


A Review on Negation Role in Twitter Sentiment Analysis

Itisha Gupta, Banasthali Vidyapith, Vanasthali, India

Nisheeth Joshi, Banasthali Vidyapith, Vanasthali, India

 <https://orcid.org/0000-0002-9256-3825>

ABSTRACT

Negation is an important linguistic phenomenon that needs to be considered for identifying correct sentiments from the opinionated data available in digital form. It has the power to alter the polarity or strength of the polarity of affected words. In this paper, the authors present a survey on the negation role that has been done until now in sentiment analysis, specifically Twitter sentiment analysis. The authors discuss the various approaches of modelling negation in Twitter sentiment analysis. In particular, their focus is on negation scope detection and negation handling methods. This article also presents some of the challenges and limits of negation accounting in the field of Twitter sentiment analysis.

KEYWORDS

Healthcare, Lexical Ambiguity, Negation Handling, Negation Modelling, Negation Scope, Reverse Polarity, Shift Polarity, Syntactic Negation, Twitter Sentiment Analysis

1. INTRODUCTION

Twitter sentiment analysis (TSA) involves sorting out and grouping of articulated texture opinions automatically on various topics. Such an analysis aids in determining the embedded sentiment and the effect of the expressed opinions intended for the readers (Cambria et al., 2017). An important subtasks of TSA then is the identification of opinionated words (sentiment carrying words). SentiWordNet (SWN) 3.0 and other lexicons avail for this purpose (Baccianella et al., 2010), but a major challenge often encountered in TSA is the presence of negation, resulting in lexical ambiguity.

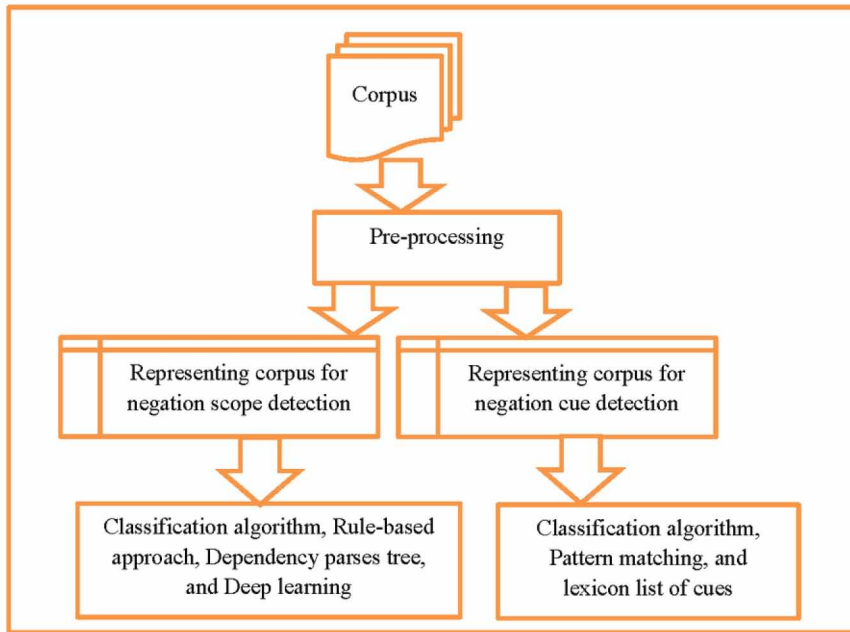
In health care, negation is an essential linguistic element that can change the semantic orientation of a piece of text. Originating from the biomedical domain (Chapman et al., 2001; Mutalik et al., 2001; Morante et al., 2008), negation plays a vital role in either written or spoken natural languages (NLs). Given the frequent use of negation in biomedicine (e.g., Taylor & Harabagiu, 2018), negation is of significance in health care. As biomedical text represents a vast depository of information feeding the healthcare system, the negation data often seen in clinical reports (Chapman et al., 2001) must be further mined to provide valuable hidden knowledge via sentiment analysis, for example, predicting the risk of dispensing certain drugs. Related to biomedical text mining, the work of Mukherjee et al. (2017), who use NegAIT parser for detecting morphological negation, double negation, and sentential negation in medical text is prominent.

To date, there is a scarcity of research focussing on the medical text such as CONLL2010 (Farkas et al., 2010) and *SEM 2012, centering on determining the negation scope (Morante & Blanco, 2012). One of the popularly used corpus for negation scope resolution has been the BioScope corpora,

DOI: 10.4018/IJHISI.20211001.0a14

This article published as an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>) which permits unrestricted use, distribution, and production in any medium, provided the author of the original work and original publication source are properly credited.

Figure 1. Different approaches used in the literature for the negation cue and scope detection in the biomedical domain



annotated with cues and scopes (Vincze et al., 2008). According to Vincze et al. (2008), one out of eight sentences in existing biomedical text has negation. A growing number of approaches have been explored over the years to address the negation modelling problem, including machine learning (ML) classifier, rule-based, and more recently, deep learning methods.

Figure 1 displays various approaches applicable in identifying negation cue and detecting scope, specifically in the biomedical domain.

In early research, several rule-based negation detection techniques have been developed and used such as NegEx (Chapman et al., 2001) and NegFinder (Mutalik et al., 2001). These techniques can identify negated findings in the clinical context and discharge summaries, but such approaches are simply no longer adequate for complex structures. Consequently, new techniques and tools have emerged to detect negation in the health care context such as DEEPEN (Mehrabi et al., 2015), MEDLEE, and other ontology-based approaches. The rule-based approach is among the firsts to be used on the Bioscope corpus for negation detection and scope resolution. A notable example is the enduring work of Chapman et al. (2001) on NegEx, a simplistic tool that requires the user to provide a list of phrases and expressions that could be negated. This work has been extensively applied in the biomedical domain.

Table 1 details earlier works in the biomedical domain for the negation modelling including the corpus, results, and negation approaches (for cue detection and scope resolution).

In this paper, various studies on negation modelling specific to sentiment analysis, especially TSA, are reviewed. The primary goal is to determine the effectiveness of modelling negation in health care; specifically, in the biomedical domain and tweets on broader health care topics given that negation and negation modelling are frequent and relevant towards achieving such a goal. The focus thus will be on various negation scope resolution and negation handling approaches used in sentiment analysis, including key challenges faced during negation modelling in TSA. The rest of the paper is organised as follows. Section 2 presents the research motivation while Section 3 describes the various forms or types of negation. Section 4 highlights the various phases of modelling negation in sentiment analysis.

Table 1. Previous work on negation modelling in the biomedical domain

Authors	Year	Corpus	Negation detection approaches	Result
Chapman et al.	2001	Discharge summaries (custom dataset)	NegEx (regular expression parser for negation cue detection)	Precision: 84.49, Recall: 77.84, and F1 score: 81.03
Elkin et al.	2005	Clinical notes	Negation ontology (based on NegEx)	-
Huang & Lowe	2007	Clinical radiology report	Hybrid approach (regular expression matching with grammatical parsing)	Precision: 99.8, recall: 92.6, and F1 score: NA
Morante et al.	2008	BioScope Corpus (Abstracts)	Machine learning approach for cue and scope detection	For negation cue detection F1 score: 91.54 For scope resolution F1: 88.4
Rokach et al.	2008	Medical narrative reports (custom dataset)	Pattern learning approach	Precision 87.8, Recall 80.99, and F1 score 84.47
Morante & Daelemann	2009	Bioscope corpus (Abstracts)	Metalearning approach for scope detection and memory-based learning algorithm (IGTREE) for cue detection	For negation cue detection f1 score: 98.68 For scope resolution F1: 90.67
Morante & Daelemann	2009	Bioscope corpus (Full papers)	Machine learning classifier for cue detection and metalearning for scope resolution	F1 score: 97.81 for cue detection F1 score: 84.71 for scope resolution
Councill et al.	2010	Bioscope corpus (Full papers)	CRF for scope resolution	F1 score: 75.5
Agarwal & Yu	2010	Biological literature (Biomedical corpus) and clinical notes	Conditional Random Fields (CRFs) for scope and cue detection	Clinical notes: F1 score: 98% for negation cue detection and F1 score: 95% for scope detection Biological literature: F1 score: 97% for cue detection and 85% for scope resolution
Diaz et al.	2012	BioScope corpus (Clinical text)	Machine learning classifiers for detecting cue and scope	Best F1 score: 97.3 for cue detection (C4.5) F1 score: 94.9 for scope detection (SVM RBF)
Mehrabi et al.	2015	Bioscope corpus (Mayo clinical dataset and pancreatic cyst concepts (IU dataset))	DEEPEN (NegEx with dependency parsing)	Mayo dataset F1 score: 80.75 IU dataset F1 score: 96.45
Ou & Patrick	2015	Biomedical corpus	Machine learning, lexicon and syntax-based approach for negation detection	Best F1 score: 82.56 with machine learning approach
Attardi et al.	2015	BioScope Corpus (clinical notes)	Dependency parse tree for Negation scope detection and classifiers for negation cue detection (linear SVM, TanI NER, and deepNL)	F1 score: 90.73 for cue detection using TanI NER and F1 score: 78.87 for scope detection
Qian et al.	2016	BioScope corpus (Abstracts)	Convolution Neural Networks for scope resolution	F1 score: 89.91
Qian et al.	2016	BioScope corpus (Full papers)	Convolution Neural Networks for scope resolution	F1 score: 83.46
Fancellu et al.	2016	Bioscope corpus (Abstracts)	BiLSTM for negation scope resolution	F1 score: 91.35
Fancellu et al.	2016	Bioscope corpus (Full papers)	BiLSTM for negation scope resolution	F1 score: 77.85
Peng et al.	2017	Bioscope corpus (Full papers)	NegBio (universal dependency patterns for negation cue detection)	F1 score: 95.9 for negation cue detection
Taylor & Harabagiu	2018	Electroencephalography reports	BiLSTM for cue and scope detection	F1 score: 88.5 for scope detection
Khandelwal & Sawant	2019	Bioscope corpus (abstract and full papers), Sherlock dataset, and SFU review corpus	BERT (transfer learning model using deep learning framework)	Bioscope Abstracts cue detection F1 score: 95.65 and Full papers cue detection F1 score: 92.42 Bioscope Abstract scope resolution F1 score: 95.68 and Full papers F1 score: 91.24
Chen	2019	Clinical notes (Electronic Health Record, EHR)	BiLSTM deep learning framework for scope detection	-
Bhatia et al.	2019	Proprietary medical condition dataset	Dep learning framework (BiLSTM encoder and LSTM decoder)	-

Here, a detailed description of various approaches used for detecting negation scope and for handling the impact of negation will be discussed. As well, this section also emphasizes the need for effective negation scope resolution in health care. Section 5 then shifts focus to the various challenges faced during negation handling in TSA while Section 6 summarizes the key contributions and limitations of the present work. Finally, the last section concludes with insights into future research directions.

2. MOTIVATION

In the sentiment analysis field, negation has been well studied (e.g., Gupta & Joshi, 2019; Muhammad et al., 2016; Councill et al., 2010; Jia et al., 2009; Wilson et al., 2005). As negation might change the text polarity or strength of polarity (Weigand et al., 2010), it is critical in sentiment analysis. The automatic process of extracting opinions from the clinical text implies the need for linguistic analysis techniques. Also, efforts to extract relevant features for lexicon or ML approach in handling such subtle phenomenon as negation should receive careful attention. In other words, negation needs to be properly accounted for in TSA.

In TSA, it is necessary to correctly evaluate the semantic orientation of a piece of text, which is affected by contextual element negation. Negation impact can be limited to the next word or sequence of words after that. For each negation, one must detect the negation word and the scope of negation (words affected by negation), for instance, consider a tweet on iPhone 7:

- “I like using the new markup features in #ios 10 on my #iphone7.”

This tweet, which is conveyed by an opinionated word “like,” clearly shows the positive opinion of a user.

Now, consider another tweet, which is the negation of the previous tweet.

- “I **don’t** [like] using the new markup features in #ios 10 on my #iphone7.”

This tweet clearly expresses a negative opinion due to the presence of negation word (don’t).

Negation cue is bold and scope is placed inside brackets. Even though the word “like” conveys a positive opinion, it is affected by the negation “don’t”, which reverses the polarity of “like”. Hence, it is crucial to capture the impact of negation in TSA; otherwise, tweets would be misclassified. However, negation presence does not necessarily mean negation (Jia et al., 2009). Thus, it is not always easy to spot expressed negative and positive opinions in a piece of text, for instance, consider the tweet:

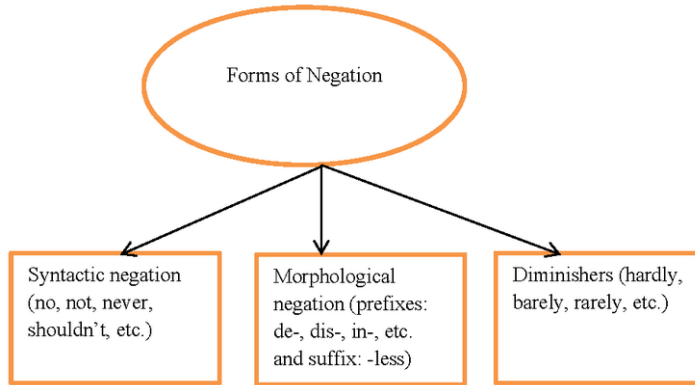
- “Isn’t that movie great?”

In this tweet negation cue “isn’t” doesn’t invert the polarity of opinionated word “great”.

Put simply, it is important to spot cases where negation is present but there is no sense of negation. Moreover, negation does not always reverse polarity or change the polarity of all opinionated words in a text. Thus, it is necessary to correctly determine the scope of negation (which words are affected by negation). Also, existing negation research has shown that negation can also be introduced via diminishers (e.g. hardly, rarely, etc.), modals (e.g. should), or even connectives (e.g. but), all of which can tone down the polarity expressed in the text as shown correspondingly in the three examples below:

- “This application hardly lags.”
- “This phone should have worked even under water.”
- “Fairly good acting but over all a disappointing movie.”

Figure 2. Different forms of negation



Such negations can be captured with the help of some linguistic rule-based methods.

However, implicit negations (as shown below) are very hard to identify due to missing explicit negation cues.

- “I don’t like this movie. Great actor, awful scenario.”

From the aforementioned examples, it is apparent that modelling negation is challenging and non-trivial in TSA. While the initial study on negation handling in computational linguistic has originated in the medical domain, with technological advances, negation has now been explored in multiple other domains, for example, product reviews (Councill et al., 2010), general tweets, and other health care-related tweets. Purportedly, Twitter has been seen as an impactful resource for health care studies (Burke et al., 2018; Doan et al., 2018; Reece et al., 2017). More recently, Gohil et al. (2018) review various tools and methods used in the healthcare domain for the TSA. Before presenting these computational methodologies developed to tackle the subtle negation phenomenon, it is important first to be familiar with various forms or types of negation.

3. FORMS OF NEGATION

Broadly categorized, three basic forms of negation emerged as depicted in Figure 2.

3.1 Syntactic Negation

Syntactic negation is the most common and well-known form of negation, which is also known as functional negation (Choi & Cardie, 2008). Identified by explicit negation cues such as “no, not, never, isn’t,” or “shouldn’t,” syntactic negation can affect the polarity of a word or phrases (a series or grouping of words).

3.2 Morphological Negation

Morphological negation is associated with one of fifteen (15) negative prefixes (“un-, dis-, ill-, di-, mis-, de-, non-, be-, counter-, a-, under-, mal-, pre-, ter-, ex-”) or with the suffix “-less”, which will modify the root word.

Also, with the negative “in-” as a prefix allomorphy, the different forms of assimilation include: (a) in- + r = irr-; (b) in- + labial consonant (b, m, p) = im-; and (c) in- + l = ill- (Bauer et al., 2013).

Hence, negation cue and opinionated word may be combined into a single word such as “dishonest” or “impolite”.

One way to handle the morphological negation is to lexicalize all such words, and, then, to obtain their polarities directly from such lexicons. With this approach, there is no need of any scope resolution.

3.3 Diminishers

With the ability to tone down the polarity of expressed opinion rather than reversing the polarity such as “hardly, barely,” or “rarely,” diminishers are also a form of negation. The words affected by the diminishers can be present anywhere in the text, for instance, “This mobile hardly lags”. Here, the word “hardly” tone down the polarity of the polar word “lags”.

Apparently, diminishers have a scope associated with their usage, which is either an adjective or a verb. Most of the earlier works such as in (Alistair & Diana, 2005) handled the diminishers by reducing the polarity of affected word by a certain amount. However, few researchers (e.g, Asghar et al., 2017; Taboada et al., 2011) attempt to study the impact of diminishers. Notwithstanding, whenever the same diminisher is used, it would be incorrect to infer that the polarity has been tone down by the same amount every time.

As the percentage value of diminishers used in existing works varies, it is a challenge for anyone when handling diminishers. Still, the percentage system (that is, a certain percentage in which each diminisher is associated with) appears to be the best way of qualifying the polarity of the affected word. Simply stated, one has to consider to what degree is the polarity of the affected word toned down based on the percentage value of that particular diminisher being used.

4. A SURVEY ON MODELLING NEGATION

Today, TAS researchers are often faced with the challenge of tackling the different natural language processing (NLP) tasks, for example, sarcasm detection, polarity detection, negation and many more (Cambria et al., 2017).

Of these, negation is a prominent aspect of TSA that needs to be modelled properly in order to get the correct classification. As elaborated in the following subsections, negation modelling is generally done in three phases:

- Identification of negation cues
- Scope resolution
- Handling the negation impact

4.1 Identifying Negation Cues

A preliminary step towards negation modelling is the identification of negation cues (such as “no, not, never,” and many more). Negation cues may be identified either by the lexicon approach (Councill et al., 2010; Qian et al., 2017), in which a list of cues is created, or by the ML approach such as in (Morante & Daelemans, 2009), whereby a classifier is employed to predict the negation cue.

In TAS, implementing a lexicon approach for cue detection is somewhat subtle due to misspelled negation cues. Thus, most of the earlier works in Twitter used the clusters created by Owoputi et al. (2013), and, then, considered the whole cluster for each matching negation cue in the lexicon list. Notwithstanding, Reitan et al. (2015) observe that using the word cluster for negation cue will lead to a high recall (less false negatives) with declining precision. Instead, the use of fine-grained clusters has been proposed, which may still increase recall without hurting the desired precision.

4.2 Negation Scope Resolution

Once a negation cue is detected, the next step is to determine which words are affected by the negation. As a negation cue may not affect all opinionated words in a text, efforts must be made to determine the impact and scope of the targeted negation.

Generally, existing works on negation scope resolution may be grouped as follows: one is linguistic rule-based methodology, second is ML, and the third is shallow semantic parsing. Among these, the traditional linguistic rule-based method is the most common; here, a varying number and/or type of words may be considered under the negation scope, for example, it could be (a) the rest of the sentence, (b) the first sentiment carrying word after a negation, or (c) a fixed static window of words. Such a methodology imitates the grammatical structure of the sentence.

Importantly, most state-of-the-art sentiment analysis systems (e.g., El Alaoui et al., 2018; Kolchyna et al., 2015; Kennedy & Inkpen 2005; Polanyi & Zaenen, 2004; Pang et al, 2002) have used the simple approach of resolving negation scope. Simply, this amounts to marking all the words between the negation cue detected and the first punctuation mark. Researchers in the SemEval tasks such as Kiritchenko et al. (2014) and Zhu et al. (2014) have used this naive approach of scope resolution, which has a major disadvantage as it leads to the modification of all words under the scope irrespective of whether it is polar or not. Moreover, few studies have compared or contrasted different linguistic rule-based methods.

Nonetheless, critics such as Farooq et al. (2017) and Gupta & Joshi (2019) have argued that the aforementioned methods might work well for simple text, but will fail in the case of compound and complex text (having more than one clause or conjunctions), owing to the presence of conjunction such as “but,” which may limit the negation scope to only a single clause. As such, the scope of negation may possibly end before a static window termination or before the occurrence of the first punctuation mark.

Apart from the simple linguistic rule-based approach, there exists few studies on the more complex approaches for scope resolution including, for example, shallow semantic parsing, ML classifiers, or reinforcement learning. The most prominent work is Jia et al. (2009) in which the concept of static v. dynamic delimiters is used together with some heuristic rules for scope resolution. First, the authors obtain the candidate scope with the help of parse tree generated by the parser; next, they prune the candidate scope via static v. dynamic delimiters; and finally, they apply heuristic rules to elucidate the exact scope. The limitation of their approach is that the scope is confined only to that clause in which the negation cue occurs. Notwithstanding, the scope may indeed be expanded to other clauses.

Use of the ML classifier, specifically, Conditional Random Field (CRF) classifier for negation scope resolution has been performed, but most of the TSA works in this line of research focus on either product reviews or biomedical corpus. For instance, Vincze et al. (2008) develop the biomedical corpora annotated with cues and scopes. In this BioScope corpora, the negation cues are always explicit and the scope is maximal. It is the first freely available negation annotated corpus that allows researchers to use the supervised classifiers for negation scope resolution. In Morante & Daelemans (2009), for example, the CRF meta-learning approach, combining the input from K-Nearest Neighbours (KNN), Support Vector Machine (SVM), and another CRF, is used for scope prediction. Here, the BioScope medical corpus is used to resolve the negation scope by participating in the CoNLL'10 task so as to achieve the best F1 score among all participants. Similarly, the work of Zhu et al. (2010) on the BioScope corpus uses a shallow semantic parsing approach for negation scope resolution.

On sentiment analysis to resolve negation scope, Moilanen & Pulman (2007) are among the firsts to develop a sentiment analysis system with a sophisticated negation scope detection approach based on the linguistic syntactic composition. More recently, Council et al. (2010), advocate the corpora for English product reviews from Goggle products and use a single CRF classifier for scope resolution. They firstly generated a parse tree via MaltParser, and, then, train the CRF on the features extracted from the parse tree. Unlike Morante & Daelemans (2009), these authors use a simple lexicon list for negation cue identification via a ML approach.

Konstantinova et al. (2012) discuss the development of the Simon Fraser University (SFU) Review corpus; unlike Councill et al. (2010), which only contains 2111 sentences, the SFU corpus contains 17000 annotated sentences. Given that these earlier approaches for scope resolution consider texts as a sequence of events (i.e., a word depends on the sequence of previously occurred words), it follows that the characterization of a negation scope may be inferred as a sequence-labelling task in which words may be labelled as (or not as) part of the negation scope (Councill et al., 2010).

In contrast to the generative probabilistic models such as CRF, Pröllochs et al. (2016) propose the use of a reinforcement learning approach for resolving the negation scope. They prove the effectiveness of their approach by comparing the rule-based methods versus the generative probabilistic models (CRF), and, then, obtain a balanced accuracy of up to 70.17%. Their results also achieve an improvement of 10.63% in sentiment analysis of financial news due to negation modelling.

Initially, when annotated corpora were unavailable in the Twitter domain, most of the TSA studies used the linguistic rule-based methods for scope resolution due to the absence of annotated corpora. Even so, annotated corpora were unavailable in the Twitter domain until 2014. Thus, teams of researchers such as NRC-Canada (Mohammad et al., 2013) use the state-of-the-art linguistic rule-based approach of scope resolution for TSA by marking all words between the negation cues up till the first punctuation mark to be under the scope.

More recently, Reitan et al. (2015) argue for the need of a Twitter corpus (annotated with cues and scopes) to train supervised learning classifiers on the negation annotated Twitter corpora. These researchers download 4000 tweets via Twitter search API, develop a web application for annotation purposes, and then, annotate the Twitter corpus with cues and scopes. Their analysis has found 539 tweets out of 4,000 containing negation (i.e. 13.5% tweets are negation tweets). With the annotated Twitter corpus, the authors could train the dependency-based binary CRF classifier to resolve the negation scope. Binary CRF classifier can predict if a token is in an affirmative or a negated context. Their CRF classifier entails a Twitter tailored version of CRF used by Councill et al. (2010), having the same feature set with an added feature (i.e., dependency distance of each token to the closest negation cue). As well, they adopt the lexicon list of explicit negation cues from Councill et al. (2010) and evaluate their Twitter trained binary CRF classifier on the SFU product review corpora (Konstantinova et al., 2012) and the BioScope corpora. Results have shown that the CRF perform reasonably well on the SFU corpora (i.e., negation annotated corpora for review domain), but not with the structured BioScope corpus.

Notwithstanding, by sharing the negation annotated Twitter corpora with the broader research community, the researchers are able to show their results to be superior to those of Councill et al. (2010), and of Morante & Daelemans (2009) on the BioScope in terms of its F1 score. Today, researchers have used deep learning approaches (i.e., recurrent neural network or NN) to resolve the negation scope with the ability of such methods to handle sequential data, for instance, the NN performance has been found to be equal, if not better than CRF, Hidden Markov Model (HMM), or any other sequence labelling models.

Finally, Gautam et al. (2018) explore various Long Short-Term Memory (LSTM) models for negation handling in tutorial dialogues. They use the LSTM for both, negation cue identification and scope resolution. They validate their models via real dialogues and their result has shown that the sequence-to-sequence tagger can handle the negation scope, cue, while staying focus. Indeed, the LSTM model has outperformed the CRF model (which needs manually crafted features). In a different work, Fancellu et al. (2016) use a NN model for negation handling and their result on *SEM2012 shared task data has shown that a simple NN model outperforms the classifier such as CRF. Most recently, Chen (2019) proposes the use of Bi-LSTM deep learning model with word embedded to detect assertion and negation in clinical notes (EHR).

Altogether, it is argued that the challenging task of negation scope resolution is rooted in the biomedical domain. Most of the state-of-the-art TSA works adopts the simple linguistic ruled-based methods for scope resolution. More generally, ML classifiers will perform better than rule-based

Table 2. Classification of negation handling approaches in sentiment analysis, specifically twitter sentiment analysis

Approaches	Characteristics	References
Negation as features	Negated words are modelled in the form of features in feature vector such as negation features, shifters features, and polarity modification features	(Alistair & Diana, 2005 ; Carrillo-de-Albornoz et al., 2018 ; Pang et al., 2002 ; Wilson et al., 2005, 2009) in sentiment analysis. (Balika & Amini, 2016 ; Gupta & Joshi, 2019; Kolchyna et al., 2015 ; Kouloumpis et al., 2011 ; Mohammad et al., 2013; Pak and Paroubek 2010; Reitan et al., 2015 ; Zhu et al., 2014) in twitter sentiment analysis.
Reverse Polarity	Word under negated context is switched to opposite polarity	(Appel et al., 2016 ; Asghar et al., 2017 ; El Alaoui et al., 2018 ; Farooq et al., 2017 ; Kennedy & Inkpen, 2006 ; Polanyi & Zaenen, 2004) in sentiment analysis. (Kolchyna et al., 2015 ; Plotnikova et al., 2015 ; Ray & Chakrabarti, 2019 ; Saif et al., 2014 ; Thelwall et al., 2012) in Twitter sentiment analysis.
Shift Polarity	Polarity of word under negated context is shifted rather than complete inversal.	Tabaoda et al. (2011) in sentiment analysis. (Gupta & Joshi, 2019; Jurek et al., 2015; Muhammad et al., 2016) in Twitter sentiment analysis.
Corpus-based statistical approach	Negated word polarity is obtained from negated context lexicon. Thus each word is having two scores, one from affirmative context lexicon and other from negated context lexicon.	(Balikas & Amini, 2016; Kiritchenko et al., 2014; Reitan et al. 2015; Zhu et al. 2014) in Twitter sentiment analysis.

systems for scope resolution, although it requires a large number of handcrafted features. Moreover, training of ML models such as CRF is computationally extensive and requires large Twitter corpora annotated with cues and scopes.

4.3 Negation Handling

The final phase of negation modelling has to do with handling the relevant words marked under the negation scope in earlier phases. Prior work done for negation handling (as shown in Table 2) include: (a) reverse polarity, (b) shift polarity, or (c) modelling the word in a negated context as a feature for ML classification. Notably, both the reverse polarity and the shift polarity approaches cannot be modelled explicitly into the ML method of sentiment analysis.

Past studies in sentiment classification have employed supervised ML classifiers such as SVM. These algorithms use the bag-of-words model with various features such as n-grams, part of speech (POS), and similar others along with domain-specific features, and, then, the classifier has to identify itself, which word is opinionated and which is not. Thus, there is no explicit knowledge of the polar expression in the standard bag-of-words representation of text. Notwithstanding, in supervised ML tasks, negated words are modelled in the form of features in the feature vector. For instance, various researchers (e.g., Gupta & Joshi, 2019; Kiritchenko et al., 2014; Kolchyna et al., 2015; Pak & Paroubek, 2010) model the negation in their supervised classification approach with the help of negation feature (i.e., if a word “w” is under a negated context; then, token “NEG” is appended so that “w_NEG” is considered a feature by classifier rather than “w”). In this way, the classifier is aided to identify the word “w” as under the negation scope; thus, separate features reflect words under affirmative and negated contexts.

Generally, research has shown that the classification performance improvement is due largely to handling negation by adding the artificial feature into bag-of-words representation (e.g., Kiritchenko et al., 2014; Kouloumpis et al., 2011; Pak & Paroubek, 2010). Furthermore, few researchers (e.g., Carrillo-de-Albornoz et al., 2018; Kiritchenko et al., 2014) have handled the negation by generating negation features in addition to adding negated words (artificial word) in bag-of-words representation. Negation features can be the presence or absence of a negation, independent of the number of negated context(s), and many more. Wilson et al. (2005) are the first to model negation in the form of three feature groups including negation, shifter, and polarity modification features on phrase-level sentiment analysis. They add those features into bag-of-words features and observe significant improvement in the performance of supervised learning classifiers. Wilson et al. (2009) then extend their previous work (Wilson et al., 2005) to determine the effectiveness of each feature group separately, and, then, observe the negation feature group to be more effective than shifter feature group. Their conclusion is consistent with Alistair & Diana (2005). Even so, the improvement achieved by Wilson et al. (2005) is more than that of Alistair & Diana (2005) as the former work is dedicated to the finer level of sentiment analysis, while the latter work is focused on a document-level.

Put together, the work of Alistair & Diana (2005), Wilson et al. (2005), and Wilson et al. (2009) are considered pioneers in the negation modelling in feature form in the field of sentiment analysis. However, in TSA, many recent studies (e.g., Balikas & Amini, 2016; Kolchyna et al., 2015; Reitan et al., 2015; Mohammad et al., 2013) too modelled negation in feature form in their supervised learning approach. Reitan et al. (2015) and Kiritchenko et al. (2014), for instance, modelled negation in the form of various features such as binary negation (presence or absence of negation), or the count of negated context features. To evaluate the effectiveness of these negation features, experiments have been conducted by the removal of the negation feature group so as to observe performance degradation due to the removal of negation features.

More recently, Carillo-De-Albornoz et al. (2018) perform an extensive analysis of the different types of features such as syntactic, network-based, semantic, and many more for sentiment analysis of patient-authored texts in e-health forums. They modelled the negation in the form of binary features only (that is presence or absence of negation). However, no contribution of binary negation features into performance improvement has been observed. The authors suggest the use of a more sophisticated mechanism for negation modelling rather than just using binary features for negation. Thus, it may be concluded that early works in TSA used the different negation features for negation modelling, resulting mostly in significant performance improvement.

The above discussion suggests that negation handling requires the knowledge of polar expression. Various lexicons, which contain the prior polarities of polar expressions or polar words, are available for this purpose. Polanyi & Zaenen (2004) are among the firsts to present a computational model that handles negation via the knowledge of polar expression. They model different forms of negation through contextual valence shifting and simply reverse the polarity of the opinionated word under the negated context. Also, they handle the diminishers (a form of negation) by shifting the polarity to a certain amount. In reality, however, their model has not been implemented; thus, their model effectiveness can only be speculated.

Alistair & Diana (2005) also present the same reverse polarity model for handling the negation. However, they differ with (Polanyi & Zaenen, 2004) in their scope detection method (negate only those words which are immediately preceded by negation cue). In light of this, the polarity reversal becomes the most common and traditional approach that has been used by most researchers (e.g., El Alaoui et al., 2018; Han et al., 2018). The polarity reversal has even been used in the most recent works in TSA (e.g., Ray and Chakrabarti, 2019; El Alaoui et al., 2018). Nonetheless, few studies (e.g., Gupta & Joshi, 2019; Muhammad et al., 2016; Kiritchenko et al., 2014) point out the inefficiency of this polarity reversal approach because sometimes intensity of polarity only changes, for instance, “not excellent” is still more positive than the phrase “not good”.

Tabaoda et al. (2011) present a lexicon-based system called SO-CAL for sentiment analysis with the incorporation of negation and intensification. They propose a shift approach for negation handling, in which polarity of a negated term is shifted to a certain amount (4) because polarity does not always translate to its opposite. However, in many cases, their shift approach disagrees with the human judgement. Kiritchenko et al. (2014) observe that most of the positive terms when negated become negative while many of the negated terms (terrible) remained negated (not terrible). Hence, two Twitter-specific automatic lexicons (NRC Hashtag and S140 lexicon) are created from the hashtag and emoticon Twitter corpus respectively for the negation handling. These authors propose a corpus-based statistical approach in which each word is having two scores: one from affirmative context lexicon and other is from negated context lexicon. Their system is among those that are top-performing (69.02% macro-averaged F score in message level subtask) vis-à-vis other participants of SemEval-2013 Task 2: sentiment analysis in Twitter. They also observe a gain of more than 8.5 percentage points in F score due to their Twitter-specific automatic lexicons developed specifically for negation handling.

Zhu et al. (2014) present an improvement over their previous system (Kiritchenko et al., 2014) in the way of handling negation. They state that each negation word can have a different impact on sentiment analysis. Thus, they discriminated against each negation word in their system presented in SemEval-2014. For instance, the word “acceptable” in a tweet: “This is never acceptable” would be marked as “acceptable_NEVER” rather than “acceptable_NOT”. They observe an improvement in F1 score due to discriminating against negation words.

According to Balikas & Amini (2016), in subsequent years of the SemEval competition tasks for handling negation, many participants have used Twitter-specific automatic lexicons that had been developed by the Kiritchenko et al. (2014). In fact, the Kiritchenko et al. (2014) automated lexicons for corpus-based statistical negation handling should be updated as the vocabulary of Twitter changes with time. In their negation modelling by shifting approach, Muhammad et al. (2016) propose a SMARTSA lexicon-based sentiment analysis system that handles both local and global context of a term (by hybridizing SWN).

Unlike the Tabaoda et al. (2011) approach, Muhammad et al. (2016) consider negation cue as both sentiment bearing and a word modifier. Thus, the score of the word under a negated context is shifted according to the score of negation cue from SWN and not a fixed shifting. The authors integrated the SWN-derived negation cue score into the sentiment aggregation process of negated words. Further, they deploy a percentage scheme for handling diminishers; this time, however, the score of diminishers is not factored into the score aggregation process. As a result, significant performance improvement has been achieved over the Thelwall et al. (2012) state-of-the-art system (SENTISTRENGTH) within the contextual sentiment analysis.

In the area of negation modelling, Xia et al. (2016) propose the use of a shift elimination method. Here, the authors first create a dictionary of antonyms of opinionated words via a corpus-based approach. Then, negation cues are identified and eliminated with the word under a negated context replaced by its antonym. Notably, to some extent, this shift elimination method parallels that of the polarity reversal approach. More recently, Gupta & Joshi (2019) address the negation in their hybrid approach of TSA via the shift method. Somewhat dissimilar to Muhammad et al. (2016), their shifting idea incorporated a few rules for handling tweets in which the negation cue does not necessarily mean negation. Using this modified shift approach incorporated with negation exception rules, the approach taken by these researchers outperforms that of Sumanth-Inkpen (2015) approach deployed on the SemEval-2013 Twitter dataset.

Today, deep learning models have become popular with recent works showing significant performance improvement in sentiment analysis by the use of the deep learning framework. Deep learning models have been applied for the various tasks such as sarcasm detection for healthcare assisted system (Chen, 2019; Mahata et al., 2018; Yadav et al., 2018), and others (e.g., Teng et al. 2016). Among other more recent works using the deep learning framework to address negation

modelling is that of Teng et al. (2016), which advocate the use of the bi-directional LSTM model. These researchers apply context-sensitive lexicon-based methods, which are based on the simple idea of the weighted sum model (i.e., the sum of prior scores of opinionated words). Here, the NN model assigns weights with evaluation conducted on different domains, including the SemEval-2013 tweet dataset, movie review dataset, and mixed domains. Significant performance improvement on the benchmark Twitter dataset over state-of-the-art recurrent neural network baselines has been observed. Moreover, the negation handling approach via deep learning model observed here appears to be much easier to implement than the proposed rules by Taboada et al. (2011).

Qian et al. (2017) perform sentence-level sentiment analysis. Using intensification and negation LSTM regularizer respectively, these authors model intensification and negation; more simply, the shifting impact of negation-intensification is being addressed via the regularized LSTMs. As negation does not always translate conversely, these authors propose a transformation matrix for each negation word with its associated matrix to be learned by a deep-learning model. Thus, for ‘m’ number of negation words, a total ‘m’ transformation matrix would be trained. This approach has been evaluated on movie review dataset (MR) and Stanford Sentiment Treebank (SST) with results outperforming the Teng et al. (2016) approach on the SST dataset. Qian et al. (2017) also conduct ablation experiments in order to evaluate the contribution of negation regularizer by its removal from the model. They observe significant performance drop on both datasets (MR and SST) due to the removal of negation regularizer.

More recently, Wang et al. (2018) propose the use of the RNN-capsule model for sentiment analysis, that is, building one capsule for each class without needing linguistic knowledge. The Wang et al. model outperforms both the Teng et al. (2016) and Qian et al. (2017) methods on MR and SST datasets. Indeed, if the model for TSA includes prior polarity lexicon, then negation modelling would be more complex as prior polarity lexicons are context-independent and knowledge of polar expression also requires domain knowledge for which the polar term appears.

5. NEGATION MODELLING LIMITATIONS

Here, we discuss some of the limitations of accounting negation in TSA. A major limitation is the handling of negation exception scenarios known as false negation cues. In this case, while the negation cues are present in text, there is no sense of negation, which may lead to a misclassification, that is, the negation presence does not necessarily mean a true negation.

Consider a tweet:

- “I’m in Petrolia. The sun’s out, isn’t that good enough for you?”

In this tweet, the negation cue “isn’t” act as a non-cue because there is no sense of negation in this tweet, and, thus, the polarity of the word “good” wouldn’t be affected by “isn’t”. Reitan et al. (2015) observe that their negation cue detector struggled in separating words used as both negation and exclamation. Often, non-cue occurs as a determiner (part of speech (POS) tag is “D”) or exclamation (“No! I am not ready”).

Basically, there are two cases where negation cue acts as non-cue:

Exception case 1: When the negation is a part of the phrase that does not have a sense of negation such as “not only”, “no wonder”, “no one”, or “by no means”, for instance, “There is no one as good as him.”

Exception case 2: When the negation occurs in a negative rhetorical question (“isn’t that movie great?”). Negative rhetoric question is identified with two heuristics: it has a question and it has negation cue in the first three words of the question.

Figure 3. False negation cues (or phrases)

no doubt	not just	not only	no wonder
without a doubt	no one	at no times	
by no means	no question		
not to mention.....			

The authors have tested the hypothesis that accounting for negation exception cases while modelling negation can lead to the performance improvement. Many researchers ignore such cases with few exceptions (e.g., Gupta & Joshi 2019; Farooq et al., 2017; Kolchyna et al., 2015; Jia et al., 2009) where such cases are described although the implementation part has not been discussed clearly. Moreover, owing to its highly unstructured nature, heuristics presented by Farooq et al. (2017) and Jia et al. (2009) for negation exception cases would not work for Twitter. In the case of the biomedical domain, few studies (e.g., De-Albornoz et al., 2012; Chapman et al., 2001) also discuss on false negation cues and negative questions to show performance improvement. The review of extant literature has proven that implementing negation exception cases properly in the Twitter and medical domain might lead to performance improvement.

Another, somewhat ironic, limitation is when there is a sense of negation implicitly with no explicit negation cue being present. For example, a piece of text may be negated implicitly such as prevent, avoid, prohibit, and so on with such an implicit negation introduced via words such as “should”, “might”, and so on (“This phone should have worked even under water”). Existing studies for negation modelling in TSA focus largely on the syntactic negation, ignoring the impact of such implicit negation. Hence, it is difficult, if at all, to detect such implicit negation via simple linguistic rule-based methods. Xia et al. (2016) call such implicit negations as “sentiment inconsistencies” and present a statistical approach for the detection of sentiment inconsistencies in document-level sentiment analysis.

6. CONCLUSION

In retrospect, while the main application area of negation scope resolution has been on original biomedical text such as clinical text and discharge summaries, with advances in technology, attention has shifted towards sentiment analysis, specifically, Twitter. Specifically, this work has contributed to extend the conceptualization of negation with respect to the negation role in sentiment analysis, more specifically, TSA. Additionally, it summarizes the various approaches for negation scope detection and negation handling. The survey presented on the topic proves the relevance of negation modelling in TSA.

Notwithstanding, when it comes to the medical domain, much research focuses on negation scope resolution and negation cue identification. Yet, negation is much more than just scope resolution. One needs to handle the words, which come under the negation scope. Today, people have started using twitter for expressing their opinions regarding many topics and events. Indeed, twitter is also being used to express sentiments regarding current viruses, health diseases, and many more related healthcare issues. It is therefore critical to handle negation carefully in such tweets; otherwise, it might lead to the false predictions of diseases.

In this work, the authors have presented a survey on the negation modelling role in TSA and suggest that tweets on healthcare have become another significant application area of negation modeling, thereby not only contributing to explore the negation role in TSA, but also realizing the root of negation handling, which came from the biomedical domain. Importantly, our work suggests that negation is not a trivial task but entails many challenges in mixing the healthcare and twitter

domains such as implicit negations and negation exception cases. Encountering these challenges would help to further research in the biomedical and twitter domains.

Overall, exploring the negation effect in the biomedical domain has aided researchers to extend its conceptualization and applications to other domains. Moreover, novel statistical approaches (machine learning and deep learning) have benefited research in negation detection; for example, while syntactic and semantic knowledge may be required for an efficient negation model, the negation model also requires handling the negation exception cases, where negation cue act as non-cue. In the future, we anticipate extending this work to explore other forms of negation such as implicit and morphological negation as negation is not only characterized by the common negation cues. Hence, the potential of deep learning framework in negation detection task is an area waiting to be explored in the near future.

Finally, the more immediate task is to inspect the impact of false negation cues (negation exception cases) on the negation model in biomedical and healthcare related tweets. In this sense, it can be argued that negation modelling is related to polar expression knowledge; thus, recognising the genuine polar expression according to the context in which it is used is still brittle and is as important as modelling negation.

REFERENCES

- Agarwal, S., & Yu, H. (2010). Biomedical negation scope detection with conditional random fields. *Journal of the American Medical Informatics Association: JAMIA*, 17(6), 696–701. doi:10.1136/jamia.2010.003228 PMID:20962133
- Alistair, K., & Diana, I. (2005). Sentiment classification of movie and product reviews using contextual valence shifters. *Proceedings of FINEXIN*.
- Appel, O., Chiclana, F., Carter, J., & Fujita, H. (2016). A hybrid approach to the sentiment analysis problem at the sentence level. *Knowledge-Based Systems*, 108, 110–124. doi:10.1016/j.knosys.2016.05.040
- Asghar, M. Z., Khan, A., Ahmad, S., Qasim, M., & Khan, I. A. (2017). Lexicon-enhanced sentiment analysis framework using rule-based classification scheme. *PLoS One*, 12(2), e0171649. doi:10.1371/journal.pone.0171649 PMID:28231286
- Attardi, G., Cozza, V., & Sartiano, D. (2015). Detecting the scope of negations in clinical notes. In *Proceedings of the second Italian conference on computational linguistics CLiC-it* (Vol. 14, pp. 130-135). Trento. doi:10.4000/books.aaccademia.1286
- Baccianella, S., Esuli, A., & Sebastiani, F. (2010, May). Sentiwordnet 3.0: an enhanced lexical resource for sentiment analysis and opinion mining. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation* (Vol. 10, pp. 2200-2204). Valletta, Malta: European Language Resources Association.
- Balikas, G., & Amini, M. R. (2016). TwiSe at semeval-2016 task 4: Twitter sentiment classification. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)* (pp. 85-91). San Diego, CA: Association for Computational Linguistics. doi:10.18653/v1/S16-1010
- Bauer, L., Lieber, R. & Plag, I. (2015). *The Oxford Reference Guide to English Morphology*. Oxford University Press.
- Bhatia, P., Celikkaya, E. B., & Khalilia, M. (2019, January). End-to-end joint entity extraction and negation detection for clinical text. In A. Shaban-Nejad & M. Michalowski (Eds.), *Precision Health and Medicine. W3PHAI 2019*. Studies in Computational Intelligence.
- Burke, D., Ibekwe, C., Jerome, M., & Tran, M. (2018). Using Twitter to identify healthcare consumer sentiment. *Annals of Physical and Rehabilitation Medicine*, 61, e503–e504. doi:10.1016/j.rehab.2018.05.1173
- Cambria, E., Poria, S., Gelbukh, A., & Thelwall, M. (2017). Sentiment analysis is a big suitcase. *IEEE Intelligent Systems*, 32(6), 74–80. doi:10.1109/MIS.2017.4531228
- Carrillo-de-Albornoz, J., Vidal, J. R., & Plaza, L. (2018). Feature engineering for sentiment analysis in e-health forums. *PLoS One*, 13(11), e0207996. doi:10.1371/journal.pone.0207996 PMID:30496232
- Chapman, W. W., Bridewell, W., Hanbury, P., Cooper, G. F., & Buchanan, B. G. (2001). A simple algorithm for identifying negated findings and diseases in discharge summaries. *Journal of Biomedical Informatics*, 34(5), 301–310. doi:10.1006/jbin.2001.1029 PMID:12123149
- Chen, L. (2019). Attention-Based Deep Learning System for Negation and Assertion Detection in Clinical Notes. *International Journal of Artificial Intelligence and Applications*, 10(1), 1–9. doi:10.5121/ijaiia.2019.10101
- Choi, Y., & Cardie, C. (2008, October). Learning with compositional semantics as structural inference for subsentential sentiment analysis. In *Proceedings of the conference on empirical methods in natural language processing* (pp. 793-801). Honolulu, HI: Association for Computational Linguistics. doi:10.3115/1613715.1613816
- Councill, I. G., McDonald, R., & Velikovich, L. (2010, July). What's great and what's not: learning to classify the scope of negation for improved sentiment analysis. In *Proceedings of the workshop on negation and speculation in natural language processing* (pp. 51-59). Uppsala, Sweden: Association for Computational Linguistics.
- De Albornoz, J. C., Plaza, L., Díaz, A., & Ballesteros, M. (2012, June). Ucm-i: A rule-based syntactic approach for resolving the scope of negation. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics (*SEM)* (pp. 282-287). Montreal, Canada: Association for Computational Linguistics.

- Díaz, N. P. C., López, M. J. M., Vázquez, J. M., & Álvarez, V. P. (2012). A machine-learning approach to negation and speculation detection in clinical texts. *Journal of the American Society for Information Science and Technology*, 63(7), 1398–1410. doi:10.1002/asi.22679
- Doan, S., Yang, E. W., Tilak, S., & Torii, M. (2018, June). Using natural language processing to extract health-related causality from Twitter messages. In *Proceedings of the 2018 IEEE International Conference on Healthcare Informatics Workshop (ICHI-W)* (pp. 84-85). New York, NY: IEEE. doi:10.1109/ICHI-W.2018.00031
- El Alaoui, I., Gahi, Y., Messoussi, R., Chaabi, Y., Todoskoff, A., & Kobi, A. (2018). A novel adaptable approach for sentiment analysis on big social data. *Journal of Big Data*, 5(1), 12. doi:10.1186/s40537-018-0120-0
- Elkin, P. L., Brown, S. H., Bauer, B. A., Husser, C. S., Carruth, W., Bergstrom, L. R., & 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. doi:10.1186/1472-6947-5-13 PMID:15876352
- Fancellu, F., Lopez, A., & Webber, B. (2016, August). Neural networks for negation scope detection. In *Proceedings of the 54th annual meeting of the Association for Computational Linguistics* (volume 1: long papers) (pp. 495-504). Berlin, Germany: Association for Computational Linguistics. doi:10.18653/v1/P16-1047
- Farkas, R., Vincze, V., Móra, G., Csirik, J., & Szarvas, G. (2010, July). The CoNLL-2010 shared task: learning to detect hedges and their scope in natural language text. In *Proceedings of the Fourteenth Conference on Computational Natural Language Learning---Shared Task* (pp. 1-12). Uppsala, Sweden: Association for Computational Linguistics.
- Farooq, U., Mansoor, H., Nongillard, A., Ouzrout, Y., & Qadir, M. A. (2017). Negation Handling in Sentiment Analysis at Sentence Level. *JCP*, 12(5), 470–478. doi:10.17706/jcp.12.5.470-478 PMID:28097676
- Friedman, C., Alderson, P. O., Austin, J. H., Cimino, J. J., & Johnson, S. B. (1994). A general natural-language text processor for clinical radiology. *Journal of the American Medical Informatics Association: JAMIA*, 1(2), 161–174. doi:10.1136/jamia.1994.95236146 PMID:7719797
- Gautam, D., Maharjan, N., Banjade, R., Tamang, L. J., & Rus, V. (2018, May). Long Short Term Memory Based Models for Negation Handling in Tutorial Dialogues. In *Proceedings of the Thirsty-First International Florida Artificial Intelligence Research Society Conference*. Association for the Advancement of Artificial Intelligence.
- Gohil, S., Vuik, S., & Darzi, A. (2018). Sentiment analysis of health care tweets: Review of the methods used. *JMIR Public Health and Surveillance*, 4(2), e43. doi:10.2196/publichealth.5789 PMID:29685871
- Gupta, I., & Joshi, N. (2019). Enhanced twitter sentiment analysis using hybrid approach and by accounting local contextual semantic. *Journal of Intelligent Systems*. .10.1515/jisys-2019-0106
- Han, H., Zhang, Y., Zhang, J., Yang, J., & Zou, X. (2018). Improving the performance of lexicon-based review sentiment analysis method by reducing additional introduced sentiment bias. *PLoS One*, 13(8), e0202523. doi:10.1371/journal.pone.0202523 PMID:30142154
- Huang, Y., & Lowe, H. J. (2007). A novel hybrid approach to automated negation detection in clinical radiology reports. *Journal of the American Medical Informatics Association: JAMIA*, 14(3), 304–311. doi:10.1197/jamia.M2284 PMID:17329723
- Jia, L., Yu, C., & Meng, W. (2009, November). The effect of negation on sentiment analysis and retrieval effectiveness. In *Proceedings of the 18th ACM conference on Information and knowledge management* (pp. 1827-1830). Association for Computational Linguistics. doi:10.1145/1645953.1646241
- Jurek, A., Mulvenna, M. D., & Bi, Y. (2015). Improved lexicon-based sentiment analysis for social media analytics. *Security Informatics*, 4(1), 9. doi:10.1186/s13388-015-0024-x
- Kennedy, A., & Inkpen, D. (2006). Sentiment classification of movie reviews using contextual valence shifters. *Computational Intelligence*, 22(2), 110–125. doi:10.1111/j.1467-8640.2006.00277.x
- Khandelwal, A., & Sawant, S. (2019). *NegBERT: A Transfer Learning Approach for Negation Detection and Scope Resolution*. arXiv preprint arXiv:1911.04211.
- Kiritchenko, S., Zhu, X., & Mohammad, S. M. (2014). Sentiment analysis of short informal texts. *Journal of Artificial Intelligence Research*, 50, 723–762. doi:10.1613/jair.4272

Kolchyna, O., Souza, T. T., Treleaven, P., & Aste, T. (2015). Twitter sentiment analysis: Lexicon method, machine learning method and their combination. In G. Mitra & X. Yu (Eds.), *Handbook of sentiment analysis in finance*. Academic Press.

Konstantinova, N., De Sousa, S. C., Díaz, N. P. C., López, M. J. M., Taboada, M., & Mitkov, R. (2012, May). A review corpus annotated for negation, speculation and their scope. In *Proceedings of the Eighth International conference on Language Resources and Evaluation* (pp. 3190-3195). Istanbul, Turkey: European Language Resources Association.

Kouloumpis, E., Wilson, T., & Moore, J. (2011, July). Twitter sentiment analysis: The good the bad and the omg! *Proceedings of the Fifth International AAAI conference on weblogs and social media*.

Mahata, D., Friedrichs, J., Shah, R. R., & Jiang, J. (2018). Detecting personal intake of medicine from twitter. *IEEE Intelligent Systems*, 33(4), 87–95. doi:10.1109/MIS.2018.043741326

Mehrabi, S., Krishnan, A., Sohn, S., Roch, A. M., Schmidt, H., Kesterson, J., Beesley, C., Dexter, P., Max Schmidt, C., Liu, H., & Palakal, M. (2015). DEEPEN: A negation detection system for clinical text incorporating dependency relation into NegEx. *Journal of Biomedical Informatics*, 54, 213–219. doi:10.1016/j.jbi.2015.02.010 PMID:25791500

Mohammad, S. M., Kiritchenko, S., & Zhu, X. (2013). NRC-Canada: Building the state-of-the-art in sentiment analysis of tweets. In *Proceedings of the seventh international workshop on Semantic Evaluation Exercises (SemEval-2013)*. Atlanta, GA: Association for Computational Linguistics.

Moilanen, K., & Pulman, S. (2007, September). Sentiment composition. In *Proceedings of the Recent Advances in Natural Language Processing International Conference* (pp. 378-382). Academic Press.

Morante, R., & Blanco, E. (2012). *SEM 2012 shared task: Resolving the scope and focus of negation. 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)* (pp. 265-274). Montreal, Canada: Association for Computational Linguistics.

Morante, R., & Daelemans, W. (2009, June). A metalearning approach to processing the scope of negation. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning* (pp. 21-29). Boulder, CO: Association for Computational Linguistics. doi:10.3115/1596374.1596381

Morante, R., Liekens, A., & Daelemans, W. (2008, October). Learning the scope of negation in biomedical texts. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing* (pp. 715-724). Association for Computational Linguistics. doi:10.3115/1613715.1613805

Muhammad, A., Wiratunga, N., & Lothian, R. (2016). Contextual sentiment analysis for social media genres. *Knowledge-Based Systems*, 108, 92–101. doi:10.1016/j.knosys.2016.05.032

Mukherjee, P., Leroy, G., Kauchak, D., Rajanarayanan, S., Diaz, D. Y. R., Yuan, N. P., & Colina, S. et al. (2017). NegAIT: A new parser for medical text simplification using morphological, sentential and double negation. *Journal of Biomedical Informatics*, 69, 55–62. doi:10.1016/j.jbi.2017.03.014 PMID:28342946

Mutalik, P. G., Deshpande, A., & 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: JAMIA*, 8(6), 598–609. doi:10.1136/jamia.2001.0080598 PMID:11687566

Ou, Y., & Patrick, J. (2015). Automatic negation detection in narrative pathology reports. *Artificial Intelligence in Medicine*, 64(1), 41–50. doi:10.1016/j.artmed.2015.03.001 PMID:25990897

Owoputi, O., O'Connor, B., Dyer, C., Gimpel, K., Schneider, N., & Smith, N. A. (2013, June). Improved part-of-speech tagging for online conversational text with word clusters. In *Proceedings of the 2013 conference of the North American chapter of the association for computational linguistics: human language technologies* (pp. 380-390). Association for Computational Linguistics.

Pak, A., & Paroubek, P. (2010, May). Twitter as a corpus for sentiment analysis and opinion mining. In *Proceedings of the International Conference on Language Resources and Evaluation* (Vol. 10, No. 2010, pp. 1320-1326). Valletta, Malta: European Language Resources Association.

- Pang, B., Lee, L., & Vaithyanathan, S. (2002, July). Thumbs up?: sentiment classification using machine learning techniques. In *Proceedings of the ACL-02 conference on Empirical methods in natural language processing- Volume 10* (pp. 79-86). Association for Computational Linguistics. doi:10.3115/1118693.1118704
- Peng, Y., Wang, X., Lu, L., Bagheri, M., Summers, R., & Lu, Z. (2018). NegBio: A high-performance tool for negation and uncertainty detection in radiology reports. *AMIA Joint Summits on Translational Science Proceedings AMIA Summit on Translational Science, 2018*, 188–196. PMID:29888070
- Polanyi, L., & Zaenen, A. (2004). Contextual valence shifters. In *Computing Attitude and Affect in Text: Theory and Applications* (Vol. 20, pp. 1–9). Springer.
- Pröllochs, N., Feuerriegel, S., & Neumann, D. (2016). Negation scope detection in sentiment analysis: Decision support for news-driven trading. *Decision Support Systems*, 88, 67–75. doi:10.1016/j.dss.2016.05.009
- Qian, Q., Huang, M., Lei, J., & Zhu, X. (2016). Linguistically regularized lstms for sentiment classification. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics* (volume 1: long papers) (pp. 1679-1689). Association for Computational Linguistics.
- Ray, P., & Chakrabarti, A. (2020, August 14). A Mixed approach of Deep Learning method and Rule-Based method to improve Aspect Level Sentiment Analysis. *Applied Computing and Informatics*. Advance online publication. doi:10.1016/j.aci.2019.02.002
- Reece, A. G., Reagan, A. J., Lix, K. L., Dodds, P. S., Danforth, C. M., & Langer, E. J. (2017). Forecasting the onset and course of mental illness with Twitter data. *Scientific Reports*, 7(1), 13006. doi:10.1038/s41598-017-12961-9 PMID:29021528
- Reitan, J., Faret, J., Gambäck, B., & 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* (pp. 99-108). Association for Computational Linguistics. doi:10.18653/v1/W15-2914
- Rokach, L., Romano, R., & Maimon, O. (2008). Negation recognition in medical narrative reports. *Information Retrieval*, 11(6), 499–538. doi:10.1007/s10791-008-9061-0
- Sumanth, C., & Inkpen, D. (2015). How much does word sense disambiguation help in sentiment analysis of micropost data? In *Proceedings of the 6th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis* (pp. 115-121). Association for Computational Linguistics. doi:10.18653/v1/W15-2916
- Taboada, M., Brooke, J., Tofiloski, M., Voll, K., & Stede, M. (2011). Lexicon-based methods for sentiment analysis. *Computational Linguistics*, 37(2), 267–307. doi:10.1162/COLI_a_00049
- Taylor, S. J., & Harabagiu, S. M. (2018). The Role of a Deep-Learning Method for Negation Detection in Patient Cohort Identification from Electroencephalography Reports. *AMIA ... Annual Symposium Proceedings - AMIA Symposium. AMIA Symposium, 2018*, 1018–1027. PMID:30815145
- Teng, Z., Vo, D. T., & Zhang, Y. (2016, November). Context-sensitive lexicon features for neural sentiment analysis. In *Proceedings of the 2016 conference on empirical methods in natural language processing* (pp. 1629-1638). Association for Computational Linguistics. doi:10.18653/v1/D16-1169
- Thelwall, M., Buckley, K., & Paltoglou, G. (2012). Sentiment strength detection for the social web. *Journal of the American Society for Information Science and Technology*, 63(1), 163–173. doi:10.1002/asi.21662
- Vincze, V., Szarvas, G., Farkas, R., Móra, G., & Csirik, J. (2008). The BioScope corpus: Biomedical texts annotated for uncertainty, negation and their scopes. *BMC Bioinformatics*, 9(11), S9. doi:10.1186/1471-2105-9-S11-S9 PMID:19025695
- Wang, Y., Sun, A., Han, J., Liu, Y., & Zhu, X. (2018, April). Sentiment analysis by capsules. In *Proceedings of the 2018 World Wide Web Conference* (pp. 1165-1174). International World Wide Web Conferences Steering Committee. doi:10.1145/3178876.3186015
- Wilson, T., Hoffmann, P., Somasundaran, S., Kessler, J., Wiebe, J., Choi, Y., & Patwardhan, S. et al. (2005, October). OpinionFinder: A system for subjectivity analysis. In *Proceedings of HLT/EMNLP 2005 Interactive Demonstrations* (pp. 34-35). Association for Computational Linguistics. doi:10.3115/1225733.1225751

- Wilson, T., Wiebe, J., & 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* (pp. 347-354). Association for Computational Linguistics. doi:10.3115/1220575.1220619
- Wilson, T., Wiebe, J., & Hoffmann, P. (2009). Recognizing contextual polarity: An exploration of features for phrase-level sentiment analysis. *Computational Linguistics*, 35(3), 399–433. doi:10.1162/coli.08-012-R1-06-90
- Xia, R., Xu, F., Yu, J., Qi, Y., & Cambria, E. (2016). Polarity shift detection, elimination and ensemble: A three-stage model for document-level sentiment analysis. *Information Processing & Management*, 52(1), 36–45. doi:10.1016/j.ipm.2015.04.003
- Yadav, S., Ekbal, A., Saha, S., & Bhattacharyya, P. (2018, May). Medical sentiment analysis using social media: towards building a patient assisted system. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. ELRA.
- Zhu, Q., Li, J., Wang, H., & Zhou, G. (2010, October). A unified framework for scope learning via simplified shallow semantic parsing. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing* (pp. 714-724). Association for Computational Linguistics.
- Zhu, X., Kiritchenko, S., & Mohammad, S. (2014, August). Nrc-canada-2014: Recent improvements in the sentiment analysis of tweets. In *Proceedings of the 8th international workshop on semantic evaluation (SemEval 2014)* (pp. 443-447). Association for Computational Linguistics. doi:10.3115/v1/S14-2077