

Quantitative comparison of polarity lexicons in sentiment analysis tasks: Using a lexicon overlap score for similarity measurement between lexicons

Felix J. M. Welter¹

Abstract: Sentiment classification is either based on sentiment lexicons or machine learning. For the construction and improvement of sentiment lexicons, several approaches and algorithms have been designed. The resulting lexicons are commonly benchmarked in different tasks and compared by their respective performance. However, this measure depends on the application domain. This work proposes a method for context-independent comparison of sentiment lexicons. Three scoring methods for similarity measurement of lexicons are explained. Furthermore, exemplarily applications of the scores are shown, including lexicon similarity analysis before and after expansion via a Distributional Thesaurus and clustering of lexicons. Adaptability and limitations of the lexicon overlap score and the demonstrated applications are discussed.

Keywords: sentiment analysis, sentiment classification, semantic orientation, sentiment lexicons, polarity lexicons, lexicon similarity, lexicon distance, lexicon overlap score

1 Introduction

Sentiment analysis is a field of Natural Language Processing (NLP) and focuses on the extraction of polarity and opinion targets in texts of any size. With the growing amount of possibilities for individuals to share content online, opinion-rich resources become available. Since they are heavily used by customers for decision making and often contain patterns and signals about customer behavior, organizations have a reasonable interest in employing advanced text mining techniques to extract valuable information. [PL08] For deeper insights, sentiment analysis can be combined with other NLP techniques such as text summarization. [To01] Application domains of sentiment analysis include, but are not limited to, political speeches [GST16], web advertising [AA16], [Qi10], and social media [VS16].

This work revolves around the subtask of sentiment classification, which involves the detection of the semantic polarity of a piece of text. Classification is most often either binary (positive or negative), ternary (positive, neutral or negative) or 5-scale ordinal (often visualized by using stars). [RFN17] It should be noted that the task of sentiment classification is generally specific to a context. The polarity of words varies between topics, domains and also changes historically. [Re05], [Ha16] Further challenges are imposed by misspellings [Mu13], negations [Po10], emoticons [Re05], multi word

¹ Universität Hamburg, Vogt-Kölln-Straße 30, 22527 Hamburg, felix.welter@studium.uni-hamburg.de

expressions [Sa02], [WWH05], idioms [Jo18], [NSW94] and language specific characteristics [ACS08].

For sentiment classification, one can generally use two approaches. While they are mostly used independently, they can also be combined. [Zh11] The first approach is based on a lexicon. This involves the creation of a list of words where each word has a semantic orientation assigned. The polarity of a piece of text is then determined by the contained words and the corresponding scores. The second approach is based on machine learning classifiers, e.g. Naïve Bayes, Support Vector Machines or Neural Networks. Features, such as contained words, part-of-speech tags, and punctuation, are extracted from the text and the classifier is trained to predict the polarity. [VS16]

The following paper focuses on the first approach. Current methods for lexicon creation contribute to an increase in lexicon size. Furthermore, the construction of lexicons becomes highly automated. [Ku16], [Ta14], [TL02], [Ve10], [YS14] This is excellent since most applications should construct a domain specific sentiment lexicon and the manual creation or construction involving crowd work is expensive. [MT10] However, this poses the challenge that many lexicons are available, which are only different in the construction algorithm, the underlying dataset and the performance in the target domain. A trivial method to compare lexicons without these context factors is missing.

This work presents a possibility to compare different sentiment lexicons solely based on their content. Three different but closely related scores are described. Lexicon size, common words and the similarity of the semantic orientation of each word are used for computation. No external resources are required, and the score is language independent. Example applications of the proposed lexicon overlap score are given.

2 Related Work

Since human-made sentiment lexicons yield poor results, [PLV02] research in lexicon construction for sentiment classification designed different algorithms for automatic lexicon construction. Hatzivassiloglou and McKeown [HM97] extracted words of the same polarity by searching for adjectives connected by the word *and*. Adjectives connected by *but* were supposedly of opposite polarity. Retrieving a multitude of these collocations from a corpus and clustering found words into two groups enabled them to construct a positive and a negative word list.

Turney [Tu02] designed a method where phrases containing adjectives were scored based on their cooccurrence with the words *excellent* and *poor*. For both cases the Point Wise Mutual Information (PMI) is calculated and the semantic polarity of a phrase is defined by the difference between the PMI with *excellent* and the PMI with *good*.

More recent work has put a focus on graph-based approaches. Relations between words or phrases are captured based on linguistic resources like WordNet [Pr10] or inferred

from corpora. Semantic orientation is then propagated between words which results in a polarity score for each word in the graph. This technique is employed in the Social Sent Project [Ha16] and by [Ve10].

Another approach to enriching sentiment lexicons is based on a Distributional Thesaurus, which is a lexical resource that contains information about semantically similar words. For highly positive or negative words or phrases, closely related terms can be found. By adding these terms to the sentiment lexicon, a higher coverage and therefore a higher sentiment classification performance can be achieved. [Ku16]

3 Lexicon overlap score

With the increase of automated lexicon creation methods for different languages and an upsurge in lexicon size a formalized metric for lexicon comparison is needed. This score should be language independent and solely taking the lexicon content into account. The proposed lexicon overlap score (LOS) takes into account lexicon size, presence of words, semantic orientation and the strength of the semantic orientation. Three different versions of the score exist. Since they are based on each other, they will be explained in order with rising complexity. Especially the first score is closely related to the widely known Jaccard index, sometimes also called Tanimoto similarity. All scores have in common that identical lexicons have a LOS of 1 and lexicons without mutual words have a score of 0. Furthermore, it is assumed that semantic orientation of a word larger than zero is interpreted as positive and semantic orientation smaller than zero interpreted as negative, as it is used by most algorithms.

3.1 Simple LOS

The simple version of the lexicons overlap score takes into account the size of both lexicons and the words they have in common. A denotes the size of the first lexicon, B the size of the second lexicon and n the count of common words, then the simple LOS is expressed by the following.

$$\frac{n}{A + B - n}$$

The simple LOS states to what extent the lexicons share a common vocabulary. It does not quantify the relation between the polarity strength of each word in the lexicons, not even about the general direction of the polarity. Depending on the lexicon creation algorithm that was used, it can be applied to find shared words of different domains. Given a fixed domain it can be used to quantify to what extent different algorithms extract a mutual vocabulary from a dataset.

3.2 Binary LOS

Building upon the simple LOS, the binary LOS additionally draws a distinction between positive and negative polarity of words. Words are only counted for the numerator if they are both positive or both negative and therefore share the same general semantic orientation. A denotes the size of the first lexicon, B the size of the second lexicon and n the count of common words. Furthermore, each word that both lexicons include, denoted by w_1, \dots, w_n , is modified by the function X which evaluates to one, if the polarities share the same sign, otherwise to zero. Given a word w the function S_a evaluates to the polarity score of the first lexicon and S_b evaluates to the polarity score of the second lexicon. The binary LOS is then expressed by the following.

$$\frac{\sum_{i=1}^n X(w_i)}{A + B - n}$$

and

$$X(w) = \begin{cases} 1, & 0 < S_a(w) * S_b(w) \\ 0, & \text{else} \end{cases}$$

Since positivity and negativity of single words can strongly vary between domains (*crazy* in sports and *crazy* in relationships [Hal6]) it should be assured that only words are counted that are similarly quantified by both lexicons. This is especially helpful for comparison of lexicons that have different underlying algorithms that in general assign strongly diverging scores. Alternatively, polarity scores can be normalized or ranked and then compared by using the general LOS, which is explained in the following part. Finally, depending on the context it may be beneficial to put a stronger punishment on words with opposite polarity. In that case, X could be modified to return -1 instead of zero for these words.

3.3 General LOS

Building upon the binary LOS, the general LOS not only takes the direction of each word polarity into account, but instead weighs them by their relative similarity. A denotes the size of the first lexicon, B the size of the second lexicon and n the count of common words. For each word present in both lexicons w_1, \dots, w_n the relative similarity is calculated by the function Y . The case where both scores are equal is individually treated to correctly process zero values. Y can evaluate to negative values, which occurs whenever in one lexicon the word in question is assigned a positive value and in the other it is assigned a negative value. S_a and S_b return the polarity for the given word from their respective lexicon. The general LOS is then expressed by the following.

$$\frac{\sum_{i=1}^n Y(w_i)}{A + B - n}$$

and

$$Y(w) = \begin{cases} 1 & , \quad S_a(w) = S_b(w) \\ \frac{S_a(w)}{S_b(w)} & , \quad |S_a(w)| < |S_b(w)| \\ \frac{S_b(w)}{S_a(w)} & , \quad \textit{else} \end{cases}$$

By utilizing the relative similarity of the polarities of each word, the general LOS is the most precise score of the ones proposed in this work. By default, it should be used for lexicon comparison. Regarding the general LOS, lexicons can also receive a negative score. This will practically occur if two lexicons contain several words that have opposite polarities assigned. However, it is more challenging to interpret compared to the simple and binary LOS. Without reference scores, the result is of low value. The following section will therefore elaborate on usage examples.

4 Example Applications

This section shows applications of the lexicon overlap score. LOS distributions are shown, and interpretation is demonstrated. For the computations numpy [vCV11], pandas [Mc10], matplotlib [Hu07] and scikit-learn [Pe11] were employed. This will additionally evaluate the utility and validity of the score for selected example usages. Generally, in this context sentiment lexicons should be perceived as entities. The lexicon overlap score represents the similarity between two entities and therefore acts as a distance measure between lexicons.

For the demonstrated applications, sentiment lexicons for different domains from the internet platform reddit are used, which are available via the SocialSent project. Reddit is a website with a focus on content generated by individuals and user interaction. Topics and domains have separated spaces where interaction takes place. Based on these so-called subreddits 250 sentiment lexicons have been created and made available by the authors of the SocialSent project. Most of these lexicons contain between 4000 and 5000 words. [Ha16]

4.1 LOS Distributions

First, distributions of the lexicon overlap scores are shown. Pairwise distances between all lexicons were computed, employing all three LOS variations. Density was estimated by Kernel Density Estimation. [Si86] In Fig. 1 it is evident that the simple LOS yields higher values than the binary LOS and the binary LOS yields higher numbers than the general LOS overall. Furthermore, negative general LOS values can be observed. As

explained, this occurs if two lexicons contain several words that have opposite polarities assigned.

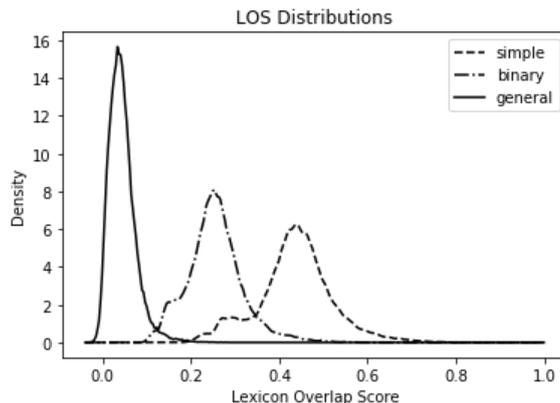


Fig. 1: Distributions of lexicon overlap scores

Selected comparisons between lexicons with all three lexicon overlap scores are displayed in Tab. 1. An extraordinary high similarity can be observed between *relationship_advice* and *relationships*, which is expectable due to strongly related topics. The second row shows a rather large common vocabulary (see simple LOS) for the subreddits *fantasyfootball* and *nba*, however the general LOS is close to zero, indicating that several common words have an opposite semantic orientation. The difference in semantic orientation between the lexicons *tipofmytongue* and *nba* is of such distinction that a negative general LOS is assigned. The last row shows two lexicons that have remarkably few words in common.

subreddit-specific lexicons		simple	binary	general
<i>relationship_advice</i>	<i>relationships</i>	0.901	0.708	0.342
<i>fantasyfootball</i>	<i>nba</i>	0.501	0.212	0.030
<i>tipofmytongue</i>	<i>nba</i>	0.343	0.145	-0.031
<i>CasualPokemonTrades</i>	<i>askscience</i>	0.190	0.102	0.008

Tab. 1: All lexicon overlap scores for selected comparisons.

4.2 Impact of lexicon enrichment via Distributional Thesaurus

Generally, the process of lexicon construction does not need to consist of only one step. Having multiple iterations for the creation of well performing lexicons is not uncommon. Kumar et al. [Ku16] used a Distributional Thesaurus (DT) to expand existing lexicons with semantically similar words and polarity scores. The lexicon overlap score can be

used to analyze the similarity of the treated lexicons before and after the process. This can answer the question, if the expansion using a Distributional Thesaurus leads to lexicons that become even more domain specific or if the enrichment process results in more similar lexicons.

To answer this question, the DT approach was applied to selected lexicons from the SocialSent project. For this purpose, the API of the JoBimText project [BR13], [G113] was accessed via the Java-based LexiExp Tool². Due to computational limitations, only a small subset of seven randomly selected lexicons was processed. This application example uses the binary LOS, since scores vary considerably between the SocialSent lexicons and the results produced by the DT approach. For each pair of lexicons, the binary LOS was calculated before and after DT expansion. In Fig. 2 both distributions are shown. The expanded lexicons are generally more similar. Analysis of the change in similarity per lexicon pair shows that the mean increase is 0.04, while the minimum is 0.02 and the maximum 0.05, expressing that the lexicon expansion increases lexicon similarity in all cases. This is supposedly since the DT expansion adds words that are semantically similar but not specifically from the same domain, therefore creating a larger common vocabulary.

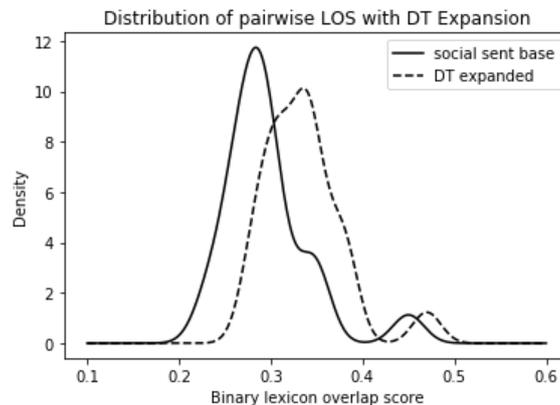


Fig. 2: Distribution of pairwise lexicon overlap scores before and after DT expansion

4.3 Lexicon clustering

Another possible application for the lexicon overlap score is the clustering of different sentiment lexicons based on their similarity. An affinity matrix can be constructed from the pairwise lexicon overlap scores which can be used as input for a clustering algorithm. The selection of the algorithm is restricted by the fact, that the lexicons have no values that indicate a location in an n-dimensional space (e.g. coordinates), which is

² <https://github.com/uhh-lt/LexiExp>

required by a fair amount of clustering approaches. In this case, Spectral Clustering [Vo07] was used.

The main parameter for Spectral Clustering is the number of clusters k . In the context of sentiment lexicons, it could be argued that the number of clusters should be set as the expected number of broader domains covered by the sentiment lexicons. This however is difficult to estimate. Therefore, the clustering algorithm was executed several times. It was found that with a high number of clusters in relation to the count of lexicons, groups of lexicons from similar domains were clustered.

Exemplarily results of the clustering ($k=100$) are shown in Tab. 2. Each cluster was manually annotated with a label which describes the commonalities of the contained lexicons. Only larger clusters are shown in the table. Lexicons in smaller clusters can be considered outliers in a sense that they do not belong in one category with other lexicons. Examples of these lexical outliers are *britishproblems*, *legaladvice*, *gameofthrones*, *nosleep*, *thewalkingdead*, *MakeupAddiction* and *BigBrother*. However, overall a negative silhouette coefficients [Ro87] suggests that clusters are overlapping and no clear distinction between lexicon clusters is possible.

Cluster Description	Contained lexicons
Trading in games	ACTrade, SVExchange, friendsafari
Technology	Android, buildapc, jailbreak, windowsphone
Family	AskWomen, Parenting, AskMen, BabyBumps
Competitive multiplayer games	CoDCompetitive, DotA2, GlobalOffensive, csgobetting, leagueoflegends, starcraft
IT jobs	programming, sysadmin, talesfromtechsupport
Fantasy multiplayer games	DestinyTheGame, Guildwars2, WildStar, archeage, elderscrollsonline, wow
Pokemon	CasualPokemonTrades, pokemon, pokemontrades, twitchplayspokemon
Religion	Christianity, DebateReligion, atheism, exmormon

Tab. 2: Clustered lexicons with human-annotated descriptions

Aside from the demonstrated examples, a range of further applications is possible. While this work focused on lexicons from different domains and constructed with the same algorithm, the comparison of lexicons from the same domain or dataset but constructed with different algorithms could yield interesting insights as well. The following chapter will discuss the possibilities of this approach in addition to the utility and expressiveness of the lexicon overlap score.

5 Discussion

The goal of the lexicon overlap score is to make sentiment lexicons comparable independent of context. The shown examples demonstrate different possibilities of applications. While the shown cases analyzed several lexicons constructed with the same approach, the lexicon overlap score can be used to compare lexicons from different approaches as well. As an example, it could be investigated if graph-based approaches yield lexicons that are more similar to each other compared to approaches based on lexical resources. Furthermore, it could be relevant to adjust the lexicon overlap score to take inflection into account when analyzing different lexicon construction methods.

As an incidental remark, a different approach to lexicon comparison could be based on the following. Lexicons can be converted to vectors, where each dimension corresponds to a word. Optionally dimensions could be reduced by application of a principal component analysis. The lexicons can then be compared by the cosine distance of the vectors. For the clustering approach as demonstrated in 4.3 this would offer coordinates instead of similarity scores, therefore enabling the usage of a wider range of clustering algorithms. However, a downside of this approach is the need to specify a mapping between words and dimensions, which reduces the ease by which new lexicons can be compared. The lexicon overlap score on the other hand side is not restricted in that way.

The sentiment lexicons used in this work were constructed in a way, that positive words have a value with a positive sign and negative words have a value with a negative sign. This does not need to be the case for every lexicon construction approach. The lexicon overlap score must be adjusted for lexicons where negative words carry a zero and positive words a one, and other cases alike. Furthermore, occasions of idioms, negations and multi word phrases have not been covered. Generally, the lexicon overlap score can be applied to sentiment lexicons which include said items as well. However, possibly it needs to be adjusted to correctly reflect cases where a word in the first lexicon occurs in a multi word phrase in the second lexicon and similar cases.

On an abstract level the lexicon overlap score can be adapted to any data structure which contains categorical items with assigned scores. Therefore, the application in different scenarios and research fields is possible as well. Generally, it should be noted that LOS values are best evaluated in comparison, since a single lexicon overlap score is challenging to interpret.

6 Conclusion

This work proposed a score for similarity measurement of sentiment lexicons for the task of sentiment classification. The three lexicon overlap scores have been implemented in python and made available as a pip package³ for easy installation and usage in python

³ <https://pypi.org/project/lexicon-overlap-score/>

projects. Several examples of applications have been shown to demonstrate the utility in different scenarios and underline the validity and significance.

The lexicon overlap score is independent of any domain, context and language and therefore can be easily applied. Aside from that, its independence from a specific set of classification items makes it faster to compute than performance benchmarks. In the end, it will help with objective quantification and more accurate descriptions of sentiment lexicons and therefore enables a better understanding of connections between lexicon construction methods as well as word polarity in different domains.

References

- [AA16] Adamov, A. Z.; Adali, E.: Opinion mining and Sentiment Analysis for contextual online-advertisement: Application of Information and Communication Technologies - AICT2016. 12-14 October 2016, Baku, Azerbaijan conference proceedings. IEEE, Piscataway, NJ, 2016; S. 1–3.
- [ACS08] Abbasi, A.; Chen, H.; Salem, A.: Sentiment analysis in multiple languages: Feature selection for opinion classification in web forums. In *ACM Transactions on Information Systems (TOIS)*, 2008, 26; S. 1–34.
- [BR13] Biemann, C.; Riedl, M.: Text: now in 2D! A framework for lexical expansion with contextual similarity. In *Journal of Language Modelling*, 2013, 1; S. 55.
- [GI13] Gliozzo, A. et al.: JoBimText Visualizer: A Graph-based Approach to Contextualizing Distributional Similarity: Proceedings of the 8th Workshop on TextGraphs in conjunction with EMNLP 2013, 2013; S. 6–10.
- [GST16] Gentzkow, M.; Shapiro, J.; Taddy, M.: *Measuring Polarization in High-Dimensional Data: Method and Application to Congressional Speech*. Working Papers, 2016.
- [Ha16] Hamilton, W. L. et al.: *Inducing Domain-Specific Sentiment Lexicons from Unlabeled Corpora*, 2016.
- [HM97] Hatzivassiloglou, V.; McKeown, K. R.: Predicting the semantic orientation of adjectives. In (Cohen, P. R.; Wahlster, W. Hrsg.): *Proceedings of the 35th annual meeting on Association for Computational Linguistics - Association for Computational Linguistics*, Morristown, NJ, USA, 1997; S. 174–181.
- [Hu07] Hunter, J. D.: *Matplotlib: A 2D Graphics Environment*. In *Computing in Science & Engineering*, 2007, 9; S. 90–95.
- [Jo18] Jochim, C. et al.: *SLIDE – a Sentiment Lexicon of Common Idioms*, 2018.
- [Ku16] Kumar, A. et al.: IIT-TUDA at SemEval-2016 Task 5: Beyond Sentiment Lexicon: Combining Domain Dependency and Distributional Semantics Features for Aspect Based Sentiment Analysis. In (Bethard, S. et al. Hrsg.): *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 2016; S. 1129–1135.

-
- [Mc10] McKinney, W.: Data Structures for Statistical Computing in Python: Proceedings of the 9th Python in Science Conference. SciPy, 2010; S. 56–61.
- [MT10] Mohammad, S.; Turney, P.: Emotions Evoked by Common Words and Phrases: Using Mechanical Turk to Create an Emotion Lexicon. In (Association for Computational Linguistics Hrsg.): Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text, 2010; S. 26–34.
- [Mu13] Muhammad, A. et al.: Contextual Sentiment Analysis in Social Media Using High-Coverage Lexicon. In (Bramer, M.; Petridis, M. Hrsg.): Research and Development in Intelligent Systems XXX. Incorporating Applications and Innovations in Intelligent Systems XXI Proceedings of AI-2013, The Thirty-third SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence. Springer International Publishing, Cham, s.l., 2013; S. 79–93.
- [NSW94] Nunberg, G.; Sag, I. A.; Wasow, T.: Idioms. In *Language*, 1994, 70; S. 491–538.
- [Pe11] Pedregosa, F. et al.: Scikit-learn: Machine Learning in Python. In *Journal of Machine Learning Research*, 2011, 12; S. 2825–2830.
- [PL08] Pang, B.; Lee, L.: Opinion Mining and Sentiment Analysis. In *Foundations and Trends® in Information Retrieval*, 2008, 2; S. 1–135.
- [PLV02] Pang, B.; Lee, L.; Vaithyanathan, S.: Thumbs up? Sentiment Classification using Machine Learning Techniques. In (Unknown Hrsg.): Proceedings of the ACL-02 conference on Empirical methods in natural language processing - EMNLP '02. Association for Computational Linguistics, Morristown, NJ, USA, 2002; S. 79–86.
- [Po10] Potts, C.: On the negativity of negation. In *Semantics and Linguistic Theory*, 2010, 20; S. 636.
- [Pr10] Princeton University: About WordNet. <https://wordnet.princeton.edu/>, 09.04.2020.
- [Qi10] Qiu, G. et al.: DASA: Dissatisfaction-oriented Advertising based on Sentiment Analysis. In *Expert Systems with Applications*, 2010, 37; S. 6182–6191.
- [Re05] Read, J.: Using Emoticons to Reduce Dependency in Machine Learning Techniques for Sentiment Classification: Proceedings of the ACL Student Research Workshop. Association for Computational Linguistics, USA, 2005; S. 43–48.
- [RFN17] Rosenthal, S.; Farra, N.; Nakov, P.: SemEval-2017 Task 4: Sentiment Analysis in Twitter. In (Bethard, S. et al. Hrsg.): Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017). Association for Computational Linguistics, Stroudsburg, PA, USA, 2017; S. 502–518.
- [Ro87] Rousseeuw, P. J.: Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. In *Journal of Computational and Applied Mathematics*, 1987, 20; S. 53–65.
- [Sa02] Sag, I. A. et al.: Multiword Expressions: A Pain in the Neck for NLP. In (Gelbukh, A. Hrsg.): Computational Linguistics and Intelligent Text Processing. Third International Conference, CICLing 2002 Mexico City, Mexico, February 17-23, 2002 Proceedings. Springer, Berlin, Heidelberg, 2002; S. 1–15.

- [Si86] Silverman, B. W.: Density estimation for statistics and data analysis. CRC press, 1986.
- [Ta14] Tang, D. et al.: Building large-scale twitter-specific sentiment lexicon: A representation learning approach: Proceedings of coling 2014, the 25th international conference on computational linguistics: Technical papers, 2014; S. 172–182.
- [TL02] Turney, P. D.; Littman, M. L.: Unsupervised Learning of Semantic Orientation from a Hundred-Billion-Word Corpus, 2002.
- [To01] Tong, R. M.: An Operational System for Detecting and Tracking Opinions in On-line Discussion: Proceedings of the Workshop on Operational Text Classification (OTC), 2001.
- [Tu02] Turney, P. D.: Thumbs up or thumbs down? In (Isabelle, P. Hrsg.): Proceedings of the 40th Annual Meeting on Association for Computational Linguistics - ACL '02. Association for Computational Linguistics, Morristown, NJ, USA, 2002; S. 417.
- [vCV11] van der Walt, S.; Colbert, S. C.; Varoquaux, G.: The NumPy Array: A Structure for Efficient Numerical Computation. In Computing in Science & Engineering, 2011, 13; S. 22–30.
- [Ve10] Velikovich, L. et al.: The Viability of Web-Derived Polarity Lexicons: Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics. Association for Computational Linguistics, USA, 2010; S. 777–785.
- [Vo07] Von Luxburg, U. Von Luxburg, U.: A Tutorial on Spectral Clustering, 2007.
- [VS16] Vishal, A.; Sonawane, S.: Sentiment Analysis of Twitter Data: A Survey of Techniques. In International Journal of Computer Applications, 2016, 139; S. 5–15.
- [WWH05] Wilson, T.; Wiebe, J.; Hoffmann, P.: Recognizing contextual polarity in phrase-level sentiment analysis: Proceedings of human language technology conference and conference on empirical methods in natural language processing, 2005; S. 347–354.
- [YS14] Yanqing Chen; Steven Skiena: Building Sentiment Lexicons for All Major Languages: ACL, 2014.
- [Zh11] Zhang, L. et al.: Combining Lexicon-based and Learning-based Methods for Twitter Sentiment Analysis, 2011.