Learning Good Opinions from Just Two Words Is Not Bad

Darius Andrei Suciu, Vlad Vasile Itu, Alexandru Cristian Cosma,
Mihaela Dinsoreanu and Rodica Potolea
Technical University of Cluj-Napoca, Cluj-Napoca, Romania

Keywords: Unsupervised Learning, Opinion Mining, NLP, Domain Independent Learning, Implementation.

Abstract: Considering the wide spectrum of both practical and research applicability, opinion mining has attracted increased attention in recent years. This article focuses on breaking the domain-dependency issues which occur in supervised opinion mining by using an unsupervised approach. Our work devises a methodology by considering a set of grammar rules for identification of opinion bearing words. Moreover, we focus on tuning our method for the best tradeoff between precision-recall, computation complexity and number of seed words while not committing to a specific input data set. The method is general enough to perform well using just 2 seed words therefore we can state that it is an unsupervised strategy. Moreover, since the 2 seed words are class representatives ("good", "bad") we claim that the method is domain independent.

1 INTRODUCTION

Information is becoming more and more abundant especially over the internet. Twitter alone reports an average of *58 million* tweets per day, this being a small fraction of the flood of free information which surges on the web. Considering we already have a free supply of information, the most important questions which can be asked are: *What can we do with it? How can we put it to good use? How can we use all of it?*

The subfield of data mining which tries to answer this question is that of Opinion Mining. Its goal is to extract useful subjective information from user generated content, like customer reviews of products, tweets, blog articles, and forum discussions.

In opinion mining, a *feature* (or target) is a topic on which opinions are expressed. Opinions without associated features would be less valuable information. As an example, in the sentence: *The camera was extraordinary* if we wouldn't know that *camera* is the target and would only have *extraordinary* as opinion, the information would not be relevant. Moreover, opinions have a polarity (or semantic orientation) which can fall into the positive, neutral or negative spectrum, depending on the context it is being used in. For example, *the actors' performance was cold* may indicate a bad performance and thus *cold* has a negative polarity. At the same time, the sentence *after installing the fan, the processor became cold* may indicate that the fan did its job, which suggests *cold* conveys a positive orientation. Therefore, context is the key.

In this paper, we focus on opinion extraction in text documents - more specifically, in customer reviews. Given a set of reviews, the goal is to identify and classify targets according to the opinion expressed toward them.

To achieve the objective, the system proposed in this paper follows a domain independent, unsupervised approach for performing feature/aspect-based opinion extraction and polarity assignment on user generated content. The starting point of the proposed method is a rule-based, iterative technique proposed in (Liu, 2012). An important problem in opinion summarization caused by domain specific opinion words is handled very well by this approach as it extracts both opinion words and features. Because the extraction process also introduces noise, we propose a set of pruning and filtering methods designed to improve performance. The proposed solution performs reliably and efficiently on cross-domain corpora while offering the possibility to fine-tune the system using a set of parameters.

2 RELATED WORK

The approaches and techniques used to perform the opinion summarization task vary and belong to

Suciu D., Itu V., Cosma A., Dinsoreanu M. and Potolea R..
Learning Good Opinions from Just Two Words Is Not Bad.
DOI: 10.5220/0005079802330241
In *Proceedings of the International Conference on Knowledge Discovery and Information Retrieval* (KDIR-2014), pages 233-241
ISBN: 978-989-758-048-2

different/complementary research areas: text mining, sentiment prediction, classification, clustering, natural language processing, usage of resource terms and so on.

In (Hatzivassiloglou, 1997) the authors extract adjectives joined by conjunction relations (*and* / *or)*, based on the concept that adjectives joined by conjunction have the same or opposite polarity and semantic value.

In (Turney, 2002) 3-grams are compared against a predefined syntactical relationship table, extracting targets and their associated opinion words along with their sematic values.

In (Hu and Liu, 2004) frequent nouns and noun phrases are used to extract product feature candidates.

The target extraction proposed in (Popescu, 2005) determines whether a noun or noun phrase is a product feature or not. A PMI score is computed between the phrase and its discriminant found by a search on the Web by using the known product class.

In (Jin et. all. 2009) lexical Hidden Markov Models are employed. A propagation module extends the previously extracted targets and opinion words. The authors expand the opinion words with synonyms *and* antonyms and expand the targets with related words combining them into bigrams. The noise is treated using weights which are assigned to the resulted bigrams.

The extraction of product features using grammar rules is described in (Zhang et. all. 2010). They also use the HITS algorithm, a link analysis algorithm for rating Web pages along with feature frequency for ranking features by relevance.

In (Liu, 2012) seed words set expansion and features identification are described. The *seed words* set, denoted also as *lexicon*, is composed of adjectives with a polarity associated – in the form of a positive, neutral or negative score. The features and opinion words are extracted in pairs, by using a dependency grammar and by exploiting the syntactic dependencies between nouns and adjectives in sentences.

Supervised and unsupervised approaches are combined for extracting opinion words and their targets in (Su Su Htay and Khin Thidar Lynn, 2013). Targets are extracted by using a training corpus, while opinion words are extracted by using grammar rules. The problem from combining approaches lies in the domain dependency given by the supervised part.

In (Hu et. all, 2013) sentiments are extracted out of the emoticons used in social texts like blogs, comments and tweets. The authors use the orthogonal nonnegative matrix tri-factorization model (ONMTF); clustering data instances based on the distribution of features, and features according to their distribution of data instances.

(Guerini et. All, 2013) tackles a polarity assignment problem, using a *posterior polarity* for achieving polarity consistency through the text. The authors also obtain better results from a framework constructed from a collection of posterior polarity calculating formulas. Their results also show the advantage of computing the average of all senses of a word over the usage of its most frequent sense.

In order to determine the opinion polarity values, in (Marrese-Taylor et. all. 2013), a lexical and a rule-based approach is proposed. A polarity lexicon and linguistic rules are used to obtain a list of words with known orientations.

Our work devises a generalized methodology by considering a comprehensive set of grammar rules for identification of opinion bearing words. Moreover, we focus on tuning our method for the best tradeoff between precision-recall, time and number of seed words. The method is general enough to perform well using just 2 seed words therefore we can state that it is an unsupervised strategy. Moreover, since the 2 seed words are class representatives ("good", "bad") we claim that the method is domain independent.

3 THE PROPOSED TECHNIQUE

The method proposed in this paper is presented in Figure 1, where the conceptual modules of our architecture together with the intermediate data produced are depicted. The architecture is composed by 3 components: 1 – Retriever Service; 2 – Feature-Opinion Pair Identification, 3 – Polarity Aggregator.

The Retriever services generate *syntactic trees* from the given input corpus. This preprocessing module handles the usual NLP tasks. The transformations applied at sentence level are: tokenization, lemmatization, part-of-speech tagging and syntactic parsing. First, each review document is segmented into sentences, which are used for discovering words in the tokenizing step. Lemmatization reduces the word to its base (root) form. Finally, the parsing step generates syntactic trees for each sentence, given the output of the previous steps. This syntactic decomposition is used as input for the second main task of the system, the identification of feature-opinion pairs.

The <feature, opinion> tuple identification component extracts the feature-opinion pairs using the double propagation algorithm. The rule-based strategy followed - double propagation - uses the extraction rules listed in (Cosma, 2014).

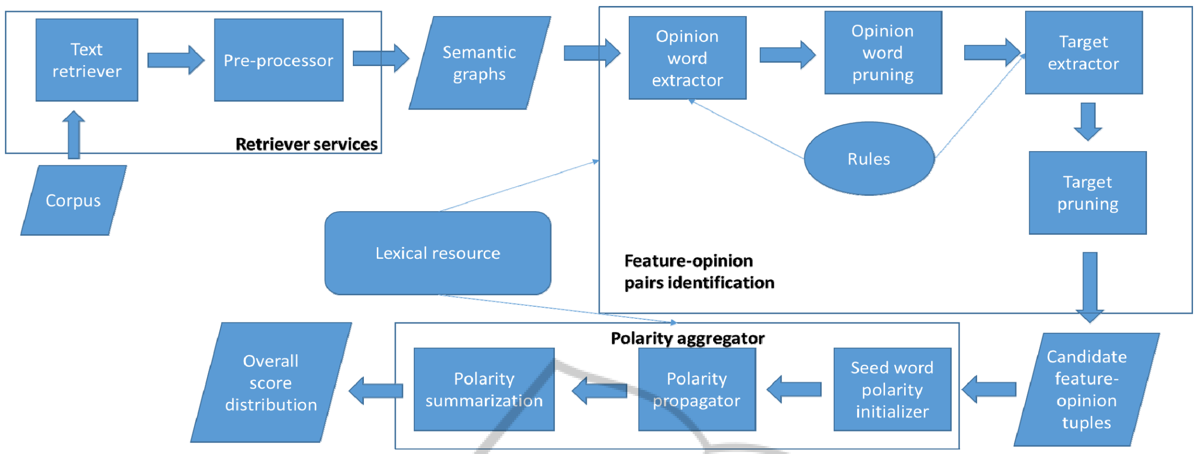

Figure 1: Overall system architecture.

The main idea of the double propagation algorithm is to boost the recognition rate in one side (opinion words) by identifying many words in the other side (targets) – back and forth. The extraction method is applied iteratively: the found adjectives and nouns are added to the input set, then new features and opinion words are extracted using the existing ones.

Based on the rules used, polarity scores are transferred from targets or opinions to the newly extracted word. The propagation ends when only few or no new entities are identified. In the end, all the polarity values of the extracted targets are aggregated to form an overall review score. The key to identifying opinion words and targets is the use of the syntactic relations defined in (Cosma 2014).

The propagation consists of four subtasks:

- Extracting targets using opinion words
- Extracting opinion words using targets
- Extracting targets using targets
- Extracting opinion words using opinion words

We propose the set of rules defined in (Cosma, 2014), which start from the set of rules defined in (Liu, 2012) along with additional constructed rules for extracting adjectives as opinion words based on (Turney, 2002) and new original ones for extracting pronouns as targets.

The process of extracting opinion words and targets from a text using syntactic dependencies introduces noise so a filter is devised to prune opinion words based on their objectivity. The filter's objective is to remove adjectives and adverbs which are not opinion words. An adjective or adverb is considered to be an opinion word if its polarity is above (in case of a positive opinion word) or below (negative) a calculated threshold, ensuring that objective words are not extracted, thus reducing the noise propagation. The finding was triggered in the initial experimental phase, when many adjectives extracted expressed a property, not an opinion (first, other, long, etc.).

The double propagation algorithm is presented in the following pseudo code:

```
   Input: Seed Word Dictionary {S},
Syntactic Trees {T}
   Output: All Features {F}, All Opinion
Words {O}
   Constant: Objectivity Threshold {Th}
   Function:
   1.  {O} = {S}
   2.  {F1} = Ø, {O1} = Ø
   3.  For each tree in T:
   4.          if( Extracted features
       not in {F})
   5.                Extract features
       {F1} using R1, R2 with {O}
   6.          endif
   7.         if( Extracted opinion
       words not in {O} and opinion
       words objectivity < {Th})
   8.               Extract opinion
       words {O1} using R3, R5 with {O}
   9.         endif
   10. endfor
   11. Set {F} = {F} + {F1},  {O} = {O} +
       {O1}
   12. For each tree in T:
   13.         if( Extracted features
       not in {F})
   14.               Extract features
       {F2} using R4 with {F1}
   15.         endif
   16.        if( Extracted opinion
       words not in {O} and opinion
       words objectivity < {Th})
   17.              Extract opinion
       words {O2} using R6, R7 with {F1}
```

```
        endif
endfor
Set {F1} = {F1} + {F2},  {O1} =
    {O1} + {O2}
Set {F} = {F} + {F2}, {O} = {O} +
    {O2}
Repeat 2 until size({F1}) = 0 and
    size({O1}) = 0
```

Take for example the following sentences: *The laptop is amazing. The processor is fast and games are amazing and fast, running them on this laptop. The display is also responsive and fast.* Considering only *amazing* as an initial seed word, at the first iteration the algorithm extracts *laptop* and *games* as targets and also extracts *fast* as an additional opinion word. At the second iteration, *processor* and *display* are extracted as targets and *responsive* is extracted as opinion word. The third iteration does not extract any new data thus ending the algorithm.

The Target pruning module filters out targets based on their occurrence frequency. Because in reviews the product and its features occur more often along opinion words than other nouns, they can be pruned after the extraction algorithm is finished by removing the ones not extracted at least *t* number of times, where *t* is a target frequency threshold. The value of t which provides the best precision/recall ration has been determined experimentally.

The third component, Polarity Aggregator, performs the task of assigning polarity values to the extracted opinion words and targets. Moreover it generates a polarity summary by aggregating the individual scores. The Polarity aggregator assigns polarities to seed words using a lexical resource described in the results section. Because a lexical resource usually contains multiple polarities for the same word, depending on the context, the resulting polarity is retrieved as the weighted average of all those polarities. The module uses the list of polarity-charged seed words to assign polarities to the entire text in two steps. In the first step, polarity-charged seed words are matched throughout the text. In the second step, the previously matched scores are *propagated* in the entire text. Polarity assignment is accomplished with respect to the following rules:

- Opinion words extracted using targets receive the same score as the target
- Targets extracted using opinion words receive the same score as the opinion word
- Targets extracted using targets receive the same score
- Opinion words extracted using opinion words receive the same score.
- If the same target is discovered using different opinion words, the resulting score is the average of the opinion words.

4 RESULTS

To evaluate our strategy we used the dataset proposed in (Hu and Liu, 2004) and adjusted it to our needs by manually annotating the opinion words and targets.

The dataset is composed of 5 subsets of documents, four of which contain multiple reviews targeting a different product, and one represents a fraction of the movie reviews from (Taboada et. all, 2006). The figures regarding each dataset document are presented in Table 1. The annotated dataset is available on our web site (the Knowledge Engineering Research Group[1]) under the DATASETS link.

In the datasets, opinion words are considered to be either adjectives or adverbs and targets either nouns or pronouns. In the case of pronouns, they are denoted as targets, but for any pronoun the actual target is the product inferred. Pronouns are used to extract inferred product features along with their corresponding opinion words.

The identification of syntactic relations between opinion words and product features was performed by making use of a syntactic parser: Stanford CoreNLP[2], from which we use the fine-grained POS tags that help identify opinion words and targets. For example, comparative and superlative adjectives are more likely to be opinion words than other kind of adjectives. For inferring the polarities of the seed words we used SentiWordNet[3] as it offers both polarity and objectivity for each word, depending on its POS tag and *context*. To achieve seed words context independency, we compute the weighted average score for each adjective considering all the possible contexts.

Table 1: Dataset details.

File	Total Words Number	Opinion Words Number	Targets Number	Sentences Number
Apex	12081	401	358	739
Canon	11543	475	405	597
Coolpix	6501	498	359	346
Nokia	9292	504	277	546
Movie	5456	138	121	248

[1] http://keg.utcluj.ro
[2] http://nlp.stanford.edu/software/
[3] http://sentiwordnet.isti.cnr.it/

The evaluation of the opinion words and targets extraction is done using an algorithm we designed which using the annotations automatically calculates the recall and precision of the solution. The result is computed by comparing each extracted target and opinion word with each lemmatized annotated word.

In order to identify different occurrences of each extracted and annotated word instance, the sentence index of each word is used. The sentence index represents the number of the sentence it belongs to, based on its order of appearance. To ensure only extracted words are evaluated, the seed words are removed from the extraction process output before the opinion words are used by the evaluation algorithm. The pseudo-code for evaluating the opinion word extraction is the following:

```
   Input: Actual Opinion Words {A},
Found Opinion Words {F}
   Output: Precision {P}, Recall {R}
   Function:
   {TP } = 0, {FP} = 0, {FN} = 0
   For each opinion word {O} in {F}:
           if ({A} contains {O})
                 {TP} = {TP} + 1
           else {FP} = {FP} + 1
         endif
   endfor
   For each opinion word {O} in {A}:
           if ({F} does not contain
   {O})
                 {FN} = {FN} + 1
         endif
   endfor
   Set {P} = {TP} / ({TP} + {FP})
   Set {R} = {TP} / ({TP} + {FN})
```

4.1 Domain Independence Evaluation

The results of the tests conducted on reviews targeting different products along with the tests conducted on movie reviews, which have a different format and belong to a different domain, are presented in Figures 2 and 3, and prove the domain independence of the proposed solution. In Figure 2, the first column from each of the four-set clusters represents the results from tests conducted on product reviews using 6785 seed words. The second column corresponds to tests conducted on movie reviews with the same amount of seed words. The equivalent columns for tests using 2 seed words are the last two of each cluster. Note that the same solution configuration was used for both product and movie reviews (a polarity threshold of 0.01 and a target frequency threshold of 1.

There are generally two types of subjective texts, one which contains only text on topic, like product reviews, and another which is more descriptive in nature, like movie reviews which also describe the plot. In the description, opinions unrelated to the actual target of the subjective text can be conveyed, which affect the extraction process. This behavior can be seen in Figure 2 on the extraction of movie reviews using 6000+ seed words.

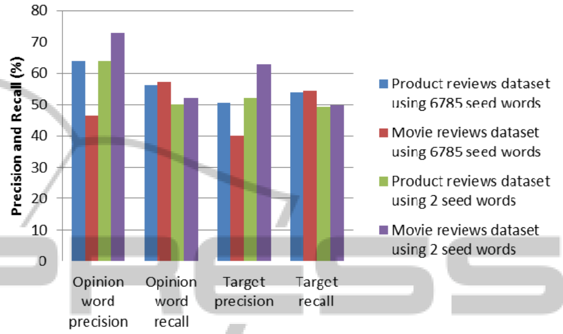

Figure 2: Cross domain evaluation (precision and recall) of opinion words and targets.

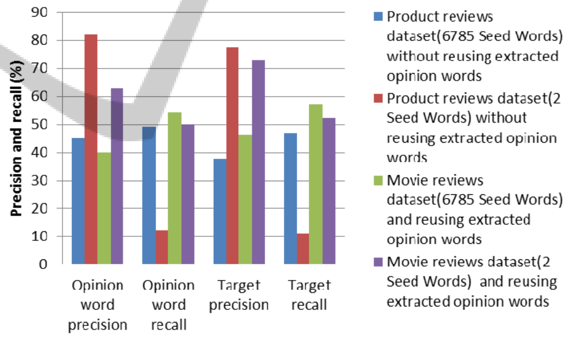

Figure 3: Influence of reusing opinion words as seed words.

The usage of only two seed words prevents this unwanted behavior, as the propagation is generally limited to related targets.

The dimension of the input data also affects the extraction process greatly when two seed words are used, as the propagation process performs poorly on a sparse data set, as can be seen in Figure 3, where the average results "without reuse" depicts the average precision and recall on 8 movie reviews, each of which contain an average of 25 opinion words and 22 targets. As can be seen on the results "with reuse", this issue is solved by reusing extracted opinion words from each text as seed words on all other texts belonging to the same domain, leading to a recall similar as when using a very large set of seed words.

4.2 Parameter Experiments and Tuning

The results of experimenting with the filtering threshold values can be seen in Figure 4. For a threshold value of 0.07, the precision increase outweighs the recall drop, but the best results are observed at a threshold value of 0.01. This is due to the fact that increasing the threshold value the number of opinion word omissions increase.

For pruning the targets, we experimented with various values of the occurrence frequency threshold, and the results can be seen in Figure 5.

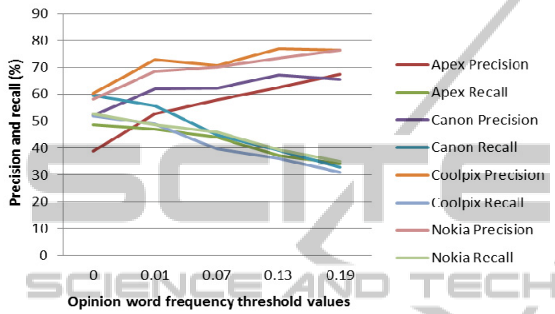

Figure 4: Opinion word polarity threshold influence on opinion word extraction results.

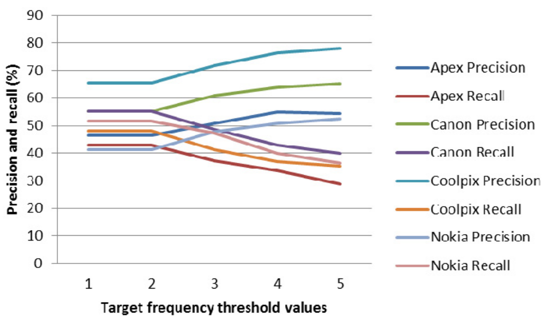

Figure 5: Target frequency threshold influence on target extraction results.

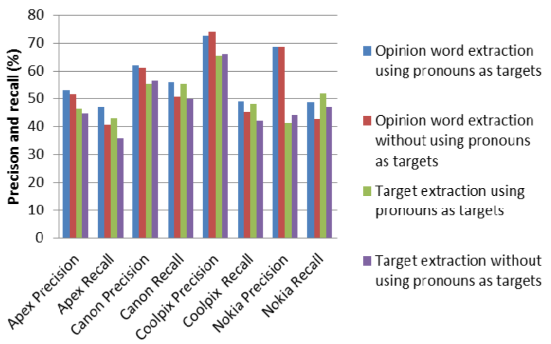

Figure 6: Influence of extracting pronouns as targets opinion word and target extraction.

In case of Figure 5, there is no best ratio of precision vs. recall with the increase in the target frequency threshold value, so the best value can be considered to be 1. Henceforth two best values for the opinion word polarity threshold and target frequency threshold are used, namely 0.01 and 1.

The rules used for extracting pronouns as targets do not have a significant impact the extraction precision for both opinion words and targets, but the increase in recall for both opinion word and target extraction is visible in Figure 6.

4.3 Seed Word Number Influence

One important finding in our experimental setting is that the number of seed words does not impact the extraction performance significantly, proven by the fact that by using only 2 seed words, i.e. *good* and *bad,* results similar to the ones using 6785 seed words were obtained. The small difference in the results presented in Figure 7 proves that no context dependent data is actually needed for a good performance. This behavior is explained by the following two facts: the number of reviews is sufficiently large; there is a high probability that the two – very common – words are used at least once to describe a product or one of its features. After at least one target is extracted, the iterative algorithm finds all the opinion words associated with it. The number of opinion words extracted in this case is close to the one found by using a very large set of seed words. Following this reasoning, we can safely state that this approach is unsupervised.

However, despite the low difference in the results induced by the number of seed words, there is a large difference in the extraction times. The number of seed words dramatically increases the processing time as can be seen in Figure 8. This is caused by the

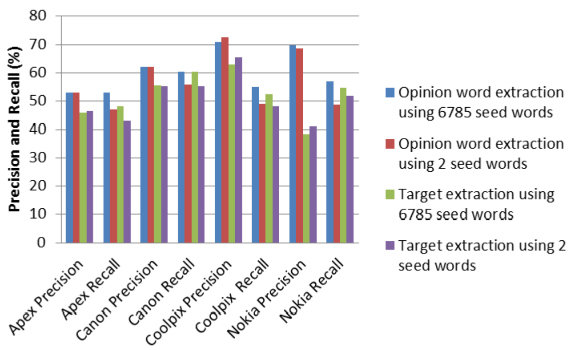

Figure 7: Seed words influence on opinion word and target extraction.

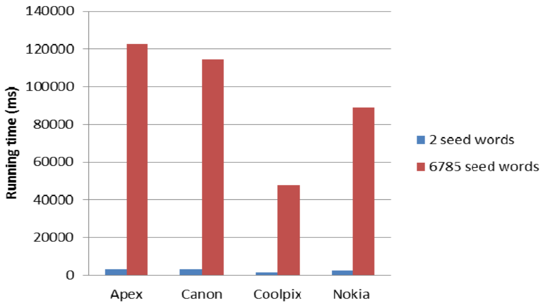

Figure 8: Run times in milliseconds based on seed words.

excessive number checks made by each rule on each possible opinion word. For extracting 400 opinion words using 2 seed words a maximum of 402 comparisons take place, but by using 6785 seed words, over 7000 comparisons are made (so, more than 1 order of magnitude).

4.4 Polarity Assignment Evaluation

The data set used for evaluating the polarity assignment consists of a selection of the first few hundred lines of textual data extracted from the *Nikon Coolpix* and *Canon G3* targeted reviews.

Initial experiments with polarity assignment were performed without taking into account the polarity consistency, that is, without averaging the scores for any of the targets. Precision in that case was just 53%. Applying the consistency rule, which averages the scores for the same target throughout the text, is justified as it improves precision and also evens out the distribution of polarity values.

Another issue that had to be tackled was achieving context independency for SentiWordNet polarity value retrieval. This is due to the fact that SentiWordNet contains multiple entries for the same word, each belonging to a different context and having a different polarity value. To fix this, the total score retrieved for a given word from SentiWordNet is the sum of the weighted averages of its occurrences. The weights decrease with the number of occurrences, as in (1), as suggested in SentiWordNet.

$$Score = \frac{1}{2} * \ First\ occurence + \frac{1}{3} * Second + \frac{1}{4} * Third + \cdots (1)$$

Further experimentation with the influence of other factors on the polarity assignment module is presented next. There are three factors that influence the precision of scoring: polarity threshold, score threshold and the number of seed words.

Figure 9 depicts the influence of the polarity threshold which has a big impact on polarity assignment precision as it has the power to sooth-out big variations in polarities and filter out inconsistent targets. Using somewhat big values, we can obtain 100% precision over non-smooth data sets. The optimal value for this value is determined to be around 0.2 for obtaining high precision values.

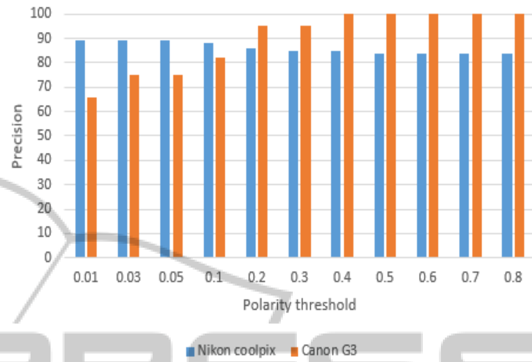

Figure 9: Influence of polarity threshold.

In Figure 10 we can see the influence of the score threshold value over the two sets of data. This threshold is necessary since true context independency is very hard to achieve and polarities tend to have variations even in the same context. Basically everything that falls whithin the value of this threshold is accepted. The polarity threshold was kept to 0.4 because this was the value for which one data set conveyed 100% precision, the target frequency was set to 2 and we used the maximum number of seed words. The optimal value was found to be 0.4. Note that a variation of 0.4 in a scale of 23 entries (-1 to +1) falls very much between most people's subjectivity measures.

Figure 11 depicts the influence of the number of seed words on the polarity assignment precision. This is by far the most interesting result and the most important one as our initial goal was to use just two seed words to obtain comparable results. It was obvious from the beginning that because of the applied rules that ensure polarity consistency

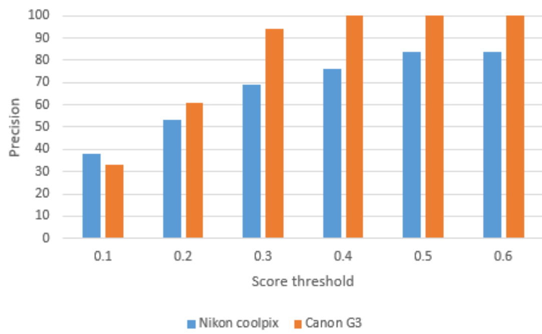

Figure 10: Influence of score threshold.

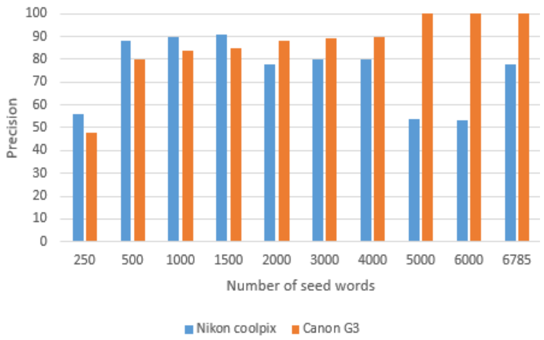

Figure 11: Influence of the number of seed words.

throughout the text using just two seed words was not possible since there must be at least one seed word for each major decimal value (0.1, 0.2, etc) and one for each decimal value in-between and so on.So theoretically, the more seed words the better, but this was not necessarily the case, as Figure 13 shows.

Because the *Nikon Coolpix* dataset contains opinion words conveying mostly the same polarities, using high numbers of seed words introduces noise by "over-averaging" polarities. So naturally, a more specific selection would be beneficial. This is not the case for the *Canon G3* dataset which contains diverse opinion words. The best compromise value is at around 1000-1500 words. Notably good results have been obtained using 500 words, out which just 10 were negative words. This is explainable by the fact that angry people tend to use the same negative words over and over again, while happy people tend to use a more elaborate vocabulary.

5 CONCLUSIONS

Our work devises a generalized methodology by considering a comprehensive set of grammar rules for better identification of opinion bearing words. We focused on creating a multidimensional configurable system for overcoming the domain-dependency issues which occur in all supervised opinion mining algorithms, by using only 2 class representative seed words. Using thorough experiments we discovered the optimal tradeoff between precision and recall, using the opinion polarity and target frequency thresholds. Furthermore, we proved that a larger amount of seed words does not yield a significant increase in recall or precision, making the approach unsupervised and domain independent.

Further work can include refining the extraction rules and increasing the preprocessing performance.

REFERENCES

Guang Qiu, Bing Liu, Jiajun Bu, Chun Chen 2012. Opinion Word Expansion and Target Extraction through Double Propagation. In *Computational Linguistics*, March 2011, Vol. 37, No. 1: 9.27.

Turney, Peter D. 2002. Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. In *Proceedings of ACL'02*, pages 417–424.

Hatzivassiloglou, Vasileios and Hathleen R. McKeown. 1997. Predicting the semantic orientation of adjectives. In *Proceedings of ACL'97,* pages 174-181. Stroudsburg, PA.

Hu, Mingqing and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of SIGKDD'04*, pages 168-177.

Popescu, Ana-Maria and Oren Etzioni. 2005. Extracting product features and opinions from reviews. In *Proceedings of EMNLP'05*, pages 339-346.

Brill, E. 1994. Some advances in transformation-based part of speech tagging. *Proceedings of the Twelfth National Conference on Artificial Intelligence* (pp.722-727). Menlo Park, CA: AAAI Press.

Jin, H. H. Ho, and R. K. Srihari, OpinionMiner: a novel machine learning system for web opinion mining and extraction, presented at the Proceedings of the 15thACM SIGKDD international conference on Knowledge discovery and data mining, Paris, France, 2009.

Htay, Su Su, and Khin Thidar Lynn, 2013. Extracting product features and opinion words using pattern knowledge in customer reviews. *The Scientific World Journal*.

Baccianella, Stefano, Andrea Esuli, and Fabrizio Sebastiani, 2010. SentiWordNet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining. *Seventh conference on Iternational Language Resources and Evaluation*.

Xia Hu, Jiliang Tang, Huiji Gao, 2013. Unsupervised Sentiment Analysis with Emotional Signals. *Proceedings of the 22nd international conference on World Wide Web*. 607-618.

Zhang, Lei, Bing Liu, Suk Hwan Lim, and Eamonn O`Brien-Strain, 2010. Extracting and Ranking Product Features in Opinion Documents. *International Conference on Computational Linquistics*. 1462-1470.

Marco Guerini, Lorenzo Gatti and Marco Turchi, 2013. Sentiment Analysis: How to Derive Prior Polarities from SentiWordNet. *arXiv Preprint,* arXiv:1309.5843.

Edison Marrese-Taylor, Juan D. Velasquez, Felipe Bravo-Marquez, 2013. OpinionZoom, a modular tool to explore tourism opinions on the Web. *ACM International Conferences on Web Intelligence and Intelligent Agent Technology*. 261-264.

Maite Taboada, Caroline Anthony and Kimberly Voll, 2006. Methods for Creating Semantic Orientation Dictionaries. *Proceedings of 5th International Conference on Language Resources and Evaluation (LREC)*. 427-432.

Christopher D. Manning, 2011. Part-of-Speech Tagging from 97% to 100%: Is It Time for Some Linguistics?. *Proceedings of the 12th International Conference on Computational Linguistics and Intelligent Text Processing,* 171-189.

Cosma Alexandru et all, 2014. Overcoming the domain barrier in opinion extraction. Accepted for publication at *10th International Conference on Intelligent Computer Communication and Processing.*