

SES: Sentiment Elicitation System for Social Media Data

Kunpeng Zhang, Yu Cheng, Yusheng Xie, Daniel Honbo
Ankit Agrawal, Diana Palsetia, Kathy Lee, Wei-keng Liao, and Alok Choudhary
*Department of Electric Engineering & Computer Science
Northwestern University
Evanston, IL 60208*

Email: {kzh980,ych133,yxi389,dkh301,ankitag,drp925,kml649,wkliao,choudhar}@eecs.northwestern.edu

Abstract—Social Media is becoming major and popular technological platform that allows users discussing and sharing information. Information is generated and managed through either computer or mobile devices by one person and consumed by many other persons. Most of these user generated content are textual information, as Social Networks(Facebook, LinkedIn), Microblogging(Twitter), blogs(Blogspot, Wordpress). Looking for valuable nuggets of knowledge, such as capturing and summarizing sentiments from these huge amount of data could help users make informed decisions. In this paper, we develop a sentiment identification system called SES which implements three different sentiment identification algorithms. We augment basic compositional semantic rules in the first algorithm. In the second algorithm, we think sentiment should not be simply classified as positive, negative, and objective but a continuous score to reflect sentiment degree. All word scores are calculated based on a large volume of customer reviews. Due to the special characteristics of social media texts, we propose a third algorithm which takes emoticons, negation word position, and domain-specific words into account. Furthermore, a machine learning model is employed on features derived from outputs of three algorithms. We conduct our experiments on user comments from Facebook and tweets from twitter. The results show that utilizing Random Forest will acquire a better accuracy than decision tree, neural network, and logistic regression. We also propose a flexible way to represent document sentiment based on sentiments of each sentence contained. SES is available online.

Keywords-Social media, sentiment, rule, machine learning

I. INTRODUCTION

Social Media has become one of the most popular platforms to allow users discussing, communicating, and sharing their interested topics without having same geo-location and same time. Information can be generated and managed through either computer or mobile devices by one person and consumed by many other persons. Different people could express different opinions on the same topic. A wide variety of topics, ranging from current events and political debate, to sports and entertainment are being actively discussed on these social forums. The power of social media as a marketing tool has been recognized, and is being actively used by governments, major organization, schools and other groups to effectively and quickly communicate with large numbers of people.

Another important metric for business to measure their online reputation is word of mouth publicity. Word of mouth is the process of spreading information from person to person, and is often done through social media networks. It also plays a major role in customer buying decisions. Some typical examples are that Facebook users could comment campaigns posted by a company or like a company on Facebook, Twitter users could send tweets with a maximum length of 140 characters to instantly share and deliver their opinions on politics, movies, sports, etc. Collecting and analyzing these data could help users or managers make informed decisions. Marketing leaders or product managers might collect and analyze feedbacks and comments on campaigns launched by themselves from Facebook aiming to adopt efficient advertising strategy and improve product quality.

Most of these user generated content are textual information. The rapid growth in volume of web texts from major social network sites like Facebook and Twitter drives us to analyze and mine the data through computational techniques. Identifying their sentiments has become an important issue and attracted many attentions. Recently, there have been a number of studies attempting to model/predict real-world events using information from social media networks. Among these, Twitter has attracted additional attention because of the huge surge in its popularity. Jansen et al [21] perform a large-scale analysis of brand sentiments on Twitter. Their study concludes that 19% of tweets contain brand references, of which nearly 20% contain sentiments about the brands.

The main research efforts on sentiment analysis done previously can be classified into 2 branches. On one hand, they take state-of-the-art sentiment identification algorithms to solve problems in real applications such as summarizing customer reviews [14], ranking products [23], finding product features that imply opinions [15]. [22] analyzes tweet sentiments about movies and attempts to predict box office revenue. The authors define different metrics to measure the popularity/sentiment of a movie and then use a linear regression model to predict box-office

revenue. Joshi et al [24] use a similar technique to predict box-office revenue of movies using review text. On the other hand, researchers put their focus on discovering new sentiment algorithms. Bag-of-Words approach produces domain-specific lexicons. There is a vast body of research which attempts to incorporate these interactions as features in a machine learning model [2], [3], [18]. Rule-based approaches has been studied by many researchers. [1] propose compositional semantics, which is based on the assumption that the meaning of a compound expression is a function of the meaning of its parts and of the syntactic rules by which they are combined. They have developed a set of composition rules to assign sentiments to individual clauses, expressions and sentences. We have augmented these rules by adding our own rules which are specific to social media texts. In addition to these rules, we require a method of assessing the impact of these on the the polarity of an expression. We also develop our version of the Compose function for computing the polarity of an expression based on Compositional Semantic rules. Our first algorithm is implemented based on this. There is one fact that most previous work provides binary polarities only (positive and negative), and the polarity of sentiment is simply reversed when a negation is detected. In this paper, we incorporate additional parse-and-paraphrase paradigms based on basic pattern (adverb-adjective-noun) proposed by Liu et al to assess the degrees of sentiments and utilize generated sentiment scores to evaluate sentiments for social media data [4]. Our second algorithm uses scores to output sentiment for a sentence. Social media texts have their own speciality that they often have emoticons and misspellings or shortened versions of common words. We observe that the presence of an emoticon almost always convey the underlying sentiment. In our third algorithm, we not only collect most common positive/negative sentiment words, negation words, but also positive/negative emoticons. Negation words do not only have simple functions of reversing the sentiment. Their position related to sentiment words in a sentence also plays an important role in deciding sentiment.

We build a web-based system called SES, which ensembles three algorithms we implemented and uses machine learning method to predict text sentiment. The system is aiming to predict sentiment on both sentence and document level. A document often contains more than one sentences. We split it into sentences which are the input of the system. In addition, we have designed a strategy to represent a overall sentiment of a document. A concept of "mixture" is imported and applied. We conduct our experiments on Facebook comments and twitter tweets using four different machine learning models: decision tree, neural network, logistic regression, and random forest. The experiment results show that random forest model reaches highest accuracy.

Due to the independence of sentiment prediction for different sentences, a simple parallel mechanism has been employed. The detailed system architecture and flow of SES will be described in section III.

II. RELATED WORK

Recently, there has been a wide range of research done on sentiment analysis, from rule-based, bag-of-words approaches to machine learning techniques. One of the main directions is sentiment classification, which classifies the whole opinion document (e.g., a product review) as positive or negative [2], [5]–[7]. In [1], they view such subsentential interactions in light of compositional semantics, and present a novel learning-based approach that incorporates structural inference motivated by compositional semantics into the learning procedure. They also find that "content-word negators" plays an important role in determining expression-level polarity. In [2], authors employ machine learning techniques to classify documents by overall sentiments and conducted their experiments on movie reviews and the results show that three machine learning methods they employed (Naive Bayes, maximum entropy classification, and support vector machines) do not perform as well on sentiment classification as on traditional topic-based categorization. Another important direction is classifying sentences as subjective or objective, and classifying subjective sentences or clauses as positive or negative [8]–[13]. In [20], authors present a linguistic analysis of conditional sentences, and build some supervised learning models to determine if sentiments expressed on different topics in a conditional sentence are positive, negative or neutral.

Several researchers also studied feature/topic-based sentiment analysis [15]–[19]. Their objective is to extract topics or product features in sentences and determine whether the sentiments expressed on them are positive or negative. In [14], authors aim to summarize all customer reviews of a product by mining the features of the product on which customers have expressed their opinions and whether the opinions are positive or negative. In [3], authors use feature-based opinion mining model to identify noun product features that imply opinions. It is mainly focusing on the problem of objective nouns and sentences with implied opinions. In [4], authors propose an approach to extracting adverb-adjective-noun phrases based on clause structure obtained by parsing sentences into a hierarchical representation. They also propose a robust general solution for modeling the contribution of adverbials and negation to the score for degree of sentiment. This is the basis of our second algorithm plus some extra phrases added by us.

III. SYSTEM DESCRIPTION

The system contains multiple tasks from the input raw data towards output the final sentiment. The first task is data preprocessing. The second task is running three algorithms

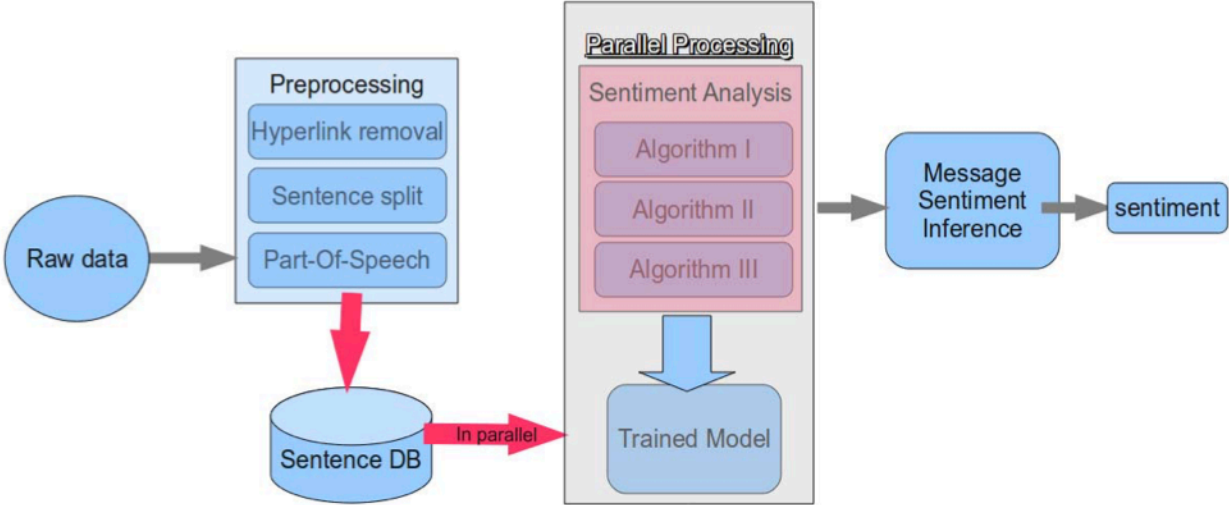


Figure 1: SES Architecture

for each sentence to get three individual outputs. Then we generate features based on these outputs and put them into trained model to output sentiment. The fourth task is the combination task which combines all sentence sentiments to generate an overall sentiment of the input document. The Figure 1 shows the architecture of the system. There is a mysql database which is used to store sentences split from the input document. Due to the independence of sentiment identification for each sentence, SES implements sentence identification in parallel. Web server is used to serve interaction with web users. The following sections will describe these tasks step by step.

A. Data Preprocessing

A document typically consists of several sentences. It is not uncommon to see multiple positive and negative opinions of a company or a person in a single comment. For example, a comment of a new released Amazon kindle may use a few sentences to praise the screen quality and other sentences to criticize its weight and wireless. To simplify this problem, we split a document into sentences. The sentences then can be assigned positive or negative sentiments. In this study, we use MAXTERMINATOR to split messages into sentences [25]. It is very common to see that most social media texts contain hyperlinks and have words or phrases misspelled. Removing hyperlinks and correcting misspellings of a sentence is required in order to improve the accuracy of sentiment identification. We have manually collected 138 pairs of misspelled and corrected words. In addition, most sentiment bearing words are adjectives, adverbs, verbs, and negation words. Rules defined by three algorithms also require us to identify part-of-speech information for all words in a sentence. In this system, CRF Tagger, a java-based conditional random field

part-of-speech(POS) tagger for English is employed to label each word [26].

B. Individual Algorithms

In this section, we will describe three algorithms of identifying sentiment for a single sentence. The first one is compositional semantic rules. A few new rules are added to the basic compositional rules proposed by [1]. The output is one of 5 integers ranging from -2 to +2 (+2 means strongly positive, -2 means strongly negative, 0 means neutral). The second algorithm is considering that sentiment does not have to be limited into three labels(positive, negative, neutral). Rather, we think that the sentiment degree should be distinguished and reflected by a numerical score. The third algorithm identifies sentiment by checking some rules defined on positive/negative words, positive/negative emoticons, the position of negation words relative to sentiment words, and some domain-specific words.

1) *Compositional Semantic Rule Algorithm*: Table I shows the compositional rules and corresponding examples. The first 7 rules are the basis for the first algorithm and have been expanded upon. Since grammar in social media texts and part-of-speech(POS) tagging is not perfect, we have also added some rules to catch some of the errors. Keep in mind that the grammar will not always be correct in these rules, for example, in the sentence “lack of killing in rural area”, killing should be a noun, words ending in “ing” are sometimes incorrectly tagged as verbs instead of nouns. In addition to these rules, we require a method of assessing the impact of these on the the polarity of an expression. Our version of the compose function for computing the polarity of an expression based on compositional semantic rules is given in Table II.

Table I: Compositional semantic rules

Rules	Example
1. Polarity(not [arg1]) = \neg polarity(arg1)	not [bad]{arg1}
2. Polarity([VP1] [NP1]) = Compose([VP], [NP])	[destroyed]{VP} the [terrorism]{NP}
3. Polarity([VP1]to[VP2]) = Compose([VP1],[VP2])	[refused]{VP1} to{to} [deceive]{VP2} the man
4. Polarity([ADJ]to[VP1]) = Compose([ADJ],[VP1])	unlikely]{ADJ} to{to} [destroy]{VP} the planet
5. Polarity([NP1]in[NP1]) = Compose([NP1],[NP2])	[lack]{NP1} offing [crime]{NP2} in rural areas
6. Polarity([NP1] [VP1]) = Compose([VP1],[NP1])	crime]{NP1} has [decreased]{VP1}
7. Polarity([NP1]be[ADJ]) = Compose([ADJ],[NP1])	[damage]{NP1} is{be} [minimal]{ADJ}
Our Rules	Example
8. Polarity([NP1]in[VP1]) = Compose([NP1],[VP1])	lack]{NP1} offing killing{VP1} in rural areas
9. Polarity(as[ADJ]as[NP]) = if(polarity(NP) !=0: return polarity(NP), else: return polarity(ADJ)	as{as} ugly{ADJ} as{as} a rock{NP}
10. Polarity(not as[ADJ]as[NP]) = -polarity(ADJ)	that was not{not} as{as} [bad]{ADJ1} as the original]{NP2}
11. If the sentence contains ‘but’, disregard all previous sentiment only take the sentiment of the sentence after ‘but’	and I’ve never liked that director, [but] I loved this movie.
12. If the sentence contains ‘despite’, only the sentiment in the previous part of the sentence is counted.	I love that movie, despite the fact that I hate the director.

Table II: The compose function is used to derive the polarity of an expression. The table lists the Compose function used in [1]. Compose2 is our version of the Compose function

Compose(arg1,arg2)	if(arg1) is a negator then \neg polarity(arg2) else if (Polarity(arg1) == Polarity(arg2)) then Polarity(arg1) else the majority polarity of data
Compose2(arg1,arg2)	if arg1 is negative: if arg2 is not neutral: return : polarity(arg2) else: return -1 else if arg1 is positive and arg2 is not neutral: return polarity(arg2) else if polarity(arg1) equals polarity(arg2): return 2 polarity(arg1) else if (arg1 is positive and arg2 is neutral) or (arg2 is positive and arg1 is neutral): return polarity(arg1) + polarity(arg2) else: return 0

2) *Numeric Sentiment Identification Algorithm:* To calculate the numerical degree of sentiment, there are two major problems to solve: 1) how to associate numeric scores with degree of textual sentiment; 2) how to combine all the scores of multiple words for one sentence. In this section, we will describe these in details.

Definition of Strength of Sentiment The approach hypothesizes that there are two kinds of phrase that can associate with numerical scores: adverb-adjective-noun(abbreviated as AAN) and verb-adverb(VA) phrase. For example, “a very good question” is the AAN type and “do not like it very much” is the VA type. Based on this, the key point is to define the strength of sentiment of words for each phrase. There have been studies on building sentiment lexicons to define the strength of word sentiment. Esuli and Sebastiani [28] constructed a lexical resource, Senti-WordNet, a WordNet-like lexicon emphasizing sentiment orientation of words and providing numerical scores of how objective, positive, and negative these words are. However, lexicon-based methods can be tedious and inefficient and may not be accurate due to the complex cross-relations in dictionaries like WordNet. Jing and Seneff’s approach [4] to sentiment scoring is to make use of collective data such

as user star ratings in reviews. It assumed that user star rating is normally consistent with the tone of the review text published by the same user. Our method is similar to this. But we argue that the sentiment scores are not only associated with user star ratings but also word appearance frequency. By associating user star ratings and frequency with each phrase extracted from review texts, we can easily associate numeric scores with textual sentiment. For both adjective and adverb-adjective phrase, we average its star ratings given by 1:

$$Score(w) = \frac{\sum_{i \in P} \frac{n_{r_i}}{N} r_i f_{r_i}}{\sum_{r_i} \frac{n_{r_i}}{N} f_{r_i}} \quad (1)$$

where P represents the set of appearances of word w , f_{r_i} is the appearance frequency in each entities, r_i represents the associated user rating in each appearance of w , N represents the number of entities (e.g., company) in the entire data set, and f_{r_i} represents the number of entities with star rating r_i . The score is averaged over all appearances, weighted by the frequency count for removing bias towards any words. We take the first AAN phrase as an example and we consider the noun words are objective without sentiment. For the VA phrase, we can calculate the scores in the same way. Table III shows the calculated scores of some adjective words.

Based on [4], for each adverb, we get a list of all

Table III: Examples of scores of some adjective words

adjective words			
Adj	score	Adj	score
Easy	4.1	hard	2.5
Good	3.9	Best	5.0
Rude	1.69	Worst	1.0
Nice	3.7		
Bad	1.5		
Great	4.5		

Table IV: Examples of scores of some adverb words

adverb words			
Adv	score	Adv	score
Super	0.41	Pretty	0.18
Not	-1.90	Most	1.0
Little	-0.16	Never	-2.0
A bit	0.03		
Pretty	0.06		
Really	0.42		

possible combinations with adjectives. Then, for each adjective in the list, we calculate the distance between the rating of adverb-adjective and the rating of the adjective. $Pol(*)$ is the index function of polarity with two possible values, +1 meaning positive, -1 meaning negative. Table IV shows the calculated scores of some adverb words.

$$Score(adv) = Pol(adj) \cdot (Score(adv, adj) - Score(adj)) \quad (2)$$

Derivation of Sentence Score After obtaining the strength rating for adverbs and adjectives, the next step is to assign the strength of sentiment to each phrase (AAN and VA) extracted by linguistic analysis, as given by 3 and 4 in a linear way:

$$Score(adv(adj(noun))) = Score(adj) + Pol(adj) \cdot Score(adv) \quad (3)$$

$$Score(adv(verb)) = Score(verb) + Pol(verb) \cdot Score(adv) \quad (4)$$

Specially for the negation-adverb-adjective-noun and negation-adverb-verb phrase, instead of treating negation as a special case, the universal model works for all negations. This allow us to handle the negation-adverb-adjective-noun phrase in this way:

$$Score(neg(adv(adj))) = Score(adj) + Pol(adj) \cdot Score(adv) + Pol(adj) \cdot Score(neg) \quad (5)$$

$$Score(neg(adv(verb))) = Score(verb) + Pol(verb) \cdot Score(adv) + Pol(verb) \cdot Score(neg) \quad (6)$$

The sentence score is calculated based on the summation of all patterns we discussed above. The algorithm is listed below(Algorithm 1). The higher score means more strongly positive.

Algorithm 1 Calculating numeric score for a sentence

1. assign scores to all adjectives and adverbs
2. extract all phrases(P) and calculate each score(PS)
3. $S = \sum_{i=1}^m PS_i$, where m is the number of phrases

3) *Bag-of-Word and Rule-based Algorithm*: Due to the special characteristics of social media texts, we define some rules to analyze sentiments. For example, most of Facebook comments and twitter tweets contain emoticon like ‘:)’ (positive sentiment) or ‘:(’ (negative sentiment) which almost always conveys the underlying sentiment. We believe that there are a very few number of cases where the underlying sentiment/polarity of the comment/tweet is different from that of the emoticon present. Moreover, such cases are notoriously difficult to classify, since they are often sarcastic. In our system, we collect 77 positive emoticons and 59 negative emoticons from wikipedia [27]. The second difference from general text is that social media text is very short and authors usually use Internet language to express their opinions, such as “1st!”, “Thank you, Obama”, “Go Bulls!”, etc. In this algorithm, a rule is introduced to process this situation: if the sentence belongs to the pattern of “[Thank you|go], [a company name(organization)|a person name]”, we label it as positive. In addition, some domain-specific keywords should be added into sentiment word sets, for instance, “Yum, Yummy” should be positive word for food comments. The detailed description of this algorithm implemented in the SES is the following(Algorithm 2).

Algorithm 2 Bag-of-word and rule-based algorithm

Require: PS: positive word set, NS: negative word set, NG: negation word set, PE: positive emoticon set, NE: negative emoticon set.

if exists an emoticon ‘e’ ∈ (PE, NE) **then**
return Polarity(e)

else

array[n] ← split sentence into words

words[m] ← sentiment and negation words of array[n]
(keeping original index order)

for $i = 0 \rightarrow m - 1$ **do**

1. pos ← count number of positive words

2. neg ← count number of negative words

3. reverse sentiment for words exactly following a negation word

4. neighboring negation words will be counteracted

end for

pos > neg return ‘P’

pos < neg return ‘N’

return ‘O’

end if

C. Machine Learning Model

In this section, we are training a predictive model which ensembles three algorithms to classify the sentiment of a sentence. It has 3 basic features: 1): the integer value generated from compositional semantic rules (S_1), 2): the float value of the second algorithm output (S_2), 3): the output of the third algorithm ('P', 'N', or 'O') (S_3), and 2 derived features: 1): $S_1 + S_2$, 2): $S_2 - S_1$. 10-cross validation is employed here. We use four different predictive models: decision tree(J48), neural network, logistic regression, and random forest.

D. Inference of Sentiment for the Entire Message

A message or a comment usually have more than one sentence. We use three individual algorithms and trained model to output sentiment for each sentence. The next task is to integrate all sentiments to represent message sentiment. A concept of "mixture" is introduced here to help us describe sentiment of a message in a more flexible way. As we mentioned before, it is very common that users use both positive and negative sentences to express their opinions within one comment. Labeling such kind of messages just as positive or negative doesn't give us an accurate and reasonable result. In this system, we define 5 different categories to label a message. If a message only have positive sentences and objective sentences, we label it as positive message. Similarly, we label a message as negative if it only has negative and objective sentences. If there are only objective sentences constituting a message, it will be classified as objective. More complicated situation is that a message contains both positive and negative sentences. In this case, if the number of positive sentences is bigger than the number of negative sentences, the message goes to a label of positive mixed. Contrarily, the message is labeled as negative mixed, otherwise it's mixed.

IV. EXPERIMENT RESULTS

This section describes the experiment conducted on Facebook comments and Twitter tweets by using four different machine learning models. We manually label 1000 twitter tweets and 2000 Facebook comments. Figure ?? shows that random forest model obtains higher weighted average accuracy than decision tree, logistic regression, and neural network. The weighted average values of 4 different measurements are shown in Figure 2 and 3.

V. SYSTEM INTERFACE

Based on what we have discussed above, we build a web-based system which allows users to identify sentiments for social media texts. The SES screenshots are shown in Figure 5. The main configuration is Apache2+MySQL+PHP5. We

use Apache2 as the webserver and MySQL to host database server. In this system, we provide users two options to interact. Users could input a document into the textarea or upload a text file.

VI. CONCLUSION AND FUTURE WORK

We implement a web-based system SES which integrates three different algorithms for individual sentence. We augment additional rules specific to social media text beyond the basic compositional semantic rules. For the second algorithm, we use numerical score to evaluate the degree of sentence sentiment. Due to the special characteristics of social media data, we bring a bag-of-word and rule-based algorithm which considers the sentiment words, emoticons, negation word position, and domain-specific words. A machine learning model is employed based on the 5 features derived from outputs generated by three algorithms. In the future, we think that the context information in a document is important. And conditional random fields is a good model to capture it. Besides these, building a comprehensive and domain-specific sentiment words will help sentiment identification. In addition, the major and minor entities in a comparative sentence will influence sentiment decision.

ACKNOWLEDGMENT

This work is supported in part by NSF award numbers CCF-0621443, OCI-0724599, CCF-0833131, CNS-0830927, IIS-0905205, OCI-0956311, CCF-0938000, CCF-1043085, CCF-1029166, and OCI-1144061, and in part by DOE grants DE-FC02-07ER25808, DE-FG02-08ER25848, DE-SC0001283, DE-SC0005309, and DE-SC0005340.

REFERENCES

- [1] Y. Choi, and C. Cardie, *Learning with compositional semantics as structural inference for subsentential sentiment analysis*. Proceedings of the Conference on Empirical Methods in Natural Language Processing, pages 793–801, 2008.
- [2] B. Pang, and L. Lee and S. Vaithyanathan, *Thumbs up? Sentiment Classification using Machine Learning Techniques*. Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 79–86, 2002.
- [3] L.Zhang, and B. Liu and S. Vaithyanathan, *Identifying noun product features that imply opinions*. Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers - Volume 2, pages 575–580, 2011.
- [4] J. Liu, and S. Seneff, *Review sentiment scoring via a parse-and-paraphrase paradigm*. Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1 - Volume 1, pages 161–169, 2009.
- [5] J. Liu, and S. Seneff, *Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews*. Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, pages 417–424, 2002.




Figure 2: The precision, recall, F-measure, ROC, and accuracy for Facebook comments




Figure 3: The precision, recall, F-measure, ROC, and accuracy for Twitter tweets



(a) Index page



(b) Result page

Figure 4: Screenshots of system SES

- [6] K. Dave, and S. Lawrence, and D. M. Pennock, *Review sentiment scoring via a parse-and-paraphrase paradigm*. Proceedings of the 12th international conference on World Wide Web, pages 519–528, 2003.
- [7] R. McDonald, and K. Hannan, and T. Neylon, and M. Wells, and J. Reynar, *Structured Models for Fine-to-Coarse Sentiment Analysis*. Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics, 2007.
- [8] J. Wiebe, and R. Bruce, and T. OHara, *Development and use of a gold standard data set for subjectivity classifications*. Proceedings of the 37th Annual Meeting of the Association of Computational Linguistics, pages 246–253, 1999.
- [9] J. Wiebe, and T. Wilson, *Learning to Disambiguate Potentially Subjective Expressions*. proceedings of the 6th conference on Natural language learning - Volume 20, pages 1–7, 2002.
- [10] H. Yu, and V. Hatzivassiloglou, *Towards answering opinion questions: separating facts from opinions and identifying the polarity of opinion sentences*. Proceedings of the 2003 conference on Empirical methods in natural language processing, pages 129–136, 2003.
- [11] T. Wilson, and J. Wiebe, and R. Hwa, *Just how mad are you? finding strong and weak opinion clauses*. Proceedings of the 19th national conference on Artificial intelligence, pages 761–767, 2004.
- [12] S. Kim, and E. Hovy, *Determining the sentiment of opinions*. Proceedings of the 20th international conference on Computational Linguistics, 2004.
- [13] E. Riloff, and J. Wiebe, *Learning extraction patterns for subjective expressions*. Proceedings of the 2003 conference on Empirical methods in natural language processing, pages 105–112, 2004.
- [14] M. Hu, and B. Liu, *Mining and summarizing customer reviews*. Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining, pages 168–177, 2004.
- [15] A. Popescu, and O. Etzioni, *Extracting product features and opinions from reviews*. Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing, pages 339–346, 2005.
- [16] Q. Mei, and X. Ling, and M. Wondra, and H. Su, and C. Zhai, *Topic sentiment mixture: modeling facets and opinions in weblogs*. Proceedings of the 16th international conference on World Wide Web, pages 171–180, 2007.
- [17] L. Ku, and Y. Liang, and H. Chen, *Opinion Extraction, Summarization and Tracking in News and Blog Corpora*. Proceedings of AAAI-CAAW-06, the Spring Symposia on Computational Approaches to Analyzing Weblogs, 2006.
- [18] X. Ding, and B. Liu, and P. S. Yu, *A holistic lexicon-based approach to opinion mining*. Proceedings of the international conference on Web search and web data mining, pages 231–240, 2008.
- [19] V. Stoyanov, and C. Cardie, *Topic identification for fine-grained opinion analysis*. Proceedings of the 22nd International Conference on Computational Linguistics - Volume 1, pages 817–824, 2008.
- [20] R. Narayanan, and B. Liu, and A. Choudhary, *Sentiment analysis of conditional sentences*. Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1 - Volume, pages 180–189, 2009.
- [21] B. Jansen, and M. Zhang, and K. Sobel, and A. Chowdury, *Twitter power: Tweets as electronic word of mouth*. J. Am. Soc. Inf. Sci. Technol., pages 2169–2188, 2009.
- [22] S. Asur, and B. A. Huberman, *Predicting the future with social media*. Arxiv preprint arXiv:1003.5699, 2010.
- [23] K. Zhang, and R. Narayanan, and W. Liao, and A. Choudhary, *Voice of the Customers: Mining Online Customer Reviews for Product Feature-based Ranking*. 3rd Workshop on Online Social Networks, 2010.
- [24] M. Joshi, and D. Das, and K. Gimpel, and N. A. Smith, *Movie reviews and revenues: An experiment in text regression*. Proceedings of NAACL-HLT, 2010.
- [25] J. C. Reynar, and A. Ratnaparkhi, *A Maximum Entropy Approach to Identifying Sentence Boundaries*. Proceedings of the fifth conference on Applied natural language processing, 1997.
- [26] X. Phan, *CRFTagger: CRF English POS Tagger*. <http://crftagger.sourceforge.net/>, 2006.
- [27] [http://http://en.wikipedia.org/wiki/List_of_emoticons](http://en.wikipedia.org/wiki/List_of_emoticons)
- [28] A. Esuli, and F. Sebastiani, *SENTIWORDNET: A Publicly Available Lexical Resource for Opinion Mining*. Proceedings of the 5th Conference on Language Resources and Evaluation, pages 417–422, 2006.