

Cues, Scope, and Focus: Annotating Negation in Spanish Corpora

Inductores, Ámbito, y Foco: La Anotación de la Negación en los Corpus de Español

Lucia Donatelli

Georgetown University
Bunn Intercultural Center, 403A 37th and O Streets, N.W.
Washington, DC 20057
led66@georgetown.edu

Abstract: The objective of NEGES Task 1 is to establish a standard for the annotation of negation in Spanish-language corpora. Specifically, the task analyzes guidelines used for such annotation in five projects over three domains: news (Sandoval and Salazar, 2013), clinical reports (Oronoz et al., 2015; Cruz et al., 2017; Marimon, Vivaldi, and Bel, 2017), and product reviews (Jiménez-Zafra et al., 2018b). Here, an assessment of these various guidelines is presented, with the goals of helping establish a standard set of guidelines for annotating negation in Spanish across domains, and of contributing to the workshop’s overall conversation about the treatment of negation in computational linguistics.

Keywords: Linguistic annotation, negation, Spanish

Resumen: El objetivo de la Tarea 1 de NEGES es establecer un estándar para anotar la negación en los corpus en español. En particular, la tarea analiza las directrices empleadas para tal anotación en tres dominios repartidas en cinco proyectos de corpus: noticias (Sandoval and Salazar, 2013), informes clínicos (Oronoz et al., 2015; Cruz et al., 2017; Marimon, Vivaldi, and Bel, 2017), y opiniones de productos (Jiménez-Zafra et al., 2018b). Aquí se presenta una evaluación de las distintas directrices, con los fines de ayudar a establecer una norma de directrices que se deben seguir para la anotación de negación en español entre dominios, y de contribuir a la conversación más amplia del taller sobre el tratamiento de la negación en la lingüística computacional.

Palabras clave: Anotación lingüística, negación, español

1 Introduction

Negation is fundamental to sentence meaning, bearing on the key questions of “what happened (and what did not)”, “who did (or did not do) what to whom,” and “what existed (or did not).” In other words, negation helps establish what is fact and what is not, due to its ability to affect the truth value of a sentence (Horn, 1989). This is important to tasks in Natural Language Processing (NLP) based on the accurate identification and representation of meaning encoded in language, such as information extraction, question answering, and sentiment analysis.

The annotation of negation is not a trivial task. Negation acts at the syntactic, semantic, and pragmatic levels (Sandoval and Salazar, 2013); and it exhibits asymmetry between a uniform semantics and a catego-

rial polyvalence, appearing as a prefix, verb, determiner, adverbial, particle, idiom, construction, and more (Bosque Muñoz and Gutiérrez-Rexach, 2009; Herburger, 2018). Often, the *scope*, or span of utterance upon which negation acts, is variable even within the same sentence; some argue that scope is different when determined at logical form (LF) (the syntax/semantics interface) as opposed to pragmatically (Moeschler, 2010), forcing resulting annotation to distinguish between underlying semantics and more general speaker meaning. Additionally, negation may carry different force depending on context; this is particularly true for Spanish, in which negative concord may offer gradations of meaning (Jiménez-Zafra et al., 2018b). This may be seen by the variety of negative expressions in the following senten-

ces, with negative linguistic elements in bold (Española-RAE, 2010):

(1) *Ella **no** dijo **nada***. ‘She did not say anything.’

(2) ***Nadie** le hacía caso*. ‘Nobody paid attention to him/her.’

(3) ***Ni** de una forma **ni** de otra consiguieron convencerla*. ‘They couldn’t convince her one way or another.’

(4) ***En mi vida** he visto cosa igual*. ‘I haven’t seen anything similar in my life.’

(5) ***No** hables tanto*. ‘Don’t talk so much.’

(6) *¿**No** son ya las dos?* ‘Isn’t it already two o’clock?’

Example (1) negates the action of speaking; *no* functions as a negation cue, and *nada* as a negative polarity item (NPI) within the scope of the cue in the form of a negative indefinite. (2) affirms the matter that no one, signaled by the indefinite pronoun *nadie*, participated in an attention-giving activity. (3) exhibits two negative conjunctions, *ni*, that describe an ineffective manner of convincing. (4) exhibits an adverbial phrase, *en mi vida*, that equates to the negative temporal adverb *nunca* ‘never’. Finally, (5) and (6) utilize negation to support speech acts: (5) as a negative command, and (6) as a leading question.

Given this complexity in mind, the process of annotating negation in Spanish is daunting. Nevertheless, it is essential to annotate corpora with such information in order to train algorithms to perform to human capacity. Ideally, an annotation framework could be designed to capture negation patterns cross-linguistically. However, since such patterns are quite varied and distinct, the task here focuses on annotating negation in Spanish as a first step towards potentially broader research.

2 Task Definition

NEGES task 1 seeks to reach an agreement on the guidelines to follow for the annotation of negation in Spanish building off previous, domain-specific guidelines used to annotate corpora built around news, clinical reports and product reviews (Jiménez-Zafra et al., 2018a). Here I present brief summaries and analyses of the guidelines presented in the five projects in question, followed by a complete evaluation of the guidelines across projects.

3 Review of Guidelines

3.1 UAM Spanish TreeBank

Sandoval and Salazar (2013) present findings from annotating the UAM Spanish Treebank (Moreno et al., 2003), which consists of 1501 sentences taken from newspapers (*El País* and *Compra Muestra*) and annotated syntactically following guidelines from the Penn Treebank (Marcus, Marcinkiewicz, and Santorini, 1993). 10.67% of the sentences were found to contain negation (160 sentences).

Given that the UAM Treebank is in xml format, the authors adapt their annotation scheme to be compatible. The authors distinguish two levels of annotation for negation: sentence-level and lexical. The former is further divided into sentential and phrase-level negation; the latter is divided into pronouns and adverbs. Annotations thus mark both negation cues and the scope of negation.

The use of a syntactically annotated corpus is helpful for the overall annotation scheme; this is especially true given the widespread knowledge of the Penn Treebank, as well as the concurrent annotation of lexical features that specify POS. Theoretically, such annotation will help identify the lexical category of negation marker as well as its syntactic scope of negation. Algorithms trained on such data will then be able to recognize patterns of use of negation for varying parts of speech and compare how syntactic scope relates to semantic and pragmatic scope.

As the authors themselves note, there is much room for improvement in the annotation scheme presented. First, the type of negation is not specified apart from <YES> and <NO> (to mark scope), and <NEG> (to mark negation cues). This fails to capture gradations of negation (i.e. whether it is an assertion or a speech act that is negated (Moeschler, 2010)), differences in intensity of negation (Jiménez-Zafra et al., 2018b), and the function of the NPI within the scope of negation (Herburger, 2018). Additionally, there is no specification of event negation, negation expressed morphologically, or negative discourse and sentence connectors.

3.2 IxaMed-GS

Ornoz et al. (2015) focus on the identification of entities and events in clinical reports with the goal of automatic extraction of adverse drug reaction events using machine

learning. Annotators were experts in pharmacology and pharmacovigilance, a notable difference from annotators trained in linguistics for other corpora in their knowledge of the clinical domain yet informal training in identifying functional linguistic elements such as negation.

The authors collected 142,154 anonymous discharge reports from the outpatient consultations of the Galdake-Usansolo Hospital from 2008 to 2012. Negation (and speculation) was annotated as a modifier of a disorder or drug, and individual cues were left unmarked. As such, the text span *sin otras alergias medicamentosas* “without other drug allergies” would possess negation on *alergias medicamentosas* while *sin* would be left bare. This practice was used to maintain consistency in the domain: a disease-entity such as *afebril* “afebrile” was also marked as negative in the absence of surrounding negative lexical material.

Four entity types were annotated: diseases, allergies, drugs, and procedures. For diseases and allergies, a distinction was made between negated entity, speculated entity and entity (for non-speculative and non-negated entities). 2,362 diseases were annotated, out of which 490 (20.75%) were tagged as negated diseases and 40 (1.69%) as speculated diseases. 404 allergy entities were identified, of which 273 (67.57%) were negated and 13 (3.22%), speculated. The quality of the annotation process was assessed by measuring the inter-annotator agreement (IAA), which was 90.53% for entities and 82.86% for events.

This annotation scheme needs to be adapted in order to extend beyond the clinical domain. First and foremost, negation needs to be treated linguistically and broken down into its components apart from the disorders and drugs it acts on. Nevertheless, the recognition that entities and events may be marked for negation in distinct ways (i.e. syntactically versus morphologically, as in the above examples), as well as the intuition that some entities and events may possess qualities of negation without being explicitly marked for it, is an important contribution to developing a comprehensive annotation scheme. Additionally, the distinction between negation and speculation is important to consider, as the linguistic interaction between negation and modality is complicated, yet merits attention.

3.3 SFU ReviewSP-NEG

Jiménez-Zafra et al. (2018b) present the SFU ReviewSP-NEG corpus, the first Spanish corpus that includes event negation as part of the annotation scheme as well as the annotation of discontinuous negation markers. The corpus was also the first to have defined a typology of patterns involving negation specific to Spanish. In this SFU ReviewSP-NEG corpus, syntactic negation, scope, focus, and event were annotated. Yet, annotations on the event and on how negation affects the polarity of the words within its scope were included for whether there is a complete change in the polarity of the span in question, or an increment or reduction of its value.

The Spanish SFU Review corpus, originally intended for work on sentiment analysis, consists of 400 reviews extracted from the website *ciao.es*. The reviews span 8 different produce areas: cars, hotels, washing machines, books, cell phones, music, computers, and movies. For each product area there are 50 positive and 50 negative reviews, which provides an informative context for how to interpret the effects of negation at the discourse level (Taboada, Anthony, and Voll, 2006).

For the SFU ReviewSP-NEG corpus, each review was automatically annotated at the token level with POS-tags and lemmas; negation cues and their corresponding scopes and events were manually annotated at the sentence level. The annotations were performed by two senior researchers with in-depth experience in corpus annotation who supervised the two trained annotators who carried out the annotation task. The final corpus is composed of 9,455 sentences, out of which 3,022 sentences (31.97%) contain at least one negation marker. The Kappa coefficient for IAA was of 0.97 for negation cues, 0.95 for negated events and 0.94 for scopes.

Similar to the UAM Spanish Treebank, the annotation scheme for SFU ReviewSP-NEG ought to be broadened to include morphological, lexical, and discourse-oriented negation. Nevertheless, the inclusion of gradient interpretations of negation within its scope captures the subtle meaning differences that negation markers, and their combinations, may produce. In fact, this gradation of meaning may be expanded to include even finer-grained distinctions of meaning in future work.

3.4 UHU-HUVR

Cruz et al. (2017) annotate a corpus composed of 604 clinical reports from the Virgen del Rocío Hospital in Sevilla, Spain. 276 of these clinical documents correspond to radiology reports and 328 to the personal history of anamnesis reports written in free text. Two domain expert annotators closely followed the Thyme corpus guidelines (Styler IV et al., 2014), developed for the annotation of English clinical records. In the anamnesis reports, 1,079 sentences (35.20%) were found to contain negations out of 3,065 sentences. On the other hand, 1,219 sentences (22.80%) out of 5,347 sentences were annotated with negations in the radiology reports. The Dice coefficient for IAA was higher than 0.94 for negation markers and higher than 0.72 for negated events.

In this corpus, all types of negation were annotated: syntactic, morphological (affixal negation), and lexical. Negation was marked both linguistically and as a modifier of a disorder or a drug, i.e. whether or not the drug was effective. Similar to (Oronoz et al., 2015), full words that expressed negative polarity were marked in their entirety (*afebril* “afebrile”) rather than just their negative affix (*a-*).

Either due to the domain of application or presentation format, the annotation guidelines presented for the UHU-HUVR corpus often seemed unclear. For example, non-clinical experts may have trouble differentiating negative test results from negative clinical events, as the annotation scheme does. Additionally, while some negation affixes are marked as negating symptoms (*a-febril* “afebrile”) others are not, considered positive symptoms unto themselves (*in-continencia urinaria* “urinary incontinence”). Finally, the authors’ treatment of coordination as a single unit of negation ought to be revised.

3.5 IULA Spanish Clinical Record

The IULA Spanish Clinical Record corpus (Marimon, Vivaldi, and Bel, 2017) contains 300 anonymized clinical records from several services of one of the main hospitals in Barcelona, Spain. The corpus was annotated with negation markers and their scopes with the ultimate goal of extracting factual knowledge from textual data; subgoals included automatic encoding of clinical records; diagnosis support; term extraction; and general study

of clinical texts. The corpus contains 3,194 sentences, out of which 1,093 (34.22%) were annotated with negation cues. In this corpus, syntactic negation and lexical negation were annotated; morphological negation was excluded. Annotators were three computational linguists and a clinician.

Annotators did not include the negation cue nor the subject in its scope as part of annotation, unless the subject was located after the verb. This practice seems to be based on linear order alone, and does not take into account the semantics of scope nor the possibility of backwards scope (Hoeksema, 2000). Additionally, it seems necessary to mark the negation cue in some manner to signal where the negation is coming from in order to better train algorithms. This aside, the annotation of scope seemed to be quite precise (for example, annotating scope over verb phrase versus just over adverb), and the project on the whole was presented very thoroughly. The authors note that they did not annotate certain verbs with negative polarity (*desaparecer*, *retirar*, *suspenderse*, *eradicar*, *negar*) on the basis that such verbs still denoted factuality. Such interactions between negation and factuality seem worth while to discuss for future annotation efforts.

Similar to (Oronoz et al., 2015) and (Cruz et al., 2017), the IULA was biased towards the clinical domain. Thus, teasing apart the effects of negative affixation (for example, in the adjective *asintomático* “asymptomatic”) will be necessary for future work to both be faithful to linguistic negation yet still express the desired level of factuality for clinical use.

4 Discussion and Preliminary Proposal

With any linguistic annotation task, striking a balance between linguistic precision and annotation feasibility is an inevitable and essential question. For the annotation of negation in Spanish, several components of the proposals discussed above may be combined into a set of complex guidelines that is both linguistically accurate and domain neutral. Here I summarize the main components I find to be worth annotating.

Most basically, the semantics of negation is represented (and ought to be annotated) through (i) the identification of the negation cue (the lexical element expressing negation); (ii) its scope (the text section that

is negated); (iii) its focus (that part of the scope that is prominently or explicitly negated); and, if present, (iv) its reinforcement (an auxiliary negation or NPI) (Altuna, Minard, and Speranza, 2017). This may be understood in an example such as the following:

(7) Juan **no** *come* [*carne*] **sino** *verduras*.

The negation cue (*no*) is represented in bold; the scope (*come carne sino verduras*) is in italics; the focus (*carner*) is in brackets; and the NPI (*sino*) in bold and italics. Negation markers that do not carry negative polarity semantic information (*nada más* “as soon as”) can be marked as such (for example, as <noneg> instead of <neg> (Jiménez-Zafra et al., 2018b).

Following (Morante, Schrauwen, and Daelemans, 2011), negation cues could be limited to just adverbs (*no, nadie, ninguno, nunca/jamás*). However, it seems that annotating morphological cues (prefixes such as *a-, in/im-, de(s)-, anti-*) as well as negative polarity verbs (*retirar, deaparecer, suspenderse*, etc.) is worth while for application to clinical domains.

This could be accomplished with both the annotation of the cues themselves and a linking to some sort of lexical definition or modal effect of the cue, as some combination of (Marimon, Vivaldi, and Bel, 2017) and (Jiménez-Zafra et al., 2018b) could produce. Figure 2 of (Jiménez-Zafra et al., 2018b) seems adequately suited to capturing the layers of complexity of negation. This, in combination with the distinction a NegPred (for (1), *comer* ‘to eat’), Negmarker (for (1), *no* ‘does not’), and NegPolItem (for (1), *sino* ‘but’) from (Marimon, Vivaldi, and Bel, 2017) could provide substantial coverage. It seems that an additional feature such as [+/-realis] may be helpful to distinguish levels of factuality of events in question, as well.

As a closing point, the Brat annotation tool (Marimon, Vivaldi, and Bel, 2017) seems suitable for any comprehensive annotation task involving negation. The multi-colored, layered format is accessible online, facilitating collaborative annotation efforts and the potential implementation of pilot annotation tasks to gauge inter-annotator agreement (IAA) as guidelines are developed.

5 Conclusion

This paper has presented an analysis and evaluation of existing guidelines for the annotation of negation in several domain-specific Spanish corpora. A preliminary proposal is given for how to combine linguistically accurate and precise annotation with more practical concerns regarding domain of application and ease of annotation. Future work points in particular towards refining the subtle meaning effects negation can have on words and phrase meaning, as well as its interaction with modality for ultimate interpretation of event factuality.

Acknowledgments

Many thanks to Elena Herburger, one of my doctoral advisors and an expert on negation. Thanks as well to Claire Bonial and the team at Army Research Lab (ARL) for supporting the work behind this paper.

References

- Altuna, B., A.-L. Minard, and M. Speranza. 2017. The Scope and Focus of Negation: A Complete Annotation Framework for Italian. In *Proceedings of the Workshop Computational Semantics Beyond Events and Roles*, pages 34–42.
- Bosque Muñoz, I. and J. Gutiérrez-Rexach. 2009. *Fundamentos de sintaxis formal*. Ediciones Akal.
- Cruz, N., R. Morante, M. J. M. López, J. M. Vázquez, and C. L. P. Calderón. 2017. Annotating negation in Spanish clinical texts. In *Proceedings of the Workshop Computational Semantics Beyond Events and Roles*, pages 53–58.
- Española-RAE, R. A. 2010. Nueva gramática de la lengua española. Manual. *Madrid: Espasa*.
- Herburger, E. 2018. What it means to be an NPI. *Linguistic Variation and Language Architecture Workshop*, June.
- Hoeksema, J. 2000. Negative polarity items: Triggering, scope and c-command. *Negation and polarity*, pages 115–146.
- Horn, L. 1989. A natural history of negation.
- Jiménez-Zafra, S. M., N. P. Cruz-Díaz, R. Morante, and M. T. Martín-Valdivia. 2018a. Tarea 1 del Taller NEGES 2018:

- Guías de Anotación. In *Proceedings of NEGES 2018: Workshop on Negation in Spanish*, volume 2174, pages 15–21.
- Jiménez-Zafra, S. M., M. Taulé, M. T. Martín-Valdivia, L. A. Ureña-López, and M. A. Martí. 2018b. Sfu Review SP-NEG: a Spanish corpus annotated with negation for sentiment analysis. A typology of negation patterns. *Language Resources and Evaluation*, 52(2):533–569.
- Marcus, M. P., M. A. Marcinkiewicz, and B. Santorini. 1993. Building a large annotated corpus of English: The Penn Treebank. *Computational linguistics*, 19(2):313–330.
- Marimon, M., J. Vivaldi, and N. Bel. 2017. Annotation of negation in the iula Spanish clinical record corpus. In *Proceedings of the Workshop Computational Semantics Beyond Events and Roles*, pages 43–52.
- Moeschler, J. 2010. Negation, scope and the descriptive/metalinguistic distinction. *Generative Grammar in Geneva*, 6:29–48.
- Morante, R., S. Schrauwen, and W. Daelemans. 2011. Annotation of negation cues and their scope: Guidelines v1. *Computational linguistics and psycholinguistics technical report series, CTRS-003*.
- Moreno, A., S. López, F. Sánchez, and R. Grishman. 2003. Developing a syntactic annotation scheme and tools for a Spanish treebank. In *Treebanks*. Springer, pages 149–163.
- Oronoz, M., K. Gojenola, A. Pérez, A. D. de Ilarraza, and A. Casillas. 2015. On the creation of a clinical gold standard corpus in spanish: Mining adverse drug reactions. *Journal of biomedical informatics*, 56:318–332.
- Sandoval, A. M. and M. G. Salazar. 2013. La anotación de la negación en un corpus escrito etiquetado sintácticamente annotation of negation in a written treebank. *Revista Iberoamericana de Linguística*, 8.
- Styler IV, W. F., S. Bethard, S. Finan, M. Palmer, S. Pradhan, P. C. de Groen, B. Erickson, T. Miller, C. Lin, G. Savova, et al. 2014. Temporal annotation in the clinical domain. *Transactions of the Association for Computational Linguistics*, 2:143.
- Taboada, M., C. Anthony, and K. D. Voll. 2006. Methods for Creating Semantic Orientation Dictionaries. In *LREC*, pages 427–432.