

Introducing a Lexicon of Verbal Polarity Shifters for English

Marc Schulder^{*}, Michael Wiegand^{*}, Josef Ruppenhofer[†], Stephanie Köser^{*}

^{*} Spoken Language Systems, Saarland University, Germany

[†] Institute for German Language, Mannheim, Germany

marc.schulder@lsv.uni-saarland.de, michael.wiegand@lsv.uni-saarland.de,

ruppenhofer@ids-mannheim.de, skoester@coli.uni-saarland.de

Abstract

The sentiment polarity of a phrase does not only depend on the polarities of its words, but also on how these are affected by their context. Negation words (e.g. *not*, *no*, *never*) can change the polarity of a phrase. Similarly, verbs and other content words can also act as *polarity shifters* (e.g. *fail*, *deny*, *alleviate*). While individually more sparse, they are far more numerous. Among verbs alone, there are more than 1200 shifters. However, sentiment analysis systems barely consider polarity shifters other than negation words. A major reason for this is the scarcity of lexicons and corpora that provide information on them. We introduce a lexicon of verbal polarity shifters that covers the entirety of verbs found in WordNet. We provide a fine-grained annotation of individual word senses, as well as information for each verbal shifter on the syntactic scopes that it can affect.

Keywords: Sentiment Analysis, Polarity Shifter, Negation, Lexical Semantics, Lexicon

1. Introduction

Polarity shifters are content words that exhibit semantic properties similar to negation. For example, the negated statement in (1) can also be achieved by the verbal shifter *fail* instead of the negation *not*, as shown in (2).

- (1) Peter did **not** pass the exam.
- (2) Peter **failed**_{shifter} to pass the exam.

As with negation words, polarity shifters change the polarity of a statement. This can happen to both positive and negative statements. In (3) the positive polarity of *scholarship* is shifted by *denied*, resulting in a negative polarity for the phrase. Conversely, the overall polarity of (4) is positive despite the negative polarity of *pain*.

- (3) She was [**denied**_{shifter} the [scholarship]⁺]⁻.
- (4) The new treatment has [**alleviated**_{shifter} her [pain]⁻]⁺.

Polarity shifting is also caused by other content word classes, such as nouns (e.g. *downfall*) and adjectives (e.g. *devoid*). However, this work focusses on verbs, due to their importance as minimal semantic units, far-reaching scopes and potential basis for nominal shifter lexicons (see §2.2.). Knowledge of polarity shifting is important for a variety of tasks, especially sentiment analysis (Wiegand et al., 2010; Liu, 2012; Wilson et al., 2005), as well as relation extraction (Sanchez-Graillet and Poesio, 2007) and textual entailment recognition (Harabagiu et al., 2006).

The majority of research into polarity shifting for sentiment analysis has focussed on negation words (Wiegand et al., 2010; Schouten and Frasincar, 2016; Pak and Paroubek, 2010). Negation words (e.g. *not*, *no*, *never*) are mostly function words, of which only a small number exists, so exhaustive coverage is comparatively simple. Content word classes, such as verbs, are considerably more difficult to cover comprehensively due to their sheer number. For example, WordNet (Miller et al., 1990) contains over 10k verbal lemmas. Most verbs are also far less frequent than

common negation words, making individual verbal shifters seem less important. However, overall, verbal shifter lemmas occur 2.6 times as often as negation words (see §4.).

Most existing resources on negation and polarity shifting cover few to no instances of verbal shifters (see §2.3.). To remedy this, we introduce a **complete lexicon of verbal shifters** with annotations of polarity shifters and their shifting scope for each word sense.

Our **contributions** are as follows:

- (i) A complete lexicon of verbal polarity shifters, covering all verbs found in WordNet 3.1.
- (ii) A fine grained annotation, labelling every sense of a verb separately.
- (iii) Annotations for shifter scope, indicating which parts of a sentence are affected by the shifting.

The entire dataset is **publicly available**.¹

2. Background

In this section we will provide a formal definition of polarity shifters (§2.1.), motivate our focus on verbal shifters (§2.2.) and discuss related work (§2.3.).

2.1. Polarity Shifters

The notion of valence or polarity shifting was brought to broad awareness in the research community by the work of Polanyi and Zaenen (2006). Those authors drew attention to the fact that the basic valence of individual lexical items may be shifted in context due to (a) the presence of certain other lexical items, (b) the genre type and discourse structure of the text and (c) cultural factors. In subsequent research, the term *shifter* has since mostly been applied to the case of lexical items that influence polarity. Further, the notion of *shifting* is most prototypically used for situations where a discrete polarity *switch* occurs between the classes *positive*, *negative* and *neutral*. However, for other authors,

¹<https://github.com/marcschulder/lrec2018>

including Polanyi and Zaenen (2006), intensification (e.g. *very disappointing*) and downtoning (e.g. *somewhat disappointing*) of polar intensity also falls within the scope of shifting. We partially follow this view in that we consider downtoning to be shifting, as it moves the polarity of a word in the opposite direction, i.e. making a positive expression less positive (e.g. *hardly satisfying*) and a negative one less negative (e.g. *slightly problematic*). We do not consider intensifiers as shifters, as they support the already existing polarity.

In most research, shifters are commonly illustrated and enumerated rather than formally defined. Polanyi and Zaenen (2006) for instance list negation words, intensifiers, modals and presuppositional items as lexical contextual polarity shifters.

Setting aside downtoners for now, the common denominator of shifting is negation. Negation marks contexts in which a situation that the speaker expected fails to occur or hold. When this situation is part of a binary opposition (*dead – alive*), one can firmly conclude that the complementary state of affairs holds (*not dead* \Rightarrow *alive*). In cases where the negation affects a scalar notion, which is common in evaluative contexts, the understanding that arises depends on which kinds of scalar inferences and default assumptions are made in the context (Paradis and Willners, 2006). Thus, *not good* denies the applicability of an evaluation in the region of good or better, but leaves open just how far in the direction of badness the actual interpretation lies: “*It wasn’t good*” may be continued with “*but it was ok*” to yield a neutral or mildly positive evaluation or with “*in fact, it was terrible*” to yield a strongly negative one.² While downtoners (e.g. *somewhat*) applied to scalar predicates such as *good* do not directly express contradiction, they do give rise to negative entailments and inferences. Moreover, the structure of scales intrinsically provides shifting. Thus, while something being *good* allows it to be even more positive (“*The movie was good. In fact, it was excellent.*”), something being *somewhat good* bounds its positiveness and opens up more negative meanings (“*The performance was somewhat good, but overall rather disappointing*”). Considering these properties of scales, one can see shifting at work even in the case of downtoning.

2.2. Verbal Shifters

While the inclusion of shifting and scalar semantics in semantic representations is not limited to lexical items of particular parts-of-speech – we also find shifter adjectives (e.g. *devoid*) and adverbs (e.g. *barely*) – we limit our work to verbal shifters for several reasons. As shown by the work of Schneider et al. (2016), verbs, together with nouns, are the most important minimal semantic units in text and thus are prime candidates for being tackled first. Verbs are usually the main syntactic predicates of clauses and sentences and thus verbal shifters can be expected to project far-reaching

²Note that the example also illustrates how distinguishing between items that induce a switch between polarities and others that affect intensity without changing overall polarity is an idealization. Simple syntactic negation of a polar adjective may influence intensity as well as polarity (e.g. *not terrible* \neq *excellent*) (Kiritchenko and Mohammad, 2016).

scopes. Most nominal shifters (e.g. *failure, loss*), on the other hand, have morphologically related verbs (e.g. *fail, lose*) and we expect that this connection can be exploited to spread shifter classification from verbs to nouns in the future. Related to this, the grammar of verbs, for instance with respect to the diversity of scope types, is more complex than that of nouns and so we expect it to be easier to project from verbs to nouns rather than in the opposite direction.

2.3. Related Work

Existing lexicons and corpora that cover polarity shifting focus almost exclusively on negation words. The most complex negation lexicon for sentiment analysis (Wilson et al., 2005) includes a mere 12 verbal shifters. In contrast, our resource covers over 1200 verbal shifter lemmas.

Corpora used as training data for negation processing, such as the Sentiment Treebank (Socher et al., 2013) or the BioScope corpus (Szarvas et al., 2008), are fairly small datasets, so only the most frequent negation words appear. The BioScope corpus, for example, contains only 6 verbal shifters (Morante, 2010). Schulder et al. (2017) show that state-of-the-art systems trained on such data do not reliably detect polarity shifting and should profit from explicit knowledge of verbal shifters.

The only work to date that covers a larger number of verbal shifters is Schulder et al. (2017), who annotate a sample of the English verbs found in WordNet for whether they exhibit polarity shifting. They start by manually annotating an initial 2000 verbs. These verbs are used to train an SVM classifier using linguistic features and common language resources. The classifier is then run on the remaining WordNet verbs to bootstrap a list of additional likely shifters. This list is then checked by a human annotator to detect false positives. Combining the initial annotation and the result of the bootstrapping process, they create a list of 3043 verbs.

While the lexicon by Schulder et al. (2017) is an important step towards full coverage of verbal polarity shifters, there are several aspects that we seek to improve upon. First of all, their lexicon covers less than a third of the verbs found in WordNet, likely missing a number of verbal shifters. Schulder et al. (2017) argue that their bootstrap process should cover the majority of shifters, however, this would mean that only 9% of all verbs are shifters.³ Their initial annotation of 2000 randomly selected verbs puts the shifter ratio at 15% instead.

Another issue with their lexicon is that it only labels lemma forms, but does not differentiate between word senses. Many verbs do not actually exhibit shifting in all of their senses, so this information will be important for contextual classification.

Lastly, they forgo the question of shifter scope, i.e. which argument of a verb can be affected by its polarity shift.

³They find 980 shifter lemmas among the 3043 verbs that are annotated by a human annotator. An additional 7538 verbs are assumed to not be shifters without human confirmation.

3. Data

We treat this annotation effort as a binary labelling task where a word can either cause polarities to shift or not. However, instead of assigning a single label to an entire verb lemma, as Schulder et al. (2017) did, we label individual word senses. We outline the rationale for this in §3.1. In addition we explicitly specify the syntactic scope of the shifting. This is motivated and explained in §3.2. §3.3. describes the annotation process. §3.4. describes the data format of our main lexicon. Based on this main lexicon we also derive two auxiliary lexicons in §3.5., providing complete labelled lists of all WordNet verb lemmas and all WordNet verb synsets respectively.

3.1. Word Senses

Many words that shift polarities only do so for some of their word senses. For example, *mark down* acts as a shifter in (5), where it has the sense of “*reducing the value of something*”, but the sense of “*writing something down to have a record of it*” in (6) causes no shifting. In our work we found that among shifter lemmas with multiple word senses, only 23% caused shifting in each of their senses. An annotation on the basis of individual word senses is therefore required.

- (5) The agency [**marked down**_{shifter} [their assets]⁺]⁻.
 (6) She [**marked down**_{no shifter} [his confession of guilt]⁻]⁻.

To differentiate the senses of a verb, we use its synset affiliations found in WordNet. Words within the same synset share a shifter label. Shifter scope, on the other hand, can differ among words of the same synset (see §3.2.). The annotation introduced in §3.3. is therefore applied to individual lemma-sense pairs to capture the best of both worlds.

3.2. Shifter Scope

A verbal shifter usually only affects the parts of a sentence that are syntactically governed by the verb through its valency. However, not every argument of a verbal shifter is subject to polarity shifting. Which argument is affected by polarity shifting depends on the verb in question. In (7), *surrender* shifts only the polarity of its subject, but does not affect the object. Conversely, *defeat* only shifts its object in (8). The polarity of the subject of *defeat* does not play a role in this, as can be seen in (9).

- (7) [[The villain]⁻ **surrendered**]⁺ [to the hero]⁺.
 (8) [The villain]⁻ [**defeated** [the hero]⁺]⁻.
 (9) Chance [**defeated** [the hero]⁺]⁻.

In the following, we present the shifting scopes we observed, their abbreviation in the annotation and examples for each:

Subject (subj): The verbal shifter affects its subject, e.g. “[*His confidence*]_{subj}⁺ [*decreased*]⁻.”

Direct Object (dobj): The verbal shifter affects its direct object, e.g. “[*The storm*]_{subj} [*ruined* [*their party*]_{dobj}⁺]⁻.”

Prepositional Object (pobj_*): The verbal shifter affects the object within a prepositional phrase. The preposition in question is included in the annotation. For example, *shield* as in “[*The wall*]_{subj} [*shielded* *them*]_{dobj} [*from the explosion*]_{pobj}⁻]⁺.” is annotated as *pobj_from* and *reimburse* as in “[*The company*]_{subj} [*reimbursed* *him*]_{dobj} [*for his expenses*]_{pobj}⁻]⁺.” as *pobj_for*.

Clausal complement (comp): The verbal shifter affects a clausal complement, such as infinitive clauses like “*He [failed [to pass the exam]]_{comp}⁺*.” or gerunds like “*She [stopped [using drugs]]_{comp}⁻*.”

The given scopes assume that verb phrases are in their active form. In passive phrases, subject and object roles are inverted. To avoid this issue, sentence structure normalization should be performed before computing shifter scope. Synsets in WordNet only capture the semantic similarity of words, but almost no syntactic properties (Ruppenhofer and Brandes, 2015). The shifter scope of a verb depends on its syntactic arguments, which can differ between verbs of the same synset. For example, *discard* and *dispose* share the sense “*throw or cast away*”, but while *discard* shifts its direct object (10), *dispose* requires a prepositional object (11). For this reason we annotate lemma-synset pairs individually, instead of assigning scope labels to an entire synset.

- (10) He [**discarded** [the evidence]_{dobj}⁺]⁻.
 (11) He [**disposed** [of the evidence]_{pobj}⁺]⁻.

We also consider cases where a verbal shifter has more than one potential scope for the same lemma-sense pair. For example, *infringe* can shift its direct object or various prepositional objects, as seen in (12) – (14). Therefore, *infringe* receives the scope labels *dobj*, *pobj_on* and *pobj_upon*.

- (12) The inquiry [**infringes** [people’s privacy]_{dobj}⁺]⁻.
 (13) The inquiry [**infringes** [on people’s privacy]_{pobj}⁺]⁻.
 (14) The inquiry [**infringes** [upon people’s privacy]_{pobj}⁺]⁻.

A verbal shifter will only ever shift the polarity of one of its scopes. Which scope is affected by the shifting depends on the given sentence.

3.3. Annotation

The entire dataset was labelled by an expert annotator with experience in linguistics and annotation work. To measure inter-annotator agreement, a second annotator re-annotated 400 word senses for their shifter label. They achieved an agreement of $\kappa = 0.73$, indicating substantial agreement (Landis and Koch, 1977).

The annotation progressed as follows: Given a complete list of WordNet verb lemmas, the annotator would inspect one lemma at a time. For this lemma, all senses were looked up. For each such lemma-sense pair, the annotator decided whether it is a shifter or not. Decisions were based on the sense definition of the synset and whether sentences using

	Shifters		Non-shifters		Total
	#	%	#	%	#
Lemmas	1220	11.53	9357	88.47	10577
Synsets	924	6.88	12502	93.12	13426
LS Pairs	2131	8.88	21855	91.12	23986

Table 1: The ratio of verbal shifters in WordNet. Lemmas are counted as shifters when at least one sense is a shifter. “LS Pairs” represents lemma-synset pairs.

this sense of the lemma cause shifting. If a word sense was labelled as a shifter, it was subsequently also annotated for its potential shifter scopes.

In cases where label conflicts between different lemma-sense pairs of the same sense were encountered, these labels were reconsidered. This introduced an additional robustness to the annotation as it let the annotator revisit challenging cases from a new perspective.

The resulting list of lemma-sense pairs provides more fine-grained information than either an annotation for only word lemmas or only synsets could (see §3.1. and §3.2.).

3.4. Main Lexicon File Format

We provide our main lexicon as a comma-separated value (csv) file in which each line represents a specific lemma-sense-scope triple of a verbal shifter. Each line follows the format “LEMMA, SYNSET, SCOPE”. The fields are defined as follows:

LEMMA: The lemma form of the verb.

SYNSET: The numeric identifier of the synset, commonly referred to as *offset* or *database location*. It consists of 8 digits, including leading zeroes (e.g. 00334568).

SCOPE: The scope of the shifting. Given as `subj` for subject position, `dobj` for direct object position and `comp` for clausal complements. Prepositional object positions are given as `pobj_*`, where `*` is replaced by the preposition in question, e.g. `pobj_from` for objects with the preposition “from” or `prep_of` for the preposition “of”.

When a lemma has multiple word senses, a separate entry is provided for each lemma-sense pair. When a lemma-sense pair has multiple potential shifting scopes, a separate entry is provided for each scope. Any combinations not provided are considered not to exhibit shifting. Take, for example, the set of entries for “blow out”:

```
(15) blow out,00436247,subj
      blow out,02767855,dobj
```

It tells us that *blow out* in the sense 00436247 (“melt, break, or become otherwise unusable”) is a shifter that affects its subject. The sense 02767855 (“put out, as of fires, flames, or lights”) also exhibits shifting, but this time affects the direct object. It is, however, not a shifter for sense 02766970 (“erupt in an uncontrolled manner”). For an example of multiple scopes for the same word sense, consider *cramp*:

```
(16) cramp,00237139,dobj
      cramp,00237139,pobj_in
```

	subj	dobj	pobj_*	comp	Total
Frequency	402	1574	212	32	2220

Table 2: Distribution of shifting scopes for individual word senses. Total is higher than number of lemma-synset pairs (Table 1) as 4% of shifters have multiple potential scopes.

Its sense 00237139 (“prevent the progress or free movement of”) can shift the polarity of either its direct object (e.g. “it cramped his progress”) or that of a prepositional object with the preposition “in” (e.g. “he was cramped in his progress”). The three other senses of *cramp* given by WordNet are not considered shifters.

3.5. Auxiliary Lexicons

Our main lexicon is labelled at the lemma-sense pair level to provide the most fine-grained level of information possible. It can, however, easily applied to more coarse-grained applications. As a convenience, we provide lemma- and synset-level auxiliary lexicons that list all WordNet lemmas and all WordNet synsets, respectively, accompanied with their shifter label. A lemma is labelled as a shifter if at least one of its senses is considered a shifter in our main lexicon. Similarly, synsets are labelled as shifters if at least one of its lemma-realizations is a shifter.

4. Statistics

In Table 1 we present the ratio of shifters among the verbs contained in WordNet. While only about 10% of verbs are shifters, this still results in 1220 lemmas and 924 synsets, more than covered in any other resource (see §2.3.).

49% of verbs in WordNet are polysemous, i.e. they have multiple meanings. Among verbal shifters, this ratio is considerably higher, reaching 73%. Of these, only 23% are shifters in all of their word senses.

To get an idea of how common verbal shifters are in actual use, we computed lemma frequencies over the *Amazon Product Review Data* corpus (Jindal and Liu, 2008), which comprises over 5.8 million reviews. We found this corpus suitable due to its size, sentiment-related content and use in related tasks (Schulder et al., 2017).

We observe 1163 different verbal shifter lemmas with an overall total of 34 million occurrences. Correcting for non-shifter senses of shifter lemmas⁴, we still estimate 13 million occurrences, accounting for 5% of all verb occurrences in the corpus. To compare, the 15 negation words found in the valence shifter lexicon by Wilson et al. (2005) occur 13 million times as well. While the frequency of individual negation (function) words is unsurprisingly higher, the total number of verbal shifter occurrences highlights that verbal shifters are just as frequent and should not be ignored.

Statistics on the distribution of shifter scopes can be found in Table 2. 74% of verbal shifters have a direct object

⁴Due to the lack of robust word-sense disambiguation tools, we estimated the likelihood of a lemma instance being a shifter based on its ratio of shifter word senses. A lemma with 3 shifter senses and 1 non-shifter sense would, therefore, have a likelihood of 0.75 to be a shifter.

scope and 10% a prepositional object scope. Among these, “from” is the most common preposition at 51%, followed by “of” with 22%. 19% shift the polarity of their subject and only 1.5% shift that of a clausal complement. This distribution shows that shifting cannot be trivially assumed to always affect the direct object and that explicit knowledge of shifter scopes will be useful for judging the polarity of a phrase.

5. Conclusion

We introduced a lexicon of verbal polarity shifters that covers the entire verb vocabulary of WordNet. Our annotation labels each individual word sense of a verb, providing more fine-grained information than annotations on the lemma-level would. In addition, we also label the syntactic scopes of each verbal shifter that can be affected by the shifting.

This is a clear improvement over the list of verbal shifters provided by Schulder et al. (2017), which only provides labels at the lemma-level rather than for individual word senses and gives no information regarding shifting scope. It also only has human expert annotation for 30% of the verb vocabulary of WordNet, as opposed to our full coverage.

We hope this resource will help improve fine-grained sentiment analysis systems by providing explicit information on where polarities may shift in a sentence.

We also hope our work will encourage the creation of similar polarity shifter lexicons for nouns and adjectives. As they are more numerous than verbs (WordNet contains 20k adjectival and 110k nominal lemmas), creating such resources will come with its own challenges, especially in the case of nouns.

Acknowledgements

The authors were partially supported by the German Research Foundation (DFG) under grants RU 1873/2-1 and WI 4204/2-1.

6. Bibliographical References

- Harabagiu, S., Hickl, A., and Lacatusu, F. (2006). Negation, Contrast and Contradiction in Text Processing. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.
- Jindal, N. and Liu, B. (2008). Opinion Spam and Analysis. In *Proceedings of the International Conference on Web Search and Data Mining*.
- Kiritchenko, S. and Mohammad, S. M. (2016). The Effect of Negators, Modals, and Degree Adverbs on Sentiment Composition. In *Proceedings of WASSA*.
- Landis, J. R. and Koch, G. G. (1977). The Measurement of Observer Agreement for Categorical Data. *Biometrics*, 33(1):159–174.
- Liu, B. (2012). Sentiment Analysis and Opinion Mining. *Synthesis Lectures on Human Language Technologies*, 5(1):1–167.
- Miller, G., Beckwith, R., Fellbaum, C., Gross, D., and Miller, K. (1990). Introduction to WordNet: An On-line Lexical Database. *International Journal of Lexicography*, 3:235–244.
- Morante, R. (2010). Descriptive Analysis of Negation Cues in Biomedical Texts. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC)*.
- Pak, A. and Paroubek, P. (2010). Twitter as a corpus for sentiment analysis and opinion mining. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC)*.
- Paradis, C. and Willners, C. (2006). Antonymy and Negation: The Boundedness Hypothesis. *Journal of Pragmatics*, 38(7):1051–1080.
- Polanyi, L. and Zaenen, A. (2006). Contextual Valence Shifters. In James G. Shanahan, et al., editors, *Computing Attitude and Affect in Text: Theory and Applications*, pages 1–10. Springer Netherlands.
- Ruppenhofer, J. and Brandes, J. (2015). Extending effect annotation with lexical decomposition. In *Proceedings of the Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA@EMNLP)*, pages 67–76, Lisboa, Portugal.
- Sanchez-Graillet, O. and Poesio, M. (2007). Negation of protein–protein interactions: analysis and extraction. *Bioinformatics*, 23(13):i424–i432.
- Schneider, N., Hovy, D., Johannsen, A., and Carpuat, M. (2016). SemEval-2016 Task 10: Detecting Minimal Semantic Units and their Meanings (DiMSUM). In *Proceedings of the International Workshop on Semantic Evaluation (SemEval@NAACL-HLT)*, pages 546–559.
- Schouten, K. and Frasincar, F. (2016). Survey on Aspect-Level Sentiment Analysis. *IEEE Transactions on Knowledge and Data Engineering*, 28(3):813–830.
- Schulder, M., Wiegand, M., Ruppenhofer, J., and Roth, B. (2017). Towards Bootstrapping a Polarity Shifter Lexicon using Linguistic Features. In *Proceedings of the International Joint Conference on Natural Language Processing (IJCNLP)*, Taipei, Taiwan.
- Socher, R., Perelygin, A., Wu, J. Y., Chuang, J., Manning, C. D., Ng, A. Y., and Potts, C. (2013). Recursive Deep Models for Semantic Compositionality over a Sentiment Treebank. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Szarvas, G., Vincze, V., Farkas, R., and Csirik, J. (2008). The BioScope Corpus: Annotation for Negation, Uncertainty and their Scope in Biomedical Texts. In *Proceedings of the Workshop on Current Trends in Biomedical Natural Language Processing (BioNLP@ACL-HLT)*.
- Wiegand, M., Balahur, A., Roth, B., Klakow, D., and Montoyo, A. (2010). A Survey on the Role of Negation in Sentiment Analysis. In *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*.
- Wilson, T., Wiebe, J., and Hoffmann, P. (2005). Recognizing Contextual Polarity in Phrase-level Sentiment Analysis. In *Proceedings of the Joint Conferences on Human Language Technology and on Empirical Methods in Natural Language Processing (HLT/EMNLP)*.