

The Utility of Linguistic Rules in Opinion Mining

Xiaowen Ding and Bing Liu

Department of Computer Science
University of Illinois at Chicago
851 S. Morgan Street
Chicago, IL 60607-0753

{xding, liub}@cs.uic.edu

ABSTRACT

Online product reviews are one of the important opinion sources on the Web. This paper studies the problem of determining the semantic orientations (positive or negative) of opinions expressed on product features in reviews. Most existing approaches use a set of *opinion words* for the purpose. However, the semantic orientations of many words are context dependent. In this paper, we propose to use some linguistic rules to deal with the problem together with a new opinion aggregation function. Extensive experiments show that these rules and the function are highly effective. A system, called *Opinion Observer*, has also been built.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval – *Information filtering*. I.2.7 [Artificial Intelligence]: Natural Language Processing – *text analysis*.

General Terms

Algorithms, Experimentation.

Keywords

Opinion mining, Sentiment analysis.

1. INTRODUCTION

In recent years, many researchers studied the problem of opinion mining from Web sources (e.g., product reviews, forum posts and blogs) [4]. In this work, we focus on product reviews. The general problem of opinion mining from product reviews is to identify product features that have been commented on by each reviewer and to determine whether the opinions are positive or negative (called *semantic (or opinion) orientation*) [2, 4]. In this paper, we only study the latter problem, i.e., determining whether an opinion is positive or negative. We assume that product features are given or have been discovered by another system [e.g., 2].

Existing techniques use opinion words such as “great”, “amazing”, “poor”, “bad”, etc, to decide the orientation of an opinion on a product feature. Although the orientations of these words are obvious, the orientations of many other words depend on context. For example, the word “small” can indicate a positive or a negative opinion on a product feature depending on the feature. We propose several linguistic rules to deal with the problem. The approach tries to infer the orientations of opinions on a product feature using context. A new function for aggregating multiple opinions in a sentence is also presented. Experiment results show that the rules and the aggregation function are very useful.

Copyright is held by the author/owner(s).
SIGIR’07, July 23–27, 2007, Amsterdam, The Netherlands.
ACM 978-1-59593-597-7/07/0007.

2. THE PROPOSED TECHNIQUE

Opinion aggregation function: A review sentence s may contain multiple product features (f_1, \dots, f_m) and multiple opinion words (w_1, w_n) . Our objective is to determine the orientation of opinion expressed on each feature f_i in s , i.e., the pair (f_i, s) . We use an opinion aggregation function to perform this task.

Given lists of positive, negative and context dependent opinion words, including phrases and idioms, the system works as follows:

- It first segments s using BUT words/phrases (e.g., “but”, “except that”, etc.). Assume f_i is in segment s_k . We compute the opinion orientation score of f_i in s_k using the following:

$$score(f_i, s_k) = \sum_{w_j \in s_k} \frac{w_j \cdot SO}{d(w_j, f_i)}, \quad (1)$$

where w_j is an opinion word in s_k , which is the sentence segment that contains the feature f_i , and $d(w_j, f_i)$ is the distance between feature f_i and opinion word w_j in s_k . $w_j \cdot SO$ is the semantic orientation of w_j . The multiplicative inverse in the formula is used to give low weights to opinion words that are far away from feature f_i . Note that a positive word is assigned the semantic orientation score of 1, and a negative word is assigned the semantic orientation score of -1 .

In [2], a simple summation is used. The new function is better because far away opinion words may not modify the current feature. However, setting a distance window within which the opinion words are considered does not perform well either because in some cases, the opinion words may be far away. The proposed new function can deal with both problems.

Note that a product feature itself can be an opinion word as it may be an adjective representing a *feature indicator*, e.g., “reliable” in the sentence “this camera is very reliable”. In this case, $score(f_i, s_k)$ is 1 or -1 depending on whether f_i (e.g., “reliable”) is positive or negative (Equation (1) is not used).

- If the final score is positive, then the opinion on the feature in the sentence s is positive. If the final score is negative, then the opinion on the feature is negative. It is neutral otherwise.

Context dependent opinion words: It was found that some linguistic rules or conventions can be used to infer opinions:

1. *Intra-sentence conjunction rule:* For example, we have the sentence, “the battery life is very long”. It is not clearly whether “long” means a positive or a negative opinion. Our algorithm tries to see whether any other reviewer said that “long” is positive (or negative). For example, another reviewer wrote “This camera takes *great* pictures and has a *long* battery life”. From this sentence, we can discover that “long” is *positive* for “battery life” because it is conjoined

with the positive opinion word “great”. We call this the *conjunction rule*, which means that a sentence only expresses one opinion orientation unless there is a BUT word which changes the direction. The following sentence is unlikely: “This camera takes *great* pictures and has a *short* battery life”

2. *Pseudo intra-sentence conjunction rule*: Sometimes, one may not use an explicit conjunction “and”. Let us use the example sentence “the battery life is long” again. We have no idea whether “long” positive or negative for “battery life”. A similar strategy can be applied. For instance, another reviewer might have written the following: “The camera has a long battery life, which is great”. The sentence indicates that the semantic orientation of “long” for “battery life” is positive due to “great”, although no explicit “and” is used.
3. *Inter-sentence conjunction rule*: The conjunction rule can also be extended to neighboring sentences. The idea is that it is common to express the same opinion in a few consecutive sentences. Opinion changes are indicated by words such as “but”, “however”, etc. For example, the following passages are natural: “The picture quality is great. The battery life is long” and “The picture quality is great. However, the battery life is short”. However, the following passage is not natural: “The picture quality is great. The battery life is short”.

Although we do not know whether “long” (or “short”) is positive or negative for “battery life”, if we know that “great” is positive then we can infer that “long” is positive and “short” is negative for “battery life”.

For rules 1 and 2, it is possible that in the reviews of a product the same opinion word for the same feature has conflicting orientations. For example, another reviewer may say: “This camera is very small, which I don’t like”. In this case, our algorithm takes the majority view. If more people indicate that “small” is positive for size, we treat it as positive and vice versa.

Synonym and Antonym Rule: If a word is found to be positive (or negative) in a context for a feature, its synonyms are also considered positive (or negative), and its antonyms are considered negative (or positive). For example, in the above sentence, we know that “long” is positive for “battery life”. Then we also know that “short” is negative for “battery life”.

We also have a set of rules to consider special words such as “but”, “however”, “not”, “too”, etc, which are also important.

Algorithm: The algorithm, which considers all the rules and the opinion aggregation function, is fairly simple to design. It works iteratively. In each iteration, the opinion orientations of some opinion words can infer the orientations of some other opinion words, which in turn will help infer yet some more subsequently. Due to space limitations, we are unable to give the details.

Related work: Our technique is similar to those in [1, 3], which use conjunction rules to find *domain opinion words* from large domain corpora. There is, however, an important difference. We believe that finding domain opinion words is still problematic because in the same domain the same word may have different orientations depending on what features it is applied to. For example, in the following review sentences in the camera domain, “the battery life is very *long*” and “it takes a *long* time to focus”, “long” is positive in the first sentence, but negative in the second. Thus, we need to consider both the feature and the opinion word rather than only the opinion word as in [1, 3]. [5] also uses similar rules to compute opinion orientations based on relaxation labeling. However, [5] produces poorer results than our method.

Table 1: Results of opinion orientation prediction

OW	Baseline			Context rules		
	recall	precision	F-score	recall	precision	F-score
20	0.52	0.85	0.65	0.73	0.76	0.74
60	0.56	0.86	0.68	0.75	0.78	0.76
100	0.58	0.86	0.69	0.76	0.79	0.77
200	0.62	0.86	0.72	0.78	0.81	0.79

Table 2: Comparison of FBS, OPINE and Opinion Observer based on the data set in [2], reviews of 5 out of our 14 products.

	Precision	Recall	F-Score
FBS	0.93	0.76	0.83
OPINE	0.86	0.89	0.87
Opinion Observer	0.92	0.91	0.91

3. RESULTS AND CONCLUSION

We carried out experiments using 742 customer reviews from 14 products: four digital cameras, one DVD player, three MP3 player, two cellular phones, two routers, one antivirus software, and one diaper. All the reviews were extracted from amazon.com.

We want to determine the semantic orientation of each opinion expressed on every given product feature. We vary the number of given adjective and adverb opinion words from 20 to 200, and see whether the linguistic rules can improve the performance. Table 1 gives the results. OW means the number of given opinion words. The performances are measured using the standard evaluation measures of precision (*p*), recall (*r*) and F-score (*F*), $F = 2pr/(p+r)$.

In the table, each baseline result is obtained without using the linguistic context rules. We can see that context rules are indeed very helpful. They are able to improve the recall dramatically without much loss in precision. The gains in F-scores are dramatic. Naturally as the number of given opinion words increases, the improvement decreases slightly.

Table 2 compares our system (called *Opinion Observer*) with two existing systems FBS [2] and OPINE [5]. Here we only used reviews of 5 products from the benchmark dataset given in [2] as there is no result for the other reviews for OPINE. For Opinion Observer, all the opinion words (1242) and linguistic rules were used. We can see that it outperforms both FBS and OPINE. Additional experiments show that both the new opinion aggregation function and linguistic rules contribute roughly equally to the improved results of Opinion Observer over FBS. We can thus conclude that both the linguistic rules and the aggregation function are highly effective.

4. REFERENCES

- [1]. V. Hatzivassiloglou, and K. McKeown, Predicting the Semantic Orientation of Adjectives. *ACL-97*, 1997.
- [2]. M. Hu and B. Liu. Mining and Summarizing Customer Reviews. *KDD-04*, 2004.
- [3]. H. Kanayama and T. Nasukawa. Fully Automatic Lexicon Expansion for Domain-Oriented Sentiment Analysis. *EMNLP-06*, 2006.
- [4]. B. Liu. *Web Data Mining: Exploring Hyperlinks, Contents and Usage Data*. Springer, 2007.
- [5]. A-M. Popescu and O. Etzioni. Extracting Product Features and Opinions from Reviews. *EMNLP-05*, 2005.