UConcordia: CLaC Negation Focus Detection at *Sem 2012

Sabine Rosenberg and Sabine Bergler
CLaC Lab, Concordia University
1455 de Maisonneuve Blvd West, Montréal, QC, Canada, H3W 2B3
`sabin_ro@cse.concordia.ca, bergler@cse.concordia.ca`

Abstract

Simply detecting negation cues is not sufficient to determine the semantics of negation, scope and focus must be taken into account. While scope detection has recently seen repeated attention, the linguistic notion of focus is only now being introduced into computational work. The *Sem2012 Shared Task is pioneering this effort by introducing a suitable dataset and annotation guidelines. CLaC's NegFocus system is a solid baseline approach to the task.

1 Introduction

Negation has attracted the attention of the NLP community and we have seen an increased advance in sophistication of processing tools. In order to assess factual information as asserted or not, it is important to distinguish the difference between

> (1) (a) *Newt Gingrich Not Conceding Race After Losing Florida Primary*
> (b) *Newt Gingrich Conceding Race After Losing Florida Primary*

This distinction is important and cannot be properly inferred from the surrounding context, *not* conceding a race after losing is in fact contrary to expectation in the original headline (1a), and the constructed (1b) is more likely in isolation.

Negation has been addressed as a task in itself, rather than as a component of other tasks in recent shared tasks and workshops. Detection of negation cues and negation scope at CoNLL (Farkas et al., 2010), BioNLP (Kim et al., 2011) and the Negation and Speculation in NLP Workshop (Morante and Sporleder, 2010) laid the foundation for the *Sem 2012 Shared Task. While the scope detection has been extended to fictional text in this task, an important progression from the newspaper and biomedical genres, the newly defined Focus Detection for Negation introduces the important question: what is the intended opposition in (1a)? The negation trigger is *not*, the scope of the negation is the entire verb phrase, but which aspect of the verb phrase is underscored as being at variance with reality, that is, which of the following possible (for the sake of linguistic argument only) continuations is the more likely one:

> (2) i *..., Santorum does.*
>   ($\neg Newt\ Gingrich$)
> ii *..., Doubling Efforts* ($\neg concede$)
> iii *..., Demanding Recount* ($\neg race$)
> iv *..., Texas redistricting at fault*
>   ($\neg Florida$)

This notion of *focus of negation* is thus a pragmatic one, chosen by the author and encoded with various means. Usually, context is necessary to determine focus. Often, different possible interpretations of focus do not change the factual meaning of the overall text, but rather its coherence. In (1 a) the imagined possible contexts (2 ii) and (2 iii) closely correspond to a simple negation of (1 b), (2 i) and (2 iv) do not feel properly represented by simply negating (1 b). This level of interpretation is contentious among people and it is the hallmark of well-written, well-edited text to avoid unnecessary guesswork while at the same time avoiding unnecessary

294

clarifying repetition. The potential for ambiguity is demonstrated by Example (3) from (Partee, 1993), where it is questionable whether the speaker in fact *has possession* of the book in question.

(3) *I didn't get that book from Mary*

Here, if the focus is *from Mary*, it would be likely that the speaker has possion of the book, but received it some other way. If the focus is *that book*, the speaker does not have possession of it.

It is important to note hat this notion of focus is not syntactically determined as shown in (3) (even though we use syntactic heuristics here to approximate it) but pragmatically and it correlates with pronunciation stress, as discussed in linguistics by (Han and Romero, 2001). More recently, focus negation has been identified as a special use (Poletto, 2008). The difference of scope and focus of negation are elaborated by (Partee, 1993), and have been used for computational use by (Blanco and Moldovan, 2011).

The *Sem 2012 Task 2 on Focus Detection builds on recent negation scope detection capabilities and introduces a gold standard to identify the focus item. Focus of negation is annotated over 3,993 sentences in the WSJ section of the Penn TreeBank marked with MNEG in PropBank. It accounts for verbal, analytical and clausal relation to a negation trigger; the role most likely to correspond to the focus was selected as focus. All sentences of the training data contain a negation. A sample annotation from the gold standard is given in (4), where PropBank semantic roles are labelled *A1, M-NEG,* and *M-TMP* and focus is underlined (until June).

(4) $\langle A\,decision_{A_1}\rangle\,is\langle n't_{M-NEG}\rangle\,expected$
$\langle\,\underline{until\,June}\;_{M-TMP}\rangle$

2 Previous Work

A recent study in combining regular pattern extraction with parse information for enhanced indexing of radiology reports showed effective detection of negated noun phrases for that corpus (Huang and Lowe, 2007). NegFinder (Mutalik et al., 2001) detects negated concepts in dictated medical documents with a simple set of corpus specific context-free rules, and they observe that in their corpus "One of the words *no, denies/denied, not,* or *without* was present in 92.5 percent of

all negations." Interestingly, several of their rules concern coordination (*and, or*) or prepositional phrase attachment patterns (*of, for*). NegEx (Chapman et al., 2001) is publicly available and maintained and updated with community-enhanced trigger lists (`http://code.google.com/p/negex/wiki/NegExTerms`). NegEx "locates trigger terms indicating a clinical condition is negated or possible and determines which text falls within the scope of the trigger terms." NegEx uses a simple regular expression algorithm with a small number of negation phrases and focuses on a wide variety of triggers but limits them to domain relevant ones. Consequently, the trigger terms and conditions are heavily stacked with biomedical domain specific terms. Outside the biomedical text community, sentiment and opinion analysis research features negation detection (Wilson, 2008). Current gold standard annotations for explicit negation as well as related phenomena include TIMEBANK (Pustejovsky et al., 2003), MPQA (Wiebe et al., 2005), and Bio-Scope (Vincze et al., 2008).

(Wiegand et al., 2010) presents a flat feature combination approach of features of different granularity and analytic sophistication, since in opinion mining the boundary between negation and negative expressions is fluid.

3 CLaC's NegFocus

CLaC Labs' general, lightweight negation module is intended to be embedded in any processing pipeline. The heuristics-based system is composed of three modules for the GATE (Cunningham et al., 2011) environment: the first component detects and annotates explicit negation cues present in the corpus, the second component detects and annotates the syntactic scope of the detected instances of verbal negation, and the third component implements focus heuristics for negation. The first two steps were developed independently, drawing on data from MPQA (Wiebe et al., 2005) and TIME-BANK (Pustejovsky et al., 2003) with validation on Bio-Scope (Vincze et al., 2008). The third step has been added based on data for the *Sem 2012 challenge and is intended to validate both, the first two "preprocessing" steps and the simple heuristic approximation of focus.

3.1 Data Preprocessing

Parser-based, our focus detection pipeline requires as input entire sentences. Therefore, the first step requires the extraction of each sentence utilizing the supplied token numbers and save them in the correct format. The system then performs standard preprocessing: sentence splitting, tokenization, parsing using the Stanford Parser (Klein and Manning, 2003; de Marneffe and Manning, 2006) and morphological preprocessing. Note that NegFocus does not use any PropBank annotations nor other provided training annotations, resulting in an independent, parser-based stand-alone module.

3.2 Detection of Negation Triggers

The Focus Detection task only considers the explicit negation cues *not, nor, never*. The first step in Neg-Focus is thus to identify these triggers in the sentences using an explicit negation trigger word list.

3.3 Syntactic Scope Detection

The Focus Detection task only considers negation of verbs. Thus, NegFocus extracts the syntactic complement of the verb to form the negated verb phrase from the dependency graphs (de Marneffe and Manning, 2006). We annotate this as the syntactic scope of the negation. Note that while we use dependency graphs, our syntactic scope is based on the parse tree and differs from the notion of scope encoded in Bio-Scope (Vincze et al., 2008) and the related format used for the *Sem 2012 Negation Scope Annotation task, which represent in our opinion the pragmatic notion of scope for the logical negation operation. Syntactic scope detection is thus considered to be a basic stepping stone towards the pragmatic scope and since the Focus Detection task does not provide scope annotations, we use syntactic scope here to validate this principle.

Our heuristics are inspired by (Kilicoglu and Bergler, 2011). In the majority of cases the dependency relation which identifies the syntactic scope is the *neg* relation. Traditionally, parse trees identify scope as lower or to the right of the trigger term, and our scope module assumes these grammatical constraints, yet includes the verb itself for the purposes of the shared task. Example 5, from the training dataset "The Hound of the Baskervilles" by Co-nan Doyle for the *Sem 2012 Negation Scope Annotation task, demonstrates our syntactic scope of the negation (underlined), in contrast with the gold standard scope annotation (in brackets). The gold annotation guidelines follow the proposal of Morante et al. (Morante et al., 2011)[1].

> (5) *[We did] not [drive up to the door] but got down near the gate of the avenue.*

3.4 Focus Heuristics

The third and final step for NegFocus is to annotate focus in sentences containing verbal negations. Using the verbal negation scope annotations of the previous step, four focus heuristics are invoked:

3.4.1 Baseline

The Baseline heuristic for this component is defined according to notions discussed in (Huddleston and Pullum, 2002), where the last constituent in the verb phrase of a clause is commonly the default location to place the heaviest stress, which we here equate with the focus. Example (6) depicts an instance where both NegFocus results (underlined) and the gold focus annotation (in brackets) match exactly. The baseline heuristic achieves 47.4% recall and 49.4% precision on the training set and 47% recall and 49.7% precision on the test set.

> (6) *NBC broadcast throughout the entire night and did not go off the air [until noon yesterday] .*

As pointed out in Section 3.3, focus is not always determined by scope (Partee, 1993). The training data gave rise to three additional heuristics.

3.4.2 Adverb

When an adverb is directly preceding and connected through an *advmod* dependency relation to the negated verb, the adverb constituent is determined as the focus of the negation.

> (7) *Although it may not be [legally] obligated to sell the company if the buy-out group can't revive its bid, it may have to explore alternatives if the buyers come back with a bid much lower than the group 's original $ 300-a-share proposal.*

296

3.4.3 Noun Subject Passive

Passives are frequent in newspaper articles and passive constructions front what would otherwise be the verb complement. Thus the fronted material should be eligible for focus assignment. Passives are flagged through the *nsubjpass* dependency, and for cases where the negated verb participates in an nsubjpass relation and has no other complement, the nsubjpass is determined as the focus.

(8) *[Billings] were n't disclosed.*

3.4.4 Negation Cue

The challenge data has cases where the negation cue itself is its own focus. These cases seem to be pragmatically determined. Error cases were reduced when determining the negation cue to be its own focus in two cases. The first case occurs when the negated verb has an empty complement (and is not a passive construction), as in Example 9.

(9) *Both said the new plan would [n't] work.*

The second case occurs when the negated verb embeds a verb that we identify as an implicit negation. We have a list of implicit negation triggers largely compiled from MPQA (Wiebe et al., 2005). Implicit negations are verbs that lexically encode a predicate and a negation, such as *reject* or *fail*.

(10) *Black activist Walter Sisulu said the African National Congress would [n't] reject violence as a way to pressure the South African government into concessions that might lead to negotiations over apartheid . . .*

4 Results

Ordering the heuristics impacts on recall. We place the most specific heuristics before the more general ones to avoid starvation effects. For example, the *adverb* heuristic followed by the *noun subject passive* heuristic achieved better results at the beginning, since they are more specific then the *negation cue* heuristic.

Table 1 shows the performance of the heuristics of NegFocus on the test set and on the development set. We observe that the heuristics are stable across the two sets with a 60% accuracy on the test set. The worst performer is the baseline, which is very coarse

for such a semantically sophisticated task: assuming that the last element of the negated verb phrase is the focus is truly a baseline.

heuristic	corr.	incorr.	acc.
Test Set			
baseline	336	238	.59
adverb	26	4	.87
nsubjpass	10	8	.56
neg. cue	33	20	.62
Development Set			
baseline	257	174	.6
adverb	15	6	.71
nsubjpass	10	6	.63
neg. cue	21	19	.53

Figure 1: Performance of NegFocus heuristics

The overall performance of the system is almost balanced between precision and recall with an f-measure of .58.

	Test Set	
Precision	60.00	[405/675]
Recall	56.88	[405/712]
F-score	58.40	
	Development Set	
Precision	59.65	[303/508]
Recall	57.06	[303/531]
F-score	58.33	

Figure 2: System Results

Our heuristics, albeit simplistic, are based on linguistically sound observations. The heuristic nature allows additional heuristics that are more tailored to a corpus or a task to be added without incurring unmanageable complexity, in fact each heuristic can be tested on the development set and can report on the test set to monitor its performance. The heuristics will also provide excellent features for statistical systems.

5 Error Analysis

We distinguish 11 classes of errors on the test set.

The classes of errors depicted in Table (3) indicates that the classes of errors and their frequencies are consistent across the different data sets. The third error class in Table (3) is of particular inter-

	Error Type	Test Set	Dev Set
1	Precision Errors: Verbal Negation Scope not found by NegFocus	37	23
2	Focus Mismatch: gold focus annotation is the neg. cue	138	112
3	Focus Mismatch: gold focus annotation is a constituent triggered by the *nsubj* dependency to the negated verb	44	16
4	Focus Mismatch: gold focus annotation is the constituent triggered by the *nsubjpass* dependency	7	12
5	Focus Mismatch: gold focus annotation is an adverb triggered by the *advmod* dependency with the verb, but is not adjacent to the verb	14	4
6	Partial Match: the spans of the gold focus annotation and NegFocus annotation overlap	6	8
7	Focus Mismatch: gold focus annotation is not contained within the NegFocus Syntactic Scope	4	5
8	NegFocus Syntactic Scope annotation error	10	9
9	Focus Mismatch: Miscellaneous errors	27	25
10	Focus Mismatch: gold focus annotation matches CLaC baseline heuristic, however another CLaC focus heuristic was chosen	3	3
11	Focus Mismatch: gold focus annotation contains two discontinuous focus annotation spans	17	11
	TOTAL	307	228

Figure 3: System Errors

est to us, as it highlights the different interpretations of verbal negation scope. NegFocus will not include the noun subject in the syntactic negation scope, and therefore the noun subject constituent is never a focus candidate as required in Example (11).

(11) *In New York, [a spokesman for American Brands] would n't <u>comment</u>.*

Similarly, the seventh error class in Table (3) contains focus annotations that are not contained in NegFocus negation scopes. Example (12) shows an error where the sentence begins with a prepositional phrase that is annotated as the gold focus.

(12) *[On some days], the Nucor plant does n't produce <u>anything</u>.*

We disagree with the gold annotations on this and similar cases: the prepositional phrase *on some days* is not negated, it provides a temporal specification for the negated statement *the Nucor plant produces something* and in our opinion, the negation negates *something*, contrasting it with

(13) *[On some days], the Nucor plant does n't produce <u>a lot</u>.*

which allows for some production, which indicates to us that without context information, low focus is warranted here.

NegFocus incorporates a focus heuristic for determining the passive noun subject constituent as the focus of the negation, however only in cases where the negated verb has an empty complement. The fourth error class contains errors in focus determination where this heuristic fails and where the passive subject is the gold focus despite the complement of the negated verb not being empty, requiring further analysis:

(14) *To simplify the calculations , [commissions on the option and underlying stock] are n't included <u>in the table</u>.*

NegFocus determines an adverb directly preceding the verb trigger as the focus of the negation, but, as described in the fifth error class, the gold focus annotations in a few cases determine adverbs to be the focus of the negation even when they don't directly precede the verb, but are linked by the *advmod* relation, as in Example (15). When we experimented with relaxing the adjacency constraint, re-

	Error Type	Test Set	Dev Set
1	NegFocus annotation is adverb	2	3
2	NegFocus annotation is passive noun subject	7	4
3	NegFocus Scope Error	7	14
4	NegFocus baseline heuristic at variance with gold annotation	122	91
	TOTAL	138	112

Figure 4: Negation cue annotation misses

sults suffered. This, too, is an area where we wish to investigate whether any general patterns are possible and what additional resources they require to be reliable.

(15) *" The intervention has been friendly, meaning that they [really] did n't have to do it, " said Maria Fiorini Ramirez, money-market economist at Drexel Burnham Lambert Inc .*

The majority of NegFocus errors occur in the second error class. Table (4) further analyzes the second error class, where the gold annotation puts the negation trigger in the focus but NegFocus finds another focus (usually in the verb complement).

The gold standard annotations place the focus of the negation of verb *v* on the negation trigger if it cannot be inferred that an action *v* occurred (Blanco and Moldovan, 2011). NegFocus will only make this assumption when the verb complement constituent is empty, otherwise the baseline focus heuristic will be triggered, as depicted in Example (16).

(16) *AMR declined to comment , and Mr. Trump did [n't] respond to requests for interviews.*

Furthermore, the CLaC system will choose to trigger the subject passive focus heuristic in the case where the verb complement constituent is empty, and the passive noun subject is present. In contrast, the gold standard annotations do not necessarily follow this heuristic as seen in Example (17).

(17) *That is n't 51 %, and the claim is [n't] documented .*

Lastly, the gold focus annotations include focus spans which are discontinuous. NegFocus will only detect one continuous focus span within one instance of a verbal negation. The eleventh error class

includes those cases where NegFocus matches one of the gold focus spans but not the other as seen in Example (18).

(18) *[The payments] aren't expected [to have an impact on coming operating results], Linear added .*

These error cases show that more analysis of the data, but also of the very notion of focus, is necessary.

6 Conclusion

We conclude that this experiment confirmed the hypothesis that negation trigger detection, syntactic scope determination, and focus determination are usefully modelled as a pipeline of three simple modules that apply after standard text preprocessing and dependency parsing. Approximating focus from a principled, linguistic point of view proved to be a quick and robust exercise. Performance on development and test sets is nearly identical and in a range around 58% f-measure. While the annotation standards as well as our heuristics warrant revisiting, we believe that the value of the focus annotation will prove its value beyond negation. The challenge data provide a valuable resource in themselves, but we believe that their true value will be shown by using the derived notion of focus in downstream applications. For initial experiments, the simple NegFocus pipeline is a stable prototype.

References

E. Blanco and D. Moldovan. 2011. Semantic representation of negation using focus detection. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL-HLT 2011)*, Portland, OR.

W. Chapman, W. Bridewell, P. Hanbury, G.F. Cooper, and B. Buchanan. 2001. A simple algorithm for identifying negated findings and diseases in discharge summaries. *Journal of Biomedical Informatics*, 34(5):301-310.

H. Cunningham, D. Maynard, K. Bontcheva, V. Tablan, N. Aswani, I. Roberts, G. Gorrell, A. Funk, A. Roberts, D. Damljanovic, T. Heitz, M.A. Greenwood, H. Saggion, J. Petrak, Y. Li, and Wim P. 2011. *Text Processing with GATE (Version 6)*. GATE (April 15, 2011).

M. de Marneffe and C.D. Manning. 2006. Generating typed dependency parses from phrase structure parses. In *LREC*.

R. Farkas, V. Vincze, G.Móra, J. Csirik, and G.Szarvas. 2010. The conll-2010 shared task: Learning to detect hedges and their scope in natural language text. In *Proceedings of the Fourteenth Conference on Computational Natural Language Learning*.

C-H. Han and M. Romero. 2001. Negation, focus and alternative questions. In K. Megerdoomian and L.A. Bar-el, editors, *Proceedings of the West Coast Conference in Formal Linguistics XX*, Somerville, MA. Cascadilla Press.

Y. Huang and H.J. Lowe. 2007. A novel hybrid approach to automated negation detection in clinical radiology reports. *Journal of the American Medical Informatics Association : JAMIA*, 14(3):304-311.

R.D. Huddleston and G.K. Pullum. 2002. *The Cambridge grammar of the English language*. Cambridge University Press, Cambridge, UK; New York.

H. Kilicoglu and S. Bergler. 2011. Effective bio-event extraction using trigger words and syntactic dependencies. *Computational Intelligence*, 27(4):583–609.

J.-D. Kim, Y. Wang, T. Takagi, and A. Yonezawa. 2011. Overview of genia event task in bionlp shared task 2011. In *Proceedings of BioNLP Shared Task 2011 Workshop at ACL-HLT*.

D. Klein and C.D. Manning. 2003. Accurate unlexicalized parsing. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*.

R. Morante and C. Sporleder, editors. 2010. *NeSp-NLP '10: Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*, Stroudsburg, PA, USA. Association for Computational Linguistics.

R. Morante, S. Schrauwen, and W. Daelemans. 2011. Annotation of negation cues and their scope. guidelines v1.0. Technical report, CLiPS, University of Antwerp.

P. 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. *Journal of the American Medical Informatics Association : JAMIA*, 8(6):598-609.

B. Partee. 1993. On the "scope of negation" and polarity sensitivity. In E. Hajicova, editor, *Functional Approaches to Language Description*.

C. Poletto. 2008. The syntax of focus negation. *University of Venice Working Papers in Linguistics*, 18.

J. Wiebe, T. Wilson, and C. Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation*, 39(2-3).

M. Wiegand, B. Roth, D. Klakow, A. Balahur, and A. 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 (NeSp-NLP 2010)*.

Th. Wilson. 2008. *Fine-Grained Subjectivity Analysis*. Ph.D. thesis, University of Pittsburgh. Intelligent Systems Program.