

Identifying Noun Product Features that Imply Opinions

Lei Zhang

University of Illinois at Chicago
851 South Morgan Street
Chicago, IL 60607, USA
lzhang3@cs.uic.edu

Bing Liu

University of Illinois at Chicago
851 South Morgan Street
Chicago, IL 60607, USA
liub@cs.uic.edu

Abstract

Identifying domain-dependent opinion words is a key problem in opinion mining and has been studied by several researchers. However, existing work has been focused on adjectives and to some extent verbs. Limited work has been done on nouns and noun phrases. In our work, we used the feature-based opinion mining model, and we found that in some domains nouns and noun phrases that indicate product features may also imply opinions. In many such cases, these nouns are not subjective but objective. Their involved sentences are also objective sentences and imply positive or negative opinions. Identifying such nouns and noun phrases and their polarities is very challenging but critical for effective opinion mining in these domains. To the best of our knowledge, this problem has not been studied in the literature. This paper proposes a method to deal with the problem. Experimental results based on real-life datasets show promising results.

1 Introduction

Opinion words are words that convey positive or negative polarities. They are critical for opinion mining (Pang et al., 2002; Turney, 2002; Hu and Liu, 2004; Wilson et al., 2004; Popescu and Etzioni, 2005; Gamon et al., 2005; Ku et al., 2006; Breck et al., 2007; Kobayashi et al., 2007; Ding et al., 2008; Titov and McDonald, 2008; Pang and

Lee, 2008; Lu et al., 2009). The key difficulty in finding such words is that opinions expressed by many of them are domain or context dependent.

Several researchers have studied the problem of finding opinion words (Liu, 2010). The approaches can be grouped into corpus-based approaches (Hatzivassiloglou and McKeown, 1997; Wiebe, 2000; Kanayama and Nasukawa, 2006; Qiu et al., 2009) and dictionary-based approaches (Hu and Liu 2004; Kim and Hovy, 2004; Kamps et al., 2004; Esuli and Sebastiani, 2005; Takamura et al., 2005; Andreevskaia and Bergler, 2006; Dragut et al., 2010). Dictionary-based approaches are generally not suitable for finding domain specific opinion words as dictionaries contain little domain specific information.

Hatzivassiloglou and McKeown (1997) did the first work to tackle the problem for adjectives using a corpus. The approach exploits some conjunctive patterns, involving *and*, *or*, *but*, *either-or*, or *neither-nor*, with the intuition that the conjoining adjectives subject to linguistic constraints on the orientation or polarity of the adjectives involved. Using these constraints, one can infer opinion polarities of unknown adjectives based on the known ones. Kanayama and Nasukawa (2006) improved this work by using the idea of coherency. They deal with both adjectives and verbs. Ding et al. (2008) introduced the concept of feature context because the polarities of many opinion bearing words are sentence context dependent rather than just domain dependent. Qiu et al. (2009) proposed a method called *double propagation* that uses dependency relations to extract both opinion words and product features.

However, none of these approaches handle nouns or noun phrases. Although Zagibalov and Carroll (2008) noticed the issue, they did not study it.

Esuli and Sebastiani (2006) used WordNet to determine polarities of words, which can include nouns. However, dictionaries do not contain domain specific information.

Our work uses the feature-based opinion mining model in (Hu and Liu, 2004) to mine opinions in product reviews. We found that in some application domains product features which are indicated by nouns have implied opinions although they are not subjective words.

This paper aims to identify such opinionated noun features. To make this concrete, let us see an example from a mattress review: “*Within a month, a valley formed in the middle of the mattress.*” Here “valley” indicates the quality of the mattress (a product feature) and also implies a negative opinion. The opinion implied by “valley” cannot be found by current techniques.

Although Riloff et al. (2003) proposed a method to extract subjective nouns, our work is very different because many nouns implying opinions are not subjective nouns, but objective nouns, e.g., “valley” and “hole” on a mattress. Those sentences involving such nouns are usually also objective sentences. As much of the existing opinion mining research focuses on subjective sentences, we believe it is high time to study objective words and sentences that imply opinions as well. This paper represents a positive step towards this direction.

Objective words (or sentences) that imply opinions are very difficult to recognize because their recognition typically requires the commonsense or world knowledge of the application domain. In this paper, we propose a method to deal with the problem, specifically, finding product features which are nouns or noun phrases and imply positive or negative opinions. Our experimental results show promising results.

2 The Proposed Method

We start with some observations. For a product feature (or feature for short) with an implied opinion, there is either no adjective opinion word that modifies it directly or the opinion word that modify it usually have the same opinion.

Example 1: No opinion adjective word modifies the opinionated product feature (“valley”):

“*Within a month, a valley formed in the middle of the mattress.*”

Example 2: An opinion adjective modifies the opinionated product feature:

“*Within a month, a **bad** valley formed in the middle of the mattress.*”

Here, the adjective “bad” modifies “valley”. It is unlikely that a positive opinion word will modify “valley”, e.g., “good valley” in this context. Thus, if a product feature is modified by both positive and negative opinion adjectives, it is unlikely to be an opinionated product feature.

Based on these examples, we designed the following two steps to identify noun product features which imply positive or negative opinions:

1. *Candidate Identification:* This step determines the surrounding sentiment context of each noun feature. The intuition is that if a feature occurs in negative (respectively positive) opinion contexts significantly more frequently than in positive (or negative) opinion contexts, we can infer that its polarity is negative (or positive). A statistical test is used to test the significance. This step thus produces a list of candidate features with positive opinions and a list of candidate features with negative opinions.
2. *Pruning:* This step prunes the two lists. The idea is that when a noun product feature is directly modified by both positive and negative opinion words, it is unlikely to be an opinionated product feature.

Basically, step 1 needs the feature-based sentiment analysis capability. We adopt the lexicon-based approach in (Ding et al. 2008) in this work.

2.1 Feature-Based Sentiment Analysis

To use the lexicon-based sentiment analysis method, we need a list of opinion words, i.e., an opinion lexicon. Opinion words are words that express positive or negative sentiments. As noted earlier, there are also many words whose polarities depend on the contexts in which they appear.

Researchers have compiled sets of opinion words for adjectives, adverbs, verbs and nouns respectively, called the *opinion lexicon*. In this paper, we used the opinion lexicon compiled by Ding et al. (2008). It is worth mentioning that our task is to find nouns which imply opinions in a specific domain, and such nouns do not appear in any general opinion lexicon.

2.1.1. Aggregating Opinions on a Feature

Using the opinion lexicon, we can identify opinion polarity expressed on each product feature in a sentence. The lexicon based method in (Ding et al. 2008) basically combines opinion words in the sentence to assign a sentiment to each product feature. The sketch of the algorithm is as follows.

Given a sentence s which contains a product feature f , opinion words in the sentence are first identified by matching with the words in the opinion lexicon. It then computes an orientation score for f . A positive word is assigned the semantic orientation (polarity) score of +1, and a negative word is assigned the semantic orientation score of -1. All the scores are then summed up using the following score formula:

$$score(f) = \sum_{w_i: w_i \in s \wedge w_i \in L} \frac{w_i \cdot SO}{dis(w_i, f)}, \quad (1)$$

where w_i is an opinion word, L is the set of all opinion words (including idioms) and s is the sentence that contains the feature f , and $dis(w_i, f)$ is the distance between feature f and opinion word w_i in s . $w_i \cdot SO$ is the semantic orientation (polarity) of word w_i . The multiplicative inverse in the formula is used to give low weights to opinion words that are far away from the feature f .

If the final score is positive, then the opinion on the feature in s is positive. If the score is negative, then the opinion on the feature in s is negative.

2.1.2. Rules of Opinions

Several language constructs need special handling, for which a set of rules is applied (Ding et al., 2008; Liu, 2010). A rule of opinion is an implication with an expression on the left and an implied opinion on the right. The expression is a conceptual one as it represents a concept, which can be expressed in many ways in a sentence.

Negation rule. A negation word or phrase usually reverses the opinion expressed in a sentence. Negation words include “no,” “not”, etc.

In this work, we also discovered that when applying negation rules, a special case needs extra care. For example, “*I am not bothered by the hump on the mattress*” is a sentence from a mattress review. It expresses a neutral feeling from the person. However, it also implies a negative opinion about “*hump*,” which indicates a product feature. We call this kind of sentences *negated feeling*

response sentences. A sentence like this normally expresses the feeling of a person or a group of persons towards some items which generally have positive or negative connotations in the sentence context or the application domain. Such a sentence usually consists of four components: a noun representing a person or a group of persons (which includes personal pronoun and proper noun), a negation word, a feeling verb, and a stimulus word. Feeling verbs include “bother,” “disturb,” “annoy,” etc. The stimulus word, which stimulates the feeling, also indicates a feature. In analyzing such a sentence, for our purpose, the negation is not applied. Instead, we regard the sentence bearing the same opinion about the stimulus word as the opinion of the feeling verb. These opinion contexts will help the statistical test later.

But clause rule. A sentence containing “but” also needs special treatment. The opinion before “but” and after “but” are usually the opposite to each other. Phrases such as “except that” and “except for” behave similarly.

Decreasing and increasing rules. These rules say that decreasing or increasing of some quantities associated with opinionated items may change the orientations of the opinions. For example, “*The drug eased my pain*”. Here “pain” is a negative opinion word in the opinion lexicon, and the reduction of “pain” indicates a desirable effect of the drug. We have compiled a list of such words, which include “decrease”, “diminish”, “prevent”, “remove”, etc. The basic rules are as follows:

Decreased Neg \rightarrow Positive

E.g: “*My problem have certainly diminished*”

Decreased Pos \rightarrow Negative

E.g: “*These tires reduce the fun of driving.*”

Neg and Pos represent respectively a negative and a positive opinion word. Increasing rules do not change opinion directions (Liu, 2010).

2.1.3. Handling Context-Dependent Opinions

As mentioned earlier, context-dependent opinion words (only adjectives and adverbs) must be determined by its contexts. We solve this problem by using the global information rather than only the local information in the current sentence. We use a conjunction rule. For example, if someone writes a sentence like “*This camera is very nice and has a long battery life*”, we can infer that

“long” is positive for “battery life” because it is conjoined with the positive word “nice.” This discovery can be used anywhere in the corpus.

2.2 Determining Candidate Noun Product Features that Imply Opinions

Using the sentiment analysis method in section 2.1, we can identify opinion sentences for each product feature in context, which contains both positive-opinionated sentences and negative-opinionated sentences. We then determine candidate product features implying opinions by checking the percentage of either positive-opinionated sentences or negative-opinionated sentences among all opinionated sentences. Through experiments, we make an empirical assumption that if either the positive-opinionated sentence percentage or the negative-opinionated sentence percentage is significantly greater than 70%, we regard this noun feature as a noun feature implying an opinion. The basic heuristic for our idea is that if a noun feature is more likely to occur in positive (or negative) opinion contexts (sentences), it is more likely to be an opinionated noun feature. We use a statistic method *test for population proportion* to perform the significant test. The details are as follows. We compute the Z-score statistic with one-tailed test.

$$Z = \frac{p - p_0}{\sqrt{\frac{p_0(1 - p_0)}{n}}} \quad (2)$$

where p_0 is the hypothesized value (0.7 in our case), p is the sample proportion, i.e., the percentage of positive (or negative) opinions in our case, and n is the sample size, which is the total number of opinionated sentences that contain the noun feature. We set the statistical confidence level to 0.95, whose corresponding Z score is -1.64. It means that Z score for an opinionated feature must be no less than -1.64. Otherwise we do not regard it as a feature implying opinion.

2.3 Pruning Non-Opinionated Features

Many of candidate noun features with opinions may not indicate any opinion. Then, we need to distinguish features which have implied opinions and normal features which have no opinions, e.g., “voice quality” and “battery life.” For normal features, people often can have different opinions. For example, for “voice quality”, people can say

“good voice quality” or “bad voice quality.” However, for features with context dependent opinions, people often have a fixed opinion, either positive or negative but not both. With this observation in mind, we can detect features with no opinion by finding direct modification relations using a dependency parser. To be safe, we use only two types of direct relations:

Type1: $O \rightarrow O-Dep \rightarrow F$

It means O depends on F through a relation $O-Dep$. E.g: “This TV has a *good* picture quality.”

Type 2: $O \rightarrow O-Dep \rightarrow H \leftarrow F-Dep \leftarrow F$

It means both O and F depends on H through relation $O-Dep$ and $F-Dep$ respectively. E.g: “The springs of the mattress are *bad*.”

Here O is an opinion word, $O-Dep / F-Dep$ is a dependency relation, which describes a relation between words, and includes *mod*, *pnmod*, *subj*, *s*, *obj*, *obj2* and *desc* (detailed explanations can be found in <http://www.cs.ualberta.ca/~lindek/minipar.htm>). F is a noun feature. H means any word. For the first example, given feature “picture quality”, we can extract its modification opinion word “good”. For the second example, given feature “springs”, we can get opinion word “bad”. Here H is the word “are”.

Among these extracted opinion words for the feature noun, if some belong to the positive opinion lexicon and some belong to the negative opinion lexicon, we conclude the noun feature is not an opinionated feature and is thus pruned.

3 Experiments

We conducted experiments using four diverse real-life datasets of reviews. Table 1 shows the domains (based on their names) of the datasets, the number of sentences, and the number of noun features. The first two datasets were obtained from a commercial company that provides opinion mining services, and the other two were crawled by us.

Product Name	Mattress	Drug	Router	Radio
# Sentences	13191	1541	4308	2306
# Noun features	326	38	173	222

Table 1. Experimental datasets

An issue for judging noun features implying opinions is that it can be subjective. So for the gold standard, a consensus has to be reached between the two annotators.

For comparison, we also implemented a baseline method, which decides a noun feature’s polarity only by its modifying opinion words (adjectives). If its corresponding adjective is positive-orientated, then the noun feature is positive-orientated. The same goes for a negative-orientated noun feature. Then using the same techniques in section 2.3 for statistical test (in this case, n in equation 2 is the total number of sentences containing the noun feature) and for pruning, we can determine noun features implying opinions from the data corpus.

Table 2 gives the experimental results. The performances are measured using the standard evaluation measures of precision and recall. From Table 2, we can see that the proposed method is much better than the baseline method on both the recall and precision. It indicates many noun features that imply opinions are not directly modified by adjective opinion words. We have to determine their polarities based on contexts.

Product Name	Baseline		Proposed Method	
	Precision	Recall	Precision	Recall
Mattress	0.35	0.07	0.48	0.82
Drug	0.40	0.15	0.58	0.88
Router	0.20	0.45	0.42	0.67
Radio	0.18	0.50	0.31	0.83

Table 2. Experimental results for noun features

Table 3 and Table 4 give the results of noun features implying positive and negative opinions separately. No baseline method is used here due to its poor results. Because for some datasets, there is no noun feature implying a positive/negative opinion, their precision and recall are zeros.

Product Name	Precision	Recall
Mattress	0.42	0.95
Drug	0.33	1.0
Router	0.43	0.60
Radio	0.38	0.83

Table 3. Features implying positive opinions

Product Name	Precision	Recall
Mattress	0.56	0.72
Drug	0.67	0.86
Router	0.40	1.00
Radio	0	0

Table 4. Features implying negative opinions

From Tables 2 - 4, we observe that the precision of the proposed method is still low, although the recalls are good. To better help the user find such

words easily, we rank the extracted feature candidates. The purpose is to rank correct noun features that imply opinions at the top of the list, so as to improve the precision of the top-ranked candidates. Two ranking methods are used:

1. rank based on the statistical score Z in equation 2. We denote this method with Z-rank.
2. rank based on negative/positive sentence ratio. We denote this method with R-rank.

Tables 5 and 6 show the ranking results. We adopt the rank precision, also called the *precision@N*, metric for evaluation. It gives the percentage of correct noun features implying opinions at the rank position N . Because some domains may not contain positive or negative noun features, we combine positive and negative candidate features together for an overall ranking for each dataset.

	Mattress	Drug	Router	Radio
Z-rank	0.70	0.60	0.60	0.70
R-rank	0.60	0.60	0.50	0.40

Table 5. Experimental results: Precision@10

	Mattress	Drug	Router	Radio
Z-rank	0.66		0.46	0.53
R-rank	0.60		0.46	0.40

Table 6. Experimental results: Precision@15

From Tables 5 and 6, we can see that the ranking by statistical value Z is more accurate than negative/positive sentence ratio. Note that in Table 6, there is no result for the Drug dataset because no noun features implying opinions were found beyond the top 10 results because there are not many such noun features in the drug domain.

4 Conclusions

This paper proposed a method to identify noun product features that imply opinions. Conceptually, this work studied the problem of objective nouns and sentences with implied opinions. To the best of our knowledge, this problem has not been studied in the literature. This problem is important because without identifying such opinions, the recall of opinion mining suffers. Our proposed method determines feature polarity not only by opinion words that modify the features but also by its surrounding context. Experimental results show that the proposed method is promising. Our future work will focus on improving the precision.

References

- Andreevskaia, A. and S. Bergler. 2006. Mining WordNet for fuzzy sentiment: Sentiment tag extraction from WordNet glosses. *Proceedings of EACL 2006*.
- Eric Breck, Yejin Choi, and Claire Cardie. 2007. Identifying Expressions of Opinion in Context. *Proceedings of IJCAI 2007*.
- Xiaowen Ding, Bing Liu and Philip S. Yu. 2008 A Holistic Lexicon-Based Approach to Opinion Mining. *Proceedings of WSDM 2008*.
- Eduard C. Dragut, Clement Yu, Prasad Sistla, and Weiyi Meng. 2010. Construction of a sentimental word dictionary. In *Proceedings of CIKM 2010*. Andrea Esuli and Fabrizio Sebastiani. 2005. Determining the Semantic Orientation of Terms through Gloss Classification. *Proceedings of CIKM 2005*.
- Andrea Esuli and Fabrizio Sebastiani. 2006. SentiWorkNet: A Publicly Available Lexical Resource for Opinion Mining. *Proceedings of LREC 2006*.
- Michael Gamon. 2004. Sentiment Classification on Customer Feedback Data: Noisy Data, Large Feature Vectors and the Role of Linguistic Analysis. *Proceedings of COLING 2004*.
- Murthy Ganapathibhotla. and Bing Liu. 2008. Mining opinions in comparative sentences. *Proceedings of COLING 2008*.
- Vasileios Hatzivassiloglou and Kathleen, McKeown. 1997. Predicting the Semantic Orientation of Adjectives. *Proceedings of ACL 1997*.
- Minqing Hu and Bing Liu. 2004. Mining and Summarizing Customer Reviews. *Proceedings of KDD 2004*.
- Jaap Kamps, Maarten Marx, Robert J. Mokken and Maarten de Rijke. 2004. *Proceedings of LREC 2004*.
- Hiroshi Kanayama, Tetsuya Nasukawa 2006. Fully Automatic Lexicon Expansion for Domain-Oriented Sentiment Analysis. *Proceedings of EMNLP 2006*.
- Nozomi Kobayashi, Kentaro Inui, and Yuji Matsumoto. 2007. Extracting aspect-evaluation and aspect-of relations in opinion mining. *Proceedings of EMLP 2007*
- Soo-Min Kim and Eduard Hovy. 2004. Determining the Sentiment of Opinions. *Proceedings of COLING 2004*.
- Lun-Wei Ku, Yu-Ting Liang, and Hsin-Hsi Chen. 2006. Opinion extraction, summarization and tracking in news and blog corpora. *Proceedings of AAAI-CAAW 2006*.
- Bing Liu. 2010. Sentiment analysis and subjectivity. A chapter in *Handbook of Natural Language Processing*, Second edition.
- Yue Lu, Chengxiang Zhai, and Neel Sundaresan. 2009. Rated aspect summarization of short comments. *Proceedings of WWW 2009*.
- Bo Pang and Lillian Lee. 2008. Opinion Mining and sentiment Analysis. *Foundations and Trends in Information Retrieval 2(1-2)*, 2008.
- Bo Pang, Lillian Lee and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. *Proceedings of EMNLP 2002*.
- Ana-Maria Popescu and Oren Etzioni. 2005. Extracting Product Features and Opinions from Reviews. *Proceedings of EMNLP 2005*.
- Guang Qiu, Bing Liu, Jiajun Bu and Chun Chen. 2009. Expanding Domain Sentiment Lexicon through Double Propagation. *Proceedings of IJCAI 2009*.
- Ellen Riloff, Janyce Wiebe, and Theresa Wilson. 2003. Learning subjective nouns using extraction pattern bootstrapping. *Proceedings of CoNLL 2003*.
- Hiroya Takamura, Takashi Inui and Manabu Okumura. 2007. Extracting Semantic Orientations of Phrases from Dictionary. *Proceedings of HLT-NAACL 2007*.
- Ivan Titov and Ryan McDonald. 2008. A joint model of text and aspect ratings for sentiment summarization. In *Proceedings of ACL 2008*. Peter D. Turney. 2002. Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews. *Proceedings of ACL 2002*.
- Janyce Wiebe. 2000. Learning Subjective Adjectives from Corpora. *Proceedings of AAAI 2000*.
- Theresa Wilson, Janyce Wiebe, Rebecca Hwa. 2004. Just how mad are you? Finding strong and weak opinion clauses. *Proceedings of AAAI 2004*.
- Taras Zagibalov and John Carroll. 2008. Unsupervised Classification of Sentiment and Objectivity in Chinese Text. *Proceedings of IJCNLP 2008*.