

SURVEY PAPER

Recent advances in processing negation

Roser Morante^{1*} and Eduardo Blanco²

¹CLTL Lab, Faculty of Humanities, VU Amsterdam, Amsterdam, The Netherlands and ²Department of Computer Science and Engineering, University of North Texas, Denton, TX, USA

*Corresponding author. E-mail: r.morantevallejo@vu.nl

(Received 7 October 2020; first published online 17 December 2020)

Abstract

Negation is a complex linguistic phenomenon present in all human languages. It can be seen as an operator that transforms an expression into another expression whose meaning is in some way opposed to the original expression. In this article, we survey previous work on negation with an emphasis on computational approaches. We start defining negation and two important concepts: scope and focus of negation. Then, we survey work in natural language processing that considers negation primarily as a means to improve the results in some task. We also provide information about corpora containing negation annotations in English and other languages, which usually include a combination of annotations of negation cues, scopes, foci, and negated events. We continue the survey with a description of automated approaches to process negation, ranging from early rule-based systems to systems built with traditional machine learning and neural networks. Finally, we conclude with some reflections on current progress and future directions.

Keywords: Processing negation; Focus; Scope; Semantics

1. Introduction

Negation is a complex linguistic phenomenon present in all human languages, allowing for the uniquely human capacities of denial, contradiction, misrepresentation, lying, and irony (Horn 1989). Broadly speaking, negation “relates an expression *e* to another expression with a meaning that is in some way opposed to the meaning of *e*” (Horn and Wansing 2017). The key challenge to understanding negation is thus to figure out the meaning that is *in some way* opposed to *e*—a semantic and highly ambiguous undertaking that comes naturally to humans in everyday communication. As described by Lawler (2010), cognitively, negation “involves some comparison between a real situation lacking some particular element and an imaginal situation that does not lack it.”

From a linguistic perspective, it is generally understood that negation has scope and focus. Scope is “the part of the meaning that is negated” and focus “the part of the scope that is most prominently or explicitly negated” (Huddleston and Pullum 2002). Consider statement (1). The falsity of any of the propositions in (1a–1d) makes statement (1) true, thus the scope of the negation includes all the elements indicated in the right column in (1a–1d):

1. This camera doesn’t take good pictures in low light

1a. Something takes something in some condition.

[verb] *takes*

1b. This camera takes something in some condition.

[AGENT of taking, *This camera*]

- 1c. Something takes good pictures in some condition. [THEME of taking, *good pictures*]
 1d. Something takes something in low light. [MANNER of knowing, *in low light*]

Determining the focus consists in pinpointing which parts of the scope are *intended to be interpreted* as true and false given the original statement. Without further context, one may consider condition (1c) or condition (1d) to be interpreted as false, and conditions (1a, 1b, and 1d) or conditions (1a–1c) to be interpreted as true, respectively. The first option yields the affirmative interpretation *This camera takes bad pictures in low light*, and the second option yields *This camera takes good pictures when there is enough light*.

Propositions where a negation is present are generally understood to carry positive meaning, or in other words, to suggest an affirmative alternative. For example, *John didn't leave the house* implicates that *John stayed inside the house*. Hasson and Glucksberg (2006) show that comprehending negation involves considering the representation of affirmative alternatives. While not fully understood, there is evidence that negation involves reduced access to the affirmative mental representation (Djokic *et al.* 2019). Orenes, Beltrán, and Santamaría (2014) provide evidence that humans switch to the affirmative alternative in binary scenarios (e.g., from *not red* to *green* when processing *The figure could be red or green. The figure is not red*). In such multi-scenarios, however, humans keep the negated representation unless the affirmative interpretation is obvious from context (e.g., humans keep *not red* when processing *The figure is red, green, yellow or blue. The figure is not red.*).

Previous work in computational linguistics has pointed out that processing negation is beneficial for some applications. In particular, sentiment analysis has a long tradition of incorporating components to process negation (Wiegand *et al.* 2010), ranging from handcrafted rules to sophisticated neural network architectures. For example, *like* generally carries positive sentiment, but not when modified by a negation cue (e.g., *don't like*). Wilson, Wiebe, and Hoffmann (2005) introduce the idea of contextual polarity, and note that negation may intensify rather than change polarity (e.g., *not good* vs. *not only good but amazing*). Jia, Yu, and Meng (2009) present a set of heuristic rules to determine sentiment when negation is present, and Councill, McDonald, and Velikovich (2010) show that information about the scope of negation is beneficial to predict sentiment. More recently, Socher *et al.* (2013) present a recursive neural network for sentiment analysis that captures sentiment polarity changes when negation is present. Reitan *et al.* (2015) show that a state-of-the-art Twitter sentiment analysis benefits from incorporating information about the scope of negation.

Outside sentiment analysis, researchers have pointed out that negation poses unsolved challenges for, among others, machine translation and natural language inference. Fancellu and Webber (2015) present a manual error analysis translating negation from Chinese to English, and Bentivogli *et al.* (2016) point out that neural machine translation struggles as much as statistical machine translation when it comes to translating negation. The task of natural language inference—determining whether a *text* entails, contradicts, or is neutral with respect to a *hypothesis*—and in particular existing benchmarks (e.g., SNLI (Bowman *et al.* 2015), MNLI (Williams, Nangia, and Bowman 2018)), do not properly account for inference in the presence of negation. A couple of insights from the literature support this claim. First, Gururangan *et al.* (2018) point out that negation cues (nobody, never, no, nothing, etc.) are a strong indicator of contradictions in MNLI. Second, creating adversarial examples by adding straightforward negations make existing systems underperform. For example, Naik *et al.* (2018) concatenate the tautology “and false is not true” at the end of hypotheses and conclude that models are not robust to negation. Similarly, Wallace *et al.* (2019) introduce universal triggers, that is, sequences of tokens that fool models when concatenated to any input. They show that concatenating negation cues such as *nobody*, *never*, *nothing*, and *none* to hypotheses decreases accuracy to almost zero when the gold label is entailment or neutral.

Processing negation is crucial when extracting information from clinical and biomedical texts, since negation is very frequent and missing negated observations can lead to misleading conclusions that have impact in health care and biomedical research. For example, the medical relevance of the sentence “shows no symptoms of COVID-19” is quite different from the sentence “shows symptoms of COVID-19”. Negation detection systems, such as the NegEx algorithm (Chapman *et al.* 2001) and its new version (Mehrabani *et al.* 2015), have shown good performance in clinical text. However, as Wu *et al.* (2014) indicate, negation is not solved because “current solutions are easily optimizable but not necessarily generalizable. Negation detection is still a challenge when considered from a practical, multi-corpus perspective, that is, one in which an algorithm is deployed in many clinical institutions and on many sources of text.” The authors conclude that in order to improve the performance in negation detection, more manually annotated data are needed in corpora that reflect the distributions of negation in different domains and genres. Cohen *et al.* (2017) find out that there are differences in the distribution of negation between clinical reports and biomedical articles. For example, there is more affixal negation in the clinical texts, whereas affixal level is more frequent in journal articles.

This special issue is devoted to advances in processing negation from a computational perspective. Section 2 provides information about existing corpora where negation is annotated and shared tasks. Section 3 presents previous work on processing negation, mainly on scope and focus identification.

2. Annotated corpora and shared tasks

That the interest in processing negation has grown in the last years is reflected in the number of corpora that have been published. Fancellu *et al.* (2017b) group corpora that contain some representation of negation into two types. The first-type corpora represent negation in a logical form, using quantifiers, predicates, and relations. One of them is the Groningen Meaning Bank (Basile *et al.* 2012),^a a collection of semantically annotated English texts with formal meaning representations. It is composed of newswire texts from Voice of America, country descriptions from the CIA Factbook, a collection of texts from the open ANC, and Aesop’s Fables. Another corpus is the DeepBank corpus (Flickinger, Zhang, and Kordoni 2012),^b which contains rich syntactic and semantic annotations for the 25 Wall Street Journal sections in the Penn Treebank (Taylor, Marcus, and Santorini 2003). The second-type corpora use string-level annotations, where negation information (cues, scope, negated event, and focus) are annotated on spans of text (e.g., BioScope (Vincze *et al.* 2008), ConanDoyle-neg (Morante and Daelemans 2012)). Jiménez-Zafra *et al.* (2020) recently surveyed corpora annotated with negation focusing on corpora that contain *string-level* negation. They describe the main features of the corpora based on the following criteria: the language, year of publication, domain, the availability, size, types of negation taken into account (syntactic and/or lexical and/or morphological), negation elements annotated (cue and/or scope and/or negated event and/or focus), and the way in which each corpus was tokenized, the annotation guidelines, and annotation scheme used. Information is provided about existing corpora for the following languages: English, Spanish, Swedish, Dutch, Japanese, Chinese, German, and Italian.

Corpora annotated with negation information usually include at least negation cues and either scope or focus. Some corpora also include event-level annotations indicating whether an event is negated. We note, however, that negation is ubiquitous in text. Jiménez-Zafra *et al.* (2020) report that the number of sentences with negation in the English corpora is between 9.37% and 32.16%, while in Spanish corpora it is between 10.67% and 34.22%, depending on the domain. Thus,

^aThe Groningen Meaning Bank is available at <http://gmb.let.rug.nl>.

^bDeepBank is available at <http://moin.delph-in.net/DeepBank>.

virtually all corpora include examples of negation. For example, parallel corpora for machine translation, product reviews annotated with sentiment, and corpora annotated with coreference chains include sentences with negation even though the negations themselves are not explicitly annotated.

Early work on annotating negation targeted the biomedical domain, for example, explicit negation of biological relationships (Pyysalo *et al.* 2007), negated events (Kim, Ohta, and Tsujii 2008), and negation cues and their scope (Vincze *et al.* 2008). Negation cues and scope have also been annotated in product reviews (Councill *et al.* 2010; Konstantinova *et al.* 2012), short stories (Morante and Daelemans 2012), and dialogs within intelligent tutoring systems (Banjade and Rus 2016). Liu, Fancellu, and Webber (2018) present a parallel corpus of English and Chinese annotated with negation. For Spanish, Sandoval and Salazar (2013) present annotations of negation cues and scope in news articles. Cruz *et al.* (2017) present similar annotations in the medical domain and include syntactic, morphological, and lexical negation. The latest and largest corpus (9455 sentences) is the SFU corpus by Jiménez-Zafra *et al.* (2018b), who annotate product reviews for a variety of product categories with, among others, negation cues and scopes. Negation markers and either events or scopes have also been annotated in Swedish (Dalianis and Velupillai 2010), Dutch (Afzal *et al.* 2014), Chinese (Zou, Zhou, and Zhu 2016), German (Cotik *et al.* 2016a), and Italian (Altuna, Minard, and Speranza 2017).

Regarding focus of negation, existing corpora work primarily with English texts. PB-FOC (Blanco and Moldovan 2011; Morante and Blanco 2012) was the first corpus to annotate the focus of negation and targets the negations marked with *argm-neg* semantic role in PropBank (Palmer, Gildea, and Kingsbury 2005). All negations annotated in PB-FOC are verbal negations, that is, the negation cues (e.g., *not*, *never*) are grammatically associated with a verb. The authors choose as focus the semantic role of the verbal negation that is most prominently negated. This decision simplifies the annotation task since usually a few semantic roles span many tokens. Anand and Martell (2012) reannotate PB-FOC and argue that (a) PB-FOC overannotates foci, and (b) alternatives arising from scalar implicatures and neg-raising predicates should be separated from those arising from focus detection. Other efforts working with English texts also refine the annotations in PB-FOC. Blanco and Moldovan (2012) define fine-grained focus as the tokens that are intended to be negated, in contrast to the coarse-grained foci annotated in PB-FOC, which are always all tokens belonging to a semantic role. They reannotate the foci in PB-FOC whose syntactic constituents are a PP, ADVP, or SBAR and discover that 68% of them can be refined. Blanco and Sarabi (2016) present an alternative annotation schema. Instead of choosing as focus the semantic role that is most prominently negated, they consider all roles and rank the likelihood of the underlying positive interpretations with a score ranging from 0 to 5. Sarabi and Blanco (2016) move away from semantic roles and work with syntactic dependencies. By selecting subtrees in dependency trees, they target foci of verbal negations—and the underlying positive interpretations—in a continuum of granularity. Finally, Sarabi and Blanco (2017) follow a similar approach manipulating syntactic dependencies, but target negation cues that modify nouns or adjectives instead of verbs.

In addition to PB-FOC and the refinements aforementioned, there are two corpora with focus annotations. Banjade and Rus (2016) work with dialogs retrieved from interactions between students and an intelligent tutoring system. Their corpus, DT-neg, contains annotations of negation cues as well as their scope and focus. The authors note that the scope and focus are often in a previous dialog turn. Kolhatkar *et al.* (2019) work with a collection of news articles and user-submitted comments online. Their corpus, the SFU Opinion and Comments Corpus (SOCC), contains annotations of negation cues as well as their scope and focus, among other annotations. To our knowledge, these two corpora are the only ones that work with user-generated content in English. Finally, to our knowledge, there is only one corpus annotating the focus of negation

in languages other than English (Matsuyoshi, Otsuki, and Fukumoto 2014). This corpus includes news articles and user reviews written in Japanese and targets negation cues and their foci.

Several shared tasks have addressed negation processing in English: in the BioNLP09 Shared Task 3 (Kim *et al.* 2009) and in the the i2b2 NLP Challenge (Uzuner *et al.* 2011), systems had to detect negated events. The *SEM 2012 Shared Task (Morante and Blanco 2012) was devoted to processing the scope and focus of negation, and in the ShARe/CLEF eHealth Evaluation Lab 2014 Task 2 (Mowery *et al.* 2014) participants had to detect whether a disorder was negated. For Spanish, there have been two editions of the NEGES task (Jiménez-Zafra *et al.* 2018a; 2019), where systems had to detect negation cues and evaluate the impact of negation in sentiment analysis.

3. Processing negation

The first attempts to process negation in English were mostly rule-based and focused on finding negated terms in clinical texts domain (Chapman *et al.* 2001; Mutalik, Deshpande, and Nadkarni 2001; Goldin and Chapman 2003; Auerbuch *et al.* 2004; Elkin *et al.* 2005; Boytcheva *et al.* 2005; Goryachev *et al.* 2006; Sanchez-Graillet and Poesio 2007; Huang and Lowe 2007; Rokach, Romano, and Maimon 2008). Some rule-based systems were developed based on lists of negations and stop words (Mitchell *et al.* 2004; Mykowiecka, Marciniak, and Kupść 2009; Harkema *et al.* 2009; Uzuner, Zhang, and Sibanda 2009; Sohn, Wu, and Chute 2012). With the surge of opinion mining, negation was studied as a marker of polarity change (Das and Chen 2001; 2005b; Polanyi and Zaenen 2006; Taboada *et al.* 2011; Jiménez-Zafra *et al.* 2017). Only with the release of the BioScope corpus (Vincze *et al.* 2008) did the work on negation receive a boost. But even so, despite the existence of several publications that focus on negation, it is difficult to find a negation processor for languages other than English. For English, some systems are available for processing clinical documents: NegEx (Chapman *et al.* 2001), ConText (Harkema *et al.* 2009), Deepen (Mehrabi *et al.* 2015), NegBio Peng *et al.* (2018), and NegTool (Enger, Velldal, and Øvreliid 2017).

3.1 Scope of negation identification

Compared to focus identification, scope identification has received substantially more attention, due to the existence of several corpora annotated with scopes. Inspired by the annotation of scopes in the BioScope corpus, the task of detecting negation scopes was introduced as a machine learning sequence labeling task (Morante, Liekens, and Daelemans 2008; Morante and Daelemans 2009). Subsequently, three main types of approaches have been applied to processing negation:

- Rule-based systems, in an attempt to improve the NegEx algorithm, such as ConText (Harkema *et al.* 2009), DEEPEN (Mehrabi *et al.* 2015), and NegMiner (Elazhary 2017). Packard *et al.* (2014) present an approach that works over explicit and formal representations of propositional semantics, deriving the scope from the structure of logical-form meaning representations.
- Machine learning techniques (Agarwal and Yu 2010; Li *et al.* 2010; Cotik *et al.* 2016b). Read *et al.* (2012) proposed an SVM-based ranking of syntactic constituents to identify the scope. Velldal *et al.* (2012) explore a combination of deep and shallow approaches to the problem of resolving the scope of speculation and negation. Reitan *et al.* (2015) apply a negation scope detection system to Twitter data.
- Deep learning approaches (Qian *et al.* 2016; Ren, Fei, and Peng 2018; Lazib *et al.* 2018). Fancellu, Lopez, and Webber (2016) present neural networks for this task, and Fancellu *et al.* (2017a) present an error analysis showing that scope is much easier to identify when delimited by punctuation.

3.2 Focus of negation identification

Automated systems to predict the focus of negation as annotated in PB-FOC are built using traditional machine learning algorithms as well as sophisticated neural networks. Blanco and Moldovan (2011) build a decision tree with a feature set emphasizing semantic role information. Rosenberg and Bergler (2012) present a rule-based approach implementing several heuristics. Zou, Zhou, and Zhu (2014) introduce a graph model for focus identification that takes into account inter-sentential information in addition to intra-sentential information. Doing so allows them to consider context from a lexical and topical perspective. Sarabi and Blanco (2019) present a simpler approach to include contextual information: they train an SVM with features derived from an off-the-shelf rhetorical structure theory (Mann and Thompson 1987) discourse parser (Surdeanu, Hicks, and Valenzuela-Escárcega 2015). To date, the best results published are presented by Shen *et al.* (2019), who obtain 70.51 accuracy. They present a BiLSTM with a CRF layer and attention mechanisms defined at the word and topic level.

There are less efforts experimenting with the other corpora annotating focus of negation. Most use an SVM trained with lexical, syntactic, and semantic features (Blanco and Moldovan 2012; Blanco and Sarabi 2016; Sarabi and Blanco 2016; Sarabi and Blanco 2017). Matsuyoshi *et al.* (2014) present 16 heuristic rules to detect the focus in Japanese. Finally, van Son *et al.* (2018) present a replication study of the corpus by Blanco and Moldovan (2012). They present results modeling the scoring class as regression and classification tasks and suggest that three labels (true, false, and uncertain) may be a better choice than a score ranging from 0 to 5.

This material is based upon work supported partly by the National Science Foundation under Grant No. 1845757. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. Roser Morante was supported by the Netherlands Organization for Scientific Research (NWO) via the Spinoza prize awarded to Piek Vossen (SPI 30-673, 2014-2019).

References

- Afzal Z., Pons E., Kang N., Sturkenboom M.C., Schuemie M.J. and Kors J.A. (2014). ContextD: an algorithm to identify contextual properties of medical terms in a Dutch clinical corpus. *BMC Bioinformatics* 15(1), 373.
- Agarwal S. and Yu H. (2010). Biomedical negation scope detection with conditional random fields. *Journal of the American Medical Informatics Association* 17(6), 696–701.
- Altuna B., Minard A.-L. and Speranza M. (2017). The scope and focus of negation: a complete annotation framework for Italian. In *Proceedings of the Workshop Computational Semantics Beyond Events and Roles*, pp. 34–42.
- Anand P. and Martell C. (2012). Annotating the focus of negation in terms of questions under discussion. In *Proceedings of the Workshop on Extra-Propositional Aspects of Meaning in Computational Linguistics*, Jeju, Republic of Korea: Association for Computational Linguistics, pp. 65–69.
- Auerbuch M., Karson T., Ben-Ami B., Maimon O. and Rokach L. (2004). Context-sensitive medical information retrieval. *Studies in Health Technology and Informatics* 107(Pt 1), 282–286.
- Banjade R. and Rus V. (2016). DT-Neg: tutorial dialogues annotated for negation scope and focus in context. In Chair N.C.C., Choukri K., Declerck T., Goggi S., Grobelnik M., Maegaard B., Mariani J., Mazo H., Moreno A., Odijk J. and Piperidis S. (eds), *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France: European Language Resources Association (ELRA).
- Barnes J., Velldal E. and Øvrelid L. (2020). Improving sentiment analysis with multi-task learning of negation. *Natural Language Engineering*.
- Basile V., Bos J., Evang K. and Venhuizen N. (2012). Developing a large semantically annotated corpus. In *LREC 2012, Eighth International Conference on Language Resources and Evaluation*, pp. 3196–3200.
- Bentivogli L., Bisazza A., Cettolo M. and Federico M. (2016). Neural versus phrase-based machine translation quality: a case study. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Texas: Association for Computational Linguistics, pp. 257–267.
- Blanco E. and Moldovan D. (2011). Semantic representation of negation using focus detection. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, Portland, Oregon, USA: Association for Computational Linguistics, pp. 581–589.
- Blanco E. and Moldovan D. (2012). Fine-grained focus for pinpointing positive implicit meaning from negated statements. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Montréal, Canada: Association for Computational Linguistics, pp. 456–465.

- Blanco E. and Sarabi Z.** (2016). Automatic generation and scoring of positive interpretations from negated statements. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, San Diego, California: Association for Computational Linguistics, pp. 1431–1441.
- Bowman S.R., Angeli G., Potts C. and Manning C.D.** (2015). A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, Lisbon, Portugal: Association for Computational Linguistics, pp. 632–642.
- Boycheva S., Strupchanska A., Paskaleva E., Tcharaktchiev D. and Str D.G.** (2005). Some aspects of negation processing in electronic health records. In *Proc. of International Workshop Language and Speech Infrastructure for Information Access in the Balkan Countries*, pp. 1–8.
- Chapman W., Bridewell W., Hanbury P., Cooper G. and Buchanan B.** (2001). A simple algorithm for identifying negated findings and diseases in discharge summaries. *Journal of Biomedical Informatics* 34, 301–310.
- Cohen K., Goss F., Zweigenbaum P. and Hunter L.** (2017). Translational morphosyntax: distribution of negation in clinical records and biomedical journal articles. In *MEDINFO 2017: Precision Healthcare through Informatics, Studies in Health Technology and Informatics*, vol. 245. IOS Press, pp. 346–350.
- Cotik V., Roller R., Xu F., Uszkoreit H., Budde K. and Schmidt D.** (2016a). Negation detection in clinical reports written in German. In *Proceedings of the Fifth Workshop on Building and Evaluating Resources for Biomedical Text Mining (BioTxtM 2016)*, pp. 115–124.
- Cotik V., Stricker V., Vivaldi J. and Rodríguez Hontoria H.** (2016b). Syntactic methods for negation detection in radiology reports in Spanish. In *Proceedings of the 15th Workshop on Biomedical Natural Language Processing, BioNLP 2016: Berlin, Germany, 12 August, 2016*. Association for Computational Linguistics, pp. 156–165.
- Councill I.G., McDonald R. and Velikovich L.** (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*. Association for Computational Linguistics, pp. 51–59.
- Cruz N., Morante R., Maña López M.J., Mata Vázquez J. and Parra Calderón C.L.** (2017). Annotating negation in Spanish clinical texts. In *Proceedings of the Workshop Computational Semantics Beyond Events and Roles*, Valencia, Spain: Association for Computational Linguistics, pp. 53–58.
- Dalianis H. and Velupillai S.** (2010). How certain are clinical assessments? Annotating Swedish clinical text for (un)certainities, speculations and negations. In *Proceedings of the Seventh Conference on International Language Resources and Evaluation (LREC'10)*. European Languages Resources Association (ELRA).
- Dalloux C., Claveau V., Grabar N., Oliveira L., Moro C., Gumiel Y. and Carvalho D.** (2020). Supervised learning for the detection of negation and of its scope in french and brazilian portuguese biomedical corpora. *Natural Language Engineering*.
- Das S. and Chen M.** (2001). Yahoo! for Amazon: Extracting market sentiment from stock message boards. In *Proceedings of the Asia Pacific Finance Association Annual Conference (APFA)*, Bangkok, Thailand, vol. 35, pp. 1–16.
- Djokic V., Maillard J., Bulat L. and Shutova E.** (2019). Modeling affirmative and negated action processing in the brain with lexical and compositional semantic models. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy: Association for Computational Linguistics, pp. 5155–5165.
- Elazhary H.** (2017). NegMiner: An automated tool for mining negations from electronic narrative medical documents. *International Journal of Intelligent Systems and Applications* 9(4), 14.
- Elkin P.L., Brown S.H., Bauer B.A., Husser C.S., Carruth W., Bergstrom L.R. and Wahner-Roedler D.L.** (2005). A controlled trial of automated classification of negation from clinical notes. *BMC Medical Informatics and Decision Making* 5(1), 13.
- Enger M., Veldal E. and Øvrelid L.** (2017). An open-source tool for negation detection: a maximum-margin approach. In *Proceedings of the Workshop Computational Semantics Beyond Events and Roles*, pp. 64–69.
- Fancellu F., Lopez A. and Webber B.** (2016). *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Berlin, Germany: Association for Computational Linguistics, pp. 495–504.
- Fancellu F., Lopez A., Webber B. and He H.** (2017a). Detecting negation scope is easy, except when it isn't. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, Valencia, Spain: Association for Computational Linguistics, pp. 58–63.
- Fancellu F., Reddy S., Lopez A. and Webber B.** (2017b). Universal dependencies to logical form with negation scope. In *Proceedings of the Workshop Computational Semantics Beyond Events and Roles*, pp. 22–32.
- Fancellu F. and Webber B.** (2015). Translating negation: a manual error analysis. In *Proceedings of the Second Workshop on Extra-Propositional Aspects of Meaning in Computational Semantics (ExProM 2015)*, Denver, Colorado: Association for Computational Linguistics, pp. 2–11.
- Flickinger D., Zhang Y. and Kordoni V.** (2012). Deepbank: A dynamically annotated treebank of the wall street journal. In *Proceedings of the 11th International Workshop on Treebanks and Linguistic Theories*, pp. 85–96.
- Goldin L. and Chapman W.** (2003). Learning to detect negation with 'Not' in medical texts. In *Proceedings of ACM-SIGIR 2003*, pp. 1–7.
- Goryachev S., Sordo M., Zeng Q.T. and Ngo L.** (2006). Implementation and evaluation of four different methods of negation detection. Technical report.

- Gururangan S., Swayamdipta S., Levy O., Schwartz R., Bowman S. and Smith N.A. (2018). Annotation artifacts in natural language inference data. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, New Orleans, Louisiana: Association for Computational Linguistics, pp. 107–112.
- Harkema H., Dowling J.N., Thornblade T. and Chapman W.W. (2009). Context: an algorithm for determining negation, experiencer, and temporal status from clinical reports. *Journal of Biomedical Informatics* 42(5), 839–851.
- Hasson U. and Glucksberg S. (2006). Does understanding negation entail affirmation?: an examination of negated metaphors. *Journal of Pragmatics* 38(7), 1015–1032.
- Horn L. (1989). *A Natural History of Negation*. Chicago: Chicago University Press.
- Horn L.R. and Wansing H. (2017). Negation. In Zalta E.N. (ed), *The Stanford Encyclopedia of Philosophy*, Spring 2017 Edn. Metaphysics Research Lab, Stanford University.
- Huang Y. and Lowe H.J. (2007). A novel hybrid approach to automated negation detection in clinical radiology reports. *Journal of the American Medical Informatics Association* 14(3), 304–311.
- Huddleston R.D. and Pullum G.K. (2002). *The Cambridge Grammar of the English Language*. Cambridge University 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*. ACM, pp. 1827–1830.
- Jiménez-Zafra S.M., Cruz Díaz N.P., Morante R. and Martín-Valdivia M.T. (2019). NEGES 2019 task: negation in Spanish. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2019)*. CEUR Workshop Proceedings, CEUR-WS, Bilbao, Spain.
- Jiménez-Zafra S.M., Díaz N.P.C., Morante R. and Martín-Valdivia M.T. (2018a). Tarea 2 del Taller NEGES 2018: Detección de Claves de Negación. *Proceedings of NEGES 2018: Workshop on Negation in Spanish*. CEUR, vol. 2174, pp. 35–41.
- Jiménez-Zafra S.M., Taulé, M., Martín-Valdivia M.T., Ureña-López L.A. and Martí M.A. (2018b). SFU ReviewSP-NEG: a Spanish corpus annotated with negation for sentiment analysis. A typology of negation patterns. *Language Resources and Evaluation* 52(2), 533–569.
- Jiménez-Zafra S.M., Valdivia M.T.M., Cámara E.M. and Ureña-López L.A. (2017). Studying the scope of negation for spanish sentiment analysis on twitter. *IEEE Transactions on Affective Computing* 10(1), 129–141.
- Jiménez-Zafra S.M., Morante R., Martín-Valdivia M.T. and Ureña-López L.A. (2020). Corpora annotated with negation: an overview. *Computational Linguistics* 46(1), 190–244.
- Kim J.-D., Ohta T., Pyysalo S., Kano Y. and Tsujii J. (2009). Overview of BioNLP'09 shared task on event extraction. In *Proceedings of the Workshop on Current Trends in Biomedical Natural Language Processing: Shared Task*. Association for Computational Linguistics, pp. 1–9.
- Kim J.-D., Ohta T. and Tsujii J. (2008). Corpus annotation for mining biomedical events from literature. *BMC Bioinformatics* 9(1), 10.
- Kolhatkar V., Wu H., Cavasso L., Francis E., Shukla K. and Taboada M. (2019). The SFU opinion and comments corpus: a corpus for the analysis of online news comments. *Corpus Pragmatics* 4(2), 155–190.
- Konstantinova N., de Sousa S.C., Cruz N.P., Maña M.J., Taboada M. and Mitkov R. (2012). A review corpus annotated for negation, speculation and their scope. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC-2012)*. European Language Resources Association (ELRA), 3190–3195.
- Lawler J. (2010). Negation and negative polarity. In Hogan P.C. (ed), *Cambridge Encyclopedia of the Language Sciences*. Cambridge, UK: CUP, pp. 554–555.
- Lazib L., Qin B., Zhao Y., Zhang W. and Liu T. (2018). A syntactic path-based hybrid neural network for negation scope detection. *Frontiers of Computer Science*, 14, 84–94.
- Li J., Zhou G., Wang H. and Zhu Q. (2010). Learning the scope of negation via shallow semantic parsing. In *Proceedings of the 23rd International Conference on Computational Linguistics*. Association for Computational Linguistics, pp. 671–679.
- Liu Q., Fancellu F. and Webber B. (2018). NegPar: a parallel corpus annotated for negation. In Chair N.C.C., Choukri K., Cieri C., Declerck T., Goggi S., Hasida K., Isahara H., Maegaard B., Mariani J., Mazo H., Moreno A., Odijk J., Piperidis S. and Tokunaga T. (eds), *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan: European Language Resources Association (ELRA), 3464–3472.
- Mann W.C. and Thompson S.A. (1987). *Rhetorical Structure Theory: A Theory of Text Organization*. Los Angeles: University of Southern California, Information Sciences Institute.
- Matsuyoshi S., Otsuki R. and Fukumoto F. (2014). Annotating the focus of negation in Japanese text. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, Reykjavik, Iceland: European Language Resources Association (ELRA), pp. 1743–1750.
- Mehrabi S., Krishnan A., Sohn S., Roch A.M., Schmidt H., Kesterson J., Beesley C., Dexter P., Schmidt C.M., Liu H. and Palakal M. (2015). DEEPEN: a negation detection system for clinical text incorporating dependency relation into negex. *Journal of Biomedical Informatics* 54, 213–219.
- Mitchell K.J., Becim M.J., Berman J.J., Chapman W.W., Gilbertson J.R., Gupta D., Harrison J., Legowski E. and Crowley R.S. (2004). Implementation and evaluation of a negation tagger in a pipeline-based system for information extraction from pathology reports. *Studies in Health Technology and Informatics, Volume 107: MEDINFO 2004*, 663–667.
- Morante R. and Blanco E. (2012). *SEM 2012 shared task: resolving the scope and focus of negation. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics (*SEM 2012)*, Montréal, Canada, pp. 265–274.

- Morante R. and Daelemans W.** (2009). A metalearning approach to processing the scope of negation. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL-2009)*, Boulder, Colorado: Association for Computational Linguistics, pp. 21–29.
- Morante R. and Daelemans W.** (2012). ConanDoyle-neg: annotation of negation cues and their scope in Conan Doyle stories. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC-2012)*. European Language Resources Association (ELRA), 563–1568.
- Morante R., Liekens A. and Daelemans W.** (2008). Learning the scope of negation in biomedical texts. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, pp. 715–724.
- Mowery D.L., Velupillai S., South B.R., Christensen L., Martinez D., Kelly L., Goeuriot L., Elhadad N., Pradhan S., Savova G. and Chapman W.W.** (2014). Task 2: ShARe/CLEF eHealth Evaluation Lab 2014. CLEF 2014 Working Notes, 1180. pp. 31–42.
- Mutalik P.G., Deshpande A. and Nadkarni P.M.** (2001). Use of general-purpose negation detection to augment concept indexing of medical documents: a quantitative study using the UMLS. *Journal of the American Medical Informatics Association* 8(6), 598–609.
- Mykowiecka A., Marciniak M. and Kupś A.** (2009). Rule-based information extraction from patients clinical data. *Journal of Biomedical Informatics* 42(5), 923–936.
- Naik A., Ravichander A., Sadeh N., Rose C. and Neubig G.** (2018). Stress test evaluation for natural language inference. In *Proceedings of the 27th International Conference on Computational Linguistics*, Santa Fe, New Mexico, USA: Association for Computational Linguistics, pp. 2340–2353.
- Orenes I., Beltrán D. and Santamaría C.** (2014). How negation is understood: evidence from the visual world paradigm. *Journal of Memory and Language*, 74, 36–45.
- Packard W., Bender E.M., Read J., Oepen S. and Drīdan R.** (2014). Simple negation scope resolution through deep parsing: a semantic solution to a semantic problem. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Baltimore, Maryland: Association for Computational Linguistics, pp. 69–78.
- Palmer M., Gildea D. and Kingsbury P.** (2005). The Proposition Bank: an annotated corpus of semantic roles. *Computational Linguistics* 31(1), 71–106.
- Peng Y., Wang X., Lu L., Bagheri M., Summers R. and Lu Z.** (2018). NegBio: a high-performance tool for negation and uncertainty detection in radiology reports. *AMIA Joint Summits on Translational Science Proceedings 2017*, 188–196.
- Polanyi L. and Zaenen A.** (2006). Contextual valence shifters. In *Computing Attitude and Affect in Text: Theory and Applications*. Springer, pp. 1–10.
- Pyysalo S., Ginter F., Heimonen J., Björne J., Boberg J., Järvinen J. and Salakoski T.** (2007). BioInfer: a corpus for information extraction in the biomedical domain. *BMC Bioinformatics* 8(1), 50.
- Qian Z., Li P., Zhu Q., Zhou G., Luo Z. and Luo W.** (2016). Speculation and negation scope detection via convolutional neural networks. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 815–825.
- Read J., Velldal E., Øvrelid L. and Oepen S.** (2012). UiO1: constituent-based discriminative ranking for negation resolution. In **SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the Main Conference and the Shared Task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, Montréal, Canada: Association for Computational Linguistics, pp. 310–318.
- Reitan J., Faret J., Gambäck B. and Bungum L.** (2015). Negation scope detection for twitter sentiment analysis. In *Proceedings of the 6th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis (WASSA 2015)*, Lisboa, Portugal: Association for Computational Linguistics, pp. 99–108.
- Ren Y., Fei H. and Peng Q.** (2018). Detecting the scope of negation and speculation in biomedical texts by using recursive neural network. In *2018 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE, pp. 739–742.
- Rokach L., Romano R. and Maimon O.** (2008). Negation recognition in medical narrative reports. *Information Retrieval* 11(6), 499–538.
- Rosenberg S. and Bergler S.** (2012). UConcordia: CLaC negation focus detection at *Sem 2012. In **SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the Main Conference and the Shared Task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, Montréal, Canada: Association for Computational Linguistics, pp. 294–300.
- Sanchez-Graillet O. and Poesio M.** (2007). Negation of protein–protein interactions: analysis and extraction. *Bioinformatics* 23(13), i424–i432.
- Sandoval A.M. and Salazar M.G.** (2013). La anotación de la negación en un corpus escrito etiquetado sintácticamente. Annotation of negation in a written treebank. *Revista Iberoamericana de Lingüística: RIL* 8, 45–60.
- Sarabi Z. and Blanco E.** (2016). Understanding negation in positive terms using syntactic dependencies. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Texas: Association for Computational Linguistics, pp. 1108–1118.
- Sarabi Z. and Blanco E.** (2017). If no media were allowed inside the venue, was anybody allowed? In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, Valencia, Spain: Association for Computational Linguistics, pp. 860–869.
- Sarabi Z. and Blanco E.** (2019). Identifying the focus of negation using discourse structure. In *Proceedings of the Thirty-Second International Florida Artificial Intelligence Research Society Conference (FLAIRS)*, pp. 448–452.

- Schulder M., Wiegand M. and Ruppenhofer J. (2020). Automatic generation of lexica for sentiment polarity shifters. *Natural Language Engineering*.
- Shen L., Zou B., Hong Y., Zhou G., Zhu Q. and Aw A. (2019). Negative focus detection via contextual attention mechanism. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Hong Kong, China: Association for Computational Linguistics, pp. 2251–2261.
- Socher R., Perelygin A., Wu J., Chuang J., Manning C.D., Ng A. and Potts C. (2013). Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Seattle, Washington, USA: Association for Computational Linguistics, pp. 1631–1642.
- Sohn S., Wu S. and Chute C.G. (2012). Dependency parser-based negation detection in clinical narratives. *AMIA Summits on Translational Science Proceedings*, 2012, 1–8.
- Surdeanu M., Hicks T. and Valenzuela-Escárcega M.A. (2015). Two practical Rhetorical Structure Theory parsers. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, Denver, Colorado: Association for Computational Linguistics, pp. 1–5.
- Sykes D., Grivas A., Grover C., Tobin R., Sudlow C., Whiteley, W. McIntosh A., Whalley H. and Alex B. (2020). Comparison of rule-based and neural network models for negation detection in radiology reports. *Natural Language Engineering*.
- Taboada M., Brooke J., Tofiloski M., Voll K. and Stede M. (2011). Lexicon-based methods for sentiment analysis. *Computational Linguistics* 37(2), 267–307.
- Taulé M., Nofre M., González M. and Martí A. (2020). Focus of negation: its identification in spanish. *Natural Language Engineering*.
- Taylor A., Marcus M. and Santorini B. (2003). The Penn TreeBank: an overview. In *Treebanks*. Springer, pp. 5–22.
- Uzuner Ö., South B.R., Shen S. and DuVall S.L. (2011). 2010 i2b2/VA challenge on concepts, assertions, and relations in clinical text. *Journal of the American Medical Informatics Association* 18(5), 552–556.
- Uzuner Ö., Zhang X. and Sibanda T. (2009). Machine learning and rule-based approaches to assertion classification. *Journal of the American Medical Informatics Association* 16(1), 109–115.
- van Son C., Morante R., Aroyo L. and Vossen P. (2018). Scoring and classifying implicit positive interpretations: a challenge of class imbalance. In *Proceedings of the 27th International Conference on Computational Linguistics*, Santa Fe, New Mexico, USA: Association for Computational Linguistics, pp. 2253–2264.
- Vellidal E., Øvrelid L., Read J. and Oepen S. (2012). Speculation and negation: rules, rankers, and the role of syntax. *Computational Linguistics* 38, 369–410.
- Vincze V., Szarvas G., Farkas R., Móra G. and Csirik J. (2008). The BioScope corpus: biomedical texts annotated for uncertainty, negation and their scopes. *BMC Bioinformatics* 9(11), S9.
- Wallace E., Feng S., Kandpal N., Gardner M. and Singh S. (2019). Universal adversarial triggers for attacking and analyzing NLP. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Hong Kong, China: Association for Computational Linguistics, pp. 2153–2162.
- Wiegand M., Balahur A., Roth B., Klakow D. and Montoyo A. (2010). 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: University of Antwerp, pp. 60–68.
- Williams A., Nangia N. and Bowman S. (2018). A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. Association for Computational Linguistics, pp. 1112–1122.
- Wilson T., Wiebe J. and Hoffmann P. (2005). Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, Vancouver, British Columbia, Canada: Association for Computational Linguistics, pp. 347–354.
- Wu S., Miller T., Masanz J., Coarr M., Halgrim S. and Clark C. (2014). Negation's not solved: generalizability versus optimizability in clinical natural language processing. *PLoS ONE* 9, 1–11.
- Zou B., Zhou G. and Zhu Q. (2014). Negation focus identification with contextual discourse information. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Baltimore, Maryland: Association for Computational Linguistics, pp. 522–530.
- Zou B., Zhou G. and Zhu Q. (2016). Research on Chinese negation and speculation: corpus annotation and identification. *Frontiers of Computer Science* 10(6), 1039–1051.