The Scope and Focus of Negation: A Complete Annotation Framework for Italian

Begoña Altuna
University of the Basque Country (UPV/EHU), Spain
begona.altuna@ehu.eus

Anne-Lyse Minard **Manuela Speranza**
Fondazione Bruno Kessler (FBK), Trento, Italy
minard,manspera@fbk.eu

Abstract

In this paper we present a complete framework for the annotation of negation in Italian, which accounts for both negation scope and negation focus, and also for language-specific phenomena such as negative concord. In our view, the annotation of negation complements more comprehensive Natural Language Processing tasks, such as temporal information processing and sentiment analysis. We applied the proposed framework and the guidelines built on top of it to the annotation of written texts, namely news articles and tweets, thus producing annotated data for a total of over 36,000 tokens.

1 Introduction

The digital era has enabled the creation of large amounts of data that can be used in many knowledge fields. These data, however, need to be "understood" to be useful. Natural Language Processing (NLP) aims at analyzing and extracting textual information that can be employed in tasks such as decision making (Demner-Fushman et al., 2009) or event forecasting (Radinsky and Horvitz, 2013) among many others.

The analysis and processing of the negation itself is relevant to areas such as event information processing (Minard et al., 2016), sentiment analysis (Jia et al., 2009) and discourse relation identification (Asr and Demberg, 2015). On the one hand, knowing whether an event is affirmed or negated is of utmost importance in the domain of temporal processing for determining the factuality of an event. The positive or negative polarity of an event, in fact, will determine its factuality value, that is to say, whether an event is said to happen or not. Knowing which part of the sentence is most directly negated, on the other hand, may help discriminate which entities participate in an event, which can be very helpful when building entity-based timelines and storylines (Laparra et al., 2015).

The task of sentiment analysis consists predominantly of determining whether a proposition has a positive or a negative polarity. In this case, the presence of a negation can revert the polarity of the proposition and, thus, its identification is essential. Finally in the domain of discourse analysis, the analysis of the expression of negation is needed when extracting relations between parts of the discourse, for example to find *chose alternative* relations or *contrast* constructions.

While affirmative sentences do not need any syntactic marker, negation is typically expressed by some kind of syntactic or lexical element that changes the polarity of the whole sentence or of some elements of the sentence. Each language has its own means to express negation. Therefore, these have to be identified and their features analyzed as a preliminary step towards the completion of an annotation framework.

We propose a complete annotation framework for the annotation of negation in Italian built on the work by Morante et al. (2011) and Blanco and Moldovan (2011). In our framework the semantics of negation is represented through the identification of the negation cue (i.e. the lexical element expressing negation), its scope (i.e. the text section that is negated), its focus (i.e. that part of the scope that is prominently or explicitly negated) and, if present, its reinforcement (i.e. an auxiliary negation). In (1) we give an example of negation and its annotation[1]. The novelty of our framework

[1] In our examples we will use the following notational conventions: if marked, a negation cue will be highlighted in bold, its reinforcement will be in italics, its focus will be included in square brackets and its scope will be underlined.

Proceedings of the Workshop Computational Semantics Beyond Events and Roles (SemBEaR), pages 34–42,
Valencia, Spain, April 4, 2017.

lies in considering all at once the annotation of the scope and the focus of a negation, which implies making some adaptation of the annotation of one with regard to the other.

(1) **Non** ha parlato [con loro].
(He didn't speak with them.)

We applied this annotation framework on sample sentences taken from written news articles and produced detailed annotation guidelines[2]. Finally, following these annotation guidelines, we annotated two different typologies of written texts, i.e. news articles and tweets, for a total of over 36,000 tokens.

The paper is divided as follows: in Section 2 we summarize the related work on negation annotation, in Section 3 we highlight the main features of the proposed annotation framework, in Section 4 we provide details about the annotation effort we have conducted and in Section 5 we explain how our annotation can be useful for other NLP tasks. Finally, in Section 6 we discuss future work.

2 Related Work

Many annotation efforts on negation have been conducted in recent years for different purposes. Some of the first attempts of negation annotation were done in biomedical text corpora for which annotation guidelines for negation were created. On one hand, the GENIA corpus manual (Kim et al., 2006) was employed for the annotation of GENIA (Kim et al., 2008), a corpus of 1,000 abstracts annotated with negated biological events and three levels of uncertainty. On the other hand, the BioInfer annotation scheme (Pyysalo et al., 2007) was used for the annotation of entities and the relations and dependencies among them in the BioInfer corpus. The absence of such relationships as in "not affected by" or "independent of" was annotated with the special predicate *NOT*.

The biomedical corpus BioScope (Vincze et al., 2008) is the first corpus in which negation was specifically targeted; it consists of 20,000 sentences, 13% of which contain some negative expression. BioScope is annotated not only with negation cues but also with negation scope, whose extent is defined as the largest syntactic unit possible.

The guidelines followed in the annotation of the BioScope corpus have been adapted to different domains. Morante et al. (2011), for example, focused on narrative texts and defined the annotation of negated events in addition to negation cues and their scope. Following this, Morante and Daelemans (2012) created and released the ConanDoyle-neg corpus, a corpus containing Sherlock Holmes' stories annotated with negation and event information, as well as coreference, semantic roles and implicit arguments. In addition, a further adaptation of the BioScope guidelines has been used to annotate the SFU Review Corpus (Konstantinova et al., 2012), a corpus consisting of 400 customer reviews.

PropBank (Palmer et al., 2005) also contains negation information: a *NEG* functional tag has been assigned to the modifiers expressing negation such as "not", "n't", "never" and "no longer", as defined in the PropBank guidelines (Bonial et al., 2010). On top of this annotation Blanco and Moldovan (2011) marked the focus of the negation, defined as the most directly and explicitly negated part of the sentence. The resulting dataset was employed, together with the ConanDoyle-neg corpus, in the *SEM 2012 Shared Task *Resolving the Scope and Focus of Negation*[3].

Among corpora containing annotations of negation and its scope it is worth mentioning the Product Review Corpus (Councill et al., 2010), which was built automatically with a system trained on the BioScope corpus.

While the work mentioned above focuses exclusively on the English language, the only work on negation in Italian we are aware of is that of Attardi et al. (2015) in the medical domain. They annotated a corpus of medical records in Italian with medical entities (diseases, drugs, etc.) and added a marker to indicate whether an entity appears in a negative context. This annotation is limited to the identification of the absence or presence of a medical entity.

However, Spanish and Italian are closely related languages and share many features, which allowed us to take into consideration work on negation in Spanish. Similar to the work on Italian medical records, Stricker et al. (2015) have anno-

[2]The guidelines for the annotation of negation cues, focus and scope in Italian are not public yet, as we are still improving the document in terms of clarity of exposition and examples, but is accessible at the following link: `https://goo.gl/kAmRwN`

[3]`http://www.clips.ua.ac.be/sem2012-st-neg/`

tated radiology reports with findings (observations and medical conditions) and they have assigned a value, "affirmed" or "negated", to each of those findings. More similar to our aim, we can cite the Spanish section of the SFU Review Corpus (Martí et al., 2016), which has been annotated with negation by Jiménez-Zafra et al. (to appear), partially following the ABSA guidelines used for Task 12 in SemEval-2015[4].

3 Annotation Specifications

We propose an annotation scheme for negation in which we have defined the elements to be annotated and their features based on (Morante et al., 2011) and (Blanco and Moldovan, 2011), thus including negation cues, negation scope and focus.

Having Italian as the target language, we relied on the Italian grammar by Serianni and Castelvecchi (1988) and on corpus observation for language-specific phenomena related to negation.

3.1 Negation Cues

Negation is a linguistic phenomenon that inverts the truth value of the proposition it is applied to (Martí et al., 2016). Negation is usually expressed by lexical and syntactic elements that are called negation cues.

Following Morante et al. (2011), only negation cues expressed by adverbs, such as *non* / "not", pronouns, as *nessuno* / "nobody", determiners, as *nessun* / "any", and prepositions, as *senza* / "without" have been taken into consideration in our annotation effort. Negation expressed by verbs or nouns, on the other hand, (as in *Rifiuto* / "I refuse", see example (2)), remains out of our scope. It is also relevant to note that we have not addressed affixal negation (e.g. the negative prefixes *in-* and *a-*, see example 3) as we do not want to go below the unit of a token.

(2) *Rifiuto* di parlare.
(I refuse to talk.)

(3) Un *im*percettibile odore inondava la stanza.
(An imperceptible smell invaded the room.)

As a result of exhaustive research on negation cues in Italian, we have compiled a list that includes both one-word constructions (e.g *non* / "not", see example (4)) and multi-word expressions (e.g. *per niente* / "(not) at all", see example (5)).

(4) Il dato **non** è ancora preciso.
(The data is not precise yet.)

(5) Era una donna **per niente** candida.
(*She was a woman not at all candid.)

Following Morante et al. (2011) we do not annotate as negation cues those negative forms that do not actually express negation, such as the expletive *non* in *non appena* / "as soon as" or *non* in fixed constructions such as *non a caso* / "not by chance".

In general, every negation cue is associated with its scope and focus. Ellipsis, unfinished sentences and other phenomena, however, may prevent it from happening as it is the case of *no* / "no" when it is a one-word answer to a question. In (6), for example, *no* is annotated as a negation cue with no focus or scope as the reference to winning is expressed in the previous sentence, while the annotation of the scope and focus of a negation does not go beyond sentence boundaries.

(6) Avete vinto? **No**.
(Did you win? No.)

3.2 Scope

As Morante et al. (2011) do, we consider the scope of a negation cue to be the extent of the text affected by the cue; more specifically, the scope of a negation corresponds to the section of text expressing the proposition whose truth value is inverted by the negation. The suggested test to determine the extent of the scope "it is not the case that" proposed by Morante et al. is also employed for Italian (*non si dà il caso*).

As a general rule, the negation cue remains out of the scope since it does not change its own polarity (7)[5]. However, an innovative feature of our framework is that negation cues which carry a richer semantic meaning than plain negation are included inside the scope; this is the case, among others, of *nessun* / "no (determiner)" (8), *mai* / "never", *nessuno* / "nobody", and *nulla* / "nothing".

We have taken the decision of including such negation cues in the scope because they convey

[4] http://alt.qcri.org/semeval2015/task12/

[5] Note that, as a consequence of this, the extent of the scope can be discontinuous, as in (7).

more than just a negative meaning. When turning negative sentences into affirmative sentences, the plain negation cues will be removed whereas richer semantic negation cues will be replaced by a positive counterpart. For example the affirmative version of the sentence (7) will be *Il presidente tratta con gli assassini* / "The president deals with the murderers", and for the sentence (8) *Qualche militare italiano é rimasto ferito* / "Some Italian soldiers have been wounded".

(7) Il presidente **non** tratta con gli assassini.
(The president does not deal with murderers.)

(8) **Nessun** militare italiano è rimasto ferito.
(No Italian soldier has been wounded.)

3.3 Focus

Focus is defined as that part of the scope that is most prominently or explicitly negated (Huddleston and Pullum, 2002); as an example, *con gli assassini* / "with the murderers" is the focus of *non* / "not" in (9).

(9) Il presidente **non** tratta [con gli assassini].
(The president does not deal with murderers.)

In some cases it is possible that the extent of the focus coincides with the negation cue; this happens with negation cues which not only express negation but carry a richer semantic meaning. For example, *mai* / "never" expresses a reference to time, while *nessuno* / "nobody" expresses a reference to human beings (10).

(10) [**Nessuno**] ha cercato di fermare l'uomo.
(Nobody tried to stop the man.)

It is worth underlining that this is perfectly in tune with the decision explained above (Section 3.2) to include these negation cues within the extent of the scope.

3.4 Reinforcement

In Italian (as in other romance languages such as Spanish), negations precede the verb (11 and 12). When the negation is moved after the verb as in (13), an auxiliary negation (reinforcement) is added to fill the position that has been left empty. In this case, we annotate the negation cue and associate it with the reinforcement, besides associating it with its focus and scope as in all other cases.

(11) I militari italiani **non** sono [rimasti feriti].
(Italian soldiers have not been wounded.)

(12) [**Nessun** militare italiano] è rimasto ferito.
(No Italian soldier has been wounded.)

(13) *Non* è rimasto ferito [**nessun** militare italiano].
(No Italian soldier has been wounded.)

3.5 Discussion

In our framework we address both scope and focus. Since, by definition, the focus is the most prominently negated part of the scope, we explicitly added the constraint according to which the focus should always be included in the scope.

In cleft sentences though, the focus is detached from the clause where the negation cue is placed; in (14), for instance, the focus is *dal 30 agosto* / "since the 30th of August". As a result, the focus would be outside the extent of the scope (the annotation of the scope, in fact, does not go over clause boundaries). To ensure that the focus is included within the scope, we decided to expand the extent of the scope to include as well the detached part of the cleft sentence.

(14) È [dal 30 agosto] che **non** si può più comprare.
(It is since the 30th of August that it is no longer possible to buy it.)

The annotation of relative pronouns and the elements they refer to (noun, pronoun or phrase) is also worth a more detailed discussion. In order to be aligned with the annotation framework proposed by Morante et al. (2011), we decided to include relative pronouns in the scope, but not their antecedents. For example, in (15), we have annotated *che hanno voluto andarci* / "who did want to go there" as the scope.

However, one might argue that the inclusion of the antecedents (*i bambini che hanno voluto andarci* / "the children who did want to go there") would have made the scope more informative and that the "it is not the case that" test (see Section 3.2) suggests to include *bambini* / "children" in the scope.

(15) Sono i bambini [che] **non** hanno voluto andarci.
(It is the children who did not want to go there.)

4 Annotating Negation

The annotation framework described above has been applied to an annotation task which included two significantly different types of texts, i.e. news articles and tweets.

4.1 Negation in Fact-Ita Bank

We annotated with negation 71 documents of Fact-Ita Bank (Minard et al., 2014), a corpus which consists of 169 news stories taken from Ita-TimeBank (Caselli et al., 2011).

From Ita-TimeBank, it inherited the annotation of events, which was performed following It-TimeML (Caselli et al., 2011), the Italian version of the TimeML annotation scheme.

For a subset of 6,958 events, Fact-Ita Bank contains the annotation of the factuality attributes (i.e. polarity, time and certainty) as defined for FactA - Factuality Annotation (Minard et al., 2016), a task which has been organised in 2016 in the context of the EVALITA evaluation campaign[6]. Fact-Ita Bank has been used as training corpus for FactA and is distributed with a CC-BY-NC license[7].

4.2 Negation in Tweets

We annotated with negation 301 tweets that were used as test set for the FactA pilot task on social media texts (Minard et al., 2016).

Also in this case, the texts contained the annotation of events (following It-TimeML) and of the event factuality attributes (as defined for the FactA task at EVALITA 2016).

4.3 Annotation Task

The annotation task has been performed using CAT[8] (Content Annotation Tool) (Lenzi et al., 2012), a web-based text annotation tool. The annotated data are in an XML-based stand-off format where different annotation layers are stored in separate document sections and are related to each other and to source data through pointers.

Four annotators were involved in the annotation task. We estimate the annotation effort to be ten working days of an expert annotator.

[6] http://www.evalita.it/

[7] http://hlt-nlp.fbk.eu/technologies/fact-ita-bank

[8] http://dh.fbk.eu/resources/cat-content-annotation-tool

4.4 Inter-Annotator Agreement

When we had completed a first version of the guidelines we tested the inter-annotator agreement (IAA) with three annotators (who had been involved in the definition of the task) over 8 news articles from Fact-Ita Bank, for a total of 47 negation cues[9] (IAA-1).

We computed the F-measure on the exact match for each annotator pair and for each markable (negation cue, scope and focus). Hripcsak and Rothschild (2005) shows that for tasks in which the number of negative cases is unknown, undefined or very large, inter-annotator agreement can be quantified using the average pairwise F-measure. The first column in Table 1 shows the average of the pairwise F-measure values obtained, which is 0.93, 0.52 and 0.55 for the negation cue, the scope and the focus, respectively.

	IAA - 1	IAA - 2
documents	8	4
# negation cues	47	30
negation cue	0.93	0.98
scope	0.52	0.67
focus	0.55	0.58

Table 1: IAA in terms of average pairwise F-measure.

As we were not completely satisfied with the results, we improved the annotation guidelines and enriched them with examples taken from the dataset used for the first test. Then, in order to evaluate the improvement, we produced a small gold standard (4 news articles from Fact-Ita Bank) annotated by two expert annotators (who had been involved in the previous test) and had it annotated by another person who had the improved version of the guidelines as its only source of information (IAA-2).

The second column in Table 1 shows the results of this experiment in terms of F-measure. The agreement on the annotation of the scope in IAA-2 is much better than in IAA-1, with a average F-measure computed on the strict match close to 0.7 (for the scope) and close to 0.6 (for the focus). Moving to a relaxed match (acceptance of one-word difference when comparing two strings) the average F-measure for the scope increases to 0.85 and for the focus it reaches 0.77.

[9] The number of negation cues was computed after the annotators completed the adjudication.

4.5 Discussion

In total we annotated 71 news articles from Fact-Ita Bank, including those used for the IAA, and 301 tweets. In Table 2 we present a quantitative description of the data. We can observe that the average size of scope and focus is bigger in news than in tweets. This is mainly due to the limitation of characters in tweets and to the writing style, which is closer to oral speech, with very short sentences. Not surprisingly, in both corpora the most frequent negation cue is *non* / "not".

	news articles	tweet corpus
docs	71	301
tokens	31,596	4,920
sentences	1,290	301
tweets/sent. w. neg.	278	59
negation cues	282	71
reinforcement	15	9
average size scopes	9.11	4.69
average size focus	3.2	1.61
non [not]	76%	80%
nessun(o/a) [no/nobody]	6%	3%
nulla/niente [nothing]	4%	8%
senza [without]	6%	4%

Table 2: Quantitative data about the annotated corpora

During the annotation and the inter-annotator agreement phases, we noticed that the annotation of the focus in written texts is a very difficult task, even for humans. Taking into account certain linguistic phenomena can help in interpreting a negative sentence to some extent. For example, the fact that a subject pronoun (which is usually omitted in Italian) is expressed in a sentence indicates that the focus is on the subject itself (e.g. in (16) the focus is on the pronoun *io* / "I"). Word order can also be used to determine the focus of a negation, but prosody is undoubtedly the most useful aspect. Since we work on written texts, and do not have this kind of data, our focus annotation strongly relies on the interpretation of the annotator, which decreases inter-annnotator agreement.

(16) [Io] **non** sono d'accordo che abbiano nominato grand Budapest hotel e il libro della vita
(I don't agree with the fact that they have nominated grand Budapest hotel and The Book of Life)

On the other hand, the annotation of the scope is a more straightforward task. In the first inter-annotator agreement phase the agreement for the scope annotation was low due to some imprecision in the guidelines, as well as small issues in the management of nested annotation by the annotation tool. The main disagreements were related to i) the inclusion or not of the negation cues in the scope, ii) the annotation of cleft sentences, and iii) the treatment of parenthetical texts. However, in the second inter-annotator agreement phase, the disagreements concerned mainly the discourse connectives which should be excluded from the scope but were not excluded by one annotator.

The annotation of tweets enabled us to observe new phenomena in the expression of negation and so to add some annotation rules. The main differences between news articles and tweets are the size of the text, and as a result the amount of context information available, and the style (which in tweets is close to that of oral speech, with the use of slang and sometimes vulgar language). In news articles, sentences are well written and often quite long; usually the reader has all the context needed to understand a piece of information. On the other hand, sentences in tweets are very short and sometimes incomplete. Incompleteness can lead to focus ambiguity and even to the absence of the focus. In (17), for example, there are dots where the focus should be. We decided to annotate the verb that is negated, but one could argue that the dots should be annotated as focus instead or that no focus should be associated with the negation cue.

(17) Il Modena ha fatto vedere buone cose ma **non** è . . .
(Modena has shown good stuff but it is not. . .)

(18) #Paritàsessi **non** è [sfoggiare ascelle pelose] o [#pisciare in un imbuto per farlo in piedi].
(* #Equalrights is not showing off hairy armpits or peeing in a funnel to do it standing up.)

Another difference between tweets and news articles is the use of non standard language. In tweets we find abbreviations, repeated words, non alpha-numeric symbols, grammar mistakes, and sometimes missing words. When annotating tweets, for instance, we added to our list a nega-

tion cue that we had not found before, *nn*, which stands for *non* / "not". Tweets also contain hashtags which are used to link tweets to some category topics and they can include a negation. For example the hashtag *#Nonbeccomailaporta* / "I never strike the target" could be decomposed and annotated with negation: "mai" will be the negation cue and the focus, "non" the reinforcement and "beccomailaporta" the scope. But at the moment, as the annotation tool does not manage the annotation of units smaller than a token, we do not annotate it. Finally, one tweet contained the only case we found of negation of coordinated phrases (e.g. in (18)).

5 Relevance of Negation Annotation

As mentioned in the introduction, our research on negation is motivated by the interest of employing it for temporal information processing and, more speciically, for the processing of events and their factuality value.

The identification of the scope may help in factuality resolution. In our corpus, which has been previously annotated with temporal information, directly negated events like *uccideranno* / "will kill" (19) are given a negative factuality value. Events like *ha rivendicato* / "has claimed responsibility" in (20), instead, have been given a positive polarity as they are not directly negated. In fact, they have a positive factuality value, although they are implicitly counterfactual, since, in this case no responsibility claim has been done.

(19) **Non** [uccideranno] il nostro futuro.
(They will not kill our future.)

(20) [**Nessuno**] ha rivendicato il sequestro.
(Nobody has claimed responsibility for the hijack.)

In example (20), the event *ha rivendicato* falls under the scope of the negation and its factuality value has changed. Either because the event is directly negated (19) or because an argument of that event is negated (20), the final factuality value of an event will be negative.

However, the scope is not enough to decide on the factuality value of an event. *sequestro* / "hijack" in (20) falls also into the scope extent but it preserves its positive polarity, since it is a subordinated event and the negation affects the main clause. Therefore, we consider scope information to be useful for factuality resolution, but it has to be complemented by other linguistic information such as sentence structure and argument information.

As far as focus is concerned, we assume that in some cases the identification of the focus may help build entity-based timelines, that is to say, timelines that gather and organize the events in which a certain entity participates. As counterfactual or non-factual events have not happened or will not happen, we will exclude those from the timeline.

When an entity is the focus of the negation, we hold that it does not take part in the event, since it is explicitly negated. As a consequence, that event will not be considered for the timeline of that entity. If we were to build a timeline from example (21) taking "Putin" as the target entity, we should include *ha detto* / "has said" and *essere* / "is" in the timeline. On the contrary, *si è recato* / "has attended" will not appear in the timeline, since it is explicitly mentioned, through the negation of the subject *lui* / "he", that he did not attend the funerals. Finally, the event "funerals" has a factual value (i.e. the funerals took place), but since Putin did not go, they will not appear in Putin's timeline.

(21) Putin ha detto di essere col cuore a Beslan, anche se [lui] **non** si è recato ai funerali.
(Putin said his heart is in Beslan, even though he did not attend the funerals.)

Although we have only worked on news documents, we expect that the processing of negation will also be useful for sentiment analysis (e.g.in movie or book reviews). The identification of negation scope may help in defining the polarity of the events in the scope, which is a highly relevant feature in these kinds of texts.

6 Conclusion

In this paper we presented our work on defining a common annotation framework for scope and focus of negations and the annotation performed on two corpora: Fact-Ita Bank, which is composed of news articles, and a corpus of tweets. We conducted this work on the Italian language but we plan to use the annotation framework to perform annotation in other languages, in particular in Spanish, French and Basque.

The interpretation of negation is an important task for detecting the factuality of events and we now have corpora annotated with both negation and factuality at our disposal. In the short term we

expect to conduct a study about the relationship between negation scope and factuality annotation.

As far as focus annotation is concerned, we will use that information for the identification of the directly negated entities in order not to include them in role structures. This information will also be used to improve the entity-based timelines and storylines, keeping the events in which negated entities participate out of candidate events to form the timeline.

Our framework does not include the annotation of negation expressed by verbs or nouns (e.g. *cancel*) and affixal negation (e.g. *illegal*). We plan to include these aspects as well and, consequently, to verify to what extent the current annotation guidelines account for their annotation of scope and focus.

Negation does not always have the same intensity and can be total or partial. Some words increase the intensity of the negation and other reduce it. For example the negation in "not all the students arrived" is partial, whereas in "he did not arrive" the negation is total. For the moment, all these cases are annotated in the same way and the different nuances are not considered, but we intend to add some markers of degree to the negation cues, so as to normalize this information.

Finally, the annotated data will be soon made available from the website of the HLT-NLP group at FBK (`http://hlt-nlp.fbk.eu/`) and used to implement and evaluate a system for negation detection in Italian.

References

Fatemeh Torabi Asr and Vera Demberg. 2015. Uniform information density at the level of discourse relations: Negation markers and discourse connective omission. In *Proceedings of the International Conference on Computation Semantics*, pages 118–128.

Giuseppe Attardi, Vittoria Cozza, and Daniele Sartiano. 2015. Annotation and extraction of relations from italian medical records. In *Proceedings of the 6th Italian Information Retrieval Workshop (IIR 2015)*.

Eduardo Blanco and Dan Moldovan. 2011. Semantic Representation of Negation Using Focus Detection. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 581–589, Portland, Oregon, USA, June. Association for Computational Linguistics.

Claire Bonial, Olga Babko-Malaya, Jinho D. Choi, Jena Hwang, and Martha Palmer. 2010. Propbank Annotation Guidelines (Version 3.0). Technical report, Center for Computational Language and Education Research, Institute of Cognitive Science, University of Colorado at Boulder.

Tommaso Caselli, Valentina Bartalesi Lenzi, Rachele Sprugnoli, Emanuele Pianta, and Irina Prodanof. 2011. Annotating Events, Temporal Expressions and Relations in Italian: the It-Timeml Experience for the Ita-TimeBank. In *Proceedings of the 5th Linguistic Annotation Workshop*, pages 143–151, Portland, Oregon, USA, June. Association for Computational Linguistics.

Isaac Councill, 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.

Dina Demner-Fushman, Wendy W. Chapman, and Clement J. McDonald. 2009. What can natural language processing do for clinical decision support? *Journal of Biomedical Informatics*, 42(5):760 – 772.

George Hripcsak and Adam S. Rothschild. 2005. Technical brief: Agreement, the f-measure, and reliability in information retrieval. *Journal of the American Medical Informatics Association*, 12(3):296–298.

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

Lifeng Jia, Clement Yu, and Weiyi Meng. 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, pages 1827–1830, New York, NY, USA. ACM.

Salud M. Jiménez-Zafra, Mariona Taulé, M. Teresa Martín-Valdivia, L. Alfonso Ureña-López, and M. Antònia Martí. to appear. SFU ReviewSP-NEG: a Spanish corpus annotated with negation for Sentiment Analysis A Typology of negation patterns. *Language Resources and Evaluation*.

Jin-Dong Kim, Tomoko Ohta, Yuka Tateisi, and Jun'ichi Tsujii. 2006. GENIA corpus manual. Technical report, Tsujiilab, University of Tokyo.

Jin-Dong Kim, Tomoko Ohta, and Jun'ichi Tsujii. 2008. Corpus annotation for mining biomedical events from literature. *BMC bioinformatics*, 9(1):1.

Natalia Konstantinova, Sheila C.M. de Sousa, Noa P. Cruz, Manuel J. Maña, Maite Taboada, and Ruslan Mitkov. 2012. A review corpus annotated for negation, speculation and their scope. In Nicoletta Calzolari, Khalid Choukri, Thierry Declerck,

Mehmet Uğur Doğan, Bente Maegaard, Joseph Mariani, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC-2012)*, pages 3190–3195, Istanbul, Turkey, May. European Language Resources Association (ELRA).

Egoitz Laparra, Itziar Aldabe, and German Rigau. 2015. From TimeLines to StoryLines: A preliminary proposal for evaluating narratives. In *Proceedings of the First Workshop on Computing News Storylines*, pages 50–55, Beijing, China, July. Association for Computational Linguistics.

Valentina Bartalesi Lenzi, Giovanni Moretti, and Rachele Sprugnoli. 2012. CAT: the CELCT Annotation Tool. In Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Mehmet Uğur Doğan, Bente Maegaard, Joseph Mariani, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC-2012)*, pages 333–338, Istanbul, Turkey, May. European Language Resources Association (ELRA).

M. Antònia Martí, M. Teresa Martín-Valdivia, Mariona Taulé, Salud María Jiménez-Zafra, Montserrat Nofre, and Laia Marsó. 2016. La negación en español: análisis y tipología de patrones de negación. *Procesamiento del Lenguaje Natural*, 57:41–48.

Anne-Lyse Minard, Alessandro Marchetti, and Manuela Speranza. 2014. Event Factuality in Italian: Annotation of News Stories from the Ita-TimeBank. In *Proceedings of CLiC-it 2014, First Italian Conference on Computational Linguistic*, pages 260–264.

Anne-Lyse Minard, Manuela Speranza, and Tommaso Caselli. 2016. The EVALITA 2016 Event Factuality Annotation Task (FactA). In Pierpaolo Basile, Franco Cutugno, Malvina Nissim, Viviana Patti, and Rachele Sprugnoli, editors, *Proceedings of the 5th Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2016)*. aAcademia University Press.

Roser Morante and Walter Daelemans. 2012. Conandoyle-neg: Annotation of negation cues and their scope in conan doyle stories. In Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Mehmet Uğur Doğan, Bente Maegaard, Joseph Mariani, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC-2012)*, pages 1563–1568, Istanbul, Turkey, May. European Language Resources Association (ELRA).

Roser Morante, Sarah Schrauwen, and Walter Daelemans. 2011. Annotation of negation cues and their scope: Guidelines v1. *Computational linguistics and psycholinguistics technical report series, CTRS-003*.

Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational linguistics*, 31(1):71–106.

Sampo Pyysalo, Filip Ginter, Juho Heimonen, Jari Björne, Jorma Boberg, Jouni Järvinen, and Tapio Salakoski. 2007. BioInfer: a corpus for information extraction in the biomedical domain. *BMC Bioinformatics*, 8(1):50.

Kira Radinsky and Eric Horvitz. 2013. Mining the web to predict future events. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 255–264. ACM.

Luca Serianni and Alberto Castelvecchi. 1988. *Grammatica italiana: italiano comune e lingua letteraria, suoni, forme, costrutti*. Utet.

Vanesa Stricker, Ignacio Iacobacci, and Viviana Cotik. 2015. Negated findings detection in radiology reports in spanish: an adaptation of negex to spanish. In *Workshop on Replicability and Reproducibility in Natural Language Processing: adaptive methods, resources and software at IJCAI 2015*.

Veronika Vincze, György Szarvas, Richárd Farkas, György Móra, and János Csirik. 2008. The bioScope corpus: biomedical texts annotated for uncertainty, negation and their scopes. *BMC bioinformatics*, 9(11):1.