

Sentiment Analysis of Conditional Sentences

Ramanathan Narayanan

Dept. of EECS

Northwestern University

ramanathan.an@gmail.com

Bing Liu*

Dept. of Computer Science

Univ. of Illinois at Chicago

liub@cs.uic.edu

Alok Choudhary

Dept. of EECS

Northwestern University

alokchoudhary01@gmail.com

Abstract

This paper studies sentiment analysis of conditional sentences. The aim is to determine whether opinions expressed on different topics in a conditional sentence are positive, negative or neutral. Conditional sentences are one of the commonly used language constructs in text. In a typical document, there are around 8% of such sentences. Due to the condition clause, sentiments expressed in a conditional sentence can be hard to determine. For example, in the sentence, *if your Nokia phone is not good, buy this great Samsung phone*, the author is positive about “*Samsung phone*” but does not express an opinion on “*Nokia phone*” (although the owner of the “*Nokia phone*” may be negative about it). However, if the sentence does not have “*if*”, the first clause is clearly negative. Although “*if*” commonly signifies a conditional sentence, there are many other words and constructs that can express conditions. This paper first presents a linguistic analysis of such sentences, and then builds some supervised learning models to determine if sentiments expressed on different topics in a conditional sentence are positive, negative or neutral. Experimental results on conditional sentences from 5 diverse domains are given to demonstrate the effectiveness of the proposed approach.

1 Introduction

Sentiment analysis (also called opinion mining) has been an active research area in recent years. There are many research directions, e.g., sentiment classification (classifying an opinion document as positive or negative) (e.g., Pang, Lee and Vaithyanathan, 2002; Turney, 2002), subjectivity classification (determining whether a sentence is subjective or objective, and its associated opinion) (Wiebe and Wilson, 2002; Yu and Hatzivassiloglou, 2003; Wilson et al, 2004; Kim and

Hovy, 2004; Riloff and Wiebe, 2005), feature/topic-based sentiment analysis (assigning positive or negative sentiments to topics or product features) (Hu and Liu 2004; Popescu and Etzioni, 2005; Carenini et al., 2005; Ku et al., 2006; Kobayashi, Inui and Matsumoto, 2007; Titov and McDonald, 2008). Formal definitions of different aspects of the sentiment analysis problem and discussions of major research directions and algorithms can be found in (Liu, 2006; Liu, 2009). A comprehensive survey of the field can be found in (Pang and Lee, 2008).

Our work is in the area of topic/feature-based sentiment analysis or opinion mining (Hu and Liu, 2004). The existing research focuses on solving the general problem. However, we argue that it is unlikely to have a one-technique-fit-all solution because different types of sentences express sentiments/opinions in different ways. A divide-and-conquer approach is needed, e.g., focused studies on different types of sentences. This paper focuses on one type of sentences, i.e., conditional sentences, which have some unique characteristics that make it hard to determine the orientation of sentiments on topics/features in such sentences. By *sentiment orientation*, we mean positive, negative or neutral opinions. By *topic*, we mean the target on which an opinion has been expressed. In the product domain, a topic is usually a *product feature* (i.e., a component or attribute). For example, in the sentence, *I do not like the sound quality, but love the design of this MP3 player*, the product features (topics) are “*sound quality*” and “*design*” of the MP3 player as opinions have been expressed on them. The sentiment is positive on “*design*” but negative on “*sound quality*”.

Conditional sentences are sentences that describe implications or hypothetical situations and their consequences. In the English language, a variety of conditional connectives can be used to form these sentences. A conditional sentence contains two clauses: the condition clause and

* This work was done when Bing Liu was on sabbatical leave at Northwestern University.

the consequent clause, that are dependent on each other. Their relationship has significant implications on whether the sentence describes an opinion. One simple observation is that sentiment words (also known as opinion words) (e.g., *great, beautiful, bad*) alone cannot distinguish an opinion sentence from a non-opinion one. A conditional sentence may contain many sentiment words or phrases, but express no opinion.

Example 1: *If someone makes a beautiful and reliable car, I will buy it* expresses no sentiment towards any particular car, although “*beautiful*” and “*reliable*” are positive sentiment words.

This, however, does not mean that a conditional sentence cannot express opinions/sentiments.

Example 2: *If your Nokia phone is not good, buy this great Samsung phone* is positive about the “*Samsung phone*” but does not express an opinion on the “*Nokia phone*” (although the owner of the “*Nokia phone*” may be negative about it). Clearly, if the sentence does not have “*if*”, the first clause is negative. Hence, a method for determining sentiments in normal sentences will not work for conditional sentences. The examples below further illustrate the point.

In many cases, both the condition and consequent together determine the opinion.

Example 3: *If you are looking for a phone with good voice quality, don’t buy this Nokia phone* is negative about the “*voice quality*” of the “*Nokia phone*”, although there is a positive sentiment word “*good*” in the conditional clause modifying “*voice quality*”. However, in the following example, the opinion is just the opposite.

Example 4: *If you want a phone with good voice quality, buy this Nokia phone* is positive about the “*voice quality*” of the “*Nokia phone*”.

As we can see, sentiment analysis of conditional sentences is a challenging problem.

One may ask whether there is a large percentage of conditional sentences to warrant a focused study. Indeed, there is a fairly large proportion of such sentences in evaluative text. They can have a major impact on the sentiment analysis accuracy. Table 1 shows the percentage of conditional sentences (sentences containing the words *if, unless, assuming*, etc) and also the total

number of sentences from which we computed the percentage in several user-forums. The figures definitely suggest that there is considerable benefit to be gained by developing techniques that can analyze conditional sentences.

To the best of our knowledge, there is no focused study on conditional sentences. This paper makes such an attempt. Specifically, we determine whether a conditional sentence (which is also called a *conditional* in the linguistic literature) expresses positive, negative or neutral opinions on some topics/features. Since our focus is on studying how conditions and consequents affect sentiments, we assume that *topics* are given, which are *product attributes* since our data sets are user comments on different products.

Our study is conducted from two perspectives. We start with the linguistic angle to gain a good understanding of existing work on different types of conditionals. As conditionals can be expressed with other words or phrases than *if*, we will study how they behave compared to *if*. We will also show that the distribution of these conditionals based on our data sets.

With the linguistic knowledge, we perform a computational study using machine learning. A set of features for learning is designed to capture the essential determining information. Note that the features here are data attributes used in learning rather than product attributes or features. Three classification strategies are designed to study how to best perform the classification task due to the complex situation of two clauses and their interactions in conditional sentences. These three classification strategies are *clause-based, consequent-based* and *whole-sentence-based*. Clause-based classification classifies each clause separately and then combines their results. Consequent-based classification only uses consequents for classification as it is observed that in conditional sentences, it is often the consequents that decide the opinion. Whole-sentence-based classification treats the entire sentence as a whole in classification. Experimental results on conditional sentences from diverse domains demonstrate the effectiveness of these classification models. The results indicate that the whole-sentence-based classifier performs the best.

Since this paper only studies conditional sentences, a natural question is whether the proposed technique can be easily integrated into an overall sentiment analysis or opinion mining system. The answer is yes because a large proportion of conditional sentences can be detected using conditional connectives. Keyword search is

Table 1: Percent of conditional sentences

Source	% of cond. (total #. of sent.)
Cellphone	8.6 (47711)
Automobile	5.0 (8113)
LCD TV	9.92 (258078)
Audio Systems	8.1 (5702)
Medicine	8.29 (160259)

thus sufficient to identify such sentences for special handling using the proposed approach. There are, however, some subtle conditionals which do not use normal conditional connectives and will need an additional module to identify them, but such sentences are very rare as Table 2 indicates.

2 The Problem Statement

The paper follows the *feature-based sentiment analysis* model in (Hu and Liu 2004; Popescu and Etzioni, 2005). We are particularly interested in sentiments on products and services, which are called objects or entities. Each object is described by its *parts* and *attributes*, which are collectively called *features* in (Hu and Liu, 2004; Liu, 2006). For example, in the sentence, *If this camera has great picture quality, I will buy it*, “*picture quality*” is a feature of the camera. For formal definitions of objects and features, please refer to (Liu, 2006; Liu, 2009). In this paper, we use the term *topic* to mean *feature* as the feature here can confuse with the feature used in machine learning. The term *topic* has also been used by some researchers (e.g., Kim and Hovy, 2004; Stoyanov and Cardie, 2008).

Our objective is to predict the sentiment orientation (positive, negative or neutral) on each topic that has been commented on in a sentence.

The problem of automatically identifying features or topics being spoken about in a sentence has been studied in (Hu and Liu, 2004; Popescu and Etzioni, 2005; Stoyanov and Cardie, 2008). In this work, we do not attempt to identify such topics automatically. Instead, we assume that they are given because our objective is to study how the interaction of the condition and consequent clauses affects sentiments. For this purpose, we manually identify all the topics.

3 Conditional Sentences

This section presents the linguistic perspective of conditional sentences.

Table 2: Percentage of sentences with some main conditional connectives

Conditional Connective	% of sentences
If	6.42
Unless	0.32
Even if	0.17
Until	0.10
As (so) long as	0.09
Assuming/supposing	0.04
In case	0.04
Only if	0.03

3.1 Conditional Connectives

A large majority of conditional sentences are introduced by the subordinating conjunction *If*. However, there are also many other conditional connectives, e.g., *even if*, *unless*, *in case*, *assuming/supposing*, *as long as*, etc. Table 2 shows the distribution of conditional sentences with various connectives in our data. Detailed linguistic discussions of them are beyond the scope of this paper. Interested readers, please refer to (Declerck and Reed, 2001). Below, we briefly discuss some important ones and their interpretations.

If: This is the most commonly used conditional connective. In addition to its own usage, it can also be used to replace other conditional connectives, except some semantically richer connectives (Declerck and Reed, 2001). Most (but not all) conditional sentences can be logically expressed in the form ‘If P then Q ’, where P is the condition clause and Q is the consequent clause. For practical purposes, we can automatically segment the condition and consequent clauses using simple rules generated by observing grammatical and linguistic patterns.

Unless: Most conditional sentences containing *unless* can be replaced with equivalent sentences with an *if* and a *not*. For example, the sentence *Unless you need clarity, buy the cheaper model* can be expressed with *If you don’t need clarity, buy the cheaper model*.

Even if: Linguistic theories claim that *even if* is a special case of a conditional which may not always imply an if-then relationship (Gauker 2005). However, in our datasets, we have observed that the usage of *even if* almost always translates into a conditional. Replacing *even if* by *if* will yield a sentence that is semantically similar enough for the purpose of sentiment analysis.

Only if, provided/providing that, on condition that: Conditionals involving these phrases typically express a necessary condition, e.g., *I will buy this camera only if they can reduce the price*. In such sentences, *only* usually does not affect whether the sentence is opinionated or not.

In case: Conditional sentences containing *in case* usually describe a precaution (*I will close the window in case it rains*), prevention (*I wore sunglasses in case I was recognized*), or a relevance conditional (*In case you need a car, you can rent one*). Identifying the conditional and consequent clauses is not straightforward in many cases. Further, in these instances, replacing *in case* with *if* may not convey the intended meaning of the conditional. We have ignored these cases in

our analysis as we believe that they need a separate study, and also such sentences are rare.

As (so) long as: Sentences with these connectives behave similarly to *if* and can usually be replaced with *if*.

Assuming/Supposing: These are a category of conditionals that behave quite differently. The participles *supposing* and *assuming* create conditional sentences where the conditional clause and the consequent clause can be syntactically independent. It is quite difficult to distinguish those conditional sentences which contain an explicit consequent clause and fit within our analysis framework. In our data, most of such sentences have no consequent, thus representing assumptions rather than opinions. We omit these sentences in our study (they are also rare).

3.2 Types of Conditionals

There are extensive studies of conditional sentences (also known as *conditionals*) in linguistics. Various theories have led to a number of classification systems. Popular types of conditionals include *actualization conditionals*, *inferential conditionals*, *implicative conditionals*, etc (Declerck and Reed, 2001). However, these classifications are mainly based on semantic meanings which are difficult to recognize by a computer program. To build classification models, we instead exploit *canonical tense patterns* of conditionals, which are often used in pedagogic grammar books. They are defined based on tense and are associated with general meanings. However, as described in (Declerck and Reed, 2001), their meanings are much more complex and numerous than their associated general meanings. However, the advantage of this classification is that different types can be detected easily because they depend on tense which can be produced by a part-of-speech tagger. As we will see in Section 5, canonical tense patterns help sentiment classification significantly. Below, we introduce the four canonical tense patterns.

Zero Conditional: This conditional form is used to describe universal statements like facts, rules and certainties. In a zero conditional, both the condition and consequent clauses are in the simple present tense. An example of such sentences is: *If you heat water, it boils.*

First Conditional: Conditional sentences of this type are also called potential or indicative conditionals. They are used to express a hypothetical situation that is probably true, but the truth of which is unverified. In the first condi-

tional, the condition is in the simple present tense, and the consequent can be either in past tense or present tense, usually with a modal auxiliary verb preceding the main verb, e.g., *If the acceleration is good, I will buy it.*

Second Conditional: This is usually used to describe less probable situations, for stating preferences and imaginary events. The condition clause of a second conditional sentence is in the past subjunctive (past tense), and the consequent clause contains a conditional verb modifier (like *would*, *should*, *might*), in addition to the main verb, e.g., *If the cell phone was robust, I would consider buying it.*

Third conditional: This is usually used to describe contrary-to-fact (impossible) past events. The past perfect tense is used in the condition clause, and the consequent clause is in the present perfect tense, e.g., *If I had bought the a767, I would have hated it.*

Based on the above definitions, we have developed approximate part-of-speech (POS) tags¹ for the condition and the consequent of each pattern (Table 3), which do not cover all sentences, but overall they cover a majority of the sentences. For those not covered cases, the problem is mainly due to incomplete sentences and wrong grammars, which are typical for informal writings in forum postings and blogs. For example, the sentence, *Great car if you need powerful acceleration*, does not fall into any category, but it actually means *It is a great car if you need powerful acceleration*, which is a zero conditional. To handle such sentences, we designed a set of rules to assign them some default types:

- If condition contains VB/VBP/VBZ → 0 conditional
- If consequent contains VB/VBP/VBS → 0 conditional
- If condition contains VBG → 1st conditional
- If condition contains VBD → 2nd conditional
- If conditional contains VBN → 3rd conditional.

Table 3: Tenses for identifying conditional types

Type	Linguistic Rule	Condition POS tags	Consequent POS tags
0	If + simple present → simple present	VB/VBP/VBZ	VB/VBP/ VBZ
1	If + simple present → will + bare infinitive	VB/VBP/VBZ /VBG	MD + VB
2	If + past tense → would + infinitive	VBD	MD + VB
3	If + past perfect → present perfect	VBD+VBN	MD + VBD

¹ The list of Part-Of-Speech (POS) tags can be found at: http://www.ling.upenn.edu/courses/Fall_2003/ling001/penn_treebank_pos.html

By using these rules, we can increase the sentence coverage from 73% to 95%.

4 Sentiment Analysis of Conditionals

We now describe our computational study. We take a machine learning approach to predict sentiment orientations. Below, we first describe features used and then classification strategies.

4.1 Feature construction

I. *Sentiment words/phrases and their locations*: Sentiment words are words used to express positive or negative opinions, which are instrumental for sentiment classification for obvious reasons. We obtained a list of over 6500 sentiment words gathered from various sources. The bulk of it is from <http://www.cs.pitt.edu/mpqa>. We also added some of our own. Our list is mainly from the work in (Hu and Liu, 2004; Ding, Liu and Yu, 2008). In addition to words, there are phrases that describe opinions. We have identified a set of such phrases. Although obtaining these phrases was time-consuming, it was only a one-time effort. We will make this list available as a community resource. It is possible that there is a better automated method for finding such phrases, such as the methods in (Kanayama and Nasukawa, 2006; Breck, Choi and Cardie, 2007). However, automatically generating sentiment phrases has not been the focus of this work as our objective is to study how the two clauses interact to determine opinions given the sentiment words and phrases are known. Our list of phrases is by no means complete and we will continue to expand it in the future.

For each sentence, we also identify whether it contains sentiment words/phrases in its condition or consequent clause. It was observed that the presence of a sentiment word/phrase in the consequent clause has more effect on the sentiment of a sentence.

II. *POS tags of sentiment words*: Sentiment words may be used in several contexts, not all of which may correspond to an opinion. For example, *I trust Motorola* and *He has a trust fund* both contain the word *trust*. But only the former contains an opinion. In such cases, the POS tags can provide useful information.

III. *Words indicating no opinion*: Similar to how sentiment words are related to opinions, there are also a number of words which imply the opposite. Words like *wondering*, *thinking*, *de-*

bating are used when the user is posing a question or expressing doubts. Thus such phrases usually do not contribute an opinion, especially if they are in the vicinity of the *if* connective. We search a window of 3 words on either side of *if* to determine if there is any such word. We have compiled a list of these words as well and use it in our experiments.

IV. *Tense patterns*: These are the canonical tense patterns in Section 3.2. They are used to generate a set of features. We identify the first verb in both the condition and consequent clauses by searching for the relevant POS tags in Table 3. We also search for the words preceding the main verb to find modal auxiliary verbs, which are also used as features.

V. *Special characters*: The presence or absence of '?' and '!'.

VI. *Conditional connectives*: The conditional connective used in the sentence (*if*, *even if*, *unless*, *only if*, etc) is also taken as a feature.

VII. *Length of condition and consequent clauses*: Using simple linguistic and punctuation rules, we automatically segment a sentence into condition and consequent clauses. The numbers of words in the condition and consequent clauses are then used as features. We observed that when the condition clause is short, it usually has no impact on whether the sentence expresses an opinion.

VIII. *Negation words*: The use of negation words like *not*, *don't*, *never*, etc, often alter the sentiment orientation of a sentence. For example, the addition of *not* before a sentiment word can change the orientation of a sentence from positive to negative. We consider a window of 3-6 words before an opinion word, and search for these kinds of words.

The following two features are singled out for easy reference later. They are only used in one classification strategy. The first feature is an indicator, and the second feature has a parameter (which will be evaluated separately).

(1). *Topic location*: This feature indicates whether the topic is in the conditional clause or the consequent clause.

(2). *Opinion weight*: This feature considers only sentiment words in the vicinity of the topic, since they are more likely to influence the opinion on the topic. A window size is used to control what we mean by vicinity. The following formula is used to assign a weight to each sentiment word, which is inversely proportional to the distance (D_{op}) of the sentiment word to the topic mention. Sentiment

value is +1 for a positive word and -1 for a negative word. *Sentwords* are the set of known sentiment words and phrases.

$$weight = \sum_{op} \frac{\pm 1}{D_{op}}, \forall op \in \{sentwords\}$$

4.2 Classification Strategies

Since we are interested in topic-based sentiment analysis, how to perform classification becomes an interesting issue. Due to the two clauses, it may not be sufficient to classify the whole sentence as positive or negative as in the same sentence, some topics may be positive and some may be negative. We propose three strategies.

Clause-based classification: Since there are two clauses in a conditional sentence, in this case we build two classifiers, one for the condition and one for the consequent.

Condition classifier: This method classifies the condition clause as expressing positive, negative or neutral opinion.

Training data: Each training sentence is represented as a feature vector. Its class is positive, negative or neutral depending on whether the conditional clause is positive, negative or neutral while considering both clauses.

Testing: For each test sentence, the resulting classifier predicts the opinion of the condition clause.

Topic class prediction: To predict the opinion on a topic, if the topic is in the condition clause, it takes the predicted class of the clause.

Consequent classifier: This classifier classifies the consequent clause as expressing positive, negative or neutral opinion.

Training data: Each training sentence is represented as a feature vector. Its class is positive, negative or neutral depending on whether the consequent clause is positive, negative or neutral while considering both clauses.

Testing: For each test sentence, the resulting classifier predicts the opinion of the consequent clause.

Topic class prediction: To predict the opinion on a topic, if the topic is in the consequent clause, it takes the predicted class of the clause.

The combination of these two classifiers is called the *clause-based classifier*. It works as follows: If a topic is in the conditional clause, the condition classifier is used, and if a topic is in the consequent clause, the consequent classifier is used.

Consequent-based classification: It is observed that in most cases, the condition clause contains no opinion whereas the consequent clause reflects the sentiment of the entire sentence. Thus, this method uses (in a different way) only the above consequent classifier. If it classifies the consequent of a testing conditional sentence as positive, all the topics in the whole sentence are assigned the positive orientation, and likewise for negative and neutral.

Whole-sentence-based classification: In this case, a single classifier is built to predict the opinion on each topic in a sentence.

Training data: In addition to the normal features, the two features (1) and (2) in Section 4.1 are used for this classifier. If a sentence contains multiple topics, multiple training instances of the same sentence are created in the training data. Each instance represents one specific topic. The class of the instance depends on whether the opinion on the topic is positive, negative or neutral.

Testing: For each topic in each test sentence, the resulting classifier predicts its opinion.

Topic class prediction: This is not needed as the prediction has been done in testing.

5 Results and Discussions

5.1 Data sets

Our data consists of conditional sentences from 5 different user forums: Cellphone, Automobile, LCD TV, Audio systems and Medicine. We obtained user postings from these forums and extracted the conditional sentences. We then manually annotated 1378 sentences from this corpus. We also annotated the conditional and consequent clauses and identified the topics (or product features) being commented upon, and their sentiment orientations. In our annotation, we observed that sentences with no sentiment words or phrases almost never express opinions, i.e., only around 3% of them express opinions. There are around 26% sentences containing no sentiment words or phrases in our data. To make the problem challenging, we restrict our attention to only those sentences that contain at least one sentiment word or phrase. We have annotated topics from around 900 such sentences. Table 4 shows the class distributions of this data. At the clause level (topics are not considered), we observe that conditional clauses contain few opinions. At the topic-level, 43.5% of the topics have positive opinions, 26.4% of the topics have negative opinions, and the rest have no opinions.

Table 4: Distribution of classes

	Positive	Negative	Neutral
Condition	6.9%	6.7%	86.4%
Consequent	49.3%	16.5%	34%
Topic-level	43.5%	26.4%	29.9%

For the annotation of data, we assume that topics are known. One student annotated the topics first. Then two students annotated the sentiments on the topics. If a student found that a topic annotation is wrong, he will let us know. Some mistakes and missing topics were found but there were mainly due to oversights rather than disagreements. The agreement on sentiment annotations were computed using the Kappa score. We achieved the Kappa score of 0.63, which indicates strong agreements. The conflicting cases were then solved through discussion to reach consensus. We did not find anything that the annotators absolutely disagree with each other.

5.2 Experimental results

We now present the results for different combinations of features and classification strategies. For model building, we used Support Vector Machines (SVM), and the LIBSVM implementation (Chang and Lin, 2001) with a Gaussian kernel, which produces the best results. All the results are obtained via 10-fold cross validation.

Two-class classification: We first discuss the results for a simpler version of the problem that involves only sentences with positive or negative orientations on some topics (at least one of the clauses must have a positive/negative opinion on a topic). Neutral sentences are not used (~28% of the total). The results of all three classifiers are given in Table 5. The feature sets have been described in Section 4.1. For all the experiments below, features (1) and (2) are only used by the whole-sentence-based classifier, but not used by the other two classifiers for obvious reasons.

{I+II}: This setting uses sentiment words and phrases, their positions and POS tags as features (we used Brill’s POS tagger). This can be seen as the baseline. We observe that both the consequent-based and whole-sentence-based classifiers perform dramatically better than the clause-based classifier. The consequent-based classifier and the whole-sentence-based classifier perform similarly (with the latter being slightly better). The precision, recall, and F-score are computed as the average of the two classes.

{I+II+III}: In this setting, the list of special non-sentiment related words is added to the fea-

ture set. All three classifiers improve slightly.

{I+II+III+IV}: This setting includes all the canonical tense based features. We see marked improvements for the consequent-based and whole-sentence-based classifiers both in term of accuracy and F-score, which are statistically significant compared to those of {I+II+III} at the 95% confidence level based on paired t-test.

All: When all the features are used, the results of all the classifiers improve further.

Two main observations worth mentioning:

1. Both the consequent-based and whole-sentence-based classifiers outperform the clause-based classifier dramatically. This confirms our observation that the consequent usually plays the key role in determining the sentiment of the sentence. This is further reinforced by the fact that the consequent-based classifier actually performs similarly to the whole-sentence-based classifier. The condition clause seems to give no help.
2. The second observation is that the linguistic knowledge of canonical tense patterns helps significantly. This shows that the linguistic knowledge is very useful.

We also noticed that many misclassifications are caused by grammatical errors, use of slang phrases and improper punctuations, which are typical of postings on the Web. Due to language irregularities (e.g., wrong grammar, missing punctuations, sarcasm, exclamations), the POS tagger makes many mistakes as well causing some errors in the tense based features.

Three-class classification: We now move to the more difficult and realistic case of three classes: positive, negative and neutral (no-opinion). Table 6 shows the results. The trend is similar except that the whole-sentence-based classifier now performs markedly better than the consequent-based classifier. We believe that this is because the neutral class needs information from both the condition and consequent clauses. This is evident from the fact that there is little or no improvement after {I+II} for the consequent-based classifier. We also observe that the accuracies and F-scores for the three-class classification are lower than those for the two-class classification. This is understandable due to the difficulty of determining whether a sentence has opinion or not. Again, statistical test shows that the canonical tense-based features help significantly.

As mentioned in Section 4.1, the whole-sentence-based classifier only considers those sentiment words in the vicinity of the topic under

Table 5: Two-class classification – positive and negative

	<i>Clause-based classifier</i>				<i>Consequent-based classifier</i>				<i>Whole-sentence-based classifier</i>			
	Acc.	Prec.	Rec.	F	Acc.	Prec.	Rec.	F	Acc.	Prec.	Rec.	F
I+II (senti. words+POS)	39.9	42.8	34.0	37.9	69.1	72.9	67.1	69.8	68.9	73.7	68.13	70.8
I+II+III (+ non-senti. words)	41.5	44.9	37.1	40.6	69.3	73.9	66.3	69.9	69.2	73.7	63.5	71.0
I+II+III+IV (+ tenses)	42.7	45.2	38.5	41.6	72.7	76.4	72.0	74.1	71.1	77.9	72.2	74.9
All	43.2	46.1	38.9	42.2	73.3	77.0	72.7	74.8	72.3	77.8	73.6	75.6

Table 6: Three-class classification – positive, negative and neutral (no opinion)

	<i>Clause-based classifier</i>				<i>Consequent-based classifier</i>				<i>Whole-sentence-based classifier</i>			
	Acc.	Prec.	Rec.	F	Acc.	Prec.	Rec.	F	Acc.	Prec.	Rec.	F
I+II (senti. words+POS)	45.2	41.3	35.1	37.9	54.6	57.7	52.9	55.2	59.1	58.1	56.4	57.2
I+II+III (+ non-senti. words)	46.9	42.8	37.8	40.1	55.3	60.0	51.3	55.3	61.4	60.1	60.8	60.4
I+II+III+IV (+ tenses)	50.3	48.7	40.9	44.5	57.3	64.0	50.0	56.1	64.6	63.3	63.9	63.6
All	53.3	49.8	44.1	46.8	58.7	64.5	50.1	56.4	67.8	66.9	65.1	66.0

Table 7: Accuracy of the whole-sentence-based classifier with varying window sizes (n)

Window size	1	2	3	4	5	6	7	8	9	10
Accuracy	66.1	62.6	64.1	64.8	65.3	65.7	66.3	67.3	66.9	66.8

investigation. For this, we search a window of n words on either side of the topic mention. To study the effect of varying n , we performed an experiment with various values of the window size and measured the overall accuracy for each case. Table 7 shows how the accuracy changes as we increase the window size. We found that a window size of 6-10 yielded good accuracies. This is because lower values of n lead to loss of information regarding sentiment words as some sentiment words could be far from the topic. We finally used 8, which gave the best results.

We also investigated ways of using the negation word in the sentence to correctly predict the sentiment. One method is to use the negation word as a feature, as described in Section 4.1. Another technique is to reverse the orientation of the prediction for those sentences which contain negation words. We found that the former technique yielded better results. The results reported so far are based on the former approach.

6 Related Work

There are several research directions in sentiment analysis (or opinion mining). One of the main directions is sentiment classification, which classifies the whole opinion document (e.g., a product review) as positive or negative (e.g., Pang et al, 2002; Turney, 2002; Dave et al, 2003; Ng et al. 2006; McDonald et al, 2007). It is clearly dif-

ferent from our work as we are interested in conditional sentences.

Another important direction is classifying sentences as subjective or objective, and classifying subjective sentences or clauses as positive or negative (Wiebe et al, 1999; Wiebe and Wilson, 2002, Yu and Hatzivassiloglou, 2003; Wilson et al, 2004; Kim and Hovy, 2004; Riloff and Wiebe, 2005; Gamon et al 2005; McDonald et al, 2007). Although these works deal with sentences, they aim to solve the general problem. This paper argues that there is unlikely a one-technique-fit-all solution, and advocates dealing with specific types of sentences differently by exploiting their unique characteristics. Conditional sentences are the focus of this paper. To the best of our knowledge, there is no focused study on them.

Several researchers also studied feature/topic-based sentiment analysis (e.g., Hu and Liu, 2004; Popescu and Etzioni, 2005; Ku et al, 2006; Carenini et al, 2006; Mei et al, 2007; Ding, Liu and Yu, 2008; Titov and R. McDonald, 2008; Stoyanov and Cardie, 2008; Lu and Zhai, 2008). Their objective is to extract topics or product features in sentences and determine whether the sentiments expressed on them are positive or negative. Again, no focused study has been made to handle conditional sentences. Effectively handling of conditional sentences can help their effort significantly.

In this work, we used many sentiment words and phrases. These words and phrases are usually compiled using different approaches (Hatzivassiloglou and McKeown, 1997; Kaji and Kitsuregawa, 2006; Kanayama and Nasukawa, 2006; Esuli and Sebastiani, 2006; Breck et al, 2007; Ding, Liu and Yu. 2008; Qiu et al, 2009). There are several existing lists produced by researchers. We used the one from the MPQA corpus (<http://www.cs.pitt.edu/mpqa>) with added phrases of our own from (Ding, Liu and Yu. 2008). In our work, we also assume that the topics are known. (Hu and Liu, 2004; Popescu and Etzioni, 2005; Kobayashi, Inui and Matsumoto, 2007; Stoyanov and Cardie, 2008) have studied topic/feature extraction.

One existing focused study is on comparative and superlative sentences (Jindal and Liu, 2006; Bos and Nissim, 2006; Fiszman et al, 2007; Ganapathibhotla and Liu, 2008). Their work identifies comparative sentences, extracts comparative relations in the sentences and analyzes comparative opinions (Ganapathibhotla and Liu, 2008). An example comparative sentence is “*Honda looks better than Toyota*”. As we can see, comparative sentences are entirely different from conditional sentences. Thus, their methods cannot be directly applied to conditional sentences.

7 Conclusion

To perform sentiment analysis accurately, we argue that a divide-and-conquer approach is needed, i.e., focused study on each type of sentences. It is unlikely that there is a one-size-fit-all solution. This paper studied one type, i.e., conditional sentences, which have some unique characteristics that need special handling. Our study was carried out from both the linguistic and computational perspectives. In the linguistic study, we focused on canonical tense patterns, which have been showed useful in classification. In the computational study, we built SVM models to automatically predict whether opinions on topics are positive, negative or neutral. Experimental results have shown the effectiveness of the models.

In our future work, we will further improve the classification accuracy and study related problems, e.g., identifying topics/features. Although there are some special conditional sentences that do not use easily recognizable conditional connectives and identifying them are useful, such sentences are very rare and spending time and effort on them may not be cost-effective at the moment.

Acknowledgements

This work was supported in part by DOE SCIDAC-2: Scientific Data Management Center for Enabling Technologies (CET) grant DE-FC02-07ER25808, DOE FASTOS award number DE-FG02-08ER25848, NSF HECURA CCF-0621443, NSF SDCI OCI-0724599, and NSF ST-HEC CCF-0444405.

References

- J. Bos, and M. Nissim. 2006. An Empirical Approach to the Interpretation of Superlatives. EMNLP-2006.
- E. Breck, Y. Choi, and C. Cardie. 2007. Identifying expressions of opinion in context, IJCAI-2007.
- C.-C. Chang and C.-J. Lin. 2001. LIBSVM: a library for support vector machines. <http://www.csie.ntu.edu.tw/~cjlin/libsvm>
- G. Carenini, R. Ng, and A. Pauls. 2006. Interactive Multimedia Summaries of Evaluative Text. IUI-2006.
- C. Gauker. 2005. Conditionals in Context. MIT Press.
- D. Dave, A. Lawrence, and D. Pennock. 2003. Mining the Peanut Gallery: Opinion Extraction and Semantic Classification of Product Reviews. WWW-2003.
- R. Declerck, and S. Reed. 2001. Conditionals: A Comprehensive Empirical Analysis. Berlin: Mouton de Gruyter.
- X. Ding, B. Liu, and P. S. Yu. 2008. A holistic lexicon-based approach to opinion mining. WSDM-2008.
- A. Esuli, and F. Sebastiani. 2006. Determining term subjectivity and term orientation for opinion mining, EACL-2006.
- M. Fiszman, D. Demner-Fushman, F. Lang, P. Goetz, and T. Rindfleisch. 2007. Interpreting Comparative Constructions in Biomedical Text. BioNLP-2007.
- M. Gamon, A. Aue, S. Corston-Oliver, S. and E. Ringger. 2005. Pulse: Mining customer opinions from free text. IDA-2005.
- G. Ganapathibhotla and B. Liu. 2008. Identifying Preferred Entities in Comparative Sentences. COLING-2008.
- V. Hatzivassiloglou, and K. McKeown, K. 1997.

- Predicting the Semantic Orientation of Adjectives. ACL-EACL-1997.
- M. Hu and B. Liu. 2004. Mining and summarizing customer reviews. KDD-2004.
- N. Jindal, and B. Liu. 2006. Mining Comparative Sentences and Relations. AAAI-2006.
- N. Kaji, and M. Kitsuregawa. 2006. Automatic construction of polarity-tagged corpus from HTML documents. ACL-2006.
- H. Kanayama, and T. Nasukawa. 2006. Fully Automatic Lexicon Expansion for Domain-Oriented Sentiment Analysis. EMNLP-2006.
- S. Kim and E. Hovy. 2004. Determining the Sentiment of Opinions. COLING-2004.
- N. Kobayashi, K. Inui and Y. Matsumoto. 2007. Extracting Aspect-Evaluation and Aspect-of Relations in Opinion Mining. EMNLP-2007.
- L.-W. Ku, Y.-T. Liang, and H.-H. Chen. 2006. Opinion Extraction, Summarization and Tracking in News and Blog Corpora. AAAI-CAAW.
- B. Liu. 2006. Web Data Mining: Exploring Hyperlinks, Content and Usage Data. Springer.
- B. Liu. 2009. Sentiment Analysis and Subjectivity. To appear in *Handbook of Natural Language Processing*, Second Edition, (editors: N. Indurkha and F. J. Damerau), 2009 or 2010.
- Y. Lu, and C. X. Zhai. 2008. Opinion integration through semi-supervised topic modeling. WWW-2008.
- R. McDonald, K. Hannan, T. Neylon, M. Wells, and J. Reynar. 2007. Structured models for fine-to-coarse sentiment analysis. ACL-2007
- Q. Mei, X. Ling, M. Wondra, H. Su, and C. X. Zhai. 2007. Topic Sentiment Mixture: Modeling Facets and Opinions in Weblogs. WWW-2007.
- V. Ng, S. Dasgupta, and S. M. Niaz Arifin. 2006. Examining the role of linguistic knowledge sources in the automatic identification and classification of reviews. ACL-2006.
- B. Pang and L. Lee. 2008. Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval* 2(1-2), pp. 1–135, 2008.
- B. Pang, L. Lee. and S. Vaithyanathan. 2002. Thumbs up? Sentiment Classification Using Machine Learning Techniques. EMNLP-2002.
- A.-M. Popescu, and O. Etzioni. 2005. Extracting Product Features and Opinions from Reviews. EMNLP-2005.
- G. Qiu, B. Liu, J. Bu and C. Chen. 2009. Expanding Domain Sentiment Lexicon through Double Propagation. IJCAI-2009.
- E. Riloff, and J. Wiebe. 2003. Learning extraction patterns for subjective expressions. EMNLP-2003.
- V. Stoyanov, and C. Cardie. 2008. Topic Identification for fine-grained opinion analysis. COLING-2008.
- I. Titov and R. McDonald. 2008. A Joint Model of Text and Aspect Ratings for Sentiment Summarization. ACL-2008.
- P. Turney. 2002. Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews. ACL-2002.
- J. Wiebe, R. Bruce, and T. O’Hara. 1999. Development and use of a gold standard data set for subjectivity classifications. ACL-1999.
- J. Wiebe, and T. Wilson. 2002. Learning to Disambiguate Potentially Subjective Expressions. CoNLL-2002.
- T. Wilson, J. Wiebe. and R. Hwa. 2004. Just how mad are you? Finding strong and weak opinion clauses. AAAI-2004.
- H. Yu, and Y. Hatzivassiloglou. 2003. Towards answering opinion questions: Separating facts from opinions and identifying the polarity of opinion sentences. EMNLP-2003.

Zero conditional sentences express general truthsâ€”situations in which one thing always causes another. When you use a zero conditional, youâ€™re talking about a general truth rather than a specific instance of something. Consider the following examples: If you donâ€™t brush your teeth, you get cavities. When people smoke cigarettes, their health suffers. There are a couple of things to take note of in the above sentences in which the zero conditional is used. First, when using the zero conditional, the correct tense to use in both clauses is the simple present tense. The following sentences illustrate a couple of the common mistakes people make when using the second conditional: If I inherit a billion dollars, I would travel to the moon. If I inherited a billion dollars, I would travel to the moon. Conditional sentences describe the result of something that might happen (in the present or future) or might have happened, but didnâ€™t (in the past). They are made using different English verb tenses with an â€”if clauseâ€™. The â€”if clauseâ€™ can go at the beginning or at the end of the conditional sentence. There are four main types of conditional sentences. type. condition. Sentiment analysis in NLP is about deciphering such sentiment from text. Is it positive, negative, both, or neither? If there is sentiment, which objects in the text the sentiment is referring to and the actual sentiment phrase such as poor, blurry, inexpensive, â€”| (Not just positive or negative.) This is also called aspect-based analysis [1]. As a technique, sentiment analysis is both interesting and useful. First, to the interesting part. A conditional Markov model (CMM) models this inference problem as one of finding the label sequence L that maximizes the conditional probability $P(L|T)$ for the given token sequence T. The Markov model makes certain assumptions which make this inference problem tractable. Specifically, $P(L|T)$ is assumed to be factorable as. Question (interrogative) sentences and conditional sentences are two important types, e.g., â€”Can you tell me which Sony camera is good?â€” and â€”If I can find a good camera in the shop, I will buy it.â€” Both these sentences contain the sentiment word â€”goodâ€”, but neither expresses a positive or negative opinion on any specific camera. However, not all conditional sentences or interrogative sentences express no sentiments, e.g., â€”Does anyone know how to repair this terrible printerâ€” and â€”If you. 12. In this sense, sentiment analysis offers a great platform for NLP researchers to make tangible progresses on all fronts of NLP with the potential of making a huge practical impact. In this book, I will describe the core problems and the current state-of-the-art algorithms. This paper studies sentiment analysis of conditional sentences. The aim is to determine whether opinions expressed on different topics in a conditional sentence are positive, negative or neutral. Conditional sentences are one of the commonly used language constructs in text. In a typical document, there are around 8% of such sentences. Due to the condition clause, sentiments expressed in a conditional sentence can be hard to determine. For example, in the sentence, if your Nokia phone is not good, buy this great Samsung phone, the author is positive about "Samsung phone" but does