

An Empirical Study on Machine Learning-Based Sentiment Classification using Polarity Clues

Ulli Waltinger

Artificial Intelligence Group,
Cognitive Interaction Technology Excellence Center (CITEC),
Bielefeld University
`ulli_marc.waltinger@uni-bielefeld.de`

Abstract. In recent years a variety of approaches in classifying the sentiment polarity of texts have been proposed. While in the majority of approaches the determination of subjectivity or polarity-related term features is at the center, the number of publicly available dictionaries is rather limited. In this paper, we investigate the performance of combining lexical resources with machine learning-based classifier for the task of sentiment classification. We systematically analyze four different English and three different German polarity dictionaries as a resources for a sentiment-based feature selection. The evaluation results show that smaller but more controlled dictionaries used for feature selection perform within a SVM-based classification setup equally good compared to the biggest available resources.

Keywords: Machine Learning, Support Vector Machine, Sentiment Analysis, Polarity Classification, Polarity Resources

1 Introduction

With the enormous growth of digital content arising in the web, document classification and categorization receives more and more interest in the information retrieval community. This relates to content-based models [1] as well as to structure-orientated approaches [2]. While a majority of approaches focusses on a thematical or topical differentiation of textual data, the task of sentiment analysis [3] refers to the (non-topical) opinion mining. This area focuses on the detection and extraction of opinions, feelings and emotions in text with respect to a certain subject. A subtask of this area, which has been extensively studied, is the sentiment categorization on the basis of certain polarities. That is, being able to distinguish between positive, neutral or negative expressions or statements of extracted textual [4–8] or spoken elements [9]. Moreover, finer-grained methods additionally explore the level or intensity of polarity inducing a rating inference (e.g. a rating scale between one and five stars) model. In the majority of approaches on sentiment polarity identification, the determination of subjectivity or polarity-related term features is in the center in order to draw conclusions

about the actual polarity-related orientation of the entire text. Since positive as well as negative expressions can occur within the same document, this task is challenging. Considering the following example of an Amazon product review:

Product-Review¹: Wonderful when it works... I owned this TV for a month. At first I thought it was terrific. Beautiful clear picture and good sound for such a small TV. Like others, however, I found that it did not always retain the programmed stations and then had to be re-programmed every time you turned it off. I called the manufacturer and they admitted this is a problem with the TV.

Although most of the polarity-related text features contribute to a positive review (e.g. wonderful, terrific, beautiful...), this user-contribution is classified as a negative review. This example clearly shows that classical text categorization approaches (e.g. bag-of-words) need to be extended or seized to the domain of sentiment analysis. Though, we consider polarity identification as a binary classification task, the determination of semantically oriented linguistic features on different structural levels (words, sentences, documents,...) is at the core of attention. With respect to the task of term feature interpretation, most of the proposed unsupervised or (semi-)supervised sentiment-related approaches make use of annotated and constructed lists of subjectivity terms.

While there are various resources and data sets proposed in the research community, only a small number are freely available to the public – most of them for the English language – most recently also with respect to the German language [10, 11]. In terms of coverage rate, the number of comprised subjectivity terms of these dictionaries varies significantly - ranging between 8,000 and 140,000 features. The questions that arise therefore are: How does the significant coverage variations of the English sentiment resources correlate to the task of polarity identification? Are there notable differences in the accuracy performance, if those resources are used within the same experimental setup? How does sentiment term selection combined with machine learning methods affect the performance? And finally, are we able to draw conclusions from the results of the experiments in building a German sentiment analysis resource?

In this paper, we investigate the effect of sentiment-based feature selection combined with machine learning algorithms in a comparative experiment, comprising the four most widely used subjectivity dictionaries. We empirically show that a sentiment-sensitive feature selection contributes to the task of polarity identification. Further, we propose based on the findings a subjectivity dictionary for the German language, that will be freely available to the public.

2 Related Work

In this section, we present related work on sentiment analysis. A focus is set on comparative studies and different algorithms applied to the task of polarity

¹ <http://www.amazon.com/>

identification. Tan and Zhang (2008) presented an empirical study of sentiment categorization on the basis of different feature selection (e.g. document frequency, chi square, subjectivity terms) and different learning methods (e.g. k-nearest neighbor, Naive Bayes, SVM) on a Chinese data set. The results indicated that the combination of sentimental feature selection and machine learning-based SVM performs best compared to other tested sentiment classifiers.

Chaovalit and Zhou (2005) published a comparative study on supervised and unsupervised classification methods in a polarity identification scenario of movie reviews. Their results confirmed also that machine learning on the basis of SVM are more accurate than any other unsupervised classification approaches. Hence, a significant amount of training and building associated models is needed.

Prabowo and Thelwall (2009) proposed a combined approach for sentiment analysis using rule-based, supervised and machine learning methods. An overview of current sentiment approaches is given, compared by their model, data source, evaluation methods and results. However, since most of the current attempts based their experiments on different setups, using mostly self-prepared corpora or subjectivity resources, a uniform comparison of the proposed algorithms is barely possible. The results of the combined approach show that no single classifier outperforms the other, and the hybrid classifier *can* result in a better effectiveness.

With respect to different methods applied to the sentiment polarity analysis, we can identify two different branches. On the one hand - rule-based approaches, as for instance counting positive and negative terms [15] on the basis of semantic lexicon, or combining it with so called discourse-based contextual valence shifters [16]. On the other hand - machine-learning approaches [17] on different document levels, such as the entire documents [4], phrases [7, 18, 19], sentences [20] or on the level of words [21], using extracted and enhanced linguistic features from internal (e.g. PoS- or text phrase information) and/or external resources (e.g. syntactic and semantic relationships extracted from lexical resources such as WordNet [22]) [23, 13]. Most notably, sentence-based models have been quite intensively studied in the past, combining machine learning and unsupervised approaches using inter-sentence information [24, 25], sentence-based linguistic feature enhancement [26] or most famous by following a sentence-based minimum cut strategy [20, 27].

In general, sentence-based polarity identification contributes to a higher accuracy performance, but induces also a higher computational complexity. Nevertheless, depending on the used methods the reported increase of accuracy of document and sentence classifier range between 2 – 10% [20, 26], mostly compared to the baseline (e.g. Naive Bayes) implementations. However, in the majority of cases, only slightly better results could be achieved [25, 26]. At the focus of almost all approaches, a set of subjectivity terms is needed, either to train a classifier or to extract polarity-related terms following a bootstrapping strategy [24].

3 Background

3.1 Modeling Opinion Orientation

Following Liu (2010)[28, pp. 5] we formally define an opinion oriented model as follow: A polarity-related document d contains a set of opinion objects $\{o_1, o_2, \dots, o_q\}$ from a set of opinion holders $\{h_1, h_2, \dots, h_p\}$. Each opinion object o_j is represented by a finite set of sentiment features, $F = \{f_1, f_2, \dots, f_n\}$. Each feature $f_i \in F$ is represented in d by a set of term or phrases $W = \{w_{i1}, w_{i2}, \dots, w_{im}\}$, which correspond to synonyms or associations of f_i and are indicated by a set of feature indicators $I_i = \{i_{i1}, i_{i2}, \dots, i_{ip}\}$ of the feature. The direct opinion of o_j is expressed through the polarity of the opinion (e.g. positive, negative, neutral) defined as oo_j with respect to the comprised set of features f_j of o_j , the opinion holder h_i and the time or position within the text t_j , an opinion is expressed. The feature indicator i_j reflects thereby the strength of the opinion (e.g. rating scale). Following this definition, contrary opinions within a text document (e.g. phrase or sentence-based) correlate to a (dis-) similarity S of two opinion objects $S(o_j, o_k)$, while a concordance of a polarity is indicated by a high similarity value. At the center of the opinion-oriented model, a mapping from the input document to the corresponding sentiment features with associated indicators ($W \mapsto F$) needs to be established. Meaning, an external resource is needed that embodies not only a set of term or phrase features, but also incorporates the polarity orientation at least as a boolean (positive, negative), preferably on a rating scale (positive, negative, neutral). We refer to these resources as *subjectivity dictionaries*. As we use machine learning classifiers, the similarity function $S(o_j, o_k)$ refers to the similarity between the supervised trained SVM-based opinion models (o_j) and the evaluation set of document opinions (o_k).

3.2 Subjectivity Dictionaries

In recent years a variety of approaches in classifying sentiment polarity in texts has been proposed. However, the number of comprised or constructed subjectivity resources are rather limited. In this section, we describe the most widely used subjectivity resources for the English language in more detail.

Adjective Conjunctions

As one of the first, Hatzivassiloglou et al. (1997) proposed a bootstrapping approach on the basis of adjective conjunctions. Thereby, a small set of manually annotated seed words (1,336 adjectives) were used in order to extract a number of 13,426 conjunctions, holding the same semantic orientation i.e. 'and' indicates an agreement of polarity (nice and comfortable) and 'but' indicates disagreement (nice but dirty). Subsequently, a clustering algorithm separated the sum of adjectives into two subsets of different sentiment orientation (positive or negative). This approach follows the notion that a pair of adjectives (e.g. conjunction in a sentence) will most likely have the same orientation (81% of the unmarked member will have the same semantic orientation as the marked member).

WordNet Distance

Maarten et al. (2004) presented an approach measuring the semantic orientation of adjectives on the basis of the linguistic resource *WordNet* [22]. A focus was set on graph-related measures on the syntactic category of adjectives. The geodesic distance is used as a measurement to extract not only synonyms but also antonyms. As a reference dataset, the manually constructed list of the General Inquirer [30] was used, comprising 1,638 polarity-rated terms. Since the evaluation focused on the intersection of both resources (General Inquirer vs. *WordNet*), no additional corpus could be gained.

WordNet-Affect

A related approach in building a sentiment resource, Strapparava and Valitutti (2004)[31] studied the synset-relations of *WordNet* with respect to their semantic orientation. Following a bootstrapping-strategy, manually classified seed words were used for constructing a list of 'reliable' relations (e.g. antonym, similarity, derived-from, also-see) out of the linguistic resource. The final dataset, *WordNet-Affect*, comprises 2,874 synsets and 4,787 words.

Subjectivity Clues

In 2005, Wiebe et al. (2005) presented the most fine-grained polarity resource. Within the Workshop on Multi-Perspective Question Answering (2002) the MPQA corpus was manually compiled. This corpus consists of 10,657 sentences comprising 535 documents. In total, 8,221 term features were not only rated by their polarity (positive, negative, both, neutral) but also by their reliability (e.g. strongly subjective, weakly subjective).

SentiWordNet

Esuli and Sebastiani (2006) introduced a method for the analysis of glosses associated to synsets of the *WordNet* data set. The proposed subjectivity resource *SentiWordNet* thereby assigns for each synset three numerical scores, describing the objective, negative, and positive polarity of interlinked terms. The used method is based on the quantitative analysis of glosses and a vectorial term representation for a semi-supervised synset classification. Overall, *SentiWordNet* comprises 144,308 terms with polarity scores assigned.

SentiSpin

Takamura et al. (2005) proposed an algorithm for extracting the semantic orientation of words using the *Ising Spin Model* [36, pp. 119]. Their approach focused on the construction of a gloss-thesaurus network inducing different semantic relations (e.g. synonyms, antonyms), and enhanced the built dataset with co-occurrence information extracted from a corpus. The construction of the gloss-thesaurus is based on *WordNet*. With respect to the co-occurrence statistics, conjunctive expressions from the Wall Street Journal and Brown corpus were used. The available subjectivity resource offers a number of 88,015 words for the English language with assigned Part-of-Speech information and a sentiment polarity orientation.

Polarity Enhancement

Waltinger (2009) proposed an approach to term-based polarity enhancement using a social network. His approach focuses on the reinforcement of polarity-related term features with respect to colloquial language. Using the entries of the *SpinModel* dataset as seed words, associated phrase and term definitions were extracted from the *urban dictionary* project. The enhanced subjectivity resource comprises 137,088 term features for the English language.

4 Methodology

With respect to the described approaches in the construction of subjectivity dictionaries, we can identify two different branches. The majority of proposals induce the lexical network *WordNet* as a foundation for either extending or extracting polarity-related semantic relations. Therefore, the constructed term set is limited to the number of entries within *WordNet*, comprising up to 144,308 polarity features. Other approaches, focused on the manual creation of a subjectivity thesaurus by inducing expert knowledge (manually annotated). These costly built resources consist of a rather small set of polarity features, inducing a dictionary size of up to 6,663 entries. The questions that arise therefore are: How does the different subjectivity resources perform within the same experimental setup of polarity identification? Does the significant difference (quantity) of used polarity features affect the performance of opinion mining applications? Our methodology focuses on the most widely used and freely available subjectivity dictionaries for the task of sentiment-based feature selection.

4.1 SVM-Classification

The method we have used for the polarity classification is a document-based hard-partition machine learning classifier [4, 13, 12, 14, 37] using Support Vector Machines (SVM) [1]. This supervised classification technique relies on training a set of polarity classifiers, each of them capable of deciding whether the input stream has a positive or negative polarity, $C = \{+1, -1\}$. The SVM predicts a hyperplane, which separates a given set into two divisions with a maximum margin (the largest possible distance) [1]. We make use of the *SVM^{Light}*^{6.01} implementation [38], using *Leave-One-Out* cross-validation, reporting *F1-Measure* as the harmonic mean between *Precision* and *Recall*. The reported *Accuracy* measures are based on a *5-fold cross-validation*. In each case of the SVM-Classifiers, *Linear-* and *RBF-Kernel* were evaluated in a comparative manner.

4.2 Subjectivity-Feature-Selection

Using SVMs for classifying the sentiment orientation, each input text needs to be converted into a vector representation. This vector consists of a set of significant term features representing the associated document. With respect to the opinion-oriented model, this task corresponds to a mapping between subjectivity features

from the particular dictionary, and the textual features of the input document. That is, only those features are selected that occur in the subjectivity lexicon. Since the polarity features can consist of single words as well as multi-word expressions, a sliding window is used, when extracting textual data from the input text. As the feature weighting function, we have used the normalized term frequency ($tf_{i,j}$), defined as

$$tf_{i,j} = \frac{f_{i,j}}{n} \quad (1)$$

$$\sum_{k=1} f_{k,j}$$

where the number of occurrences of feature i in document j is normalized by the total number of features n in j . While various subjectivity resources have been proposed in recent years, only a few of them are freely available. In this paper, we evaluate the four most widely used and available resources (see Table 2):

1. Subjectivity Clues [32]
2. SentiSpin [35]
3. SentiWordNet [34]
4. Polarity Enhancement [37]

4.3 German Subjectivity Resource

As described in Section 3.2, the majority of subjectivity resources are based on the English language. Different to the approach of [39], by translating a German input text into the English language (*SentiWordNet* as a resource), we rather focused on building a new German dictionary by translating polarity features only. Since a goal of this paper is to evaluate the correlation between the size of subjectivity dictionaries and the accuracy performance, we have built three different German polarity resources of different size [40]. The construction of these dictionaries is based on a semi-supervised translation of existing English polarity term-sets. That is, we automatically translated each polarity feature into the German language – using an English-to-German translation software², and manually reviewed the translation quality. We have used the following English datasets: First, the *Subjectivity Clues* [32, 7, 33] dictionary, comprising 9,827 term features, further called *German Subjectivity Clues*. Second, we translated the dataset of *SentiSpin* [35], comprising 105,561 polarity features. We will refer to this resource as the *German SentiSpin* dictionary. While there are in many cases more than one possible translations available, we decided to take a maximum number of three translations for dictionary construction into account. Therefore, the size of the built German resources differ to their English pendant. With respect to polarity feature weights, each aggregated German feature has inherited the sentiment orientation score (e.g. positive, negative, neutral) of the initial seed word from the English resource (e.g. English: "brave" — "positive" \mapsto German: "mutig" — "positive"). This approach clearly leads to a problem of

² We have used the online service of dict.leo.org for the translation of term features.

Id:	Feature	PoS	A(+)	A(-)	A(o)	B(+)	B(-)	B(o)
5653	Begründung	NN	0	0	1	0	0.5	0.5
7573	Katastrophe	NN	0	1	0	0	0.68	0.32
7074	ideal	ADJD	1	0	0	0.76	0.13	0.11

Table 1. Overview of the GermanPolarityClues data schema by (A) automatic- and (B) corpus-based polarity orientation rating after [40].

term ambiguity. We therefore decided to compile in a third step the *GermanPolarityClues* dictionary, by *manually* assessing each individual term feature of the *German Subjectivity Clues* dataset by their sentiment orientation (See Table 1). In addition, we added to this resource a number of 290 German negation-phrases (e.g. "nicht schlecht" = "not bad") and the most frequent positive and negative synonyms of existing term features, which previously had not been in there³ - inducing a total size of 10, 141 polarity features. All resources are freely available for research purposes⁴.

5 Experiments

5.1 Corpora

We have used two different datasets for the experiments. For the English language we conducted the polarity identification classification using the movie review corpus initially compiled by [4]. This corpus consists of two polarity categories (positive and negative), each category comprises 1000 articles with an average of 707.64 textual features. With respect to the German language, we manually created a reference corpus by extracting review data from the Amazon.com website. Reviews at Amazon.com correspond to human-rated product reviews with an attached rating scale from 1 (worst) to 5 (best) stars. For the experiment, we have used 1000 reviews for each of the 5 ratings, each comprising 5 different categories. All category and star label information but also the name of the reviewers were removed from the documents. All textual data (term features in the document) were passed through a pre-processing component, that is lemmatized and tagged by a PoS-Tagger. The average number of term features of the comprised reviews is 109.75. With respect to the experiments on the German corpus, we evaluated different "Star" combinations as positive and negative categories (e.g classifying Star1 against Star5, but also Star1 and Star2 against Star 4 and Star 5).

5.2 Results

With respect to the English polarity experiment (see Table 3 and 4), we have used not only the published accuracy results of [4], using the Naive Bayes (NB),

³ Note, these features were extracted from the dataset of the de.wiktionary.org project.

⁴ The constructed resources can be accessed at: <http://hudesktop.hucompute.org/>

Resource:	Subject. Clues	Senti Spin	Senti WordNet	Polarity Enhance	German SentiSpin	German Subject.	German Polarity Clues
No. of Features:	6,663	88,015	144,308	137,088	105,561	9,827	10,141
Positive-AMean:	76.83	236.94	241.36	239.25	53.63	27.70	26.66
Positive-StdDevi:	30.81	84.29	85.61	84.98	6.90	4.59	5.01
Negative-AMean:	69.72	218.46	223.11	221.25	50.18	25.68	24.14
Negative-StdDevi:	26.22	74.08	75.37	74.68	10.40	5.88	5.41
Text-AMean:	707.64	707.64	707.64	707.64	109.75	109.75	109.75
Text-StdDevi:	296.94	296.94	296.94	296.94	24.52	24.52	24.52

Table 2. The standard deviation (StdDevi) and arithmetic mean (AMean) of subjectivity features by resource, text corpus (Text) and polarity category (Positive, Negative).

Sentiment-Method	Accuracy
Naive Bayes - unigrams [4]	78.7
Maximum Entropy - top 2633 unigrams [4]	81.0
SVM - unigrams+bigrams [4]	82.7
SVM -unigrams [4]	82.9
Polarity Enhancement - PDC (without feature enhancement) [37]	81.9
Polarity Enhancement - PDC (with feature enhancement) [37]	83.1
Subjectivity-Clues SVM Linear-Kernel	84.1
Subjectivity-Clues SVM RBF-Kernel	83.5
SentiWordNet SVM Linear-Kernel	83.9
SentiWordNet SVM RBF-Kernel	82.3
SentiSpin SVM Linear-Kernel	83.8
SentiSpin SVM RBF-Kernel	82.5

Table 3. Accuracy results comparing four subjectivity resources and four baseline approaches.

the Maximum Entropy (ME) and the N-Gram-based SVM implementation, but also the results of [37], a feature-enhanced SVM implementation as corresponding baselines. As Table 3 shows, the smallest resource, Subjectivity Clues, performs best with $acc = 84.1$. However, SentiWordNet ($acc = 83.9$), SentiSpin ($acc = 83.8$) but also the Polarity Enhancement ($acc = 83.1$) dataset used for feature selection, perform almost within the same accuracy. It can be stated that all subjectivity feature selection resources clearly outperform not only the well known NB and ME classifier but also the N-Gram-based SVM implementation. Not surprisingly, with respect to the feature coverage of the used subjectivity resources (see Table 2), we can argue that the size of the dictionary clearly correlates to the coverage (arithmetic mean of polarity-features selected varies between 76.83 – 241.36). Interestingly, the biggest dictionary with the highest coverage property does not outperform the resource with the lowest number of polarity-features. In contrast, we can state that operating in the present settings, on 6,663 term features (in contrast to 144,308 of *SentiWordNet*), seem

Resource	Model	F1-Positive	F1-Negative	F1-Average
Subjectivity Clues	SVM-Linear	.832	.823	.828
	SVM-RBF	.828	.823	.826
SentiWordNet	SVM-Linear	.832	.828	.830
	SVM-RBF	.816	.812	.814
SentiSpin	SVM-Linear	.831	.827	.829
	SVM-RBF	.815	.811	.813
Polarity Enhancement	PDC	.828	.827	.828
	SVM-Linear	.841	.837	.839

Table 4. F1-Measure evaluation results of an English subjectivity feature selection using SVM.

Resource	Model	F1-Positive	F1-Negative	F1-Average
German SentiSpin	SVM-Linear	.827	.828	.828
Star1+2 vs. Star4+5	SVM-RBF	.830	.830	.830
German SentiSpin	SVM-Linear	.857	.861	.859
Star1 vs. Star5	SVM-RBF	.855	.858	.857
German Subjectivity	SVM-Linear	.810	.813	.811
Star1+2 vs. Star4+5	SVM-RBF	.804	.803	.803
German Subjectivity	SVM-Linear	.841	.842	.841
Star1 vs. Star5	SVM-RBF	.834	.834	.834
GermanPolarityClues	SVM-Linear	.875	.730	.803
Star1+2 vs. Star4+5	SVM-RBF	.866	.661	.758
GermanPolarityClues	SVM-Linear	.875	.876	.876
Star1 vs. Star5	SVM-RBF	.855	.850	.853

Table 5. F1-Measure evaluation results of a German subjectivity feature selection using SVM.

to be a sufficient number for the task of document-based polarity identification. This claim is also supported by the evaluation F1-Measure results as shown in Table 4. All subjectivity resources nearly perform equally well (F1-Measure results range between 82.9 – 83.9). In this *Leave – One – Out* estimation, the polarity-enhanced implementation performs with a touch better than the other resources. Table 5 shows the results of the new build German subjectivity resources, used for the document-based polarity identification. With respect to the correlation of subjectivity dictionary size and classification performance, similar results can be achieved. Using the *German SentiSpin* version, comprising 105,561 polarity features, lets us gain a promising F1-Measure of 85.9. The *German Subjectivity Clues* dictionary, comprising 9,827 polarity features, performs with an F1-Measure of 84.1 almost at the same level. However, the *German-PolarityClues* dictionary, comprising 10,141 polarity features, outperforms with

an F1-Measure of 87.6 all other German resources. In addition, with respect to the number of polarity features actually used within the Amazon-based SVM-classification experiments (see Table 6), we can identify that a number of 2,700 features only within the *GermanPolarityClues* dictionary exhibits the best performance. It seems that this newly created sentiment resource, which induces a rather small feature size (10-times smaller than the *German SentiSpin*), is due to its manual controlled vocabulary and its introduced negation- and synonym-pattern, of high-quality for the task of polarity identification. In general, in terms of Kernel-Methods, we can argue that RBF-Kernel are inferior to the Linear-Kernel SVM implementation, though only to a minor extend. With reference to the coverage of subjectivity dictionaries for a polarity-based feature selection - *size does matter*. However, the classification accuracy results indicate - for both languages - that a smaller but controlled dictionary contributes to the accuracy performance (almost equally to big-sized data) of opinion mining systems.

German SentiSpin:	10,802
German Subjectivity:	2,657
German Polarity Clues:	2,700

Table 6. Number of polarity features used for the SVM-Classification by comprised resources.

6 Conclusion

This paper proposed an empirical study to machine learning-based sentiment analysis. We systematically analyzed the four most widely used English-based subjectivity resources for the task of sentiment polarity identification. In addition, we proposed and evaluated three new polarity resources for the German language. The evaluation results showed that the size of subjectivity dictionaries does not correlate with classification accuracy. Smaller but more controlled dictionaries used for a sentiment feature selection perform within a SVM-based classification setup equally good compared to the biggest available resources. We can conclude, that combining a polarity-based feature selection with machine learning, SVMs using Linear-Kernel exhibit the best performance ($acc = 84.1$, $f1 = 83.9$). The results of the German polarity identification experiments, with an F1-Measure of 80.3 - 87.6 are quite promising.

Acknowledgment

This paper is based on work that was first reported in [41]. The GermanPolarityClues resource was first reported in [40]. We gratefully acknowledge financial support of the German Research Foundation (DFG) through the EC 277 *Cognitive Interaction Technology* at Bielefeld University.

References

1. Thorsten Joachims. *Learning to Classify Text Using Support Vector Machines: Methods, Theory and Algorithms*. Kluwer Academic Publishers, Norwell, MA, USA, 2002.
2. Alexander Mehler, Peter Geibel, and Olga Pustyl'nikov. Structural classifiers of text types: Towards a novel model of text representation. *Journal for Language Technology and Computational Linguistics (JLCL)*, 22(2):51–66, 2007.
3. Bo Pang and Lillian Lee. *Opinion Mining and Sentiment Analysis*. Now Publishers Inc, July 2008.
4. Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. Thumbs up?: sentiment classification using machine learning techniques. In *EMNLP '02: Proceedings of the ACL-02 conference on Empirical methods in natural language processing*, pages 79–86, Morristown, NJ, USA, 2002. Association for Computational Linguistics.
5. Kushal Dave, Steve Lawrence, and David M. Pennock. Mining the peanut gallery: opinion extraction and semantic classification of product reviews. In *WWW '03: Proceedings of the twelfth international conference on World Wide Web*, pages 519–528. ACM Press, 2003.
6. Mingqing Hu and Bing Liu. Mining and summarizing customer reviews. In *KDD '04: Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177, New York, NY, USA, 2004. ACM.
7. Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. Recognizing contextual polarity in phrase-level sentiment analysis. In *HLT '05: Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, pages 347–354, Morristown, NJ, USA, 2005. Association for Computational Linguistics.
8. Michelle Annett and Grzegorz Kondrak. A comparison of sentiment analysis techniques: Polarizing movie blogs. In *Canadian Conference on AI*, pages 25–35, 2008.
9. Christian Becker-Asano and Ipke Wachsmuth. Affective computing with primary and secondary emotions in a virtual human. *Autonomous Agents and Multi-Agent Systems*, 2009.
10. Simon Clematide and Manfred Klenner. Evaluation and extension of a polarity lexicon for german. In *Proceedings of the 1st Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA)*, 2010.
11. Robert Remus, Uwe Quasthoff, and Gerhard Heyer. Sentiws – a publicly available german-language resource for sentiment analysis. In *Proceedings of the 7th International Language Resources and Evaluation (LREC'10)*, pages 1168–1171, 2010.
12. Songbo Tan and Jin Zhang. An empirical study of sentiment analysis for chinese documents. *Expert Syst. Appl.*, 34(4):2622–2629, 2008.
13. Pimwadee Chaovalit and Lina Zhou. Movie review mining: a comparison between supervised and unsupervised classification approaches. *Hawaii International Conference on System Sciences*, 4:112c, 2005.
14. Rudy Prabowo and Mike Thelwall. Sentiment analysis: A combined approach. *J. Informetrics*, 3(2):143–157, 2009.
15. Peter D. Turney and Michael L. Littman. Unsupervised learning of semantic orientation from a hundred-billion-word corpus. *CoRR*, cs.LG/0212012, 2002.
16. Alistair Kennedy and Diana Inkpen. Sentiment classification of movie reviews using contextual valence shifters. *Computational Intelligence*, 22(2):110–125, 2006.

17. Peter D. Turney. Thumbs up or thumbs down?: semantic orientation applied to unsupervised classification of reviews. In *ACL '02: Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, pages 417–424, Morristown, NJ, USA, 2001. Association for Computational Linguistics.
18. Maite Taboada, Julian Brooke, and Manfred Stede. Genre-based paragraph classification for sentiment analysis. In *Proceedings of the SIGDIAL 2009 Conference*, pages 62–70, London, UK, September 2009. Association for Computational Linguistics.
19. Apoorv Agarwal, Fadi Biadisy, and Kathleen McKeown. Contextual phrase-level polarity analysis using lexical affect scoring and syntactic n-grams. In *EACL2009*, Athens, Greece, 2009.
20. Pang and Lee. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *In Proceedings of the ACL*, pages 271–278, 2004.
21. Jaap Kamps Maarten, Maarten Marx, Robert J. Mokken, and Maarten De Rijke. Using wordnet to measure semantic orientations of adjectives. In *National Institute for*, pages 1115–1118, 2004.
22. Christiane Fellbaum, editor. *WordNet. An Electronic Lexical Database*. The MIT Press, 1998.
23. Tony Mullen and Nigel Collier. Sentiment analysis using support vector machines with diverse information sources. In Dekang Lin and Dekai Wu, editors, *Proceedings of EMNLP 2004*, pages 412–418, Barcelona, Spain, July 2004. Association for Computational Linguistics.
24. Hong Yu and Vasileios Hatzivassiloglou. Towards answering opinion questions: Separating facts from opinions and identifying the polarity of opinion sentences. In *Proceedings of EMNLP'03*, 2003.
25. Satoshi Sekine Kugatsu Sadamitsu and Mikio Yamamoto. Sentiment analysis based on probabilistic models using inter-sentence information. In Bente Maegaard Joseph Mariani Jan Odjik Stelios Piperidis Daniel Tapias Nicoletta Calzolari (Conference Chair), Khalid Choukri, editor, *Proceedings of the Sixth International Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco, may 2008. European Language Resources Association (ELRA).
26. Michael Wiegand and Dietrich Klakow. The role of knowledge-based features in polarity classification at sentence level.
27. Bo Pang and Lillian Lee. Seeing stars: exploiting class relationships for sentiment categorization with respect to rating scales. In *ACL '05: Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pages 115–124, Morristown, NJ, USA, 2005. Association for Computational Linguistics.
28. Bing Liu. Sentiment analysis and subjectivity. *Handbook of Natural Language Processing*, 2:568, 2010.
29. Vasileios Hatzivassiloglou and Kathleen R. McKeown. Predicting the semantic orientation of adjectives. In *Proceedings of the eighth conference on European chapter of the Association for Computational Linguistics*, pages 174–181, Morristown, NJ, USA, 1997. Association for Computational Linguistics.
30. Philip J. Stone, Dexter C. Dunphy, Marshall S. Smith, and Daniel M. Ogilvie. *The General Inquirer: A Computer Approach to Content Analysis*. MIT Press, 1966.
31. C. Strapparava and A. Valitutti. WordNet-Affect: an affective extension of WordNet. In *Proceedings of LREC*, volume 4, pages 1083–1086, 2004.
32. Janyce Wiebe, Theresa Wilson, and Claire Cardie. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation*, 1(2):0, 2005.

33. Janyce Wiebe and Ellen Riloff. Creating subjective and objective sentence classifiers from unannotated texts. In *Proceeding of CICLing-05, International Conference on Intelligent Text Processing and Computational Linguistics.*, volume 3406 of *Lecture Notes in Computer Science*, pages 475–486, Mexico City, MX, 2005. Springer-Verlag.
34. Andrea Esuli and Fabrizio Sebastiani. Sentiwordnet: A publicly available lexical resource for opinion mining. In *In Proceedings of the 5th Conference on Language Resources and Evaluation (LREC06)*, pages 417–422, 2006.
35. Hiroya Takamura, Takashi Inui, and Manabu Okumura. Extracting semantic orientations of words using spin model. In *ACL '05: Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pages 133–140, Morristown, NJ, USA, 2005. Association for Computational Linguistics.
36. David Chandler. *Introduction to Modern Statistical Mechanics*. Oxford University Press, 1987.
37. Ulli Waltinger. Polarity reinforcement: Sentiment polarity identification by means of social semantics. In *Proceedings of the IEEE Africon 2009, September 23-25, Nairobi, Kenya*, 2009.
38. T. Joachims. SVM light, <http://svmlight.joachims.org>, 2002.
39. Kerstin Denecke. Using sentiwordnet for multilingual sentiment analysis. In *ICDE Workshops*, pages 507–512. IEEE Computer Society, 2008.
40. Ulli Waltinger. Germanpolarityclues: A lexical resource for german sentiment analysis. In Bente Maegaard Joseph Mariani Jan Odjik Stelios Piperidis Mike Rosner Daniel Tapias Nicoletta Calzolari (Conference Chair), Khalid Choukri, editor, *Proceedings of the Seventh conference on International Language Resources and Evaluation (LREC'10)*, Valletta, Malta, may 2010. European Language Resources Association (ELRA).
41. Ulli Waltinger. Sentiment analysis reloaded: A comparative study on sentiment polarity identification combining machine learning and subjectivity features. In *Proceedings of the 6th International Conference on Web Information Systems and Technologies (WEBIST 2010), Valencia*, 2010.