

Predictive Features for Detecting Indefinite Polar Sentences

Michael Wiegand, Dietrich Klakow

Spoken Language Systems, Saarland University
D-66123 Saarbrücken, Germany

Michael.Wiegand@lsv.uni-saarland.de, Dietrich.Klakow@lsv.uni-saarland.de

Abstract

In recent years, text classification in sentiment analysis has mostly focused on two types of classification, the distinction between objective and subjective text, i.e. *subjectivity detection*, and the distinction between positive and negative subjective text, i.e. *polarity classification*. So far, there has been little work examining the distinction between *definite* polar subjectivity and *indefinite* polar subjectivity. While the former are utterances which can be categorized as either positive or negative, the latter cannot be categorized as either of these two categories. This paper presents a small set of domain independent features to detect indefinite polar sentences. The features reflect the linguistic structure underlying these types of utterances. We give evidence for the effectiveness of these features by incorporating them into an unsupervised rule-based classifier for sentence-level analysis and compare its performance with supervised machine learning classifiers, i.e. Support Vector Machines (SVMs) and Nearest Neighbor Classifier (kNN). The data used for the experiments are web-reviews collected from three different domains.

1. Introduction

There has been a growing interest in the analysis of opinionated texts in natural language processing within the last few years. As far as text classification is concerned most research has focused on two types of classification, the distinction between objective and subjective text, i.e. *subjectivity detection* (Riloff and Wiebe, 2003; Wiebe et al., 2004; Dias et al., 2009), and the distinction between positive and negative subjective text, i.e. *polarity classification* (Pang et al., 2002; Turney, 2002; Wilson et al., 2005; Ng et al., 2006). So far, there has been little work examining the distinction between definite polar subjectivity and indefinite polar subjectivity.

Sentences (1) and (2) are definite polar utterances since these sentences can be categorized as either positive or negative:

- (1) She's always the best of the best!
- (2) That product is so bad, it should be illegal.

Sentences (3) - (5) are examples of indefinite polar utterances:

- (3) That first record was amazing but then they fell off really fast.
- (4) She has an average voice.
- (5) I'm not hellishly impressed.

These utterances have in common that they are subjective and express a value judgment. None of these statements can be categorized as definite positive or negative. The indefiniteness is achieved either by stating both positive and negative aspects (Sentence (3)), by using polar expressions not denoting definite polarity (*average* in Sentence (4)) or by diminishing/negating definite polar phrases (Sentence (5)). This paper presents a small set of domain independent features to detect indefinite polar sentences. The features reflect the linguistic structure underlying these types of utterances. Since indefinite utterances or even entire indefinite

reviews are part of a realistic review collection, those features might be helpful for an accurate text classification.

We give evidence for the effectiveness of these features by incorporating them into an unsupervised rule-based classifier for sentence-level analysis and compare its performance with supervised machine learning classifiers. We restrict ourselves to sentence-level analysis since we are primarily interested in basic utterances for which sentences are a suitable approximation.

2. Related Work

Koppel and Schler (2006) present a machine learning approach to polarity classification where also reviews with indefinite polarity are considered. A binary classifier for positive and negative polarity is learned using bag-of-words features. Reviews being predicted with a low confidence are classified as indefinite polar reviews. The paper does not address features specifically designed for detecting indefinite polar reviews.

Zhao et al. (2008) consider a CRF-based model for sentence-level polarity classification of reviews also taking into consideration indefinite polar sentences as a separate class. Again, there is no discussion about what predictive features are for this class.

Wilson et al. (2005) present polarity classification of news text on phrase level. Apart from positive and negative polar phrases, phrases with both polarities and neutral polarity are considered. However, our task differs greatly from theirs. Wilson et al. (2005) carry out classification of phrases whereas this work deals with sentence-level classification. Moreover, this paper addresses another text type being online reviews whereas Wilson et al. (2005) deal with news texts. As all four polar classes are classified within the same classifier, it is not clear which features are predictive for the indefinite polar classes.

Wilson et al. (2004) present features for distinguishing strong from weak opinion clauses. Weak opinion clauses bear some resemblance to the class of indefinite polar expressions. However, the paper does not address polarity. Moreover, the same differences as the one mentioned

Domain	Number of Sentences
Person	1914
Sports	980
Food	1618

Table 1: Size of the different datasets.

to Wilson et al. (2005) (i.e. level of granularity and text type) also apply to Wilson et al. (2004).

3. Data

We extracted a set of reviews from *Rate-It-All*¹ from the domains *person*, *sports*, and *food*. Since we want to classify sentences, we restricted our choice to reviews which only comprise one sentence. For definite polar utterances, we extracted reviews rated with 1 or 5 stars and for indefinite reviews, we extracted reviews rated with 3 stars. Of the latter subset, some reviews were manually removed, since they were deemed definite polar utterances. As it is fairly difficult to have a realistic estimate of what the underlying class distribution is, we generated a balanced dataset via random-sampling. Table 1 lists the size of the resulting datasets.

4. Feature Set

Table 2 lists all the features that we use. The feature set can be divided into the subset indicating indefinite polarity and the subset indicating definite polarity. We will discuss each of these features individually in the forthcoming subsections. Several of the features require the knowledge of polar expressions (e.g. *PosInPast* or *PolarSuper*). For their detection we use the *Subjectivity Lexicon* (Wilson et al., 2005) from the MPQA project. This lexicon is well suited for our experiments since it contains a binary intensity feature dividing entries into *weak* polar expressions (e.g. *valid* or *bulky*) and *strong* polar expressions (e.g. *wonderful* or *hideous*). We make use of this distinction in one of our features (*NegStrongPol*). In order to increase the coverage of the polarity lexicon, we add adjectives from the *Macquarie Semantic Orientation Lexicon* (Mohammad et al., 2009)². All these entries are categorized as *weak* polar expressions.

4.1. Indefinite Polarity Features

The following subsections describe features indicative of indefinite polar opinions.

4.1.1. Concessive Conjunctions (*ConcConj*)

In Section 1., we pointed out that one way of expressing indefinite polarity is to state both a positive and a negative opinion in a sentence. An intuitive heuristic to look for utterances in which both positive and negative polar expressions occur is not very effective. We ascribe it to the fact that the detection of polar opinions is very error prone. The pertaining polar expressions may not be detected if they are not included in the polarity lexicon, and even if they can be

detected, their contextual polarity may be computed incorrectly. Contextual polarity comprises many linguistic phenomena, such as *negation* or *irony*, which are difficult to model computationally.

We found, however, that there is another feature which most often co-occurs with this type of utterance. Concessive conjunctions, such as *but* or *although*, indicate that two clauses represent semantically opposed propositions. In our dataset this is usually a juxtaposition of two polar opinions. Thus, such a conjunction is also indicative of a sentence with an overall indefinite polarity:

(6) A nice⁺ wine, *but* definitely [not worth]⁻ the price.

4.1.2. Concessive Conjunctions Preceded by a Polar Expression (*ConcAndPolar*)

Even though concessive conjunctions may be detected more easily than two contrasting polar opinions, the concessive conjunction may itself be an ambiguous word. For instance, *but* in the following sentence is not a concessive conjunction:

(7) They are nothing *but* an untalented stain on the music world ... totally atrocious music.

We found, however, that a co-occurrence of a polar expression preceding the potential concessive conjunction is a fairly reliable way of disambiguating these words.

4.1.3. Detensifiers (*Detens*)

Another way of expressing indefinite polarity is to diminish polar phrases. Therefore, a further cue may be diminishing expressions, or so-called *detensifiers*, such as *almost*, *slightly* or *less*:

(8) Terry is *almost* as good as Robert Jordan, his stories are *slightly less* word encompassing.

For detensifiers, we mainly adhere to the list presented in Jason (1988).

4.1.4. Negated Strong Polar Expressions (*NegStrongPol*)

In traditional polarity classification negated polar expressions are interpreted as if the polarity of the polar expressions were reversed (Kennedy and Inkpen, 2005; Klenner et al., 2009). We argue that for the detection of indefinite polarity negated polar expressions should not be equated with unnegated polar expressions with the opposite polarity. Instead, they should be treated as a separate category. In particular, negated strong polar expressions (Sentence (9)) may similarly convey indefinite polarity as detensified polar expressions (Sentence (10)):

(9) They are *not bad*.

(10) They are *quite good*.

4.1.5. Negation Expressions (*NegExp*)

NegStrongPol is a fairly complex feature in which several properties may have to co-occur, i.e. the sentence must contain a polar expression which has to be of strong intensity and it has to be within the scope of a negation. The computation of such a feature is error-prone as the negation

¹<http://www.rateitall.com>

²We found that other entries are too noisy for our application.

Feature	Abbreviation	Indef nite Polarity Feature	Def nite Polarity Feature	Example(s)
Concessive Conjunctions	<i>ConcConj</i>	✓		but, although, however
Concessive Conjunctions Preceded by a Polar Expression	<i>ConcAndPolar</i>	✓		he is nice but . . .
Detensifiers	<i>Detens</i>	✓		rather, kind of, slightly, almost
Negated Strong Polar Expressions	<i>NegStrongPol</i>	✓		not excellent, not bad
Negation Expressions	<i>NegExp</i>	✓		not, never, nothing
Middle-of-the-road Polar Expressions	<i>MiddleExp</i>	✓		solid, average, ordinary
Positive Polar Expressions in Past Tense Clause	<i>PosInPast</i>	✓		he used to be funny
Polar Superlatives	<i>PolarSuper</i>		✓	best, funniest, worst
Emphatic Cues	<i>EmphCues</i>		✓	yeah, ah, grrreeaaaat, !

Table 2: Description of the feature set.

may not be correctly computed or the strong polar expression may be overlooked as it is not specified in the polarity lexicon. Therefore, we add a feature just recognizing negations. Admittedly, this feature is not equivalent to the previous feature but its computation should be much more reliable and, often, it should coincide with *NegStrongPol*.

4.1.6. Middle-of-the-road Polar Expressions (*MiddleExp*)

Indefinite polarity may not only be conveyed by the use of certain linguistic constructions, be it on discourse level (*ConcConj*) or on syntax level (*Detens* or *NegStrongPol*). It can also be lexically realized by so-called *middle-of-the-road polar expressions*, such as *ok*:

- (11) This beer brand is *ok* ... really far away of the Paulaner Heffeweissen.

We compiled a list of such expressions by starting with a couple of manually defined seed words which were expanded using semantic resources, such as WordNet (Miller et al., 1990). Moreover, we also manually selected a subset of *weak* polar expressions from the *Subjectivity Lexicon*. Note that middle-of-the-road polar expressions differ quite substantially from the polar expressions marked as *both* (e.g. think, believe) or *neutral* (e.g. demand, brag) in that lexicon, though the category names may suggest otherwise. *MiddleExp* always implies a value judgment whereas the two categories in the *Subjectivity Lexicon* usually do not have that property. Besides, these two types of expressions did not show any noticeable predictiveness on our datasets.

4.1.7. Positive Polar Expressions in Past Tense Clause (*PosInPast*)

We observed that in many indefinite polar reviews, people tend to recall positive aspects concerning the topic they review which they experienced in the past and contrast them with negative aspects they presently perceive. We found that this behavioural pattern can be automatically identified by detecting a positive polar expression uttered in a past tense clause. Reviews are usually written in present tense and we found that if a clause occurs in past tense, then this will most often be accompanied by a switch in tense:

- (12) [I *used_{Past}* to *like⁺* those chips a lot *better⁺* some years ago], now the only way I eat them is with sour cream.

4.2. Definite Polarity Features

The following subsections describe features indicative of definite polar opinions.

4.2.1. Polar Superlatives (*PolarSuper*)

Definite polar opinions may often be conveyed by a polar superlative:

- (13) He's the *best* actor.

Intuitively, the polar intensity of a polar superlative (e.g. *best*) is stronger than the intensity of a polar positive (e.g. *good*) or comparative (e.g. *better*). Though polar superlatives are similar to strong polar expressions, such as *excellent*, or intensified polar expressions, such as *very good*, we found in our initial experiments that they are far less predictive for our task than the polar superlative.

4.2.2. Emphatic Cues (*EmphCues*)

Often, emphatic cues, such as interjections (*yeah, ah* etc.), co-occur with definite polar sentences. A feature detecting such cues may help since in our dataset there are many definite polar sentences in which – apart from the emphatic cue – there is no other feature that could be that easily computed. For instance, in the following sentence the polar opinion is pragmatic, i.e. it is not lexicalized. However, there are three exclamation marks whose occurrence is interpreted as an emphatic cue:

- (14) I can eat this peanut butter on anything!!!

For the implementation of this feature, we mainly relied on exclamation marks and the part-of-speech tag indicating interjections, i.e. *UH*. In addition, we formulated regular expressions capturing irregular spelling as in *suuuper* or *grrreeaaaat*.

5. Rule-based Classifier

The features from Section 4. can be used as a rule-based classifier. For each test instance, the occurrences of features indicating definite and indefinite polar utterances are

counted. We assign the instance the class with the majority of feature occurrences. In case of ties the instance is classified as definite polar since we have fewer features formulated for that class.

6. Experiments

Table 3 displays the individual performance of the different features used as a rule-based classifier (as formulated in Section 5). We test for each feature whether it is significantly different from a random baseline. We report statistical significance on the basis of a χ^2 test.

Each of the features is at least significantly better than the baseline when the entire dataset is considered. It is very striking that among the best performing features are *ConcConj* and *NegExp* which are features describing different types of closed-word classes. Their advantage is that they comprise words frequently occurring across all domains.

The features that fail to be significantly better than the baseline on each domain, i.e. *PolarSuper*, *NegStrongPol*, and *PosInPast*, are more complex than most of the other better performing features. They all describe a co-occurrence of separate properties, e.g. *PosInPast* is a polar expression that also happens to be positive and occurs in a past tense clause. We assume that the reason for these features performing less well lies in the sparsity of their occurrence.

Table 4 compares the performance of the unsupervised rule-based classifier using all features with supervised classifiers on 10-fold crossvalidation. We compare Support Vector Machines (SVMs) using *SVMLight*³ and a k Nearest Neighbor Classifier (kNN) using *TiMBL*⁴. For *SVMLight* we use the standard configuration and for *TiMBL* we use the 5 nearest neighbors. This setting produces the best overall performance on all domains. All words contained in the training sets are used as features for the supervised classifiers. Following the insights of Pang et al. (2002), features indicate presence within an instance and not its frequency. The inclusion of our novel high-level features (Table 2) did not improve performance of these classifiers when they were added to the bag of words. For the rule-based classifier, we also considered subsets of the features, but no significant improvement towards the entire feature set could be achieved. SVMs achieve best performance. Both kNN and the rule-based classifier are significantly worse than SVMs. Surprisingly, the rule-based classifier is as robust as kNN. There is no significant difference between the rule-based classifier and kNN⁵.

Figure 1 shows the average performance of the different classifiers with varying amounts of labeled training data. For each configuration, we randomly sampled n training instances from the domain corpus and use the remaining instances as test data. We sampled 20 times and report the averaged result. Even for SVMs, it takes more than 400 labeled data instances to achieve a significantly better accuracy than the unsupervised rule-based classifier. For less

Type	Person	Sports	Food	Average
Rule-based	76.18	78.06	77.32	77.19
kNN	78.00	77.55	75.59	77.05
SVMs	81.19	81.02	80.22	80.81

Table 4: Comparison of accuracy of the different classifiers.




Figure 1: Average accuracy of the different classifiers using different amounts of labeled training data.

robust supervised classifiers, such as kNN, more than 800 labeled data instances are required to achieve the same performance as the rule-based classifier.

7. Error Analysis and Future Work

We manually inspected the data instances being incorrectly classified by the rule-based classifiers. Thus, we hope to get an idea of what the shortcomings of the proposed feature set are. Unfortunately, we found no systematic error that could be solved by adding another linguistic feature.

Our results suffer from the low writing quality of many reviews that is common for this type of web data. Various spelling mistakes and grammatical errors have a significant impact on the quality of our feature extraction. They often cause features not to be detected, be it due to the fact that words are assigned incorrect part-of-speech tags or they are misspelled and cannot be matched with the pertaining entries in the lexicons we use.

We also found that among the definite polar sentences which were automatically extracted from 1 and 5 star reviews were also some actual indefinite polar reviews. Thus, manually filtering these data might result in a better dataset. Some of our features depend on lexical resources, such as *MiddleExp* or *NegStrongPol*. The resources, we use are domain independent but usually have a lower coverage on informal texts, such as web-reviews that are used in this work, than on formal texts, such as news documents that are predominant in research in natural language processing. Thus, adapting these resources to that register might also result in an improvement of our proposed rule-based approach.

³<http://svmlight.joachims.org>

⁴<http://ilk.uvt.nl/timbl>

⁵Statistical significance is again reported on the basis of a χ^2 test with significance level $p < 0.001$.

Type	Person	Sports	Food	All
<i>ConcConj</i>	72.99***	71.53***	73.24***	72.76***
<i>ConcAndPolar</i>	65.94***	62.76***	66.25***	65.36***
<i>NegExp</i>	58.99***	60.92***	61.37***	60.26***
<i>EmphCues</i>	59.98***	57.86***	60.88***	59.84***
<i>MiddleExp</i>	59.14***	58.06***	59.77***	59.13***
<i>Detens</i>	55.28**	54.90*	55.56**	55.30***
<i>PolarSuper</i>	52.46	57.65***	53.58*	54.56***
<i>NegStrongPol</i>	52.72	54.08	54.39*	53.73***
<i>PosInPast</i>	53.29*	52.65	50.74	52.23*

Table 3: Accuracy of the different features on the different domains. Statistical significance is reported on the basis of a χ^2 test with significance levels $p < 0.05$ (*), $p < 0.01$ (**) and $p < 0.001$ (***)

8. Conclusion

In this paper, we presented a set of discriminative features for the detection of indefinite polar sentences. We showed that these features can be used as an unsupervised rule-based classifier which provides as good performance as supervised machine learning classifiers, such as kNN trained on bag-of-words. When only small amounts of training data are available (i.e. less than 300 sentences), the unsupervised approach even outperforms more robust supervised classifiers, such as SVMs. Since the feature set uses domain-independent features the classifier works equally well throughout different domains.

Acknowledgements

Michael Wiegand was funded by the German research council DFG through the International Research Training Group “IRTG” between Saarland University and University of Edinburgh.

9. References

- Gaël Dias, Dinko Dimchev Lambov, and Veska Noncheva. 2009. High-level Features for Learning Subjective Language across Domains. In *Proceedings of the Third International AAAI Conference on Weblogs and Social Media (ICWSM)*.
- Gary Jason. 1988. Hedging as a Fallacy of Language. *Informal Logic*, 10(3):169–175.
- Alistair Kennedy and Diana Inkpen. 2005. Sentiment Classification of Movie Reviews Using Contextual Valence Shifters. In *Workshop on the Analysis of Formal and Informal Information Exchange during Negotiations*, volume 22.
- Manfred Klenner, Stefanos Petrakis, and Angela Fahrni. 2009. Robust Compositional Polarity Classification. In *Proceedings of the Conference on Recent Advances in Natural Language Processing (RANLP)*, Borovets, Bulgaria.
- Moshe Koppel and Jonathan Schler. 2006. The Importance of Neutral Examples for Learning Sentiment. *Computational Intelligence*, 22(2):100–109.
- George Miller, Richard Beckwith, Christiane Fellbaum, Derek Gross, and Katherine Miller. 1990. Introduction to Wordnet: An On-line Lexical Database. *International Journal of Lexicography*, 3:235–244.
- Saif Mohammad, Cody Dunne, and Bonnie Dorr. 2009. Generating High-Coverage Semantic Orientation Lexicons from Overtly Marked Words and a Thesaurus. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Vincent Ng, Sajib Dasgupta, and S. M. Niaz Arifin. 2006. Examining the Role of Linguistic Knowledge Sources in the Automatic Identification and Classification of Reviews. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 611–618.
- Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment Classification Using Machine Learning Techniques. In *Proceedings of Empirical Methods in Natural Language Processing (EMNLP)*, pages 79–86.
- Ellen Riloff and Janyce Wiebe. 2003. Learning extraction patterns for recognizing subjective expressions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Peter Turney. 2002. Thumbs up or Thumbs down?: Semantic Orientation Applied to Unsupervised Classification of Reviews. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 417–424, Philadelphia, Pennsylvania.
- Janyce Wiebe, Theresa Wilson, Rebecca Bruce, Matthew Bell, and Melanie Martin. 2004. Learning Subjective Language. *Computational Linguistics*, 30(3).
- Theresa Wilson, Janyce Wiebe, and Rebecca Hwa. 2004. Just how mad are you? Finding strong and weak opinion clauses. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI)*.
- Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing Contextual Polarity in Phrase-level Sentiment Analysis. In *Proceedings of Human Language Technologies Conference/Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP)*, Vancouver, Canada.
- Jun Zhao, Kang Liu, and Gen Wang. 2008. Adding Redundant Features for CRFs-based Sentence Sentiment Classification. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.