

Learning to Shift the Polarity of Words for Sentiment Classification

Daisuke Ikeda[†] Hiroya Takamura[‡] Lev-Arie Ratinov^{††} Manabu Okumura[‡]

[†]Department of Computational Intelligence and Systems Science, Tokyo Institute of Technology
ikedalr.pi.titech.ac.jp

^{††}Department of Computer Science, University of Illinois at Urbana-Champaign
ratinov2@uiuc.edu

[‡]Precision and Intelligence Laboratory, Tokyo Institute of Technology
{takamura, oku}@pi.titech.ac.jp

Abstract

We propose a machine learning based method of sentiment classification of sentences using word-level polarity. The polarities of words in a sentence are not always the same as that of the sentence, because there can be polarity-shifters such as negation expressions. The proposed method models the polarity-shifters. Our model can be trained in two different ways: word-wise and sentence-wise learning. In sentence-wise learning, the model can be trained so that the prediction of sentence polarities should be accurate. The model can also be combined with features used in previous work such as bag-of-words and n-grams. We empirically show that our method almost always improves the performance of sentiment classification of sentences especially when we have only small amount of training data.

1 Introduction

Due to the recent popularity of the internet, individuals have been able to provide various information to the public easily and actively (e.g., by weblogs or online bulletin boards). The information often includes opinions or sentiments on a variety of things such as new products. A huge amount of work has been devoted to analysis of the information, which is called *sentiment analysis*. The sentiment analysis has been done at different levels including words, sentences, and documents. Among them, we focus on the sentiment classification of sentences, the task

to classify sentences into “positive” or “negative”, because this task is fundamental and has a wide applicability in sentiment analysis. For example, we can retrieve individuals’ opinions that are related to a product and can find whether they have the positive attitude to the product.

There has been much work on the identification of sentiment polarity of words. For instance, “beautiful” is positively oriented, while “dirty” is negatively oriented. We use the term *sentiment words* to refer to those words that are listed in a predefined polarity dictionary. Sentiment words are a basic resource for sentiment analysis and thus believed to have a great potential for applications. However, it is still an open problem how we can effectively use sentiment words to improve performance of sentiment classification of sentences or documents.

The simplest way for that purpose would be the majority voting by the number of positive words and the number of negative words in the given sentence. However, the polarities of words in a sentence are not always the same as that of the sentence, because there can be polarity-shifters such as negation expressions. This inconsistency of word-level polarity and sentence-level polarity often causes errors in classification by the simple majority voting method. A manual list of polarity-shifters, which are the words that can shift the sentiment polarity of another word (e.g., negations), has been suggested. However, it has limitations due to the diversity of expressions.

Therefore, we propose a machine learning based method that models the polarity-shifters. The model can be trained in two different ways: *word-wise*

and *sentence-wise*. While the word-wise learning focuses on the prediction of polarity shifts, the sentence-wise learning focuses more on the prediction of sentence polarities. The model can also be combined with features used in previous work such as bag-of-words, n-grams and dependency trees. We empirically show that our method almost always improves the performance of sentiment classification of sentences especially when we have only small amount of training data.

The rest of the paper is organized as follows. In Section 2, we briefly present the related work. In Section 3, we discuss well-known methods that use word-level polarities and describe our motivation. In Section 4, we describe our proposed model, how to train the model, and how to classify sentences using the model. We present our experiments and results in Section 5. Finally in Section 6, we conclude our work and mention possible future work.

2 Related Work

Supervised machine learning methods including Support Vector Machines (SVM) are often used in sentiment analysis and shown to be very promising (Pang et al., 2002; Matsumoto et al., 2005; Kudo and Matsumoto, 2004; Mullen and Collier, 2004; Gamon, 2004). One of the advantages of these methods is that a wide variety of features such as dependency trees and sequences of words can easily be incorporated (Matsumoto et al., 2005; Kudo and Matsumoto, 2004; Pang et al., 2002). Our attempt in this paper is not to use the information included in those substructures of sentences, but to use the word-level polarities, which is a resource usually at hand. Thus our work is an instantiation of the idea to use a resource on one linguistic layer (e.g., word level) to the analysis of another layer (sentence level).

There have been some pieces of work which focus on multiple levels in text. Mao and Lebanon (2006) proposed a method that captures local sentiment flow in documents using isotonic conditional random fields. Pang and Lee (2004) proposed to eliminate objective sentences before the sentiment classification of documents. McDonald et al. (2007) proposed a model for classifying sentences and documents simultaneously. They experimented with joint classification of subjectivity for sentence-level,

and sentiment for document-level, and reported that their model obtained higher accuracy than the standard document classification model.

Although these pieces of work aim to predict not sentence-level but document-level sentiments, their concepts are similar to ours. However, all the above methods require annotated corpora for all levels, such as both subjectivity for sentences and sentiments for documents, which are fairly expensive to obtain. Although we also focus on two different layers, our method does not require such expensive labeled data. What we require is just sentence-level labeled training data and a polarity dictionary of sentiment words.

3 Simple Voting by Sentiment Words

One of the simplest ways to classify sentences using word-level polarities would be a majority voting, where the occurrences of positive words and those of negative words in the given sentence are counted and compared with each other. However, this majority voting method has several weaknesses. First, the majority voting cannot take into account at all the phenomenon that the word-level polarity is not always the same as the polarity of the sentence. Consider the following example:

I have not had any distortion problems with this phone and am more pleased with this phone than any I've used before.

where negative words are underlined and positive words are double-underlined. The example sentence has the positive polarity, though it locally contains negative words. The majority voting would misclassify it because of the two negative words.

This kind of inconsistency between sentence-level polarity and word-level polarity often occurs and causes errors in the majority voting. The reason is that the majority voting cannot take into account negation expressions or adversative conjunctions, e.g., "I have not had any ..." in the example above. Therefore, taking such polarity-shifting into account is important for classification of sentences using a polarity dictionary. To circumvent this problem, Kennedy and Inkpen (2006) and Hu and Liu (2004) proposed to use a manually-constructed list of polarity-shifters. However, it has limitations due to the diversity of expressions.

Another weakness of the majority voting is that it cannot be easily combined with existing methods that use the n-gram model or tree structures of the sentence as features. The method we propose here can easily be combined with existing methods and show better performance.

4 Word-Level Polarity-Shifting Model

We assume that when the polarity of a word is different from the polarity of the sentence, the polarity of the word is shifted by its context to adapt to the polarity of the sentence. Capturing such polarity-shifts will improve the classification performance of the majority voting classifier as well as of more sophisticated classifiers.

In this paper, we propose a word polarity-shifting model to capture such phenomena. This model is a kind of binary classification model which determines whether the polarity is shifted by its context. The model assigns a score $s_{shift}(x, S)$ to the sentiment word x in the sentence S . If the polarity of x is shifted in S , $s_{shift}(x, S) > 0$. If the polarity of x is not shifted in S , $s_{shift}(x, S) \leq 0$. Let \mathbf{w} be a parameter vector of the model and ϕ be a pre-defined feature function. Function s_{shift} is defined as

$$s_{shift}(x, S) = \mathbf{w} \cdot \phi(x, S). \quad (1)$$

Since this model is a linear discriminative model, there are well-known algorithms to estimate the parameters of the model.

Usually, such models are trained with each occurrence of words as one instance (word-wise learning). However, we can train our model more effectively with each sentence being one instance (sentence-wise learning). In this section, we describe how to train our model in two different ways and how to apply the model to a sentence classification.

4.1 Word-wise Learning

In this learning method, we train the word-level polarity-shift model with each occurrence of sentiment words being an instance. Training examples are automatically extracted by finding sentiment words in labeled sentences. In the example of Section 3, for instance, both negative words (“distortion” or “problems”) and a positive word (“pleased”) appear in a positive sentence. We regard “distortion”

and “problems”, whose polarities are different from that of the sentence, as belonging to the *polarity-shifted* class. On the contrary, we regard “pleased”, whose polarity is the same as that of the sentence, as not belonging to *polarity-shifted* class.

We can use the majority voting by those (possibly polarity-shifted) sentiment words. Specifically, we first classify each sentiment word in the sentence according to whether the polarity is shifted or not. Then we use the majority voting to determine the polarity of the sentence. If the first classifier classifies a positive word into the “polarity-shifted” class, we treat the word as a negative one. We expect that the majority voting with polarity-shifting will outperform the simple majority voting without polarity-shifting. We actually use the weighted majority voting, where the polarity-shifting score for each sentiment word is used as the weight of the vote by the word. We expect that the score works as a confidence measure.

We can formulate this method as follows. Here, N and P are respectively defined as the sets of negative sentiment words and positive sentiment words. For instance, $x \in N$ means that x is a negative word. We also write $x \in S$ to express that the word x occurs in S .

First, let us define two scores, $score_p(S)$ and $score_n(S)$, for the input sentence S . The $score_p(S)$ and the $score_n(S)$ respectively represent the number of votes for S being positive and the number of votes for S being negative. If $score_p(S) > score_n(S)$, we regard the sentence S as having the positive polarity, otherwise negative. We suppose that the following relations hold for the scores:

$$score_p(S) = \sum_{x \in P \cap S} -s_{shift}(x, S) + \sum_{x \in N \cap S} s_{shift}(x, S), \quad (2)$$

$$score_n(S) = \sum_{x \in P \cap S} s_{shift}(x, S) + \sum_{x \in N \cap S} -s_{shift}(x, S). \quad (3)$$

When either a polarity-unchanged positive word ($s_{shift}(x, S) \leq 0$) or a polarity-shifted negative word occurs in the sentence S , $score_p(S)$ increases. We can easily obtain the following relation between two scores:

$$score_p(S) = -score_n(S). \quad (4)$$

Since, according to this relation, $score_p(S) > score_n(S)$ is equivalent to $score_p(S) > 0$, we use only $score_p(S)$ for the rest of this paper.

4.2 Sentence-wise Learning

The equation (2) can be rewritten as

$$\begin{aligned} score_p(S) &= \sum_{x \in S} s_{shift}(x, S)I(x) \\ &= \sum_{x \in S} \mathbf{w} \cdot \phi(x, S)I(x) \\ &= \mathbf{w} \cdot \left\{ \sum_{x \in S} \phi(x, S)I(x) \right\}, \end{aligned} \quad (5)$$

where $I(x)$ is the function defined as follows:

$$I(x) = \begin{cases} +1 & \text{if } x \in N, \\ -1 & \text{if } x \in P, \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

This $score_p(S)$ can also be seen as a linear discriminative model and the parameters of the model can be estimated directly (i.e., without carrying out word-wise learning). Each labeled sentence in a corpus can be used as a training instance for the model.

In this method, the model is learned so that the predictive ability for sentence classification is optimized, instead of the predictive ability for polarity-shifting. Therefore, this model can remain indecisive on the classification of word instances that have little contextual evidence about whether polarity-shifting occurs or not. The model can rely more heavily on word instances that have much evidence.

In contrast, the word-wise learning trains the model with all the sentiment words appearing in a corpus. It is assumed here that all the sentiment words have relations with the sentence-level polarity, and that we can always find the evidence of the phenomena that the polarity of a word is different from that of a sentence. Obviously, this assumption is not always correct. As a result, the word-wise learning sometimes puts a large weight on a context word that is irrelevant to the polarity-shifting. This might degrade the performance of sentence classification as well as of polarity-shifting.

4.3 Hybrid Model

Both methods described in Sections 4.1 and 4.2 are to predict the sentence-level polarity only with

the word-level polarity. On the other hand, several methods that use another set of features, for example, bag-of-words, n-grams or dependency trees, were proposed for the sentence or document classification tasks. We propose to combine our method with existing methods. We refer to it as *hybrid model*.

In recent work, discriminative models including SVM are often used with many different features. These methods are generally represented as

$$score'_p(X) = \mathbf{w}' \cdot \phi'(X), \quad (7)$$

where X indicates the target of classification, for example, a sentence or a document. If $score'_p(X) > 0$, X is classified into the target class. $\phi'(X)$ is a feature function. When the method uses the bag-of-words model, ϕ' maps X to a vector with each element corresponding to a word.

Here, we define new score function $score_{comb}(S)$ as a linear combination of $score_p(S)$, the score function of our sentence-wise learning, and $score'_p(S)$, the score function of an existing method. Using this, we can write the function as

$$\begin{aligned} score_{comb}(S) &= \lambda score_p(S) + (1 - \lambda) score'_p(S) \\ &= \lambda \sum_{x \in S} \mathbf{w} \cdot \phi(x, S)I(x) + (1 - \lambda) \mathbf{w}' \cdot \phi'(S) \\ &= \mathbf{w}_{comb} \cdot \left\langle \lambda \sum_{x \in S} \phi(x, S)I(x), (1 - \lambda) \phi'(S) \right\rangle. \end{aligned} \quad (8)$$

Note that $\langle \rangle$ indicates the concatenation of two vectors, \mathbf{w}_{comb} is defined as $\langle \mathbf{w}, \mathbf{w}' \rangle$ and λ is a parameter which controls the influence of the word-level polarity-shifting model. This model is also a discriminative model and we can estimate the parameters with a variety of algorithms including SVMs. We can incorporate additional information like bag-of-words or dependency trees by $\phi'(S)$.

4.4 Discussions on the Proposed Model

Features such as n-grams or dependency trees can also capture some negations or polarity-shifters. For example, although “satisfy” is positive, the bigram model will learn “not satisfy” as a feature correlated with negative polarity if it appears in the training data. However, the bigram model cannot generalize the learned knowledge to other features such

Table 1: Statistics of the corpus

	customer	movie
# of Labeled Sentences	1,700	10,662
Available	1,436	9,492
# of Sentiment Words	3,276	26,493
Inconsistent Words	1,076	10,674

as “not great” or “not disappoint”. On the other hand, our polarity-shifter model learns that the word “not” causes polarity-shifts. Therefore, even if there was no “not disappoint” in training data, our model can determine that “not disappoint” has correlation with positive class, because the dictionary contains “disappoint” as a negative word. For this reason, the polarity-shifting model can be learned even with smaller training data.

What we can obtain from the proposed method is not only a set of polarity-shifters. We can also obtain the weight vector w , which indicates the strength of each polarity-shifter and is learned so that the predictive ability of sentence classification is optimized especially in the sentence-wise learning. It is impossible to manually determine such weights for numerous features.

It is also worth noting that all the models proposed in this paper can be represented as a kernel function. For example, the hybrid model can be seen as the following kernel:

$$K_{comb}(S_1, S_2) = \lambda \sum_{x_i \in S_1} \sum_{x_j \in S_2} K((x_i, S_1), (x_j, S_2)) + (1 - \lambda)K'(S_1, S_2). \quad (9)$$

Here, K means the kernel function between words and K' means the kernel function between sentences respectively. In addition, $\sum_{x_i} \sum_{x_j} K((x_i, S_1), (x_j, S_2))$ can be seen as an instance of *convolution kernels*, which was proposed by Haussler (1999). Convolution kernels are a general class of kernel functions which are calculated on the basis of kernels between substructures of inputs. Our proposed kernel treats sentences as input, and treats sentiment words as substructures of sentences. We can use high degree polynomial kernels as both K which is a kernel between substructures, i.e. sentiment words, of sentences, and K' which is a kernel between sentences to make the

classifiers take into consideration the combination of features.

5 Evaluation

5.1 Datasets

We used two datasets, customer reviews ¹ (Hu and Liu, 2004) and movie reviews ² (Pang and Lee, 2005) to evaluate sentiment classification of sentences. Both of these two datasets are often used for evaluation in sentiment analysis researches. The number of examples and other statistics of the datasets are shown in Table 1.

Our method cannot be applied to sentences which contain no sentiment words. We therefore eliminated such sentences from the datasets. “Available” in Table 1 means the number of examples to which our method can be applied. “Sentiment Words” shows the number of sentiment words that are found in the given sentences. Please remember that sentiment words are defined as those words that are listed in a predefined polarity dictionary in this paper. “Inconsistent Words” shows the number of the words whose polarities conflicted with the polarity of the sentence.

We performed 5-fold cross-validation and used the classification accuracy as the evaluation measure. We extracted sentiment words from General Inquirer (Stone et al., 1996) and constructed a polarity dictionary. After some preprocessing, the dictionary contains 2,084 positive words and 2,685 negative words.

5.2 Experimental Settings

We employed the Max Margin Online Learning Algorithms for parameter estimation of the model (Crammer et al., 2006; McDonald et al., 2007). In preliminary experiments, this algorithm yielded equal or better results compared to SVMs. As the feature representation, $\phi(x, S)$, of polarity-shifting model, we used the local context of three words to the left and right of the target sentiment word. We used the polynomial kernel of degree 2 for polarity-shifting model and the linear kernel for oth-

¹<http://www.cs.uic.edu/~liub/FBS/FBS.html>

²<http://www.cs.cornell.edu/people/pabo/movie-review-data/>

Table 2: Experimental results of the sentence classification

methods	customer	movie
Baseline	0.638	0.504
BoW	0.790	0.724
2gram	0.809	0.756
3gram	0.800	0.762
Simple-Voting	0.716	0.624
Negation Voting	0.733	0.658
Word-wise	0.783	0.699
Sentence-wise	0.806	0.718
Hybrid BoW	0.827	0.748
Hybrid 2gram	0.840	0.755
Hybrid 3gram	0.837	0.758
Opt	0.840	0.770

ers, and feature vectors are normalized to 1. In hybrid models, the feature vectors, $\sum_{x \in S} \phi(x, S)I(x)$ and $\phi'(S)$ are normalized respectively.

5.3 Comparison of the Methods

We compared the following methods:

- **Baseline** classifies all sentences as positive.
- **BoW** uses unigram features. **2gram** uses unigrams and bigrams. **3gram** uses unigrams, bigrams, and 3grams.
- **Simple-Voting** is the most simple majority voting with word-level polarity (Section 3).
- **Negation Voting** proposed by Hu and Liu (2004) is the majority voting that takes negations into account. As negations, we employed *not*, *no*, *yet*, *never*, *none*, *nobody*, *nowhere*, *nothing*, and *neither*, which are taken from (Polanyi and Zaenen, 2004; Kennedy and Inkpen, 2006; Hu and Liu, 2004) (Section 3).
- **Word-wise** was described in Section 4.1.
- **Sentence-wise** was described in Section 4.2.
- **Hybrid BoW, hybrid 2gram, hybrid 3gram** are combinations of sentence-wise model and respectively **BoW**, **2gram** and **3gram** (Section 4.3). We set $\lambda = 0.5$.

Table 2 shows the results of these experiments. Hybrid 3gram, which corresponds to the proposed method, obtained the best accuracy on customer review dataset. However, on movie review dataset, the proposed method did not outperform 3gram. In Section 5.4, we will discuss this result in details. Comparing word-wise to simple-voting, the accuracy increased by about 7 points. This means that the polarity-shifting model can capture the polarity-shifts and it is an important factor for sentiment classification. In addition, we can see the effectiveness of sentence-wise, by comparing it to word-wise in accuracy.

“Opt” in Table 2 shows the results of hybrid models with optimal λ and combination of models. The optimal results of hybrid models achieved the best accuracy on both datasets.

We show some dominating polarity-shifters obtained through learning. We obtained many negations (e.g., no, not, n’t, never), modal verbs (e.g., might, would, may), prepositions (e.g., without, despite), comma with a conjunction (e.g., “, but” as in “the case is strong and stylish, but lacks a window”), and idiomatic expressions (e.g., “hard resist” as in “it is hard to resist”, and “real snooze”).

5.4 Effect of Training Data Size

When we have a large amount of training data, the n-gram classifier can learn well whether each n-gram tends to appear in the positive class or the negative class. However, when we have only a small amount of training data, the n-gram classifier cannot capture such tendency. Therefore the external knowledge, such as word-level polarity, could be more valuable information for classification. Thus it is expected that the sentence-wise model and the hybrid model will outperform n-gram classifier which does not take word-level polarity into account, more largely with few training data.

To verify this conjecture, we conducted experiments by changing the number of the training examples, i.e., the labeled sentences. We evaluated three models: sentence-wise, 3gram model and hybrid 3gram on both customer review and movie review.

Figures 1 and 2 show the results on customer review and movie review respectively. When the size of the training data is small, sentence-wise outper-




Figure 1: Experimental results on customer review

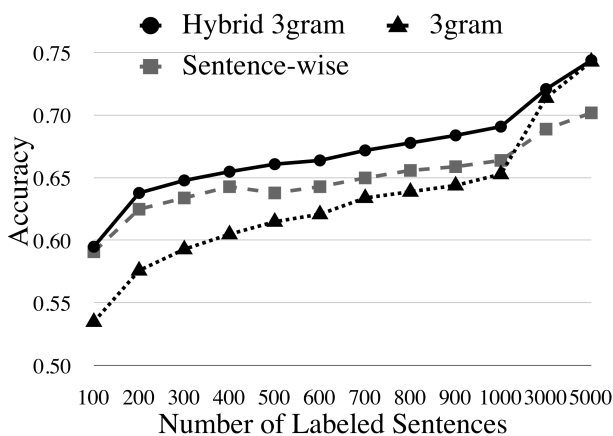


Figure 2: Experimental results on movie review

forms 3gram on both datasets. We can also see that the advantage of sentence-wise becomes smaller as the amount of training data increases, and that the hybrid 3gram model almost always achieved the best accuracy among the three models. Similar behaviour was observed when we ran the same experiments with 2gram or BoW model. From these results, we can conclude that, as we expected above, the word-level polarity is especially effective when we have only a limited amount of training data, and that the hybrid model can combine two models effectively.

6 Conclusion

We proposed a model that captures the polarity-shifting of sentiment words in sentences. We also presented two different learning methods for the model and proposed an augmented hybrid classifier

that is based both on the model and on existing classifiers. We evaluated our method and reported that the proposed method almost always improved the accuracy of sentence classification compared with other simpler methods. The improvement was more significant when we have only a limited amount of training data.

For future work, we plan to explore new feature sets appropriate for our model. The feature sets we used for evaluation in this paper are not necessarily optimal and we can expect a better performance by exploring appropriate features. For example, dependency relations between words or appearances of conjunctions will be useful. The position of a word in the given sentence is also an important factor in sentiment analysis (Taboada and Grieve, 2004). Furthermore, we should directly take into account the fact that some words do not affect the polarity of the sentence, though the proposed method tackled this problem indirectly. We cannot avoid this problem to use word-level polarity more effectively. Lastly, since we proposed a method for the sentence-level sentiment prediction, our next step is to extend the method to the document-level sentiment prediction.

Acknowledgement

This research was supported in part by Overseas Advanced Educational Research Practice Support Program by Ministry of Education, Culture, Sports, Science and Technology.

References

- Koby Crammer, Ofer Dekel, Joseph Keshet, Shai Shalev-Shwartz, and Yoram Singer. Online Passive-Aggressive Algorithms. In *Journal of Machine Learning Research*, Vol.7, Mar, pp.551–585, 2006.
- Michael Gamon. Sentiment classification on customer feedback data: noisy data, large feature vectors, and the role of linguistic analysis. In *Proceedings of the 20th International Conference on Computational Linguistics (COLING-2004)*, pp.841–847, 2004.
- David Haussler. Convolution Kernels on Discrete Structures, Technical Report UCS-CRL-99-10, University of California in Santa Cruz, 1999.
- Minqing Hu and Bing Liu. Mining Opinion Features in Customer Reviews. In *Proceedings of Nineteenth National Conference on Artificial Intelligence (AAAI-2004)*, pp.755–560, San Jose, USA, July 2004.

- Alistair Kennedy and Diana Inkpen. Sentiment Classification of Movie and Product Reviews Using Contextual Valence Shifters. In *Workshop on the Analysis of Formal and Informal Information Exchange during Negotiations (FINEXIN-2005)*, 2005.
- Taku Kudo and Yuji Matsumoto. A Boosting Algorithm for Classification of Semi-Structured Text. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2004)*, pp.301–308, 2004.
- Yu Mao and Guy Lebanon. Isotonic Conditional Random Fields and Local Sentiment Flow. In *Proceedings of the Neural Information Processing Systems (NIPS-2006)*, pp.961–968, 2006.
- Shotaro Matsumoto, Hiroya Takamura, and Manabu Okumura. Sentiment Classification using Word Sub-Sequences and Dependency Sub-Trees. In *Proceedings of the 9th Pacific-Asia International Conference on Knowledge Discovery and Data Mining (PAKDD-2005)*, pp.301–310, 2005.
- Ryan McDonald, Kerry Hannan, Tyler Neylon, Mike Wells, and Jeff Reynar. Structured Models for Fine-to-Coarse Sentiment Analysis. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics (ACL-2007)*, pp.432–439, 2007.
- Tony Mullen and Nigel Collier. Sentiment analysis using support vector machines with diverse information sources. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2004)*, pp.412–418, 2004.
- Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. Thumbs up? Sentiment Classification using Machine Learning Techniques. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2002)*, pp.76–86, 2002.
- Bo Pang and Lillian Lee. A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts. In *Proceedings of the 42th Annual Meeting of the Association for Computational Linguistics (ACL-2004)*, pp.271–278, 2004.
- Bo Pang and Lillian Lee. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL-2005)*, pp.115–124, 2005.
- Livia Polanyi and Annie Zaenen. Contextual Valence Shifters. In *AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications (AAAI-EAAT2004)*, 2004.
- Philip J. Stone, Dexter C. Dunphy, Marshall S. Smith, and Daniel M. Ogilvie. *The General Inquirer: A Computer Approach to Content Analysis*. The MIT Press, 1996.
- Maite Taboada and Jack Grieve. Analyzing Appraisal Automatically. In *AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications (AAAI-EAAT2004)*, pp.158–161, 2004.