A Semantic Approach to Negation Detection and Word Disambiguation with Natural Language Processing

Izunna Okpala
okpalaiu@mail.uc.edu
University of Cincinnati
Cincinnati, Ohio, USA

Guillermo Romera Rodriguez
gkr5144@psu.edu
Penn State University
Philadelphia, Pennsylvania, USA

Andrea Tapia
axh50@psu.edu
Penn State University
Philadelphia, Pennsylvania, USA

Shane Halse
halsese@ucmail.uc.edu
University of Cincinnati
Cincinnati, Ohio, USA

Jess Kropczynski
kropczjn@ucmail.uc.edu
University of Cincinnati
Cincinnati, Ohio, USA

ABSTRACT

This study aims to demonstrate the methods for detecting negations in a sentence by uniquely evaluating the lexical structure of the text via word-sense disambiguation. The proposed framework examines all the unique features in the various expressions within a text to resolve the contextual usage of all tokens and decipher the effect of negation on sentiment analysis. The application of popular expression detectors skips this important step, thereby neglecting the root words caught in the web of negation and making text classification difficult for machine learning and sentiment analysis. This study adopts the Natural Language Processing (NLP) approach to discover and antonimize words that were negated for better accuracy in text classification using a knowledge base provided by an NLP library called WordHoard. Early results show that our initial analysis improved on traditional sentiment analysis, which sometimes neglects negations or assigns an inverse polarity score. The SentiWordNet analyzer was improved by 35%, the Vader analyzer by 20% and the TextBlob by 6%.

CCS CONCEPTS

• **Computing methodologies → Natural language processing**.

KEYWORDS

Negation Detection, Artificial Intelligence, Machine Learning, Text analysis, Word Sense Disambiguation

1 INTRODUCTION

Negation is an inescapable variable when it comes to human communication. To convey spoken words or written texts, negations need not to be neglected because they are integral to everyday expression [6, 13]. Popular text mining techniques have been studied over the years as to how they can automatically detect emotions, attitudes, sentiments, and perceptions, to name a few [15]. Some notable features that have received less attention in text analysis are negations and multiple negatives [4]. As a result, the overall accuracy of popular sentiment analyzers has decreased. Understanding negation cues is essential for looping every word in a sentence together and capturing the elements that were affected.

The term "negation" is a linguistic feature that negates the meaning of its closest neighbor in a sentence; i.e., it is the semantic opposite of a phenomenon [36]. This feature (negation) often poses a challenge to text mining techniques, leading to incorrect classification. This is because, generally, text mining techniques and sentiment analysis programs classify texts using a combined factor of words that are commonly associated with emotions [33]. However, most techniques or algorithms never factor in the effect of negations on the cluster of words. They only concentrate on the evaluation of the polarity score for a given sentence, thus creating a potential scenario for misclassification if the sentence contains negations. The typical way the computer handles negation is to invert the polarity of the lexical item that lies next to the negator in a phrase [18] (e.g., good: 0.5 and not good: -0.5). In this study, we propose a different method. Rather than just inverting the value, we recommended developing a function that computes the polarity of the negated word based on a list of antonyms from 5 different dictionaries, as well as the overall mean score of the polarity of all the antonyms.

This study explores a novel approach to detecting and analyzing negations within sentences. This was achieved using a mixed approach: the detection of negation and word sense disambiguation. The detection of negation occurs through the use of pre-defined labels representing the negation signals, while the processing is applied through the use of sequence labeling and an antonymization function that averages out the polarity of all the antonyms, putting Parts of Speech (POS) tagging into perspective. Early results showed that our approach improved on conventional sentiment analysis, which sometimes ignores word negations or assigns an inverse polarity score without considering the context

of the negated words. The SentiWordNet analyzer was enhanced by 35%, the Vader analyzer by 20%, and the TextBlob analyzer by 6%.

2 BACKGROUND

Negations are extremely important in all human languages. The highly complex but subtle expression of negation in natural language, as conveyed in parts of speech, defies the basic syntactic structure of logical negation (adverbs, verbs, adjectives, quantifiers) [11]. Negations in literature are words that systemically negate another expression. Double negatives, on the other hand, a terminology often misinterpreted to mean negations, are two negative words occurring concurrently in a sentence [26, 36]. For example, if we have the text, *"the prisoner was a dreaded terrorist",* The two words *"dreaded"* and *"terrorists"* occurring in the same sentence are referred to as *double negatives or multiple negatives.* An example of negation, however, is when we have *"The warden is not good"*; the word *"not"* negates the phrase, *"good."* To shorten the sentence *"the warden is not good,"* we have, *"the warden is bad."* This form of the sentence structure contains a negation clause "not" and a positive adjective "good." For better interpretation by a non-native speaker or even a machine, the negating factor needs to cancel out the adjective and derive an alternative word [26] i.e., when a negation and a positive are used together, the sentence is usually transformed into a negative one, vice versa; mathematically we would have $-x * -x = x^2$. The square of x (x^2) shows that there are some permutations that took place after the removal of the negation sign.

In some instances where negations are not inverted, i.e., whenever the relational expression is not a phrase that modifies or qualifies an adjective, verb, or other adverb, or a phrase naming an attribute, attached to or grammatically related to a noun to change or describe it, the negation stands. A good example is in the text *"Soldiers could **never attend** a parade."* The negation signal *"never"* could not cancel out the word *"attend"* because it does not modify or qualify an adjective, nor does it name an attribute. But a negation signal can negate a positive clause that modifies or qualifies and adjective; e.g., *"The boy is not dirty"* becomes *"The boy is clean".* Mathematically, we represent this as: $-x * x = -x^2$. Having negation in a sequence seems proper, but this burdens text classification and can be misleading [20]. Nonetheless, they are occasionally used in everyday casual conversation, and numerous examples can be found on popular social media platforms [28].

2.1 Previous Work

Research on negation detection is critical for building intelligent systems that use texts as the foundation of their computational process. Academics have studied the detection of negatives and negations. Some of the approaches used are rule-based, machine learning, and conditional random field (CRF) methods [19]. What these have in common is that they methodically addressed the subject of negation, which can be considered the foundation of this research. Rule-based techniques handle negation detection using static, non-dynamic rules [5, 25]. Machine learning methods focus on teaching a machine how to distinguish negations. It detects negations by automatically generating regular expression patterns

[30], while the CRF follows the natural order of events in a sentence [2].

2.2 The Theory of Double Negatives

The concept of words with negative connotation in languages is the foundation of double negatives [16]. In standard English, each subject-predicate composition should only have one negative form for ease of interpretation, but double negatives are unavoidable due to the unconventional statement structure used in social discussions [28]. Though it's common to believe that double negatives are just an artificial deviation from the norm, in some languages, like Spanish, anything other than a double negative is grammatically improper, i.e., they utilize double negatives as a proper expression [29]. This is not so in English. However, the use of double negatives is prevalent in casual forums, discussion boards, and social media platforms [28]. Often times, data is collected from those mediums for predictive analytics.

2.3 Lexicon-Based Analysis for Dealing with Negations

In order to understand the meaningful bits in a text, lexicon-based analysis can be used for negative context detection. This technique employs about 6300 words [18], and was constructed manually using SentiWordNet (a lexical resource for text mining) [3] as the baseline. Each word in the lexicon has a rating ranging from 100 (the most negative) to 100 (the most positive). This study also found evidence that some words appear to have a neutral rating when handled alone, but when put in between other words, the rating changes. This is based on the construction or lexical structure of the sentence. Depending on the sentiment analysis tool to be used, this rating differs; some use a ratio of from -1 to +1 [3, 12], while others use 0 to +1 rating [3]. Those ratings denote the score given to each word in their data dictionary. The lexicon-based analysis adopts a probabilistic model. To make this work, a labelled dataset of over a million data points was used [18]. As demonstrated by [18], the approach for calculating the probabilities follows a recursive mechanism by going through the words represented in the sentence.

2.4 Natural Language Processing (NLP)

The growing complexity of text data has made NLP applications increasingly vital in analyzing large volumes of data [10, 14]. NLP solutions are one of the backbones of some intelligent models like BERT, XLNet, and GPT, which have advanced the cause of sentiment analysis [21, 32], and machine language translation [27].

NLP-enabled models propel the understanding of language structure and interpretation. The role of NLP, therefore, is to autonomously interpret human languages to return contextual patterns [34]. This technology involves methods for extracting meaning from text data [31]. Another notable application of NLP in the research space is in the monitoring of fake news or the detection of spam emails, which has proven to be successful [23]. The applicability of NLP comes in a variety of forms. In our everyday discourse, such as social media conversations, negation signals are employed. In practice, some NLP techniques used in this space are Word2Vec and GloVe [38]. BERT and XLNet deliver state-of-the-art outcomes as well [7]. However, each model has its

own set of constraints. Most algorithms underperform when confronted with negation, which is why it is critical to identify novel approaches to overcome this issue in the NLP domain.

2.5 Word Negation and Sequence Labeling

The accuracy of NLP models such as sentiment or perception analysis depends on word negation and sequence labeling. Some NLP models typically work by analyzing each word, or sequence of words, independently. These algorithms deconstruct text into its minimum units called tokens [35]. Once these tokens have been cleaned up, the algorithms read each word or tokenized entity independently and come up with a result based on the different tokens found in the sentence or text. However, one of the limitations of most sentiment analysis algorithms is that they do not make sense out of multiple words, but just one word at a time [1]. This can create a problem where sentences are mislabeled or miscategorized, causing problems down the line [9]. For example, a sentence that reads *"Vaccines are **not** bad"* could be mislabeled because the presence of the word *not*, but we know that sentence does not carry a negative sentiment.

While popular opinion mining tools like sentiment analyzers try to find the negatives and positives, it is important to make a clear distinction between the terms "negative" and "negation." The most common negation signals are *not, non, never, and neither* [22]. Since these negation signals are considered to be stopwords in the Natural Language Toolkit (NLTK), spacy and sklearn [22], it's crucial to understand how they affect NLP tasks. There have been different attempts to overcome this drawback. The most recent and state-of-the-art approach was Bidirectional Encoder Representations from Transformers (BERT). Unlike other approaches, the novelty of BERT resides in its bidirectional training capacity, which means that sentences are read in both directions to get the complete context of the sentence. The trick here is also that tokens are read at the same time and not in order (either left-to-right or right-to-left). This allows BERT to take the context of each word into account. Similar in its construction, XLNet [32] improved its masking mechanism with peculiar assumptions during its pre-training stage, and improved over the work done by BERT. Despite these advances, studies [8] have shown that even these approaches have still struggled with negation. Which is why this study aims to examine the problem in a holistic manner.

3 METHODOLOGY

This paper addressed two research problems by utilizing Natural Language Processing (NLP) techniques, a lexicon-based approach, and sequence labeling. To guide the design of the aforementioned automated tool, we answered the following two questions:

- **RQ1:** *How do we detect negation in a given text?*
- **RQ2:** *How can a automated system apply negation to appropriate words to improve sentiment analysis?*

NLP was used mainly because it guides the extraction of information from texts in their natural form. The most frequent lexicon-based technique for negation is to reverse the polarity of the object that is affected by the negator in a sentence [18]. This study presents an alternative approach by developing a

disambiguation function that can be used to average out the polarity scores of the negated word antonyms using five dictionaries. Beyond sentiment detection, we propose using antonyms to construct a human-readable semantic construction of the entire sentence, taking a context-based approach into consideration. First, detecting negation in a sentence requires the use of an automated decontraction system, as some negations may be in contraction mode e.g., *"isn't = **is not**"*. With the deconstruction pattern shown in algorithm 1, words that are merged together as contractions are separated so that the system can read the negations appropriately. An annotated dictionary containing negations is created to mirror the words that should be matched when performing tokenized search on a list or text data. The tokenized search is implemented using the list comprehension technique. This method returns a list containing the index of all the negations in a text dataset.

Second, after detecting the negations, the researchers used NLP's Parts of Speech (POS) and Sequential Labeling algorithms to label all of the keywords in the text. With this in place, the machine can readily interpret the POS in each word as to whether the nearest neighbor to a negating signal is a phrase that modifies or qualifies an adjective, verb, or other adverb, or a phrase naming an attribute attached to or grammatically related to a noun to change or describe it. If the POS is an adjective (ADJ) or an adverb (ADV), the system eliminates the negation and inverts the expression. Otherwise, the negation stands.

3.1 Data Collection

Our data collection phase was comprised of two stages. The first stage was searching for appropriate datasets that could allow our model to be trained. The complexity in this stage was finding appropriate training datasets that contained negation and, at the same time, were validated by prior literature. The most suitable archive that matched our approach was the Stanford Contradiction Corpora dataset . The second stage was the collection of the Stanford Contradiction Corpora dataset from (https://nlp.stanford.edu/projects/contradiction/) to test our algorithm. This stage was very critical as it was the basis for demonstrating the efficacy of the algorithm and setting the pace for testing with a larger dataset.

3.2 Data Pipeline

Considering the nature of the study, our pipelining architecture follows a sequential order in combining various functions down to the disambiguation function as represented in Figure 1.

3.2.1 Text Decontraction. Words or syllables are often condensed in contractions. By deleting key letters, these shorter versions or contractions of words are generated. In the case of English contractions, one of the vowels is removed from the word. A simple example is when *do not* is shortened to *don't*. The text decontraction function helps in converting each contraction to its enlarged form to ease machine processing. The pseudocode in algorithm 1, demonstrates a procedure for the decontraction of *n't* and *'t* phrases in a sentence to *not*. This procedure can be implemented in any programming language, but the Python programming language was chosen for this study because of its
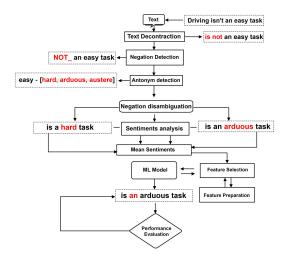
Figure 1: Data Pipeline

extensive library for machine learning and natural language processing [24].

Algorithm 1 Text Decontraction

procedure DECONTRACT(*text*)
    if [n't, 't] IN *text* **then**
        text ← [*not, not*]
        return *text*
    end if
end procedure

3.2.2 Negation Detection. Negations, as discussed in the literature review, are not necessarily negative phrases but words that negate the meaning of another. To detect negation, a tokenized search is performed on a list of textual data. An annotated dictionary including negations is constructed to mirror the terms that should be matched. The tokenized search is implemented using the list comprehension approach. This method returns a list with the indexes of all negations in a text data set. The procedure described in algorithm 2, brings negation detection into play. The reference list is first instantiated, displaying the negations in a sentence that should be identified. The procedure then searches through a given text to detect the negation index (key), after the text must have been decontrated using an algorithm 1.

3.2.3 Text Antonymization. The manual process of antonymizing a given phrase in its natural form involves checking words and their opposites in a dictionary. To do this automatically, we need to make use of some standard dictionaries. In this section, we will be making use of five dictionaries: the Collins Dictionary, the Merriam-Webster Dictionary, synonym.com, thesaurus.com, and NLTK's wordnet. These dictionaries, when combined, provide us with rich content as opposed to using only the popular WordNet database. Furthermore, a Python package called *wordhoard* provides an elaborate library for combining all of these dictionaries [17]. Algorithm 3 demonstrates the process of antonymization. In some cases, when an antonym

---

Algorithm 2 Detect Negation

procedure DETECT_NEG(*text*)
    ref ← ["*not*","*nor*","*never*","*neither*"]
    text ←text.*split*()
    for *ref* IN *text* **do**
        if (any *item* IN *ref*) **then**
            st = set(*ref*)
            if (*e* IN *st*) **then**
                for *i, e* IN enumerate(*text*) **do**
                    return *i*
                end for
            end if
        end if
    end for
end procedure

---

is not found, the procedure queries the synonym database, then repeats the process to get the antonym. With this in place, we can easily move ahead to carry out sequence labeling, part of speech tagging, and negation disambiguation based on some pre-defined criteria.

Algorithm 3 Get Antonymns

dictionary[0] ← *collins_dictionary*()
dictionary[1] ← *merriam_webster*()
dictionary[2] ← *synonym_com*()
dictionary[3] ← *thesaurus_com*()
dictionary[4] ← *wordnet*()
error ← *notfound*
procedure GET_ANTONYMS(*word*)
    antonyms = Antonyms(*word*).find_antonyms(*dict*[*i*])
    if *error* NOT IN *antonyms* **then**
        return *antonymns*
    else
        synonym ← *get_synonyms*(*word*)
        key ← *synonym*[0].*find_antonyms*(*dict*[*i*])
        antonyms = *Antonyms*(*key*)
        return *antonym*
    end if
end procedure

3.2.4 Sequence Labeling, Parts of Speech (POS) Tagging, and Negation Disambiguation. Sequence labeling in this phase is a form of pattern recognition in machine learning that includes assigning a categorical label to each component of a sequence of observed values using an algorithm ([37]). Part of speech (POS) tagging, which attempts to assign a part of speech to each word in an input sentence or document, is a frequent example of a sequence labeling problem. The POS of a word in a sentence is dependent on its neighbors, thus the use of "sequence" in the term "sequence labeling." The decontract and negation detection procedures play a crucial role in this procedure. Algorithm 4 demonstrates the process of tagging each word in a sentence using the NLTK POS Tagging function. The text input is converted to a list and

sequentially labeled using the NLTK library. The label is based on the POS function, which detects distinct parts of speech in text. The tokenized system detects tags associated with **JJ** (Adjective), **VBG** (Verb Gerund), and **VBN** (Verb Past Participle) that have been negated. The main structure of the pseudocode shows that the procedure only removes the negation and inverts the phrase if the POS is an adjective, verb, or adverb. If this is not true, the negation remains in effect. With this in place, the machine can quickly determine whether the neighbor of a negating signal needs to be inverted or not, and it takes three steps towards fulfilling this.

Algorithm 4 Word Disambiguation

---

    procedure DISAMB($sentence$)
        $token \leftarrow nltk.word\_tokenize(sentence)$
        $tags \leftarrow nltk.pos\_tag(token)$
        if $tags[found] = NEGATED$ **then**
            $token[int\_index + 1] = get\_antonyms(token)$
            $token.pop(int\_index)$
            return $token$
        else
            return $token$
        end if
    end procedure

In training a model to understand the patterns for context purposes, we used the BERT model and a recurrent neural network (RNN) to build a dictionary of tags associated with the sequence of words in a sentence. As the antonymization takes place, the neural network model stores the POS of the negated text and an array of the POS of the sentence that accompanies the negation. This value was then plugged into the BERT model and RNN to put the negation into context. The neural network studies the whole POS associated with the sentence that was just analyzed and stores such a pattern in its model. It does not store the exact word because the next sentence that comes might not have the same wording. POS tagging plays a major role in helping the machine learning model learn specific parts of speeches and how to work when similar situations come into play.

4 RESULTS

The result from the algorithm was tested with various sentiment analyzers to see if there was a difference in prediction. It is important to point out that most statements that were negated were not necessarily positive or negative statements. A large portion of the dataset consisted of neutral statements. This will also help to test the efficacy of our algorithm because the fact that some statements contain negation does not make the statement a negative sentence, e.g., *"Samuel L. Husk does not work for the Council of Great City Schools."* This statement is neutral, and when returned by our algorithm, it should maintain a neutral stance.

Early results from the Stanford Contradiction Corpora dataset show improvement in predicting sentiments in statements. We will be demonstrating this through the use of three sentiment analyzers. Most importantly, we will be making use of the polarity scores and not the binary output. The result from Figure 2 which was

derived from a 30% randomized sample of the original data, shows a slight similarity in output. The randomized sample was used in order not to cluster the figure. The slight similarity in output with Figure 2 was because Vader already implements a system that assigns a negative polarity score to sentences that contain negation. The job of Vader is not to look for a dictionary antonym to use in place of the negated words, but rather to apply a negative score to such a statement. Our algorithm does that by using the exact part of speech tagging against an antonymized word. Vader's improvement is remarkable, but additional improvement is required in terms of using the actual antonyms and not necessarily negating the polarity scores. The other area where Vader played a crucial role was maintaining neutral scores whenever a negation was used instead of going in an extreme direction.
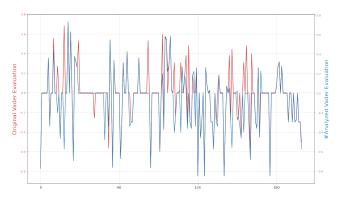

Figure 2: Vader test on the original and antonymized text

The TextBlob tool works in the same fashion as its Vader counterpart, but with higher accuracy, as we can see from Figure 3. There is a higher similarity to that of the original text. Though this tool does a great job, the slight difference is very key to making a crucial decision about a delicate situation with data.
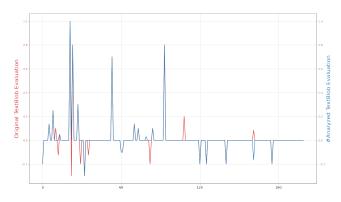

Figure 3: TextBlob test on the original and antonymized text

The tool that performed the least was the SentiWordNet. This tool needs data to be preprocessed first before passing it through the analyzer, but one key point to note is that the word "not" is a stop word and will end up being eliminated when such analyzers are used. If this is the case, the valuable power derived from such

a system will be lost, leaving the user with values that can be misleading. The graph representation of the output is shown in Figure 4.
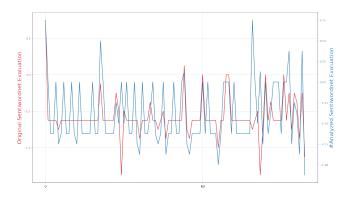

Figure 4: Sentiwordnet test on the original the antonymized text

To compare all the outputs together, we make use of a correlogram heatmap shown in Figure 5. The heatmap is presented in the form of a 2D correlation matrix with colored cells and a monochromatic scale *('RdBu_r')*. The number of measurements that match the dimensional values determines the color value of the cells. This allows us to easily discover patterns of occurrence and spot anomalies. Here is the linear relationship between all the different sentiment analyses. From the matrix, TextBlob gathered a 94% correlation between the original text and the analyzed text, while Vader returned 80%. It is also important to note the similarities between the different tools. The TextBlob and Vader output remained unchanged at 27% while SentiWordNet fluctuated, showing its weaknesses in terms of correlation.
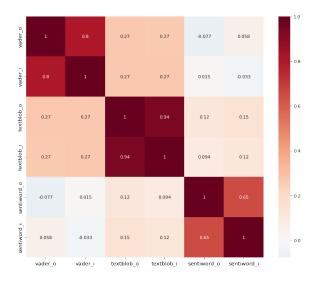

Figure 5: Correlation Heatmap of all the sentiments

A negation cue is a token that triggers the semantic opposite of a phenomenon in a sentence. As a result, negation connects one

statement, *x* to another expression whose meaning is diametrically opposed to *x*'s. This relationship can be expressed in a variety of ways, both syntactically and pragmatically. There are various types of semantic opposition. From both a historical and systematic standpoint, semantic opposition is primarily concerned with negation and opposition in natural language.

Testing our analyzed dataset with the three sentiment analyzers gave a considerable amount of improvement. With a 94% correlation between the original text and the studied text by Textblob and 80% correlation with Vader, there was a considerable amount of improvement. The highest impact of our program was on SentiWordNet, which had a correlation score of 65% showing that our algorithm improved the SentiWordNet output by 35%.

Furthermore, since the application of artificial neural networks in this research allows for the model to work in a progressive manner, similar to the human mind, the solution shows promise in increasing the accuracy of traditional sentiment analyzers, thereby improving future research in domains that have to do with sentiment analysis, such as crisis informatics.

4.1 Limitations

Negation has been studied as a truth-functional operator, among other things. The precise relationship between negation and words makes negation disambiguation a complex task. Statements caught in the web of negation are also psychologically more complex and difficult to process. This imbalance has been attributed to logic or semantics by several linguists and psychologists, such as the assertion that every negation implies a matching affirmative but not the other way around. This means that there is no special rule for negation disambiguation; rather, the context of the analyzed sentence plays a huge role in understanding which antonym fits for a particular negation signal and not only POS tagging. Since the analysis didn't utilize machine learning for checking grammar, identifying grammar clarity, and checking delivery mistakes, further work needs to be done in this section to improve the neural network that was trained based on POS tagging.

4.2 Future Work

Further research is needed to evaluate the system with a larger dataset, especially with data from social media aside from the prepared data that was analyzed in this study. A larger dataset spanning numerous annotations and not necessarily negations can be explored for extensive search and also for improving the neural network algorithm used to store POS tags assigned to analyzed sentences.

The incorporation of a natural grammar-checking program based on deep learning methods will have a significant impact on the analyzed sentences in terms of clarity and error corrections. When dealing with text data, deep learning's capacity to process vast amounts of features makes it incredibly strong. Exploring this technique will improve the work further and, in turn, build a large corpus for automated language understanding and grammar checking.

5 CONCLUSION

This study examined and demonstrated methods for detecting negations in a phrase by examining the distinguishing properties of textual expressions in a sequence and resolving whether or not to invert a word that was negated in a sentence. Three sentiment analyzers were used to test the algorithm, and the results show an improvement in prediction. The improvement was solely because of the way those tools ignored the negation variable, sometimes discarding it as a stop word, resulting in the misclassification of sentences and thereby making accuracy in data prediction with sentiment analyzers lower. This study applied the natural language processing (NLP) approach to find and invert negated terms for proper text classification, resolving the research questions that were asked in the introduction. This approach acts as a lens, reading through a given word sequence utilizing an intelligent NLP algorithm to discover negation signals and invert the keywords that were negated.

Early results suggest that our detailed analysis improved standard sentiment analysis, which sometimes ignores word negations or assigns an inverse polarity score without considering the context of the negated words. The SentiWordNet analyzer was enhanced by 35%, the Vader analyzer by 20%, and the TextBlob analyzer was improved by 6%. Though this study utilized the NLP process, training a machine learning model seemed promising in terms of being able to circumvent the arduous stages necessary in the NLP process. Going by our findings, negation detection and the antonymization of negated words are helpful for researchers involved in crisis evaluation using text analysis. This will allow academics and/or industry practitioners who derive insights from textual data to conduct their analyses with greater accuracy than was previously possible.

REFERENCES

[1] AM Abirami and V Gayathri. 2017. A survey on sentiment analysis methods and approach. In *2016 Eighth International Conference on Advanced Computing (ICoAC)*. IEEE, 72–76.

[2] Shashank Agarwal and Hong Yu. 2010. Biomedical negation scope detection with conditional random fields. *Journal of the American medical informatics association* 17, 6 (2010), 696–701.

[3] Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. 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 (LREC'10)*.

[4] C Leroy Baker. 1969. Double negatives. *Research on Language & Social Interaction* 1, 1 (1969), 16–40.

[5] Wendy W Chapman, Will Bridewell, Paul Hanbury, Gregory F Cooper, and Bruce G Buchanan. 2001. A simple algorithm for identifying negated findings and diseases in discharge summaries. *Journal of biomedical informatics* 34, 5 (2001), 301–310.

[6] W Timothy Coombs. 2015. The value of communication during a crisis: Insights from strategic communication research. *Business horizons* 58, 2 (2015), 141–148.

[7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).

[8] Allyson Ettinger. 2020. What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models. *Transactions of the Association for Computational Linguistics* 8 (2020), 34–48.

[9] Xianghua Fu, Jingying Yang, Jianqiang Li, Min Fang, and Huihui Wang. 2018. Lexicon-enhanced LSTM with attention for general sentiment analysis. *IEEE Access* 6 (2018), 71884–71891.

[10] Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. Allennlp: A deep semantic natural language processing platform. *arXiv preprint arXiv:1803.07640* (2018).

[11] Laurence R. Horn and Heinrich Wansing. 2020. Negation. In *The Stanford Encyclopedia of Philosophy* (Spring 2020 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University.

[12] Clayton Hutto and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media*, Vol. 8. 216–225.

[13] Nwakanma Ifeanyi, Oluigbo Ikenna, and Okpala Izunna. 2014. Text–To–Speech Synthesis (TTS). *International Journal of Research in Information Technology (IJRIT)* 2, 5 (2014), 154–163.

[14] Okpala Izunna, Rodriguez Guillermo, Romera, Zheng Weibing, Halse Shane, and Kropczynski Jess. 2022. Perception Analysis: Pro- and Anti- Vaccine Classification with NLP and Machine Learning. *Proceedings of the 55th Hawaii International Conference on System Sciences* (2022), 2981–2990.

[15] Yiru Jiao and Qing-Xing Qu. 2019. A proposal for Kansei knowledge extraction method based on natural language processing technology and online product reviews. *Computers in Industry* 108 (2019), 1–11.

[16] SHEN Jiaxuan. 2010. Division of negatives and noun/verb division in English and Chinese [J]. *Studies of the Chinese Language* 5 (2010).

[17] johnbumgarner. 2021. wordhoard. https://github.com/johnbumgarner/wordhoard [Online; accessed 30. Nov. 2021].

[18] Anna Jurek, Maurice D Mulvenna, and Yaxin Bi. 2015. Improved lexicon-based sentiment analysis for social media analytics. *Security Informatics* 4, 1 (2015), 1–13.

[19] Aditya Khandelwal and Suraj Sawant. 2019. Negbert: A transfer learning approach for negation detection and scope resolution. *arXiv preprint arXiv:1911.04211* (2019).

[20] Joki Kimberly. 2021. Double Negatives: 3 Rules You Must Know. https://www.grammarly.com/blog/3-things-you-must-know-about-double-negatives [Online; accessed 23. Jul. 2021].

[21] Olga Kovaleva, Saurabh Kulshreshtha, Anna Rogers, and Anna Rumshisky. 2021. BERT Busters: Outlier Dimensions that Disrupt Transformers. *arXiv preprint arXiv:2105.06990* (2021).

[22] Xu Liang. 2020. Treat Negation Stopwords Differently According to Your NLP Task. *Medium* (May 2020). https://towardsdatascience.com/treat-negation-stopwords-differently-according-to-your-nlp-task-e5a59ab7c91f

[23] Syed Ishfaq Manzoor, Jimmy Singla, et al. 2019. Fake news detection using machine learning approaches: A systematic review. In *2019 3rd International Conference on Trends in Electronics and Informatics (ICOEI)*. IEEE, 230–234.

[24] Andreas C Müller and Sarah Guido. 2016. *Introduction to machine learning with Python: a guide for data scientists.* " O'Reilly Media, Inc.".

[25] Pradeep G Mutalik, Aniruddha Deshpande, and Prakash M Nadkarni. 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 (2001), 598–609.

[26] Luis Alonso Ovalle and Elena Guerzoni. 2004. Double negatives, negative concord and metalinguistic negation. In *CLS 38.1: The main session. Proceedings from the main session of the 38th meeting of the Chicago Linguistic Society*. Citeseer, 15–31.

[27] Madhura Mandar Phadke and Satish R Devane. 2017. Multilingual Machine translation: An analytical study. In *2017 International Conference on Intelligent Computing and Control Systems (ICICCS)*. IEEE, 881–884.

[28] Ellen Quilty. 2019. University Library: AMA Referencing (Vancouver): Internet & Social Media. (2019).

[29] Vanesa del Río Zamora et al. 2015. Comparative Study of the Use of Double Negatives by Native English Speakers and Spanish Learners of English. (2015).

[30] Lior Rokach, Roni Romano, and Oded Maimon. 2008. Negation recognition in medical narrative reports. *Information Retrieval* 11, 6 (2008), 499–538.

[31] Antoine J-P Tixier, Matthew R Hallowell, Balaji Rajagopalan, and Dean Bowman. 2016. Automated content analysis for construction safety: A natural language processing system to extract precursors and outcomes from unstructured injury reports. *Automation in Construction* 62 (2016), 45–56.

[32] M Onat Topal, Anil Bas, and Imke van Heerden. 2021. Exploring transformers in natural language generation: Gpt, bert, and xlnet. *arXiv preprint arXiv:2102.08036* (2021).

[33] Okpala Izunna Udebuana, Ijioma Patricia Ngozi, and Emejulu Augustine Obiajulu. 2019. Analysis of Evaluated Sentiments; a Pseudo-Linguistic Approach and Online Acceptability Index for Decision-Making with Data: Nigerian Election in View. *Computing* 7, 2 (2019), 39–44.

[34] Sudha Verma, Sarah Vieweg, William Corvey, Leysia Palen, James Martin, Martha Palmer, Aaron Schram, and Kenneth Anderson. 2011. Natural language processing to the rescue? extracting" situational awareness" tweets during mass emergency. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 5.

[35] Jonathan J Webster and Chunyu Kit. 1992. Tokenization as the initial phase in NLP. In *COLING 1992 Volume 4: The 14th International Conference on Computational Linguistics*.

[36] Dominic Widdows. 2003. Orthogonal negation in vector spaces for modelling word-meanings and document retrieval. In *Proceedings of the 41st annual meeting of the association for computational linguistics*. 136–143.

[37] Sam Wiseman and Karl Stratos. 2019. Label-agnostic sequence labeling by copying nearest neighbors. *arXiv preprint arXiv:1906.04225* (2019).

[38] Liang-Chih Yu, Jin Wang, K Robert Lai, and Xuejie Zhang. 2017. Refining word embeddings using intensity scores for sentiment analysis. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 26, 3 (2017), 671–681.