomir Radev. 2010. Identifying text polarity using random walks. In *ACL*.

Bing Liu. 2012. *Sentiment Analysis and Opinion Mining*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.

Yue Lu, Malu Castellanos, Umeshwar Dayal, and ChengXiang Zhai. 2011. Automatic construction of a context-aware sentiment lexicon: an optimization approach. In *WWW*, pages 347–356. ACM.

Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *ACL*, pages 142–150.

Saif Mohammad, Cody Dunne, and Bonnie Dorr. 2009. Generating high-coverage semantic orientation lexicons from overtly marked words and a thesaurus. In *EMNLP*.

B. Pang and L. Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *ACL*.

Jan Svartvik, editor. 1991. *Directions in Corpus Linguistics*. Nobel Symposium 82, Mouton de Gruyter.

M. Taboada and J. Grieve. 2004. Analyzing appraisal automatically. In *AAAI Spring Symposium*.

Hiroya Takamura, Takashi Inui, and Manabu Okumura. 2005. Extracting semantic orientations of words using spin model. In *ACL*, pages 133–140.

T. Wilson, J. Wiebe, and P. Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *HLT/EMNLP*.

Yunfang Wu and Miaomiao Wen. 2010. Disambiguating dynamic sentiment ambiguous adjectives. In *COLING*, pages 1191–1199.

Hongliang Yu, Zhi-Hong Deng, and Shiyinxue Li. 2013. Identifying sentiment words using an optimization-based model without seedwords. In *ACL*.