

ConanDoyle-neg: Annotation of negation in Conan Doyle stories

Roser Morante and Walter Daelemans

CLiPS - University of Antwerp
Prinsstraat 13, B-2000 Antwerp, Belgium
{Roser.Morante,Walter.Daelemans}@ua.ac.be

Abstract

In this paper we present ConanDoyle-neg, a corpus of stories by Conan Doyle annotated with negation information. The negation *cues* and their *scope*, as well as the event or property that is negated have been annotated by two annotators. The inter-annotator agreement is measured in terms of F-scores at scope level. It is higher for cues (94.88 and 92.77), less high for scopes (85.04 and 77.31), and lower for the negated event (79.23 and 80.67). The corpus is publicly available.

Keywords: Negation, scopes, corpus annotation

1. Introduction

In this paper we present ConanDoyle-neg, a corpus of Conan Doyle stories annotated with negation cues and their scope. The annotated texts are *The Hound of the Baskervilles* (HB) and *The Adventure of Wisteria Lodge* (WL). The original texts are freely available from the Gutenberg Project at <http://www.gutenberg.org/browse/authors/d/#a37238>. The main reason to choose this corpus is that part of it has been annotated with coreference, semantic roles and null instantiations of semantic roles for the SemEval Task *Linking Events and Their Participants in Discourse* (Ruppenhofer et al., 2010). In this way, negation complements the annotations undertaken for the shared task. Another reason is that there is a lack of corpora annotated with negation information for texts from domains other than the biomedical domain.

Negation is a grammatical category that comprises devices used to reverse the truth value of propositions. In natural language, negation functions as an operator along with quantifiers and modals. The main characteristic of operators is that they have a scope: elements to which negative, modals and quantifiers refer are in the scope of the negative operator.

The study of negation from a philosophical point of view started with Aristotle and nowadays is still a topic that generates a considerable number of publications in philosophy, logic, psycholinguistics, and linguistics. Horn (1989) provides an extensive description of negation from a historic perspective and an analysis of negation in relation to semantic and pragmatic phenomena. Tottie (1991) studies negation in English from a descriptive and quantitative point of view, based on the analysis of empirical material. She defines two main types of negation in natural language, *rejections* of suggestions and *denials* of assertions, which can be *explicit* and *implicit*.

Languages have devices to negate entire propositions (*clausal negation*) or constituents of clauses (*constituent negation*). Most languages have several grammatical devices to express clausal negation, which are used with different purposes such as negating existence versus negating facts, or negating different aspects, modes or speech acts (Payne, 1997). Negation is not only a grammatical phe-

nomenon present in all languages. As (Lawler, 2010) puts it, “negation is a linguistic, cognitive, and intellectual phenomenon. Ubiquitous and richly diverse in its manifestations, it is fundamentally important to all human thought”. Negation is a frequent phenomenon in language. Tottie reports that negation is twice as frequent in spoken text (27,6 per 1000 words) as in written text (12,8 per 1000 words). Councill et al. (2010) annotate a corpus of product reviews with negation information and they find that 19% of the sentences contain negations (216 out of 1135). In the ConanDoyle-neg corpus 22.49% of sentences contain at least one negation.

The interest in automatically processing negation first originated in the medical domain as a response to the need of automatically processing and indexing clinical reports and discharge summaries. For this task it is very relevant to find negated symptoms, signs, treatments, and outcomes. Interest in the biomedical text mining community to extract accurate information about biological relations has boosted the research on negation processing. The release of the BioScope corpus (Vincze et al., 2008) has allowed to develop negation scope resolvers for biomedical texts. The corpus gathers clinical free-texts, biological full papers, and biological paper abstracts annotated with negation cues, i.e., words that express negation, and their scope. Blanco and Moldovan (2011) take a different approach by annotating the *focus*, “that part of the scope that is most prominently or explicitly negated”, in the 3,993 verbal negations signaled with MNEG in the PropBank corpus. According to the authors, the annotation of the focus allows to derive the implicit positive meaning of negated statements. For example, in (1) the focus of the negation is on *until 2008*, and the implicit positive meaning is ‘They released the UFO files in 2008’.

(1) They didn’t release the UFO files until 2008.

However, there is a lack of resources annotated with negation information for general domain texts. We consider this to be of great importance because negation adds information about an extra-propositional aspect of meaning. Processing negation is essential in order to know whether an event is presented as factual or counterfactual. The ex-

amples in (2) show two sentences with the same propositional meaning, but with opposite truth values. A standard propositional representation of the sentences such as the produced by a semantic role labeler would not capture this difference.

- (2) a. Watson knows who the assassin is.
b. Watson does not know who the assassin is.

The ConanDoyle-neg corpus as well as the complete guidelines are available for download at <http://www.clips.ua.ac.be/BiographTA/corpora.html>. The corpus is in TIGER/SALSA xml format (Erk and Padó, 2004), which is the format of the corpus of the *SemEval Task Linking Events and Their Participants in Discourse*.

In Section 2. we present examples of negation in the ConanDoyle-neg corpus. In Section 3. we summarize the annotation guidelines that have been followed to annotate the corpus. Section 4. describes the annotation process and presents inter-annotator agreement scores. Finally, in Section 5., we put forward some conclusions.

2. Negation in ConanDoyle-neg

Wouden (1997) shows that negation can be expressed by a variety of lexical and syntactic categories. We list in (3) some examples of negation in the ConanDoyle-neg corpus, which are expressed with different lexical categories. If there is more than one cue in the sentence, we mark only the information about the cue that is the focus of attention in the example. The scope is marked with square brackets, the cue in bold, and the negated event, if there is one, is underlined.

- (3) a. **Verbs:** [I] **fail** to [see how you could have done more].
b. **Adverbs:** [It was] suggested, but **never** [proved, that the deceased gentleman may have had valuables in the house, and that their abstraction was the motive of the crime].
c. **Prepositions and prepositional phrases:** [The woman's disappearance counts] **for nothing**, since in that extraordinary household any member of it might be invisible for a week.
d. **Determiners:** [To us there is] **no** [fiend in hell like Juan Murillo], and no peace in life while his victims still cry for vengeance.
e. **Pronouns:** [The various bedrooms and sitting-rooms had yielded] **nothing** [to a careful search].
f. **Conjunctions:** It wasn't black, sir, **nor** [was it white], nor any colour that I know but a kind of queer shade like clay with a splash of milk in it.

Negation can also be expressed by affixes, as the examples in (4) show.

- (4) a. I could look straight through the **un**[curtained window].
b. The whole **in**[explicable tangle] seemed to straighten out before me.

Negation cues can occur in a variety of syntactic constructions, as shown in (3). However, the presence of a negation cue in certain constructions does not imply that the events within the scope of the cue are negated. We do not consider the events *think* in (5.a), *mean* in (5.b), and *return* in (5.c) to be negated because they do not occur in factual contexts.

- (5) a. **Imperatives:** Pray [do]**n't** [think it a liberty if I give you a word of friendly warning].
b. **Questions:** [You do]**n't** [mean that Baynes has got him]?
c. **Conditionals:** But the attempt was a dangerous one, and if [Garcia did] **not** [return by a certain hour] it was probable that his own life had been sacrificed.

Additional examples of events that are not negated despite being within the scope of a negation cue are shown in (6). In (6.a) we do not consider the event *help* to be negated by the cue *without* because it is under the scope of the modal *should* and because from the context we know that the event *solved* has not taken place yet. In (6.b) the event *come* is not negated by the cue *not* because of the third conditional construction, where the presence of the negation cue indicates that the event happened. The same case applies to the event *looking out* in (6.c), which is under the scope of *not*.

- (6) a. I should be glad to be able to say afterwards that I had solved it **without** [your help].
b. I call it luck, but [it would] **not** [have come my way had I not been looking out for it].
c. I call it luck, but it would not have come my way [had I] **not** [been looking out for it].

Information about the corpus is presented in Table 1. In HB there are 3640 sentences out of which 850 contain negations. There are 1.15 negation cues per sentence. In WL there are 783 sentences out of which 145 contain negations. There are 1.2 negation cues per sentence.

Story	# sent	# neg sent	#cues	# scopes	# negated
Hound	3640	850	985	888	617
Wisteria	783	145	175	170	123

Table 1: Corpus statistics.

3. Annotation guidelines

In this section we provide a brief summary of the annotation guidelines (Morante et al., 2011) that annotators followed to annotate the corpus. The annotation tasks consist in annotating the negation *cues* and their *scope*, as well as the event or property that is negated. The cues are the words that express negation and the scope is the part of a sentence that is affected by the negation words. The final goal of annotating negation cues and their scope is to determine which events in the sentence are affected by the negation. We will use the term *event* in a very general way, since we take an event to be a process, an action, or a state. (7) shows an annotated sentence.

- (7) [I do]**n't** [know what made me look up], but there was a face looking in at me through the lower pane.

Only factual events can be annotated as negated. As defined in Saurí and Pustejovsky (2009), a factual event corresponds to a fact in the world. In (8.a) the event SAY is a fact and is negated, whereas in the rest of examples in (8) the event SAY is not a fact. In (8.b) the event is presented in

the future, in (8.b) it is presented as a condition for another event, and in (8.d) it is within the scope of a modality cue. In these cases, although we mark the negation cue and the scope, we do not mark the event as a negated fact.

- (8) a. [He had] **never** [said as much before]
 b. [He will] **never** [say as much]
 c. If [he does] **not** [say anything], we will consider the case closed
 d. It is unclear whether [he did] **not** [say as much before]

The current annotation focuses on sentential negation, but in the future it will be extended to intersentential negation as in (9), where the negation cue *No* of sentence (9.b) expresses a rejection of the statement in (9.a).

- (9) a. Don't blame me.
 b. No, sir; I believe you mean well by me.

The annotation style is inspired by the guidelines of the BioScope corpus¹, but there are several differences between the two annotations. A main difference is that the BioScope corpus does not annotate the negated event. Another important difference is that in the ConanDoyle-neg corpus the scope model is different than the model in the BioScope corpus. The cue is not considered to be part of the scope, the scope can be discontinuous, and all arguments of the event being negated are considered to be within the scope, including the subject, which is kept out of the scope in the BioScope corpus. Additionally, elided elements that belong to the scope are recovered from the context, as shown in (10), where *I could* is marked as part of the scope of *nor*. Another difference is that affixal negation is annotated in the ConanDoyle-neg corpus whereas it is not annotated in the BioScope corpus.

- (10) I tell you, sir, [I **could**]n't move a finger, **nor** [get my breath], till it whisked away and was gone.

Because the guidelines that we present here are based on the annotation of Conan Doyle stories, some constructions that are characteristic of other domains are left out, such as constructions that express absence of an entity, which are very frequent in clinical notes. An example of this construction from the BioScope corpus is shown in (11.a). Other expressions of negation that occur in the BioScope corpus and do not occur in the Conan Doyle corpus are shown in (11.b,c).

- (11) a. **No** [focal consolidation to suggest pneumonia].
 b. [Right middle lobe abnormalities suggest] airways disease **rather than** [bacterial pneumonia].
 c. **It is not the case that** [right middle lobe abnormalities suggest bacterial pneumonia].

3.1. Negation cues

Negation cues can be single words, multiwords (12), prefixes, such as *im-* in (13.a), or suffixes, such as *-less* in (13.b).

- (12) The story of the Stapletons could **no longer** be withheld from him, but he took the blow bravely when he learned the truth about the woman whom he had loved.

- (13) a. You saw me, perhaps, on the night of the convict hunt, when I was so **imprudent** as to allow the moon to rise behind me?
 b. In that impassive colour**less** man, with his straw hat and his butterflynet, I seemed to see something terrible ... [· · ·]

Multiword negation cues are often fixed expressions, for instance signifying the gradation of negation like *by no means* in (14).

- (14) But that was all.
 No, no, my dear Watson, not all - **by no means** all.

Words that have the same form as negation cues but different function are not marked as negation cues. Some examples are shown in (15).

- (15) a. He would talk of *nothing but* art, of which he had the crudest ideas, from our leaving the gallery until we found ourselves at the Northumberland Hotel.
 b. "Why about Sir Henry in particular? *I could not help* asking.
 c. "Only a joke, *as like as not*."
 d. You have been inside the house, *have you not*, Watson?
 e. "*Don't* tell me that it is our friend Sir Henry!"
 f. Partly it came *no doubt* from his own masterful nature, which loved to dominate and surprise those who were around him

3.2. Scope

The annotation of the scope has several characteristics. First, the negation cue is not marked within the scope. We aim at distinguishing clearly what part of the sentence is affected by the change of polarity. Since a cue does not change its own polarity, it should be kept out of the scope. Second, the longest relevant scope of the negation cue is marked. Third, the scope can be discontinuous. The parts of the sentence that are not affected by the negation cue should be left out of the scope. For example, discourse markers such as *yet* in (16) are never included in the scope because they are not affected by the negation:

- (16) a. "Then you use me, and yet [do] **not** [trust me]!" I cried with some bitterness.
 b. "And yet [have been] **un**[able to save him]!"

Finally, the annotation of negation cues and their scope should reflect the meaning of statements. This is why constructions that express the same meaning should be annotated similarly. For example, the sentences in (17) are analysed similarly. In both cases the negated event is *Watson was prudent*. (17.a) uses a negative prefix attached to the adjective *prudent*, whereas (17.b) uses the adverb *not* modifying the verb.

- (17) a. [Watson was] **im**[prudent].
 b. [Watson was] **not** [prudent].

Active and passive constructions, which in the annotation of the BioScope corpus are treated differently, are annotated following the same criteria. The sentence in (18.1) has active voice, and the sentence in (18.2) passive, but in both cases all the arguments of the verb are considered to be within the scope of the negation cue.

¹The annotation guidelines of BioScope are available at <http://www.inf.u-szeged.hu/rgai/project/nlp/bioscope/Annotation%20guidelines2.1.pdf>

- (18) a. [Watson did] **not** [solve the case].
 b. [The case was] **not** [solved by Watson].

In order to know what part of the sentence is within the scope of a negation cue annotators are asked to apply a paraphrase test using the “It is not the case” construction followed by the part of the sentence that is within the scope. For example, the sentence in (18) should be paraphrased as the sentence in (19).

- (19) a. It is not the case that [Watson did solve the case].
 b. It is not the case that [the case was solved by Watson].

When a verb is negated, like in (20), the full clause is within the scope of the negation. In (20.a) *not* scopes over the clause *we did drive up to the door*, which includes the subject and complement of the verb *drive*, as shown by the paraphrase in (20.b)

- (20) a. [We did] **not** [drive up to the door].
 b. **It is not the case that** [we drove up to the door].

If the negated verb is the main verb of the sentence, the entire sentence is within the scope of the cue. In (21) the verb *be able* is negated, which is the main verb. Prepositional complements *in half an hour* and *in front of us* have to be included in the scope.

- (21) [In half an hour we wo]n’t [be able to see our hands in front of us].

In coordinate clauses, negation cues scope always only over their clause, as in (22.a), and in subordinate clauses (22.b), also, excluding the subordinate particle.

- (22) a. [We did] **not** [drive up to the door] but got down near the gate of the avenue.
 b. I suppose that [it does] **not** [always run loose upon the moor].

When the subject is negated, the negation scopes over all the clause, as in (23).

- (23) **No** [sign could we see of the man whom we were chasing].

When the object of a verb is negated, the negation scopes over the clause headed by the verb. In (24.a) the cue *no* is syntactically attached to the object (*reason*), but it scopes over the clause. We interpret the sentence in (24.a) as expressing the same meaning as the one in (24.b).

- (24) a. [I see] **no** [reason for further concealment].
 b. [I do] **not** [see any reason for further concealment].

As for negation scoping over adjectives (25), if the adjective is embedded in a noun phrase, the full noun phrase should be annotated as scope.

- (25) a. “And if we take this as a working hypothesis we have a fresh basis from which to start our construction of [this] **un**[known visitor].”
 b. “As it is, by [an] **in**[discreet eagerness], which was taken advantage of with extraordinary quickness and energy by our opponent, we have betrayed ourselves and lost our man.”

When the adjective is the attribute of a copulative sentence, the negation scopes over the entire clause, as in (26).

- (26) a. You saw me, perhaps, on the night of the convict hunt, [when I was so] **im**[prudent as to allow the moon to rise behind me]?

3.3. Negated event or property

The negated element is the main event or property actually negated by the negation cue, for example the verb *said* in (27).

- (27) [He had] **never** [said as much before] . . .

The negated element has to be a factual event. The determiners, modifiers or auxiliaries of negated events are never marked, in order to minimize the marked elements for negated event annotation and only the head of the phrase to which the negated event belongs is marked. In attributive constructions (28), the attribute is marked as negated instead of the verb “to be”.

- (28) He declares that [he] heard cries but [is] **un**[able to state from what direction they came].

4. Annotation process

The annotation was made using the Salto Tool (Burchardt et al., 2006). A *Negation* frame is created, which is assigned to the negation cue. The frame has two elements: SCOPE and NEGATED_EVENT. In this way the same resources used to annotate semantic frames for the SemEval Task *Linking Events and Their Participants in Discourse* could be used to annotate negation. Figure 1 presents a screenshot of the annotation interface.

The corpus was annotated by two annotators, a master student and a researcher, both of them with a background in linguistics. The researcher produced the annotation guidelines in collaboration with the master student. The two annotators based their annotation on the annotation guidelines. The researcher was also in charge of producing the adjudicated version of the annotation. A preliminary version of the guidelines was produced by analyzing negation in Chapter 10 of *The Hound of the Baskervilles*. The researcher and the master student annotated this chapter separately based on the preliminary guidelines. Disagreements were discussed, as well as new cases and cases that were not properly described in the guidelines. A second preliminary version of the guidelines was produced and then the same annotators and another experienced researcher who had not seen the guidelines before annotated Chapter 11, in order to test their robustness. After a second process of discussion, the guidelines were modified and the first version of the guidelines was produced. Based on this version the full corpus was annotated. The adjudicated version was produced by the researcher. In cases of disagreement between annotators the closest solution to the guidelines indications was selected.

Inter-annotator agreement (IAA) has been calculated in terms of the F-measures, precision, recall and F1. These measures were also reported to calculate IAA at scope level for the BioScope corpus. We provide several scores:

- Cue-level F-measures (‘Cues’). In order to count a true positive, all tokens of the cue have to be correctly identified.
- Scope-level F-measures (‘Scopes’) matching cues and scopes: a scope will be counted true positive if all tokens of the cue and of the scope have been correctly

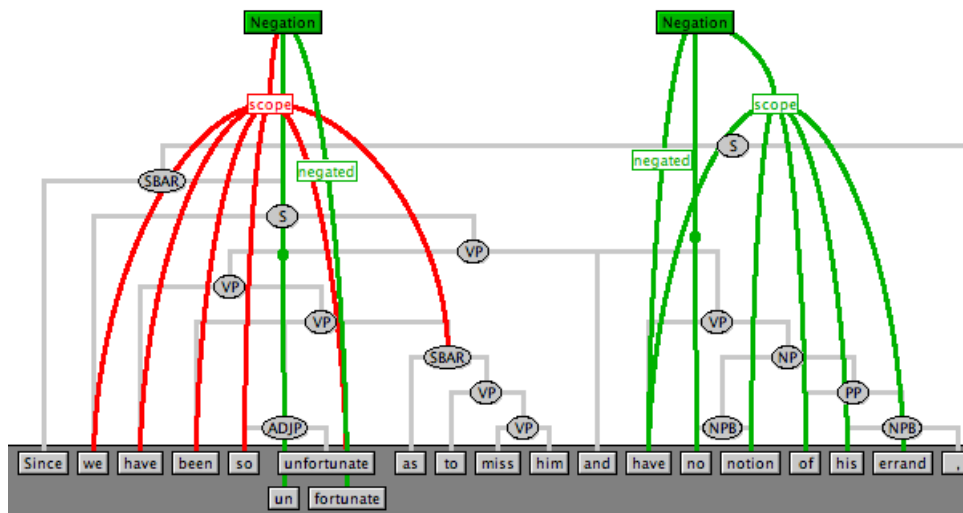


Figure 1: Annotation interface.

identified. An incorrect scope is not counted as false positive, but only as false negative, if at least a token of the cue has been correctly identified.

- F-measures over negated events ('Negated'), computed independently from cues and from the scopes.
- Token-level F-measures ('Scope tokens') to evaluate agreement in scopes.

Table 2 presents IAA when the annotations of the researcher annotator are evaluated against the adjudicated version. Table 3 presents IAA when the annotations of the student annotator are evaluated against the adjudicated version. Finally, Table 4 presents IAA when the annotations of the student annotator are evaluated against the annotations of the researcher annotator.

Hound	precision	recall	F1
Cues	99.38	97.66	98.51
Scopes	99.39	92.45	95.79
Negated	93.56	92.95	93.25
Scope tokens	98.80	96.97	97.88
Wisteria	precision	recall	F1
Cues	100.00	91.43	95.52
Scopes	100.00	78.82	88.16
Negated	94.87	90.24	92.50
Scope tokens	97.15	91.17	94.07

Table 2: Inter-annotator agreement: the annotations of the researcher annotator are evaluated against the adjudicated version.

The IAA scores of the researcher annotator evaluated against the adjudicated version is higher than the IAA scores of the student annotator. This is to be expected, since the researcher has more experience with this type of tasks and has been in charge of defining the annotation guidelines. However, the scores of the student annotator evaluated against the adjudicated version are still high. Given the complexity of the task, the agreement between the two annotators is also high. It is higher for HB than for WL, which might

Hound	precision	recall	F1
Cues	96.15	96.24	96.19
Scopes	94.98	83.11	88.65
Negated	81.82	88.21	84.89
Scope tokens	92.70	94.25	93.47
Wisteria	precision	recall	F1
Cues	97.69	96.57	97.13
Scopes	97.20	81.76	88.81
Negated	85.04	92.31	88.53
Scope tokens	94.38	92.33	93.34

Table 3: Inter-annotator agreement: the annotations of the student annotator are evaluated against the adjudicated version.

Hound	precision	recall	F1
Cues	94.02	95.76	94.88
Scopes	91.98	79.08	85.04
Negated	76.21	82.49	79.23
Scope tokens	89.95	93.17	91.53
Wisteria	precision	recall	F1
Cues	90.59	95.06	92.77
Scopes	87.20	69.43	77.31
Negated	75.59	86.49	80.67
Scope tokens	85.88	90.31	88.04

Table 4: Inter-annotator agreement: the annotations of the student annotator are evaluated against the annotations of the researcher annotator.

be explained by the fact that the annotation guidelines were developed by analyzing Chapter 10 of HB.

The IAA scores are very high for cues, as could be expected since identifying negation cues is a simpler task. The IAA scores reported for the BioScope corpus are also higher for cues in the annotation of full articles (Table 5). The maximum IAA reported for the BioScope corpus is 91.46 for abstracts. The scores obtained for scopes are lower than for cues both in the ConanDoyle-neg corpus and in the full articles section of the BioScope corpus. Since annotating scopes involves making more decisions than annotating

cues, this is to be expected. There is also an effect of error accumulation, since for a scope to be counted as correct, the cue also has to be correctly identified.

	Hound	Wisteria	BioScope abstracts	BioScope full articles
Cues	94.88	92.77	91.46	79.42
Scopes	85.04	77.31	92.46	70.86

Table 5: Inter-annotator agreement between annotators in ConanDoyle-neg compared to inter-annotator agreement in BioScope in terms of F1.

Tagging the negated event has the lowest IAA between annotators in HB. Annotators need to make two decisions: (i) Is there a negated event? (ii) If there is one, which token in the sentence should be marked as such?. To make the first decision annotators need to take into account the context in which the potential negated event occurs and failing to make the proper analysis of the context might be the cause of the errors.

5. Conclusions

We have presented the ConanDoyle-neg, a corpus of stories by Conan Doyle annotated with negation information. The negation *cues* and their *scope*, as well as the event or property that is negated have been annotated by two annotators. The inter-annotator agreement in HB and WL is higher for cues (94.88 and 92.77), less high for scopes (85.04 and 77.31), and lower for the negated event (79.23 and 80.67). In general, the IAA are high, which indicates that the annotation task has been well defined. As future research we will analyze what aspects of annotating negation are more difficult for annotators by making an error analysis of the annotations of the two annotators. We aim at identifying negation constructions that are more difficult to analyze in order to improve the annotation guidelines for future annotations. Additionally, the corpus will be annotated by a third annotator who has not been involved in the process of developing the guidelines. New IAA scores will be computed and compared to the existing ones.

6. Acknowledgements

The research reported in this paper was made possible through financial support from the University of Antwerp (GOA project BIOGRAPH). We are grateful to Sarah Schrauwen for her collaboration in the annotation guidelines and for her annotation work.

7. References

Eduardo Blanco and Dan Moldovan. 2011. Semantic representation of negation using focus detection. In *Proceedings of 49th Annual Meeting of the Association for Computational Linguistics*, pages 19–24, Portland, Oregon, USA.

Aljoscha Burchardt, Katrin Erk, Anette Frank, Andrea Kowalski, and Sebastian Pado. 2006. Salto a versatile multi-level annotation tool. In *Proceedings of LREC 2006*.

Isaac Council, Ryan McDonald, and Leonid Velikovich. 2010. What’s great and what’s not: learning to classify the scope of negation for improved sentiment analysis. In *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*, pages 51–59, Uppsala, Sweden, July. University of Antwerp.

K. Erk and S. Padó. 2004. A powerful and versatile XML format for representing role-semantic annotation. In *Proceedings of LREC-2004*.

Laurence R. Horn. 1989. *A natural history of negation*. Chicago University Press, Chicago.

John Lawler. 2010. Negation and negative polarity. In P. C. Hogan, editor, *Cambridge Encyclopedia of the Language Sciences*, pages 554–555. CUP, Cambridge, UK.

Roser Morante, Sarah Schrauwen, and Walter Daelemans. 2011. Annotation of negation cues and their scope. guidelines v1.0. Technical Report Series CTR-003, CLiPS, University of Antwerp, Antwerp, April.

Thomas E. Payne. 1997. *Describing morphosyntax*. Cambridge University Press, Cambridge, UK.

Josef Ruppenhofer, Caroline Sporleder, Roser Morante, Collin Baker, and Martha Palmer. 2010. Semeval-2010 task 10: Linking events and their participants in discourse. In *Proc. of the 5th International Workshop on Semantic Evaluation*, pages 45–50, Uppsala, Sweden, July. ACL.

Roser Saurí and James Pustejovsky. 2009. FactBank: A corpus annotated with event factuality. *Language Resources and Evaluation*, 43(3):227–268.

Gunnel Tottie. 1991. *Negation in English speech and writing: a study in variation*. Academic Press, New York.

Ton van der Wouden. 1997. *Negative contexts: collocation, polarity, and multiple negation*. Routledge, London.

Veronika Vincze, György Szarvas, Richard Farkas, György Móra, and János Csirik. 2008. The BioScope corpus: biomedical texts annotated for uncertainty, negation and their scopes. *BMC Bioinformatics*, 9((Suppl 11)):S9.