Negation Scope Delimitation in Clinical Text Using Three Approaches: NegEx, PyConTextNLP and SynNeg

Hideyuki Tanushi[1], Hercules Dalianis[1], Martin Duneld[1], Maria Kvist[1,2],
Maria Skeppstedt[1], Sumithra Velupillai[1]

(1) Dept. of Computer and Systems Sciences (DSV), Stockholm University
Forum 100, 164 40 Kista, Sweden
(2) Dept. of clinical immunology and transfusion medicine, Karolinska University Hospital
171 76 Stockholm, Sweden

`hide-tan@dsv.su.se`, `hercules@dsv.su.se`, `xmartin@dsv.su.se`,

`maria.kvist@karolinska.se`, `mariask@dsv.su.se`, `sumithra@dsv.su.se`

ABSTRACT

Negation detection is a key component in clinical information extraction systems, as health record text contains reasonings in which the physician excludes different diagnoses by negating them. Many systems for negation detection rely on negation cues (e.g. *not*), but only few studies have investigated if the syntactic structure of the sentences can be used for determining the scope of these cues. We have in this paper compared three different systems for negation detection in Swedish clinical text (NegEx, PyConTextNLP and SynNeg), which have different approaches for determining the scope of negation cues. NegEx uses the distance between the cue and the disease, PyConTextNLP relies on a list of conjunctions limiting the scope of a cue, and in SynNeg the boundaries of the sentence units, provided by a syntactic parser, limit the scope of the cues. The three systems produced similar results, detecting negation with an F-score of around 80%, but using a parser had advantages when handling longer, complex sentences or short sentences with contradictory statements.

KEYWORDS: clinical text, negation detection, syntactic analysis.

1 Introduction

Clinical text is written by physicians and nurses in patient records. Most of these notes are noisy, written in telegraphic style, containing e.g. incomplete sentences with missing subjects. Moreover clinical text contains much more spelling errors than ordinary text. All this together makes it difficult for ordinary text processing tools to work properly.

In a clinical setting, it is important to record the symptoms and signs of diseases that the patient reports as a base for the following reasoning and speculation of possible diagnoses. Possible diseases that can be ruled out, as well as symptoms that the patient does not experience, are therefore frequently mentioned in the health record narrative. In one study, some of the most frequently used clinical observations were found to be negated the majority of the time (Chapman et al., 2001a). Negation detection in a clinical setting is essential for, e.g., text summarisation of clinical text, or for information retrieval of symptoms and diagnoses. Such systems would be very valuable for clinical decision support, knowledge extraction and supporting clinicians in their daily work.

The aim of this study is to determine if syntactic information is useful for negation scope detection in clinical text, compared to fixed or flexible rule-based scope detection.

Some of the problems with negation detection and determining the scope of what is negated, particularly in clinical texts, are linked with the characteristics mentioned above: the texts are often telegraphic and may contain information dense sentences where negated and affirmed diagnostic expressions are in close proximity, long enumerations, and complex sentences, where current negation detection approaches produce errors. Such problems may be better handled if syntactic information is used.

We therefore compared three different approaches for determining negation delimitation scope of diagnostic expressions:

1. Using the heuristic that the scope of a negation cue is six tokens, that is, if a diagnostic expression is mentioned in the proximity of six tokens from a negation cue it is considered as negated. This heuristic has previously been used for Swedish (Skeppstedt 2011).

2. Using the sentence as the scope for the negation cues, but limiting the scope if a conjunction is present in the sentence. This approach has previously been used for Swedish (Velupillai et al., 2013)

3. Using morphological and syntactic labels provided by a syntactic parser for determining the scope of a negation trigger.

2 Related Research

There are several studies on negation detection in English clinical text, both rule-based methods (e.g. Chapman et al., 2001b, Mutalik et al., 2001) and machine learning based methods (e.g. Rokach et al., 2010). The rule-based studies use a cue phrase that indicates negation, e.g. *not, no, rule out*. The most simple approaches use a heuristically determined scope, negating all symptoms and diseases that occur in close proximity to a cue phrase (Chapman et al., 2001b), whereas there are also approaches using other cue words, such as conjunctions, for limiting the scope of a negation cue phrase (Elkin et al., 2005, Chapman et al., 2011). Morante and Daelemans (2009) have studied the span of the scope to the left and to the right of the negated term in the BioScope Clinical Corpus and found that over 97 percent of the scopes are to the

right of the negation and are on average 6.33 tokens long, while the ones to the left are only 4.84 tokens long. Aronow et al. (1999) describe a negation detection system, NegExpander, where they checked the scope of conjunctions to decide scope of negations, i.e. a sort of basic syntactic analysis. In Goryachev et al. (2006) this approach was compared with other approaches such as NegEx (see below) and machine learning based approaches (SVM and Naïve Bayes).

There are also studies in which parsers have been used for detecting the scope of negation cues. The *SEM shared task (Morante and Blanco, 2012) provided a phrase structure parsed literary corpus, annotated for negation cues and their scope, and had as one of the tasks to detect this scope. Most participating groups based the scope detection on machine learning, and the two most successful both used conditional random fields, one with features from the syntactic tree and the other also using lexical features.

In the clinical domain, Huang and Lowe (2007) categorized negation cues, based on their syntactic category and on typical phrase patterns they occur in, and constructed a grammar for how negation is expressed for each category, using phrase structures given by the Stanford syntactic parser. This resulted in a precision of 98.6% and a recall of 92.6% for detecting negated noun phrases in radiology reports.

Another rule-based system, using the dependency parser Minipar, was constructed by Ballesteros et al. (2012). This system determines scope by traversing the dependency graph towards the terminals, starting from the negation cue, if the cue is a verb, and otherwise with the verb that is affected by the negation cue. This system detected negation cues and their scope with a precision of 95.8% and a recall of 90.6% on the clinical part of the BioScope corpus.

Zhu et al. (2010) also used the BioScope corpus and employed techniques developed for shallow semantic parsing for detecting scope. The negation cue was regarded as a predicate and its constituents, i.e. words included in its scope, were found by a machine learning system among candidates given by a phrase structure parser. This approach resulted in a precision of 82.2% and a recall of 80.6% for determining negation scope in the clinical part of BioScope, when using automatic cue and parse tree detection.

Velldal et al. (2012) combined the output of one rule-based and one machine-learning system for detecting scope of negation in the BioScope corpus. For the rule-based system, rules were constructed for determining the scope of the cue, given the output of the dependency parser MaltParser, trained on the Penn Tree-bank. Separate rules were constructed for each cue part-of-speech category, with a special rule for the word *none*. For the machine learning system, a phrase structure grammar was adapted to identify phrase structure constituents identical to those in the annotated scopes. For the scientific texts in the BioScope corpus, this method gave substantially improved results compared to a baseline of using the entire sentence as the scope. For the clinical texts in the corpus, however, the baseline scope detection gave slightly better results with an F-score of 91.4, compared to 90.7 (using gold standard cues). It was found that 40% of the scope errors in the clinical domain were due to parse errors.

3 Method

We here compare three approaches – fixed size context window, lexical cues and syntactic tags – to determine the scope of a negated diagnostic expression in a Swedish clinical corpus. All approaches rely on predefined trigger lists of negation cues. In order to study the effect of scope delimitation in negation detection we have used the same trigger lists for all systems.

3.1 NegEx

NegEx (Chapman et al., 2001b) is a widely used system for negation detection in clinical text that is built on three different lists of cue phrases that trigger negation; pre-negations, post-negations and pseudo-negations. The pre-negation list consists of cue phrases indicating that a clinical condition, i.e. a disease or a symptom, following the cue is negated, whereas the post-negation list consists of cue phrases indicating that a clinical condition preceding them is negated. Pseudo-negation phrases, on the other hand, are phrases that should not trigger a negation of a clinical condition, even though they contain a negation cue, e.g. *not only*. The first version of NegEx uses a heuristically determined number of tokens for the scope of a negation cue. A clinical condition is thus here negated if it is in the range of one to six tokens from a post- or pre-negation trigger. This version of NegEx has been adapted to Swedish, through a translation of English negation cues into Swedish[1], and through retaining the six token heuristics for determining the scope of the triggers used in the English study (Skeppstedt, 2011). Later versions of NegEx for English use a list of termination terms, e.g. conjunctions, to limit the scope of what clinical conditions a negation cue negates.

3.2 PyConTextNLP

PyConTextNLP (Chapman et al. 2011) is an extension of the NegEx algorithm that includes modifications of the scoping rules, how lexical triggers are matched, and more functionalities for defining user- and task-specific rules. It works on a sentence level and takes as input a database containing the texts or sentences to be analyzed, user-defined targets (here: diagnostic expressions) and user-defined lexical triggers. These triggers are represented as four-tuples: a literal (the lexical cue/trigger), a category (what the trigger represents), a regular expression (to capture variants of the trigger) and a rule that defines the direction a trigger modifies a target (forward, backward or bidirectional). For handling scopes, instead of a six-token window, the algorithm operates on the whole sentence, unless it finds user-defined conjunctions. PyConTextNLP has been ported to Swedish (Velupillai et al., 2013).

3.3 SynNeg

SynNeg is a negation scoping tool that uses morphological and syntactic information provided by the MaltParser (Nivre et al., 2007). The MaltParser is a language-independent and data-driven dependency parsing system which relies on inductive learning from treebank data for the analysis of new sentences. In Hassel et al. (2011) it was shown that the MaltParser trained on general Swedish text worked sufficiently on Swedish clinical text (92.4% accuracy for part-of-speech tagging and a labeled attachment score of 76.6% for the syntactic dependency parsing). Therefore, it is likely that the information provided by a parser trained on general Swedish text can be useful also for negation detection in clinical text.

The basic idea of SynNeg is that a negation scope does not cross the boundary of a sentence unit (i.e. subject + verb phrase). It uses a list of negation triggers in order to identify negative expressions in a sentence and tries to delimit a sentence unit to which a negation cue belongs, through finding another sentence unit in the sentence. The MaltParser assigns the ES (= logical subject), FS (= dummy subject) or SS (=other subject) DEPREL (Dependency Relation) tag to a subject of a sentence unit. When a negation cue is found, SynNeg checks the DEPREL

---

[1]This trigger list is used for all three systems. It contains a total of 42 negation cues: five post-negations, nine pseudo-negations and the remaining are pre-negations.

tags of either the following token (pre-negation) or the preceding token (post-negation) from the negation cue to find a subject DEPREL tag. It also checks the part-of-speech (POS) tag for coordinating conjunction (KN), minor delimiter, e.g. comma (MID), and subordinating conjunction (SN) in order to set the negation span boundary position. Every time one of these POS tags is found, the position of the token is stored as a boundary candidate. Once a subject DEPREL tag is found, the nearest boundary candidate from the subject DEPREL tag is set as the negation span boundary.

3.4 Data

A Swedish clinical corpus[2] annotated for uncertainty and negation on a diagnostic expression level was used in these experiments. The corpus consists of assessment entries from a medical emergency ward in a Swedish hospital (Velupillai et al., 2011) and was annotated by a physician. The annotations, originally for six levels of uncertainty and negation, were collapsed into two annotation classes: negated and non-negated. The data set consists of 8874 sentences (average sentence length 9.56 tokens), of which 2189 sentences contain diagnostic expressions (one sentence contains the same diagnostic expression mentioned twice). 421 of these diagnostic expressions were tagged as *negated* during the annotation of this data set, while 1769 were tagged as *non-negated*. The distribution of negation cues in the data set is shown in table 1.

Pre-negation			Post-negation		
Negation cue	n	Percentage	Negation cue	n	Percentage
ingen (no, common gender)	363	21.25%	*negativt* (negative for)	13	0.76%
ej (not)	344	20.14%	*osannolikt* (unlikely)	9	0.53%
inga (no, plural)	240	14.05%	*inget avvikande* (no abnormal)	6	0.35%
inte (not)	209	12.24%			
utan (without)	120	7.03%	*saknas* (absence of)	1	0.06%
utesluta (rule out)	87	5.09%	Pseudo-negation		
inga tecken (no signs of)	76	4.45%	Negation cue	n	Percentage
inget (no, neuter gender)	69	4.04%	*kan inte uteslutas* (cannot be ruled out)	7	0.41%
uteslutas (be ruled out)	27	1.58%			
inte har (not have, reversed word order)	20	1.17%	*inte utesluta* (not rule out)	3	0.18%
kan inte (cannot)	17	1.00%	*ingen förändring* (no change)	2	0.12%
aldrig (never)	13	0.76%			
har inte (not have)	13	0.76%	*vet inte* (do not know)	2	0.12%
icke (non-, not)	13	0.76%			
inte kan (cannot, reversed word order)	13	0.76%			
inget som (nothing that)	12	0.70%			
utan tecken (without sign of)	10	0.59%			
förnekar (denies)	7	0.41%			
inte visar (not demonstrate)	6	0.35%			
avsaknad av (absence of)	4	0.23%			
kunde inte (cannot, past tense)	1	0.06%			
utan några (with no)	1	0.06%			

Table 1: Frequency and percentage of negation cues in the data set. n is the number of instances of each cue. The total number of cues in the data was 1708.

The MaltParser was used to parse the texts and used directly as input for the SynNeg algorithm. This parsed data was also used to construct data in the input formats used by NegEx and PyConTextNLP, respectively. As the MaltParser requires POS tagged data as input, we first tagged the clinical text with the Granska tagger (Knutsson et al., 2003).

4 Results

Table 2 summarizes the result of negation detection for the three systems. As the data was annotated on a diagnostic expression level, the figures show precision and recall for detecting whether these diagnostic expressions are negated or not. Overall results are very similar for NegEx, PyConTextNLP and SynNeg. NegEx produces slightly higher precision results for *negated* (79.6), while SynNeg results in slightly higher recall (82.9) and F-Score (79.9).

	NegEx	PyConTextNLP	SynNeg
Precision	79.6	78.1	77.0
Recall	79.6	81.2	82.9
F-Score	79.6	79.6	79.9

Table 2: Results of negation detection for NegEx, PyConTextNLP and SynNeg.

Some examples of how diagnostic expressions are classified by the three systems are presented in figures 1 and 2 and table 3. Figure 1 shows an example where a non-negated instance is within 6 tokens from a negation cue and a conjunction is missing. A sentence with a negated instance beyond a 6-token window is presented in figure 2. Table 3 demonstrates a long sentence containing an adversative expression and a subordinate conjunction.
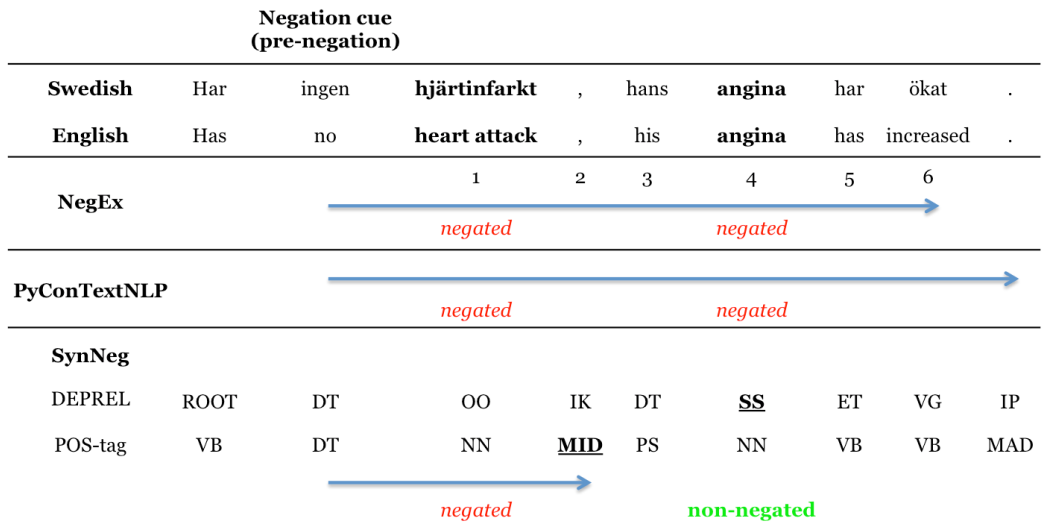

Figure 1: Sentence with non-negated instance within 6 tokens and lack of conjunction. Negated instances and non-negated instances are illustrated with red-colored italic and green-colored bold, respectively. DEPREL (Dependency Relation) and POS-tag as given by MaltParser and Granska, respectively.

Figure 2 table:

Negation cue (pre-negation)											
Swedish	Han	har	inte	någon	på	röntgen	synlig	**fissur**	eller	**fraktur**	.
English	He	does	not	have	any	x-ray	visible	fissure	or	fracture	.

NegEx — 1 2 3 4 5 6 → *negated* / **non-negated**

PyConTextNLP → *negated* / *negated*

SynNeg

DEPREL	SS	ROOT	NA	DT	ET	PA	AT	CJ	OO	CJ	IP
POS-tag	PN	VB	AB	PN	PP	NN	JJ	NN	KN	NN	MAD

negated / *negated*

Figure 2: Sentence with negated instance beyond a 6-token window. Negated instances and non-negated instances are illustrated with red-colored italic and green-colored bold, respectively. DEPREL (Dependency Relation) and POS-tag as given by MaltParser and Granska, respectively.

Vid ankomst opåv, ingen *astma*, inga hållpunkter för *lunginflammation*, dock **UVI** och anamnestisk klåda utan att man kan se *urticaria* eller annan *dermatit*.
(At arrival unaff, no *asthma*, no signs of *pneumonia*, but **UTI** and anamnestical itching without any visible *urticharia* or other *dermatitis*.)

Swedish	English	Annotation	NegEx	PyConTextNLP	SynNeg
astma	asthma	*negated*	*negated*	*negated*	*negated*
lunginflammation	pneumonia	*negated*	*negated*	*negated*	*negated*
UVI	UTI	**non-neg**	*negated*	**non-neg**	*negated*
urticharia	urticharia	*negated*	*negated*	*negated*	**non-neg**
dermatit	dermatitis	*negated*	**non-neg**	*negated*	**non-neg**

Table 3: Sentence with an adversative expression and a subordinate conjunction. Negation cue, negated instance and non-negated instance are illustrated with underlined, red-colored italic and green-colored bold, respectively.

5 Discussion

Given the relatively small sample size, the differences in results between the three systems are difficult to draw conclusions from. Employing a simple, rule-based approach with a small amount of negation triggers and a fixed context window for determining scope is very efficient and useful, if results around 80% F-score are sufficient for a given purpose.

However, to better handle sentences such as the examples in figures 1 and 2 and table 3, other approaches are necessary. In the example in figure 1, (*Has no heart attack, his angina has increased.*), the lack of a conjunction makes it difficult to specify a cue phrase for delimiting the scope in PyConTextNLP, while this can be captured through the syntactic rules employed in SynNeg. On the other hand, table 3, (*At arrival unaff, no asthma, no signs of pneumonia, but UTI and anamnestical itching without any visible urticharia or other dermatitis.*), shows an

example where PyConTextNLP correctly classifies all instances, while the other two systems produce errors. The scoping window of six tokens in NegEx is problematic (*UTI* incorrectly classified as negated and *dermatitis* incorrectly classified as non-negated), and SynNeg produces an error due to the subordinating conjunction *att* following *utan* (without) (*UTI*, *urthicaria* and *dermatitis*).

Both PyConTextNLP and SynNeg can correctly classify more complex cases, such as sentences with enumerations beyond the 6-token window (see figure 2) and short sentences with conjunctions that contradict instances. For improving the results of PyConTextNLP, it is essential to define lexical cue phrases, such as conjunctions, that delimit the scope of the negation trigger. For SynNeg, on the other hand, no lexical information is needed, only correct Part-of-Speech tags and syntactic labels. The advantage of the latter approach is that it generalizes over lexical terms.

The current version of SynNeg does not utilize the dependency tree produced by the parser, i.e. no traversal is done as in e.g. Ballesteros et al. (2012). This could be useful and will be studied further in the development of SynNeg.

The results for all three systems are in line with previous studies on negation detection for Swedish, e.g. Skeppstedt (2011). The results for most English studies are higher, e.g. Huang and Lowe (2007), but it is difficult to directly compare the findings, for instance given differences in used data.

One limitation in this study is that no evaluation on the performance of the MaltParser on this data has been conducted. A previous study shows that applying the MaltParser on Swedish clinical text yields promising results (Hassel et al., 2011), but that the parser produces errors in sentences with complicated conjunctional, conditional and prepositional constructions. As these are important factors for the correct identification of negation scope, this needs to be studied further.

The fixed context window approach (NegEx) is efficient and simple, but limited in its approach and will not be able to handle longer, complex sentences or short sentences with contradictory statements. With the more flexible approach employed in PyConTextNLP, it is possible to improve results if lexical phrases that define boundaries for the negation scope are defined. However, this approach will be problematic for cases that are ambiguous. Moreover, finding these lexical phrases may be time-consuming, and there is a risk of overfitting. Using syntactic information (SynNeg) instead may prove fruitful, as it is more generalizable. Also, with such an approach, it will be easier to port the system to another domain or language. However, it is crucial that the syntactic parser produces correct labels, and, of course, that a syntactic parser is available for that domain and language.

We plan to further develop SynNeg and evaluate it on larger datasets. In particular, we intend to incorporate information from traversing the dependency tree, in order to handle cases where the current heuristics are limited (double negation is, for example, currently not handled by any of the systems). Ongoing studies on automatically generating trigger terms both for negation and conjunctions for improving PyConTextNLP for Swedish clinical texts will also be of importance, and we will continue to compare the two approaches in order to further understand how well these types of systems can perform, along with comparisons of current state-of-the-art machine learning based approaches. We hypothesize that a combination of lexical- and syntactic-based systems may be very powerful, as they may complement each other.

6 Concluding Remarks

In this study we compare three different approaches for determining negation delimitation scope in Swedish clinical text. The approaches are rule-based, one of them relying on a fixed context window (NegEx) and one relying on full sentences or predefined lexical cue phrases for delimitation (PyConTextNLP). The third approach utilizes Part-of-Speech and syntactic labels (SynNeg). All three systems depend on the same trigger list of negation cues. Results show that although we are at the initial stages of developing SynNeg we achieve similar results compared to established lexical rule-based negation detection approaches: around 80% F-score. Using syntactic information may also prove more generalizable, facilitating porting to other domains or languages.

References

Aronow, D. B., Fangfang, F., and Croft, W. B. (1999). Ad Hoc Classification of Radiology Reports. *J Am Med Inform Assoc*, 5(6):393–411.

Ballesteros, M., Francisco, V., Díaz, A., Herrera, J., and Gervás, P. (2012). Inferring the scope of negation in biomedical documents. In *13th International Conference on Intelligent Text Processing and Computational Linguistics (CICLING 2012)*, New Delhi. Springer, Springer.

Chapman, B., Lee, S., Kang, H., and Chapman, W. (2011). Document-level classification of CT pulmonary angiography reports based on an extension of the ConText algorithm. *J Biomed Inform*, 44(5):728–737.

Chapman, W., Bridewell, W., Hanbury, P., Cooper, G., and Buchanan, B. (2001a). Evaluation of negation phrases in narrative clinical reports. In *Proceedings AMIA*, pages 105–109.

Chapman, W. W., Bridewell, W., Hanbury, P., Cooper, G. F., and Buchanan, B. G. (2001b). A simple algorithm for identifying negated findings and diseases in discharge summaries. *J Biomed Inform*, pages 34–301.

Elkin, P. L., Brown, S. H., Bauer, B. A., Husser, C. S., Carruth, W., Bergstrom, L. R., , and Wahner-Roedler, D. L. (2005). A controlled trial of automated classification of negation from clinical notes. *BMC Med Inform Decis Mak*, 13(5).

Goryachev, A., Sordo, M., Ngo, L., and Zeng, Q. (2006). Implementation and evaluation of four different methods of negation detection. Technical report, DSG.

Harkema, H., Dowling, J. N., Thornblade, T., and Chapman, W. W. (2009). Context: An algorithm for determining negation, experiencer, and temporal status from clinical reports. *J. of Biomedical Informatics*, 42(5):839–851.

Hassel, M., Henriksson, A., and Velupillai, S. (2011). Something Old, Something New – Applying a Pre-trained Parsing Model to Clinical Swedish. In *Proc. 18th Nordic Conf. on Comp. Ling. - NODALIDA '11*.

Huang, Y. and Lowe, H. J. (2007). A Novel Hybrid Approach to Automated Negation Detection in Clinical Radiology Reports. *JAMIA*, 3(14):304–311.

Knutsson, O., Bigert, J., and Kann, V. (2003). A robust shallow parser for Swedish. In *Proc. 14th Nordic Conf. on Comp. Ling. – NODALIDA '03*.

Morante, R. and Blanco, E. (2012). *sem 2012 shared task: resolving the scope and focus of negation. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics - Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation*, SemEval '12, pages 265–274, Stroudsburg, PA, USA. Association for Computational Linguistics.

Morante, R. and Daelemans, W. (2009). Learning the scope of hedge cues in biomedical texts. In *Proceedings of the Workshop on BioNLP*, pages 28–36.

Mutalik, P, Deshpande, A., and Nadkarni, P. (2001). Use of general-purpose negation detection to augment concept indexing of medical documents a quantitative study using the umls. *JAMIA*, 8(6):598–609.

Nivre, J., Hall, J., Nilsson, J., Chanev, A., Eryigit, G., Kübler, S., Marinov, S., and Marsi, E. (2007). MaltParser: A language-independent system for data-driven dependency parsing. *Natural Language Engineering*, 13(2):95–135.

Rokach, L., Romano, R., and Maimon, O. (2010). Negation recognition in medical narrative reports. *Information Retrieval*, 11(6):499–538.

Skeppstedt, M. (2011). Negation detection in Swedish clinical text: An adaption of NegEx to Swedish. *Journal of Biomedical Semantics*, 2(Suppl 3):S3.

Velldal, E., Øvrelid, L., Read, J., and Oepen, S. (2012). Speculation and negation: Rules, rankers, and the role of syntax. *Comput. Linguist.*, 38(2):369–410.

Velupillai, S., Dalianis, H., and Kvist, M. (2011). Factuality Levels of Diagnoses in Swedish Clinical Text. In Moen, A., Andersen, S. K., Aarts, J., and Hurlen, P, editors, *Proc. XXIII International Conference of the European Federation for Medical Informatics – User Centred Networked Health Care*, pages 559 – 563, Oslo. IOS Press.

Velupillai, S., Skeppstedt, M., Kvist, M., Mowery, D., Chapman, B., Dalianis, H., and Chapman, W. (2013). Porting a rule-based assertion classifier for clinical text from English to Swedish. In *Proc. 4th Louhi 2013*, Sydney, Australia.

Zhu, Q., Li, J., Wang, H., and Zhou, G. (2010). A unified framework for scope learning via simplified shallow semantic parsing. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, EMNLP '10, pages 714–724, Stroudsburg, PA, USA. Association for Computational Linguistics.