

Adapting Sentiment Lexicons using Contextual Semantics for Sentiment Analysis of Twitter

Hassan Saif,¹ Yulan He,² Miriam Fernandez¹ and Harith Alani¹

¹ Knowledge Media Institute, The Open University, United Kingdom
{h.saif, m.fernandez, h.alani}@open.ac.uk

² School of Engineering and Applied Science, Aston University, United Kingdom
y.he@cantab.net

Abstract. Sentiment lexicons for sentiment analysis offer a simple, yet effective way to obtain the prior sentiment information of opinionated words in texts. However, words' sentiment orientations and strengths often change throughout various contexts in which the words appear. In this paper, we propose a lexicon adaptation approach that uses the contextual semantics of words to capture their contexts in tweet messages and update their prior sentiment orientations and/or strengths accordingly. We evaluate our approach on one state-of-the-art sentiment lexicon using three different Twitter datasets. Results show that the sentiment lexicons adapted by our approach outperform the original lexicon in accuracy and F-measure in two datasets, but give similar accuracy and slightly lower F-measure in one dataset.

Keywords: Sentiment Analysis, Semantics, Lexicon Adaptation, Twitter

1 Introduction

Sentiment analysis on Twitter has been attracting much attention recently due to the rapid growth in Twitter's popularity as a platform for people to express their opinions and attitudes towards a great variety of topics. Most existing approaches to Twitter sentiment analysis can be categorised into machine learning [7, 11, 13] and lexicon-based approaches [2, 8, 15, 6].

Lexicon-based approaches use lexicons of words weighted with their sentiment orientations to determine the overall sentiment in texts. These approaches have shown to be more applicable to Twitter data than machine learning approaches, since they do not require training from labelled data and therefore, they offer a domain-independent sentiment detection [15]. Nonetheless, lexicon-based approaches are limited by the sentiment lexicon used [21]. Firstly, because sentiment lexicons are composed by a generally static set of words that do not cover the wide variety of new terms that constantly emerge in the social web. Secondly, because words in the lexicons have fixed prior sentiment orientations, i.e. each term has always the same associated sentiment orientation independently of the context in which the term is used.

To overcome the above limitations, several lexicon bootstrapping and adaptation methods have been previously proposed. However, these methods are either supervised [16], i.e., they require training from human-coded corpora, or based on studying the statistical, syntactical or linguistic relations between words in general textual corpora (e.g., The Web) [17, 19] or in static lexical knowledge sources (e.g., WordNet) [5]

ignoring, therefore, the specific textual context in which the words appear. In many cases, however, the sentiment of a word is implicitly associated with the semantics of its context [3].

In this paper we propose an unsupervised approach for adapting sentiment lexicons based on the contextual semantics of their words in a tweet corpus. In particular, our approach studies the co-occurrences between words to capture their contexts in tweets and update their prior sentiment orientations and/or sentiment strengths in a given lexicon accordingly.

As a case study we apply our approach on Thelwall-Lexicon [15], which, to our knowledge, is the state-of-the-art sentiment lexicon for social data. We evaluate the adapted lexicons by performing a lexicon-based polarity sentiment detection (positive vs. negative) on three Twitter datasets. Our results show that the adapted lexicons produce a significant improvement in the sentiment detection accuracy and F-measure in two datasets but gives a slightly lower F-measure in one dataset.

In the rest of this paper, related work is discussed in Section 2, and our approach is presented in Sections 3. Experiments and results are presented in Sections 4. Discussion and future work are covered in Section 5. Finally, we conclude our work in Section 6.

2 Related Work

Existing approaches to bootstrapping and adapting sentiment lexicons can be categorised into dictionary and corpus-based approaches. The dictionary-based approach [5, 14] starts with a small set of general opinionated words (e.g., good, bad) and lexical knowledge base (e.g., WordNet). After that, the approach expands this set by searching the knowledge base for words that have lexical or linguistic relations to the opinionated words in the initial set (e.g., synonyms, glosses, etc).

Alternatively, the corpus-based approach measures the sentiment orientation of words automatically based on their association to other strongly opinionated words in a given corpus [17, 14, 19]. For example, Turney and Littman [17] used *Pointwise Mutual Information* (PMI) to measure the statistical correlation between a given word and a balanced set of 14 positive and negative paradigm words (e.g., good, nice, nasty, poor). Although this work does not require large lexical input knowledge, its identification speed is very limited [21] because it uses web search engines in order to retrieve the relative co-occurrences of words.

Following the aforementioned approaches, several lexicons such as MPQA [20] and SentiWordNet [1] have been induced and successfully used for sentiment analysis on conventional text (e.g., movie review data). However, on Twitter these lexicons are not as compatible due to their limited coverage of Twitter-specific expressions, such as abbreviations and colloquial words (e.g., “l000v”, “luv”, “gr8”) that are often found in tweets.

Quite few sentiment lexicons have been recently built to work specifically with social media data, such as Thelwall-Lexicon [16] and Nielsen-Lexicon [8]. These lexicons have proven to work effectively on Twitter data. Nevertheless, such lexicons are similar to other traditional ones, in the sense that they all offer fixed and context-insensitive word-sentiment orientations and strengths. Although a training algorithm has been proposed to update the sentiment of terms in Thelwall-Lexicon[16], it requires to be trained from human-coded corpora, which is labour-intensive to obtain.

Aiming at addressing the above limitations we have designed our lexicon-adaptation approach in away that allows to (i) work in unsupervised fashion, avoiding the need for labelled data, and (ii) exploit the contextual semantics of words. This allows capturing their contextual information in tweets and update their prior sentiment orientation and strength in a given sentiment lexicon accordingly.

3 A Contextual Semantic Approach to Lexicon Adaptation

The main principle behind our approach is that the sentiment of a term is not static, as found in general-purpose sentiment lexicons, but rather depends on the context in which the term is used, i.e., it depends on its contextual semantics.³ Therefore, our approach functions in two main steps as shown in Figure 1. First, given a tweet collection and a sentiment lexicon, the approach builds a contextual semantic representation for each unique term in the tweet collection and subsequently uses it to derive the term’s contextual sentiment orientation and strength. The SentiCircle representation model is used to this end [10]. Secondly, rule-based algorithm is applied to amend the prior sentiment of terms in the lexicon based on their corresponding contextual sentiment. Both steps are further detailed in the following subsections.




Fig. 1. The systematic workflow of our proposed lexicon adaptation approach.

3.1 Capturing Contextual Semantics and Sentiment

The first step in our pipeline is to capture the words contextual semantics and sentiment in tweets. To this end, we use our previously proposed semantic representation model, SentiCircle [10].

Following the distributional hypothesis that words that co-occur in similar contexts tend to have similar meaning [18], SentiCircle extracts the contextual semantics of a word from its co-occurrence patterns with other words in a given tweet collection. These patterns are then represented as a geometric circle, which is subsequently used to compute the contextual sentiment of the word by applying simple trigonometric identities on it. In particular, for each unique term m in a tweet collection, we build a two-dimensional geometric circle, where the term m is situated in the centre of the circle, and each point around it represents a context term c_i (i.e., a term that occurs with m in the same context). The position of c_i , as illustrated in Figure 2, is defined jointly by its Cartesian coordinates x_i, y_i as:

$$x_i = r_i \cos(\theta_i * \pi) \quad y_i = r_i \sin(\theta_i * \pi)$$

Where θ_i is the polar angle of the context term c_i and its value equals to the prior sentiment of c_i in a sentiment lexicon before adaptation, r_i is the radius of c_i and its value represents the degree of correlation (tdoc) between c_i and m , and can be computed as:

$$r_i = tdoc(m, c_i) = f(c_i, m) \times \log \frac{N}{N_{c_i}}$$

³ We define context as a textual corpus or a set of tweets.

where $f(c_i, m)$ is the number of times c_i occurs with m in tweets, N is the total number of terms, and N_{c_i} is the total number of terms that occur with c_i . Note that all terms' radii in the SentiCircle are normalised. Also, all angles' values are in radian. The trigonometric properties of the SentiCircle allows us to encode the contextual semantics of a term as sentiment orientation and sentiment strength. Y-axis defines the sentiment of the term, i.e., a positive y value denotes a positive sentiment and vice versa. The X-axis defines the sentiment strength of the term. The smaller the x value, the stronger the sentiment.⁴ This, in turn, divides the circle into four sentiment quadrants. Terms in the two upper quadrants have a positive sentiment ($\sin \theta > 0$), with upper left quadrant representing stronger positive sentiment since it has larger angle values than those in the top right quadrant. Similarly, terms in the two lower quadrants have negative sentiment values ($\sin \theta < 0$). Moreover, a small region called the “Neutral Region” can be defined. This region, as shown in Figure 2, is located very close to X-axis in the “Positive” and the “Negative” quadrants only, where terms lie in this region have very weak sentiment (i.e., $|\theta| \approx 0$).




Fig. 2. SentiCircle of a term m . Neutral region is shaded in blue.

Calculating Contextual Sentiment In summary, the Senti-Circle of a term m is composed by the set of (x, y) Cartesian coordinates of all the context terms of m . An effective way to compute the overall sentiment of m is by calculating the geometric median of all the points in its SentiCircle. Formally, for a given set of n points (p_1, p_2, \dots, p_n) in a SentiCircle Ω , the 2D geometric median g is defined as: $g = \arg \min_{g \in \mathbb{R}^2} \sum_{i=1}^n \|p_i - g\|_2$. We call the geometric median g the **SentiMedian** as its position in the SentiCircle determines the final contextual-sentiment orientation and strength of m .

Note that the boundaries of the neutral region can be computed by measuring the density distribution of terms in the SentiCircle along the Y-axis. In this paper we use similar boundaries to the ones used in [10] since we use the same evaluation datasets.

3.2 Lexicon Adaptation

The second step in our approach is to update the sentiment lexicon with the terms' contextual sentiment information extracted in the previous step. As mentioned earlier, in this work we use Thelwall-Lexicon [16] as a case study. Therefore, in this section we first describe this lexicon and its properties, and then introduce our proposed adaptation method.

Thelwall-Lexicon consists of 2546 terms coupled with integer values between -5 (very negative) and +5 (very positive). Based on the terms' prior sentiment orientations and strengths (SOS), we group them into three subsets of 1919 negative terms ($SOS \in [-2, -5]$), 398 positive terms ($SOS \in [2, 5]$) and 229 neutral terms ($SOS \in \{-1, 1\}$).

The adaptation method uses a set of antecedent-consequent rules that decides how the prior sentiment of the terms in Thelwall-Lexicon should be updated according to the positions of their SentiMedians (i.e., their contextual sentiment). In particular, for a term m , the method checks (i) its prior SOS value in Thelwall-Lexicon and (ii) the SentiCircle

⁴ This is because $\cos \theta < 0$ for large angles.

quadrant in which the SentiMedian of m resides. The method subsequently chooses the best-matching rule to update the term’s prior sentiment and/or strength.

Table 1 shows the complete list of rules in the proposed method. As noted, these rules are divided into *updating rules*, i.e., rules for updating the existing terms in Thelwall-Lexicon, and *expanding rules*, i.e., rules for expanding the lexicon with new terms. The *updating rules* are further divided into rules that deal with terms that have similar prior and contextual sentiment orientations (i.e., both positive or negative), and rules that deal with terms that have different prior and contextual sentiment orientations (i.e., negative prior, positive contextual sentiment and vice versa).

Although they look complicated, the notion behind the proposed rules is rather simple: *Check how strong the contextual sentiment is and how weak the prior sentiment is \rightarrow update the sentiment orientation and strength accordingly.* The strength of the contextual sentiment can be determined based on the sentiment quadrant of the SentiMedian of m , i.e., the contextual sentiment is strong if the SentiMedian resides in the “*Very Positive*” or “*Very Negative*” quadrants (See Figure 2). On the other hand, the prior sentiment of m (i.e., $prior_m$) in Thelwall-Lexicon is weak if $|prior_m| \leq 3$ and strong otherwise.

Updating Rules (Similar Sentiment Orientations)		
Id	Antecedents	Consequent
1	$(prior \leq 3) \wedge (SentiMedian \notin StrongQuadrant)$	$ prior = prior + 1$
2	$(prior \leq 3) \wedge (SentiMedian \in StrongQuadrant)$	$ prior = prior + 2$
3	$(prior > 3) \wedge (SentiMedian \notin StrongQuadrant)$	$ prior = prior + 1$
4	$(prior > 3) \wedge (SentiMedian \in StrongQuadrant)$	$ prior = prior + 1$
Updating Rules (Different Sentiment Orientations)		
5	$(prior \leq 3) \wedge (SentiMedian \notin StrongQuadrant)$	$ prior = 1$
6	$(prior \leq 3) \wedge (SentiMedian \in StrongQuadrant)$	$prior = -prior$
7	$(prior > 3) \wedge (SentiMedian \notin StrongQuadrant)$	$ prior = prior - 1$
8	$(prior > 3) \wedge (SentiMedian \in StrongQuadrant)$	$prior = -prior$
9	$(prior > 3) \wedge (SentiMedian \in NeutralRegion)$	$ prior = prior - 1$
10	$(prior \leq 3) \wedge (SentiMedian \in NeutralRegion)$	$ prior = 1$
Expanding Rules		
11	$SentiMedian \in NeutralRegion$	$(contextual = 1) \wedge AddTerm$
12	$SentiMedian \notin StrongQuadrant$	$(contextual = 3) \wedge AddTerm$
13	$SentiMedian \in StrongQuadrant$	$(contextual = 5) \wedge AddTerm$

Table 1. Adaptation rules for Thelwall-Lexicon, where $prior$: prior sentiment value, $StrongQuadrant$: very negative/positive quadrant in the SentiCircle, Add : add the term to Thelwall-Lexicon.

For example, the word “*revolution*” in Thelwall-Lexicon has a weak negative sentiment ($prior=-2$) while it has a neutral contextual sentiment since its SentiMedian resides in the neutral region ($SentiMedian \in NeutralRegion$). Therefore, rule number 10 is applied and the term’s prior sentiment in Thelwall lexicon will be updated to neutral ($|prior| = 1$). In another example, the words “*Obama*” and “*Independence*” are not covered by the Thelwall-Lexicon, and therefore, they have no prior sentiment. However, their SentiMedians reside in the “*Positive*” quadrant in their SentiCircles, and therefore rule number 12 is applied and both terms will be assigned with a positive sentiment strength of 3 and added to the lexicon consequently.

4 Evaluation Results

We evaluate our approach on Thelwall-Lexicon using three adaptation settings: (i) the *update* setting where we update the prior sentiment of existing terms in the lexicon, (ii) The *expand* setting where we expand Thelwall-Lexicon with new opinionated terms, and (iii) the *update+expand* setting where we try both aforementioned settings together. To

this end, we use three Twitter datasets OMD, HCR and STS-Gold. Numbers of positive and negative tweets within these datasets are summarised in Table 2, and detailed in the references added in the table. To evaluate the adapted lexicons under the above settings, we perform binary polarity classification on the three datasets. To this end, we use the sentiment detection method proposed with Thelwall-Lexicon [15]. According to this method a tweet is considered as positive if its aggregated positive sentiment strength is 1.5 times higher than the aggregated negative one, and negative vice versa.

Dataset	Tweets	Positive	Negative
<i>Obama-McCain Debate (OMD)</i> [4]	1081	393	688
<i>Health Care Reform (HCR)</i> [12]	1354	397	957
<i>Stanford Sentiment Gold Standard (STS-Gold)</i> [9]	2034	632	1402

Table 2. Twitter datasets used for the evaluation

Applying our adaptation approach to Thelwall-Lexicon results in dramatic changes in it. Table 3 shows the percentage of words in the three datasets that were found in Thelwall-Lexicon with their sentiment changed after adaptation. One can notice that on average 9.61% of the words in our datasets were found in the lexicon. However, updating the lexicon with the contextual sentiment of words resulted in 33.82% of these words flipping their sentiment orientation and 62.94% changing their sentiment strength while keeping their prior sentiment orientation. Only 3.24% of the words in Thelwall-Lexicon remained untouched. Moreover, 21.37% of words previously unseen in the lexicon were assigned with contextual sentiment by our approach and added to Thelwall-Lexicon subsequently.

	OMD	HCR	STS-Gold	Average
Words found in the lexicon	12.43	8.33	8.09	9.61
Hidden words	87.57	91.67	91.91	90.39
Words flipped their sentiment orientation	35.02	35.61	30.83	33.82
Words changed their sentiment strength	61.83	61.95	65.05	62.94
Words remained unchanged	3.15	2.44	4.13	3.24
New opinionated words	23.94	14.30	25.87	21.37

Table 3. Average percentage of words in the three datasets that had their sentiment orientation or strength updated by our adaptation approach

Table 4 shows the average results of binary sentiment classification performed on our datasets using (i) the original Thelwall-Lexicon (*Original*), (ii) Thelwall-Lexicon induced under the *update* setting (*Updated*), and (iii) Thelwall-Lexicon induced under the *update+expand* setting.⁵ The table reports the results in accuracy and three sets of precision (P), recall (R), and F-measure (F1), one for positive sentiment detection, one for negative, and one for the average of the two.

From these results in Table 4, we notice that the best classification performance in accuracy and F1 is obtained on the STS-Gold dataset regardless the lexicon being used. We also observe that the negative sentiment detection performance is always higher than the positive detection performance for all datasets and lexicons.

As for different lexicons, we notice that on OMD and STS-Gold the adapted lexicons outperform the original lexicon in both accuracy and F-measure. For example, on OMD the adapted lexicon shows an average improvement of 2.46% and 4.51% in accuracy and F1 respectively over the original lexicon. On STS-Gold the performance improvement is

⁵ Note that in this work we do not report the results obtained under the expand setting since no improvement was observed comparing to the other two settings.

Datasets	Lexicons	Accuracy	Positive Sentiment			Negative Sentiment			Average		
			P	R	F1	P	R	F1	P	R	F1
OMD	Original	66.79	55.99	40.46	46.97	70.64	81.83	75.82	63.31	61.14	61.4
	Updated	69.29	58.89	51.4	54.89	74.12	79.51	76.72	66.51	65.45	65.8
	Updated+Expanded	69.2	58.38	53.18	55.66	74.55	78.34	76.4	66.47	65.76	66.03
HCR	Original	66.99	43.39	41.31	42.32	76.13	77.64	76.88	59.76	59.47	59.6
	Updated	67.21	42.9	35.77	39.01	75.07	80.25	77.58	58.99	58.01	58.29
	Updated+Expanded	66.99	42.56	36.02	39.02	75.05	79.83	77.37	58.8	57.93	58.19
STS-Gold	Original	81.32	68.75	73.1	70.86	87.52	85.02	86.25	78.13	79.06	78.56
	Updated	81.71	69.46	73.42	71.38	87.7	85.45	86.56	78.58	79.43	78.97
	Updated+Expanded	82.3	70.48	74.05	72.22	88.03	86.02	87.01	79.26	80.04	79.62

Table 4. Cross comparison results of original and the adapted lexicons

less significant than that on OMD, but we still observe 1% improvement in accuracy and F1 comparing to using the original lexicon. As for the HCR dataset, the adapted lexicon gives on average similar accuracy, but 1.36% lower F-measure. This performance drop can be attributable to the poor detection performance of positive tweets. Specifically, we notice from Table 4 a major loss in the recall on positive tweet detection using both adapted lexicons. One possible reason is the sentiment class distribution in our datasets. In particular, one may notice that HCR is the most imbalanced amongst the three datasets. Moreover, by examining the numbers in Table 3, we can see that HCR presents the lowest number of new opinionated words among the three datasets (i.e., 10.61% lower than the average) which could be another potential reason for not observing any performance improvement.

5 Discussion and Future Work

We demonstrated the value of using contextual semantics of words for adapting sentiment lexicons from tweets. Specifically, we used Thelwall-Lexicon as a case study and evaluated its adaptation to three datasets of different sizes. Although the potential is palpable, our results were not conclusive, where a performance drop was observed in the HCR dataset using our adapted lexicons. Our initial observations suggest that the quality of our approach might be dependent on the sentiment class distribution in the dataset. Therefore, a deeper investigation in this direction is required.

We used the SentiCircle approach to extract the contextual semantics of words from tweets. In future work we will try other contextual semantic approaches and study how the semantic extraction quality affects the adaptation performance.

Our adaptation rules in this paper are specific to Thelwall-Lexicon. These rules, however, can be generalized to other lexicons, which constitutes another future direction of this work.

All words which have contextual sentiment were used for adaptation. Nevertheless, the results conveyed that the prior sentiments in the lexicon might need to be unchanged for words of specific syntactical or linguistic properties in tweets. Part of our future work is to detect and filter those words that are more likely to have stable sentiment regardless the contexts in which they appear.

6 Conclusions

In this paper we proposed an unsupervised approach for sentiment lexicon adaptation from Twitter data. Our approach extracts the contextual semantics of words and uses them to update the words' prior sentiment orientations and/or strength in a given sentiment lexicon. The evaluation was done on Thelwall-Lexicon using three Twitter datasets. Results showed that lexicons adapted by our approach improved the sentiment classification performance in both accuracy and F1 in two out of three datasets.

Acknowledgment

This work was supported by the EU-FP7 project SENSE4US (grant no. 611242).

References

1. Baccianella, S., Esuli, A., Sebastiani, F.: Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In: Seventh conference on International Language Resources and Evaluation, Malta. Retrieved May. Valletta, Malta (2010)
2. Bollen, J., Mao, H., Zeng, X.: Twitter mood predicts the stock market. *Journal of Computational Science* 2(1), 1–8 (2011)
3. Cambria, E.: An introduction to concept-level sentiment analysis. In: *Advances in Soft Computing and Its Applications*, pp. 478–483. Springer (2013)
4. Diakopoulos, N., Shamma, D.: Characterizing debate performance via aggregated twitter sentiment. In: *Proc. 28th Int. Conf. on Human factors in computing systems*. ACM (2010)
5. Esuli, A., Sebastiani, F.: Determining term subjectivity and term orientation for opinion mining. In: *EACL*. vol. 6, p. 2006 (2006)
6. Hu, X., Tang, J., Gao, H., Liu, H.: Unsupervised sentiment analysis with emotional signals. In: *Proceedings of the 22nd World Wide Web conf* (2013)
7. Kouloumpis, E., Wilson, T., Moore, J.: Twitter sentiment analysis: The good the bad and the omg! In: *Proceedings of the ICWSM*. Barcelona, Spain (2011)
8. Nielsen, F.Å.: A new anew: Evaluation of a word list for sentiment analysis in microblogs. *arXiv preprint arXiv:1103.2903* (2011)
9. Saif, H., Fernandez, M., He, Y., Alani, H.: Evaluation datasets for twitter sentiment analysis a survey and a new dataset, the sts-gold. In: *Proceedings, 1st ESSEM Workshop*. Turin, Italy (2013)
10. Saif, H., Fernandez, M., He, Y., Alani, H.: Senticircles for contextual and conceptual semantic sentiment analysis of twitter. In: *Proc. 11th Extended Semantic Web Conf. (ESWC)*. Crete, Greece (2014)
11. Saif, H., He, Y., Alani, H.: Semantic sentiment analysis of twitter. In: *Proc. 11th Int. Semantic Web Conf. (ISWC)*. Boston, MA (2012)
12. Speriosu, M., Sudan, N., Upadhyay, S., Baldrige, J.: Twitter polarity classification with label propagation over lexical links and the follower graph. In: *Proceedings of the EMNLP First workshop on Unsupervised Learning in NLP*. Edinburgh, Scotland (2011)
13. Suttles, J., Ide, N.: Distant supervision for emotion classification with discrete binary values. In: *Computational Linguistics and Intelligent Text Processing*, pp. 121–136. Springer (2013)
14. Takamura, H., Inui, T., Okumura, M.: Extracting semantic orientations of words using spin model. In: *Proc. 43rd Annual Meeting on Association for Computational Linguistics* (2005)
15. Thelwall, M., Buckley, K., Paltoglou, G.: Sentiment strength detection for the social web. *J. American Society for Information Science and Technology* 63(1), 163–173 (2012)
16. Thelwall, M., Buckley, K., Paltoglou, G., Cai, D., Kappas, A.: Sentiment strength detection in short informal text. *J. American Society for Info. Science and Technology* 61(12) (2010)
17. Turney, P., Littman, M.: Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems* 21, 315–346 (2003)
18. Turney, P.D., Pantel, P., et al.: From frequency to meaning: Vector space models of semantics. *Journal of artificial intelligence research* 37(1), 141–188 (2010)
19. Velikovich, L., Blair-Goldensohn, S., Hannan, K., McDonald, R.: The viability of web-derived polarity lexicons. In: *Human Language Technologies: ACL* (2010)
20. Wilson, T., Wiebe, J., Hoffmann, P.: Recognizing contextual polarity in phrase-level sentiment analysis. In: *Proc. Empirical Methods in NLP Conf. (EMNLP)*. Vancouver, Canada (2005)
21. Xu, T., Peng, Q., Cheng, Y.: Identifying the semantic orientation of terms using s-hal for sentiment analysis. *Knowledge-Based Systems* 35, 279–289 (2012)