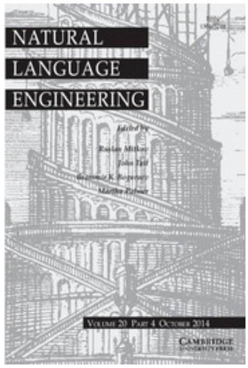

Retrieving implicit positive meaning from negated statements

EDUARDO BLANCO and DAN MOLDOVAN

Natural Language Engineering / Volume 20 / Issue 04 / October 2014, pp 501 - 535
DOI: 10.1017/S1351324913000041, Published online: 26 February 2013

Natural Language Engineering **20** (4): 501–535. © Cambridge University Press 2013
doi:10.1017/S1351324913000041

Retrieving implicit positive meaning from negated statements

EDUARDO BLANCO and DAN MOLDOVAN
Human Language Technology Research Institute, The University of Texas at Dallas
Richardson, TX 75080 USA
email: {eduardo,moldovan}@hlt.utdallas.edu

(*Received 27 September 2011; revised 14 January 2013; accepted 16 January 2013; first published online 26 February 2013*)

Abstract

This paper introduces a model for capturing the meaning of negated statements by identifying the negated concepts and revealing the implicit positive meanings. A negated sentence may be represented logically in different ways depending on what is the scope and focus of negation. The novel approach introduced here identifies the focus of negation and thus eliminates erroneous interpretations. Furthermore, negation is incorporated into a framework for composing semantic relations, proposed previously, yielding a richer semantic representation of text, including hidden inferences. Annotations of negation focus were performed over PropBank, and learning features were identified. The experimental results show that the models introduced here obtain a weighted f-measure of 0.641 for predicting the focus of negation and 78 percent accuracy for incorporating negation into composition of semantic relations.

1 Introduction

Capturing the meaning of text is key to text understanding and reasoning. These tasks could potentially improve the performance of several natural language processing applications and help tools requiring inferences, e.g. recognizing textual entailment. Even though philosophers and linguists have proposed several theories and models to represent the meaning of text, the state of the art is far from doing so automatically.

Negation is present in all human languages and it is always the case that statements are affirmative by default. Negation is marked and typically signals something unusual or an exception. It may be present in several units of language, e.g. words (*<u>in</u>credible*) and clauses (*He does<u>n't</u> have friends*). Negation and its correlates (truth values, lying, irony, false or contradictory statements) are exclusive characteristics of humans (Horn 1989; Horn and Kato 2000).

Negation is well understood in grammars (Quirk *et al.* 1985; Huddleston, Rodney and Pullum 2002) and the valid ways to express a negation are extensively documented. However, within computational linguistics, there has not been much work on detecting it, and more importantly, on representing its meaning. Semantic role labelers, e.g. systems trained over PropBank (Palmer, Gildea and Kingsbury 2005) or FrameNet (Baker *et al.* 1998), and participants in the SemEval tasks regarding relation detection, e.g. Hendrick *et al.* (2009), ignore the semantics of negation. Scope detectors are a step forward, but they are far from fully representing the meaning of negation.

At first glance, one would think that interpreting negation could be reduced to finding negative keywords, detecting their scope using syntactic analysis and reversing the polarity. Actually, it is far more problematic. Negation plays a crucial role in text understanding and poses considerable challenges.

Detecting the scope of negation in itself is challenging: *All vegetarians do not eat meat* means that vegetarians do not eat meat, and yet *All that glitters is not gold* means that it is not the case that all that glitters is gold (so out of all things that glitter, some are gold and some are not). In the former example, the quantifier *all* has scope over the negation; in the latter, the negation has scope over *all.*

In classical logic, two negatives always cancel each other out. On the other hand, in language this is not always the case: *she is <u>not</u> <u>un</u>happy* does not mean that *she is happy*; it means that *she is not fully unhappy, but she is not happy either*.

Some negated statements carry implicit positive meaning. For example, *cows do <u>not</u> eat meat* implicitly states that *cows eat something other than meat*. Otherwise, the speaker would have stated *cows do not eat*. A clearer example is the correct and yet puzzling statement *tables do <u>not</u> eat meat*. This sentence sounds unnatural because of the underlying positive meaning *tables eat something other than meat*.

Negation can express *less than* or *in between* when used in a scalar context. For example, *John does <u>not</u> have three children* probably means that he has either one or two children (as opposed to no children at all). Contrasts may use negation to disagree about a statement and not to negate it, e.g. *That place is <u>not</u> big, it is massive* defines the place as *massive*, and therefore, *big*.

This paper introduces a model that thoroughly interprets negation by surfacing implicit positive meaning. The rationale behind this relies on detecting the focus of negation. The main contributions are: (a) interpretation of negation using focus detection; incorporation of negation to (b) semantic relations and (c) composition of semantic relations (CSR) and (d) focus of negation annotation over all PropBank verbal negations. We also report experimental results on focus detection and composing semantic relations when negation is present. The work presented here builds on our first proposal to semantically represent negation (Blanco and Moldovan 2011b).

1.1 Negation in natural language

Unlike affirmative statements, negation is marked by words (e.g. *not*, *no*, *never*) or affixes (e.g. -*n't*, *un*-, *dis*-). Negation interacts with other words in special ways. For example, negated clauses use different connective adjuncts than positive clauses do: *neither*, *nor* instead of *either*, *or*. The so-called *negatively oriented polarity-sensitive items* (Huddleston and Pullum 2002) include, among many others, words starting with *any*- (*anybody*, *anyone*, *anywhere* etc.), the modal auxiliaries *dare* and *need* and the grammatical units *at all*, *much* and *till*. Negation in verbs usually requires an auxiliary; if none is present, the auxiliary *do* is inserted (*I read the paper* vs. *I didn't read the paper*).

Huddleston and Pullum (2002) distinguish four contrasts for negation. The first two have to do with the expression of negation, while the second two have to do with meaning or interpretation:

Verbal or nonverbal. Verbal if the marker of negation is grammatically associated with the verb (e.g. *I did not see anything at all*); nonverbal if it is associated with a dependent of the verb (e.g. *I saw nothing at all*).

Analytic or synthetic. Analytic if the sole function of the negated mark is to mark negation (e.g. *Bill did not go*); synthetic if it has some other function as well (e.g. *[Nobody]*$_{\text{AGENT}}$ *went to the meeting*).

Clausal or subclausal. Clausal if the negation yields a negative clause (e.g. *She didn't have a large income*); subclausal otherwise (e.g. *She had a not inconsiderable income*).

Ordinary or metalinguistic. Ordinary if it indicates that something is not the case, e.g. (a) *She didn't have lunch with my old man: he couldn't make it*; metalinguistic if it does not dispute the truth but rather reformulates a statement, e.g. (b) *She didn't have lunch with your 'old man': she had lunch with your father*. Note that in (a) the lunch never took place, whereas in (b) a lunch did take place.

In this paper we focus on verbal, analytic, clausal and both metalinguistic and ordinary negation.

2 Previous work

Negation has been widely studied outside of computational linguistics. In traditional logic, it is usually the simplest unary operator and reverses the truth value. The seminal work on negation by Horn (1989) presents the main thoughts in philosophy and psychology. We follow him in the next two paragraphs.

Two of the most basic philosophical laws put forth by Aristotle are the Law of Contradiction (LC, it is impossible to be and not be at the same time) and the Law of Excluded Middle (LEM, in every case we must either affirm or deny). The Law of Excluded Middle is not always applicable to statements involving negation (e.g. one can deny *being cold* and *not being cold*). Philosophers realized that a negative statement can have latent positive meaning, e.g. *Socrates is not well* presupposes that *Socrates is alive*. They differentiate between contradictory and contrary statements: two statements are contradictories if one or the other must be true, but both cannot be true or false at the same time; they are contraries if only one can be true, but both can be false at the same time. For example, *The plant is alive* and *The plant is dead* are contradictories; *The flower is white* and *The flower is blue* are contraries.

Psychology researchers have studied the constructs, usage and cognitive processing of negation. They note that negated statements are not on equal footing with positive statements; they are a different and subordinate kind of statements. Interestingly, evidence suggests that children learn first to communicate using positive statements and acquire negation later (Horn 1989). Psychology also confirms the intuitive thought that humans normally communicate in positive terms and reserve negation to describe unusual or unexpected situations (Boucher and Osgood 1969).

2.1 Linguistics

Linguists have found negation a highly complex phenomenon. *The Cambridge Grammar of the English Language* (Huddleston and Pullum 2002) dedicates over sixty pages to negation, covering scope and focus, verbal and non-verbal negation, polarity items (e.g. *already*, *any*) and multiple negation. This paper borrows from them the descriptions of negation types, definitions of scope and focus and several examples. Huddleston and Pullum (p. 798) also point out that the choice of focus reveals positive implicatures.

A Comprehensive Grammar of the English Language (Quirk *et al.* 1985) covers negation in different sections of the book. Negative forms, scope and focus are covered in Sections 2.54–2.56, 10.64 and 10.65. Although very briefly, Quirk *et al.* (p. 86) provide an example of focus and implicated meaning: *I don't drink instant coffee* (focus is *instant*) should be interpreted as *I drink some coffee, but not instant*. Negative pronouns are also detailed (Section 6.62), e.g. *None of the students have failed*, *Neither accusation is true*, as well as words negative in form and meaning (e.g. *not*, *no*), words negative in meaning but not in form (e.g. *seldom*, *little*), local and double negation and negative intensification (Sections 10.54–10.70).

Rooth (1985) presented a theory of focus in his dissertation, *Association with Focus* and posterior publications, e.g. Rooth (1992). He proposes an *alternative semantics* (i.e. semantics of alternatives) and studies not only the focus of negated statements but also the focus of affirmative ones and questions. His alternative semantics (e.g. *they didn't order the right parts* implies that some alternative of the form *they ordered X* is true) was an inspiration for this work. In this paper we follow the insights on scope and focus of negation by Huddleston and Pullum (2002) rather than Quirk *et al.* (1985) and Rooth (1985).

Among others, negation interacts with quantifiers and anaphora (Hintikka 2002). For example, the reference to *They* is easier to solve in (a) than in (b): (a) *Some of the students passed the exam. They must have studied hard*; (b) *Not all the students failed the examination. They must have studied hard*. Negation also influences reasoning (Sánchez Valencia 1991; Dowty 1994). In the simplest scenarios (Dowty 1994), one can perform upward (but not downward) monotone inference with positive statements (e.g. *She has a bulldog, bulldog is a dog; therefore she has a dog*), and downward (but not upward) monotone inference with negative statements (e.g. *She didn't give him a flower; a rose is a flower; therefore, she didn't give him a rose*).

Ladusaw (1996) analyzes natural language negation (intra-domain and inter-domain negation, strengths of negation etc.), focusing on clausal negation, polarity items and negative concord (i.e. multiple negation). Zeijlstra (2007) investigates the way different languages position and form negative elements (basically negative verbs for clausal negation or negative particles for sentential negation), as well as the interpretation of negative concord. Sandu (1994) extends first-order logic to represent negation, allowing for *strong or dual* and *weak or contradictory* negation. Löbner (2000) studies the nature of polarity in natural languages. He formally defines the concept of falsity (p. 224) and differentiates between *syntactic* (when a lexically negative element is added to the sentence or is substituted for a positive element) and *semantic* negation (i.e. polarity counterparts, A and B are polarity counterparts of each other iff: A is true iff B is false).

2.2 Computational linguistics

Within natural language processing, negation has drawn attention mainly in the biomedical domain and sentiment analysis. Recently, two events have dealt directly with negation, specifically scope detection. Most contributions to the Negation and Speculation in Natural Language Processing workshop (Morante and Sporleder 2010) focused in the aforementioned subfields. The CoNLL-2010 Shared Task (Farkas *et al.* 2010) aimed at the detection of hedges and their scope.

There have been several proposals to detect the scope of negation, most of them are within the biomedical domain and use the BioScope corpus. The first approach to model the problem as a classification task was made by Morante, Liekens and Daelemans (2008). They present a supervised scope detector in two phases: hedge cue identification, and scope finding. They use standard machine learning algorithms and a very specialized feature set for scope detection. Özgür and Radev (2009) present a similar two-phase approach. They use more sophisticated features for detecting keywords, but decide scopes based upon manually defined syntactic rules. Øvrelid, Velldal and Oepen (2010) propose another two-step scope detector for uncertainty: first, they use a maximum entropy classifier for identifying cue words; second, they define hand-crafted rules over dependency representations to decide their scope. Councill, McDonald and Velikovich (2010) present a supervised scope detector using their own corpus. They manually annotated the scope of negations contained in product reviews and apply their scope detector to sentiment classification. Li *et al.* (2010) approach scope detection from a novel perspective. Using ideas from traditional semantic role labeling, they treat negation clues as predicates and their scopes as predicate arguments. Their models outperform scope detectors based on chunking techniques.

Negation has been incorporated in systems performing sentiment and opinion analysis. The main goal of these efforts is to improve sentiment classification and related tasks by introducing some basic treatment of negation. They do not consider focus or implicit positive meaning, and their approaches to scope resolution use heuristics. Wiegand *et al.* (2010) survey the role of negation in sentiment analysis. Pang, Lee and Vaithyanathan (2002) mark as negated all words between a negation word (*not*, *isn't*, *didn't* etc.) and the first punctuation mark following that word. Hu and Liu (2004) flip the opinion orientation of words within a five-word window of a negation word. Jia, Yu and Meng (2009) propose several heuristics and exceptions to determine scope of negation for sentiment analysis.

Some natural language processing applications deal indirectly with negation, e.g. machine translation (van Munster 1988) and text classification (Rose *et al.* 2003). These applications only treat negation to improve the performance of a particular application and do not offer a model to semantically represent it. Boxer (2008), an off-the-shelf semantic parser, accounts for negation and detects its scope, but disregards focus.

A phenomenon related to negation is factuality, that is, the quality of being actual or based on a fact. Computational approaches to event factuality aim at recognizing whether events are presented as corresponding to real situations in the world, situations that have not happened or situations of uncertain status (Saurí and Pustejovsky 2008). Determining the factuality of events has been used for recognizing textual entailment (Saurí and Pustejovsky 2007).

Regarding corpora, the BioScope corpus (Szarvas *et al.* 2008) annotates negation marks and linguistic scopes exclusively on biomedical texts. BioScope does not annotate focus of negation and purposely ignores negations such as *the reactions in NK3.3 cells are not always identical* (Szarvas *et al.* 2008), which carry the kind of positive meaning this work aims at extracting (the reactions in NK3.3 cells *are sometimes* identical). Morante (2010) depicts twenty-nine common negation clues found in BioScope and discusses their ambiguity.

Even though scope detectors trained using BioScope achieve a relatively high performance, it remains unknown how well they will perform in other domains. There is an ongoing effort that will allow the community to answer this question soon enough. Scope of negation is being annotated over the novel *The Hound of the Baskervilles* by Conan Doyle (2011). Authors point out several interesting phenomena that greatly influence negation and are not found in BioScope due to its limited domain. For example, '*negation words in exclamative particles do not have a negation function, e.g. Don't tell that it is our friend Sir Henry!*'.

PropBank (Palmer *et al.* 2005) adds a layer of predicate-argument information, or semantic role labels, to the syntactic structures of the Penn Treebank. When it comes to negation, the only information provided for a negation mark is the verb it attaches to, labeled with MNEG. Predicting this label is rather simple, the best performing system (Koomen *et al.* 2005) participating in CoNLL-2005 Shared Task (Carreras and Màrquez 2005) achieved an f-measure of 97.61. MNEG, though, does not help much capturing the meaning of verbal negation (Section 3.1).

FrameNet (Baker *et al.* 1998) is based on frame semantics (Fillmore 1976). It provides a set of frames, their participants (*frame elements* using their jargon) and tokens that can instantiate the frame. The resource has been widely used and proven useful to encode the meaning of text, but it ignores negation.

FactBank (Saurí and Pustejovsky 2009) annotates degrees of factuality for event mentions. It considers several factuality values: *certainly, probably or possibly positive*, and *certainly, probably or possibly negative*. In addition, an event can be labeled *underspecified* (Saurí and Pustejovsky 2008).

None of the above references aim at detecting or annotating the focus of negation in natural language. Neither do they aim at carefully representing the meaning of negated statements nor extracting implicit positive meaning from them. To the best of our knowledge, the work described in this paper is the first attempt to automatically tackle these problems within computational linguistics.

3 The semantics of negation

Simply put, negation is a process that turns part of a statement into its opposite. However, pinpointing what is negated and what is implicitly positive is challenging.

3.1 Negation and implicit positive meaning

State-of-the-art semantic role labelers, e.g. the ones trained over PropBank (Kingsbury, Palmer and Marcus 2002) or NomBank (Meyers *et al.* 2004), do not accurately represent the meaning of negated statements. Given *John didn't build a house to impress Mary*, they encode AGENT(*John*, *build*) & THEME(*a house*, *build*) & PURPOSE(*to impress Mary*, *build*) & NEGATION(*n't*, *build*). This representation corresponds to the interpretation *it is not the case that John built a house to impress Mary*, ignoring that it is implicitly stated that *John built a house, but not to impress Mary*. Another option to interpret the above representation is (a) *it is not the case that John built,* (b) *it is not the case that a house was built and* (c) *it is not the case that a building took place to impress Mary*. Note that both options are not logically equivalent and none detect any positive meaning.

Several examples of negated statements carrying implicit positive meaning are shown in Table 1. For all statements *s*, role labelers would only encode *it is not the case that* s. However, the negations in examples (1–8) carry positive meaning underneath the direct meaning:

- Examples (1–4) are simple, in both cases humans would skip the last prepositional phrase unless they intend to convey the proposed positive meaning.
- Regarding (5), encoding that the UFO files *were released in 2008* is crucial to fully interpret the statement.
- Statements (7–9) show that different verb arguments modify the interpretation and even signal the existence of positive meaning.
- Examples (6 and 10) further illustrate the difficulty of the task; they are very similar (both have AGENT, THEME and MANNER) and their interpretation is altogether different.

Note that the negations in statements (9 and 10) do not carry any positive meaning, but the statements as a whole do, e.g. statement (9) implies *he has a new job*. Even though the interpretations of statements (9 and 10) do not contain a verbal negation, the meaning remains negative. Some examples could be interpreted differently depending on the context (Section 5.1). For example, if the next sentence to example (1) from Table 1 were *He built a castle*, we would have *John didn't build a house to impress Mary* (i.e. *John built something to impress Mary, but not a house*);

Table 1. *Examples of negated statements and their interpretation considering underlying positive meaning. Wavy underlines indicate the focus of negation (Section 3.2); the negations in examples (9 and 10) do not carry any positive meaning*

No.	Statement and interpretation
1.	John didn't build a house to impress Mary. – John built a house for another purpose.
2.	The cow didn't eat grass with a fork. – The cow ate grass, but not with a fork.
3.	I don't have a watch with me. – I have a watch, but it is not with me.
4.	We don't have an evacuation plan for flooding. – We have an evacuation plan for something else (e.g. fire).
5.	They didn't release the UFO files until 2008. – They released the UFO files in 2008.
6.	John doesn't know exactly how they met. – John knows how they met, but not exactly.
7.	His new job doesn't require driving. – His new job has requirements, but not driving.
8.	His new job doesn't require driving yet. – His new job requires driving in the future.
9.	His new job doesn't require anything. – His new job has no requirements.
10.	A panic on Wall Street doesn't exactly inspire confidence. – A panic on Wall Street discourages confidence.

if it were *He acquired a mansion*, we would have *John didn't build a house to impress Mary* (i.e. *a building did not take place*).

3.2 Scope and focus of negation

Negation has both scope and focus and these are extremely important to capture its semantics. 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).

The two concepts are interconnected. Scope refers to all elements whose individual falsity would make the negated statement strictly true. Focus is the element of the scope that is *intended* to be interpreted as false to make the overall negative true. Huddleston and Pullum (2002) provide the following description for focus:

> In all but the most trivial negative clauses there are several different conditions whose failure to hold would cause the clause to be strictly true. Which condition is intended can be indicated by a speaker through the device of stressing the most closely associated word. A constituent marked by stress as being crucial to the way in which an instance of negation should be understood is called the focus of that negation.

Consider statement (1) and its positive counterpart (2):

1. Cows don't eat meat.
2. Cows eat meat.

The truth conditions of statement (2) are: (a) somebody eats something; (b) cows are the ones who eat; and (c) meat is what is eaten. In order for statement (2) to be true, conditions (a–c) have to be true.

On the other hand, the falsity of any of the conditions (a–c) is sufficient to make statement (1) true. In other words, statement (1) would be true if *nobody eats*, *cows don't eat* or *meat is not eaten*. Therefore, all three conditions (a–c) are inside the scope of statement (1).

The focus is usually more difficult to identify, especially without knowing stress or intonation. Text understanding is often needed and context plays a key role. The most probable focus for statement (1) is *meat*, which corresponds to the interpretation *cows eat something else than meat*. Another possible but not likely focus is *cows*, which yields *someone eats meat, but not cows*.

Both scope and focus are primarily semantic, highly ambiguous and context-dependent. More examples can be found in Tables 1 and 5, and Huddleston and Pullum (2002). In this paper, a wavy underline indicates the focus of negation.

4 Semantic representation of negation

Negation does not stand on its own. To be useful, it should be added as part of another existing semantic representation. In this section we present a model to incorporate negation into semantic relations (Section 4.1). We also extend a framework to compose semantic relations with negation (Section 4.2).

4.1 A model to incorporate negation into semantic relations

Semantic relations capture connections between concepts and label them according to their nature. We denote a semantic relation R holding between two concepts x and y as R(x, y). R(x, y) could be read '*x is* R *of y*'. For example, AGENT(*John*, *bought*) encodes *John is the* AGENT *of bought*. The semantic representation of text via relations and a specific set of relations were presented in Blanco and Moldovan (2011c), but our treatment of negation is not tied to any relation inventory.

Our model to incorporate negation into semantic relations has the following characteristics:

- It aims at detecting implicit positive meaning from negated statements.
- It aims at selecting the smallest number of negative concepts in order to retrieve the largest amount of implicit positive meaning.
- It encodes the meaning of negated statements in context. The goal is to obtain the meaning of text taking into account negation, not to obtain a detailed representation of negated statements in isolation.
- It is inspired by theoretical works on negation, but does not strictly follow any of them.

Table 2. *Possible semantic representations for statement (1),* The cow didn't eat grass with a fork. *Options (ii–iv) reveal implicit positive meaning within the verbal negation*

No.	Representation				
i	AGENT(*the cow*, ∼*ate*)	&	THEME(*grass*, ∼*ate*)	&	INSTRUMENT(*with a fork*, ∼*ate*)
ii	AGENT(∼*the cow*, *ate*)	&	THEME(*grass*, *ate*)	&	INSTRUMENT(*with a fork*, *ate*)
iii	AGENT(*the cow*, *ate*)	&	THEME(∼*grass*, *ate*)	&	INSTRUMENT(*with a fork*, *ate*)
iv	AGENT(*the cow*, *ate*)	&	THEME(*grass*, *ate*)	&	INSTRUMENT(∼*with a fork*, *ate*)

We propose to incorporate the symbol '∼' to indicate an argument of a relation that must be negated. R(∼x, y) is interpreted *[not x] is* R *of y*, and R(x, ∼y) is interpreted *x is* R *of [not y]*. This way, the semantic representation for a negated statement explicitly states which concepts are positive and which are negative.

We found that applying '∼' to the first or second argument of a relation is enough for representing verbal, analytical and clausal negation:

- R(∼x, ∼y) would encode *[not x] is* R *of [not y]*. This is rather useless since knowing that a relation R exists between two concepts and that those concepts are not x and y is vague. For example, AGENT(∼*Mary*, ∼*run*) would encode *someone (but not Mary) is the* AGENT *of something (but not run)*.
- ∼R(x, y) would encode *x is [not* R*] of y*, i.e. there is a semantic relation between x and y, and that relation is not R. Consider, for example, *Bob did not steal the bicycle*. In a usual first-order logic, this sentence may be represented as ∼STEAL(*Bob*, *bicycle*). However, our representation is AGENT(*Bob*, ∼*steal*) & THEME(*bicycle*, ∼*steal*), which falls under the negated argument case.
- ∼R(∼x, y) and ∼R(x, ∼y). These two cases are also not possible in our representation, they fall under the negated argument case.

Given statement (1) *The cow didn't eat grass with a fork*, the state of the art encodes AGENT(*the cow*, *eat*) & THEME(*grass*, *eat*) & INSTRUMENT(*with a fork*, *eat*) & NEGATION(*n't*, *eat*). This representation fails to detect implicit positive meaning and only differs on the last relation from the positive counterpart. Its interpretation is *it is not the case that the cow ate grass with a fork*.

Several options arise to thoroughly represent statement (1). First, we find it useful to consider the semantic representation of the affirmative counterpart, (1') *The cow ate grass with a fork*: (1'a) AGENT(*the cow*, *ate*) & (1'b) THEME(*grass*, *ate*) & (1'c) INSTRUMENT(*with a fork*, *ate*).

Table 2 depicts four semantic representations for statement (1). Option (i) negates the verb, and it corresponds to the interpretation *the cow is the* AGENT *of [not eat]*, *grass is the* THEME *of [not eat]* and *a fork is the* INSTRUMENT *of [not eat]*. This option, like typical semantic roles, does not detect positive meaning and does not encode that an *eating* event took place. Options (ii–iv) negate the first argument of a single relation and their interpretations surface implicit positive meaning. Henceforth, 'negates the R' must be read *negates the first argument of* R*(x, y)*:

- Option (ii) negates the AGENT, it encodes *grass was eaten with a fork, but not by a cow.*
- Option (iii) negates the THEME, it corresponds to *the cow ate with a fork, but not grass.*
- Option (iv) negates the INSTRUMENT, encoding *the cow ate grass, but not with a fork.*

Option (iv) is preferred since it captures the best implicit positive meaning. It corresponds to the semantic representation of the affirmative counterpart (1'a–c) after negating the argument corresponding to the focus of the negation (i.e. *with a fork*). This fact and the examples in Tables 1 and 5 justify and motivate the importance of considering the focus of negation.

Consider statement (2), *A panic in Wall Street doesn't exactly inspire confidence.* The proposed representation is (2a) AGENT(*A panic in Wall Street*, ∼*inspire*) & (2b) MANNER(*exactly*, ∼*inspire*) & (2c) THEME(*confidence*, ∼*inspire*). This example does not contain implicit positive meaning, (2a–c) simply negate the verb.

4.1.1 Wide focus and multiple foci

It is out of the scope of this paper to review the extensive literature in linguistics about focus-sensitive phenomena, but we provide examples to show that our model can cope with wide focus and multiple foci (Rooth 1985). Simply put, our proposal selects an argument of one semantic relation for representing verbal, analytical and clausal negation. Similar examples are discussed by Jackendoff (1972).

Example 1. Consider the following sentences:

1. John didn't build a house to impress Mary.
2. He bought a mansion.

Recall that our model aims at extracting the meaning of text taking into account negation, not to represent negated statements in isolation. Therefore, we aim at obtaining a representation encoding what a human would understand after reading (1, 2), i.e. *John bought a mansion with the purpose of impressing Mary*. Note that the meaning of a negated statement in isolation is ambiguous in general and can only be determined when context is taken into account (see examples 2 and 3).

Following our proposal, we select as focus of (1) *build* and obtain the following representations:

- (1a) AGENT(*John*, ∼*build*) & (1b) THEME(*a house*, ∼*build*) & (1c) PURPOSE(*to impress Mary*, ∼*build*)
- (2a) AGENT(*He*, *bought*) & (2b) THEME(*a mansion*, *bought*)

Relations (1a–c) encode that no building took place. After resolving that *He* in (2) refers to *John* in (1) relations (2a–b) encode that *John bought a mansion.* We believe that detecting that the purpose of buying the mansion was *to impress Mary* can be done by ellipsis resolution, i.e. detecting that *to impress Mary* was omitted in (2) simply to avoid unnecessary repetition. How to detect this kind of ellipsis is beyond the scope of this paper; the problem has been approached before (Nielsen 2004) and annotation is available (Bos and Spenader 2011).

Another option to represent the meaning of (1) is grounded on selecting a wider focus, *John didn't build a house to impress Mary*. This option is legitimate and cannot be encoded with our proposal because the focus maps to two relation arguments (*build* and *a house*). However, we hypothesize that our proposal combined with ellipsis resolution should extract the same meaning while simplifying the task of focus detection since we restrict focus to a single argument.

Example 2. Let us consider the following statements and proposed representations:

1. John didn't build a house to impress Mary.
   (1a) AGENT(*John*, $build_1$) & (1b) THEME(∼*a house*, $build_1$) & (1c) PURPOSE(*to impress Mary*, $build_1$)
2. He built a castle.
   (2a) AGENT(*He*, $built_2$) & (2b) THEME(*a castle*, $built_2$)

The interpretation for (1, 2) is *John built a castle to impress Mary*. Relations (1a–c) encode that *John* $built_1$ *something to impress Mary, but not a house*; relations (2a–b) encode that *John* $built_2$ *a castle*. Note that neither $build_1$ nor $built_2$ are negated, and an event coreference system should detect that both of them actually refer to the same building: $build_1$ indicates the PURPOSE, $built_2$ the THEME and both indicate the AGENT. We believe that detecting that $build_1$ did occur even though it is negated by *n't* should help resolving the event coreference.

Example 3. A simple, theoretical example of multiple foci is the following:

1. John didn't cheat_1 on Mary.
2. Mary cheated_2 on John.

Sentences (1, 2) are interpreted as *Mary cheated on John* and further imply that John would never do such a thing by preceding (2) with (1) (detecting the implicature is outside the scope of this paper). One can rightfully consider multiple foci and select as focus for (1) *John didn't cheat on Mary*. However, we argue that our proposal, without allowing multiple foci, captures the right meaning of (1, 2) as well:

- (1a) AGENT(*John*, $\sim cheat_1$) & (1b) THEME(*Mary*, $\sim cheat_1$)
- (2a) AGENT(*Mary*, $cheated_2$) & (2b) THEME(*John*, $cheated_2$)

The representation in (1a–b, 2a–b) successfully encodes that $cheat_1$ did not occur and $cheat_2$ occurred with *Mary* as AGENT and *John* as THEME. In this example, both instances of cheated cannot refer to the same event since the first one did not occur and the second did.

Note that the sentence following (1) is key to determining the meaning of (1, 2):

- (2′) *John* $cheated_{2'}$ *on Sue*.
  Both *cheated* are positive, (1, 2′) provide the AGENT and (2′) the THEME.
- (2″) *Bill* $cheated_{2''}$ *on Mary*.
  Both *cheated* are positive, (1, 2″) provide the THEME and (2″) the AGENT.
- (2‴) *Bill* $cheated_{2'''}$ *on Sue*.
  Only $cheated_{2'''}$ occurred, all positive meaning is encoded by (2‴).

4.1.2 Interpreting '∼'

Rewriting a negated argument into its positive counterpart is not always easy. Sometimes, however, the task can be solved by checking a dictionary of antonyms. Consider again the statement, *A panic in Wall Street doesn't exactly inspire confidence* and its proposed representation AGENT(*A panic in Wall Street*, ∼*inspire*) & MANNER(*exactly*, ∼*inspire*) & THEME(*confidence*, ∼*inspire*). One can easily realize that ∼*inspire* is semantically equivalent to *discourage*.

Consider now example (6, Table 1), *John doesn't know exactly how they met*. The suggested representation is (3a) AGENT(*John*, *know*) & (3b) MANNER(∼*exactly*, *know*) & (3c) THEME(*how they met*, *know*). In this case it is easy to rewrite (3b) as MANNER(*incompletely*, *know*).

However, negating an argument is often not straightforward. Take the statement (3, Table 1), *I don't have a watch with me*. The proposed representation is: (4a) AGENT(*I*, *have*) & (4b) THEME(*a watch*, *have*) & (4c) LOCATION(∼*with me*, *have*). Rewriting (4c) without including the symbol '∼' is difficult, even if we took into account the full context and meanings associated with its concepts. The watch might be *at home*, *in the car*, *with my son* or in many other places. Unless context states the actual location or gives a hint, the location probably cannot be specified. Relation (4c) does state that a possible location is not *with me*, but it does not specify any precise area. Note, though, that unlike typical semantic roles, (4a–c) do encode that *I have a watch*, i.e. positive meaning carried by the verbal negation.

Similarly, the proposed representation for example (2, Table 1) *The cow didn't eat grass with a fork* is (1a) AGENT(*The cow*, *ate*) & (1b) THEME(*grass*, *ate*) & (1c) INSTRUMENT(∼*with a fork*, *ate*) (Table 2, Option iv). Commonsense knowledge tells us that the cow (probably) used its mouth, i.e. INSTRUMENT(*its mouth*, *ate*). However, how to determine that this relation is the positive counterpart of (1c) is not simple.

To complicate things further, rewriting a negated argument into its positive equivalent requires deep understanding of text and world knowledge. Decisions must be made based on the full meaning of the current sentence, and potentially a larger context and external knowledge. For illustration purposes, take the following examples:

- *The kids didn't eat soup with a fork*; INSTRUMENT(∼*with a fork*, *ate*) could be rewritten as INSTRUMENT(*with a spoon*, *ate*).
- *The kids didn't eat sandwiches with a fork*; INSTRUMENT(∼*with a fork*, *ate*) could be interpreted as INSTRUMENT(*with their hands*, *ate*).

We believe that even without fully specifying the meaning of a negated argument, the model described here is a significant improvement on semantic representation of negation. The main advantage is to recognize implicit positive meaning, a task currently ignored by the state of the art.

4.2 Incorporating negation into composition of semantic relations

So far we have looked at incorporating negation into semantic relations. In this section we study the impact of negation when composing relations. First, we summarize a model to compose semantic relations (Blanco and Moldovan 2011a, 2011c), then we enhance the model to compose relations when an argument is negated.

4.2.1 Composition of semantic relations

Composition of semantic relations is a framework that aims at extracting relations currently ignored by semantic parsers. It is not coupled to a particular relation set and automatically obtains inference axioms to compose relations. In CSR, a semantic relation R is defined using an extended definition including (a) DOMAIN(R) and RANGE(R), i.e. semantic restrictions on the sorts of concepts that can be the first and second argument; and (b) P_{R}, an array corresponding to the values for a set of semantic primitives. CSR was inspired by previous work within knowledge bases (Cohen and Loiselle 1988; Huhns and Stephens 1989).

Primitives capture elemental semantic properties of relations. Each relation takes a value for each primitive from the set $V = \{-, 0, +\}$, indicating if the primitive does not hold, does not apply or holds. For example, the primitive *temporal* indicates if the first argument must precede the second. We have $P^{temporal}_{\mathrm{CAUSE}} = +$ (a cause must precede its effect) and $P^{temporal}_{\mathrm{MANNER}} = 0$ (there is no temporal precedence between an event and the manner in which it occurs). Using primitives to define relations has an advantage: key semantic properties are explicit, allowing for automatic reasoning.

Axioms are denoted $\mathrm{R}_1(x, y) \circ \mathrm{R}_2(y, z) \rightarrow \mathrm{R}_3(x, z)$, where R_1 and R_2 are the premises, '$\circ$' is the composition operator and R_3 is the conclusion. Given the premises, an axiom infers the conclusion, adding a new link between the ends of the chain formed by the premises. In order to chain two relations $\mathrm{R}_1(x, y)$ and $\mathrm{R}_2(y, z)$, they must have an argument in common, y. We define the inverse of a relation to facilitate chaining relations, and denote the inverse of R as R^{-1}. Given $\mathrm{R}_1(x, y)$, $\mathrm{R}_1{}^{-1}(y, x)$ always holds. Consider sentence *[They]$_x$ [came]$_y$ to [talk]$_z$ about the issue*, where AGENT(*They*, *came*) and PURPOSE(*talk*, *came*) are extracted by an existing semantic parser. Instantiating axiom AGENT$(x, y) \circ$ PURPOSE$^{-1}(y, z) \rightarrow$ AGENT(x, z) with these two relations, we obtain AGENT(*They*, *talk*), a relation probably ignored by some semantic parsers: AGENT(*They*, *came*) $\circ$ PURPOSE^{-1}(*came*, *talk*) $\rightarrow$ AGENT(*They*, *talk*).

Composition of semantic relations automatically obtains axioms using the extended definition and a manually defined algebra for composing primitives. The algebra determines rules for composing primitives, e.g. *the relation resulting from the composition of two relations that hold temporal also holds temporal*. For each primitive, the algebra defines the result of composing the nine possible combinations of values of primitives ($|V \times V| = 9$).

Obtaining inference axioms is reduced to two analytical tasks easy to automate: (a) Enforce domain and range compatibility to find pairs of premises to be used in composition, and (b) for each pair of possible premises, find a relation that fits as the conclusion using the algebra for composing primitives. The following pseudo code (henceforth *CSR algorithm*) loops over all unique combinations of relations and performs both tasks (Blanco and Moldovan 2011c):

Table 3. *Inference axioms obtained using CSR over PropBank semantic roles.* CAU *stands for* CAUSE, PRP *for* PURPOSE, AGT *for* AGENT, THM *for* THEME, LOC *for* LOCATION *and* TMP *for* TIME

No.	Axiom
1.	CAU$\circ$AGT^{-1} $\rightarrow$ AGT^{-1}, the agent of an action is inherited by its cause.
2.	CAU$\circ$LOC^{-1} $\rightarrow$ LOC^{-1}, the spatial context of an action is inherited by its cause.
3.	CAU$\circ$TMP^{-1} $\rightarrow$ TMP^{-1}, the temporal context of an action is inherited by its cause.
4.	PRP$\circ$AGT^{-1} $\rightarrow$ AGT^{-1}, the agent of an action is inherited by its purpose.
5.	PRP$\circ$THM^{-1} $\rightarrow$ THM^{-1}, the theme of an action is inherited by its purpose.
6.	PRP$\circ$LOC^{-1} $\rightarrow$ LOC^{-1}, the spatial context of an action is inherited by its purpose.
7.	PRP$\circ$TMP^{-1} $\rightarrow$ TMP^{-1}, the temporal context of an action is inherited by its purpose.
8.	PRP$\circ$MNR^{-1} $\rightarrow$ MNR^{-1}, the manner of an action is inherited by its purpose.

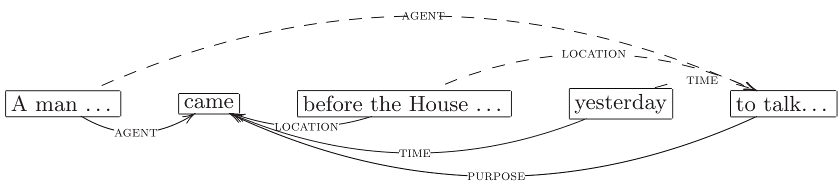

Fig. 1. Example of instantiation of inference axioms. Solid arrows indicate annotation in PropBank, discontinuous arrows inferred relations using axioms 4, 6 and 7 from Table 3.

CSR algorithm(R)
Input: R, set of semantic relations defined using the extended definition
Output: *inference_axioms*, list of axioms using $R_1 \in R$ and $R_2 \in R$ as their premises
Repeat for $(R_1, R_2) \in R \times R$:
Repeat for $(R_i, R_j) \in [(R_1, R_2), (R_1{}^{-1}, R_2), (R_2, R_1), (R_2, R_1{}^{-1})]$:

1. **Domain and range compatibility.**
   If $\text{RANGE}(R_i) \cap \text{DOMAIN}(R_j) = \emptyset$, **break**
2. **Conclusion match.**
   Using the algebra for composing primitives, calculate $P_{R_i} \circ P_{R_j}$
   Repeat for $R_3 \in R$:
   (a) **If** $\text{DOMAIN}(R_3) \cap \text{DOMAIN}(R_i) = \emptyset$ **or** $\text{RANGE}(R_3) \cap \text{RANGE}(R_j) = \emptyset$, **break**
   (b) **If** *consistent*$(P_{R_3}, P_{R_i} \circ P_{R_j})$,
   inference_axioms.append$(R_i(x, y) \circ R_j(y, z) \rightarrow R_3(x, z))$

Composition of semantic relations is applicable to any relation inventory. Blanco and Moldovan (2011a) applied it to PropBank and obtained eight inference axioms (Table 3). Instantiating these axioms improves the basic representation that PropBank provides by adding relations currently ignored. Consider the following sentence from PropBank (file wsj_0134, sentence 0):

[A man from the Bush administration]$_{z,\text{AGENT}}$ *[came]*$_{y,verb}$ *[before the House Agriculture Committee]*$_{z',\text{LOCATION}}$ *[yesterday]*$_{z'',\text{TIME}}$ *[to talk about the U.S.'s intention to send some \$100 million in food aid to Poland, with more to come from the EC]*$_{x,\text{PURPOSE}}$.

Figure 1 shows PropBank annotation for the verb *came* with solid arrows. By instantiating axioms 4, 6 and 7 from Table 3 with this basic annotation, one can infer the three relations shown with discontinuous arrows.

4.2.2 Instantiating axioms with negated arguments

As illustrated in Figure 1, instantiating an axiom $\text{R}_1(x, y) \circ \text{R}_2(y, z) \rightarrow \text{R}_3(x, z)$ over relations without a negated argument is reduced to finding a chain $\text{R}_1(x, y) \circ \text{R}_2(y, z)$ and adding $\text{R}_3(x, z)$. However, when a premise has a negated argument, the conclusion may not have that same argument negated. In this section, we illustrate this phenomenon.

Consider statement (5) *John didn't build a house to impress Mary*. The proposed representation is: (5a) AGENT(*John*, *build*) & (5b) THEME(*a house*, *build*) & (5c) PURPOSE(∼*to impress Mary*, *build*). Instantiating $\text{PURPOSE}(x, y) \circ \text{AGENT}^{-1}(y, z) \rightarrow \text{AGENT}^{-1}(x, z)$ over relations (5c, 5a), we must conclude AGENT^{-1}(∼*to impress Mary*, *John*), i.e. *John is the agent of [not to impress Mary]*. In other words, *John is the agent of some action, but not of impressing Mary*. This example suggests that instantiating an inference axiom with a negated argument might be as simple as maintaining the negated mark in the conclusion: PURPOSE(∼*to impress Mary*, *build*) ∘ AGENT^{-1}(*build*, *John*) → AGENT^{-1}(∼*to impress Mary*, *John*).

Now consider sentence (6) *In Frankfurt, stocks didn't open for the first 45 minutes because of order imbalances* (wsj_0719, 18). Its meaning is *In Frankfurt, stocks opened, but not for the first 45 minutes, because of order imbalances*, and its representation is (6a) AGENT(*stocks*, *open*) & (6b) LOCATION(*In Frankfurt*, *open*) & (6c) TIME(∼*for the first 45 minutes*, *open*) & (6d) CAUSE(*because of order imbalances*, *open*). Instantiating $\text{CAUSE}(x, y) \circ \text{TIME}^{-1}(y, z) \rightarrow \text{TIME}^{-1}(x, z)$ over (6d, 6c), we must get TIME^{-1}(*because of order imbalances*, *for the first 45 minutes*), i.e. *the order imbalances happened during the first 45 minutes*. Note that in this example a premise has a negated argument and yet the conclusion has *all arguments positive*: CAUSE(*because of order imbalances*, *open*) ∘ TIME^{-1}(*open*, ∼*for the first 45 minutes*) → TIME^{-1}(*because of order imbalances*, *for the first 45 minutes*).

4.2.3 Enhancing CSR to incorporate negation

As exemplified in the previous section, instantiating an axiom when considering negation is not straightforward. In this section, we enhance the general CSR framework presented in Section 4.2.1 to take into account negation, a scenario not studied before. Composing semantic relations is broken down into two steps: axiom extraction and axiom instantiation.

The extraction of inference axioms is done using the CSR algorithm exactly as before. Given any set of relations defined using the extended definition, this algorithm extracts inference axioms. The instantiation of an axiom $\text{R}_1(x, y) \circ \text{R}_2(y, z) \rightarrow \text{R}_3(x, z)$, though, is more complicated than finding a chain of relations matching the premises, in order to determine if the conclusion has a negated argument, we must examine whether the premises have a negated argument as well as the semantics of the premises.

We introduce the primitive *negation* to incorporate negation into CSR. *Negation* indicates if the instance of a relation has any argument negated. '−' indicates that the first argument is negated, '+' indicates that the second argument is negated and '0' indicates that neither argument is negated. Regarding the inverse relation, if $\text{R}(\sim x, y)$ (*negation* = −), then $\text{R}^{-1}(y, \sim x)$ (*negation* = +). For example, TIME(∼*for the first 45 minutes*, *open*) and AGENT(*stocks*, *open*) from sentence (6) *In Frankfurt, stocks didn't open for the first 45 minutes because of order imbalances* take '−' and '0' respectively; MANNER(*exactly*, ∼*encourages*) from (2) *A panic in Wall street doesn't exactly encourage confidence* takes '+'. Note that assigning a value to *negation* indicates whether the instance has an argument negated.

The algebra to compose primitives is extended with the *negation* primitive. Given the values two relation instances take for *negation* ($P_{\text{R}_1}^{neg}$, $P_{\text{R}_2}^{neg}$), it can be observed that the following rules depict the values for their composition ($P_{\text{R}_1}^{neg} \circ P_{\text{R}_2}^{neg}$):

$P_{\text{R}_1}^{neg}$	$P_{\text{R}_2}^{neg}$	$P_{\text{R}_1}^{neg} \circ P_{\text{R}_2}^{neg}$
−	−	−
−	0	−
−	+	0
0	−	−
0	0	0
0	+	+
+	−	0
+	0	+
+	+	+

As the examples in Section 4.2.2 show, deciding if the conclusion of instantiating an axiom $\text{R}_1(x, y) \circ \text{R}_2(y, z) \rightarrow \text{R}_3(x, z)$ must have an argument negated ($P_{\text{R}_3}^{neg}$) depends not only on whether or not the premises have an argument negated ($P_{\text{R}_1}^{neg}$ and $P_{\text{R}_2}^{neg}$), but also on the semantics of R_3. We found that the primitives *separable* and *intrinsic* are essential.

Separable is defined as *x can be temporally or spatially separated from y, thus x can exist independently of y* (Winston, Chaffin and Herrmann 1987). Actions and their purposes can exist independently of each other, therefore $P_{\text{PURPOSE}}^{separable} = +$. *Intrinsic* is defined as *relation is an attribute of the essence/stufflike nature of x or y* (Huhns and Stephens 1989). Consider CAUSE(x, y), where the occurrence of y is due to x. The relation between x and y is an attribute of the essence of y (its existence), therefore $P_{\text{CAUSE}}^{intrinsic} = +$. On the other hand, the relation between an action y and its purpose x is not in the essence of y (y may have multiple purposes with different motivations, each purpose may or may not happen in the future), so we have $P_{\text{PURPOSE}}^{intrinsic} = -$.

If *intrinsic* holds for the first premise and *separable* does not apply to the second premise, then *negation* does not apply to the conclusion R_3 ($P^{neg}_{\text{R}_3} = 0$). Otherwise $P^{neg}_{\text{R}_3}$ is calculated according to the above rules for $P^{neg}_{\text{R}_1} \circ P^{neg}_{\text{R}_2}$.

The procedure for instantiating inference axioms considering negation is depicted in the Instantiate-Axiom algorithm. For each pair of relations ($\text{R}_1(x, y)$, $\text{R}_2(y, z)$) instantiating each axiom, a new relation $\text{R}_3(x, z)$ is inferred. Depending on the value assigned to $P^{neg}_{\text{R}_3}$, which is calculated over $P^{neg}_{\text{R}_1}$, $P^{neg}_{\text{R}_2}$, $P^{intrinsic}_{\text{R}_1}$ and $P^{separable}_{\text{R}_2}$, the conclusion may or may not have an argument negated.

The following pseudo-code presents the Instantiate-Axiom algorithm:

Instantiate-Axioms algorithm$(A, rels)$
Input: A, set of inference axioms $\text{R}_1(x, y) \circ \text{R}_2(y, z) \rightarrow \text{R}_3(x, z)$ extracted by CSR algorithm
rels, list of relation instances
Output: *inferred_relations*, list of relation instances inferred with A over *rels*
Repeat for all *axiom* $\text{R}_i(x, y) \circ \text{R}_j(y, z) \rightarrow \text{R}_k(x, z) \in A$
Repeat for all relation instances $(\text{R}_1(x, y), \text{R}_2(y, z)) \in rels$ instantiating *axiom*:

1. **If** $P^{intrinsic}_{\text{R}_1} =$ '+' **and** $P^{separable}_{\text{R}_2} =$ '0', $P^{neg}_{\text{R}_3} \leftarrow$ '0'
2. **else** $P^{neg}_{\text{R}_3} \leftarrow P^{neg}_{\text{R}_1} \circ P^{neg}_{\text{R}_2}$

inferred_relations.append$(\text{R}_3(x, z))$

An example. Consider axiom (4) $\text{PURPOSE}(x, y) \circ \text{AGENT}^{-1}(y, z) \rightarrow \text{AGENT}^{-1}(x, z)$, automatically extracted using the CSR algorithm, and the following sentence (7) *He ultimately became so well-known for cutting compensations, however, that clients didn't seek him out for anything else*. A partial representation of (7) is (7a) AGENT(*clients*, *seek*) & (7b) THEME(*him*, *seek*) & (7c) PURPOSE(∼*for anything else*, *seek*), encoding *clients sought him for the same (i.e. for cutting compensations)*.

The above axiom (4) can be instantiated using as premises (7c, 7a): PURPOSE(∼*for anything else*, *seek*) ∘ AGENT^{-1}(*seek*, *clients*). AGENT^{-1} holds the primitive *separable* ($P^{separable}_{\text{AGENT}^{-1}} = +$ (Blanco and Moldovan 2011a)), so we must follow case 2 of Instantiate-Axiom algorithm. Following the rules for composing *negated*, the conclusion takes '−' for this primitive ('−' ∘ '0' = '−'). Thus, we infer AGENT^{-1}(∼*for anything else*, *clients*). Paraphrasing, clients are the agents of cutting compensations.

Now consider the previous example (6) *In Frankfurt, stocks didn't open for the first 45 minutes because of order imbalances*. Axiom $\text{CAUSE}(x, y) \circ \text{TIME}^{-1}(y, z) \rightarrow \text{TIME}^{-1}(x, z)$ can be instantiated with relations (6d, 6c): CAUSE(*because of order imbalances*, *open*) ∘ TIME^{-1}(*open*, ∼*for the first 45 minutes*). Because the primitive *separable* does not apply to TIME^{-1} ($P^{separable}_{\text{TIME}^{-1}} = 0$) and $P^{intrinsic}_{\text{CAUSE}} = +$ (Blanco and Moldovan 2011a), we must follow case 1 of Instantiate-Axiom algorithm. The conclusion takes '0' as value for *negated*, inferring TIME^{-1}(*for the first 45 minutes*, *because of order imbalances*), i.e. the order imbalances happened during the first 45 minutes. Even though *for the first 45 minutes* appears negated in the first premise, the procedure correctly deletes the negated mark in the conclusion.

Table 4. *Argument modifiers in PropBank*

MLOC:	location	MNEG:	negation marker	MPNC:	purpose
MEXT:	extent	MMOD:	modal verb	MMNR:	manner
MDIS:	discourse connective	MCAU:	cause	MDIR:	direction
MADV:	general-purpose	MTMP:	time		

5 Annotating the focus of negation

Section 4 introduced a novel representation for negation grounded on focus detection. This representation requires automatic detection of focus of negation, a task never undertaken before. In this section, we describe our annotation effort for focus of negation (Blanco and Moldovan 2011b) and in Section 6 we present experimental results. The annotations are publicly available at `http://www.clips.ua.ac.be/sem2012-st-neg/` (pb-foc corpus) and were used in the *SEM 2012 Shared Task (Morante and Blanco 2012).

Because of the lack of corpora containing annotation for focus of negation, new annotations are needed. An obvious option is to add it to any text collection. However, building on top of publicly available resources is a better approach: they are known by the community, they contain useful information for detecting the focus of negation and tools have already been developed to predict their annotations.

We decided to work over PropBank. Unlike other resources (e.g. FrameNet), gold syntactic trees are available. Compared to the BioScope corpus, PropBank provides semantic annotations and is not limited to the biomedical domain. The additional annotations can be readily used by any system working with PropBank, quickly incorporating interpretation of negation to them. On top of that, there has been active research on predicting PropBank roles for years, including systems participating in the CoNLL-2004 and CoNLL-2005 Shared Tasks (Carreras and Màrquez 2004, 2005), the special issue on 'Semantic Role Labeling' of *Computational Linguistics* (Màrquez *et al.* 2008), and many workshop and conference papers, e.g. Dang and Palmer (2005), Koomen *et al.* (2005), Zapirain, Agirre and Màrquez (2008), Merlo and Van der Plas (2009) and Lang and Lapata (2010).

5.1 Annotation guidelines

PropBank annotates exclusively semantic roles, i.e. semantic relations between a verb and its arguments. Each verb is annotated with up to six numeric core arguments (A0, A1, A2, A3, A4, A5) and argument modifiers (MTMP, MLOC etc.). Numeric arguments do not have a uniform meaning across verbs, but A0 generally exhibits features of a prototypical AGENT, while A1 is a prototypical PATIENT or THEME (Palmer *et al.* 2005). Argument modifiers have a fixed meaning across verbs summarized in Table 4. For a discussion on the creation of the corpus and the semantics of each label, refer to Palmer *et al.* (2005) and the annotation guidelines (`http://verbs.colorado.edu/ mpalmer/projects/ace/PBguidelines.pdf`).

We targeted verbal negations involving a verb v and having a semantic role of the form MNEG(x, v). The focus is resolved as follows:

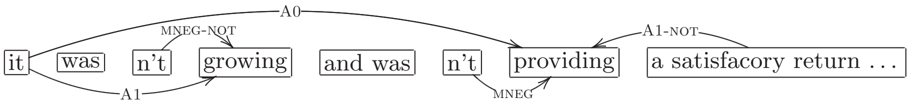

Fig. 2. Example of focus annotation (marked with NOT).The new representation explicitly states that *it (the company) was not growing and was providing a not satisfactory return on invested capital.*

- If it cannot be inferred that an action v occurred (i.e. focus is v), mark as focus MNEG.
- Otherwise, mark as focus the role that is most prominently negated.

All decisions are made considering the context of the previous and next sentence. The mark -NOT is used to indicate the focus. Consider the following statements (wsj_2282, sentences 15 and 16):

- Applied, then a closely held company, was stagnating under the management of its controlling family.
- [While profitable]$_{\text{MADV}_{1,2}}$, [it]$_{\text{A1}_1,\text{A0}_2}$ 'was[n't]$_{\text{MNEG}_1}$ [growing]$_{\text{V}_1}$ and was[n't]$_{\text{MNEG}_2}$ [providing]$_{\text{V}_2}$ [a satisfactory return on invested capital]$_{\text{A1}_2}$,' he says.

Regarding the first verb (*growing*), one cannot infer that anything was growing, so focus is MNEG. For the second verb (*providing*), it is implicitly stated that the company was providing *a not satisfactory return on investment*, therefore focus is A1. Figure 2 shows the new annotation.

The guidelines assume that the focus corresponds to a single role or the verb. In cases where more than one role could be selected, the most likely focus is chosen; context and text understanding help resolving ambiguities. We define the most likely focus as the one that yields the most meaningful implicit information.

The task of selecting the focus is highly ambiguous. For illustration purposes, consider the following sentences (wsj_1856, sentences 27 and 28; also in Table 5):

1. Kathie Huff, a respondent in the Journal survey from Spokane, Wash., says her husband is adamant about eating only Hunt's ketchup.
2. [He (her husband)]$_{\text{A0}}$ [simply]$_{\text{MMNR}}$ [ca]$_{\text{MMOD}}$[n't]$_{\text{MNEG}}$ [stomach]$_{\text{V}}$ [the taste of Heinz]$_{\text{A1}}$, she says.

Four options for the focus of the verbal negation in (2) are:

(a) [He]$_{\text{A0}}$, encoding *there are people who simply can stomach the taste of Heinz, but not he.*
(b) [simply]$_{\text{MMNR}}$, encoding *He can stomach the taste of Heinz, but with difficulties.*
(c) [stomach]$_{\text{V}}$, encoding *he simply can do a few things with the taste of Heinz, but not stomach it.*
(d) [the taste of Heinz]$_{\text{A1}}$, encoding *he simply can stomach any ketchup but Heinz.*

Taking into account the context, the best option is A1, *the taste of Heinz*. Note that we are purposely limiting the choice of focus to all the words belonging to a labeled role from PropBank. This way we ease the task and make it compatible with the original PropBank. A possible finer grained focus is the word *Heinz*, but we select A1.

Several examples of annotated statements from PropBank are depicted in Table 5. The negation in example (1) does not carry any positive meaning, the focus is the verb. In examples (2–10) the verb must be interpreted as affirmative, as well as all roles except the one marked with '⋆' (i.e. the focus). For each example, we provide PropBank annotation (top), the new annotation (i.e. the focus, bottom right) and its interpretation (bottom left). Some examples in Table 5 are ambiguous, e.g. in example (5) there is ambiguity between A0 and A2; and in example (9) between MTMP and A4.

Selecting a semantic role, as opposed to specific words, has the drawback of potentially selecting too coarse of a focus. For example, in *the company wasn't providing a satisfactory return on investment*, choosing *satisfactory* as focus would help determining the positive counterpart (i.e. *an unsatisfactory return on investment*). On the other hand, selecting a role has the advantage of simplifying the task of predicting the focus and allows existing role labelers to make this prediction without severe modifications: the only requirement is to retrain with the new labels.

5.1.1 Phenomena influencing focus of negation

Annotating the focus of negation is a semantic task influenced by many phenomena commonly found in text. Context plays an important role and deep understanding of text is critical. The following examples illustrate some interesting recurrent phenomena that we found during the annotation process.

Shorter clauses tend to have as their focus the verb, e.g. *'[I]$_{A0}$'m not [interested]$_v$,' said Dallas investor Harold Simmons*. These statements typically contain a limited number or roles and they should often be interpreted by negating the verb.

Context is key for deciding the focus, especially the previous sentence. Consider the following example, describing the buying of Farmers (an American company) by Axa (a French company).

- Claude Bebear, chairman and chief executive officer, of Axa-Midi Assurances, pledged to retain employees and management of Farmers Group Inc., including Leo E. Denlea Jr., chairman and chief executive officer, if Axa succeeds in acquiring Farmers.
- Mr. Bebear added that [the French insurer]$_{A0}$ would keep Farmers' headquarters in Los Angeles and '[will]$_{MMOD}$ not [send]$_v$ [French people]$_{A1}$ [to run the company]$_{MPNC}$'.

Given the context, it is clear that no one will send anyone for any purpose (i.e. no *sending* event will take place); if Axa ends up buying Farmers, they will let current Farmers employees run that division of the company.

Some verbs are rarely the focus, e.g. *seek*, *want*, *believe*. These verbs usually signal that the focus is A1(x, y), corresponding to the interpretation that what is sought, wanted or believed is *[not x]*.

Table 5. *Negated statements from PropBank and their interpretation considering implicit positive meaning. '✓' indicates that the role is present and '★' indicates that it corresponds to the focus. The verbal negation in example (1) does not carry positive meaning*

No.	Example	NEG	A0	A1	A2	A4	TMP	MNR	ADV	LOC	PNC	EXT	DIS	MOD
1.	Even if [that deal]$_{\text{A1}}$ isn't [revived]$_{\text{V}}$, NBC hopes to find another. – Even if that deal is suppressed, NBC hopes to find another one.	★	-	✓	-	-	-	-	-	-	-	-	-	-
2.	[He]$_{\text{A0}}$ [simply]$_{\text{MDIS}}$ [ca]$_{\text{MMOD}}$n't [stomach]$_{\text{V}}$ [the taste of Heinz]$_{\text{A1}}$, she says. – He simply can stomach any ketchup but Heinz's.	✓	✓	★	-	-	-	-	-	-	-	-	✓	✓
3.	[A decision]$_{\text{A1}}$ isn't [expected]$_{\text{V}}$ [until some time next year]$_{\text{MTMP}}$. – A decision is expected at some time next year.	✓	-	✓	-	-	★	-	-	-	-	-	-	-
4.	[. . .] it told the SEC [it]$_{\text{A0}}$ [could]$_{\text{MMOD}}$n't [provide]$_{\text{V}}$ [financial statements]$_{\text{A1}}$ [by the end of its first extension]$_{\text{MTMP}}$ '[without unreasonable burden or expense]$_{\text{MMNR}}$'. – It could provide them by that time with a huge overhead.	✓	✓	✓	-	-	✓	★	-	-	-	-	-	✓
5.	[For example]$_{\text{MDIS}}$, [P&G]$_{\text{A0}}$ [up until now]$_{\text{MTMP}}$ hasn't [sold]$_{\text{V}}$ [coffee]$_{\text{A1}}$ [to airlines]$_{\text{A2}}$ and does only limited business with hotels and large restaurant chains. – Up until now, P&G has sold coffee, but not to airlines.	✓	✓	✓	★	-	✓	-	-	-	-	-	✓	-
6.	[Decent life [. . .]]$_{\text{A1}}$ [wo]$_{\text{MMOD}}$n't be [restored]$_{\text{V}}$ [unless the government reclaims the streets [. . .]]$_{\text{MADV}}$. – It will be restored if the government reclaims the streets.	✓	-	✓	-	-	-	-	★	-	-	-	-	✓
7.	But [quite a few money managers]$_{\text{A0}}$ aren't [buying]$_{\text{V}}$ [it]$_{\text{A1}}$. – Very little managers are buying it.	✓	★	✓	-	-	-	-	-	-	-	-	-	-
8.	[When]$_{\text{MTMP}}$ [she]$_{\text{A0}}$ isn't [performing]$_{\text{V}}$ [for an audience]$_{\text{MPNC}}$, she prepares for a song by removing the wad of gum from her mouth, and indicates that she's finished by sticking the gum back in. – She prepares in that way when she is performing, but not for an audience.	✓	✓	-	-	-	✓	-	-	-	★	-	-	-
9.	[It]$_{\text{A1}}$ [can]$_{\text{MMOD}}$not [fall]$_{\text{V}}$ [below \$185 million]$_{\text{A4}}$ [after the dividends are issued]$_{\text{MTMP}}$. – It can fall after the dividends are issued, but not below \$185 million.	✓	-	✓	-	★	✓	-	-	-	-	-	-	✓
10.	Mario Gabelli, an expert at [. . .], says that [takeovers]$_{\text{A1}}$ aren't [totally]$_{\text{MEXT}}$ [gone]$_{\text{V}}$. – Mario Gabelli says that takeovers are partially gone.	✓	-	✓	-	-	-	-	-	-	-	★	-	-

- But he noted that [speculators]$_{\text{A0}}$ [apparently]$_{\text{MMNR}}$ don't [believe]$_v$ [there is much more of a decline in store for cocoa]$_{\text{A1}}$.
  – He noted that they apparently believe *there is little or not decline in store for cocoa.*

Verb modifiers play a role, especially adverbs. For example, *even* typically signals a strong negative verb polarity. Thus, this adverb usually indicates that the negation does not carry positive meaning and the focus is the verb.

- *[The company's conduct]*$_{\text{A0}}$ *'does not [even]*$_{\text{MADV}}$ *[raise]*$_v$ *[a question of wrongful corporate intent, ratification or cover-up]*$_{\text{A1}}$*', GE's brief asserts.*
  – There are no issues with the company.

Conditionals attached to a verbal negation often carry implicit positive meaning. For example, the following statements are equivalent:

- The Los Angeles investor can't [buy]$_v$ UAL stock unless he makes a formal offer of $300 a share or UAL accepts an offer below $300.
  – The Los Angeles investor can buy UAL stock *if he makes a formal offer of $300 a share or UAL accepts an offer below $300.*

Certain quantifiers in A0 usually signal that the focus is A0. This role corresponds to the prototypical AGENT (Palmer *et al.* 2005). Generally, if one says that *some/most/part of x* does not *y*, it is implicitly saying that *others/few/another part of x* does *y*. The following example illustrates this:

- But John LaWare, a Fed governor, told the subcommittee the evidence is mixed and that the Fed's believes [the vast majority of banks]$_{\text{A0}}$ aren't [discriminating]$_v$.
  – He told them that he believes that *very few banks* are discriminating.

Modifiers in A1 usually signal that the focus is A1. This role corresponds to the prototypical THEME (Palmer *et al.* 2005). If it contains adjectives such as *sufficient* and *necessary*, the event they attach to actually occurs with a negated A1:

- 'We would be the first to admit that [we]$_{\text{A0}}$ have not [devoted]$_v$ [the necessary amount of emphasis]$_{\text{A1}}$ [over the past several years]$_{\text{MTMP}}$' [to developing examinations for discrimination]$_{\text{MPNC}}$, said Jonathan Fiechter, a top official of the Office of Thrift Supervision.
  – 'We would be the first to admit that we have devoted *insufficient amount of emphasis* over the past several years' to developing examinations for discrimination said Jonathan Fiechter, a top official of the Office of Thrift Supervision.

Certain roles are more likely to be the focus. Particularly roles that do not occur frequently are often the focus when present. For example, MMNR is frequently the focus, but not always. Compare the following statements:

- Seeing all those millions in action, I was just so relieved that [Ms. Gruberova, gawky thing that she is]$_{\text{A0}}$, didn't [accidentally]$_{\text{MMNR}}$ [smother]$_v$ [herself]$_{\text{A1}}$ [in

a drape]$_{A2}$.
– [...] I was just so relieved that *nobody smothered himself.*

- Some critics say [they]$_{A1}$ [wo]$_{MMOD}$n't be [quickly]$_{MMNR}$ [embraced]$_{v}$ [by consumers]$_{A0}$ [because of the high price]$_{MCAU}$.
  – Some critics say that they will be embraced by consumers, *but only in the future*, because of the high price.

Temporal information signaled by *until* repeatedly indicates that the focus is MTMP, but not always. Compare the following statements:

- [The union]$_{A0}$ [wo]$_{MMOD}$n't [respond]$_{v}$ [to the USX statement]$_{A1}$ [until Mr. Williams has studied it]$_{MTMP}$, the spokesman said.
  – The union will respond to the USX statement *after Mr. Williams has studied it*, the spokesman said.
- We urge [our people]$_{A1}$ not to [wait]$_{v}$ [until they have to fight for their own nation]$_{MTMP}$.
  – We urge our people to [not wait] (i.e. our people should act now).

5.2 Normalization of -NOT

According to the proposed guidelines, NOT is attached to the role corresponding to the focus of negation or MNEG if the focus is the verb. The guidelines were specifically designed to maintain the original PropBank format; for each instance only one label is changed. One can easily normalize this enhanced PropBank annotation to the model proposed in Section 4.1:

- If MNEG-NOT(x, y), all roles R(x', y) are normalized as R(x', $\sim y$).
- If any other role R(x, y) is marked with NOT, that role is normalized as R($\sim x$, y) and all other roles remain as in the original PropBank.

After normalization, MNEG (with or without NOT) can be ignored since it is overwritten by '$\sim$'. For example, the normalized representation for the example in Figure 2 is MADV(*While profitable*, $\sim$*growing*) & A1(*it*, $\sim$*growing*) & MADV(*while profitable*, *providing*) & A0(*it*, *providing*) & A1($\sim$*a satisfactory return on investment*, *providing*).

5.3 Annotation process

We annotated 3,993 verbal negations signaled with MNEG in PropBank, contained in 3,779 sentences. Before annotation began, all semantic information was removed by mapping all role labels to ARG. This step is necessary to ensure that focus selection is not biased by the semantic labels provided by PropBank.

As an annotation tool, we use Jubilee (Choi, Bonial and Palmer 2010). This Java-based application required minimal modifications to configure files and proved very convenient for our focus annotation task. It provides multi-annotator capabilities and a simple methodology to resolve disagreements. For each instance (i.e. predicate from the original PropBank with MNEG), annotators decide the focus given the parse

Table 6. *Roles, total instantiations and counts corresponding to focus over training and held-out instances*

		Focus				Focus	
Role	#Instances	#	(%)	Role	#Instances	#	(%)
A1	2,930	1,194	(40.75)	MLOC	114	22	(19.30)
MNEG	3,196	1,109	(34.70)	MEXT	25	22	(88.00)
MTMP	609	246	(40.39)	A4	26	22	(84.62)
MMNR	250	190	(76.00)	A3	48	18	(37.50)
A2	501	179	(35.73)	MDIR	35	13	(37.14)
MADV	466	94	(20.17)	MPNC	87	9	(10.34)
A0	2,163	73	(3.37)	MDIS	287	6	(2.09)

tree, as well as the previous and next sentence. A post-processing step incorporates focus annotation to the original PropBank by adding NOT to the corresponding role.

In the first round, 50 percent of instances were annotated twice by two graduate students in computational linguistics. Inter-annotator agreement was 72 percent, calculated as the percentage of annotations that were a perfect match. After careful examination of the disagreements, they were resolved and annotators were given clearer instructions. Annotators often disagreed about whether the verb or a semantic role was the focus. The remaining instances were annotated once.

Table 6 depicts counts for each role over the training and held-out splits.

6 Experimental results

In this section we present experimental results on predicting the focus of negation (Section 6.1) and the methodology to incorporate negation into CSR (Section 6.2). All experiments were carried out using gold annotations, and therefore it is expected that using automatic annotations will lower the performance.

6.1 Focus detection

We have implemented four baselines and more complicated models to predict the focus of negation. These systems show that it is feasible to obtain the proposed representation for negated statements (Section 4.1) automatically.

Each of the 3,993 verbal negations annotated (Section 5.3) was divided into training (70 percent), held-out (10 percent) and test (20 percent). The held-out portion is used to tune the feature set and results are reported for the test split only, i.e. using unseen instances. Because PropBank adds semantic role annotation on top of the Penn Treebank (Marcus, Santorini and Marcinkiewicz 1994), we have available gold syntactic annotation and semantic role labels for all instances.

We implemented four baselines to measure the difficulty of the task:

- `A1`: Select the most likely focus to be the role, A1. If A1 is not present, then select MNEG.

Table 7. *Accuracies over test split for the four baselines*

System	Accuracy (%)
A1	42.11
FIRST	7.00
LAST	58.39
BASIC	61.38

Table 8. *Full set of features. Features (1–5) are extracted for all roles, and features (7, 8) for all POS tags and keywords considered*

	Feature	Values	Explanation
1.	role-present	{yes, no}	is role present?
2.	role-f-pos	{DT, NNP, ...}	first POS tag of role
3.	role-f-word	{This, to, overseas, ...}	first word of role
4.	role-length	ℕ	number of words in role
5.	role-posit	ℕ	position within the set of roles
6.	A1-top	{NP, SBAR, PP, ...}	syntactic node of A1
7.	A1-postag	{yes, no}	does A1 contain the tag *postag*?
8.	A1-keyword	{yes, no}	does A1 contain the word *keyword*?
9.	first-role	{A0, A1, MTMP, ...}	label of the first role
10.	last-role	{A0, A1, MTMP, ...}	label of the last role
11.	verb-word	{appear, describe, ...}	main verb
12.	verb-postag	{VBN, VBZ, ...}	POS tag main verb
13.	VP-words	{were-n't, be-quickly, ...}	sequence of words of VP until verb
14.	VP-postags	{VBP-RB-RB-VBG, ...}	sequence of POS tags of VP until verb
15.	VP-has-CC	{yes, no}	does the VP until verb contain a CC?
16.	VP-has-RB	{yes, no}	does the VP until verb contain a RB?
17.	predicate	{rule-out, come-up, ...}	predicate
18.	them-role-A0	{preparer, assigner, ...}	thematic role for A0
19.	them-role-A1	{effort, container, ...}	thematic role for A1
20.	them-role-A2	{audience, loaner, ...}	thematic role for A2
21.	them-role-A3	{intensifier, collateral, ...}	thematic role for A3
22.	them-role-A4	{beneficiary, end point, ...}	thematic role for A4

- FIRST: Select the first role, i.e. the one whose first content word starts the earliest within the sentence.
- LAST: Select the last role, i.e. the one whose first content word starts the latest within the sentence.
- BASIC: Same as FOC-DET (Section 6.1.1) but only using features *last-role* and flags indicating presence of roles (*role-present*).

The accuracy for each baseline, ranging from 7.00 to 61.38, is depicted in Table 7. In Section 6.1.2 we provide detailed results.

6.1.1 Selecting features

BASIC yields an accuracy of 61.38. We improved this baseline by considering an enhanced feature set (Table 8). All features are fairly simple and specifically target A1 and MNEG, the roles that most frequently correspond to the focus of negation (Table 6).

Features (1–5) are extracted for all roles present in PropBank, regardless of whether or not they correspond to the focus in our annotations. These features capture their presence, first word and part-of-speech (POS) tag, length and position within the roles present for that instance. Features (6–8) further characterize A1. A1-postag is extracted for the following POS tags: DT, JJ, PRP, CD, RB, VB and WP; A1-keyword for the following words: *any*, *anybody*, *anymore*, *anyone*, *anything*, *anytime*, *anywhere*, *certain*, *enough*, *full*, *many*, *much*, *other*, *some*, *specifics*, *too* and *until*. The above lists were extracted after manual examination of training examples and aim at signaling whether A1 corresponds to the focus.

Examples of A1 corresponding to the focus and including one of the POS tags or keywords are as follows:

- *[Apparently]*$_{\text{MADV}}$, *[the respondents]*$_{A0}$ *don't think [that an economic slowdown would harm the major investment markets very*^{RB} *much]*$_{\text{A1}}$. (i.e. the responders think it would harm the investments little).
- *[The oil company]*$_{\text{A0}}$ *does n't anticipate [any*^{keyword} *additional charges]*$_{\text{A1}}$ (i.e. *the company anticipates no additional charges*).
- *[Money managers and other bond buyers]*$_{\text{A0}}$ *haven't [shown]*$_{\text{V}}$ *[much*^{keyword} *interest in the Refcorp bonds]*$_{\text{A1}}$ (i.e. *they have shown little interest in the bonds*).
- *He concedes H&R Block is well-entrenched and a great company, but says '[it]*$_{\text{A1}}$ *doesn't [grow]*$_{\text{V}}$ *[fast enough*^{keyword} *for us]*$_{\text{A1}}$*'* (i.e. it is growing too slow for us).
- [We]$_{A0}$ don't [see]$_{\text{V}}$ [a domestic source for some^{keyword} of our HDTV requirements]$_{\text{A1}}$, and that's a source of concern [...] (i.e. we see a domestic source for some other of our HDTV requirements).

Features (9, 10) indicate the first and last role respectively. Features (11–16) characterize the main verb. VP-postag (VP-words) corresponds to the full sequence of POS tags (words) from the beginning of the VP until the main verb. Features (15–16) check for POS tags as the presence of certain tags usually signal that the verb is not the focus of negation (e.g. *[Thus]*$_{\text{MDIS}}$, *he asserts, [Lloyd's]*$_{\text{A0}}$ *[[ca]*$_{\text{MMOD}}$*n't [react]*$_{v}$ *[quickly*^{RB}*]*$_{MMNR}$ *[to competition]*$_{\text{A1}}$*]*$_{VP}$).

Features (17–22) tackle the predicate, which includes the main verb and may include other words (typically prepositions). We consider the words in the predicate, as well as the specific thematic roles for each numbered argument. This is useful since PropBank uses the same numbered arguments for different thematic roles depending on the frame (e.g. A3 is used as PURPOSE in *authorize.01* and as INSTRUMENT in *avert.01*).

The total number of features, taking into account that features (1–5) are extracted for eighteen roles, features (7, 8) for seven POS tags and seventeen keywords

Table 9. *Per-class precision, recall and f-measure using baselines* A1*,* FIRST *and* LAST*.* OTHER *stands for all roles that are never predicted by the baseline*

Baseline	Role	Focus (%)	Precision	Recall	f-measure
A1	A1	37.35	0.410	1.000	0.582
	MNEG	34.69	0.535	0.136	0.217
	OTHER	27.96	0.000	0.000	0.000
	Weighted avg.		0.339	0.421	0.375
FIRST	A1	37.35	0.114	0.079	0.093
	MNEG	34.69	0.760	0.061	0.113
	MTMP	7.69	0.123	0.051	0.072
	A2	5.60	0.024	0.011	0.015
	MADV	2.94	0.093	0.151	0.115
	A0	2.28	0.018	0.415	0.035
	MLOC	0.69	0.125	0.167	0.143
	OTHER	8.76	0.000	0.000	0.000
	Weighted avg.		0.321	0.070	0.115
LAST	A1	37.35	0.649	0.802	0.717
	MNEG	34.69	0.727	0.281	0.405
	MTMP	7.69	0.566	0.746	0.644
	MMNR	5.94	0.807	0.588	0.680
	A2	5.60	0.473	0.851	0.608
	MADV	2.94	0.397	0.585	0.473
	A0	2.28	0.342	0.317	0.329
	MLOC	0.69	0.181	0.500	0.266
	MEXT	0.69	0.800	0.333	0.470
	A4	0.69	0.818	0.750	0.783
	A3	0.56	0.526	1.000	0.689
	MDIR	0.41	0.385	0.714	0.500
	MPNC	0.28	0.132	1.000	0.233
	MDIS	0.19	0.067	0.333	0.112
	Weighted avg.		0.649	0.584	0.615

respectively, is 129. We created a system (FOC-DET) training with bagging over standard C4.5 decision trees and using the full set of features. We used the implementation in the Weka software pacakage (Hall *et al.* 2009).

6.1.2 Detailed results

Table 9 provides per class precision, recall and f-measure using baselines A1, FIRST and LAST. A1 only predicts A1 or MNEG as focus, and the overall performance is low (f-measure 0.375). FIRST is not a sound baseline (f-measure 0.115), and simply choosing the last role (LAST baseline) yields an f-measure of 0.615. LAST and BASIC obtain very similar performance (Table 10).

Table 10 provides per class precision, recall and f-measure using BASIC and FOC-DET. These systems obtain weighted f-measures of 0.611 and 0.641 respectively.

Table 10. *Per-class precision, recall and f-measure using* `BASIC` *and* `FOC-DET`

		BASIC			FOC-DET		
	Focus (%)	Precision	Recall	f-measure	Precision	Recall	f-measure
A1	37.35	0.643	0.804	0.714	0.658	0.837	0.736
MNEG	34.69	0.604	0.385	0.470	0.634	0.407	0.496
MTMP	7.69	0.565	0.725	0.635	0.640	0.797	0.710
MMNR	5.94	0.738	0.972	0.839	0.814	0.981	0.890
A2	5.60	0.496	0.663	0.568	0.536	0.802	0.643
MADV	2.94	0.523	0.434	0.474	0.563	0.340	0.424
A0	2.28	0.385	0.122	0.185	0.333	0.122	0.179
MLOC	0.69	0.250	0.167	0.200	0.500	0.250	0.333
MEXT	0.69	0.875	0.583	0.700	0.667	0.167	0.267
A4	0.69	0.818	0.750	0.783	0.750	0.750	0.750
A3	0.56	0.538	0.700	0.609	0.615	0.800	0.696
MDIR	0.41	0.250	0.286	0.267	0.444	0.571	0.500
MPNC	0.28	0.000	0.000	0.000	0.000	0.000	0.000
MDIS	0.19	0.000	0.000	0.000	0.000	0.000	0.000
Weighted avg.		0.606	0.615	0.611	0.636	0.646	0.641

MPNC and MEXT, the two less frequent roles corresponding to the focus (Table 6), are never predicted (together they only correspond to 0.47 percent of instances). On the other hand, `FOC-DET` successfully predicts A1, the most likely focus, with an f-measure of 0.736 (precision: 0.658, recall: 0.837). MNEG, the second most likely role to correspond to the focus obtains an f-measure of 0.496 (precision: 0.634, recall: 0.407).

6.2 Incorporating negation into composition of semantic relations

In this section we evaluate the Instantiate-Axiom algorithm (Section 4.2.3). As a set of axioms A to be instantiated, we use the eight inference axioms automatically obtained by CSR algorithm over PropBank (Table 3). As relation instances *rels*, we use annotation present in PropBank for verbs y, which have a relation of the form MNEG(x, y). In order to avoid inferring a relation that might already be present in the annotation provided by PropBank, we do not instantiate an axiom $\text{R}_1(x, y) \circ \text{R}_2(y, z) \rightarrow \text{R}_3(x, z)$ if a relation of the form $\text{R}_3(x', z)$ is already present.

Even in a vast corpus like PropBank (112,917 annotated predicates), the number of instantiations found is small. Out of the 3,993 predicates marked with MNEG, only 103 contain PURPOSE and 123 CAUSE. Taking into account the above constraint, there are forty-three instantiations of axioms involving CAUSE (axioms 1–3) and 114 of axioms involving PURPOSE (axioms 4–8).

Accuracies for instantiating each axiom are depicted in Table 11. The remainder of this Section illustrates real inferences using the Instantiate-Axiom algorithm. For each example, we provide the sentence, partial representation and inferred relation:

Table 11. *Evaluation of the Instantiate-Axiom algorithm (Section 4.2.3)*

No.	Axiom	Instances	Accuracy
1.	CAUSE $\circ$ AGENT^{-1} $\rightarrow$ AGENT^{-1}	16	0.69
2.	CAUSE $\circ$ LOCATION^{-1} $\rightarrow$ LOCATION^{-1}	6	0.83
3.	CAUSE $\circ$ TIME^{-1} $\rightarrow$ TIME^{-1}	21	0.86
1–3	CAUSE $\circ$ R_2 $\rightarrow$ R_3	43	0.79
4.	PURPOSE $\circ$ AGENT^{-1} $\rightarrow$ AGENT^{-1}	26	0.77
5.	PURPOSE $\circ$ THEME^{-1} $\rightarrow$ THEME^{-1}	64	0.78
6.	PURPOSE $\circ$ LOCATION^{-1} $\rightarrow$ LOCATION^{-1}	1	1.00
7.	PURPOSE $\circ$ TIME^{-1} $\rightarrow$ TIME^{-1}	19	0.79
8.	PURPOSE $\circ$ MANNER^{-1} $\rightarrow$ MANNER^{-1}	4	0.75
4–8	PURPOSE $\circ$ R_2 $\rightarrow$ R_3	114	0.78
1–8	All	157	0.78

- *If you're homeless, you don't sleep for fear of being robbed or murdered.*
  – AGENT^{-1}($\sim$*sleep*, *you*) & CAUSE(*for fear of being robbed [. . .]*, $\sim$*sleep*)
  Axiom CAUSE$\circ$AGENT^{-1} $\rightarrow$ AGENT^{-1} yields AGENT^{-1}(*for fear of being robbed or murdered*, *you*). [$P^{intrinsic}_{\text{CAUSE}} = +$ and $P^{separable}_{\text{AGENT}^{-1}} = +$, so $P^{neg}_{R_3} = P^{neg}_{\text{CAUSE}} \circ P^{neg}_{\text{AGENT}^{-1}} = 0$].
- *But such operations typically aren't performed because there is no sign right after an injury that surgery would be beneficial.*
  – TIME^{-1}(*performed*, $\sim$*typically*) & CAUSE(*because there is no sign right after an injury that . . . , performed*)
  Axiom CAUSE $\circ$ TIME^{-1} $\rightarrow$ TIME^{-1} yields TIME^{-1}(*because there is no sign right after . . . , typically*). [$P^{intrinsic}_{\text{CAUSE}} = +$ and $P^{separable}_{\text{TIME}^{-1}} = 0$, so $P_{R_3} = 0$].
- *Workers, except for senior management, were asked not to report for work yesterday.*
  – PURPOSE(*for work*, *report*) & TIME^{-1}(*report*, $\sim$*yesterday*)
  Axiom PURPOSE $\circ$ TIME^{-1} $\rightarrow$ TIME^{-1} yields TIME^{-1}(*for work*, $\sim$*yesterday*). [$P^{intrinsic}_{\text{PURPOSE}} = -$ and $P^{separable}_{\text{TIME}^{-1}} = 0$, so $P^{neg}_{R_3} = P^{neg}_{\text{PURPOSE}} \circ P^{neg}_{\text{TIME}^{-1}} = +$].

7 Discussion

In this paper we have presented a model to thoroughly represent the meaning of negation. This model surfaces implicit positive meaning from negated statements and is grounded on detecting the focus of negation, a problem not previously considered. Negation is incorporated into semantic relations using the symbol '$\sim$', indicating if an argument of a relation must be negated. We also devise a methodology to incorporate negation to a previous framework to compose semantic relations. This approach yields a richer representation of text combining negation with semantic relation composition in a unified manner.

The proposed model goes beyond the state of the art. Given *The UFO files weren't released until 1998*, semantic role labelers simply indicate that the verb *released* is negated by *n't*. Scope detectors are a step forward, but by no means encode the correct interpretation of the statement, i.e. *The UFO files were released, but not until 1998*. The proposed representation encodes this interpretation using '∼': THEME(*The UFO files*, *released*) & TIME(∼*until 1998*, *released*).

Using this novel representation is conceptually simple but requires detecting the focus of negation. Because of the lack of corpora, new annotation over PropBank is presented. This annotation can be readily used by any existing semantic role labeler trained over PropBank, quickly incorporating interpretation of negation.

Deciding the focus of negation is a highly ambiguous task. Our annotation guidelines were designed to maintain the original PropBank annotation and select a role as focus. Examples show that deep understanding of text is necessary. Several phenomena influencing this task, most notably context, are also exemplified.

Simple baselines and a more complex feature set to predict the focus of negation have been proposed, showing that the model is feasible. The Instantiate-Axiom algorithm provides a procedure to instantiate an inference axiom when a premise may have an argument negated. It is simple and obtains high accuracy.

One issue remains open: Rewriting a negated argument into its positive counterpart is often difficult. For example, statement s: *[when]*$_{\text{TIME}}$ *[she]*$_{\text{AGENT}}$ *isn't [performing]*$_v$ *[for an audience]*$_{\text{PURPOSE}}$, *she prepares for a song by removing the wad of gum from her mouth, and [...]* is interpreted as *she prepares in that manner when she is performing, but not for an audience*. PURPOSE(∼*for an audience*, *performing*) could be rewritten as PURPOSE(*for herself*, *performing*), but doing so automatically is not without its challenges. The proposed model explicitly encodes *that kind of preparation happens when she is performing (but not for an audience)*, but stating the positive counterpart of *not for an audience* deserves more work.

References

Baker, Collin F., Fillmore, Charles J., and Lowe, John B. 1998. The Berkeley framenet project. In *Proceedings of the 17th International Conference on Computational Linguistics*, Montreal, Canada, pp. 86–90.

Blanco, E., and Moldovan, D. 2011a. A model for composing semantic relations. In *Proceedings of the 9th International Conference on Computational Semantics (IWCS 2011)*, Oxford, UK, pp. 45–54.

Blanco, E., and Moldovan, D. 2011b. 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, pp. 581–589.

Blanco, E., and Moldovan, D. 2011c. Unsupervised learning of semantic relation composition. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL-HLT 2011)*, Portland, OR, pp. 1456–1465.

Bos, J. 2008. Wide-coverage semantic analysis with boxer. In J. Bos, and R. Delmonte (eds.), *Semantics in Text Processing. STEP 2008 Conference Proceedings*, vol. 1. *Research in Computational Semantics*, pp. 277–86. London: College Publications.

Bos, J., and Spenader, J. 2011. An annotated corpus for the analysis of VP ellipsis. *Language Resources and Evaluation* **45**(2): 1–32.

Boucher, J., and Osgood, Charles E. 1969, February. The pollyanna hypothesis. *Journal of Verbal Learning and Verbal Behavior*, **8**(1): 1–8.

Carreras, X., and Màrquez, L. 2004, May. Introduction to the CoNLL-2004 shared task: semantic role labeling. In Hwee T. Ng, and E. Riloff (eds.), *HLT-NAACL 2004 Workshop:*

Eighth Conference on Computational Natural Language Learning (CoNLL-2004), Boston, MA, pp. 89–97. Stroudsburg, PA: Association for Computational Linguistics.

Carreras, X., and Màrquez, L. 2005. Introduction to the CoNLL-2005 shared task: semantic role labeling. In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CONLL '05)*, Morristown, NJ, pp. 152–64. Stroudsburg, PA: Association for Computational Linguistics.

Choi, Jinho D., Bonial, C., and Palmer, M. 2010. Propbank instance annotation guidelines using a dedicated editor, Jubilee. In *Proceedings of the Seventh Conference on International Language Resources and Evaluation (LREC'10)*, Valletta, Malta, pp. 1871–1875.

Cohen, Paul R., and Loiselle, Cynthia L. 1988. Beyond ISA: structures for plausible inference in semantic networks. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, St. Paul, MN, pp. 415–420.

Councill, I., McDonald, R., and Velikovich, L. 2010, July. What's great and what's not: learning to classify the scope of negation for improved sentiment analysis. In *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*, Uppsala, Sweden, pp. 51–9. Antwerp, Belgium: University of Antwerp.

Dang, Hoa T., and Palmer, M. 2005, June. The role of semantic roles in disambiguating verb senses. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, Ann Arbor, MI, pp. 42–9. Stroudsburg, PA, USA: Association for Computational Linguistics.

Dowty, D. 1994. The role of negative polarity and concord marking in natural language reasoning. In *Proceedings of Semantics and Linguistics Theory (SALT) 4*, Rochester, NY, pp. 114–44.

Farkas, R., Vincze, V., Móra, G., Csirik, J., and Szarvas, G. 2010, July. 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*, Uppsala, Sweden, pp. 1–12. Stroudsburg, PA, USA: Association for Computational Linguistics.

Fillmore, C. J. 1976, October. Frame semantics and the nature of language. *Annals of the New York Academy of Sciences* (Origins and Evolution of Language and Speech), **280**: 20–32.

Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., and Witten, Ian H. 2009, November. The weka data mining software: an update. *SIGKDD Exploration Newsletter*, **11**(1): 10–18.

Hendrickx, I., Kim, Su N., Kozareva, Z., Nakov, P., O' Seaghdha, D., Padó, S., Pennacchiotti, M., Romano, L., and Szpakowicz, S. 2009, June. SemEval-2010 task 8: multi-way classification of semantic relations between pairs of nominals. In *Proceedings of the Workshop on Semantic Evaluations: Recent Achievements and Future Directions (SEW-2009)*, Boulder, Colorado, pp. 94–9. Stroudsburg, PA: Association for Computational Linguistics.

Hintikka, J. 2002. Negation in logic and in natural language. *Linguistics and Philosophy*, **25**(5/6): 585–600.

Horn, Laurence R. 1989, June. *A Natural History of Negation.* Chicago, IL: University of Chicago Press.

Horn, Laurence R., and Kato, Y. (eds.) 2000, October. *Negation and Polarity – Syntactic and Semantic Perspectives (Oxford Linguistics)*, New York: Oxford University Press.

Hu, M., and Liu, B. 2004. Mining and summarizing customer reviews. In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (KDD '04), New York, NY, pp. 168–77. New York, NY: ACM.

Huddleston, Rodney D., and Pullum, Geoffrey K. 2002, April. *The Cambridge Grammar of the English Language*. Cambridge, UK: Cambridge University Press.

Huhns, Michael N., and Stephens, Larry M. 1989. Plausible inferencing using extended composition. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence (IJCAI'89)*, San Francisco, CA, pp. 1420–5. Burlington, MA: Morgan Kaufmann.

Jackendoff, R. 1972. *Semantic Interpretation in Generative Grammar*. Cambridge, MA: MIT Press.

Jia, L., Yu, C., and Meng, W. 2009. The effect of negation on sentiment analysis and retrieval effectiveness. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, *(CIKM '09)*, New York, NY, pp. 1827–30. New York, NY: ACM.

Kingsbury, P., Palmer, M., and Marcus, M. 2002. Adding semantic annotation to the Penn TreeBank. In *Proceedings of the Human Language Technology Conference*, San Diego, CA.

Koomen, P., Punyakanok, V., Roth, D., and Yih, Wen T. 2005, June. Generalized inference with multiple semantic role labeling systems. In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*, Ann Arbor, MI, pp. 181–4. Stroudsburg, PA: Association for Computational Linguistics.

Ladusaw, William A. 1996. Negation and polarity items. In S. Lappin (ed.), *Handbook of Contemporary Semantic Theory*, pp. 321–41. Malden MA: Blackwell.

Lang, J., and Lapata, M. 2010, June. Unsupervised induction of semantic roles. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, Los Angeles, CA, pp. 939–47. Stroudsburg, PA: Association for Computational Linguistics.

Li, J., Zhou, G., Wang, H., and Zhu, Q. 2010, August. Learning the scope of negation via shallow semantic parsing. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, Beijing, China, pp. 671–9. Stroudsburg, PA: ACL (Coling 2010 Organizing Committee).

Löbner, S. 2000, June. Polarity in natural language: predication, quantification and negation in particular and characterizing sentences. *Linguistics and Philosophy*, **23**(3): 213–308.

Marcus, M., Santorini, B., and Marcinkiewicz, Mary A. 1994. Building a large annotated corpus of English: the Penn Treebank. *Computational Linguistics*, **19**(2): 313–30.

Màrquez, L., Carreras, X., Litkowski, Kenneth C., and Stevenson, S. 2008, June. Semantic role labeling: an introduction to the special issue. *Computational Linguistics*, **34**(2): 145–59.

Merlo, P., and Van der Plas, L. 2009, August. Abstraction and generalisation in semantic role labels: PropBank, VerbNet or both? In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, Suntec, Singapore, pp. 288–96. Stroudsburg, PA: Association for Computational Linguistics.

Meyers, A., Reeves, R., Macleod, C., Szekely, R., Zielinska, V., Young, B., and Grishman, R. 2004. Annotating noun argument structure for NomBank. In *Proceedings of LREC-2004*, Lisbon, Portugal, pp. 803–806.

Morante, R. 2010, May. Descriptive analysis of negation cues in biomedical texts. In *Proceedings of the Seventh Conference on International Language Resources and Evaluation (LREC'10)*, Valletta, Malta. Paris, France: European Language Resources Association (ELRA), pp. 1429–1436.

Morante, R., and Blanco, E. 2012. *SEM 2012 shared task: resolving the scope and focus of negation. In Proceedings of **SEM 2012: The First Joint Conference on Lexical and Computational Semantics*, Montréal, Canada, 7–8 June, pp. 265–74. Stroudsburg, PA: Association for Computational Linguistics.

Morante, R., Liekens, A., and Daelemans, W. 2008, October. Learning the scope of negation in biomedical texts. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, Honolulu, Hawaii, pp. 715–24. Stroudsburg, PA: Association for Computational Linguistics.

Morante, R., Schrauwen, S., and Daelemans, W. 2011. Corpus-based approaches to processing the scope of negation cues: an evaluation of the state of the art. In *Proceedings of the 9th International Conference on Computational Semantics (IWCS 2011)*, Oxford, UK, pp. 350–354.

Morante, R., and Sporleder, C. (eds.) 2010, July. *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*. Antwerp, Belgium: University of Antwerp.

Nielsen, Leif A. 2004. Verb phrase ellipsis detection using automatically parsed text. In *Proceedings of the 20th International Conference on Computational Linguistics (COLING '04)*, Stroudsburg, PA, Stroudsburg, PA: Association for Computational Linguistics, pp. 1093–1099.

Øvrelid, L., Velldal, E., and Oepen, S. 2010, August. Syntactic scope resolution in uncertainty analysis. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, Beijing, China, pp. 1379–87. Stroudsburg, PA: ACL (Coling 2010 Organizing Committee).

Özgür, A., and Radev, Dragomir R. 2009, August. Detecting speculations and their scopes in scientific text. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, Singapore, pp. 1398–407. Stroudsburg, PA: Association for Computational Linguistics.

Palmer, M., Gildea, D., and Kingsbury, P. 2005. The proposition bank: an annotated corpus of semantic roles. *Computational Linguistics*, **31**(1): 71–106.

Pang, B., Lee, L., and Vaithyanathan, S. 2002, July. Thumbs up? sentiment classification using machine learning techniques. In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing*, pp. 79–86. Stroudsburg, PA: Association for Computational Linguistics.

Quirk, R., Greenbaum, S., Leech, G., and Svartvik, J. 1985. *A Comprehensive Grammar of the English Language*. Harlow, UK: Longman.

Rooth, M. 1985. *Association with Focus*, PhD thesis, Univeristy of Massachusetts, Amherst.

Rooth, M. 1992. A theory of focus interpretation. *Natural Language Semantics*, **1**: 75–116.

Rose, Carolyn P., Roque, A., Bhembe, D., and Vanlehn, K. 2003. A hybrid text classification approach for analysis of student essays. In *Building Educational Applications Using Natural Language Processing*, Stroudsburg, PA: Association for Computational Linguistics, pp. 68–75.

Sánchez Valencia, V. 1991. *Studies on Natural Logic and Categorial Grammar*, PhD thesis, University of Amsterdam, Netherlands.

Sandu, G. 1994. Some aspects of negation in English. *Synthese*, **99**: 345–60.

Saurf, R., and Pustejovsky, J. 2007. Determining modality and factuality for text entailment. In *Proceedings of the International Conference on Semantic Computing, (ICSC '07)*, Washington, DC, pp. 509–16. Washington, DC: IEEE Computer Society.

Saurí, R., and Pustejovsky, J. 2008. From structure to interpretation: a double-layered annotation for event factuality. In *Proceedings of the 2nd Linguistic Annotation Workshop (The LAW II), LREC 2008*, Marrakech, Morocco, pp. 1–8.

Saurí, R., and Pustejovsky, J. 2009, September. FactBank: a corpus annotated with event factuality. *Language Resources and Evaluation*, **43**(3): 227–68.

Szarvas, G., Vincze, V., Farkas, R., and Csirik, J. 2008, June. The bioscope corpus: annotation for negation, uncertainty and their scope in biomedical texts. In *Proceedings of the Workshop on Current Trends in Biomedical Natural Language Processing*, Columbus, OH, pp. 38–45. Stroudsburg, PA: Association for Computational Linguistics.

van Munster, E. 1988. The treatment of scope and negation in Rosetta. In *Proceedings of the 12th International Conference on Computational Linguistics*, Budapest, Hungary, pp. 442–447.

Wiegand, M., Balahur, A., Roth, B., Klakow, D., and Montoyo, A. 2010, July. A survey on the role of negation in sentiment analysis. In *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*, Uppsala, Sweden, pp. 60–8. Antwerp, Belgium: University of Antwerp.

Winston, Morton E., Chaffin, R., and Herrmann, D. 1987. A taxonomy of part-whole relations. *Cognitive Science*, **11**(4): 417–44.

Zapirain, Be N., Agirre, E., and Màrquez, L. 2008, June. Robustness and generalization of role sets: PropBank vs. VerbNet. In *Proceedings of ACL-08: HLT*, Columbus, OH, pp. 550–558. Stroudsburg, PA: Association for Computational Linguistics.

Zeijlstra, H. 2007. Negation in natural language: on the form and meaning of negative elements. *Language and Linguistics Compass*, **1**(5): 498–518.