

Polarity Reinforcement: Sentiment Polarity Identification By Means Of Social Semantics

Ulli Waltinger
Text Technology
Bielefeld, Germany
ulli_marc.waltinger@uni-bielefeld.de

Abstract—We propose a combination of machine learning and socially constructed concepts for the task of sentiment polarity identification. Detecting words with polarity is difficult not only due to limitations in current sentiment dictionaries but also due to the colloquial terms that are often used. Current approaches disregard the dynamics of language, i.e. that new words are often created comprising different polarities. In fact, the online community is very creative in coining terms about certain subjects such as "tweetup" (a request by a user to meet with friends via Twitter) or "whack" (Street slang, meaning bad). Our approach utilizes a user generated dictionary of urban term definitions as a resource for polarity concepts. Therefore, we are not only able to map newly created words to their respective polarity but also enhance common expressions with additional features and reinforce the polarity, strengthening our initial finding. We empirically show that the use of polarity reinforcement improves the sentiment classification.

I. INTRODUCTION

In recent years, sentiment analysis has received considerable attention in the research community. The task of polarity or emotion identification may contribute to a broad variety of possible applications such as recommender systems, the collection of opinions in product reviews [1], [2], [3], [4], in financial news [5] and also in the domain of human-computer-interaction [6], [7], [8]. In general, sentiment analysis refers to the task of extracting and assessing sentiments written in text documents. Sentiment classification focuses on the assignment of polarity labels (e.g. positive, neutral, negative) for a given text fragment. Sophisticated systems additionally introduce polarity rating scales (e.g. 1 to 5 stars), referred to as *rating inference* [9]. The task of sentiment analysis typically determines the semantic orientation on the level of words [10], sentences [11], phrases [12] or on entire documents [1]. Analysis on different textual levels is needed, because positive as well as negative expressions and opinions may occur within a document. Most existing approaches focus thereby on counting, summing or selecting certain lexical characteristics conducting Maximum Entropy (ME), Naive Bayes (NB) or Support Vector Machine implementations (SVM) [13], [9]. Most unsupervised and semi-supervised methods make use of annotated lexicons as a resource for feature interpretations. The most prominent examples are SentiWordNet [14] and *WordNet-Affect* [15]. Both resources offer a set of words with associated semantic orientations. That is, positive, neutral and negative polarity scores are given. In a different approach, Takamura (2005) et al. [16] presented a method for extracting the semantic

orientations of words using the so called *spin model*. This method automatically creates polarity-related word lists out of a dictionary or a thesaurus. However, sentiment thesauri, dictionaries and resources are mostly limited to a small set of synset relations or single tokens and are obviously very 'expensive' to build, since expert knowledge is involved. That is, they rely on the work of lexicographers with a focus on standard language or even high level language. Therefore, the coverage of such dictionaries differs due to the diverse user vocabularies, especially when handling direct user input as for example by conversational agents [17], chat conversations, or with respect to the various interactive online services (e.g. twitter¹, weblogs or message boards). In fact, the online community is very creative in coining new terms with semantic orientation. For example expressions such as "take a digger" (commonly used to describe an ungraceful fall), "whack" (meaning bad) or "douche bag" (someone you do not like) are in the majority of cases not covered by the expert-built resources stated above. However, we can infer the semantic orientation of these phrases and associated definitions. The questions that arise therefore are: which additional resource(s) are suitable for the extraction of words with semantic orientation with regards to colloquial language? How to circumvent the mainly static property of existing sentiment resources? How to detect the sentiment polarity of colloquial phrases? Does the consideration of colloquial language improve the classification of sentiment polarity? In this paper, we propose a method for enhancing current sentiment resources with additional colloquial-language-based phrases using a web-based dictionary. Thereby, colloquial phrases are interlinked to terms with semantic orientation labels. These connections are further used for a polarity-related enhancement. By expanding the existing document representation using terms with associated semantic orientation, a reinforcement of polarity features is performed. For illustration purposes consider the following example: if we detect the word 'wicked' in our document, we add the terms 'excellent', 'awesome' and 'good' to it. If the term 'lame' is found, we add 'stupid', 'boring' and 'bad' to it ... and so on. This attempt supports the task of polarity exaggeration (an overstatement of positive and negative po-

¹Twitter refers to a social messaging and micro-blogging service, which enables users to send and read text-based messages (*tweets*) in real-time. <http://twitter.com/>

larities by terms with similar sentiment orientation). Frequent appearances of negative words in a product or film review – identified through a list of words with semantic relevance – are therefore reinforced, resulting in a truly negative text. The goal of this paper is to propose a method for automatically enhancing words of polarity with associated terms and phrases out of a human-created resource. Doing this, we apply a polarity approximation by reinforcement. We empirically show that feature enhancement contributes to sentiment analysis. We apply a SVM and N-gram based methods for the task of polarity identification using a corpus of 2000 movie reviews.

II. RELATED WORK

Sentiment polarity identification refers to the task of text classification. Hence, texts are classified based on the sentimental orientation of opinions they contain instead of topics. Therefore, different features selection and enhancement methods have to be applied. See [18] for an overview of the computational treatment of feature selection in the domain of sentiment analysis in texts. In classical text categorization scenarios, feature enhancement has been shown to be successful [19] - improving the classification accuracy. Hence, traditional resources for feature enhancement like *WordNet* [20], *Open Directory Project*² [21] or most recently the *Wikipedia*³ dataset [22], [23], [24] are mostly used in order to overcome the data sparseness problem in generating topic-related features. However, enhancing polarity terms differs to recent approaches and resources since newly-created descriptors need to have sentiment orientation information assigned. With regards to the creation of sentiment-analysis-related resources for feature generation, only few are available. Most of them form the basis on the famous linguistic resources *WordNet* [25]. Hatzivasiloglou et al. (1997) and Maarten et al. (2004) [26], [10] used *WordNet* for measuring semantic orientations of adjectives. As an extension of *WordNet*, Strapparava et al. (2004) [15] presented *WordNet-Affect* – a resource of polarity values for words within *WordNet*. Similar to it, Esuli et al.(2006) [14] presented *SentiWordNet* as an additional sentiment resource. Also on the basis of *WordNet*, Takamura et al. (2005) [16] presented an approach for semantic term orientations – offering a list of around 80,000 words for the English language. Even though these resources have been built by experts, the size of these semantic dictionaries are rather low – not covering a quite significant number of the most recent expressions (e.g. "That's The Shizzle" = That is good). One might think that expressions like that are used in oral communication by younger people only. This might be true, but is also exactly the point we are trying to make. The language used in e.g. chat forums, email message or talks e.g. conversational agents differs from those in a written article in the e.g. Wikipedia dataset or newspaper contributions.

²<http://www.dmoz.org/>

³<http://www.wikipedia.org/>

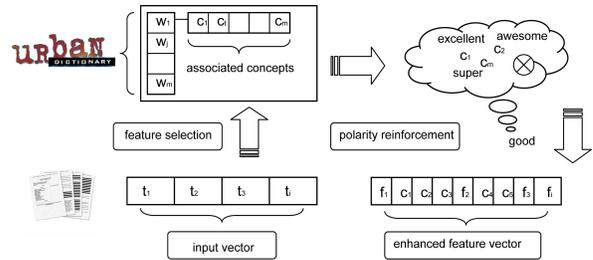


Fig. 1. Sentiment polarity reinforcement architecture

III. METHODOLOGY

In order to incorporate additional polarity related term definitions, we have developed a process that is able to map dictionary defined words to associated terms with sentiment orientation labels. Note that our focus is to enhance the existing document representation by polarity-related terms building *feature clouds*. The sentiment classification procedure (see Figure 1) comprises four steps described in this section.

Knowledge Base Processing: as a key point, we have converted a web-based dictionary offering a number of different slang words and phrases with associated definitions. A detailed description about the preprocessing of this resource will be given in Section III-A. We will refer to the preprocessed web-dictionary as our knowledge base (KB). Each entry (word or phrase) of the dictionary is defined as our polarity concept.

Feature Candidate Selection: Since the polarity concepts consist of single words as well as multi-word expressions, the mapping of term candidates of an input text cannot be computed by querying KB directly. We therefore introduce a sliding window, when extracting textual information of our input text. That is, we consider single tokens as well as multi-word expressions within the given window size of 5. For each feature candidate, a search query is sent to KB.

Feature Candidate Enhancement: Having detected an entry in KB for an expression, we request those related polarity concepts that are on the basis of KB connected to the search entry. A fixed number of related polarity concepts are then added to the concept representation (CR).

Polarity Identification: The last step in the process is to use the new enhanced concept representation CR for the polarity classification (see Section IV). Note, we used only the features of CR for the classification task.

A. Polarity Concepts by Urban Dictionary

For the purpose of feature enhancement, a knowledge base (KB) is needed. In this paper, we make use of the dataset provided by the online project Urban Dictionary (UD)⁴. This resource launched in 1999 and is currently offering 3,933,862 definitions contributed by users. In short, UD allows users to introduce new term definitions – mostly slang words – or to rate existing definitions (see Figure 2). Due to its online

⁴<http://www.urbandictionary.com>



Fig. 2. Urban Dictionary Screenshot

dynamic, this dictionary is constantly growing and entries are rated in a 'thumbs up/down' manner. Each dictionary entry is described through multiple descriptions (see Table I) contributed by the online community. Each description is

TABLE I
URBAN DICTIONARY CONCEPT WITH ASSOCIATED DESCRIPTORS.

Seed Word	Concept Description
Tight	Stylish, cool, having everything together; Scrooge cool or hip
	Original use: the mood or spirit between members of a music group who are having a particularly excellent performance.
	Modern use: describes something particularly enjoyable or awesome.

rated on the basis of its relevance to the source term. Most importantly, UD offers a list of the most frequently-associated concepts on the basis of their assigned descriptors (See Table II for an example of associated concepts for a seed word). Most of these associated concepts (AC) are based on semantic relation (synonymy, meronymy...). However, it should be noted that quite a few words are abusive.

B. Feature Enhancement For Polarity Identification

Our focus is the identification of sentiment polarities by enhanced polarity concepts. The task of accumulation is done by embracing the UD data set. In terms of feature selection, for a given word or multi-word a search query is sent to KB - the UD data set. In order to compute this feature selection efficiently, we build an inverted index representation of KB. In doing this, we make use of an existing list of words with polarity values assigned. Note that each new generated feature of KB needs to have a sentiment orientation. Therefore, we also utilize the data set of Takamura et al. (2005) [16], which assembled a list of 80,000 words including their semantic polarity orientation, defined as T . Each entry of T poses as a

TABLE II
SEED WORDS ENRICHED BY URBAN DICTIONARY BASED ASSOCIATED CONCEPTS.

Seed Word	Associated Concepts
awesome	sound cool, amazing, sweet, great, rad, sexy, hot, sick, fantastic, wicked ...
marvelous	wonderful, superb, excellent, stunning, splendid, magnificent, sweet, terrific, mind-blowing ...
shiznit	shizzle, tight, sweet, poop, the bomb, foshizzle, awesome, badass, aiiight ...
crappy	shitty, lame, stupid, horrible, sucky, useless, rubbish, janky, cruddy, cheesy

TABLE III
SEED WORD DISAMBIGUATION BY URBAN DICTIONARY BASED ASSOCIATED CONCEPTS.

Seed Word	Associated Concepts
is shit	crap, poop, fuck, ass, damn, dookie...
is the shit	cool, awesome, amazing, sweet, thigh, great ...

seed word w_i in requesting new associated concepts of KB. That is, for each seed word we retrieved the term definitions – represented as a set of urban concepts – of UD. Let s_i be a vector of a polarity seed word w_i , its polarity value k_i and its associated urban concepts $u_j = \{u_t, \dots, u_M\}$. M defines the total number of considered concepts. Each u_t inherits the polarity value k_i of w_i – the predefined semantic orientation. That is, the size of S equals the size of T . In order to remove concepts with bi-polarities (terms assigned to positive and negative seed words), we need to compute the polarity probability of a given concept u_t , given a positive class C_{pos}

$$p(u_t|C_{pos}) = \prod_i p(w_i|C_{pos}) \quad (1)$$

and negative class C_{neg}

$$p(u_t|C_{neg}) = \prod_i p(w_i|C_{neg}) \quad (2)$$

respectively. We define the unique polarity class C_{pol} as the final membership of concept u_t as:

$$p(u_t|C_{pol}) = \max\{p(u_t|C_{pos}), p(u_t|C_{neg})\} \quad (3)$$

Therefore, we are able to request for each word or multi-word of a given text fragment a set of additional polarity-related concepts. This feature enhancement method forms the basis of the polarity reinforcement used for the sentiment classification task.

IV. EXPERIMENTAL SETUP

In order to evaluate the performance of feature enhancement in the domain of polarity identification, we conducted two

Feature	Frequency	Rank
worthwhile	532	1
witty	412	2
chancellery	298	3
closely	268	4
goodness	143	5
...

TABLE IV
FEATURE RANKING OF TOKEN N-GRAMS SORTED BY RANK AND FEATURE FREQUENCY.

different classification procedures. First, we used a support vector machine implementation, the most commonly used approach in text categorization and sentiment analysis [11], [27], [28], [29]. Secondly, an n-gram classification [30], [1], [31] is applied to the domain of sentiment analysis. Both settings are conducted with and without feature enhancement. In the next two sections we briefly sketch the different settings of the experimental setup.

A. Polarity SVM Classification (SVM-C)

In the first experiment, we tested the support vector machine (SVM) implementation. Following Pang et al. (2002), we used unigram-presence features for the classification. Different to them, we enhanced each training and evaluation document with additional features on the basis of KB (as described in the previous section). We conducted the feature selection for each input document on the basis of our inverted index as proposed in the previous section. Therefore, only features with a direct semantic orientation are used for the classification task. For each polarity status (positive and negative), a SVM is trained using the SVM-light implementation by Joachims [32] with a linear kernel.

B. Polarity Difference Classification (PDC)

With regards to the n-gram classification, we merely follow Waltinger et al.(2009)[31] in building a feature frequency ranking of different n-gram models using the feature difference coefficient. Therefore, for each polarity status a classifier is trained comprising different textual n-gram models. What is different in our approach is the fact that we are additionally conducting the feature enhancement technique before model creation. In general, we build three different n-gram models comprising:

- Token n-grams, occurrences of n-gram using a sequence of n word forms.
- Letter N-Grams, occurrences of n-gram using a sequence of n characters
- Part of Speech Tag Frequency, occurrences of n-gram using a sequence of n Part of Speech Tags.

Previous approaches additionally comprised structural features, hence preliminary studies showed that those features had no impact on the performance of the system. The parameter setting for the corresponding n-gram information was set to $n = 1$ to $n = 5$. Document classification of unknown polarity

TABLE VI
ACCURACY RESULTS COMPARING BASELINE APPROACHES, SVM AND PDC IMPLEMENTATION.

Sentiment-Method	Accuracy
Naive Bayes - unigrams [1]	78.7
Maximum Entropy - top 2633 unigrams [1]	81.0
SVM - unigrams+bigrams [1]	82.7
SVM - unigrams [1]	82.9
PDC without feature enhancement	81.9
PDC with feature enhancement	83.1

TABLE VII
EVALUATION RESULTS OF FEATURE ENHANCED CLASSIFICATION ON THE BASIS OF PRECISION, RECALL AND F1-MEASURE.

Model	Precision	Recall	F-Measure
SVM: Positive-Model	.850	.832	.841
SVM: Negative-Model	.828	.847	.837
SVM - Overall	.839	.840	.839
PDC: Positive-Model	.850	.807	.828
PDC: Negative-Model	.810	.844	.827
PDC - Overall	.830	.826	.828

is done measuring the distance between input and trained n-gram profiles (see Table IV) on the basis of their frequency profiles (see Table V).

C. Corpus

For the evaluation of our approach, we conducted experiments using the movie review corpus compiled by Pang et al. (2002) [1] – the most popular dataset in sentiment analysis. The corpus consists of 1000 positive and 1000 negative reviews. The corpus was compiled by 312 authors written before 2002. With regards to the PDC approach, we conducted a three-fold cross-validation, and reporting Precision, Recall, F1-Measure and Accuracy. With regards to the SVM approach, we used the leave-on-out accuracy measure by SVM-light [33].

V. RESULTS

From Table VI we can see that only mediocre results can be achieved using the normal PDC implementation. That is, only tokens of the actual documents are used for the n-gram classification and no feature enhancement is conducted. However, using the polarity reinforcement technique the average accuracy rises to 83.1 outperforming the existing baseline approaches proposed by Pang et al. (2002) [1]. As we can see in Table VII comparing the F1-Measure of PDC-Enhanced and the corresponding SVM-Enhanced implementation, we are able to boost the results to an F-Measure of 0.839. Both enhanced methods perform very well. We noted that, the performance of the very simple method of n-gram classification is especially astonishing. Using only n-gram frequencies outperforms the best SVM baseline. Comparing the non-enhanced SVM with their reinforced version we can also see an improvement of .10. The main contribution presented here is the incorporation of polarity concepts as features into a n-gram and a SVM classifier. In contrast to previous approaches, which relied more or less entirely on and therefore were limited to the

Model	Frequency	Rank (r_m)	Document	Frequency	Rank (r_d)	Distance
succeed	634	1	playful	223	1	max
exonerate	477	2	exonerate	212	2	0
mate	398	3	positivity	134	3	2
playful	378	4	perfect	98	4	max
positivity	245	5	cool	87	5	max
...
Distance (D_{md})						2+max

TABLE V
FEATURE DISTANCE RANKING OF LETTER N-GRAMS

bag-of-words [34] model, our method circumvents this issue by adding associated polarity concepts that might not have occurred before. In general, we can argue that polarity feature enhancement contributes to the task of sentiment analysis. To our knowledge, feature enhancement by polarity concepts in the area of sentiment analysis has not been done before. Hence, we could examine whether polarity or feature reinforcement improves evaluation accuracy on the basis of two different classification implementations. However, the level of improvement is lower than expected but nevertheless outperforms the well-known baseline. For the future it would be interesting to analyze the performance within different settings as for example in a chat or human-computer-interaction environment.

VI. CONCLUSION

We examined the impact of feature generation for the task of sentiment polarity identification, showing that polarity reinforcement improves the evaluation accuracy over previous experiments. Previous approaches based their methods mostly on a limited sentiment dictionary comprising only features of the training document collection – disregarding the dynamic of language, i.e. that new words are often created comprising different polarities. We proposed a combination of machine learning and polarity feature enhancement for the task of sentiment polarity identification. We presented a method for detecting the sentiment polarity of colloquial language using the dataset of the Urban Dictionary project. Therefore, we were not only able to map newly-created words to its polarity but also enhance common expressions with additional features and reinforce the polarity. We empirically showed that the performance of polarity reinforcement improves the sentiment classification.

ACKNOWLEDGMENT

We gratefully acknowledge financial support of the German Research Foundation (DFG) through the EC 277 *Cognitive Interaction Technology*, the SFB 673 *Alignment in Communication (X1)*, the Research Group 437 *Text Technological Information Modeling*, the DFG-LIS-Project *P2P-Agents for Thematic Structuring and Search Optimization in Digital Libraries* and the *Linguistic Networks* project funded by the German Federal Ministry of Education and Research (BMBF) at Bielefeld University.

REFERENCES

- [1] B. Pang, L. Lee, and S. 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*. Morristown, NJ, USA: Association for Computational Linguistics, 2002, pp. 79–86.
- [2] K. Dave, S. Lawrence, and D. 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*. ACM Press, 2003, pp. 519–528.
- [3] M. Hu and B. Liu, “Mining and summarizing customer reviews,” in *KDD '04: Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*. New York, NY, USA: ACM, 2004, pp. 168–177.
- [4] M. Annett and G. Kondrak, “A comparison of sentiment analysis techniques: Polarizing movie blogs,” in *Canadian Conference on AI*, 2008, pp. 25–35.
- [5] A. Devitt and K. Ahmad, “Sentiment polarity identification in financial news: A cohesion-based approach,” in *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*. Prague, Czech Republic: Association for Computational Linguistics, June 2007, pp. 984–991. [Online]. Available: <http://www.aclweb.org/anthology/P/P07/P07-1124>
- [6] M. A. Arbib and J.-M. Fellous, “Emotions: from brain to robot,” *Trends in Cognitive Sciences*, vol. 8, no. 12, pp. 554–561, December 2004.
- [7] S. Kopp, C. Becker, and I. Wachsmuth, “The virtual human max - modeling embodied conversation,” in *KI 2006 - Demo Presentation, Extended Abstracts*, 2006, pp. 21–24.
- [8] C. Becker-Asano and I. Wachsmuth, “Affective computing with primary and secondary emotions in a virtual human,” *Autonomous Agents and Multi-Agent Systems*, 2009.
- [9] B. Pang and L. 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*. Morristown, NJ, USA: Association for Computational Linguistics, 2005, pp. 115–124.
- [10] J. K. Maarten, M. Marx, R. J. Mokken, and M. D. Rijke, “Using wordnet to measure semantic orientations of adjectives,” in *National Institute for*, 2004, pp. 1115–1118.
- [11] Pang and Lee, “A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts,” in *In Proceedings of the ACL*, 2004, pp. 271–278.
- [12] T. Wilson, J. Wiebe, and P. 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*. Morristown, NJ, USA: Association for Computational Linguistics, 2005, pp. 347–354.
- [13] P. 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*. Morristown, NJ, USA: Association for Computational Linguistics, 2001, pp. 417–424.
- [14] A. Esuli and F. Sebastiani, “Sentiwordnet: A publicly available lexical resource for opinion mining,” in *In Proceedings of the 5th Conference on Language Resources and Evaluation (LREC06)*, 2006, pp. 417–422.
- [15] C. Strapparava and A. Valitutti, “WordNet-Affect: an affective extension of WordNet,” in *Proceedings of LREC*, vol. 4, 2004, pp. 1083–1086.
- [16] H. Takamura, T. Inui, and M. Okumura, “Extracting semantic orientations of words using spin model,” in *ACL '05: Proceedings of the*

43rd Annual Meeting on Association for Computational Linguistics. Morristown, NJ, USA: Association for Computational Linguistics, 2005, pp. 133–140.

- [17] C. Becker, S. Kopp, and I. Wachsmuth, “Why emotions should be integrated into conversational agents,” in *Conversational Informatics: An Engineering Approach*. Wiley, 2007, pp. 49–68.
- [18] B. Pang and L. Lee, “Opinion mining and sentiment analysis,” *Foundation and Trends in Information Retrieval*, vol. 2, no. 1-2, pp. 1–135, 2008.
- [19] E. Gabrilovich and S. Markovitch, “Feature generation for text categorization using world knowledge,” in *Proceedings of The Nineteenth International Joint Conference for Artificial Intelligence*, Edinburgh, Scotland, 2005, pp. 1048–1053.
- [20] S. B. Andreas and A. Hotho, “Boosting for text classification with semantic features,” in *In Proceedings of the MSW 2004 Workshop at the 10th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2004, pp. 70–87.
- [21] Gabrilovich and Markovitch, “Overcoming the brittleness bottleneck using wikipedia: Enhancing text categorization with encyclopedic knowledge,” *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, Boston, MA, 2006.
- [22] A. Mehler, “A short note on social-semiotic networks from the point of view of quantitative semantics,” in *Proceedings of the Dagstuhl Seminar on Social Web Communities, September 21-26, Dagstuhl*, H. Alani, S. Staab, and G. Stumme, Eds., 2008.
- [23] U. Waltinger, A. Mehler, and R. Gleim, “Social semantics and its evaluation by means of closed topic models: An svm-classification approach using semantic feature replacement by topic generalization,” in *Proceedings of the GSCL-Conference, Potsdam (DE), 2009.*, 2009.
- [24] U. Waltinger and A. Mehler, “Social semantics and its evaluation by means of semantic relatedness and open topic models,” in *Proceedings of the 2009 IEEE/WIC/ACM International Conference on Web Intelligence*, 2009.
- [25] C. Fellbaum, Ed., *WordNet. An Electronic Lexical Database*. The MIT Press, 1998.
- [26] V. Hatzivassiloglou and K. R. McKeown, “Predicting the semantic orientation of adjectives,” in *Proceedings of the eighth conference on European chapter of the Association for Computational Linguistics*. Morristown, NJ, USA: Association for Computational Linguistics, 1997, pp. 174–181.
- [27] T. Mullen and N. Collier, “Sentiment analysis using support vector machines with diverse information sources,” in *Proceedings of EMNLP 2004*, D. Lin and D. Wu, Eds. Barcelona, Spain: Association for Computational Linguistics, July 2004, pp. 412–418.
- [28] S. Tan and J. Zhang, “An empirical study of sentiment analysis for chinese documents,” *Expert Syst. Appl.*, vol. 34, no. 4, pp. 2622–2629, 2008.
- [29] A. Abbasi, H. Chen, and A. Salem, “Sentiment analysis in multiple languages: Feature selection for opinion classification in web forums,” *ACM Trans. Inf. Syst.*, vol. 26, no. 3, pp. 1–34, 2008.
- [30] W. B. Cavnar and J. M. Trenkle, “N-gram-based text categorization,” 1994, pp. 161–175.
- [31] U. Waltinger and A. Mehler, “The feature difference coefficient: Classification by means of feature distributions,” in *Proceedings of Text Mining Services (TMS), March 23-25, Leipzig, Germany*, 2009.
- [32] T. Joachims, *Learning to Classify Text Using Support Vector Machines: Methods, Theory and Algorithms*. Norwell, MA, USA: Kluwer Academic Publishers, 2002.
- [33] —, “SVM light, <http://svmlight.joachims.org>,” 2002.
- [34] G. Salton, *Automatic Text Processing: The Transformation, Analysis, and Retrieval of Information by Computer*. Reading, Massachusetts: Addison Wesley, 1989.