

# Sentiment Composition of Words with Opposing Polarities

Svetlana Kiritchenko and Saif M. Mohammad

National Research Council Canada

{svetlana.kiritchenko, saif.mohammad}@nrc-cnrc.gc.ca

## Abstract

In this paper, we explore sentiment composition in phrases that have at least one positive and at least one negative word—phrases like *happy accident* and *best winter break*. We compiled a dataset of such opposing polarity phrases and manually annotated them with real-valued scores of sentiment association. Using this dataset, we analyze the linguistic patterns present in opposing polarity phrases. Finally, we apply several unsupervised and supervised techniques of sentiment composition to determine their efficacy on this dataset. Our best system, which incorporates information from the phrase’s constituents, their parts of speech, their sentiment association scores, and their embedding vectors, obtains an accuracy of over 80% on the opposing polarity phrases.

## 1 Introduction

The Principle of Compositionality states that the meaning of an expression is determined by the meaning of its constituents and by its grammatical structure (Montague, 1974). By extension, sentiment composition is the determining of sentiment of a multi-word linguistic unit, such as a phrase or a sentence, based on its constituents. In this work, we study sentiment composition in phrases that include at least one positive and at least one negative word—for example, phrases such as *happy accident*, *couldn’t stop smiling*, and *lazy sundays*. We refer to them as *opposing polarity phrases*. Such phrases present a particular challenge for automatic sentiment analysis systems that often rely on bag-of-word features.

Word–sentiment associations are commonly captured in *sentiment lexicons*. However, most existing manually created sentiment lexicons include only single words. Lexicons that include sentiment associations for *multi-word phrases* as well as their constituent words can be very useful in studying sentiment composition. We refer to them as *sentiment composition lexicons (SCLs)*.

We created a sentiment composition lexicon for opposing polarity phrases and their constituent words (Kiritchenko and Mohammad, 2016c).<sup>1</sup> Both phrases and single words were manually annotated with real-valued sentiment association scores using an annotation scheme known as Best–Worst Scaling.<sup>2</sup> We refer to the created resource as the *Sentiment Composition Lexicon for Opposing Polarity Phrases (SCL-OPP)*. The lexicon includes entries for 265 trigrams, 311 bigrams, and 602 unigrams.

In this paper, we use SCL-OPP to analyze regularities present in different kinds of opposing polarity phrases. We calculate the extent to which different part-of-speech combinations result in phrases of positive and negative polarity. We also show that for most phrases, knowing the parts of speech and polarities of their constituents is not enough to reliably predict the sentiment of the phrase.

We apply several unsupervised and supervised techniques of sentiment composition to determine their efficacy on predicting the sentiment of opposing polarity phrases. Our experiments indicate that

<sup>1</sup>[www.saifmohammad.com/WebPages/SCL.html#OPP](http://www.saifmohammad.com/WebPages/SCL.html#OPP)

<sup>2</sup>Best–Worst Scaling has been shown to produce reliable real-valued sentiment association scores (Kiritchenko and Mohammad, 2016a).

the sentiment of the last unigram or the sentiment of the most polar unigram in the phrase are not strong predictors of the overall sentiment of the phrase. Similarly, adjectives and verbs do not always dominate the sentiment in such phrases. Finally, we show that the constituent words, their parts of speech, their sentiment association scores, and their embedding vectors are all useful features—a supervised sentiment composition system that incorporates them obtains accuracies over 80% on both bigram and trigram opposing polarity phrases.

## 2 Related Work

A number of approaches have been proposed to address sentiment composition, which include manually derived syntactic rules (Moilanen and Pulman, 2007; Neviarouskaya et al., 2010), combination of hand-written rules and statistical learning (Choi and Cardie, 2008), and machine learning approaches (Nakagawa et al., 2010; Yessenalina and Cardie, 2011; Dong et al., 2015). Much work has been devoted to model the impact of negators and (to a lesser degree) intensifiers, words commonly referred to as contextual valence shifters, on sentiment of words they modify (Polanyi and Zaenen, 2004; Kennedy and Inkpen, 2005; Liu and Seneff, 2009; Wiegand et al., 2010; Taboada et al., 2011; Kiritchenko et al., 2014). Kiritchenko and Mohammad (2016b) created a sentiment composition lexicon for negators, modals, and adverbs (SCL-NMA) through manual annotation and analyzed the effect of these groups of modifiers on sentiment in short phrases. Recently, recursive deep model approaches have been proposed for handling sentiment of syntactic phrases through sentiment composition over parse trees (Socher et al., 2013; Zhu et al., 2014; Irsay and Cardie, 2014; Tai et al., 2015). In this work, we apply several unsupervised and supervised techniques of sentiment composition for a specific type of phrases—opposing polarity phrases.

## 3 Creating a Sentiment Lexicon for Opposing Polarity Phrases

This section summarizes how we created a sentiment composition lexicon for opposing polarity phrases using the Best–Worst Scaling annotation technique. For more details we refer the reader

Term	Sentiment score
best winter break	0.844
breaking free	0.172
isn't long enough	-0.188
breaking	-0.500
heart breaking moment	-0.797

**Table 1:** Example entries in SCL-OPP.

to (Kiritchenko and Mohammad, 2016c). Table 1 shows a few example entries from the lexicon.

**Term selection:** We polled the Twitter API (from 2013 to 2015) to collect about 11 million tweets that contain emoticons: ‘:)’ or ‘:(’. We will refer to this corpus as the *Emoticon Tweets Corpus*. From this corpus, we selected bigrams and trigrams that had at least one positive word and at least one negative word. The polarity labels (positive or negative) of the words were determined by simple look-up in existing sentiment lexicons: Hu and Liu lexicon (Hu and Liu, 2004), NRC Emotion lexicon (Mohammad and Turney, 2010; Mohammad and Turney, 2013), MPQA lexicon (Wilson et al., 2005), and NRC’s Twitter-specific lexicon (Kiritchenko et al., 2014; Mohammad et al., 2013).<sup>3</sup> In total, 576 opposing polarity  $n$ -grams (bigrams and trigrams) were selected. We also chose for annotation all unigrams that appeared in the selected set of bigrams and trigrams. There were 602 such unigrams. Note that even though the multi-word phrases and single-word terms were drawn from a corpus of tweets, most of the terms are used in everyday English.

**Best–Worst Scaling Method of Annotation:** Best–Worst Scaling (BWS), also sometimes referred to as Maximum Difference Scaling (MaxDiff), is an annotation scheme that exploits the comparative approach to annotation (Louviere and Woodworth, 1990; Louviere et al., 2015). Annotators are given four items (4-tuple) and asked which term is the Best (highest in terms of the property of interest) and which is the Worst (least in terms of the property of interest). Responses to the BWS questions can then be translated into real-valued scores through a simple counting procedure: For each term, its score is calculated as the percentage of times the term was

<sup>3</sup>If a word was marked with conflicting polarity in two lexicons, then that word was not considered as positive or negative.

SCP	Occ.	# phrases
▽adj. + △adj. → △phrase	0.76	17
▽adj. + △noun → ▽phrase	0.59	68
△adj. + ▽noun → ▽phrase	0.53	73
△adverb + ▽adj. → ▽phrase	0.89	18
△adverb + ▽verb → ▽phrase	0.91	11
▽noun + △noun → △phrase	0.60	10
△noun + ▽noun → ▽phrase	0.52	25
▽verb + det. + △noun → ▽phrase	0.65	17
▽verb + △noun → ▽phrase	0.82	17

**Table 2:** Sentiment composition patterns (SCPs) in SCL-OPP. △denotes a positive word or phrase, ▽denotes a negative word or phrase. ‘Occ.’ stands for occurrence rate of an SCP.

chosen as the Best minus the percentage of times the term was chosen as the Worst (Orme, 2009). The scores range from -1 to 1.

We employ Best–Worst Scaling for sentiment annotation by providing four (single-word or multi-word) terms at a time and asking which term is the most positive (or least negative) and which is the least positive (or most negative). Each question was answered by eight annotators through a crowdsourcing platform, CrowdFlower.<sup>4</sup> We refer to the resulting lexicon as the *Sentiment Composition Lexicon for Opposing Polarity Phrases (SCL-OPP)*.

Portions of the created lexicon have been used as development and evaluation sets in SemEval-2016 Task 7 ‘Determining Sentiment Intensity of English and Arabic Phrases’ (Kiritchenko et al., 2016).<sup>5</sup> The objective of that task was to test different methods of automatically predicting sentiment association scores for multi-word phrases.

## 4 Sentiment Composition Patterns

SCL-OPP allows us to explore sentiment composition patterns in opposing polarity phrases. We define a *Sentiment Composition Pattern (SCP)* as a rule that includes on the left-hand side the parts of speech (POS) and the sentiment associations of the constituent unigrams (in the order they appear in the phrase), and on the right-hand side the sentiment association of the phrase. Table 2 shows examples. SCPs that have a positive phrase on the right-hand side will be called *positive SCPs*, whereas SCPs that

have a negative phrase on the right-hand side will be called *negative SCPs*. Below are some questions regarding SCPs and opposing polarity phrases that we explore here:

- Which SCPs are common among opposing polarity phrases?
- With the same left-hand side of an SCP, how often is the composed phrase positive and how often is the composed phrase negative? For example, when negative adjectives combine with a positive noun, how often is the combined phrase negative?
- Are some parts of speech (of constituent words) more influential in determining the sentiment of a phrase than others?

To answer these questions, each of the entries in SCL-OPP is marked with the appropriate SCP. The part-of-speech sequence of a phrase is determined by looking up the most common part-of-speech sequence for that phrase in the Emoticon Tweets Corpus.<sup>6</sup> Next, for every left-hand side of an SCP, we determine the ratio of ‘how often occurrences of such combinations in SCL-OPP resulted in a positive phrase’ to ‘how often such combinations were seen in total’. We will refer to these scores as the *occurrence rates* (‘Occ.’) of *positive SCPs*. The *occurrence rates of negative SCPs* are calculated in a similar manner.

Table 2 presents all SCPs with the left-hand side combination appearing at least ten times in SCL-OPP, and whose occurrence rate is equal to or greater than 50%. For example, the second row tells us that there are 68 bigrams in SCL-OPP such that the first word is a negative adjective and the second word is a positive noun. Out of these 68 bigrams, 59% are negative, and the remaining 41% are positive, so the occurrence rate of this pattern is 0.59.

The most common SCPs in our lexicon are “△adj. + ▽noun → ▽phrase” (73) and “▽adj. + △noun → ▽phrase” (68). Observe that the occurrence rates of the patterns are spread over the entire range from 52% to 91%. Only two patterns have very high occurrence rates (around 90%): “△adverb + ▽adj. → ▽phrase” and “△adverb + ▽verb → ▽phrase”.

<sup>4</sup>Let *majority answer* refer to the option most chosen for a question. 81% of the responses matched the majority answer.

<sup>5</sup><http://alt.qcri.org/semeval2016/task7/>

<sup>6</sup>The corpus was automatically POS tagged using the CMU Tweet NLP tool (Gimpel et al., 2011).

Thus, for most opposing polarity phrases, their sentiment cannot be accurately determined based on the POS and sentiment of the constituents alone.

Both SCPs with high occurrence rates include adverbs that serve as intensifiers—words that increase or decrease the degree of association of the following word with positive (negative) sentiment (e.g., *incredibly slow*, *dearly missed*). Only the degree of association for the next word is changed while its polarity (positive or negative) is often preserved. Some adjectives can also play the role of an intensifier when combined with another adjective (e.g., *crazy talented*) or a noun (e.g., *epic fail*). For example, the adjective *great*, often considered highly positive, becomes an intensifier when combined with some nouns (e.g., *great loss*, *great capture*). Other adjectives determine the polarity of the entire phrase (e.g., *happy tears*, *bad luck*). Therefore, the occurrence rates of patterns like “ $\nabla$ adj. +  $\triangle$ noun  $\rightarrow$   $\nabla$ phrase” are low. Overall, even though adjectives and verbs are frequently the primary source of sentiment in the phrase, some nouns can override their sentiment as in *new crisis* or *leave a smile*. SCL-OPP includes phrases corresponding to many different kinds of sentiment composition patterns, and therefore, it is a useful resource for studying linguistic underpinnings of sentiment composition as well as for evaluating sentiment composition algorithms for opposing polarity phrases.

## 5 Automatically Predicting Sentiment

We now investigate whether accurate models of sentiment composition for opposing polarity phrases can be learned. We conduct experiments with several baseline unsupervised classifiers as well supervised techniques using features, such as unigrams, POS, sentiment scores, and word embeddings.

The problem of sentiment composition can be formulated in two different ways: a binary classification task where the system has to predict if the phrase is positive or negative; and a regression task where the system has to predict the real-valued sentiment association score of the phrase. We evaluate binary classification with simple accuracy (*acc.*) and the regression task with Pearson correlation coefficient (*r*). Learning and evaluation are performed separately for bigrams and trigrams.

### 5.1 Baseline Classifiers

The oracle ‘majority label’ baseline assigns to all instances the most frequent polarity label in the dataset. The ‘last unigram’ baseline returns the sentiment score (or the polarity label) of the last unigram in the phrase. For the regression task, we use the real-valued sentiment score of the unigram whereas for the binary classification task we use the polarity label (positive or negative). The ‘most polar unigram’ baseline assigns to the phrase the sentiment score (or the polarity label) of the most polar word in that phrase, i.e., the word with the highest absolute sentiment score. The ‘part-of-speech (POS) rule’ baseline assigns sentiment as follows:

1. If the phrase has an adjective, return the sentiment score (polarity) of the last adjective;
2. Else, if the phrase has a verb, return the sentiment score (polarity) of the last verb;
3. Else, return the sentiment score (polarity) of the most polar word.

### 5.2 Supervised Classifiers

We train a Support Vector Machines classifier with RBF kernel for the binary classification task and a Support Vector regression model with RBF kernel for the regression task using the LibSVM package (Chang and Lin, 2011). For both tasks, the models are trained using different combinations of the following features obtained from the target phrase: all unigrams, POS tag of each unigram, sentiment label of each unigram, sentiment score of each unigram, and the word embedding vector for each unigram. The word embeddings are obtained by running word2vec software (Mikolov et al., 2013) on the Emoticon Tweets Corpus. We use the skip-gram model with the default parameters and generate 200-dimensional vectors for each unigram present in the corpus. For each task, ten-fold cross-validation is repeated ten times, and the results are averaged.

### 5.3 Results

The results for all baseline and supervised methods are presented in Table 3. The ‘majority label’, ‘last unigram’, ‘most polar unigram’, and ‘POS rule’ baselines are shown in rows a to d. Observe that the sentiment association of the last unigram is not very

Features	Binary (Acc.)		Regression (Pearson $r$ )	
	2-gr	3-gr	2-gr	3-gr
<b>Baselines</b>				
a. majority label	56.6	60.8	-	-
b. last unigram	57.2	59.3	0.394	0.376
c. most polar unigram	66.9	69.8	0.416	0.551
d. POS rule	65.6	63.8	0.531	0.515
<b>Supervised classifiers</b>				
e. POS + sent. label	65.7	64.2	-	-
f. POS + sent. score	74.9	74.8	0.662	0.578
g. row f + uni	82.0	81.3	0.764	0.711
h. row f + emb(avg) + emb(max)	78.2	79.5	0.763	0.710
i. row f + emb(conc)	80.2	76.5	0.790	0.719
j. row f + emb(conc) + uni	<b>82.6</b>	<b>80.9</b>	<b>0.802</b>	<b>0.753</b>
k. POS + emb(conc) + uni	76.3	80.2	0.735	0.744

**Table 3:** Performance of the automatic systems on SCL-OPP. Features used: unigrams (uni), part-of-speech of a unigram (POS), sentiment binary label of a unigram (sent. label), sentiment real-valued score of a unigram (sent. score), embeddings (emb). ‘emb(conc)’ is the concatenation of the embedding vectors of the constituent unigrams; ‘emb(avg)’ is the average vector of the unigram embeddings; ‘emb(max)’ is maximal vector of the unigram embeddings.

predictive of the phrase’s sentiment (row b).<sup>7</sup> Both the ‘most polar unigram’ and the ‘POS rule’ classifiers perform markedly better than the majority baseline. Interestingly, the ‘most polar unigram’ classifier outperforms the slightly more sophisticated ‘POS rule’ approach on most tasks. Also, we found that within bigram phrases that contain adjectives or verbs, the adjective or verb constituents are the most polar words in only about half of the instances (and even less so in trigrams). This indicates that adjectives and verbs do not always dominate the sentiment in a phrase.

The results obtained using supervised techniques with various feature combinations are presented in rows e to k (Table 3). Using only POS and binary sentiment labels of the constituent unigrams, the supervised learning algorithm does not perform much better than our ‘POS rule’ baseline (the accuracies in row e are just slightly higher than those

<sup>7</sup>Note that the results for the ‘last unigram’ baseline are still better than the results of random guessing ( $\text{acc} = 50$ ,  $r = 0$ ). For the majority of  $n$ -grams in SCL-OPP, the polarity of the first unigram is opposite to the polarity of the last unigram. Thus, the results for a similar ‘first unigram’ baseline (not shown here) are worse than those obtained by the ‘last unigram’ baseline.

in row d). With access to real-valued sentiment scores of unigrams much more accurate models can be learned (row f). Furthermore, the results show that the sentiment of a phrase depends on its constituent words and not only on the sentiment of the constituents (row g shows markedly better performance than row f; all the differences are statistically significant,  $p < .01$ ). Concatenating word embeddings was found to be more effective than averaging. (Averaging is common when creating features for sentences). Having access to both unigrams and word embedding features produces the best results. (The differences between the scores in row i and row j are statistically significant,  $p < .01$ .) Row k shows results of the model trained without the gold sentiment scores of the unigrams. Observe that for bigrams, there is a substantial drop in performance compared to row j (6.3-point drop in accuracy on the binary task, 6.7-point drop in Pearson correlation on the regression task) whereas for trigrams the performance is not affected as much (less than 1-point change on both tasks). Thus, having access to sentiment scores of constituents is particularly useful for determining sentiment of bigram phrases.

## 6 Conclusions

We created a real-valued sentiment composition lexicon for opposing polarity phrases and their constituent words, through manual annotation. We analyzed patterns of sentiment composition across phrases formed with different POS combinations. Further, we applied several unsupervised and supervised techniques of sentiment composition to determine their efficacy on opposing polarity phrases. We showed that for most phrases the sentiment of the phrase cannot be reliably predicted only from the parts of speech and sentiment association of their constituent words, and that the constituent words, their parts of speech, their sentiment scores, and their embedding vectors are all useful features in supervised sentiment prediction on this dataset.

We intend to use SCL-OPP in the following applications: (1) to automatically create a large coverage sentiment lexicon of multi-word phrases and apply it in downstream applications such as sentence-level sentiment classification, and (2) to investigate how the human brain processes sentiment composition.

## References

- Chih-Chung Chang and Chih-Jen Lin. 2011. LIBSVM: A Library for Support Vector Machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27.
- Yejin Choi and Claire Cardie. 2008. Learning with compositional semantics as structural inference for subsentential sentiment analysis. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 793–801.
- Li Dong, Furu Wei, Shujie Liu, Ming Zhou, and Ke Xu. 2015. A statistical parsing framework for sentiment classification. *Computational Linguistics*.
- Kevin Gimpel, Nathan Schneider, Brendan O’Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, and Noah A. Smith. 2011. Part-of-speech tagging for Twitter: Annotation, features, and experiments. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 168–177, New York, NY, USA.
- Ozan Irsoy and Claire Cardie. 2014. Deep recursive neural networks for compositionality in language. In *Advances in Neural Information Processing Systems*, pages 2096–2104.
- Alistair Kennedy and Diana Inkpen. 2005. Sentiment classification of movie and product reviews using contextual valence shifters. In *Proceedings of the Workshop on the Analysis of Informal and Formal Information Exchange during Negotiations (FINEXIN)*, Ottawa, Ontario, Canada.
- Svetlana Kiritchenko and Saif M. Mohammad. 2016a. Capturing reliable fine-grained sentiment associations by crowdsourcing and best–worst scaling. In *Proceedings of The 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, San Diego, California.
- Svetlana Kiritchenko and Saif M. Mohammad. 2016b. The effect of negators, modals, and degree adverbs on sentiment composition. In *Proceedings of the Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis (WASSA)*.
- Svetlana Kiritchenko and Saif M. Mohammad. 2016c. Happy accident: A sentiment composition lexicon for opposing polarity phrases. In *Proceedings of 10th edition of the the Language Resources and Evaluation Conference (LREC)*, Portorož, Slovenia.
- Svetlana Kiritchenko, Xiaodan Zhu, and Saif M. Mohammad. 2014. Sentiment analysis of short informal texts. *Journal of Artificial Intelligence Research*, 50:723–762.
- Svetlana Kiritchenko, Saif M. Mohammad, and Mohammad Salameh. 2016. SemEval-2016 Task 7: Determining sentiment intensity of English and Arabic phrases. In *Proceedings of the International Workshop on Semantic Evaluation (SemEval)*, San Diego, California, June.
- Jingjing Liu and Stephanie Seneff. 2009. Review sentiment scoring via a parse-and-paraphrase paradigm. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 161–169.
- Jordan J. Louviere and George G. Woodworth. 1990. Best-worst analysis. Working Paper. Department of Marketing and Economic Analysis, University of Alberta.
- Jordan J. Louviere, Terry N. Flynn, and A. A. J. Marley. 2015. *Best-Worst Scaling: Theory, Methods and Applications*. Cambridge University Press.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *Proceedings of Workshop at ICLR*.
- Saif M. Mohammad and Peter D. Turney. 2010. Emotions evoked by common words and phrases: Using Mechanical Turk to create an emotion lexicon. In *Proceedings of the NAACL-HLT Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, LA, California.
- Saif M. Mohammad and Peter D. Turney. 2013. Crowdsourcing a word–emotion association lexicon. *Computational Intelligence*, 29(3):436–465.
- Saif M. Mohammad, Svetlana Kiritchenko, and Xiaodan Zhu. 2013. NRC-Canada: Building the state-of-the-art in sentiment analysis of tweets. In *Proceedings of the International Workshop on Semantic Evaluation*, Atlanta, Georgia.
- Karo Moilanen and Stephen Pulman. 2007. Sentiment composition. In *Proceedings of Recent Advances in Natural Language Processing (RANLP)*, volume 7, pages 378–382.
- Richard Montague. 1974. *Formal Philosophy; Selected papers of Richard Montague*. Yale University Press.
- Tetsuji Nakagawa, Kentaro Inui, and Sadao Kurohashi. 2010. Dependency tree-based sentiment classification using CRFs with hidden variables. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 786–794.
- Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. 2010. Recognition of affect, judgment,

- and appreciation in text. In *Proceedings of the International Conference on Computational Linguistics*, pages 806–814.
- Bryan Orme. 2009. Maxdiff analysis: Simple counting, individual-level logit, and HB. Sawtooth Software, Inc.
- Livia Polanyi and Annie Zaenen. 2004. Contextual valence shifters. In *Proceedings of the Exploring Attitude and Affect in Text: Theories and Applications (AAAI Spring Symposium Series)*.
- Richard Socher, Alex Perelygin, Jean Y. Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Seattle, USA.
- Maite Taboada, Julian Brooke, Milan Tofiloski, Kimberly Voll, and Manfred Stede. 2011. Lexicon-based methods for sentiment analysis. *Computational Linguistics*, 37(2):267–307.
- Kai Sheng Tai, Richard Socher, and Christopher D Manning. 2015. Improved semantic representations from tree-structured long short-term memory networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 1556–1566, Beijing, China.
- Michael Wiegand, Alexandra Balahur, Benjamin Roth, Dietrich Klakow, and Andrés Montoyo. 2010. A survey on the role of negation in sentiment analysis. In *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing (NeSp-NLP)*, pages 60–68, Stroudsburg, PA, USA.
- Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of the Joint Conference on HLT and EMNLP*, pages 347–354, Stroudsburg, PA, USA.
- Ainur Yessenalina and Claire Cardie. 2011. Compositional matrix-space models for sentiment analysis. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 172–182.
- Xiaodan Zhu, Hongyu Guo, Saif Mohammad, and Svetlana Kiritchenko. 2014. An empirical study on the effect of negation words on sentiment. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pages 304–313, Baltimore, Maryland, June.