

Negation Modeling for German Polarity Classification

Michael Wiegand and Maximilian Wolf

Spoken Language Systems
Saarland University

D-66123 Saarbrücken, Germany

michael.wiegand@lsv.uni-saarland.de

mwolf@coli.uni-saarland.de

Josef Ruppenhofer

Institute for German Language
D-68161 Mannheim, Germany

ruppenhofer@ids-mannheim.de

Abstract

We present an approach for modeling German negation in open-domain fine-grained sentiment analysis. Unlike most previous work in sentiment analysis, we assume that negation can be conveyed by many lexical units (and not only common negation words) and that different negation words have different scopes.

Our approach is examined on a new dataset comprising sentences with mentions of polar expressions and various negation words. We identify different types of negation words that have the same scopes. We show that already negation modeling based on these types largely outperforms traditional negation models which assume the same scope for all negation words and which employ a window-based scope detection rather than a scope detection based on syntactic information.

1 Introduction

Negation is one of the most central linguistic phenomena. Therefore, negation modeling is essential to various common tasks in natural language processing, such as relation extraction (Sanchez-Graillet and Poesio, 2007), recognition of textual entailment (Harabagiu et al., 2006) and particularly sentiment analysis (Wiegand et al., 2010). In the latter task, negation typically inverts the polarity of polar expressions. For example, in (1), the negated positive polar expression *like* conveys negative polarity.

- (1) I do [not [like]⁺]⁻ this new Nokia model.

While most research on negation has been carried out on English language data, little research has looked into the behaviour of negation in German. This is surprising since German negation

is even harder to handle than English negation. For example, since German displays a more flexible word order than English, the German negation word *nicht* (*not*) may appear both left (2) or right (3) of a polar expression it modifies. In English, however, there is a strong tendency of a negation word to precede the polar expression it negates (1).

- (2) Der Kuchen ist [nicht [köstlich]⁺]⁻.
(The cake is not delicious.)
(3) Ich [[mag]⁺ den Kuchen nicht]⁻.
(I do not like the cake.)

To make the task even more difficult, there are not only function words, such as the particle *nicht* (*not*), to express negation but also content words, such as verbs (4), nouns (5) or adjectives (6). (2)-(6) also show that these different negation word types have different scopes.

- (4) [[Dieses Bemühen]⁺ scheiterte^{verb}]⁻.
(These efforts failed.)
(5) [Das Scheitern^{noun} [dieser Bemühungen]⁺]⁻ war vorhersehbar.
(The failure of these efforts was foreseeable.)
(6) Angesichts [dieser gescheiterten^{adj} [Bemühungen]⁺]⁻ ist nun ein Umdenken erforderlich.
(These failed efforts now require a change of thinking.)

In this paper, we follow a **rule-based approach** to negation modeling for fine-grained sentiment analysis that largely draws information from lexicons. We focus on the **task** of identifying the scope of negation words with regard to polarity classification. In other words, given a mention of a negation word and a polar expression, we want to automatically determine whether the negation word negates the polar expression.

We do *not* claim to have full knowledge of all German negation words. (Given that content words can perform implicit negation, we assume the overall vocabulary of negation words to be fairly large.) Instead, we propose a typology of negation words and assign a characteristic scope to each type. Therefore, we provide a *formalism* that is able to compute the respective scope of every possible negation word, once the negation word has been assigned to its respective type.

Our approach heavily relies on syntactic knowledge, particularly information contained in a dependency parse. We demonstrate that the analyses that state-of-the-art parsers produce for German are insufficient for our task and require further normalization.

The **contributions** of this paper are:

- We present the first comprehensive study on German negation modeling for fine-grained sentiment analysis.
- Instead of having one generic scope for all types of negation words, we formulate different types of scopes for different types of negation words.
- We substantially go beyond negation (function) words, that is, we also consider negation verbs, nouns and adjectives.
- We introduce a new dataset¹ comprising German sentences in which negation words are manually annotated with respect to the polar expressions they negate.
- We publicly release a tool¹ for fine-grained German sentiment analysis that implements our proposed approach.

2 Data & Annotation

In order to evaluate negation in context, we built a small focused dataset comprising sentences with negation. In order to keep the annotation effort manageable, we extracted those sentences in which a negated polar expression is likely. We therefore extracted from a corpus only those sentences in which both some negation word co-occurs with at least one polar expression according to the **sentiment lexicon** of the PolArt-system (Klenner et al., 2009). In order not to bias the scope of the negation in those sentences we did not impose any restriction regarding the relation between negation words and polar expressions. In order to recognize negation words, we also created a **negation lexicon**. For that we used several resources. On the one hand, we used all negation expressions from the PolArt system. In addition, we translated a large list of English negation verbs to German and also manually added morphologically related nouns and adjectives if existent, e.g. for the verb *stagnieren* (*stagnate*) we would also add the noun *Stagnation* (*stagnation*).

¹Available under <https://github.com/artificial-max/polcla>

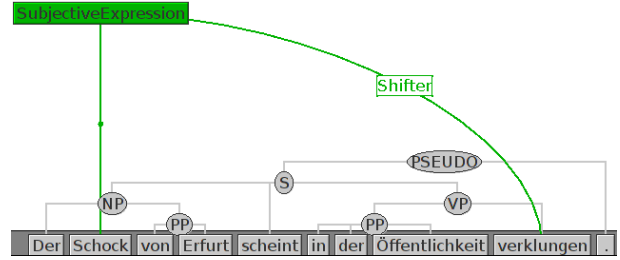


Figure 1: Example sentence annotation from dataset (translation: *The shock of Erfurt seems to have faded away in the public*). Polar expressions (e.g. *Schock*) evoke a frame *SubjectiveExpression*; the word that negates a polar expression (e.g. *verklungen*) is assigned the frame element *Shifter* of that frame.

In total, we sampled 500 sentences from the *DeWaC*-corpus (Baroni et al., 2009). We manually annotated *every* polar expression in those sentences. (Note that we did not only annotate those polar expressions we could automatically identify with the help of the PolArt sentiment lexicon.) We also marked every negation word in case it negates a polar expression. The dataset comes in *TIGER/SALSA format* (Erk and Padó, 2004). Figure 1 illustrates the annotation of our dataset. Polar expressions evoke a frame *SubjectiveExpression*. If a polar expression is negated, then its negation word is labeled as a frame element *Shifter* of its frame.²

Of the 500 sentences, we removed 67 sentences which contained obvious errors (i.e. misspellings, grammatical mistakes or incorrect sentence boundaries). We excluded those sentences since the methods we are going to examine rely on a correct syntactic parse. Erroneous sentences are likely to produce spurious syntactic analyses.

On a sample of 200 sentences, we measured an interannotation agreement of $\kappa = 0.87$ which can be considered substantial (Landis and Koch, 1977).

Table 1 provides some statistics of our dataset. Even though every sentence contains at least one polar expression (in most cases there is more than

²In our annotation, we refer to negation words as *shifters* since, in computational linguistics, this is the preferred term. A shifter need not fully invert the polarity of a polar opinion but just *shifts* it into the opposite direction (Polanyi and Zaenen, 2006). For example, *faded optimism* does not mean that *optimism* is completely absent but it means that the current amount has substantially decreased – which can be interpreted as a negative opinion.

Property	Freq
number of sentences	433
number of polar expressions	979
number of sentences with negated polar expression	282
number of negation words <i>left</i> of polar expression	142
number of negation words <i>right</i> of polar expression	140

Table 1: Statistics of negation detection dataset.

one) and a negation word, there are only 282 cases in which a polar expression is within the scope of the negation word, i.e. it is actually negated. This shows that it is not trivial to determine whether a polar expression has been negated. It is also worth pointing out that a negation is as likely to precede the polar expression it negates as to follow it.

3 Methods

3.1 Baseline I: Window-based Scope

Our first baseline applies a simple window-based approach for the scope detection of negation. It is inspired by various works from polarity classification on English language data (Wilson et al., 2005; Wiegand et al., 2013). One considers as scope a span of n tokens around the negation word. While on English data it typically suffices to scan only the tokens succeeding the negation word, on German data we check three different window types: one that assumes the polar expression to succeed the negation word, one that assumes the polar expression to precede the negation word and one in which both directions are examined.

Figure 2 shows the performance of those different window-based scopes on our dataset. It shows that for German negation one needs to look into both directions. (This is in line with our statistic from Table 1.) All of the three window types have their maximum at $n=4$. In our forthcoming experiments, we use the best window-based scope (i.e. using both directions at $n=4$) as a baseline.

3.2 Baseline II: Clause-based Scope

Our second baseline models the scope on the basis of syntactic information. Instead of using a window of fixed size, we scan all words in the same clause in which the negation word occurs for a polar expression. Typically the scope of a negation never exceeds clause boundaries. For example, in (7) the negation word *niemand* (*nobody*) does not negate the polar expression *entsetzlich* (*appalling*) in the subordinate clause. From a linguistic perspective, this scope is more adequate




Figure 2: Illustration of window-based scope using different window sizes.

than the window-based approach.

- (7) [Niemand wird etwas zu dem Ereignis sagen wollen]_{main clause},
[weil es sich dabei um eine [entsetzliche] Angelegenheit
handelt]_{subordinate clause}.
(Nobody will want comment on this incident, since it is an appalling
affair.)

3.3 Our Approach

Our approach is fundamentally different to the previous baselines in the sense that we define individual scopes for different types of words. Our framework allows arbitrary scopes to be defined for every possible negation word. A scope is defined in terms of a grammatical relation. For instance, we could formulate that the *subject* of the negation word *aufhören* (*subside*) is the expression that is negated as in (8).

- (8) [[Die Schmerzen] - hören auf]_{verb}⁺.
(The pain subsides.)

We do not have the knowledge to explicitly enumerate the scope for every possible negation word from our negation lexicon (§2). Instead, we grouped words with similar scope characteristics (§3.3.1-§3.3.2) and assigned one scope which satisfies the entire group of words.

Our framework allows the specification of a **priority scope list** for a negation word, i.e. a list with more than one argument position (see also Table 2). We process such a list from left to right and apply the first argument that matches for the specific negation word in some sentence. The advantage of such a list is that there are negation words that may negate different arguments. The flexibility we gain with priority lists is essential for identifying the correct scope of negation content words as we will explain in §3.3.2.

We do not claim that our proposed approach perfectly models the scope of every German negation word. But we show that with relatively little lexical knowledge we can largely outperform a traditional approach that treats all negation words in the same way. Therefore, our proposed method should be regarded as a strong baseline for future research.

Table 2 summarizes the different negation word types that we discuss in detail below.

3.3.1 Scope for Negation Function Words


The type of lexical units most commonly associated with negation are negation particles, such *nicht* (*not*), negation adverbs, such as *niemals* (*never*), indefinite pronouns, such as *kein* (*no*), and a few prepositions, such as *ohne* (*without*). Even though these negation words only constitute a handful of lexical units, they are known to have a large impact. This is due to the fact that these words are function words which entails that they occur frequently. We call these words **negation function words**. Regarding the scope of those words, we distinguish between three types.

Negation Adverbs and Indefinite Pronouns.

These negation words exhibit similar behaviour (in terms of scope) as sentential adverbs. As a consequence, these negation words have a wide scope. It is the entire clause in which they are embedded (9) & (10). We use the same definition as we applied for our baseline in §3.2.

- (9) [Noch **nie** wollte Kiew [Frieden]⁺]⁻.
(Kiev never wanted peace.)
- (10) [**Kein** Mensch möchte sie dabei [unterstützen]⁺]⁻.
(No one wants to support them with that.)

Negation Particle. The particle *nicht* (*not*) has a narrow scope. We only include the word which it *governs* in the dependency graph (11).



Negation Prepositions. Negation prepositions also have a narrow scope. However, unlike the negation particle, their scope does not include the words which they govern but which are their dependents (hence the reverse relation), e.g. *Hass* (*hatred*) in (12).

- (12) Wir bauen eine Welt ganz [**ohne**^{prep} [Hass]^{dependent}]⁺.
(We create a world without hatred.)

3.3.2 Scope for Negation Content Words

In the following, we describe the remaining words, all of which are content words. We therefore refer to these words as **negation content words**.

Negation Nouns. Negation nouns typically reverse the polarity of two types of dependents, either a genitive modifier (13) or a prepositional object (14). Note that we leave the preposition underspecified so that it can match any potential preposition.

- (13) Das Gericht beschloss [die **Aufhebung**^{noun} [der Strafe]^{gmod}]⁺.
(The court decided to lift the sentence.)
- (14) Qi Gong dient auch zur [**Vorbeugung**^{noun} [vor Krankheiten]^{objp-vor}]⁺.
(Qi gong is also used for preventing diseases.)

Negation Adjectives. There are two different major constructions in which adjectives may occur. Adjectives may be used *predicatively* or *attributively*. Negation adjectives may occur in both constructions. Therefore, polar expressions negated by an adjective may be in two different argument positions, namely a noun in subjective position in the predicative case (15) or a noun that is modified by an attributive adjective (16).

- (15) [[Diese Bemühungen]^{subj} sind **gescheitert**^{pred-adj}]⁻.
(These efforts failed.)
- (16) Das sind alles [**korrigierbare**^{attr-adj} [Fehler]^{attr-rev}]⁺.
(These are recoverable errors.)

Negation Verbs. For this study, we distinguish between two major types of verb groups, *transitive verbs* and *intransitive verbs*. In the case of transitive negation verbs, it is the object that is negated (17), while for intransitive verbs, it is the subject that is negated (18).

- (17) Dieses Medikament [**lindert**^{transitive-verb} [die Schmerzen]^{obja}]⁺.
(This drug cures the pain.)
- (18) [[Die Schmerzen]^{subj} **hören auf**^{intransitive-verb}]⁺.
(The pain subsides.)

Note that by transitive verbs, we understand all verbs that have *at least* two arguments. By arguments we do not only mean the subject and (direct) accusative object but all other types of objects, for instance, a dative object (19), a prepositional object (20) or object clause (21).

- (19) Die Menschheit [**entging**^{transitive-verb} [einer Katastrophe]^{objd}]⁺.
(Mankind averted disaster.)
- (20) Wir [**kämpfen**^{transitive-verb} [gegen dieses Problem]^{objp-gegen} **an**]⁺.
(We fight against this problem.)
- (21) Ich [**bezweifle**^{transitive-verb}, [dass dies eine gute Idee ist]^{objc}]⁻.
(I doubt that this is a good idea.)

For sentences where negation verbs have more than one object, the ordering of the objects on our *priority scope list* decides which type of object is

Type	Example Negation Words	Priority Scope List
negation function words		
negation adverbs and indefinite pronouns	<i>nie, niemals, kein, kaum</i>	clause
negation particle	<i>nicht</i>	governor
negation prepositions	<i>ohne, gegen</i>	dependent
negation content words		
negation adjectives	<i>weniger, gescheitert, korrigierbar</i>	subj, attr-rev
negation nouns	<i>Abschaffung, Linderung, Zerstörung</i>	gmod, objp-*
negation verbs intransitive	<i>enden, nachlassen, verschwinden</i>	subj
negation verbs transitive	<i>ablegen, lindern, senken, vermindern</i>	objg, obja, objd, objc, obji, s, objp-*, subj

Table 2: The different negation word types and their scopes.

given priority. For example, in case of ditransitive verbs, the accusative object is more likely to be negated than the dative object (22).

- (22) Das [ersparte^{transitive_verb} [uns]_{objd} [viel Ärger]_{obja}]⁺.
(This saved us a lot of trouble.)

In principle, the arguments of verbs to be negated could be most adequately described in terms of *semantic roles*. In the terminology of *FrameNet* (Baker et al., 1998), we are basically looking for THEME or PATIENT; in the terminology of *PropBank* (Palmer et al., 2005) it is *A1*. Unfortunately, automatic semantic role labeling for German is still in its infancy. As a consequence, we need to approximate semantic roles with dependency relations.

Using a priority scope list also partly allows us to model sense ambiguity. Some verbs may be used both intransitively and transitively. We simply add the subject position at the end of the priority list. This allows the German negation verb *abnehmen* (take from/decrease) to negate its accusative object in (23) while it negates its subject in (24).

- (23) Sie [nahm^{transitive_verb} ihm [ein große Last]_{obja} ab]⁺.
(She took a great burden from him.)
(24) [[Seine Wut]_{subj} nahm^{intransitive_verb} deutlich ab]⁺.
(His anger notably decreased.)

3.3.3 Normalization of the Dependency Graph


Our previous examples have shown that in order to model scope, we largely rely on a syntactic analysis, particularly on a dependency parse. We employ *ParZu* (Sennrich et al., 2009). We chose that particular parser because of its fine-grained label inventory which is essential for our approach. Still, our rules cannot be immediately applied to the original output of that parser.


Our rules are defined for active-voice constructions. The parse for passive-voice constructions

would be misleading since *ParZu* provides dependency structures that describe the surface structure. For example, in (25) we would not be able to correctly establish the scope of *bremsen* over *Fortschritt*, since it is marked as the *surface subject*. By normalizing the dependency relation labels to active voice (i.e. the deep structure), as indicated by (26), however, our rules work correctly since *Fortschritt* becomes an accusative object. It would be uneconomical to directly operate on the surface representation as it would mean writing redundant rules for negation scopes.

- (25) [[Der Fortschritt]_{subj-surface} wurde [von der Kirche]_{objp-surface} stets gebremst^{verb}]⁻.
(Progress was held off by the church.)
(26) [[Der Fortschritt]_{obja-deep} wurde [von der Kirche]_{subj-deep} stets gebremst^{verb}]⁻.
(Progress was held off by the church.)


Another major problem is that for several tensed predicates, such as *wird versiegt sein* (will be faded away) in (27), *ParZu* adds several auxiliary edges accommodating the auxiliary verbs of the predicate. As a consequence, a full verb and its arguments may no longer be directly related. For instance, in (27) the negation verb *versiegen* (dry up) is not directly related to its polar subject *Zuversicht* (confidence). Neither is the adjective *korrigierbar* (recoverable) in (29) directly related to its polar subject *Fehler* (error). In a further normalization step we, therefore, remove the edges involving the auxiliary verbs so that the full verb and its argument (28) or the predicate adjective and its argument (30) are directly connected.





3.3.4 Scope Expansion

Most of our negation rules assume the negation word and the polar expression it negates to be in a direct syntactic relation (31). However, there are also cases of negation in which there is no such direct relationship. For example, in (32) the polar expression that is negated is not the accusative object of the negation verb but its attributive adjective. To account for this, we implemented a *scope expansion* where also indirect relationships are allowed (i.e. we include the dependents of the words that match the direct syntactic relation).



4 Experiments

4.1 Intrinsic Evaluation on Negation Dataset

In this section, we evaluate on the dataset we specially created for the task of German negation detection for fine-grained sentiment analysis (§2). The task is to identify for each polar expression the negation word in whose scope it falls.

Since the focus of our work is neither to automatically detect polar expressions nor to detect negation words, in our first set of experiments, we consider them as given. That is, we read them off from the gold standard. The specific task therefore becomes to decide whether a *given* polar expression is negated by a *given* negation word.

Table 3 compares the different negation detection approaches. It clearly shows that our proposed method outperforms the two baseline methods, that is, the window-based approach (§3.1) and the clause-based approach (§3.2). Table 3 also displays the performance of our proposed method with some components, i.e. scope expansion (§3.3.4) or normalization (§3.3.3) switched off. The table shows that, clearly, both functionalities have a beneficial effect. With regard to the

Approach	Prec	Rec	F1
window-based (baseline I)	42.13	55.97	48.08
clause-based (baseline II)	38.89	60.07	47.21
proposed method	67.22	60.45	63.65
proposed meth. w/o scope expansion	69.65	52.24	59.70
proposed meth. w/o any normalization	71.54	34.70	46.73
proposed meth. w/o active-voice norm	66.80	60.07	63.26

Table 3: Comparison of different approaches.

Approach	Prec	Rec	F1
1 rule: <i>obja</i>	78.00	43.66	55.98
1 rule: <i>objg,obja,objd,objc,obji,s,objp-*,subj</i>	67.29	53.73	59.75
1 rule: <i>objg,obja,objd,objc,obji,s,objp-*,subj</i>	64.82	61.19	62.96
2 atomic rules: <i>obja</i> for trans.; <i>subj</i> for intrans.	75.78	45.52	56.88
2 verb rules as proposed in Table 2	67.22	60.45	63.65

Table 4: Impact of different **verbs** rules.

normalization, however, the active-voice conversion only contributes a minor share to the overall performance. So it is the conflation of relation edges in the dependency graph (27)-(30) that has the biggest impact.

Table 4 compares different verb rules. First, we evaluate a set of single verb rules, that is, we ignore the distinction between transitive and intransitive verbs. The performance is fairly competitive if we use the largest possible priority list *objg,obja,objd,objc,obji,s,objp-*,subj*. If we distinguish between transitive and intransitive verbs but only have two atomic rules and have no priority list (i.e. *obja* vs. *subj*), then this is worse than having only one rule but a priority list (i.e. *objg,obja,objd,objc,obji,s,objp-*,subj*). From that we conclude that many negated polar expressions are realized as a type of object but not necessarily as an accusative object (i.e. *obja*). Accounting for intransitive verbs has a relatively marginal impact, since the scores of *1 rule obja* and *2 atomic rules: obja for trans.; subj for intrans.* are not that far apart. We assume that the reason for this is that (deep) subjects are relatively rarely negated.

Our previous experiments all assumed knowledge of polar expressions in a sentence as given. We now want to examine how performance changes if we detect all polar expressions automatically. For this experiment, we employ the sentiment lexicon of the PolArt system (Klenner et al., 2009). The detection of polar expressions based on a lexicon has two disadvantages. Firstly, all existing sentiment lexicons only have a limited coverage. Secondly, lexicon look-up does not account for word-sense ambiguity, that is, some

Approach	Manual Detection			Automatic Detection		
	Prec	Rec	F1	Prec	Rec	F1
window-based (baseline I)	42.1	56.0	48.1	26.0	35.5	30.0
clause-based (baseline II)	38.9	60.1	47.2	23.2	37.7	28.7
proposed method	67.2	60.5	63.7	46.8	35.8	40.6

Table 5: Comparison of manual and automatic detection of polar expressions.

words may only convey subjectivity in certain contexts (Akkaya et al., 2009).

Table 5 compares the performance of our two baselines and our proposed method based on the manual detection of polar expressions and the automatic detection of those expressions. It comes as no surprise that the performance of classifiers based on the automatic detection is lower than that using manual detection. However, by and large, the difference between our three approaches to determine the scope of negation is similar on both detection types. In other words, no matter how the polar expressions are detected, our proposed method to determine negation always largely outperforms the two baseline classifiers.

We refrain from carrying out a similar experiment by detecting negation words automatically, since our dataset is biased towards the negation words we know. Moreover, inspection of our data revealed that polar expressions tend to be much more ambiguous than negation words.

4.2 Extrinsic Evaluation on Sentence-Level Polarity Classification

In this section, we evaluate our negation modeling approach on the task of sentence-level polarity classification. The task is to correctly classify the overall polarity of a given sentence.

We consider two datasets: the *Multi-layered Reference Corpus for German Sentiment Analysis (MLSA)* (Clematide et al., 2012) and the *Heidelberg Sentiment Treebank (HeiST)* (Haas and Versley, 2015). MLSA contains 270 sentences from the DeWaC Corpus (Baroni et al., 2009) which is a collection of German-language documents of various genres obtained from the web. HeiST contains 1184 sentences from German movie reviews.

We run two types of evaluations: a **three-class setting** in which the sentences are to be labeled as either *positive*, *negative* or *neutral*, and a **two-class setting** where we remove the neutral instances and the classifier just has to distinguish between *positive* and *negative* polarity. For HeiST,

Expression	Score	Score of Negation
exzellent (excellent)	+1.0	-0.3
ausreichend (sufficient)	+0.5	-0.8
umstritten (controversial)	-0.5	+0.8
miserable (abysmal)	-1.0	+0.3

Table 6: Illustration of negation scores.

Approach	2 Classes			3 Classes		
	Prec	Rec	F1	Prec	Rec	F1
w/o negation	65.3	52.9	58.4	50.8	50.8	50.8
with negation	67.3	54.7	60.3	52.0	51.9	52.0

Table 7: Polarity classification on HeiST.

we remove 253 (neutral) sentences in the two-class setting while for MLSA, we remove 91 sentences.

The polarity classification algorithm we follow is kept simple. For each sentence we sum the scores associated with the polar expressions occurring in that sentence according to the sentiment lexicon of the PolArt-system. In case a polar expression is within the scope of a negation, we move the polarity score in the opposite direction by the absolute value of 1.3. This is an adhoc-value, however, it complies with the recent elicitation study from Kiritchenko and Mohammad (2016) in that the score of a negated polar expression should not be represented as its inverse. This is since a negated polar expression (e.g. *not excellent*) is less polar intense than a (plain) polar expression of the opposite polarity with the same polar intensity (e.g. *abysmal*). Our scoring is illustrated in Table 6. The final sentence-level polarity is derived from the sign of the sum of scores that we computed.

In our experiments, we examine two different configurations, one where no negation modeling is considered and another where our proposed negation modeling is incorporated. We evaluate in terms of macro-average precision, recall and F-score. Table 7 shows the evaluation on HeiST while Table 8 shows the evaluation on the MLSA. In both cases our proposed negation modeling outperforms the polarity classifier in which no negation modeling is incorporated.

Approach	2 Classes			3 Classes		
	Prec	Rec	F1	Prec	Rec	F1
w/o negation	78.7	74.5	76.6	51.1	50.6	50.8
with negation	80.9	77.4	79.1	51.0	51.6	51.3

Table 8: Polarity classification on MLSA.

5 Related Work

The most notable work dealing with different types of negation is Wilson et al. (2005) who point out that there are other words expressing negation than the commonly associated negation (function) words *not*, *never*, *no* etc. Since that work is carried out on English data, the scope modeling is kept simple using a window-based approach. Recently, Socher et al. (2013) proposed the usage of Recursive Neural Tensor Network (RNTN) for sentiment analysis. RNTN is a compositional sentence-level polarity classifier providing polarity values for each node in a constituency parse of a sentence. The authors claim that this method allows learning negation directly from labeled training data without explicit knowledge of negation words and their scopes. However, there has been no empirical examination of how reliably RNTN actually models negation. Moreover, that approach only produced results inferior to conventional SVMs trained on bag of words on German data (Haas and Versley, 2015). For a detailed summary of negation modeling in sentiment analysis, we refer the reader to Wiegand et al. (2010).

Next to sentiment analysis, negation modeling has also been studied in the biomedical domain. Most of this work focuses on supervised classification on the (English) BioScope corpus (Szarvas et al., 2008), such as Morante et al. (2008) or Zou et al. (2013). The approach which is mostly related to ours is, however, the descriptive work by Morante (2010) who analyzes the individual negation words within the BioScope corpus and their scopes. This is one of the very few prominent research efforts that explicitly enumerates the different scopes of different negation words.

As far as German NLP is concerned, we are only aware of two research efforts that address negation. PolArt (Klenner et al., 2009) is a system that carries out sentence-level polarity classification. It matches polar expressions from a sentiment lexicon and then computes the sentence-level polarity compositionally on the basis of rules operating on syntactic constituents. This algorithm also incorporates negation modeling. However, the underlying lexicon only includes 22 polar shifters of which the majority are negation function words. The scope detection is further restricted by the fact that syntactic information is drawn from a chunk parser which only produces very flat output structures. Our work substantially

differs from Klenner et al. (2009) in that we devised a framework for negation scope detection that is able to handle many more types of negation words and allows the specification of individual scopes. Moreover, unlike Klenner et al. (2009), we employ a dependency parser and further normalize its output. So, we exploit much more accurate syntactic information.

Cotik et al. (2016) propose a method for negation detection in clinical reports. This approach is an adaptation of *NegEx* (Chapman et al., 2001) which is simple negation detection algorithm that operates on a set of negation cues embedded in lexical patterns (i.e. word token sequences). This method operates on the string level, that is, unlike our approach no form of syntactic parsing is considered. Cotik et al. (2016) consider a set of 167 German negation phrases. These are highly domain-specific phrases most of which include a common negation function word, e.g. *trifft fuer den Patienten nicht zu* (does **not** apply for the patient) or *keine Beschwerden ueber* (**no** complaints of). Due to the domain specificity of that approach, the negation cues and the scope detection mechanism cannot be applied to our dataset. Our approach also differs from Cotik et al. (2016) in that it is aimed at processing *unrestricted* text.

6 Conclusion

We presented an approach for modeling German negation in open-domain fine-grained sentiment analysis. Unlike most previous work in sentiment analysis, we assume that negation can be conveyed by many lexical units and that different negation words have different scopes.

We examined our approach on a new dataset comprising sentences with mentions of polar expressions and various negation words. We identify different types of negation words that have similar scopes. We showed that negation modeling based on these types largely outperforms traditional negation models assuming the same scope for all negation words no matter whether a window-based or clause-based scope is employed.

Our proposed method is only a first approximation of a more advanced negation handling for German. By making our implementation publicly available, we hope to stimulate further research in that direction using our new tool as a basis.

Acknowledgements

The authors would like to thank Stephanie Köser for annotating the dataset presented in this paper. The authors were partially supported by the German Research Foundation (DFG) under grants RU 1873/2-1 and WI 4204/2-1.

References

- Cem Akkaya, Janyce Wiebe, and Rada Mihalcea. 2009. Subjectivity Word Sense Disambiguation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 190–199, Singapore.
- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet Project. In *Proceedings of the International Conference on Computational Linguistics and Annual Meeting of the Association for Computational Linguistics (COLING/ACL)*, pages 86–90, Montréal, Quebec, Canada.
- Marco Baroni, Silvia Bernardini, Adriano Ferraresi, and Eros Zanchetti. 2009. The WaCky Wide Web: A Collection of Very Large Linguistically Processed Web-Crawled Corpora. *Language Resources and Evaluation*, 43(3):209–226.
- Wendy Webber Chapman, Will Bridewell, Paul Hanbury, Gregory F. Cooper, and Bruce G. Buchanan. 2001. A Simple Algorithm for Identifying Negated Findings and Diseases in Discharge Summaries. *Journal of Biomedical Informatics*, 34:301 – 310.
- Simon Clematide, Stefan Gindl, Manfred Klenner, Stefanos Petrakis, Robert Remus, Josef Ruppenhofer, Ulli Waltinger, and Michael Wiegand. 2012. MLSA – A Multi-layered Reference Corpus for German Sentiment Analysis. In *Proceedings of the Conference on Language Resources and Evaluation (LREC)*, pages 3551–3556, Istanbul, Turkey.
- Viviana Cotik, Roland Roller, Feiyu Xu, Hans Uszkoreit, Klemens Budde, and Danilo Schmidt. 2016. Negation Detection in Clinical Reports Written in German. In *Proceedings of the COLING-Workshop on Building and Evaluating Resources for Biomedical Text Mining (COLING-BioTxtM)*, pages 115–124, Osaka, Japan.
- Katrin Erk and Sebastian Padó. 2004. A powerful and versatile xml format for representing role-semantic annotation. In *Proceedings of the Conference on Language Resources and Evaluation (LREC)*, pages 799–802, Lisbon, Portugal.
- Michael Haas and Yannick Versley. 2015. Subsentential Sentiment on a Shoestring: A Crosslingual Analysis of Compositional Classification. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the ACL (HLT/NAACL)*, pages 694–704, Denver, CO, USA.
- Sanda Harabagiu, Andrew Hickl, and Finley Lacatusu. 2006. Negation, Contrast and Contradiction in Text Processing. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 755–762, Boston, MA, USA.
- Svetlana Kiritchenko and Saif M. Mohammad. 2016. The Effect of Negators, Modals, and Degree Adverbs on Sentiment Composition. In *Proceedings of the Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA)*, San Diego, CA, USA.
- Manfred Klenner, Angela Fahrni, and Stefanos Petrakis. 2009. PolArt: A Robust Tool for Sentiment Analysis. In *Proceedings of the Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 235–238, Odense, Denmark.
- J. Richard Landis and Gary G. Koch. 1977. The Measurement of Observer Agreement for Categorical Data. *Biometrics*, 33(1):159–174.
- Roser Morante, Anthony Liekens, and Walter Daelemans. 2008. Learning the Scope of Negation in Biomedical Texts. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 715–724, Honolulu.
- Roser Morante. 2010. Descriptive analysis of negation cues in biomedical texts. In *Proceedings of the Conference on Language Resources and Evaluation (LREC)*, Valletta, Malta, pages 1429–1436.
- Martha Palmer, Dan Gildea, and Paul Kingsbury. 2005. The Proposition Bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106.
- Livia Polanyi and Annie Zaenen. 2006. Contextual Valence Shifters. In *Computing Attitude and Affect in Text: Theory and Applications*, pages 1–10. Springer-Verlag.
- Olivia Sanchez-Graillet and Massimo Poesio. 2007. Negation of proteinprotein interactions: analysis and extraction. *Bioinformatics*, 23(13):i424–i432.
- Rico Sennrich, Gerold Schneider, Martin Volk, and Martin Warin. 2009. A New Hybrid Dependency Parser for German. In *Proceedings of the German Society for Computational Linguistics and Language Technology (GSCL)*, pages 115–124, Potsdam, Germany.
- 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)*, pages 1631–1642, Seattle, WA, USA.
- György Szarvas, Veronika Vincze, Richárd Farkas, and János Csirik. 2008. The BioScope Corpus: Annotation for Negation, Uncertainty and Their Scope

in Biomedical Texts. In *Proceedings of Current Trends in Biomedical Natural Language Processing (BioNLP)*, pages 38–45, Columbus, OH, USA.

Michael Wiegand, Alexandra Balahur, Benjamin Roth, Dietrich Klakow, and Andres 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*, pages 60–68, Uppsala, Sweden.

Michael Wiegand, Manfred Klenner, and Dietrich Klakow. 2013. Bootstrapping polarity classifiers with rule-based classification. *Language Resources and Evaluation*, 47(4):1049–1088.

Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing Contextual Polarity in Phrase-level Sentiment Analysis. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing (HLT/EMNLP)*, pages 347–354, Vancouver, BC, Canada.

Bowei Zou, Guodong Zhou, and Qiaoming Zhu. 2013. Tree Kernel-based Negation and Speculation Scope Detection with Structured Syntactic Parse Features. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 968–976, Seattle, WA, USA.