Using Data Mining Techniques for Sentiment Shifter Identification

Samira Noferesti, Mehrnoush Shamsfard
University of Sistan and Baluchestan, Shahid Beheshti University
Iran
E-mail: snoferesti@ece.usb.ac.ir, m-shams@sbu.ac.ir

Abstract

Sentiment shifters, i.e., words and expressions that can affect text polarity, play an important role in opinion mining. However, the limited ability of current automated opinion mining systems to handle shifters represents a major challenge. The majority of existing approaches rely on a manual list of shifters; few attempts have been made to automatically identify shifters in text. Most of them just focus on negating shifters. This paper presents a novel and efficient semi-automatic method for identifying sentiment shifters in drug reviews, aiming at improving the overall accuracy of opinion mining systems. To this end, we use weighted association rule mining (WARM), a well-known data mining technique, for finding frequent dependency patterns representing sentiment shifters from a domain-specific corpus. These patterns that include different kinds of shifter words such as shifter verbs and quantifiers are able to handle both local and long-distance shifters. We also combine these patterns with a lexicon-based approach for the polarity classification task. Experiments on drug reviews demonstrate that extracted shifters can improve the precision of the lexicon-based approach for polarity classification 9.25 percent.

Keywords: sentiment analysis, sentiment shifters, association rule mining

1. Introduction

Opinion mining, also called sentiment analysis, is the task of extracting and analyzing opinions, sentiments, evaluations or feelings from user-generated contents such as reviews, discussion groups and blogs. Due to its wide range of applications such as analysis of customer reviews (Hu & Liu, 2004) and reputation management (Wiegand et al., 2010), this field has received considerable attention both in industrial and academic research areas.

One of the main subtasks of opinion mining is polarity classification, which aims at classifying opinions into predefined classes (usually positive and negative). Existing approaches to polarity classification can be grouped into two main categories: lexicon-based and machine learning approaches. Lexicon-based approaches mainly rely on linguistic resources containing polar terms and concepts such as SentiWordNet (Esuli & Sebastiani, 2006), General Inquirer (Stone et al., 1966) and Subjectivity Lexicon (Wilson et al., 2009). For example, “This drug is amazing.” is a positive sentence since the term “amazing” is positive in sentiment lexicons. However, these resources are not sufficient since the polarity classification is a challenging task that needs to tackle many subtle phenomena such as sentiment shifters.

Sentiment shifters, also called valence shifters, are words and expressions that affect the polarity of an opinion by changing its magnitude or its direction. For example, in the sentence “I do not like this drug.”, the shifter word “not” before the positive word “like” changes the text polarity to negative. Therefore, ignoring sentiment shifters can lead to noticeable decline in overall accuracy of opinion mining systems.

There are two types of shifter words or shifter trigger words: words that reverse the polarity of the given text (e.g., “no” and “never”), and words that change sentiment values by a constant amount (e.g., “severe” and “mild”). In this paper, we only focus on the first type, i.e., reversing words. Reversing words are not limited to negation words. Some kinds of verbs (like “reduce”) and quantifiers (like “less”) can act as sentiment shifters.

From another perspective, sentiment shifters can be classified into two main groups: local shifters indicating shifter words which are directly applied to polar words (e.g., “Accutane doesn’t help”), and long-distance shifters which allow longer distance dependencies between the shifter words and the polar words (e.g., “No one likes this drug”).

Although sentiment shifter identification plays a fundamental role in recognizing polarity of textual expressions, it has not been completely solved. Existing approaches to shifter identification can be classified into two main categories: rule-based and machine learning approaches. Applying shifter rules is a two-step procedure: identifying shifter words and determining their scope (i.e., part of the sentence that is affected by the shifter). For the first step, most of previous works have mainly relied on a list of common shifter words which is built manually (Huang et al., 2014; Marrese-Taylor et al., 2014). The main limitation of these methods is that such lists in any language may be incomplete, and hence, there is always a need to propose a way to deal with words that are not in the lists. Furthermore, due to the language dependency nature of shifter words, it is difficult to adapt these lists to other languages.

For the second step, some researchers have proposed simple heuristic rules that define the scope of a shifter word using a window of fixed size (Hu & Liu, 2004; Heerschop et al., 2011; Huang et al., 2014). In (Shaikh et al., 2007; Asmi & Ishaya, 2012) the scope of negation has been identified by using dependency tree which indicates how a negation word interacts by other words of the sentence.

Some researchers have used shifter words and their scopes as a feature for polarity classification using machine learning approaches (Pang et al., 2002; Kennedy & Inkpen, 2006; Jia et al., 2009; Wilson et al., 2009; Morante & Blanco, 2012). These approaches can capture some aspects of the shifters effectively. However, they depend upon the availability of an annotated corpus in which shifter words and their scopes are tagged. Manual construction of such corpora is a tedious, expensive and time consuming task.

This paper presents a novel and efficient semi-automatic approach to extract sentiment shifter patterns from a domain-specific corpus of polarity-tagged sentences. The proposed approach is based on dependency relations between words of a sentence. It is able to handle both local and long-distance shifter words. It extracts patterns for different kinds of shifter words such as negation structures (e.g., "no" and "not"), shifter verbs (e.g., "decrease", and "eliminate"), and shifter quantifiers, i.e., words which express a decreased/increased value of quantity (e.g., "less"), while most of the existing approaches just focus on negation words. In addition, the proposed approach is language-independent and hence, although we tested it on English, it can be used for other languages as well. We also incorporate the extracted patterns into a lexicon-based method for polarity classification. Experimental results show that the extracted shifters improve the performance of the lexicon-based method.

2. The Proposed Approach

As mentioned before, sentiment shifters can reverse the polarity of the given text, and hence are vital for polarity classification. In particular, in the drug domain, most of medical terms such as "pain" and "depression" are negative, but they occur frequently in positive sentences. Consider the following examples,

"Accutane eliminated my cystic acne."
"It reduced my pain."
"No pain"
"Less acne"

We assume that when the polarity of a sentence is different from the polarity of the majority of its words, that sentence may have a valence shifter. In the above examples, the valence shifters "eliminated", "reduced", "no" and "less" invert the polarity of the corresponding polar terms. Therefore, capturing such shifters will improve the performance of polarity classification. Given this insight, the idea behind our approach is to extract frequent dependency patterns from such sentences as shifter patterns.

Our proposed approach for sentiment shifter identification consists of two steps: candidate sentence extraction and frequent shifter pattern mining. In the first step, we extract candidate sentences from a domain-specific corpus. Then, in the second step, we mine frequent dependency patterns in a set of candidate sentences. In the following subsections, we describe each of these steps in detail.

2.1 Extracting Candidate Sentences

To extract candidate sentences for shifter identification we use a corpus of polarity-tagged sentences. Here, P and N denote the sets of positive and negative sentences of the corpus, respectively. We divide each set into two subsets: positive sentences with negative terms (PN), positive sentences with positive terms (PP), negative sentences with negative terms (NN), and negative sentences with positive terms (NP). If a sentence includes both positive and negative terms, we use term-counting method. For example, PN includes positive sentences which have more negative terms than positive ones. This is the case for other sets as well. To determine polar terms we use a sentiment lexicon of 5330 words in medical domain which was built manually. Finally, we select PN and NP sets as candidate sets, i.e., sentences including sentiment shifters.

2.2 Mining Frequent Dependency Patterns

The second step is to extract dependency relations that appear significantly more frequently in the PN (or NP) than other sets. The idea is that these relations represent shifter patterns since they are frequent in sentences with shifters but not frequent in sentences without shifters.

In order to extract shifter patterns, we use weighted association rule mining (WARM). ARM is one of the key data mining techniques that have been used to tackle a variety of applications (Agraval & Srikant, 1994). ARM consists of two subtasks. The first subtask is frequent itemset mining which generates all items whose supports are higher than a predefined threshold called minimum support. The second subtask generates association rules which satisfy the minimum support and minimum confidence thresholds. WARM generalizes the classical model to the case where different items have different weights to reflect their different importance. To extract shifter patterns, we present a two-step procedure. First, we extract important dependency relations of the sentences, and then we adopt WARM to find frequent shifter patterns.

2.2.1. Extracting Important Dependency Relations of the Sentences

To extract important dependency relations of a sentence, we perform the following steps:

- Extracting dependency relations of the sentence: For each sentence, a set of dependency relations is obtained from the Stanford dependency parser[1]. Each dependency relation represents a relation between two words. We show a dependency relation with a triplet *r(relation-name, word$_1$, word$_2$)*.
- Removing less important relations: Less important relations i.e., the dependency relations containing very common words (stop words) are stripped out.
- Stemming: For each remaining dependency relation, we use Stanford stemmer to reduce different forms of a word to one canonical form.
- Assigning word classes: For each dependency relation, we replace the polar words involved in the dependency with their classes. The class indicates the part of speech (POS) tag and the polarity of that word. For example, the class "A_POS" is assigned to positive adjectives like "good". In this way, we generalize the dependency relations and as a result the extracted shifter patterns. Generalized shifter patterns have higher coverage than specific ones, and so have a greater chance of matching a context.
- Assigning weights: In this step, each sentence is described by a set of dependency relations which is represented as a vector $v = \{(r_1, w_1), (r_2, w_2), \ldots, (r_m, w_m)\}$, where r_i is a dependency relation and w_i is its weight. m is the number of dependency relations in the sentence. The weights can be determined in a number of ways. In this paper, we simply use one of the most widely used weighting approaches called TF-IDF [2]. In this approach, the weight of relation r_i in sentence j is defined as follows,

$$w_{ij} = tf_{ij} . \log_2(\frac{n}{df_{ij}})$$

[1] http://nlp.stanford.edu/software/corenlp.shtml
[2] Term frequency-inverse document frequency

Where, tf_{ij} is the number of occurrences of relation r_i in the PN (or NP), df_{ij} is the total number of occurrences of relation r_i and n is the total number of sentences. In fact, TF-IDF is intended to reflect the importance of a dependency relation as a shifter pattern.

2.2.2. Extracting Shifter Patterns

In this step, we use an Apriori-like algorithm to explore frequent weighted relations (Zhang & Zhang, 2002). Apriori is a well-known algorithm for ARM (Agraval & Srikant, 1994). Given a set of transactions, where each transaction is a set of items, Apriori algorithm aims to find frequent item sets, i.e., item sets whose occurrences are greater than a user-specified minimum support. In the first step, Apriori finds the frequent individual items, and in each next step, it extends each subset with one item at a time to generate frequent groups of items.

We use a modified implementation of an Apriori-like method to mine frequent shifter patterns. Dependency relations and sentences become "items" and "transactions", respectively, in the frequent item set mining framework. The first scan finds weighted frequent individual dependency relations whose supports (weights) are greater than the minimum support threshold. In the first scan, we impose a restriction; we only mine weighted frequent individual dependency relations that contain at least one polar word. This set is called 1-relation set. Each subsequent scan k starts with the set of frequent relation sets found in the previous scan. This set is used to generate a set of new potential shifter patterns. Candidates whose weights are greater than the threshold form the set of newly found shifter patterns, called k-relation set. The algorithm terminates when no candidate relation set can be generated or no candidate pattern can be found.

Among extracted frequent patterns (relation sets), we only select those whose confidences are higher than a threshold, called minimum confidence. The confidence of a pattern presents its accuracy, i.e., the ratio of correct shifters detected by this pattern in a set of instances matching it. The confidence is computed as follows:

$$\text{confidence} = \frac{\text{No. of correctly detected instances}}{\text{No. of instances match the pattern}}$$

We employ particle swarm optimization (PSO) (Kennedy, 2010) to adjust the values of minimum support and minimum confidence parameters. In the area of association rule mining, PSO is successfully used for determination of these threshold values (Kuo et al., 2011). PSO tries to find the best values with which we gained the best performance in shifter identification on the development set. In this way, we can have different values for minimum supports in each iteration. Finally, the extracted frequent relation sets represent shifter patterns.

2.3 Incorporating Shifter Patterns into Lexicon-based Approaches for Polarity Classification

To incorporate the extracted shifter patterns into a lexicon-based approach for polarity classification, we first tag the polarity of the given sentence using a sentiment lexicon. Then we parse the sentence with the Stanford parser and make the vector of dependency relations for it. For each word in a dependency relation, we perform lemmatization and assign corresponding word classes to polar words. Finally, if the vector of dependency relations matches with a shifter pattern the polarity of the sentence will be reversed.

3. Experiments

To extract candidate sentences for shifter identification we used a corpus of polarity-tagged sentences which were collected from the *www.druglib.com* website. This corpus contains 2776 reviews for 85 drugs. Sometimes different parts of a compound sentence have different polarities. Thus, to achieve more accurate shifter patterns, compound sentences were broken down into simple units. Splitting was done by exploiting dependency tree and conjunction structure of the sentence (De Marneffe et al., 2006). Applying the proposed approach for shifter identification, we extracted a set of 826 shifter patterns.

Table 1 depicts some examples of the extracted shifter patterns. The second column of Table 1 shows a shifter pattern. The third column illustrates an example sentence for each shifter pattern. As can be seen from Table 1, the proposed approach is able to handle some kinds of shifter verbs (e.g., "reduce" and "go away" in examples 2 and 6, respectively). Likewise, there are some patterns (e.g., example 7) to detect shifter quantifiers (e.g., "less"). Furthermore, the proposed approach can detect some kinds of long-distance shifters (e.g., examples 3 and 5).

No.	Shifter patterns	Example sentences or phrases
1	{det(n_neg,any), dobj(experience,n_neg), neg(experience,not)}	I did not experience any side effects from taking this medication.
2	{nsubjpass(reduce,n_neg), auxpass(reduce,be)}	My leg pain was reduced.
3	{nsubj(v_pos,one), det(one,no)}	No one likes this drug.
4	{prep(lack,n_pos)}	Lack of energy
5	{advmod(n_neg,longer), advmod(longer,no)}	I no longer have panic attacks.
6	{nsubj(go,n_neg), prt(go,away)}	Also joint swelling and pain in my legs have gone away
7	{amod(n_neg,less), dobj(have,n_neg)}	I have less stress

Table 1: Examples of shifter patterns

In addition, as can be seen in Table 1, most of the extracted patterns are not domain-specific. Thus, although the shifter identification method was tested on drug review domain, the extracted patterns can be used in any other domain as well. However, in order to have a more general pattern set, we can extract shifter patterns from several domains for those polarity-tagged corpora are available. Therefore, as future work, we would like to address the issues of extracting general patterns (i.e., domain-independent patterns) for shifter identification and provide an efficient method to alleviate these issues.

To assess the effectiveness of incorporating shifter patterns into lexicon-based methods for polarity classification, we first determined the polarity of each sentence in a test set of 1500 sentences which were collected from *www.druglib.com* and *www.askapatient.com* review sites, using two sentiment lexicons: a domain-specific and a general-purpose (i.e., SentiWordNet) lexicon. Then, if that sentence matched a shifter pattern, its polarity would be inverted. Table 2 illustrates the performance of the proposed approach and compares it with lexicon-based methods without shifter identification. As we might expect, including shifter patterns has the marked effect on the performance of polarity classification.

Method	Precision (%)
Using SentiwordNet	58
Using domain-specific lexicon	61.05
The proposed approach	67.25

Table 2: Comparison of the precision of the proposed approach for polarity classification with lexicon-based approaches

In the next experiment, we compared the performance of the proposed approach with four methods of shifter identification (Figure 1): a baseline method, where each appearance of valence shifters inverts the polarity of text, NegEx algorithm (Chapman et al., 2001), using a window of fixed size (Huang et al., 2014), and a rule-based approach (Pang et al., 2002). NegEx, is a negation detection algorithm in biomedical texts that is based on regular expressions and a dictionary of medical terms. NegEx usually correctly detects negated terms; however, it is not able to detect other kinds of shifters such as shifter quantifiers and shifter verbs.

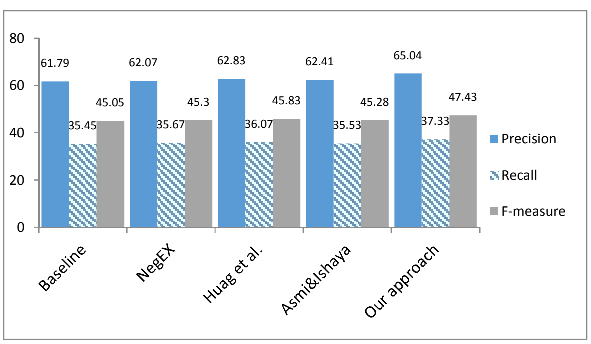

Figure 1: Comparison of the proposed approach with other methods

Figure 1 shows that the proposed approach for shifter identification outperforms NegEx algorithm (a domain-specific method) and general purpose methods.

4. Conclusion

This paper proposed a novel semi-automatic method for sentiment shifter identification. First, we employed data mining techniques to mine sentiment shifter patterns from a domain-specific corpus of polarity-tagged sentences. These patterns include different kinds of sentiment shifters such as negation structures (e.g., “no”, “not”, "no longer"), shifter verbs (e.g., “reduce”, “eliminate”) and quantifiers (e.g., “less”), and can detect both local and long-distance shifters. Then we incorporated the extracted shifter patterns into lexicon-based approaches for polarity classification. Experimental results showed that the proposed approach improves the performance of lexicon-based approaches significantly.

Furthermore, we compared the performance of the proposed approach on a dataset of drug reviews to that of a baseline method and other approaches for shifter identification. Experimental results indicated that the proposed approach outperforms other methods. Although the shifter identification method was tested on drug review domain, its results can be used in any other domain as well.

We concluded that our approach is appropriate for sentiment shifter identification, although extra knowledge is required to increase the performance. Therefore, future work aims at extracting lexico-semantic shifter patterns by replacing medical entities with their semantic classes.

We also plan to present a new and efficient method for extracting domain-independent shifter patterns.

5. References

Agraval, R., Srikant, R. (1994). Fast algorithms for mining association rules in large data bases. *20th Inter-national Conference on Very Large Databases*, Santiagom, vol. 1215, pp. 487--499.

Asmi, A., Ishaya, T. (2012). Negation Identification and Calculation in Sentiment Analysis. *IMMM 2012, The Second International Conference on Advances in Information Mining and Management*, pp. 1--7.

Chapman, W.W., Bridewell, W., Hanbury, P., Cooper, G.F., Buchanan, B.G. (2001). A simple algorithm for identifying negated findings and diseases in discharge summaries. *Journal of biomedical informatics*, 34(5), pp. 301--310.

De Marneffe, M.-C., MacCartney, B., Manning, C.D. (2006). Generating typed dependency parses from phrase structure parses. *Proceedings of LREC*, 6(2006), pp. 449--454.

Esuli, A., Sebastiani, F. (2006). Sentiwordnet: A publicly available lexical resource for opinion mining. *Proceedings of LREC*, vol. 6, pp. 417--422.

Heerschop, B., Iterson, P.van, Hogenboom, A., Frasincar, F., Kaymak, U. (2011). Analyzing sentiment in a large set of web data while accounting for negation. *Advances in Intelligent Web Mastering–3*, Springer: pp. 195--205.

Hu, M., Liu, B. (2004). Mining and summarizing customer reviews. *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 168--177.

Huang, S., Niu, Z., Shi, C. (2014). Automatic construction of domain-specific sentiment lexicon based on constrained label propagation. *Knowledge-Based Systems*, 56, pp. 191--200.

Kuo, R.J., Chao, C.M., Chiu, Y. (2011). Application of particle swarm optimization to association rule mining. *Applied Soft Computing*, 11(1), pp. 326--336.

Marrese-Taylor, E., Velásquez, J.D., Bravo-Marquez, F. (2014). A novel deterministic approach for aspect-based opinion mining. *Expert Systems with Applications*, 41(17), pp. 7764--7775.

Pang, B., Lee, L., Vaithyanathan, S. (2002). Thumbs up?: sentiment classification using machine learning techniques. *Proceedings of the ACL-02 conference on Empirical methods in natural language processing*,vol. 10, pp. 79--86.

Shaikh, M.A.M., Prendinger, H., Mitsuru, I. (2007). Assessing sentiment of text by semantic dependency and contextual valence analysis. *Affective Computing and Intelligent Interaction*, Springer, pp. 191--202.

Stone, Ph.J., Dunphy, D.C., Smith, M.S. (1966). The General Inquirer: A Computer Approach to Content Analysis.

Wiegand, M., Balahur, A., Roth, R., Klakow, D., Montoyo, A. (2010). A survey on the role of negation in sentiment analysis. *Proceedings of the workshop on negation and speculation in natural language processing*, Association for Computational Linguistics, pp. 60--68.

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.