

UCM-I: A Rule-based Syntactic Approach for Resolving the Scope of Negation

Jorge Carrillo de Albornoz, Laura Plaza, Alberto Díaz and Miguel Ballesteros

Universidad Complutense de Madrid
C/ Prof. José García Santesmases, s/n
28040 Madrid (Spain)

{jcalbornoz, lplazam, albertodiaz, miballes}@fdi.ucm.es

Abstract

This paper presents one of the two contributions from the Universidad Complutense de Madrid to the *SEM Shared Task 2012 on Resolving the Scope and Focus of Negation. We describe a rule-based system for detecting the presence of negations and delimitating their scope. It was initially intended for processing negation in opinionated texts, and has been adapted to fit the task requirements. It first detects negation cues using a list of explicit negation markers (such as *not* or *nothing*), and infers other implicit negations (such as affixal negations, e.g. *undeniable* or *improper*) by using semantic information from WordNet concepts and relations. It next uses the information from the syntax tree of the sentence in which the negation arises to get a first approximation to the negation scope, which is later refined using a set of post-processing rules that bound or expand such scope.

1 Introduction

Detecting negation is important for many NLP tasks, as it may reverse the meaning of the text affected by it. In information extraction, for instance, it is obviously important to distinguish negated information from affirmative one (Kim and Park, 2006). It may also improve automatic indexing (Mutalik et al., 2001). In sentiment analysis, detecting and dealing with negation is critical, as it may change the polarity of a text (Wiegand et al., 2010). However, research on negation has mainly focused on the biomedical domain, and addressed the problem of

detecting if a medical term is negated or not (Chapman et al., 2001), or the scope of different negation signals (Morante et al., 2008).

During the last years, the importance of processing negation is gaining recognition by the NLP research community, as evidenced by the success of several initiatives such as the Negation and Speculation in Natural Language Processing workshop (NeSp-NLP 2010)¹ or the CoNLL-2010 Shared Task², which aimed at identifying hedges and their scope in natural language texts. In spite of this, most of the approaches proposed so far deal with negation in a superficial manner.

This paper describes our contribution to the *SEM Shared Task 2012 on Resolving the Scope and Focus of Negation. As its name suggests, the task aims at detecting the scope and focus of negation, as a means of encouraging research in negation processing. In particular, we participate in Task 1: scope detection. For each negation in the text, the negation cue must be detected, and its scope marked. Moreover, the event or property that is negated must be recognized. A comprehensive description of the task may be found in (Morante and Blanco, 2012).

For the sake of clarity, it is important to define what the organization of the task understands by negation cue, scope of negation and negated event. The words that express negation are called negation cues. *Not* and *no* are common examples of such cues. Scope is defined as the part of the meaning that is negated, and encloses all negated concepts. The negated event is the property that is

¹<http://www.clips.ua.ac.be/NeSpNLP2010/>

²www.inf.u-szeged.hu/rgai/conll2010st/

negated by the cue. For instance, in the sentence: *[Holmes] did not [say anything]*, the scope is enclosed in square brackets, the negation cue is underlined and the negated event is shown in bold. More details about the annotation of negation cues, scopes and negated events may be found in (Morante and Daelemans, 2012).

The system presented to the shared task is an adaptation of the one published in (Carrillo de Albornoz et al., 2010), whose aim was to detect and process negation in opinionated text in order to improve polarity and intensity classification. When classifying sentiments and opinions it is important to deal with the presence of negations and their effect on the emotional meaning of the text affected by them. Consider the sentence (1) and (2). Sentence (1) expresses a positive opinion, whereas that in sentence (2) the negation word *not* reverses the polarity of such opinion.

(1) *I liked this hotel.*

(2) *I didn't like this hotel.*

Our system has the main advantage of being simple and highly generic. Even though it was originally conceived for treating negations in opinionated texts, a few simple modifications have been sufficient to successfully address negation in a very different type of texts, such as Conan Doyle stories. It is rule-based and does not need to be trained. It also uses semantic information in order to automatically detect the negation cues.

2 Methodology

As already told, the UCM-I system is a modified version of the one presented in (Carrillo de Albornoz et al., 2010). Next sections detail the modifications performed to undertake the present task.

2.1 Detecting negation cues

Our previous work was focused on explicit negations (i.e., those introduced by negation tokens such as *not*, *never*). In contrast, in the present work we also consider what we call *implicit negations*, which includes affixal negation (i.e., words with prefixes such as *dis-*, *un-* or suffixes such as *-less*; e.g., *impatient* or *careless*), infixal negation (i.e., *pointlessness*, where the negation cue *less* is in the middle of the noun phrase). Note that we did not

Table 1: Examples of negation cues.

Explicit negation cues			
no	not	non	nor
nobody	never	nowhere	...
Words with implicit negation cues			
unpleasant	unnatural	dislike	impatient
fearless	hopeless	illegal	...

have into account these negation cues when analyzing opinionated texts because these words themselves usually appear in affective lexicons with their corresponding polarity values (i.e., *impatient*, for instance, appears in SentiWordNet with a negative polarity value).

In order to detect negation cues, we use a list of predefined negation signals, along with an automatic method for detecting new ones. The list has been extracted from different previous works (Councill et al., 2010; Morante, 2010). This list also includes the most frequent contracted forms (e.g., *don't*, *didn't*, etc.). The automated method, in turn, is intended for discovering in text new affixal negation cues. To this end, we first find in the text all words with prefixes *dis-*, *a-*, *un-*, *in-*, *im-*, *non-*, *il-*, *ir-* and the suffix *-less* that present the appropriate part of speech. Since not all words with such affixes are negation cues, we use semantic information from WordNet concepts and relations to decide. In this way, we retrieve from WordNet the synset that correspond to each word, using WordNet::SenseRelate (Patwardhan et al., 2005) to correctly disambiguate the meaning of the word according to its context, along with all its antonym synsets. We next check if, after removing the affix, the word exists in WordNet and belongs to any of the antonym synsets. If so, we consider the original word to be a negation cue (i.e., the word without the affix has the opposite meaning than the lexical item with the affix).

Table 1 presents some examples of explicit negation cues and words with implicit negation cues. For space reasons, not all cues are shown. We also consider common spelling errors such as the omission of apostrophes (e.g., *isnt* or *nt*). They are not likely to be found in literary texts, but are quite frequent in user-generated content.

This general processing is, however, improved with two rules:

Table 2: Examples of false negation cues.

no doubt	without a doubt	not merely	not just
not even	not only	no wonder	...

1. **False negation cues:** Some negation words may be also used in other expressions without constituting a negation, as in sentence (3). Therefore, when the negation token belongs to such expressions, this is not processed as a negation. Examples of false negation cues are shown in Table 2.

(3) ... *the evidence may implicate not only your friend Mr. Stapleton but his wife as well.*

2. **Tag questions:** Some sentences in the corpora present negative tag questions in old English grammatical form, as it may shown in sentences (4) and (5). We have implemented a specific rule to deal with this type of constructions, so that they are not treated as negations.

(4) *You could easily recognize it , could you not?*

(5) *But your family have been with us for several generations , have they not?*

2.2 Delimiting the scope of negation

The scope of a negation is determined by using the syntax tree of the sentence in which the negation arises, as generated by the Stanford Parser.³ To this end, we find in the syntax tree the first common ancestor that encloses the negation token and the word immediately after it, and assume all descendant leaf nodes to the right of the negation token to be affected by it. This process may be seen in Figure 1, where the syntax tree for the sentence: *[Watson did] not [solve the case]* is shown. In this sentence, the method identifies the negation token *not* and assumes its scope to be all descendant leaf nodes of the common ancestor of the words *not* and *solve* (i.e., *solve the case*).

This modeling has the main advantage of being highly generic, as it serves to delimit the scope of negation regardless of what the negated event is (i.e., the verb, the subject, the object of the verb, an adjective or an adverb). As shown in (Carrillo de Al-

³<http://nlp.stanford.edu/software/lex-parser.shtml>




Figure 1: Syntax tree of the sentence: *Watson did not solve the case.*

bornoz et al., 2010), it behaves well when determining the scope of negation for the purpose of classifying product reviews in polarity classes. However, we have found that this scope is not enough for the present task, and thus we have implemented a set of post-processing rules to expand and limit the scope according to the task guidelines:

1. **Expansion to subject.** This rule expands the negation scope in order to include the subject of the sentence within it. In this way, in sentence (6) the appropriate rule is fired to include “This theory” within the negation scope.

(6) *[This theory would] not [work].*

It must be noted that, for polarity classification purposes, we do not consider the subject of the sentence to be part of this scope. Consider, for instance, the sentence: *The beautiful views of the Eiffel Tower are not guaranteed in all rooms.* According to traditional polarity classification approaches, if the subject is considered as part of the negation scope, the polarity of the positive polar expression “beautiful” should be changed, and considered as negative.

2. **Subordinate boundaries.** Our original negation scope detection method works well with coordinate sentences, in which negation cues scope only over their clause, as if a “boundary” exists between the different clauses. This occurs, for instance, in the sentence:

Table 3: List of negation scope delimiters.

Tokens	POS
so, because, if, while	IN
until, since, unless	IN
before, than, despite	IN
what, whose	WP
why, where	WRB
however	RB
“;” , - , : , ; , (,) , ! , ? , .	-

(7) [It may be that you are] not [yourself luminous], but you are a conductor of light.

It also works properly in subordinate sentences, when the negation occurs in the subordinate clause, as in: *You can imagine my surprise when I found that [there was] no [one there].*

However, it may fail in some types of subordinate sentences, where the scope should be limited to the main clause, but our model predict both clauses to be affected by the negation. This is the case for the sentences where the dependent clause is introduced by the subordinate conjunctions in Table 3. An example of such type of sentence is (8), where the conjunction token *because* introduces a subordinate clause which is out of the negation scope. To solve this problem, the negation scope detection method includes a set of rules to delimit the scope in those cases, using as delimiters the conjunctions in Table 3. Note that, since some of these delimiters are ambiguous, their part of speech tags are used to disambiguate them.

(8) [Her father] refused [to have anything to do with her] **because** she had married without his consent.

- Prepositional phrases:** Our original method also fails to correctly determine the negation scope when the negated event is followed by a prepositional phrase, as it may be seen in Figure 2, where the syntax tree for the sentence: *[There was] no [attempt at robbery]* is shown. Note that, according to our original model, the phrase “at robbery” does not belong to the negation scope. This is an error that was not detected before, but has been fixed for the present task.




Figure 2: Syntax tree for the sentence: *There was no attempt at robbery.*

2.3 Finding negated events

We only consider a single type of negated events, so that, when a cue word contains a negative affix, the word after removing the affix is annotated as the negated event. In this way, “doubtedly” is correctly annotated as the negated event in sentence (9). However, the remaining types of negated events are relegated to future work.

(9) [The oval seal is] undoubtedly [a plain sleeve-link].

3 Evaluation Setup

The data collection consists of a development set, a training set, and two test sets of 787, 3644, 496 and 593 sentences, respectively from different stories by Conan Doyle (see (Morante and Blanco, 2012) for details). Performance is measured in terms of recall, precision and F-measure for the following subtasks:

- Predicting negation cues.
- Predicting both the scope and cue.
- Predicting the scope, the cue does not need to be correct.
- Predicting the scope tokens, where not a full scope match is required.
- Predicting negated events.
- Full evaluation, which requires all elements to be correct.

Table 4: Results for the development set.

Metric	Pr.	Re.	F-1
<i>Cues</i>	92.55	86.13	89.22
<i>Scope (cue match)</i>	86.05	44.05	58.27
<i>Scope (no cue match)</i>	86.05	44.05	58.27
<i>Scope tokens (no cue match)</i>	88.05	59.05	70.69
<i>Negated (no cue match)</i>	65.00	10.74	18.43
<i>Full negation</i>	74.47	20.23	31.82

4 Evaluation Results

The results of our system when evaluated on the development set and the two test sets (both jointly and separately), are shown in Tables 4, 5, and 6.

It may be seen from these tables that our system behaves quite well in the prediction of negation cues subtask, achieving around 90% F-measure in all data sets, and the second position in the competition. Performance in the scope prediction task, however, is around 60% F-1, and the same results are obtained if the correct prediction of cues is required (*Scope (cue match)*). This seems to indicate that, for all correct scope predictions, our system have also predicted the negation cues correctly. Obviously these results improve for the *Scope tokens* measure, achieving more than 77% F-1 for the Card-board data set. We also got the second position in the competition for these three subtasks. Concerning detection of negated events, our system gets poor results, 22.85% and 19.81% F-1, respectively, in each test data set. These results affect the performance of the full negation prediction task, where we get 32.18% and 32.96% F-1, respectively. Surprisingly, the result in the test sets are slightly better than those in the development set, and this is due to a better behavior of the WordNet-based cue detection method in the formers than in the later.

5 Discussion

We next discuss and analyze the results above. Firstly, and regarding detection of negation cues, our initial list covers all explicit negations in the development set, while the detection of affixal negation cues using our WordNet-based method presents a precision of 100% but a recall of 53%. In particular, our method fails when discovering negation cues such as *unburned*, *uncommonly* or *irreproachable*, where the word after removing the affix is a derived

form of a verb or adjective.

Secondly, and concerning delimitation of the scope, our method behaves considerably well. We have found that it correctly annotates the negation scope when the negation affects the predicate that expresses the event, but sometimes fails to include the subject of the sentence in such scope, as in: *[I know absolutely] nothing [about the fate of this man]*, where our method only recognizes as the negation scope the terms *about the fate of this man*.

The results have also shown that the method frequently fails when the subject of the sentence or the object of an event are negated. This occurs, for instance, in sentences: *I think, Watson, [a brandy and soda would do him] no [harm]* and *No [woman would ever send a reply-paid telegram]*, where we only point to “harm” and “woman” as the scopes.

We have found a further category of errors in the scope detection tasks, which concern some types of complex sentences with subordinate conjunctions where our method limits the negation scope to the main clause, as in sentence: *[Where they came from, or who they are,] nobody [has an idea]*, where our method limits the scope to “has an idea”. However, if the negation cue occurs in the subordinate clause, the method behaves correctly.

Thirdly, with respect to negated event detection, as already told our method gets quite poor results. This was expected, since our system was not originally designed to face this task and thus it only covers one type of negated events. Specifically, it correctly identifies the negated events for sentences with affixal negation cues, as in: *It is most improper, most outrageous*, where the negated event is “proper”. However, it usually fails to identify these events when the negation affects the subject of the sentence or the object of an event.

6 Conclusions and Future Work

This paper presents one of the two contributions from the Universidad Complutense de Madrid to the *SEM Shared Task 2012. The results have shown that our method successes in identifying negation cues and performs reasonably well when determining the negation scope, which seems to indicate that a simple unsupervised method based on syntactic information and a reduced set of post-processing rules

Table 5: Results for the test sets (jointly).

Metric	Gold	System	Tp	Fp	Fn	Precision	Recall	F-1
Cues	264	278	241	29	23	89.26	91.29	90.26
Scopes (cue match)	249	254	116	24	133	82.86	46.59	59.64
Scopes (no cue match)	249	254	116	24	133	82.86	46.59	59.64
Scope tokens (no cue match)	1805	1449	1237	212	568	85.37	68.53	76.03
Negated (no cue match)	173	33	22	11	151	66.67	12.72	21.36
Full negation	264	278	57	29	207	66.28	21.59	32.57

Table 6: Results for the Cardboard and Circle test sets.

Metric	Cardboard set			Circle set		
	Pr.	Re.	F-1	Pr.	Re.	F-1
Cues	90.23	90.23	90.23	88.32	92.37	90.30
Scope (cue match)	83.33	46.88	60.00	82.35	46.28	59.26
Scope (no cue match)	83.33	46.88	60.00	82.35	46.28	59.26
Scope tokens (no cue match)	84.91	72.08	77.97	85.96	64.50	73.70
Negated (no cue match)	66.67	13.79	22.85	66.67	11.63	19.81
Full negation	68.29	21.05	32.18	64.44	22.14	32.96

is a viable approach for dealing with negation. However, detection of negated events is the main weakness of our approach, and this should be tackled in future work. We also plan to improve our method for detecting affixal negations to increment its recall, by using further WordNet relations such as “derived from adjective”, and “pertains to noun”, as well as to extend this method to detect infixal negations.

Acknowledgments

This research is funded by the Spanish Ministry of Science and Innovation (TIN2009-14659-C03-01) and the Ministry of Education (FPU program).

References

- Jorge Carrillo de Albornoz, Laura Plaza, and Pablo Gervás. 2010. A hybrid approach to emotional sentence polarity and intensity classification. In *Proceedings of the 14th Conference on Computational Natural Language Learning (CoNLL 2010)*, pages 153–161.
- W. W. Chapman, W. Bridewell, P. Hanbury, G. F. Cooper, and B.G. Buchanan. 2001. A simple algorithm for identifying negated findings and diseases in discharge summaries. *J Biomed Inform*, 34:301–310.
- Isaac Council, Ryan McDonald, and Leonid Velikovich. 2010. What’s great and what’s not: learning to classify the scope of negation for improved sentiment analysis. In *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*, pages 51–59.
- Jung-Jae Kim and Jong C. Park. 2006. Extracting contrastive information from negation patterns in biomedical literature. *ACM Trans. on Asian Language Information Processing*, 5(1):44–60.
- Roser Morante and Eduardo Blanco. 2012. Sem 2012 shared task: Resolving the scope and focus of negation. In *Proceedings of the 1st Joint Conference on Lexical and Computational Semantics (*SEM 2012)*.
- Roser Morante and Walter Daelemans. 2012. Conandoyle-neg: Annotation of negation in conandoyle stories. In *Proceedings of the 8th International Conference on Language Resources and Evaluation*.
- Roser Morante, Anthony Liekens, and Walter Daelemans. 2008. Learning the scope of negation in biomedical texts. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 715–724.
- Roser Morante. 2010. Descriptive Analysis of Negation Cues in Biomedical Texts. In *Proceedings of the 7th International Conference on Language Resources and Evaluation*.
- A.G. Mutalik, A. Deshpande, and P.M. Nadkarni. 2001. Use of general-purpose negation detection to augment concept indexing of medical documents. A quantitative study using the UMLS. *J Am Med Inform Assoc*, 8(6):598–609.
- Siddharth Patwardhan, Satanjeev Banerjee, and Ted Pedersen. 2005. SenseRelate::TargetWord: a generalized framework for word sense disambiguation. In *Proceedings of the ACL 2005 on Interactive poster and demonstration sessions*, pages 73–76.
- Michael Wiegand, Alexandra Balahur, Benjamin Roth, Dietrich Klakow, and Andrés Montoyo. 2010. A survey on the role of negation in sentiment analysis. In *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*, pages 60–68.