Processing Sentiments and Opinions in Text: A Survey

Gong Tianxia
Department of Computer Science,
School of Computing,
National University of Singapore
gongtian@comp.nus.edu.sg

Abstract

Much of the current research is now focusing on the area of sentiment analysis as the amount of internet documents (especially online reviews like product reviews, movie reviews and etc.) are increasing rapidly. In sentiment analysis, the usual tasks include sentiment word or phrase identification, sentiment orientation identification and sentiment sentence or document classification. Sentiment analysis are helpful in a wide range of applications, including message filtering (Spertus, 1997), product review summarization (Hu and Liu, 2004a), movie review recommender systems (Pang et al., 2002; Terveen et al., 1997; Tatemura, 2000). Some challenging aspects of sentiment analysis exist in use of other types of words, sentiment lexicon construction, dealing with negation expressions, degrees of sentiment, complexity of sentence/document, words in different contexts, heterogenous text sentiment classification etc.

1 Introduction

With the rapid growth of internet documents, especially online reviews (product reviews, movie reviews and etc.), much of the current research has been focused on the area of sentiment analysis, also known as opinion mining, which is a recent discipline involving information retrieval (IR), natural language processing (NLP) and computational linguistics. In sentiment analysis, the system often tries to find out the words or phrases that indicate the sentiment, determine to orientation or polarity of the sentiment words or phrases, and then classify the sentence, the paragraph, or the whole document involving the processed sentiment based on its orientation. Therefore, the usual tasks for sentiment analysis include sentiment word or phrase identification, sentiment orientation identification and sentiment sentence or document classification. Different from genre classification or topic classification, sentiment classification requires more understanding of the text in the sentiment sense. The main challenging aspects of sentiment analysis exist in use of other types of words, sentiment lexicon construction, dealing with negation expressions, degrees of sentiment, complexity of sentence/document, words in different contexts, heterogenous text sentiment classification etc. Some approaches have been proposed to solve these issues; however, more future research could be dedicated to these challenges.

2 Aspects of Current Research

From a few years ago, sentiment mining has become a hot subject among NLP and IR researchers. Large amount of effort has been put into the research on this field that quite a number of papers have been published and systems for various applications using sentiment mining have been developed and put into commercial use. Though the works on sentiment mining all have different focuses, emphasizes and objectives; nevertheless, they can generally be categorized into three main aspects: sentiment words or phrases identification, sentiment orientation identification and sentiment sentence or document classification.

2.1 Sentiment Words/Phrases Identification

Sentiment words or phrases are those that are primarily used to express the write's sentiment, emotion and subjective opinion on the matter he or she is discussing about. Sentiment words or phrases extraction, also called opinion words or phrases identification by some papers, are usually the first phase in a sentiment mining project, as these words or phrases are the keys for later sentiment orientation identification and classification.

As current work on sentiment analysis focus on content words (nouns, verbs, adjectives and adverbs), most of the work use part-of-speech (POS) tagging to extract them (Hu and Liu, 2004a; Turney, 2002). POS tagging is also used for word sense disambiguation, which can help to identify the sentiment orientation of the words/phrases at later stage. Other natural language processing technique like stopwords removal, stemming and fuzzy matching are also used in the preprocessing stage to extract sentiment words/phrases.

2.2 Sentiment Words/Phrases Orientation Identification

The semantic orientation of a word, also called polarity sometimes, indicated the direction of the word deviates from the norm for its sematic group. In sentiment analysis, the orientation of the sentiment words or phrases directly shows the sentiment or opinion of the text writer. The main approaches to identify the semantic orientation of a sentiment word/phrases are statistical-based or lexicon-based.

2.2.1 Using adjectives and adverbs

Current works to identify sentiment words and their sentiment orientation or polarity are mainly focused on adjectives and adverbs, as they are often considered as the most obvious indication of subjectivity (Hatzivassiloglou and McKeown, 1997; Bruce and Wiebe, 1999; Wiebe and Bruce, 1999).

Hu and Liu (2004a) apply POS tagging and some natural language processing techniques to the text as discussed in Subsection 2.1 to extract the adjectives as sentiment words. Then they identify the features of the product reviews as their approach focus on feature-based sentiment classification. They define a sentence containing one or more product features and one or more opinion words as an opinion sentence. Then for each sentence in the review database, they extract all the adjective words as opinion words. And for each feature in the sentence, they label the nearby adjective (the adjacent adjective that modifies the noun/noun phrase that is a frequent feature) as its effective opinion. The experiment result of their opinion sentence extraction has a precision of 64.2% and a recall of 69.3%.

They use WordNet (Fellbaum, 1998) to determine whether the extracted adjective has a positive or negative polarity. In WordNet, adjectives are organized into bipolar clusters, where The second half cluster is headed by the antonym of the first cluster. Each half cluster is headed by a head synset, followed by the satellite synsets, which represent sematic meanings that are similar to those of the head adjective. In contrast to the lexicon-based approaches, they used the semantic orientation of synonyms and antonyms to predict the orientation of the adjectives. They start with a seed list which consists of 30 manually selected common adjectives. Then they use WordNet to predict the orientation of all the adjectives in the extracted sentiment word list by searching through its bipolar cluster to find out whether its synonyms or antonyms are in the seed list or not. Once the adjective's orientation is predicted, it is added to the seed list and can be used to determine other adjectives' orientation. In the process, the seed list gradually grows as more adjectives' orientation have been identified; and when it stops grows, i.e. the size of seed list is the same as the size of opinion word list, then all adjectives' orientation have been identified and the process ends.

Hatzivassiloglou and McKeown (1997) used a method to automatically retrieve semantic orientation information using indirect information collected from a large corpus as they pointed out that dictionaries such as WordNet do not include sematic orientation information and there lacks direct association between antonyms and synonyms especially when they are domain dependent. They first extract all conjunctions of adjectives are extracted from the corpus with relevant morphological relations. Then they use a log-linear regression model and combines the information from the different conjunctions to determine if each two conjoined adjectives are of same or different orientation. The adjectives are represented in a graph with hypothesized same or different orientation links and are then separated into two subsets of different orientation using a clustering algorithm. Lastly they compare the average frequencies in each adjective group and label the higher frequency group as positive.

Turney (2002) used mutual information between two words to classify the adjective or adverb's orientation in reviews of different domains. Prior to word sentiment classification, they use POS tagging to extract adjectives and adverbs based on Brill (1994)'s approach. The pointwise mutual information (PMI) (Church and Hanks, 1989; Turney, 2001) which is a measure of the strength of semantic association between two words is used. The semantic orientation of a word/phrase x is then calculated as PMI(x, "excellent")−PMI(x, "poor")–the word/phrase is classified as positive if it is more strongly associated with "excellent", and negative otherwise. They chose the words "excellent" and "poor" because these two words are commonly used to express the two ends of sentiments in reviews.

2.2.2 Using Verbs

Adjectives and adverbs play an important role in sentiment analysis and have been the dominant type of words for sentiment words extraction and orientation identification in current research; however, other word types, e.g. verbs, are also used to express emotions and opinions in texts.

Besides adjectives and adverbs, Nasukawa and Yi (2003) also consider verbs as sentiment expressions for their sentiment analysis. As they point out that the polarity of sentiment denoted by sentiment verbs may depend on relationships with their arguments in contrast to adjectives, adverbs or nouns, they classify sentiment-related verbs into two types. The first type directs either positive or negative sentiment toward their arguments like "beat" in "X beats Y". The second type does not denote sentiment by themselves, but transfers sentiments among their arguments like "is" in "X is good".

They use Markov-model-based POS tagger (Manning and Schutze, 1999) and rule-based shallow parsing (Neff et al., 2003) for preprocessing. Then they analyze the syntactic dependencies among the phrases and look for phrases with a sentiment term that modifies or is modified by a subject term. The polarity is identified by looking into the definition in the sentiment dictionary when the sentiment term is a verb.

2.2.3 Using Collocations

Collocations are also an important type of subjectivity clue (Wiebe et al., 2001). Their approach shows a straightforward method for learning potentially subjective collocations from corporate. The method is first applied directly to the words in the text, then applied again to the same data but with all words that appear just once replaced by a single word (UNIQUE). The two applications yield different results, and both show increased classification accuracy. They also show consistency in performance with subjectivity clues identified in previous work. Together with clues in previous work of collocation extraction, their approach shows promise for recognizing opinionated documents. In addition to classifying editorials in news reports that are not previously marked with opinions, the method can also be applied to filtering the results of search engines. Although many potentially subjective collocations are not identified by their approach; however, their work have laid a foundation in this aspects for future research. Similar works on sentiment orientation identification using collocations also include Tong (2001)'s approach.

2.3 Sentiment Sentence/Document Classification

Sentiment sentence or document classification is to classify a sentence or document according to its polarity into different sentiment categories–positive or negative, the category neutral is also added sometimes. Sentiment sentence classification is useful when each individual sentence's polarity is important, where as sentiment document classification is useful when the general opinion tendency of the overall document is emphasized. Sentiment word/phrase orientation identification is used for sentence/document classification as in (Hu and Liu, 2004a), whereas other works (Pang et al., 2002) classifies sentiment sentences/documents without the knowledge of each sentiment words.

Hu and Liu (2004a) predict the orientation of the opinion sentence in their study of customer reviews.

For their work is focused on the overall opinion for a particular product feature, the sentiment classification for each individual opinion sentence comes to use as they need to group the opinion sentences for the same feature to find out the overall opinion towards that product feature. They use the dominant orientation of the opinion words by counting them according to their polarity to determine the orientation of the sentence. So if the number of positive/negative opinion words exceeds the other one, then the sentence is classified as positive/negative. In case of a draw, the average orientation of closest opinion word for the product feature or the orientation to the previous opinion sentence is used for classification. Their sentence orientation accuracy is 84.2%.

Turney (2002) used a simple unsupervised algorithm to classify reviews in different domains as recommended or not recommended. and then do sentiment words/phrase extraction based on Hatzivassiloglou and McKeown (1997)'s approach and orientation identification based on Turney (2001)'s approach. Then they calculate the average semantic orientation of the extracted words and phrases of the entire review–if the average is positive, then the review is classified as recommended, otherwise not recommended. The average classification accuracies of the reviews in different domains is 74.39%. The last step is similar to Hu and Liu (2004a)'s approach except that they classify the whole review document instead of one sentence.

In contrast to Turney (2002)'s approach, Pang et al. (2002) used supervised machine learning to classify movie reviews. Without classifying individual sentiment words or phrases, they extract different features from the review and use machine learning algorithms Naïve Bayes (NB), Maximum Entropy (ME) and Support Vector Machine (SVM) to classify the reviews. The features include a single item or a combination of the following: presence of unigrams, frequency of unigrams, bigrams, POS tag, adjectives, top 2633 unigrams and position of the word in text. They achieved accuracies between 78.7% and 82.9%.

3 Challenges and Possible Future Work

The main challenging aspects of sentiment analysis exist in use of other types of words, sentiment lexicon construction, dealing with negation expressions, degrees of sentiment, complexity of sentence/document, words in different contexts, heterogenous text sentiment classification etc. Some approaches have been proposed to solve these issues; however, more future research could be dedicated to these challenges.

3.1 Other Word Types

Most of the work done for sentiment analysis so far has been focused on content words as nouns, verbs, adjectives and adverbs, especially the latter two types. Although some works that focus on classifying the document as a whole use all types of words or collocations as in (Pang et al., 2002; Wiebe et al., 2001; Tong, 2001), they do not make use of a particular word type other than content words semantically. However, other word types could also have affect sentiment classification. For example, conjunctions as "but" connects two parts of a sentence together but emphasize on the part following "but" (just as shown in this very sentence!). "The movie is good but difficult to understand" would be classified as neutral if we simply count the number of positive sentiment words ("good") and negative ones ("difficult"). It may even be classified as positive if we look at the opinion word ("good") closest to the feature ("the movie") as in (Hu and Liu, 2004a). However, if we make use the conjunction "but" and give a higher weight to the sentence part following "but", in this case "difficult", the sentence would be classified correctly as negative. Besides "but", other words like "although", "nevertheless" and "still" etc. can also be utilized for sentence sentiment classification.

3.2 Sentiment Lexicon

Many works like (Nasukawa and Yi, 2003; Bruce and Wiebe, 1999; Wiebe and Bruce, 1999) manually define sentiment expressions in their own sentiment lexicon by denoting the polarity, POS tagging, canonical form and argument type (subject or object) for the corresponding sentiment words. However, this method is tedious and inefficient and may not be accurate. Also due to fact that these independent sentiment lexicons are not publicly available, the sentiment analysis community can not collaborate and thus are unable to provide a standard lexicon for future use. The situation did not improve until recently when some work has focused on solving this problem. Esuli and Sebastiani (2006) construct a public available lexical resource for opinion mining–SentiWordNet, a WordNet (Fellbaum, 1998) like lexicon but having emphasis on sentiment orientation of the words. They associate each synset s in WordNet to three numerical scores Obj(s), Pos(s) and Neg(s) to describe how objective, positive, and negative the terms contained in the synset are. Their method to construct the SentiWordNet is based on the quantitative analysis of the glosses associated to the synsets, and on the use of the resulting vectorial term representations for semi-supervised synset classification. Their approach and related approaches have extended the dimension of current work and provided a new tool in sentiment analysis; however, these lexicons are yet to be developed. Many features of the terms in the sense of subjectivity and sentiment orientation are yet to be added to solve the remaining problems in this field.

3.3 Negations

When negative expressions are associated Nasukawa and Yi (2003) simply reverse the polarity. Hu and Liu (2004a) also use the opposite orientation if negation word appears closely around the opinion word in the sentence. However, unlike the operations in mathematics where the negation of positive is negative and vice versa, the adding negations to a word or phrase in real world text does not equal to the effect of putting a minus sign in front of a number. In other words, the negation of positive may not be negative for words and phrases. For example, ”late” is negative; but adding a ”not” in front does not make ”not late” to be positive as ”not late” is not equal to ”early” which is the opposite of ”late”. This problem may also due to the fact that many approaches use only two values positive and negative (sometimes three values including neutral) to represent the sentiment orientation whereas words/phrases could have a degree of polarity as we discuss in Subsection 3.4.

3.4 Degrees of Sentiment

As pointed out in (Esuli and Sebastiani, 2006), there are several tasks concerning the degree of sentiment level in opinion mining. The first is called the SO-polarity, which is to determine whether the opinion is subjective (S) or objective (O) (Pang and Lee, 2004; Yu and Hatzivassiloglou, 2003). The second subtask is called the PN-polarity, which is to determine whether the opinion is positive (P) or negative (N) as described in many current approaches (Hatzivassiloglou and McKeown, 1997; Hu and Liu, 2004a; Pang et al., 2002; Pang and Lee, 2004; Turney and Littman, 2003; Turney and Littman, 2002).

As we discuss in Subsection 3.3, classifying sentiment orientation to mere positive and negative may pose some problems, the degree or strength of the polarity should also be analyzed. In (Pang and Lee, 2005; Wilson et al., 2004), opinions are assigned a degree of polarity, e.g. a positive opinion can be weakly positive, mildly positive, or strongly positive. This approach models sentiment orientation to more reality.

However, in addition to classifying words into different degrees of sentiment orientations as mentioned in above approaches, some modifying words (especially adverbs like ”very”, ”a little”, ”great”, ”somewhat” etc.) can also be utilized to determine the sentence sentiment orientation level. The phrase ”very satisfied” and ”somewhat satisfied” will be classified into strongly positive and weakly positive properly if ”very” and ”somewhat” are analyzed and used to determine the degree of polarity.

3.5 Sentence/Documents Complexity

Sentence complexity is quite an issue for current research in sentiment sentence classification. Approaches like (Nasukawa and Yi, 2003) only attempt to classify sentences with simple structures. For complex sentences, they could only partition the sentence and solve the parts individually. Without analyzing the whole sentence structure, the overall sentiment may be classified wrongly and the classification accuracy is hence affected.

Document complexity can pose bigger problem for sentiment document classification. Turney (2002) points out that in movie reviews, there could be different opinions in different sections of the text regarding to different aspects of the movie. Then it is rather difficult for the sentiment classification system to determine the true intension and sentiment of the reviewer towards the movie. A even worse case happens when the writer uses a lot of paragraphs to describe how he/she hates one of the movie actors but uses only a small paragraph to express how he/she still loves the movie after all. When encountered such situations, the system may very well be fooled to classify the review as negative. One possible approach to solve this problem is to differentiate the various aspects of a document, and classify the sentiment with respect to the aspects rather than classifying the whole document. Hu and Liu (2004a) used this approach to differentiate the sentiments towards different features of the same product in the same document.

3.6 Contextual Sentiment

Current works for sentiment orientation identification of words have not considered much of the context environment.

Same words in different contexts can have the different POS tags and different meanings. In this case, POS tagging is a good preprocessing tool to help the word sentiment orientation identification. Approaches mentioned in Subsection 2.1 indeed employ this method to help improve the performance.

However, Same words in different contexts can have the same POS tag but different meaning and different sentiment orientations. The word "poor" is "the system performance is poor" is has a negative sentiment orientation; but in "we should help the poor people", "poor" is neutral. When such cases are encountered, word sense disambiguation may help to solve the problem. Once we know from word sense disambiguation that "poor" in first sentence means inferior in quality while in second sentence it means having little money, the sentiment orientation identification for "poor" is easy to solve.

The worst case can still occur when same words in different contexts have the same POS tag and meaning but different sentiment orientations. For example, "old" in "The man is old" is neutral; whereas in "The technique is old" it is negative. In this case, unfortunately, neither POS tagging nor word sense disambiguation could help. More algorithms and techniques should be developed to tackle this problem.

3.7 Heterogenous Documents

Documents of different types, or even the documents of the same type but different domains are not of the same difficulty level to do sentiment classification. As the results reported in Turney (2002)'s paper, the classification accuracy differ for reviews of different domains-reviews on automobiles and banks have higher accuracies (84.00% and 80.00% respectively) whereas reviews on movies and travel destinations have lower accuracies (65.83% and 70.53% respectively). This may be also due to the different sentence/document complexity level for different types/domains of documents as discussed in Subsection 3.5. In future work, natural language processing techniques can be developed aiming at specific type/domain of document, sentiment classifier could also be trained with respect to a particular type/domain to improve the accuracy for this particular type/domain.

3.8 Other issues

Besides the problems discussed in previous Subsections (3.1 to 3.7), there are other issues like reference resolution etc. remain unsolved for sentiment analysis. Future work may take these issues into account as well in order to improve the accuracy.

4 Conclusion

In this survey, we have compared various approaches on sentiment analysis, the recent hot topic developed with the fast growth of internet documents, especially online reviews (product reviews, movie reviews and etc.). Much of the current research has been focused on this filed of study, also known as opinion mining, which is a recent discipline involving information retrieval (IR), natural language processing (NLP) and computational linguistics. In sentiment analysis, the usual tasks include sentiment word or phrase identification, sentiment orientation identification and sentiment sentence or document classification. Different from genre classification or topic classification, sentiment classification requires more understanding of the text in the sentiment sense. The applications of opinion mining include message filtering (Spertus, 1997), product review summarization (Hu and Liu, 2004a), movie review recommender systems (Pang et al., 2002; Terveen et al., 1997; Tatemura, 2000) etc. There have been various approaches and methods on each of the above tasks for sentiment analysis where each approach or method has its own focus, strength and limitations. Although the techniques and algorithms used for opinion mining are advancing fast; however, a lot of problems in this field of study remain unsolved. The main challenging aspects exist in use of other types of words, sentiment lexicon construction, dealing with negation expressions, degrees of sentiment, complexity of sentence/document, words in different contexts, heterogenous text sentiment classification etc. More future research could be dedicated to these challenges.

References

Alina Andreevskaia and Sabine Bergler. 2006. *Mining WordNet for Fuzzy Sentiment: Sentiment Tag Extraction from WordNet Glosses*. Proceedings of 11th Conference of the European Chapter of the Association for Computational Linguistics (EACL-06). 2006, Trento, Italy.

Philip Beineke, Trevor Hastie and Shivakumar Vaithyanathan. 2004. *The Sentimental Factor: Improving Review Classification via Human-Provided Information*. Proceedings of the 42th Annual Meeting of the Association for Computational Linguistics (ACL-04). 2004, Barcelone, Spain.

Eric Brill. 1994. *Some Advances in Transformation-Based Part of Speech Tagging*. Proceedings of the 12th National Conference on Artificial Intelligence. 1994, Menlo Park, CA: AAAI Press, pp722–727.

Rebecca F. Bruce and Janyce M. Wiebe. 1999. *Recognizing subjectivity: a case study in manual tagging*. Natural Language Engineering 5(2):187–205.

Kenneth Ward Church, Patrick Hanks. 1989. *Word association norms, mutual information and lexicography*. Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics. 1989, Vancouver, B.C., Canada, pp76–83.

Das and Chen. 2001. *Yahoo for amazon: Extracting market sentiment from stock message boards*. Proceedings of the 8th Asia Pacific Finance Association Annual Conference.

Kushal Dave, Steve Lawrence and David M. Pennock 2003. *Mining the peanut gallery: opinion extraction and semantic classification of product reviews*. Proceedings of the 12th international conference on World Wide Web. 2003, Budapest, Hungary, pp519–528.

Andrea Esuli and Fabrizio Sebastiani 2006. *SentiWordNet: A Publicly Available Lexical Resource for Opinion Mining*. Proceedings of the 5th Conference on Language Resources and Evaluation (LREC-06). 2006, Genova, Italy.

Christiane Fellbaum . 1998. *WordNet: an electronic lexical database*. MIT Press.

Vasileios Hatzivassiloglou and Kathleen R. McKeown. 1997. *Predicting the Semantic Orientation of Adjectives*. Proceedings of the 8th conference on European chapter of the Association for Computational Linguistics. 1997, Madrid, Spain.

Vasileios Hatzivassiloglou and Janyce M. Wiebe. 2000. *Effects of adjective orientation and gradability on sentence subjectivity*. Proceedings of the 8th conference on European chapter of the Association for Computational Linguistics. 1997, Madrid, Spain.

Minqing Hu and Bing Liu. 2004a. *Mining and summarizing customer reviews*. Proceedings of the 10th ACM SIGKDD international conference on Knowledge discovery and data mining. Aug. 22–25, 2004, Seattle, WA, USA.

Minqing Hu and Bing Liu. 2004b. *Mining opinion features in customer reviews*. Proceedings of 9th National Conference on Artificial Intellgience. Jul. 2004, San Jose, USA.

Chris Manning and Hinrich Schutze. 1999. *Foundations of Statistical Natural Language Processing*. MIT Press, Cambridge, MA.

Tetsuya Nasukawa and Jeonghee Yi. 2003. *Sentiment analysis: capturing favorability using natural language processing*. Proceedings of the 2nd international conference on Knowledge capture. Oct. 23–25, 2003, Sanibel Island, FL, USA.

Mary S. Neff, Roy J. Byrd, and Branimir K. Boguraev. 2003. *The Talent System: TEXTRACT Architecture and Data Model*. Proceedings of the HLT-NAACL 2003 Workshop on Software Engineering and Architecture of Language Technology systems (SEALTS).

Bo Pang, Lillian Lee and Shivakumar Vaithyanathan. 2002. *Thumbs up? Sentiment classification using machine learning techniques*. Proceedings of the 7th Conference on Empirical Methods in Natural Language Processing (EMNLP-02).

Bo Pang and Lillian Lee. 2004. *A sentiment education: sentiment analysis using subjectivity summrarization based on minimum cuts*. Proceedings of the Conference of the Association for Computational Linguistics (ACL-04).

Bo Pang and Lillian Lee. 2005. *Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales*. Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL-05). 2005, Ann Arbor, USA, pp133–140.

Ellen Spertus. 1997. *Smokey: automatic recognition of hostile messages*. Proceedings of Innovative Applications of Artificial Intelligence (IAAI). pp1058–1065.

Junichi Tatemura. 2000. *Virtual reviewers for collaborative exploration of movie reviews*. Proceedings of the 5th International Conference on Intelligent User Interface. pp272–275.

Loren Terveen, Will Hill, Brian Amento, David McDonald and Josh Creter. 1997. *PHOAKS: a system for sharing recommendations*. Communications of the ACM, 40(3). pp272–275.

Tong. 2001. *An operational system for detecting and tracking opinions in on-line discussion*. Proceedings of SIGIR Workshop on Operational Text Classification. 2001, New York, pp1–6.

Peter Turney and Michael Littman. 2002. *Unsupervised Learning of Semantic Orientation from a Hundred-Billion-Word Corpus*. Technical Report NRC Technical Report ERB-1094. Institute for Information Technology, National Research Council Canada.

Peter Turney. 2001. *Mining the Web for synonyms: PMI-IR versus LSA on TOEFL*. Proceedings of the 12th European Conference on Machine Learning. 2001, Berlin: Spinger-Verlag, pp. 491–502.

Peter Turney. 2002. *Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews*. Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL-02). Jun. 2002, Philadelphia, PN, USA, pp. 417–424.

Peter Turney and Michael Littman. 2003. *Measuring praise and criticism: Inference of semantic orientation from association*. ACM Tranctions on Information Systems, 21(4):315-346.

Janyce M. Wiebe, Rebecca F. Bruce, Thomas O'Hara. 1999. *Development and use of a gold standard data set for subjectivity classifications*. Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics (ACL-99). Jun. 20–26 1999, University of Maryland, ML, USA.

Janyce M. Wiebe, Theresa Wilson and Matthew Bell. 2001. *Identifying collocations for recognizing opinions*. Proceedings of the ACL/EACL Workshop on Collocation. 2001.

Theresa Wilson, Janyce M. Wiebe and Rebecca Hwa. 2004. *Just how mad are you? Finding strong and weak opinion clauses*. Proceedings of the 21st Conference of the American Association for Artificial Intelligence (AAAI-04). 2004, San Jose, USA, pp761–769

Jeonghee Yi and Wayne Niblack. 2005. *Sentiment mining in WebFountain*. Proceedings of the 21st International Conference on Data Engineering. Apr. 2005.

Hong Yu and Vasileios Hatzivassiloglou. 2003. *Towards answering opinion questions: Seperating facts from opinions and identifying the polarity of opinion sentences*. Proceedings of the 8th Conference on Empirical Methods in Natural Language Processing (EMNLP-03). 2003, Sapporo, Japan, pp129–136.