Extracting Domain-Dependent Semantic Orientations of Latent Variables for Sentiment Classification

Yeha Lee, Jungi Kim, and Jong-Hyeok Lee

Division of Electrical and Computer Engineering
Pohang University of Science and Technology
San 31, Hyoja-Dong, Nam-Gu, Pohang, 790–784, Republic of Korea
{sion,yangpa,jhlee}@postech.ac.kr

Abstract. Sentiment analysis of weblogs is a challenging problem. Most previous work utilized semantic orientations of words or phrases to classify sentiments of weblogs. The problem with this approach is that semantic orientations of words or phrases are investigated without considering the domain of weblogs. Weblogs contain the author's various opinions about multifaceted topics. Therefore, we have to treat a semantic orientation domain-dependently. In this paper, we present an unsupervised learning model based on aspect model to classify sentiments of weblogs. Our model utilizes domain-dependent semantic orientations of latent variables instead of words or phrases, and uses them to classify sentiments of weblogs. Experiments on several domains confirm that our model assigns domain-dependent semantic orientations to latent variables correctly, and classifies sentiments of weblogs effectively.

Keywords: sentiment classification, sentiment analysis, information extraction, text mining.

1 Introduction

An increasing number of internet users are expressing not only factual data but also their subjective thoughts and opinions on various topics in weblogs. A lot of recent research tries to access the author's opinions and sentiments hidden in the weblogs. There are many applications that can be exploited when we can access weblogs' opinions or sentiments. For instance, information extraction and question-answering systems could flag statements and queries regarding opinions rather than facts [1]. Also, it has proven useful for companies, recommender systems and editorial sites to create summaries of people's experiences and opinions that consist of subjective expressions extracted from reviews (as is commonly done in movie ads) or even just a review's *polarity* – positive ("thumbs up") or negative ("thumbs down") [2].

Most previous work [3,4,5,6] investigated semantic orientations of words or phrases to analyze sentiments of sentences or documents. However, semantic orientations of words or phrases are investigated without reflecting properties of the domain in which documents are included. Weblogs deal with various topics, and they contain both facts

W. Li and D. Mollá-Aliod (Eds.): ICCPOL 2009, LNAI 5459, pp. 201–212, 2009.

and opinions, and the same word might possess different semantic orientations in different domains. Therefore, we have to assign different semantic orientations of words or phrases according to the domain of the weblogs.

In this paper, we present an unsupervised learning model to classify sentiments of weblogs. Therefore our model does not require sentiment tagged corpus to be trained. Our model, which is based on an aspect model, assigns semantic orientations of latent factors depending on the domain of the weblogs. Our model uncovers semantic orientations of latent factors, which are domain dependent, and use them for sentiment classification. The experiments on three domains (movie, automobiles and digital cameras) show that our model is effective at accessing a domain dependent semantic orientation and classifying weblogs' sentiments.

The paper is organized as follows. Section 2 introduces related work. In Section 3, we formulate the sentiment classification problem of the weblogs and show the detail of our approach. The experimental results are presented in Section 4. Finally, we conclude the paper and discuss future work in Section 5.

2 Related Work

In the case of opinion analysis, there has been a lot of previous research such as the development of linguistic resource, subjectivity detection and sentiment classification.

In developing linguistic resources, some previous works have focused on learning adjectives or adjectival phrases [3,4,5]. Riloff et al. [7] developed a system that can distinguish subjective sentences from objective sentences using lists of subjective nouns learned by bootstrapping algorithms. Wilson et al. [8] presented a two-step contextual polarity classification to analyze phrase–level sentiment. Kamps et al. [9] proposed measures that determine the semantic orientation of adjectives using WordNet's synonymy relation. Takamura et al. [10] used latent variables to find semantic orientations of phrases. Their approach is similar to our model in that it uses latent variables. However our approach is different from Takamura et al. in that our approach finds semantic orientations of latent variables and uses them to classify the sentiments of the weblogs.

Turney [5] classified the sentiment of a document using the average semantic orientation of phrases that was assigned using the PMI method. Pang et al. [11] employed three standard machine learning techniques (Naive Bayes, Maximum Entropy, and SVM) to classify the sentiment of a document. Pang and Lee [2] presented cut-based classification. In their approach, sentences in a document are labeled as either subjective or objective, and a standard machine learning classifier that classifies document-level sentiment is applied to subjective sentences. Whitelaw et al. [12] presented a method for sentiment classification which is based on analysis of appraisal groups.

Applications using sentiment analysis of weblogs or news have been proposed [13, 14]. Mei et al. [15] proposed a probabilistic model to capture the mixture of topics and sentiments simultaneously. Liu et al. [16] investigated ways to use sentiment information from blogs for predicting product sales performance. In order to mine sentiment information from blogs, they presented a S-PLSA model based on PLSA. They used sentiment information captured by S-PLSA as features of an auto regression model. Our approach, on the other hand, presents a general framework to classify sentiments of the weblogs using semantic orientations of latent variables.

3 Weblogs Classification

In this section, we propose a novel probabilistic approach to classify the sentiments of weblogs, which is an unsupervised learning model.

The weblogs contain various opinions of the author about many fields. This property of the weblog causes a word to have different semantic orientations according to the domain in which it is used. For example, the adjective "unpredictable" may have a negative orientation in an automotive review in a phrase such as "unpredictable steering", but it could have a positive orientation in a movie review in a phrase such as "unpredictable plot" [5]. This problem results from the use of domain-independent semantic orientations of words or phrases without considering the properties of the domain in which they are used. In general, the semantic orientation of a word is affected by the surrounding context as well as the domain. For example, words following a negation word can have the opposite semantic orientation. However, in this study, we deal with the problem that results from the property of the weblogs, exemplified in the above case. In order to analyze complicated properties of weblogs, we regard them as results generated by the mixture model of the latent factors. We expect that those latent factors would correspond to the author's sentiments presented in the weblogs. We use semantic orientations of latent factors, instead of words or phrases, to classify the sentiments of weblogs.

To this end, we present a Sentiment Aspect Model (SentiAM) based on the aspect model. Aspect model [17] is a latent variable model for co-occurrence data which associates an unobserved class variable $z \in Z = \{z_1, \cdots, z_K\}$ with each observation. SentiAM adds a semantic orientation factor to the aspect model. We regard the semantic orientation factor as having a dependency relationship with the latent variable. Figure 1 (a) shows a graphical representation of the statistical dependencies of SentiAM. If our model can capture domain-dependent semantic orientations of latent variables, it will yield better sentiment classification accuracy compared to the model that uses domain-independent semantic orientations of words or phrases.

3.1 Sentiment Aspect Model

In this section, we formally present SentiAM to classify sentiments of weblogs.

Let $D = \{d_1, d_2, \cdots, d_N\}$ be a set of weblogs, with words from a vocabulary $W = \{w_1, w_2, \cdots, w_M\}$. The weblog data set can be represented as a $M \times N$ matrix $R = (n(d_i, w_j))_{ij}$, where $n(d_i, w_j)$ denotes the number of times w_j occurs in weblog d_i. And

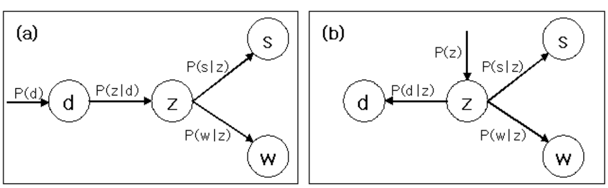

Fig. 1. Graphical model representation of the Sentiment Aspect Model. **S**: semantic orientation factor.

let $S = \{P, N, F\}$ be a set of semantic orientations. P, N and F represent positive, negative and neutral orientations respectively. Let $Z = \{z_1, z_2, \cdots, z_K\}$ be a set of latent variables.

Suppose that the weblog d, the semantic orientation s and the word w are conditionally independent given the latent variable z (corresponding graphical model representation is depicted in Figure 1 (b)). Then, we can view the problem of classifying the sentiment of the weblog d as follows:

$$S^* = \arg\max_{s \in \{P,N\}} P(s|d) = \arg\max_{s \in \{P,N\}} \sum_{z_k \in Z} P(s|z_k)P(z_k|d) \tag{1}$$

Because probability $P(d)$ does not have any effect on the decision of the weblog's sentiment, we can transform the equation (1) by using Bayes rule to the following:

$$S^* = \arg\max_{s \in \{P,N\}} \sum_{z_k \in Z} P(s|z_k)P(d|z_k)P(z_k) \tag{2}$$

In equation (1), $P(z|d)$ represents how much a latent variable z occupies in the sentiment of a weblog d, intuitively, and $P(s|z)$ represents a semantic orientation of a latent variable z.

For the sentiment classification of the weblogs, the parameters of SentiAM are learnt through two steps. In the first step, the generative probabilities ($P(d|z)$, $P(w|z)$ and $P(z)$) are calculated based on the aspect model. Next step finds the probabilities of semantic orientations $P(s|z)$ of latent variables.

3.2 Generative Probabilities

According to Figure 1 (b), assuming that the weblog d and the word w are conditionally independent given the latent variable z, the generative process of weblog-word pair (d, w) is defined by the mixture:

$$P(d,w) = \sum_{z \in Z} P(z)P(d|z)P(w|z) \tag{3}$$

We use Expectation-Maximization [18] algorithm to estimate the parameters, $P(d|z)$, $P(w|z)$ and $P(z)$, which maximizes the below likelihood function.

$$P(D,W) = \prod_{d \in D} \prod_{w \in W} P(d,w)^{n(d,w)} \tag{4}$$

EM algorithm involves an iterative process with alternating steps, expectation step and maximization step, in order to learn the parameters which maximize complete likelihood.

In the expectation step (E-Step), the posterior probabilities are computed for the latent variables:

$$P(z|d,w) = \frac{P(z)P(d|z)P(w|z)}{\sum_{z' \in Z} P(z')P(d|z')P(w|z')} \tag{5}$$

In the maximization step (M-Step), the parameters are updated:

$$P(w|z) = \frac{\sum_{d\in D} n(d,w)P(z|d,w)}{\sum_{d\in D}\sum_{w'\in W} n(d,w')P(z|d,w')} \tag{6}$$

$$P(d|z) = \frac{\sum_{w\in W} n(d,w)P(z|d,w)}{\sum_{d'\in D}\sum_{w\in W} n(d',w)P(z|d',w)} \tag{7}$$

$$P(z) = \frac{\sum_{d\in D}\sum_{w\in W} n(d,w)P(z|d,w)}{\sum_{d\in D}\sum_{w\in W} n(d,w)} \tag{8}$$

The parameters, $P(d|z)$ and $P(z)$, learnt by the EM algorithm are used to classify sentiments of the weblogs, and the parameter $P(w|z)$ is used to find the probabilities of semantic orientations of latent variables $P(s|z)$.

3.3 Finding Domain-Dependent Semantic Orientations of Latent Variables

We use a lexicon resource tagged with positive and negative semantic orientations in order to find domain-dependent semantic orientations of the latent variables.

In our approach, SentiWordNet [19] is used as a polarity tagged lexicon. SentiWordNet is a lexical resource in which each synset of WORDNET (version 2.0) is associated to three numerical scores $Obj(s)$, $Pos(s)$ and $Neg(s)$. And it describes how Objective, Positive and Negative the terms contained in the synset are. For each word in SentiWordNet, positive, negative and neutral scores are defined as probabilities ($Obj+Pos+Neg=1$). Therefore we can readily find probabilities of semantic orientations of words, and we use them to find semantic orientations of latent variables.

The semantic orientations of the latent variables are calculated as the following:

$$P(S=P|z) = \sum_{w\in W} P_{IND}(S=P|w)P(w|z) \tag{9}$$

$$P(S=N|z) = \sum_{w\in W} P_{IND}(S=N|w)P(w|z) \tag{10}$$

Given a word, $P_{IND}(\cdot|w)$ represents probability of domain-independent semantic orientation defined in SentiWordNet. SentiWordNet defines semantic orientation of a word according to its sense. Therefore, semantic orientation of a word could become different if its sense is changed. We use the first sense of the word to simplify the experiment, because WSD (Word Sense Disambiguation) is another issue and out of the scope of this paper.

In equation (9) and (10), the probability of semantic orientation of latent variable z is the expectation of the probabilities of semantic orientations of all words which are generated by the latent variable z. The expectation makes the semantic orientation of the latent variable to be domain-dependent.

In the aspect model, different latent factors correspond with different topics, and words generated with high probability by a particular latent factor are related with topics captured by the latent factor. Due to these properties, different meanings (dependent on the context) of a polysemous word are identified by different latent factors. Similarly, different latent variables of SentiAM correspond with different opinions and sentiments contained in weblogs, and words generated with high probability by a particular latent factor have similar semantic orientations. For example, “unpredictable” is defined as negative semantic orientation with 0.625 probability by SentiWordNet. On the other hand, “unpredictable” is used with positive semantic orientation in the movie domain. Therefore a latent variable which generates “unpredictable” with high probability also generates many other words which have positive semantic orientation. Finally, the use of the expectation of semantic orientations of words makes the latent variable to assign the semantic orientation domain-dependently. Through experimenting on several different domains, we can verify the effect of the expectation of the semantic orientations.

3.4 Features Selection

In reality, SentiAM gives poor results when we use all words in the weblogs to train parameters of SentiAM. Weblogs contain not only the writers’ opinions but also factual descriptions, and this damages the effectiveness of the modeling. This is because the opinionated words and non-opinionated words may co-occur with each other, thus they will not be separated by the EM algorithm. This inhibits latent variables from properly modeling weblogs’ opinions. We consider feature selection to solve the problem. Instead of considering all words present in the weblogs, we attempt to seek out appraisal words to describe weblogs.

To this end, we choose subjective words in SentiWordNet ($Pos+Neg > 0.0$) as candidate words. To investigate the influence of feature selection on sentiment classification accuracy, we make four feature types using candidate words.

Type A consists of adjectives only. **Type B** is comprised of adjectives and verbs. **Type C** consists of adjectives and adverbs. **Type D** is comprised of all feature candidates, ie. nouns, verbs, adjectives and adverbs.

4 Experiments and Results

We conducted several experiments on three different domains. We tried to verify that latent variables are able to model domain-dependent semantic orientation correctly in each domain and that the use of the latent variables are effective for sentiment classification of the weblogs.

4.1 Data Sets

We used reviews in the domains of Movie, Automobiles and Digital Cameras. As a movie review dataset, the polarity dataset of Pang and Lee [2] was used. It contains 1000 positive and 1000 negative reviews all written before 2002 with a cap of 20 reviews per author per category. For an automobile review dataset, Epinion’s[1] result set from ”Cars & Motorsports” section for a keyword automobile was used. For a digital cameras dataset, Epinion’s result set from ”Digital Cameras” section for keywords Canon, Nikon and Pentax was used. Among the reviews, short reviews with less than 50 words were excluded. Epinions review system displays authors’ rating with five stars.

[1] http://epinions.com

Table 1. Test Data Sets. # Type A: the number of Type A (adjectives); # Type B: the number of Type B (adjectives+verbs); # Type C: the number of Type C (adjectives+adverbs); # Type D: the number of Type D (nouns+verbs+adjectives+adverbs); Auto.: Automobiles domain; Dig. Cam.: Digital Cameras domain.

Domain	# Pos.	# Neg.	# Avg. words	# Type A	# Type B	# Type C	# Type D
Movie	1000	1000	746.33	4174	5391	5406	9721
Auto.	130	130	613.28	1114	1528	1511	2616
Dig. Cam.	200	200	131.78	518	706	704	1213

We regard reviews with more than 3 stars (≥ 4) as positive reviews and reviews with less than 3 stars (≤ 2) as negative reviews. The number of positive reviews obtained from epinions significantly outnumbers the number of negative reviews. Therefore, we randomly chose positive reviews to match the number of negative reviews. In general, the use of a different number of positive and negative reviews as test dataset can lead to a biased result. Once the two groups are selected in equal numbers, the reviews are parsed for html-tags removal and POS tagging by Stanford Tagger. No stemming and stopword lists were used. All feature words between a negation word (”not“, ”isn’t“, etc.) and the first punctuation mark following the negation word is marked with NOT tag [11]. Table 1 shows information of the dataset which we used for experiment.

4.2 Baseline

For the baseline of the experiment, we classified weblogs’ sentiments using semantic orientations of subjective words within SentiWordNet. For each word, SentiWordNet defines its semantic orientation as a score (probability). We regard a review as a positive review when the sum of positive scores of subjective words within a review is greater than that of negative scores, otherwise we regard it as a negative review.

The baseline precisions of the movie domain, the automobile domain and the digital camera domain are 0.44, 0.63 and 0.62, respectively. We can see that the precision was worse than that of random selection in movie domain. This result shows the difficulty of sentiment classification in the movie domain. This result is because there are many words that have a different semantic orientation compared to domain-independent semantic orientation defined at SentiWordNet such as “unpredictable” in the movie domain. This result shows the problem when using domain-independent semantic orientations of words or phrases. Therefore domain-dependent semantic orientation of latent variables can remedy this problem.

4.3 Semantic Orientation of Latent Variables

We wanted to verify that latent variables capture domain-dependent semantic orientation correctly. To do this, we used words such as “unpredictable” that have different semantic orientations depending on the domain where they are used. In SentiWordNet, “unpredictable” has a negative semantic orientation with its 0.625 score. For each domain, we verified semantic orientations of latent variables which generate “unpredictable” with high probability. Table 2 shows semantic orientations of two

Table 2. Semantic orientation of the two latent variables that produces the word "unpredictable" with the highest probability

		Movie	Auto.	Dig. Cam.
1st	$P(Pos\|z)$	0.27	0.23	0.23
	$P(Neg\|z)$	0.23	0.27	0.23
2nd	$P(Pos\|z)$	0.26	0.23	0.23
	$P(Neg\|z)$	0.24	0.24	0.22

latent variables which generate "unpredictable" with the highest probability. The results of movie and automobile datasets correspond to our expectation. Even though the domain-independent semantic orientation of the word "unpredictable" is negative, the two latent variables that generate the word the most have positive semantic orientations in the movie domain.

We expected that latent variables generating the word "unpredictable" with high probability would have a negative semantic orientation in the digital cameras dataset. To the contrary of our expectation, positive semantic orientation is higher or equal to negative semantic orientation. We think this is due to the small size we had for SentiAM to learn parameters $P(z)$, $P(d|z)$, $P(w|z)$ and $P(s|z)$ through the EM algorithm. As shown in the Table 1, the size of the digital cameras dataset is noticeably smaller than the other two. In fact the word, "unpredictable", occurs only once in 400 reviews in the digital cameras dataset upon a closer inspection. This might lead to improper parameter training of SentiAM.

Table 3. Words with different semantic orientations in the movie domain. IND. SO: (positive/negative) semantic orientation pair defined at SentiWordNet; DEP. SO: (positive/negative) semantic orientation pair of latent variables which generate the word with the highest probability; CONTEXT: a segment of text in which the word appears.

WORD	IND. SO	DEP. SO	CONTEXT
otherwordly	0.00/0.88	0.31/0.21	was a striking mood piece with an otherworldly feel
unconventional	0.00/0.88	0.27/0.23	that movie worked because it took an unconventional story
ad-lib	0.00/0.75	0.31/0.21	by people with little gift for ad-lib
sarcastic	0.00/0.75	0.27/0.23	succeeds in giving his character a sarcastic sense of humor
unexpected	0.00/0.63	0.26/0.23	there is plenty of action and an unexpected dose of humour
unpredictable	0.00/0.63	0.27/0.23	the ending is unpredictable and incredibly satisfying

We found words in the movie domain with semantic orientations different to domain-independent ones. Table 3 lists some of them with their semantic orientations and contexts.

The experiment shows that SentiAM assigns domain-dependent semantic orientation to latent variables correctly. The experiment presented in the next section shows that these latent variables also outperform other types of features such as words and phrases in the sentiment classification of the weblogs.

Table 4. Best Overall Precision. Prec.: precision; Rec.: recall.

Domain	Type	Prec.	Rec.	Feature
Movie	Positive	0.67	0.91	Type A
	Negative	0.85	0.56	
	Overall	0.74	0.74	
Auto.	Positive	0.61	0.85	Type A
	Negative	0.75	0.45	
	Overall	0.65	0.65	
Dig. Cam.	Positive	0.59	0.80	Type A
	Negative	0.62	0.42	
	Overall	0.62	0.61	

4.4 Sentiment Classification

In this section, we investigate the influence of various feature sets (Type A, Type B, Type C, and Type D) and the number of latent variables on accuracy of weblogs' sentiment classification.

We used the same number of positive reviews as negative reviews. Therefore, a random guessing sentiment of the reviews will have 50% accuracy.

Table 4 shows the best performance we achieved for each domain and the feature set used. In the movie domain, we obtained 24% higher precision than random guessing and 30% higher precision than the baseline. In the automobile domain, we also obtained 15% and 2% higher precision than random guessing and the baseline. In the digital camera domain, we obtained 12% better precision than random guessing but the same precision with the baseline.

We think that the small improvement of accuracy in automobiles and digital cameras compared to that in the movie domain is due to the small size of the two domains' test corpora. As we saw in section 4.3, the bigger the size of the dataset, the better SentiAM assigns domain-dependent semantic orientations of latent variables, thus resulting in higher accuracy.

Feature Type A which consists of adjectives only, resulted in the best accuracy in all domains. Many previous works [5, 20, 3] related to sentiment analysis dealt with adjectives heavily. Adjectives are less ambiguous with respect to sentiment than other word classes. Therefore, they can convey the author's opinion more correctly than any other parts of speech.

Figure 2 (a) shows the influence of varying feature types on sentiment classification of the weblogs in the movie domain. The two other domains yield similar results and the following analysis on the movie domain also applies to them.

Type A and Type B produced results with higher accuracy than the others. This is because adjectives and verbs reflect the authors' opinions less ambiguously with respect to sentiment than other word classes. We obtained the worst accuracy when using Type D which includes nouns. Nouns are likely to be non-opinionated or less-opinionated. Thus including them in the feature set might generate noise when latent variables separate the different opinions and sentiments of the weblogs. We have not yet investigated the influence of nouns on sentiment classification. We leave this issue for future work.

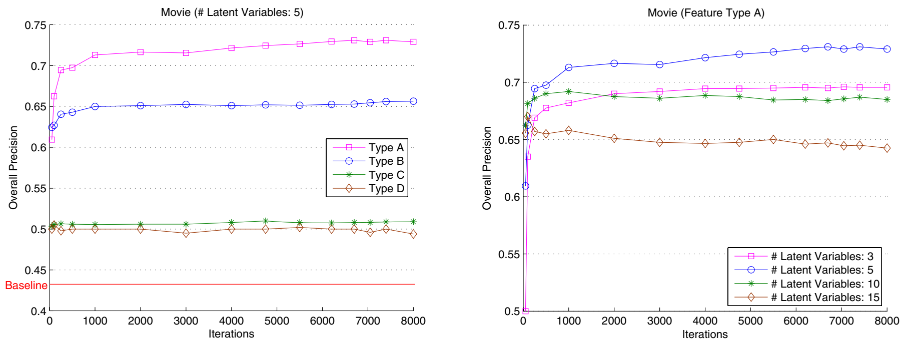

(a) Feature types (b) Latent variables

Fig. 2. In the movie doman, overall precision according to feature types and latent variables

Figure 2 (b) shows the accuracy of sentiment classification according to different numbers of latent variables, $z = 3, 5, 10, 15$. We obtained the best accuracy using five latent variables, the worst accuracy using fifteen latent variables. Too many latent variables overfits the data set and does not produce either positive or negative polarity. On the other hand, too few latent variables might not be able to capture semantic orientation with enough accuracy and thus reduces the accuracy of sentiment classifier as well.

5 Conclusions and Future Work

In this paper, we have presented a novel probabilistic approach to classify sentiments of weblogs. Our model, SentiAM, is an unsupervised learning model based on the aspect model. SentiAM assigns domain-dependent semantic orientations to latent variables using resources in which semantic orientations of words or phrases are tagged domain-independently. Experimental results confirm that semantic orientations of latent variables are effective at classifying the sentiments of weblogs.

We have presented a general framework to capture domain-dependent semantic orientations of latent variables. Though we used SentiWordNet as a polarity tagged lexicon resource, any resource can be applied. Therefore, semantic orientations of lexicons exploited by previous work can be also used by our framework.

We used semantic orientation of the first sense of a word in SentiWordNet. Knowledge about word sense can improve accuracy of sentiment classification. We could use WSD or any means to improve word sense detection, thus improving the accuracy of sentiment classification.

We use semantic orientations of latent variables to classify sentiments of weblogs. Our framework could, of course, also be applied to analyzing semantic orientations of more fine units such as sentences and words. We leave this issue for our future work.

Acknowledgement. This work was supported in part by MKE & IITA through IT Leading R&D Support Project and also in part by the BK 21 Project in 2008.

References

1. Cardie, C., Wiebe, J., Wilson, T., Litman, D.: Combining low-level and summary representations of opinions for multiperspective question answering. In: Working Notes of the 2003 AAAI Spring Symposium on New Directions in Question Answering, pp. 20–27 (2003)
2. Pang, B., Lee, L.: A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In: Proceedings of the 42nd Meeting of the Association for Computational Linguistics (ACL 2004), Barcelona, Spain, pp. 271–278 (July 2004)
3. Hatzivassiloglou, V., McKeown, K.R.: Predicting the semantic orientation of adjectives. In: Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics, Madrid, Spain, pp. 174–181. Association for Computational Linguistics (July 1997)
4. Wiebe, J.M.: Learning subjective adjectives from corpora. In: Proceedings of AAAI 2000, pp. 735–740 (2000)
5. Turney, P.: Thumbs up or thumbs down? semantic orientation applied to unsupervised classification of reviews. In: Proceedings of 40th Annual Meeting of the Association for Computational Linguistics, Philadelphia, Pennsylvania, USA, pp. 417–424. Association for Computational Linguistics (July 2002)
6. Kim, S.M., Hovy, E.: Automatic detection of opinion bearing words and sentence. In: Proceedings of IJCNLP 2005, pp. 61–66 (2005)
7. Riloff, E., Wiebe, J., Wilson, T.: Learning subjective nouns using extraction pattern bootstrapping. In: Proceedings of CoNLL 2003, pp. 25–32 (2003)
8. Wilson, T., Wiebe, J., Hoffmann, P.: Recognizing contextual polarity in phrase-level sentiment analysis. In: Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing, Vancouver, British Columbia, Canada, pp. 347–354. Association for Computational Linguistics (October 2005)
9. Kamps, J., Marx, M., Mokken, R.J., de Rijke, M.: Using wordnet to measure semantic orientations of adjectives. In: Proceedings of LREC 2004, pp. 1115–1118 (2004)
10. Takamura, H., Inui, T., Okumura, M.: Latent variables models for semantic orientations of phrase. In: Proceedings of EACL 2006, pp. 201–208 (2006)
11. Pang, B., Lee, L., Vaithyanathan, S.: Thumbs up? sentiment classification using machine learning techniques. In: Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing, pp. 79–86. Association for Computational Linguistics (July 2002)
12. Whitelaw, C., Garg, N., Argamon, S.: Using appraisal groups for sentiment analysis. In: Proceedings of CIKM 2005, pp. 625–631. ACM, New York (2005)
13. Morinaga, S., Yamanishi, K., Tateishi, K., Fukushima, T.: Mining product reputations on the web. In: KDD 2002: Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 341–349. ACM, New York (2002)
14. Devitt, A., Ahmad, K.: Sentiment polarity identification in financial news: A cohesion-based approach. In: Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics, Prague, Czech Republic, pp. 984–991. Association for Computational Linguistics (June 2007)
15. Mei, Q., Ling, X., Wondra, M., Su, H., Zhai, C.: Topic sentiment mixture: modeling facets and opinions in weblogs. In: WWW 2007: Proceedings of the 16th international conference on World Wide Web, pp. 171–180. ACM, New York (2007)
16. Liu, Y., Huang, X., An, A., Yu, X.: Arsa: a sentiment-aware model for predicting sales performance using blogs. In: Proceedings of SIGIR 2007, pp. 607–614. ACM, New York (2007)

17. Hofmann, T., Puzicha, J., Jordan, M.I.: Unsupervised learning from dyadic data. In: Advances in Neural Information Processing Systems (1999)
18. Dempster, A.P., Laird, N., Rubin, D.: Maximum likelihood from incomplete data via the em algorithm. Journal of the Royal Statistical Society 39(B), 1–38 (1997)
19. Esuli, A., Sebastiani, F.: Sentwordnet: A publicly available lexical resource for opinion mining. In: Proceedings of LREC 2006, pp. 417–422 (2006)
20. Hatzivassiloglou, V., Wiebe, J.: Effects of adjective orientation and gradability on sentence subjectivity. In: Proceedings of COLING 2000, pp. 299–305 (2000)