The Effect of Negation on Sentiment Analysis and Retrieval Effectiveness

Lifeng Jia
Department of Computer Science
University of Illinois at Chicago
Chicago, IL 60607, USA

ljia@cs.uic.edu

Clement Yu
Department of Computer Science
University of Illinois at Chicago
Chicago, IL 60607, USA

yu@cs.uic.edu

Weiyi Meng
Department of Computer Science
SUNY at Binghamton
Binghamton, NY 13902, USA

meng@cs.binghamton.edu

ABSTRACT

We investigate the problem of determining the polarity of sentiments when one or more occurrences of a negation term such as "not" appear in a sentence. The concept of the scope of a negation term is introduced. By using a parse tree and typed dependencies generated by a parser and special rules proposed by us, we provide a procedure to identify the scope of each negation term. Experimental results show that the identification of the scope of negation improves both the accuracy of sentiment analysis and the retrieval effectiveness of opinion retrieval.

Categories and Subject Descriptors

H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval – *retrieval models, selection process*; I.2.7 [**Artificial Intelligence**]: Natural Language Processing – *text analysis*;

General Terms: Algorithms, Experimentation, Languages

Keywords: Scope of Negation, Candidate Scope of Negation, Sentiment Analysis, Retrieval Effectiveness, Opinion Retrieval

1. INTRODUCTION

In opinion retrieval, an opinionated document satisfies two conditions: it is relevant to the query and has an opinion about the query [10]. The TREC 2007 blog track [11] introduced a new "polarity classification" task. The task is to provide sentiment analysis on opinionated documents, i.e. it is to determine whether a given opinionated document carries positive, negative or mixed (both positive and negative) opinions. To determine the polarity of an opinionated document, we first classify the polarities of individual sentences and then aggregate the sentence level results into a document level polarity. The polarity of a sentence is very often recognized by certain sentimental words or phrases within it. However, their contextual polarities are dependent on the scope of each negation word or phrase preceding them, because their polarities might be flipped by negation words or phrases.

Existing research [12, 3, 4, 17, 6, 2, 1] has been conducted on determining the impact of negation words or phrases on the

sentimental polarity of a sentence. In [12], the scope of a negation word or phrase is assumed to be those words between that negation and the first punctuation mark following it. [3, 4] suggest that the scope of a negation term to be its next 5 words. In [17], the polarity of a sentimental term is flipped within the vicinity of negation, which implies that the scope of negation is several words to its right. Besides these heuristics in identifying the scope of a negation word or phrase, some research evaluated the impact of negation differently. [6] introduces the concept of contextual valence shifter, which consists of negation, intensifier and diminisher. Contextual valence shifters have an impact of flipping the polarity, increasing or decreasing the degree to which a sentimental term is positive or negative. [2] categorizes negations into function negations, such as "not", and contextual negation, such as "eliminate". Both kinds of negations can flip the polarity of sentimental terms. The same problem but in the medical/health domain was investigated in [1]. However the precision involving the negative word "not" is very low, at 63%. In this paper, we assume that sentimental terms are either individual word or multi-word phrases whose polarities have been pre-determined by methods such as [15, 16, 18] and negation terms are either individual negation words or negation phrases. We concentrate on the impact of negation terms on sentiment analysis. Negation terms are not restricted to "not". The most common negation words are: no, not (or its contraction n't), never, less, without, barely, hardly and rarely; the most common negation phrases are: no longer, no more, no way, no where, by no means, at no time, not ... anymore. We restrict our analysis to the set of negative terms given above in this paper.

The objectives of the paper are to determine the polarities of the parts of a sentence, which may be affected by each occurrence of a negation term and then utilize the polarity information to improve retrieval effectiveness. Our study has the following contributions. (a) We introduce the concept of the scope of a negation term and provide a methodology to determine it. To our knowledge, this study is the first one in which the scope of a negation term is defined and a non-trivial procedure is provided for its computation. (b) We study different methods of determining the polarity of the candidate scope acted on by each occurrence of a negation term in a sentence, where a candidate scope represents a logical unit of a sentence containing the scope. Experiments are performed to compare the effectiveness of these methods. (c) We incorporate our technique of determining polarity into an opinion retrieval system [18, 19] and compare it against other existing techniques. Experimental results show that our technique outperforms the other techniques in retrieval effectiveness.

2. PRELIMINARY DEFINITIONS

To identify the scope of a negation term, t, our strategy is to first compute a candidate scope, which is a minimal "logical unit" of the sentence containing the scope. Then, we prune those words that are within the candidate scope but not within the scope. Clearly, the candidate scope of t is a subset of the words appearing after t in the sentence. The "logical unit" is the set of descendant terminal (leaf) nodes of a non-terminal node in a parse tree of the sentence. To ensure that the candidate scope is minimal, we restrict the candidate scope not to extend to another independent clause of the sentence. We give the following computational procedure to approximate the candidate scope.

Procedure ComputedCandidateScope: Suppose a negation term t occurs in a sentence S. Obtain a parse tree for S, say using Stanford's parser [7]. Find the least common ancestor, LCA, of the node representing t and the node representing the word, say t', immediately after t. Then, all descendant leaf nodes of LCA starting from t' and extending to its right hand side form the candidate scope of t. Intuitively, it is the set of words whose polarities may be flipped by t. The logical unit of the sentence, which is the descendant leaf nodes of LCA with the exception of t and its preceding words, is the computed candidate scope.

In order to precisely locate the scope of a negation term in a sentence, we need to analyze the sentence syntactically. The parse tree and typed dependencies [8] of a sentence provide helpful assistance in syntax analysis. A parse tree is an ordered and rooted tree that represents the syntactic structure of the sentence according to some formal grammar. The definitions of typed dependency and five concrete typed dependencies, namely "conjunction", "copula" "open clausal complement", "direct object" and "indirect object" are given in [8, 9] and are omitted due to limitation of space. They will be utilized in determining the scope.

3. IDENTIFY THE SCOPE OF NEGATION

In this section, we discuss how to recognize the scope of a negation term from the candidate scope, utilizing the concepts defined in the last section.

3.1 Delimiters

The candidate scope of a negation term is not always its exact scope. Therefore, after obtaining the candidate scope, we need to identify its actual scope. In order to do it, the concepts of a "delimiter" and a "conditional word delimiter" are introduced. A delimiter d when encountered in the candidate scope CS of a negation term eliminates certain words, including d and the words after it from CS, whereas a conditional delimiter behaves as a delimiter, if certain conditions are satisfied.

Definition 3.1 Delimiter: A delimiter has the capability to eliminate some words from the candidate scope of a negation term. The set of words eliminated by the delimiter is given by either one of the two following rules:

(1)   All words including the delimiter and those after it are eliminated, if there is no conjunction, satisfying rule (2).

(2)   Only a portion of the succeeding words is eliminated by the delimiter. Let w be a delimiter. If w', a succeeding word of w forms a typed dependency of conjunction by either "AND" or "OR" with some word w" that precedes w and w" is not eliminated from the candidate scope, then

the words from w to w' (including w but excluding w') are eliminated.

Examples of delimiters are "when", "whenever", "whether", "because", "unless", "until", "since" and "hence".

Definition 3.2 Conditional Word Delimiter: A conditional word delimiter may or may not serve as a delimiter. It serves as a delimiter and eliminates a subset of words from the candidate scope if it satisfies some specific conditions. Examples of conditional word delimiters include "so", "as", "which", "who", "why", "where", "for", "like" and quotation marks. If a conditional delimiter serves as a delimiter, one of the above two rules (1) and (2), which is applicable to a delimiter, is applied to eliminate some words from the candidate scope. The types of conditions, which make a conditional delimiter a delimiter, include:(a) the part of speech of the conditional delimiter; (b) the location of the negation term relative to the conditional delimiter; and (c) the word leading to an adjective clause. Due to space limitations, only a small number of conditions are listed above.

3.2 Heuristic Rules for Scope Detection

Besides certain specific words, which are word delimiters and conditional word delimiters, we also propose rules involving sentimental verbs, sentimental adjectives and sentimental nouns such that the word immediately after one of these sentimental terms acts as a delimiter. Furthermore, a heuristic rule concerning double objects is proposed.

Sentimental Verb Rule: Whenever a negation term in a sentence negates a sentimental verb, the word w immediately after the verb serves as a delimiter.

Sentimental Adjective Rule: Whenever a sentimental adjective forms a "cop" or "xcomp" typed dependency with the closest preceding copula or verb, which is negated by a negation term, the term immediately after this adjective serves as a delimiter.

Sentimental Noun Rule: Whenever a sentiment noun acts as the object of a verb, which is negated by a negation term, the term immediately after this noun is a delimiter.

Double Object Rule: Whenever a negation term negates a verb taking double objects, only the direct object should be in the scope and the indirect object should be excluded.

3.3 Exceptions of Scope of Negation

For a sentence with a negation term, we have introduced various kinds of methods in identifying the scope of that negation term. However, sometimes a negation term in a sentence does not have any scope. In this section, we summarize several situations when a negation term does not have a scope.

Exception Situation 1: Whenever a negation term is a part of some special phrase without any negation sense, there is no scope for this negation term. Examples of these special phrases include "not only", "not just", "not to mention" and "no wonder".

Exception Situation 2: A negation term does not have a scope when it occurs in a negative rhetorical question. A negative rhetorical question is identified by the following heuristic. (1) It is a question; and (2) it has a negation term within the first three words of the question.

Exception Situation 3: A negation term does not have a scope when the sentence itself is a "restricted comparative sentence". Such a sentence is approximated by the pattern: modal word (such

as "can") immediately followed by a negation term, immediately followed by a copular verb (or get) or another tense of the verb and followed by a comparative word.

3.4 Scope Identification Procedure

After introducing various techniques to identify the scope of a negation term, we now present a procedure, which identifies the scope of each occurrence of a negation term. This procedure takes a sentence with one or more occurrences of negation terms as input and outputs their scopes within the sentence. If a sentiment term is within the scopes of i negation terms, then its polarity is flipped i times.

For each occurrence of a negation term t within a sentence S,

Case 1: If the occurrence of t in S satisfies Exception Situations 1, 2 or 3, then the entire candidate scope of t is discarded.

Case 2: If all three conditions in Case 1 fail, obtain the candidate scope, CS, of t according to the parse tree of the sentence S and then identify the scope from CS by processing according to the following cases.

Case 2.1: There are no word delimiters within CS and none of the sentimental verb, adjective and noun rules and the double object rule is satisfied. In this case, the scope is the candidate scope.

Case 2.2: When a word delimiter or a conditional word delimiter satisfying the conditions to serve as a delimiter is encountered, the scope is obtained from CS by applying delimiter rules (1) or (2) in Definition 3.1.

Case 2.3: When the conditions in any of the three sentimental rules are satisfied, the scope is obtained from CS by applying the delimiter rules (1) or (2) in Definition 3.1.

Case 2.4: When the condition in the double object rule is satisfied, the candidate scope is modified by discarding the indirect object.

It is possible that different parts of a candidate scope of a negation term satisfy the conditions of Cases 2.2, 2.3 or 2.4. For each satisfied part, the candidate scope is modified accordingly.

4. SENTIMENT ANALYSIS

4.1 Sentiment Analysis on Candidate Scope

To analyze the polarity of a candidate scope CS, if CS contains some sentimental terms, the contextual polarities of sentimental terms are first determined by taking into the consideration of their predetermined polarities and the scopes of negation terms (if exist) preceding them; CS is then classified to be positive or negative if all contextual polarities of sentimental terms are of the same polarity; otherwise, CS is classified to be mixed. If no sentimental terms occur in CS, the polarity of CS is neutral. Let this simple combination method be denoted by **SM**.

Besides **SM** above, we also employ decision tree to determine the polarity of CS. The decision tree has 8 independent features and requires training examples. These features indicate the syntactic roles of the sentimental terms within CS and are explained below. Each of these features can take on one of the four values: {positive, negative, mixed, neutral}. The 8 features are: (1)**SI:** the sentimental polarity of the subject of an independent clause. (2)**PI:** the sentimental polarity of the predicate of an independent clause. (3)**OI:** the sentimental polarity of the object of an independent clause. (4)**MI:** the sentimental polarity of the modifier of an independent clause. (5)**SD:** the sentimental polarity

of the subject of a dependent clause. (6)**PD:** the sentimental polarity of the predicate of a dependent clause. (7)**OD:** the sentimental polarity of the object of a dependent clause. (8)**MD:** the sentimental polarity of the modifier of a dependent clause. For a sentence, all feature values are assigned accordingly. In general, for each candidate scope, CS, a vector of 8 features is computed and fed into Quinlan's C4.5 decision tree program [13] to generate the classifier. Let the decision tree method be denoted by **DT**.

4.2 Retrieval Effectiveness on TREC Collection

The scope identification technique described in Section 3 is incorporated into an opinion retrieval system [18, 19] to improve the retrieval effectiveness of polarity classification of TREC documents for given queries. A brief description is as follows. Given a query, opinionated relevant documents are first retrieved by an opinion retrieval system [18, 19] and then opinionative sentences are classified into to be positively or negatively opinionative ones by a polarity classifier, which uses features consisting of positive and negative sentimental terms. If a feature is within the scopes of an odd number of negation terms, its negated feature is used instead. All features which are present in the remaining part of the sentence are not modified. The classifier used to classify each sentence is SVM-Light [5], which produces either a positive score or a negative score for each sentence. Each opinionative document is assigned a positive score and a negative score by summing its positive/negative scores of its opinionative sentences. If a document has a certain proportion of positive score to negative score, then it is classified to be positive, negative or mixed. Two ranked lists of opinionated documents are produced, one for the positively ranked documents, and the other for the negatively ranked documents.

5. EXPERIMENTS

To evaluate the effectiveness of our method on polarity determination, we conduct two sets of experiments. In each set of experiments, we compare our method with other methods, which handle the negation terms in different ways. The first set of experiments involves the accuracy of computing the polarity of a sentence. The second set of experiments involves the ranking of positively and negatively ranked opinionated documents retrieved from 3.2 million TREC documents with respect to 150 TREC queries.

5.1 Experimental Results

5.1.1 Sentiment Analysis on Candidate Scope

We now evaluate the accuracy of our method, **SCT**, in identifying the scope of negation by a dataset that consists of 1000 sentences. These sentences are randomly sampled from the review corpus crawled from Rateitall.com. Instead of computing the scope of negation, T, using our proposed method, three heuristics in identifying T are as follows: (a) T is within K words to the right of the negation term [3, 4, 17]. For (a), we test values of K = 3, 4 and 5 and K = 4 gives the best results. Thus, we report results for K = 4 only in this Section. This method (a) is denoted by **SC4**. (b) T is the set of words containing the first sentimental term to the right of the negation word. This method (b) is denoted by **SC1st**. (c) T is the set of all the words within CS and this method is denoted by **SCCS.** The polarity of the candidate scope can be determined automatically using either **SM** or **DT**. For example,

SCT+DT represents that the scope of negation is first identified by **SCT** and then the polarity of candidate scope is determined by **DT**.

Table 1: The Accuracies of Various Methods.

Methods	Accuracy	Methods	Accuracy
SCT+DT	**88.4%**	SC1st+DT	82.8%
SCT+SM	**85.8%**	SC1st+SM	79.8%
SC4+DT	83.6%	SCCS+DT	82.1%
SC4+SM	80.7%	SCCS+SM	79.2%

5.1.2 Retrieval Effectiveness on TREC Collection

In the second set of experiments, we rank the positively and negatively opinionative documents in the TREC blogosphere collection including all 150 queries released from 2006 to 2008. In the blog track of TREC 2008, the key measure to evaluate the retrieval effectiveness is the Mean Average Precision (MAP). Thus, we utilize the same measure here. Our method to rank these two sets of documents for each query has been described in Section 4.2. It is denoted by **SCT**, and is compared against other methods. Specifically, the methods to be compared against are listed below. (1) The method that is utilized in [19] is denoted by **SCBL**. It only flips the polarity of the closest sentimental term. (2) The scope of each negation term is within K word to the right of the negation term. Two methods with K = 4 and K = 5 are proposed in [3, 4, 17] and are denoted as **SC4** and **SC5**. (3) Two methods [2] have been proposed to determine the polarity of an expression within a sentence. An author in [2] suggested that we utilize the method denoted by **SCNegEx** [2], which achieved the best accuracy of sentiment analysis at the sentence level in the corpus of SemEval-07 [14]. We follow this suggestion. The gold standard provided by TREC is utilized. Table 2 shows the effectiveness of the various methods in ranking positive and negative documents.

Table 2: MAP scores of 5 methods on all TREC queries

	150 TREC Queries: 851-950 and 1001-1050			
	Positive	Improvement By SCT	Negative	Improvement By SCT
SCBL	0.1596	2.9%	0.0779	11.3%
SC4	0.1634	0.5%	0.0805	9.8%
SC5	0.1630	0.5%	0.0812	8.9%
SCNegEx	0.1487	10.4%	0.0823	7.4%
SCT	**0.1642**	-	**0.0884**	-

6. CONCLUSION

We study the impact of each occurrence of a negation term in a sentence on its polarity. We introduce the concept of scope of the negation term t, which is precisely the sequence of words after t and is affected by t. Techniques are provided to compute it. Two sets of experiments are performed to compare our method against other existing methods. Experimental results show that our method outperforms other methods in both the accuracy of sentiment analysis and the retrieval effectiveness of polarity classification in opinion retrieval.

7. ACKNOWLEDGEMENTS

8. REFERENCES

[1] WW. Chapman, W. Bridewell, P. Hanbury P, GF. Cooper and BG. Buchanan. A simple algorithm for identifying negated findings and diseases in discharge summaries. J Biomed Inform. 2001 Oct. 34(5):301-10.

[2] Yejin Choi and Claire Cardie. Learning with Compositional Semantics as Structural Inference for Subsentential Sentiment Analysis. In Proc. of EMNLP 2008.

[3] G. Grefenstette, Y. Qu, J. Shanahan and D. Evans. Coupling Niche Browsers and Affect Analysis for an Opinion Mining Application. In Proc. of RIAO 2004.

[4] Minqing Hu and Bing Liu. Mining and summarizing customer reviews. In Proc. of SIGKDD 2004.

[5] T Joachims. Making large-scale SVM learning practical. Advances in Kernel Methods:Support Vector Learning. 1999.

[6] Alistair Kennedy, Diana Inkpen. Sentiment Classification of Movie Reviews Using Contextual Valence Shifters. Computational Intelligence, Vol. 22, (2006), pp. 110-125.

[7] Dan Klein and Christopher D. Manning. Accurate Unlexicalized Parsing. In Proc. of ACL 2003, pp. 423-430.

[8] Marie-Catherine de Marneffe, Bill MacCartney and Christopher D. Manning. Generating Typed Dependency Parses from Phrase Structure Parses. In Proc. of LREC 2006.

[9] Marie-Catherine de Marneffe and Christopher D. Manning. Stanford typed dependencies manual, September 2008.

[10] I. Ounis, M. Rijke, C. Macdonald, G. Mishne, I. Soboroff. Overview of the TREC-2006 Blog Track. In TREC 2006.

[11] I. Ounis, C. Macdonald and I. Soboroff. Overview of the TREC-2007 Blog Track. In TREC 2007.

[12] Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. Thumbs up? Sentiment classification using machine learning techniques. In Proc. of EMNLP 2002.

[13] J. R. Quinlan. C4.5: Programs for Machine Learning. Morgan-Kaufman, 1993.

[14] Carlo Strapparava, Rada Mihalcea. Semeval-2007 task 14: Affective text. In Proc. of SemEval, 2007.

[15] M. Taboada, J. Grieve. Analyzing appraisal automatically. In AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications. 2004, pp 158-161.

[16] T. Wilson, J. Wiebe, P. Hoffmann. Recognizing contextual polarity in phrase-level sentiment analysis. In Proc. of HLT/EMNLP 2005.

[17] K. Yang. WIDIT in TREC 2008 Blog Track: Leveraging Multiple Sources of Opinion Evidence. In TREC 2008.

[18] W. Zhang, C. Yu and W. Meng. Opinion Retrieval from Blogs. In Proc. of CIKM 2007.

[19] W. Zhang, L. Jia, C. Yu and W. Meng. Improve the Effectiveness of the Opinion Retrieval and Opinion Polarity Classification. In Proc. of CIKM 2008.