Polarity preference of verbs: What could verbs reveal about the polarity of their objects?

Klenner, Manfred ; Petrakis, Stefanos

Abstract: The current endeavour focuses on the notion of positive versus negative polarity preference of verbs for their direct objects. This preference has to be distinguished from a verb's own prior polarity - for the same verb, these two properties might even be inverse. Polarity preferences of verbs are extracted on the basis of a large and dependency- parsed corpus by means of statistical measures. We observed verbs with a relatively clear positive or negative polarity preference, as well as cases of verbs where positive and negative polarity preference is balanced (we call these bipolar-preference verbs). Given clear-cut polarity preferences of a verb, nouns, whose polarity is yet unknown, can now be classified. We reached a lower bound of 81% precision in our experiments, whereas the upper bound goes up to 92%.

DOI: https://doi.org/10.1007/978-3-642-31178-9_4

Polarity preference of verbs: What could verbs reveal about the polarity of their objects?

Manfred Klenner and Stefanos Petrakis

Institute for Computational Linguistics, University of Zurich, Switzerland,
{klenner,petrakis}@cl.uzh.ch,
http://www.cl.uzh.ch/

Abstract. The current endeavour focuses on the notion of positive versus negative polarity preference of verbs for their direct objects. This preference has to be distinguished from a verb's own prior polarity - for the same verb, these two properties might even be inverse. Polarity preferences of verbs are extracted on the basis of a large and dependency-parsed corpus by means of statistical measures. We observed verbs with a relatively clear positive or negative polarity preference, as well as cases of verbs where positive and negative polarity preference is balanced (we call these bipolar-preference verbs). Given clear-cut polarity preferences of a verb, nouns, whose polarity is yet unknown, can now be classified. We reached a lower bound of 81% precision in our experiments, whereas the upper bound goes up to 92%.

Keywords: sentiment analysis, polarity inference, data mining

1 Introduction

There are two major classes of polarity lexicons: those that tag polarity at the word level (mostly done semi-automatically), e.g. the Subjectivity Lexicon by Wilson et al. (described in [1]), and those that determine the polarity on the basis of lexical resources, e.g. the SentiWordNet [2]. Accordingly, positive, negative and neutral polarity is attached to word *senses*. Fortunately, quite a number of words do have a clear-cut prior polarity that is invariant over the word's senses. A carefully designed word-level prior-polarity lexicon thus is a valuable resource for sentiment analysis.

The work reported here strives to semi-automatically augment the publicly available word-level prior-polarity lexicon for German by Clematide and Klenner [3]. We base our approach on a large text corpus for German, the DeWaC corpus [4], and try to automatically identify new polar nouns and add them to the lexicon after manual inspection.

The focal point of the proposed method is the verb-object relationship. Verbs play a crucial role in clause-level sentiment determination. Some verbs (e.g. to love) not only have a prior polarity, but they also seem to have a specific polarity 'disposition' (or 'preference') as far as the polarity on their objects is concerned (e.g. I love books). In the case of a human subject, some authors, namely Wiebe

et al. in [5], speak of a positive or negative attitude of the opinion holder, the subject, towards the opinion target, the object. Our goal was to automatically identify verbs with such a polarity preference for their objects based on corpus material. Starting from a prior-polarity lexicon for German (see [3]) and the noun polarities defined there, all verbs that have a polar object (a single noun) have been collected. Although there are clear-cut cases, where the prior polarities of all nouns of a verb have a single orientation, most of the time, both polarities, positive and negative, occur with some frequency. Moreover, there are cases, where both polarities occur with (almost) the same frequency. Those verbs, which we name bipolar-preference verbs, do not impose any restriction on the polarity of their objects. In order to operationalize the choice of whether a verb actually has a clear preference or not, we first use a statistical test. Then, after having obtained a list of polarity preference verbs, we measured how well these verbs could actually predict the polarity of (unseen) nouns occurring as their objects. We did this in a held-out setting, where we used nouns from the prior-polarity lexicon. But instead of letting one verb predict the polarity of these unseen nouns, we rather used a Naive Bayes assumption to let all verbs that co-occur with a noun vote for its polarity. It turned out that our approach was able to reproduce from 81% and up to 92% of the true noun polarities.
The paper is organized as follows. First, we introduce the resources being used, then we discuss our approach to separate bipolar-preference from single-preference verbs. We then show how the classification of prior polarities of nouns can be achieved on the basis of the polar verbs and a Naive Bayes model. We finally discuss related work and conclude.

2 Resources

Our experiments were based on a freely available German polarity lexicon comprising 8'400 entries, 3'400 nouns, 1'600 verbs and 3'800 adjectives (see [3]). The word polarities for nouns and verbs were manually tagged, whereas 1000 of the 3800 adjectives were semi-automatically derived.
Our primary data source was the DeWaC corpus (see [4]), a large collection of web-retrieved documents in the German language. We worked with a subset of the DeWaC corpus consisting of 20 Million sentences. We fed this dateset to a dependency parser developed by Sennrich et al. [6], in order to obtain dependency information relating verbs with their objects. We subsequently labeled this automatically parsed dataset with polarity information originating from the lexicon. From this preprocessed collection of sentences we proceeded to extract verbs and their polar objects by applying the following restrictions:

- based on the available dependency information, select all verbs and their objects
- allowed objects are noun phrases consisting of exactly one noun and an optional preceding determiner (an article in most cases)
- the core noun of such a noun phrase must be labeled as positive or negative

Based on these restrictions we produced 22'155 triples, where each triple comprises a verb, a noun being the verb's object and a count of the observed co-occurences of the given verb-object pair. The total number of noun types is 1'781, they co-occur with 1'776 verbs. Fig. 1 gives a couple of examples[1].

verb	noun	translation(en)	freq
abbauen	Stress	stress	62
abbauen	Übergewicht	overweight	18
abbauen	Unsicherheit	unease	5
auslösen	Ekel	disgust	1
auslösen	Empörung	indignation	45
auslösen	Enthusiasmus	enthusiasm	1

Fig. 1. Input Data

3 Verb's Preference Determination

As mentioned above, we extracted triples consisting of verbs, their objects and the triple's observed frequency. From this set of triples, a list of verbs can be further generated with new frequency counts: the number of positive nouns and the number of negative nouns they have as their direct objects. We distinguish three such frequency patterns:

1. verbs, where one polarity is clearly prevailing
2. verbs, where the frequency of a single polarity orientation is higher than the frequency of the opposite polarity, but where the frequency counts do not indicate a clear-cut preference
3. verbs, where - on average - both polarities occur with the same frequency

In all these cases, neutral NPs are to be expected as well, but we can not calculate their frequency as the polarity lexicon we use contains no entries with neutral polarity. With verbs of type (1), these neutral NPs are likely to adopt the polarity preference (cf. Fig. 2, example 1). It is not entirely clear, how verbs of type (2) and type (3) do influence and can define these neutral NPs and if they do it at all. (cf. Fig. 2, example 2 and 3). Moreover, verbs with a clear preference sometimes occur with a NP that has a prior polarity that conflicts with the verbs polarity preference. For instance, the sentence *"He approves the war"*, where *"approve"* has a preference for positive polarity (one normally approves of something positive, which indicates that a neutral object or an object of unknown polarity should be regarded as something positive) and *"war"* is clearly negative. Such conflicts are to be expected given a large text corpus. There are

[1] *"abbauen"*: *"to reduce"*, *"auslösen"*: *"to trigger"*

1. She loves books/skating.
2. He expresses interest/contempt.
3. She feels happy/angry.

Fig. 2. Polar and Bipolar Verbs

cases where a real conflict was intended by the writer/speaker, but noise can also occur, e.g. stemming from preprocessing (e.g. parsing).
The presence of neutral nouns cannot tell anything about a verb's preference. So we do not have to care about our inability to classify them as neutral. Also, verbs with a balanced frequency pattern (same number of positive and negative nouns) seem to be less interesting, since one cannot predict anything for a noun occurring as an object of such verbs (we call these verbs bipolar). This is not to say that they are of completely no interest, since some of them might be shifters - turning the polarity from positive to negative and vice versa (e.g. *"lose control"* = negative, *"lose fear"*= positive).

3.1 Hypothesis testing

As previously described, we are interested in those verbs that could help us predict the polarity of unseen nouns. It poses, however, the question of how to operationalize distinguishing between verbs that exhibit a clear-cut polarity preference in regards to their direct objects and those verbs that tend to be 'indifferent' towards the polarity of their direct objects.

$$\chi^2_{(k-1)} = \sum_{j=1}^{k} \frac{(f_j - e_j)^2}{e_j}$$

where:
e_j = expectation
f_j = frequency
k = number of independent trials

Fig. 3. Chi Square

The chi-square test (see. Fig. 3) used as a hypothesis test allows for evaluation of a postulated prior probability of events in the light of frequency counts from a sample. Based on the expected frequencies, the difference to the real frequencies is used to derive the so-called p-value, which is a value taken from the chi-square distribution. If the likelihood of the p-value and any value greater than it is below a significance threshold, the so-called null hypothesis is rejected and the alternative hypothesis is adopted. The null hypothesis in our case could either be that a) a verb has a clear polarity preference or b) that a verb is a bipolar-preference verb. In either case, we first have to define what we mean by these distinctions. We define it like this:

1. A verb has a clear-cut preference, it is polar, if the probability of one orientation is 1 and the probability of the other orientation is 0.
2. A verb is bipolar, if the probability for both polarities is 0.5

Definition (1) does not lead to a valid null hypothesis, since on one hand the expected frequencies of the null class, namely 0, would lead to a division by zero (one of the denominators) and on the other hand they would not be in line with the requirement that $e_j > 5$. So we had to continue with definition (2). We thus set up the test hypothesis shown in Fig. 4. H_0 states that a verb is bipolar, H_1

$$h_0 : p(verb = POS) = p(verb = NEG)$$

$$h_1 : p(verb = POS) \neq p(verb = NEG)$$

Fig. 4. Hypothesis testing

that is not bipolar. But does being non-bipolar actually implies that the verb is polar, i.e. has a clear polarity preference? Clearly, this test cannot tell us much about the strength of such a preference. Thus, we use conditional probabilities to quantify the strength parameter. Although H_0 seems to be rather strict, defining bipolarity by an equal probability, the chi-square distribution allows for some deviation from such a strict requirement.

For every verb v, we set the null hypothesis to be that v is bipolar. We then took the verb's positive and negative noun distribution (*ndist*) and determined a p-value (one degree of freedom). If the p-value (*pval*) was less or equal to 0.05 we rejected H_0 and adopted H_1, meaning that verb was categorized as polar, i.e.:

$$pval(v, ndist) = chi^2_{(1)} =< 0.05 \rightarrow polar(v)$$

Since we do not have a gold standard of preclassified polar and bipolar verbs, we decided to test our approach by finding out how well as being polar classified verbs predict the prior polarity of unseen nouns. This is described in the next section. We here give some examples of polar verbs derived by this criterion (see Fig. 5). From the total of 1776 verbs, 420 were classified as polar, while 1356 were classified as bipolar. Each row in Fig. 5 shows a polar verb, the conditional probability that a noun occurring with the verb is negative and positive, respectively and the p-value (on the 0.05 significance level). There are clear cases, e.g. *"tilgen"* (to extinguish) with a polarity of 1 for negative polarity and 0 for positive polarity. But also weaker candidates are among the entries, e.g. *"kompensieren"* (to compensate), where the probability of a negative noun

verb	translation(en)	P(n=POS|v)	P(n=NEG|v)	p-value
kompensieren	to compensate	0.133	0.866	0.004509
hinnehmen	here: to bear	0.125	0.875	0.0027
zusagen	to promise	0.875	0.125	0.03389
eindämmen	to stem	0.111	0.888	0.00096
gewährleisten	to guarantee	0.9444	0.0555	0.00016
beheben	to mend	0.076	0.923	0.002282
erbitten	to ask for	1	0	0.00031
hassen	to hate	0.1	0.9	0.00034
befriedigen	to satisfy	1	0	0.0081
vorgaukeln	to simulate	1	0	0.00091
tilgen	to extinguish	0	1	0.01431
zusichern	to assure	1	0	0.000532

Fig. 5. Polarity preference of verbs

is 0.133 and 0.866 for a positive noun. Before we turn to the question how to evaluate this list, we give the formulas to determine the conditional probabilities from Fig. 5. We use the relative frequency of a verb and the nouns that are its objects. Fig. 6 gives the estimation for the conditional probability of a positive orientation POS for a noun given the verb. The formula for a negative orientation is defined accordingly.

$$P(POS|verb) = \frac{\#pos\_nouns}{\#pos\_nouns + \#neg\_nouns}$$

Fig. 6. Conditional Probability Estimation

4 Classification of Prior Noun Polarities

Having classified verbs according to their positive or negative polarity preference is crucial for the polarity classification of unseen nouns (i.e. nouns that are not in the polarity lexicon). For instance, if the noun *"stability"* is unknown, but occurs in a phrase like *"it guarantees the stability"*, according to the polarity preference of *"to guarantee"* it should be classified as POS, since $P(POS|guarantee) > P(NEG|guarantee) = 0.9444 > 0.0555$. However, with *"to promise stability"* this is not so clear, since $P(NEG|promise)$ is 0.125. Moreover, *"stability"* might occur with a verb that has a negative polarity preference, namely *"to bear"* or *"to compensate"*. That is, the polarity of an unseen noun should be the result of a voting among all verbs that have the noun as their direct object. We can operationalize this under a Naive Bayes Assumption, see Fig. 7. The probability

$$P(n_j = POS|v_1, ..v_n) = \prod_i P(POS|v_i)$$

Fig. 7. Naive Bayes Assumption

of a noun's polarity given its co-occurrence as an object of a number of verbs, v_i, is approximated by making an independence assumption and computing the product of the probabilities for the probability of a polarity orientation given that the noun occurs with a single verb. Fig. 8 gives the full definition of the polarity determination for an unseen noun. Strictly spoken, we can only infer

$$pol(n_j) = \begin{cases} POS \;\; if \;\; P(n_j = POS|v_1, ..v_n) > P(n_j = NEG|v_1, ..v_n) \\ NEG \;\; if \;\; P(n_j = POS|v_1, ..v_n) < P(n_j = NEG|v_1, ..v_n) \\ undef \; else \end{cases}$$

Fig. 8. Noun Polarity

the absence of the inverse polarity of the polarity that wins. If *"stability"* was classified as being positive, then, if we want to be precise, we could only say that it is not negative. The reason is, it could be positive *or* neutral. As we mentioned earlier, nouns co-occurring with a polarity preference verb inherit the preference if they are neutral. But this does not mean they necessarily have this as a prior polarity. For instance, if one tells us that he prefers coke (over beer), we would not classify coke as positive in general, but as being positive only for the person (group or institution) that approved to such a statement. However, this is true only as long as we consider a single statement. But if a noun co-occurs many times with different polarity preference words and possibly different opinion holders, we start becoming more confident in deducing a noun polarity that is based on common sense, rather than on personal preference. In the way our approach works we aim to derive a prior polarity of nouns.

5 Evaluation

How can we evaluate our approach? One way is to have a closer look at the upper bound for precision. A simple way to measure this was to use the prior-polarity lexicon and the preclassified nouns therein and let the system try to recover their real orientation. This is described in the next section. A second way for evaluation was to test the learned verb preferences on nouns with unknown prior polarity, and check the predictions of our method against human evaluators. Both of these ways are described in the following two subsections.

5.1 Experiments Based on a Polarity Lexicon

The preference polarity of a verb is determined in our approach based on the prior polarities of the nouns it has as objects in a large corpus. We have sampled over 22'000 triples, where each triple specifies the frequency a single noun co-occurs with the given verb as its direct object. For instance, the verb *"abbauen"* (here *"to reduce"*) has 65 times the negative noun *"Stress"* (stress) as a direct object. But *"Stress"* might occur with other verbs as well, actually is occurs with 91 different verbs with a total frequency of 312. In order to measure how well our approach reproduces the prior polarities of the nouns co-occuring with the verbs, we adopted a held-out scenario. We removed a single noun from our triples (i.e. every triple that counts the frequency of the noun given a single verb) and subsequently ran our algorithm (i.e. trained the system) and classified the heldout noun with the learned statistical model. Fig. 9 shows the core algorithm. A verb is polar, if it passes the threshold set by chi-square. There are however

```
forall heldout ∈ nouns
    ho_verbs = all verbs that have heldout as an object
    forall hov ∈ ho_verbs
        npos(hov) = number of nouns types with positive prior polarity (not counting heldout)
        nneg(hov) = number of nouns types with negative prior polarity (not counting heldout)
        pval(hov) = chi(1) based on npos(hov), nneg(hov) and p=q=0.5
        if pval(hov) < 0.05 do (i.e. the verb is polar)
            P(POS|hov) = npos(hov) /(npos(hov) +nneg(hov))
            P(NEG|hov) = nneg(hov) /(npos(hov) +nneg(hov))
            pred_pol(heldout)= POS if P(POS|hov) > P(NEG|hov)
            pred_pol(heldout)= NEG if P(POS|hov) < P(NEG|hov)
            pred_pol(heldout)= no_prediction if P(POS|hov) == P(NEG|hov)
            hit+= if pred_pol(heldout) == true_polarity_of(heldout)
```

Fig. 9. Core Algorithm

other parameters that influence the precision of the algorithm:

1. minfvn: the number of different verbs the heldout noun occurs with
2. minfn: the frequency of the heldout noun
3. minfv: the frequency of the verb, the heldout noun occurs with

It turned out that the precision ranges from 81.97% up to 92.20 % by varying these parameters. Fig. 10 gives the details of various runs with different settings for minfvn, minfn and minfv. By introducing these thresholds, nouns that do not occur in a context that fulfills the parameter settings are left unclassified, i.e. no prediction is made and, thus, the total number of classifications drops the more restrictive the setting is (see column # all). For instance, the last line

minfvn	minfv	minfn	prec	#hit	#all
1	1	1	81.97	1210	1476
1	5	2	83.32	1228	1475
2	5	2	85.52	768	898
3	5	2	85.54	497	581
4	5	2	86.68	358	413
4	10	5	87.62	347	396
4	10	10	87.62	347	396
5	10	5	88.53	224	253
5	20	5	91.06	214	235
5	**25**	**5**	**93.34**	**217**	**235**
5	30	5	89.16	210	235
6	10	5	92.43	159	173
6	25	5	91.66	154	168

Fig. 10. Varying Thresholds

is the most restrictive setting, here only 168 out of the total number of 1781 nouns are classified. 312 nouns are not classified in any setting, they occur only with bipolar verbs. Please note that we can control this situation, that is, we can actually decide to run for high precision by simply setting the parameters accordingly.
The best result in terms of precision is achieved with minfvn=5, minfv=25 and minfn=5. Here, 235 nouns are classified, 217 were classified correctly (93.34 %), This is the upper bound precision our approach reaches given the corpus at hand. We deliberately have not measured recall, we wanted to fix the parameter setting that gives us the best precision if applied to unseen examples (see the next section) For illustration purposes, we give some examples of misclassified nouns in Fig. 11. Please note that among these conflicts are also cases where

verb v	English	predicted	true
Kompliment	compliment	NEG	POS
Schauder	shiver	POS	NEG
Entwarnung	all-clear	NEG	POS
Begehrlichkeit	greediness	NEG	POS
Spannung	suspense	NEG	POS
Unterwerfung	repression	POS	NEG
Offenheit	frankness	NEG	POS

Fig. 11. Wrong Predictions

the polarity value of a word according to the prior-polarity lexicon seems to be wrong. For example, *"Begehrlichkeit"* (roughly, greediness) is positive according to the lexicon, but negative according to the algorithm's prediction. We actually

would argue that the prediction is right, and the lexicon entry wrong. We have not systematically explored the possibility to automatically detect such false entries.
It is an interesting question where the other conflicts arise from. Noise might be a reason, e.g. the wrong parse trees (i.e. wrong objects). Another and quite to be expected factor is word sense ambiguity. Take *"Spannung"*: it is positive if suspense is meant, and negative if it is used in the sense of strain (Anspannung); it is even neutral in the (electrical) sense of voltage. The mix-up of these different polarities for the word might explain the misclassification. If might even be the case that in our corpus actually *"Spannung"* is mostly used negative in the sense of strain. So another function of our approach could be to detect words that are highly ambiguous and should better not be part of a prior-polarity lexicon. Conflicting polarities could also indicate, if the corpus is domain-specific, that the polarity of a word in that domain is different from the *"normal"* one. Further work is needed to explore these possibilities. Finally, there is also the possibility of real conflicts in the data (e.g. someone approves/likes something negative), especially given low frequency cases, where no majority vote can compensate for extreme attitudes towards an object.

5.2 The Case of Unseen Nouns

In order to evaluate our approach not only on heldout nouns that have a known prior positive and negative polarity, but to also see how well it works in the presence of neutral nouns, we extracted 200 unseen nouns that have no entry in the polarity lexicon. We took nouns that occur with a high frequency with the polar verbs our approach have identified. We then ran the system with the parameter setting that proved best (see Fig. 10) and produced a list of 200 polarity classifications of unseen nouns. Fig. 12 shows some examples. Two

noun	pred	true
Aufenthaltserlaubnis	POS	POS
Urlaub	POS	POS
Bereitschaft	POS	POS
Hindernis	NEG	NEG
Ausbildungsplatz	POS	POS
Geburtstag	NEG	POS
Wertung	NEG	NEUT
Tisch	POS	NEUT
Kommentar	NEG	NEUT
Rad	NEG	NEUT

Fig. 12. Nouns' Prior-Polarity Predictions

annotators have then evaluated this list. The labels were (1) agree, (2) disagree,

(3) neutral. That is, if the noun actually was neutral, no disagreement was (and should not be) measured, since the nouns appeared as objects of verbs with a polarity preference (the annotators did not see the verb or the whole sentence). So in the context of their verbs and sentences, the labeling as polar must not necessarily be wrong (the neutral noun inherits the preference by definition). For instance, *"Rad"* (bicycle) was, according to Fig. 12, labelled as negative - so it occurred with negative verbs mostly (e.g. *"I hate bicycles"*). In that context, *"bicycle"* is negative. We simply have not checked this in the present investigation. So neutrals cannot be evaluated properly by just looking at our result list. The annotators agreed on 24 cases (86%) and disagreed on 4 cases, so we found 28 polar nouns, the rest, 144, were neutral. What was somewhat disappointing is the fact, that only few nouns with a prior polarity turned up - most nouns were neutral.

6 Related Work

It is not the first time that a corpus-driven approach has been used for predicting the semantic orientation of a specific word class. Such approaches have in general been widely applied in the field of Sentiment Analysis over the last decade. Exploiting syntactic relations between specific word-classes and patterns of co-occurrences has proven quite effective. The word classes that mostly receive the focus of such approaches are the two usual suspects: adjectives as treated by Clematide and Klenner in [3] and by Hatzivassiloglou and McKeown in [7] and nouns as treated by Rillof et al. in [8]. Especially the latter attempt comes quite close to the idea behind our method. It is however a bootstrapping method while we are equipped with a complete polarity lexicon.
The role that verbs play in sentiment analysis, especially their semantic role, is a very important point and has been well defined in the work of Chesley et al. in [9] and also by Neviarouskaya et al. in [10]. Our method is not incompatible with these ideas, but we choose to neglect any semantic information in favor of simplicity. The semantics of verbs are extremely important, especially in compositional, subsentential frameworks. We choose an engineering approach trying to see how far minimal linguistic knowledge can lead in a prediction task.

7 Conclusion and Future Work

We introduced a purely corpus-based method to both, detecting verbs that have a polarity preference towards their objects and - based on that - a method to derive the prior polarity of nouns co-occurring with these verbs. The polarity preference of a verb allows e.g. to deduce the attitude that an opinion holder has towards the opinion target. So our approach fits well into existing lines of research. Moreover, it demonstrates how to learn the polarity preference of verbs instead of manually assigning it. The advantage is that these polarity preferences are empirically licensed. A statistical measure was used to operationalize the distinction between such *polar* verbs and verbs without a clear-cut tendency,

which we called *bipolar* verbs. As our experiments with heldout data showed, the polarity preference of verbs can be reliably learned.
The polarity preference expressed in a single statement based on the basis of a single polar verb does not justify the classification of an opinion (noun) object as bearing a prior, positive or negative polarity. However, if a single noun occurs with different polar verbs produced by different opinion holders in different statements, then we no longer talk about personal preferences but enter the realm of common sense consensus. However, as discussed, noise and ambiguity at the word level interfere with this tendency, a fully automatic identification of nouns with a prior polarity is therefore quite challenging.
We have discussed various applications for our approach, e.g. one could use it to carry out domain adaption (detecting words that have a domain-specific polarity) or to check the entries of a polar lexicon on the basis of a large corpus for false entries. It might be better to remove highly ambiguous words from the lexicon - our method could help to identify such words.

References

1. Wilson T., Wiebe J., Hoffmann P.: Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing, 347–354 (2005)
2. Esuli, A., Sebastiani, F.: Sentiwordnet: A publicly available lexical resource for opinion mining. Proceedings of the fifth Conference on Language Resources and Evaluation, 417–422 (2006)
3. Clematide S., Klenner M.: Evaluation and extension of a polarity lexicon for German. Proceedings of the First Workshop on Computational Approaches to Subjectivity and Sentiment Analysis, 7–13 (2010)
4. Baroni M., Bernardini S., Ferraresi A., Zanchetta E.: The WaCky Wide Web: A collection of very large linguistically processed Web-crawled corpora. Language Resources and Evaluation, 209–226 (2009)
5. Wiebe J., Wilson T., Cardie C.: Annotating expressions of opinions and emotions in language. Language Resources and Evaluation 39(2), 165–210 (2005)
6. Sennrich R., Schneider G., Volk M., Warin M.: A new hybrid dependency parser for German. Proceedings of the German Society for Computational Linguistics and Language Technology, 115–124 (2009)
7. Hatzivassiloglou V., McKeown K. R.: Predicting the semantic orientation of adjectives. Proceedings of the eighth conference on European chapter of the Association for Computational Linguistics, 174–181 (1997)
8. Riloff E., Wiebe J., Wilson T.: Learning subjective nouns using extraction pattern bootstrapping. Proceedings of the seventh conference on Natural language learning at HLT-NAACL 2003, (4): 25–32 (2003)
9. Chesley P., Vincent B., Xu L., Srihari R.K.: Using Verbs and Adjectives to Automatically Classify Blog Sentiment. Training, 580 (263): 233 (2006)
10. Neviarouskaya A., Prendinger H., Ishizuka M.: Semantically distinct verb classes involved in sentiment analysis. Proceedings of the International Conference on Applied Computing,pp. 27–34 (2009)