

# Fine-Grained Focus for Pinpointing Positive Implicit Meaning from Negated Statements

Eduardo Blanco and Dan Moldovan

Lymba Corporation

Richardson, TX 75080 USA

{eduardo,moldovan}@lymba.com

## Abstract

Negated statements often carry positive implicit meaning. Regardless of the semantic representation one adopts, pinpointing the positive concepts within a negated statement is needed in order to encode the statement's meaning. In this paper, novel ideas to reveal positive implicit meaning using focus of negation are presented. The concept of granularity of focus is introduced and justified. New annotation and features to detect fine-grained focus are discussed and results reported.

## 1 Introduction

Semantic representation of text is an important step towards text understanding. Current approaches are based on relatively shallow representations and ignore pervasive linguistic phenomena such as negation and metaphor. Despite these weaknesses, shallow representations have been proven useful for several tasks, e.g., coreference resolution (Kong et al., 2009), machine translation (Wu and Fung, 2009).

Consider statement (1) *The company won't ship the new product to the United States until next year.* Existing approaches to represent the meaning of (1) either indicate that the verb *ship* is negated or disregard the negation altogether. Semantic role labelers trained over PropBank would link *n't* to *ship* with MNEG (i.e., negate the verb); any system based on FrameNet and more recent unsupervised proposals (Poon and Domingos, 2009; Liang et al., 2011; Titov and Klementiev, 2011) ignore negation.

In order to represent the meaning of (1), one must first ascertain that the negation mark *n't* is actually negating the TEMPORAL context linked to *ship* and

not the verb per se; more specifically, *n't* is negating exclusively the preposition *until*. Only doing so one can aim at representing the actual meaning of (1): *The company will ship the new product to the United States during next year.* Note that the verb *ship*, and its AGENT, THEME and LOCATION (i.e., *The company, the new product* and *to the United States*) are positive, as well as the temporal anchor *next year*.

Regardless of the semantic representation one favors (logic forms, predicate calculus, semantic relations, semantic frames, etc.), we argue that pinpointing the numerous words that contribute to implicit positive meanings within a negated statement is a required subtask to obtain it. This paper aims at extracting specific positive implicit meaning from negated statements. The main contributions are: (1) interpretation of negation using fine-grained focus; (2) fine-grained focus of negation annotation over a subset of PropBank; (3) feature set to detect fine-grained focus of negation; and (4) model to retrieve *precise* positive implicit meaning from negated statements.

## 2 Related Work

Negation has been widely studied from a theoretical point of view. The seminal work by Horn (1989) presents the main thoughts in philosophy and psychology. Work in linguistics has studied the interaction of negation with quantifiers and anaphora (Hintikka, 2002), as well as the role in reasoning (Sánchez Valencia, 1991; Dowty, 1994): one can perform downward (but not upward) monotone inference with negative statements. Zeijlstra (2007) analyzes the position and form of negative ele-

ments and negative concords; concepts such as intra and inter-domain negation and strength of negation (Ladusaw, 1996), syntactic and semantic negation (Löbner, 2000) have been discussed in the extensive literature, although we do not use them.

In computational linguistics, negation has mainly drawn attention in sentiment analysis (Wilson et al., 2009; Wiegand et al., 2010) and the biomedical domain. Recently, two events (Morante and Sporleder, 2010; Farkas et al., 2010) targeted negation mostly on those subfields. Among many others, Morante and Daelemans (2009) and Li et al. (2010) propose scope detectors using the BioScope corpus. Considering scope is indeed a step forward, but focus must also be taken into account to represent negated statements and detect their positive implicit meanings.

Regarding corpora, BioScope annotates negation marks and linguistic scopes exclusively on biomedical texts. It does not annotate focus and it purposely disregards negations such as *the reactions in NK3.3 cells are not always identical* (Vincze et al., 2008), which carry the kind of positive meaning we aim at extracting (the reactions in NK3.3 cells are sometimes identical). Recently, Morante et al. (2011) present scope annotation in two Conan Doyle works, but they dismiss focus and positive meaning extraction. As stated before, PropBank (Palmer et al., 2005) treats negation superficially and FrameNet (Baker et al., 1998) regrettably disregards negation.

Blanco and Moldovan (2011) introduce a semantic representation of negation using focus detection. They target verbal negation and work on top of PropBank, selecting as focus the role that corresponds to the focus of negation. Simply put, they propose that all roles but the one corresponding to the focus are actually positive. Their approach, however, has a major drawback: selecting the whole role often yields too coarse of a focus and the positive implicit meaning is not fully specified (Section 3.1).

**Focus-Sensitive Phenomena.** The literature uses the term *focus* for widely distinct phenomena; space permits only a cursory review. Within functional generative grammars, focus is defined as *what is being asserted about the topic* (Hajičová et al., 1995). The term is also used in pragmatics (Glanzberg, 2005), and in phonetics and phonology (Xu and Xu, 2005; Beaver et al., 2007).

In linguistics, *focus* is largely associated with the theory presented in Mats Rooth's dissertation (1985) and posterior publications (Rooth, 1992). He analyzes the effect of focus in diverse phenomena, e.g., questions and answers, reasons and counterfactuals, conversational implicature, bare remnant ellipsis. His alternative semantics (e.g., *they didn't order the right parts* implies that some alternative of the form *they ordered X* is true) (Rooth, 1997) was an inspiration for this work. However, Rooth does not discuss how to detect focus of negation or its granularity and only provides simple made-up examples.

### 3 Scope and Focus

Negation has both scope and focus and they are key to capture its meaning. Scope is the part of the meaning that is negated. Focus is that part of the scope that is most prominently or explicitly negated (Huddleston and Pullum, 2002). All elements whose individual falsity would make the negated statement strictly true belong to the scope. Focus is the element of the scope that is *intended* to be interpreted as false to make the overall negative true.

Consider (1) *We didn't get an offer for more than \$40* and its positive counterpart (2) *We got an offer for more than \$40*. The truth conditions of (2) are: (a) somebody got something; (b) we got; (c) an offer was gotten; and (d) the offer was for more than \$40. In order for (2) to be true, (a–d) have to be true. Conversely, the falsity of any of them is sufficient to make (1) true: (1) would be true if *nobody got anything, we didn't get, an offer wasn't gotten or the offer wasn't for more than \$40*. Thus, all four statements (a–d) are inside the scope of (1).

The focus is often more difficult to identify. Text understanding is needed and context plays an important role. The most probable focus for (1) is *more than*, which corresponds to the interpretation *we got an offer for \$40 or less*. Another possible focus is *for more than \$40*, which yields *we got an offer, but not for more than \$40*. A third possible focus is *an offer for more than \$40*, which yields *we got something, but not an offer for more than \$40*. Section 3.1 discusses coarse versus fine-grained focus.

Both scope and focus are primarily semantic, highly ambiguous and context-dependent. More examples can be found in Table 1 and 3, and (Huddleston and Pullum, 2002, Chap. 9).

No.	Statement	Interpretation
1	People don't <u>always</u> follow instructions.	People sometimes follow instructions.
2	The new group isn't doing <u>any better than the old one</u> .	The new group is doing equal or worse than the old one.
3	The first two games didn't finish <u>in the top 10</u> .	The first two games finished below the top 10.
4	They don't sell to as <u>many clients as Maryland Club</u> .	They sell to less clients than Maryland Club.
5	She said she is not going home <u>until The Word Series is over</u> .	She said she is going home when The Word Series is over.
6	People don't believe I <u>want to give this money away</u> .	People believe I want to keep this money.
7	I cannot see <u>how this news doesn't benefit them</u> .	I can see how this news benefits them.
8	I don't believe <u>in this business you can be totally laissez-faire because of the high degree of public interest</u> .	I believe in this business you can be only partially laissez-faire because of the high degree of public interest.

Table 1: Examples of negated statements and their interpretation using fine-grained focus (regular underline). Using coarse-grained focus (wavy underline) would yield a much more generic, less preferred interpretation.

### 3.1 Granularity of Focus

In this paper, we refer to the focus considered by Blanco and Moldovan (2011) as *coarse-grained* and indicate it with a wavy underline; we refer to the focus we work with as *fine-grained* and indicate it with a regular underline, e.g., *We didn't get an offer for more than \$40*. Whereas coarse-grained focus is restricted to include all words belonging to a verb argument (as per their definition and annotation, focus is the full text of a semantic role in PropBank), fine-grained focus is not. This allows us to narrow down the actual negative meaning and pinpoint more positive implicit meaning.

Considering fine-grained focus is a substantial step towards a comprehensive semantic representation of negation. Following with the example above, encoding that *we got something, but not an offer for more than \$40* (coarse-grained) is useful, but encoding *we got an offer for \$40 or less* (fine-grained) is preferred. Several examples of coarse versus fine-grained focus and the benefits of using the latter over the former are provided in Table 1. In all statements, using coarse-grained focus yields an interpretation with all words underlined with a wavy underline negative and the rest positive, e.g., statement (8) would be interpreted as *I believe in something because of the high degree of public interest, but not that in this business you can be totally laissez-faire*.

Selecting the elements that belong to the fine-grained focus is a difficult task. In example (1), both coarse and fine-grained foci are the same and yield the same interpretation. In the rest of examples and in the vast majority of negations we annotated (Section 4), fine-grained focus comprises fewer words and yields more specific interpretations.

The coarse-grained focus in statements (1, 2) is an adverbial phrase. In (1) coarse-grained focus is a single word and thus fine-grained focus is trivially that word. In statement (2), fine-grained focus allows us to keep the comparison between the *new* and *old group* in the interpretation.

Examples (3, 4) correspond to statements whose coarse-grained focus is a prepositional phrase. Simple rules based on part-of-speech tags are not suitable here, deep understanding of text is needed. The fine-grained focus in example (3) is the preposition, but that is not the case in (4). Fine-grained focus in these statements allows us to obtain more complete interpretations, namely spell out the location (metaphorically speaking) where the games ended in (3) and the quantity sold in (4).

Examples (5–8) correspond to statements whose coarse-grained focus is a subordinate clause. Note that a verb is contained in the coarse-grained focus in these examples. In statement (5), the fine-grained focus is the first word, a preposition. However, that is not the case in (8), where the MANNER of the verb within the subordinate clause (i.e., *totally*) is selected as fine-grained focus. In (6), the phrasal verb *give away* is the fine-grained focus. Statement (7) is specially interesting because it contains a double negation and fine-grained focus is the negation mark within the coarse-grained focus.

Note that interpreting statements using coarse-grained focus is by no means wrong, but it is not optimal. The interpretation using fine-grained focus entails the one using coarse-grained focus. For example, in (2), *The new group is doing equal or worse than the old one* (fine) entails *The new group is doing, but not any better than the old one* (coarse).

Node	# Negations	% Negations
NP	1,051	39.93
PP	570	21.65
ADVP	415	15.75
SBAR	323	12.30
S	202	7.67
ADJP	33	1.26
Other	38	1.43

Table 2: Syntactic nodes for coarse-grained focus.

## 4 Annotating Fine-Grained Focus

We have annotated fine-grained focus of negation on top of the coarse-grained focus annotated by Blanco and Moldovan (2011). In this paper, we concentrate on negations whose coarse-grained focus is a prepositional phrase (PP), adverbial phrase (ADVP) or subordinate clause (SBAR). Excluding cases in which the verb is the coarse-grained focus, these syntactic nodes correspond to 49.70% of negations (Table 2). When a verb is the coarse-grained focus, it is not advantageous to consider fine-grained focus because both of them are always the same. e.g., *We urge our citizens not to wait until it is too late* [interpretation: we urge out citizens to *act*]. An example of NP being coarse-grained focus is *They realized they didn't order the right parts*.

We chose PP, ADVP and SBAR over noun phrases (NP, the most common syntactic realization) because they offer a variety of lexical and syntactic realizations, and thus allow us to tackle the task of fine-grained focus prediction in an assortment of constructions (as opposed to target exclusively NP). As we shall see, ADVP are shorter and easier, whereas PP and SBAR often contain complex syntactic (and semantic) structures and are tougher.

Annotation is done at the word level. Each word belonging to the coarse-grained focus is marked if it also belongs to the fine-grained focus. This allows us to narrow down the actual negative meaning and reveal the most positive implicit meaning. In some cases (32%, Table 4), coarse and fine-grained foci include the same words (e.g., *It doesn't always hurt* [interpretation: *it hurts sometimes*]). However, fine-grained focus usually (68%) comprises fewer words.

Annotators were first trained with examples similar to the ones in Table 1. In a first round, they were asked to select as fine-grained focus the words within the coarse-grained focus that they be-

lieved were intended to be negated. These instructions were purposely vague to analyze disagreements and allow us to define detailed guidelines. Inter-annotator agreement (exact match) was 41%. This number is low, but the task is challenging and a mismatch of one token (potentially a noncontent word (*the, a, etc.*) or even a punctuation mark (comma, dash, etc.) is counted as disagreement.

Conflicts were resolved and their causes analyzed. In a second round, sentences were annotated following the improved guidelines (Section 4.1). In both rounds, annotators were presented with plain text; they did not have access to any other information.

### 4.1 Annotation Guidelines

We aim at annotating fine-grained focus in order to pinpoint the numerous positive concepts within a negated statement. All concepts but the ones belonging to the fine-grained focus should be interpreted positive. Our annotation criteria is succinctly summarized by the following principles:

1. We annotate fine-grained focus of negation to reveal specific positive implicit meaning; we do not strictly follow any theory of focus.
2. We assume that fine-grained focus is contained within the coarse-grained focus.
3. Decisions are made taking into account the current sentence and context. Context is limited to the previous and next sentence.
4. World knowledge is taken into account. Thus, sentences are fully interpreted to identify positive implicit meaning.
5. In case of ambiguity, we prioritize:
  - (a) fine-grained focus that yields novel meaning over foci yielding meaning already stated elsewhere;
  - (b) narrow over wide fine-grained focus. The narrower the focus, the more specific the positive meaning revealed.
  - (c) the fine-grained focus that reveals the most obvious positive implicit meaning, i.e., meaning requiring the least world knowledge and assumptions to hold.
6. If there are two options for fine-grained focus yielding semantically equivalent positive implicit meanings, we select the fine-grained focus occurring earlier within the sentence.

No.	Example
1	The plan indeed raises from 40% to 50% the number of freshmen applicants admitted strictly by academic criteria. But that doesn't mean "half of the students attending" will be admitted this way.
2	"[...] and tied it to the stake with a chain," he says proudly. "And you can't cut this chain with bolt cutters".
3	Although other parties have stated they have no complaints, it is not growing fast enough for us.
4	Mr. Katz happily agreed, sliding over the fact that California's roads and bridges aren't funded by property taxes but by state and federal gasoline taxes.
5	[...] in a criminal case, a prosecutor can not comment on a defendant's failure to testify [...].
6	You think you can go out and turn things around. The reason doesn't relate to your selling skills.
7	Respondents don't think that an economic slowdown would harm the major investment markets very much.
8	The first two games of the World Series between [...] didn't finish in the top 10 [...]

Table 3: Examples of annotation (and relevant context) exemplifying the annotation guidelines.

## 4.2 Examples of Annotation

In this section, we exemplify our annotation guidelines with the statements in Table 3. When example (1) is interpreted in context [criterion 3], we obtain *at most half of the students will be admitted strictly by academic criteria*. Word knowledge [criterion 4] allows us to determine that if 40–50% of students are admitted a certain way, at most half of students attending will be admitted this way (a student admitted may not enroll). Word knowledge is also used in example (2): however strong the chain is, one could cut it with a stronger tool than bolt cutters.

Statement (3) implicitly states that *it is growing fast enough for other parties*. Thus, we choose *enough* [interpretation: it is growing insufficiently fast for us] since it reveals novel positive meaning [criterion 5a]. Another option discarded is *us* [interpretation: it is growing fast enough for someone, but not us]. Note that revealing novel positive implicit meaning is not always possible, e.g., statement (4).

There are several options for statement (5): (5a) *a defendant's failure to testify* [interpretation: a prosecutor can comment, but not on a defendant's failure to testify]; (5b) *a defendant's* [a prosecutor can comment on somebody's failure to testify, but not the defendant's]; and (5c) *testify* [a prosecutor can comment on the defendant's failures to do something, but not to testify]. We prefer (5c) since it reveals the most specific positive meaning [criterion 5b]. Note that narrowing down the coarse-grained focus is not always possible as exemplified in example (6): one cannot tell if the reason relates to another skill or to something else (e.g., economy, weather).

In example (7), we choose the fine-grained focus that reveals the most obvious implicit positive mean-

	#FGF	%(CGF = FGF)	#FGF/#CGF
PP	5.53	1.17%	0.44
ADVP	1.38	89.19%	0.94
SBAR	9.79	14.79%	0.32
All	5.25	32.41%	0.57

Table 4: Numeric analysis: average number of words in fine-grained focus, percentage of negations in which coarse and fine-grained focus are the same and average ratio of words in fine versus coarse-grained focus.

ing [criterion 5c], *very much* [interpretation: an economic slowdown would harm the major investment markets *a little*]. Another option is *slowdown*, yielding the plausible but less felicitous interpretation *responders think that an economic recession/turmoil (but not a slowdown) would harm the major investment markets very much*. A third option is *major* [responders think that an economic slowdown would harm *minor* investment markets very much]. The last two options are plausible but less likely.

Finally, statement (8), there are two semantically equivalent options: (8a) *in* [interpretation: the games finished below the top 10] and (8b) *10* [interpretation: the games finished in the top X, where X is larger than 10]. We choose the former since it occurs earlier in the sentence [criterion 6].

## 4.3 Annotation Analysis

The three syntactic realizations of coarse-grained focus we aim at narrowing down have significantly different characteristics. Table 4 summarizes some basic numeric analysis. Intuitively, ADVPs are fairly easy (they are short and coarse-grained and fine-grained foci are often the same). On the other hand, PP and SBAR are longer and only 44% and 32% of words belonging to the coarse grained focus belong to the fine-grained focus respectively.

Baseline		P	R	F
COARSE	PP	1.96	1.89	1.92
	ADVP	92.86	92.86	92.86
	SBAR	15.38	13.33	14.29
	All	29.52	27.93	28.70
FIRST-WORD	PP	33.33	32.08	32.69
	ADVP	92.86	92.86	92.86
	SBAR	35.29	20.00	25.53
	All	51.04	44.14	47.34
FIRST-JJ	PP	29.82	32.08	30.91
	ADVP	92.86	92.86	92.86
	SBAR	15.38	13.33	14.29
	All	52.34	42.34	42.34
BASIC	PP	54.17	49.06	51.49
	ADVP	92.86	92.86	92.86
	SBAR	45.00	30.00	36.00
	All	<b>63.54</b>	<b>54.95</b>	<b>58.94</b>

Table 5: Precision, recall and f-measure of baselines.

## 5 Learning Fine-Grained Focus

We follow a standard supervised learning approach. Each token from each annotated negation becomes an instance. The decision to be made is whether or not an instance is part of the fine-grained focus. The annotated sentences (comprising several instances) were divided into training (70%), held-out (15%) and test (15%). The held-out portion was used to tune the feature set and results are reported for the test split only, i.e., using unseen instances.

Detecting fine-grained focus is similar to text chunking. Text chunking consists of dividing text into syntactically related nonoverlapping groups of words (Tjong Kim Sang and Buchholz, 2000). On the other hand, we aim at dividing the words within a negated statement into belonging or not belonging to the fine-grained focus. Our problem can be redefined as detecting one type of chunk indicating the fine grained focus (FGF). We use the standard BIO notation, in which the first element of a chunk is prefixed by B- (beginning) and other elements of the chunk are preceded by I- (inside). The label O is used to indicate tokens outside any FGF chunk.

**Baselines.** We have implemented four baselines to predict fine-grained focus from the elements within the coarse-grained focus:

- COARSE: select all words.
- FIRST-WORD: select the first word.
- FIRST-JJ: select the first adjective; if none is found, apply FIRST-WORD.

- BASIC: same as system in Section 5.2 but using features *POS-tag*, *word* and *coarse-chunk*.

Table 5 shows the performance of these baselines. All of them obtain the same performance for ADVPs, and BASIC yields the best results. FIRST-WORD successfully predicts fine-grained focus mostly in cases in which the fine-grained focus is a preposition positioned at the beginning of the coarse-grained focus (e.g., Table 3, statement 8).

### 5.1 Selecting Features

We use a mixture of features proposed for standard text chunking, semantic role labeling and novel features characterizing negation (Table 6). We only provide more details for the non-obvious ones.

Features 1–5 characterize the current token with an emphasis on negation. *Neg-prefix* indicates if a word is an adjective, starts with a negation prefix and the remainder of it is a valid adjective. We consider the following negation prefixes: *a-*, *an-*, *anti-*, *dis-*, *il-*, *im-*, *in-*, *ir-*, *non-* and *un-* and check whether the remainder is a valid adjective querying WordNet. This successfully allows us to detect *irrelevant* (prefix *ir-*; *relevant* is a valid adjective) and disregard adjectives that just happen to start with a negated prefix, e.g., *artistic*, *intelligent*. *Any-prefix* indicates if a word starts with *any* (e.g., *anytime*). Huddleston and Pullum (2002, p.823) refer to these words as “*any class of items*” and include them in the negatively-oriented polarity-sensitive items (NPIs). Features signaling other NPIs (until, dare, yet, etc.) did not bring an improvement on the development set. *Ly-suffix* typically signals an adverb indicating the manner in which something happened.

Features 6–18 describe the coarse-grained focus. *Coarse-path* corresponds to four features indicating paths of length 1–4 from *coarse-node* to the token. Including the full path did not yield an improvement on the development set. *Coarse-head* is calculated following (Collins, 1999).

Finally, features *coarse-verb* and *sem-role* are useful in cases in which the token is not only part of the semantic role corresponding to the coarse-grained focus (i.e., a role of verb *pred-word*), but also a role of a verb within the coarse-grained focus (i.e., a role of verb *coarse-verb*). For example in Table 3, example (7), for token *slowdown* we have *word* = *slowdown*, *pred-word* = *think*, *coarse-role* = *AI*, *coarse-verb* = *harm* and *sem-role* = *A0*.

No.	Feature	Values	Explanation
1–2	POS-tag and word	{NN, VBD, ...}	POS tag and text of current token
3	neg-prefix (PP, SBAR)	{yes, no}	does word start with a negation prefix?
4	any-prefix (PP, SBAR)	{yes, no}	does word start with prefix <i>-any</i> ?
5	ly-suffix (SBAR)	{yes, no}	does word end with suffix <i>-ly</i> ?
6–7	coarse-{node, parent}	{S, PP, ...}	syntactic node of coarse-grained focus and parent
8–9	coarse-{left, right}	{NP, VP, ...}	syntactic node of coarse-node left and right siblings
10	coarse-struct	{IN=NP, IN=S, ...}	syntactic nodes of of coarse-node daughters
11	coarse-length	$\mathbb{N}$	length of coarse-grained focus
12–15	coarse-path (PP, SBAR)	{PP, PP-NP, ...}	paths of length 1–4 from coarse-node to token
16	coarse-role	{ARG1, MTMP, ...}	semantic role of coarse-grained focus
17	coarse-head (PP, SBAR)	{clock, detail, ...}	head of coarse-grained focus
18	coarse-verb (SBAR)	{think, predict, ...}	first verb within coarse-grained focus
19	pred-word	{affected, go, ...}	predicate text
20	pred-POS	{VB, VBN, ...}	predicate POS tag
21	sem-role (SBAR)	{ARG1, MLOC, ...}	semantic role this token belongs to wrt coarse-verb
22	coarse-chunk	{B-CFG, I-CFG, O}	coarse-grained annotation using BIO

Table 6: Feature set used to predict fine-grained focus of negation. If a feature is especially useful for a particular syntactic node, we indicate so between parenthesis in the right hand side of column 1 (otherwise it is useful for all).

## 5.2 Experiments and Results

We have carried our experiments using Yamcha (Kudoh and Matsumoto, 2000), a generic, customizable, and open source text chunker<sup>1</sup> implemented using TinySVM<sup>2</sup>. Following Yamcha’s design, we distinguish between static and dynamic features. Static features are the ones depicted in Table 6 for a fixed size window. Dynamic features are the predicted classes for a fixed set of previous instances. Whereas values for static features are considered correct, values for dynamic features are predictions of previous instances and therefore may contain errors. Varying window size effectively varies the number of features considered, the larger the window the more local context is taken into account.

Window sizes are defined using ranges between instances. The instance to be predicted has index ‘0’, the previous one ‘-1’, the next one ‘1’, and so on. The range  $[i..j]$  indicates we take into account from the  $i$ th to the  $j$ th instances to predict the current instance. Ranges for dynamic features can only contain instances preceding the current one.

The best performing system was obtained using a window including the current and two previous instances, and taking into account dynamic features. This system uses a total of 68 features: 66 static features ( $22 \times 3 = 66$ , 22 features per instance, window contains 3 instances) and 2 dynamic features.

<sup>1</sup><http://chasen.org/taku/software/yamcha/>

<sup>2</sup><http://chasen.org/taku/software/TinySVM/>

Window Size		P	R	F
static	dynamic			
[-1..0]	none	59.20	66.67	62.71
	[-1..-1]	68.27	63.96	66.05
[-1..1]	none	66.04	63.06	64.52
	[-1..-1]	70.10	61.26	65.38
[0..1]	none	57.85	63.06	60.34
	[-1..-1]	63.92	55.86	59.62
[-2..0]	none	60.00	62.16	61.06
	[-2..-1]	<b>71.15</b>	<b>66.67</b>	<b>68.84</b>
[-2..2]	none	62.96	61.26	62.10
	[-2..-1]	68.42	58.56	63.11
[0..2]	none	60.00	59.46	59.73
	[-2..-1]	64.21	54.95	59.22
[-3..0]	none	55.65	62.16	58.72
	[-3..-1]	68.93	63.96	66.36
[-3..3]	none	62.62	60.36	61.47
	[-3..-1]	67.01	58.56	62.50
[0..3]	none	57.80	56.76	57.27
	[-3..-1]	64.13	53.15	58.13

Table 7: Results using different window sizes.

Table 7 provides results on the test split for several window sizes considering and not considering dynamic features. The best performing system obtains precision 71.15, recall 66.67 (f-measure 68.84). In general, windows encompassing the  $i$  previous instances (e.g.,  $[-2..0]$ ) perform better than windows encompassing the  $i$  next instances (e.g.,  $[0..2]$ ). Windows not considering the  $i$  next instances yield better performance when using dynamic features (i.e.,  $[-i..0]$  is superior to  $[-i..i]$ ). Also, including dy-

Phrase	P	R	F
PP	64.71	62.26	63.46
ADVP	92.86	92.86	92.86
SBAR	60.00	50.00	54.55
All	71.15	66.67	68.84

Table 8: Detailed results per phrase using the best window size of features (in bold in Table 7).

namic features is favorable for almost all window sizes (the only exceptions are [0..1] and [0..2] by a negligible margin). Larger and discontinuous windows (e.g., [-4..-3, -1..-1]) did not bring an improvement during development and were discarded.

Finally, we report detailed results for the best performing system in Table 8.

## 6 Limitations and Future Work

The work presented here effectively extracts specific positive implicit meaning from negated statements. We depict below some limitations and shortcomings that could be targeted as future work.

**Types of negation.** We only targeted verbal, clausal and analytic negation (Huddleston and Pullum, 2002). Analyzing other types (e.g., synthetic, non-verbal: *I ate nothing, Nobody liked the party*) is needed for a more comprehensive approach.

**All positive meanings.** Not all implicit positive meanings are always detected. For example, *If the payment isn't received by today, an eviction notice will be send out* [interpretation: If the payment is received after today, an eviction notice will be send out]. Our proposal fails to detect that if no payment is received, the notice will also be send. Allowing multiple fine-grained foci seems a valid solution.

**Fine-grained within coarse-grained.** In a few examples, interpreting a negated statement using fine-grained focus requires modifications in other parts of the sentence as well. For example, *That increase in the money supply would not have happened without the consent of the Federal Reserve*. The interpretation is *That increase would have happened with the consent of the Federal Reserve*. This is not wrong, but a better option is to remove the modal *would* in the positive interpretation: the increase *did happen (with the consent of the Federal Reserve)*.

**Overall Interpretation.** A complete semantic representation for a statement (not only the verbal negation within) may require the same concept with two polarities. Consider *[In the past]<sub>TEMPORAL</sub>,*

*[you]<sub>AGENT</sub> just wore an unknown brand and didn't [care]<sub>verb</sub>*. The verbal negation is correctly interpreted *now you care*, but *in the past* remains as is (i.e., positive) for the verb *wore* [interpretation: in the past you just wore an unknown brand]. Strictly speaking, this is not a limitation but something to take into account to obtain a semantic representation of the whole statement. Our proposal successfully retrieves positive implicit meaning.

## 7 Conclusions

In this paper, we have argued that negated statements often carry positive implicit meaning and that its detection is key in order to capture their semantics, regardless of the semantic representation one favors (e.g., predicate calculus, semantic relations).

We have introduced the concept of granularity of focus of negation. Going beyond previous work, considering *fine*-grained focus allows us to reveal *narrow* positive implicit meaning. In the majority of cases (68%, Table 4) we are able to detect more positive implicit meaning than previous work considering a *coarse*-grained focus. We do not impose any syntactic restriction on which parts of a sentence might belong to the fine-grained focus. The annotation was done selecting words without taking into account any syntactic or semantic information. This approach effectively marks only the words that should be negated, but arguably makes prediction more difficult since fine-grained focus often does not correspond to a single node in the syntactic tree.

We have approached the task of fine-grained focus detection as a chunking problem in which we predict one chunk, FGF. The best model obtains an f-measure of 68.84, calculated by considering exact matches between chunks. In other words, unless the model predicts as fine-grained focus exactly the actual fine-grained focus, it is considered wrong when calculating performance. We believe this is the honest way of evaluating performance, even though partial matches could be useful for an actual application. For example, in *The U.S.'s largest suppliers haven't been filling their quotas to the full extent* [interpretation: they have been fulfilling their quotas to a partial extent], if the model predicts *full* as the only word belonging to fine-grained focus we count it wrong even though it successfully detects the most important part of it, i.e., the adjective *full*.



## References

- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet Project. In *Proceedings of the 17th international conference on Computational Linguistics*, Montreal, Canada.
- David Beaver, Brady Z. Clark, Edward Flemming, T. Florian Jaeger, and Maria Wolters. 2007. When Semantics Meets Phonetics: Acoustical Studies of Second-Occurrence Focus. *Language*, 83(2):245–276.
- Eduardo Blanco and Dan 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*, pages 581–589, Portland, Oregon, USA, June. Association for Computational Linguistics.
- Michael Collins. 1999. Head-Driven Statistical Models for Natural Language Parsing. *University of Pennsylvania*.
- David Dowty. 1994. The Role of Negative Polarity and Concord Marking in Natural Language Reasoning. In *Proceedings of Semantics and Linguistics Theory (SALT) 4*, pages 114–144.
- Richárd Farkas, Veronika Vincze, György Móra, János Csirik, and György 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*, pages 1–12, Uppsala, Sweden, July. Association for Computational Linguistics.
- Michael Glanzberg. 2005. Focus: A Case Study on the semanticspragmatics Boundary. In Zoltan G. Szabo, editor, *Semantics versus Pragmatics*, chapter 3. Oxford University Press, USA, first edition edition, February.
- Eva Hajičová, Petr Sgall, and Hana Skoumalová. 1995. An automatic procedure for topic-focus identification. *Comput. Linguist.*, 21:81–94, March.
- Jaakko Hintikka. 2002. Negation in Logic and in Natural Language. *Linguistics and Philosophy*, 25(5/6).
- Laurence R. Horn. 1989. *A Natural History of Negation*. University Of Chicago Press, June.
- Rodney D. Huddleston and Geoffrey K. Pullum. 2002. *The Cambridge Grammar of the English Language*. Cambridge University Press, April.
- Fang Kong, GuoDong Zhou, and Qiaoming Zhu. 2009. Employing the Centering Theory in Pronoun Resolution from the Semantic Perspective. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 987–996, Singapore, August. Association for Computational Linguistics.
- Taku Kudoh and Yuji Matsumoto. 2000. Use of support vector learning for chunk identification. In *Proceedings of the 4th conference on Computational natural language learning (CoNLL-2000)*, ConLL '00, pages 142–144, Stroudsburg, PA, USA. Association for Computational Linguistics.
- William A. Ladusaw. 1996. Negation and polarity items. In Shalom Lappin, editor, *Handbook of Contemporary Semantic Theory*, pages 321–341. Blackwell.
- Junhui Li, Guodong Zhou, Hongling Wang, and Qiaoming Zhu. 2010. Learning the Scope of Negation via Shallow Semantic Parsing. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 671–679, Beijing, China, August. Coling 2010 Organizing Committee.
- Percy Liang, Michael Jordan, and Dan Klein. 2011. Learning Dependency-Based Compositional Semantics. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 590–599, Portland, Oregon, USA, June. Association for Computational Linguistics.
- Sebastian Löbner. 2000. Polarity in Natural Language: Predication, Quantification and Negation in Particular and Characterizing Sentences. *Linguistics and Philosophy*, 23(3):213–308–308, June.
- Roser Morante and Walter Daelemans. 2009. Learning the Scope of Hedge Cues in Biomedical Texts. In *Proceedings of the BioNLP 2009 Workshop*, pages 28–36, Boulder, Colorado, June. Association for Computational Linguistics.
- Roser Morante and Caroline Sporleder, editors. 2010. *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*. University of Antwerp, Uppsala, Sweden, July.
- Roser Morante, Sara Schrauwen, and Walter Daelemans. 2011. Annotation of negation cues and their scope. Guidelines v1.0. Technical report, CLiPS, University of Antwerp.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The Proposition Bank: An Annotated Corpus of Semantic Roles. *Computational Linguistics*, 31(1):71–106.
- Hoifung Poon and Pedro Domingos. 2009. Unsupervised Semantic Parsing. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 1–10, Singapore, August. Association for Computational Linguistics.
- Mats Rooth. 1985. *Association with Focus*. Ph.D. thesis, Uiveristy of Massachusetts, Amherst.
- Mats Rooth. 1992. A Theory of Focus Interpretation. *Natural Language Semantics*, 1:75–116.

- Mats Rooth. 1997. Focus. In Shalom Lappin, editor, *The Handbook of Contemporary Semantic Theory*, Blackwell Handbooks in Linguistics, chapter 10. Wiley-Blackwell, December.
- Victor Sánchez Valencia. 1991. *Studies on Natural Logic and Categorical Grammar*. Ph.D. thesis, University of Amsterdam.
- Ivan Titov and Alexandre Klementiev. 2011. A Bayesian Model for Unsupervised Semantic Parsing. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1445–1455, Portland, Oregon, USA, June. Association for Computational Linguistics.
- Erik F. Tjong Kim Sang and Sabine Buchholz. 2000. Introduction to the CoNLL-2000 Shared Task: Chunking. In Claire Cardie, Walter Daelemans, Claire Nedellec, and Erik Tjong Kim Sang, editors, *Proceedings of CoNLL-2000 and LLL-2000*, pages 127–132. Lisbon, Portugal.
- Veronika Vincze, Gyorgy Szarvas, Richard Farkas, Gyorgy Mora, and Janos Csirik. 2008. The BioScope corpus: biomedical texts annotated for uncertainty, negation and their scopes. *BMC Bioinformatics*, 9(Suppl 11):S9+.
- 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, Uppsala, Sweden, July. University of Antwerp.
- Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2009. Recognizing Contextual Polarity: An Exploration of Features for Phrase-Level Sentiment Analysis. *Computational Linguistics*, 35(3):399–433.
- Dekai Wu and Pascale Fung. 2009. Semantic Roles for SMT: A Hybrid Two-Pass Model. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Short Papers*, pages 13–16, Boulder, Colorado, June. Association for Computational Linguistics.
- Y. Xu and C. Xu. 2005. Phonetic realization of focus in English declarative intonation. *Journal of Phonetics*, 33(2):159–197, April.
- H. Zeijlstra. 2007. Negation in Natural Language: On the Form and Meaning of Negative Elements. *Language and Linguistics Compass*, 1(5):498–518.