

Augmenting WordNet with Polarity Information on Adjectives

Alekh Agarwal and Pushpak Bhattacharyya

Dept. of Comp Science & Eng.

I.I.T. Bombay Mumbai 400076

alekh@cse.iitb.ac.in

Abstract

Polarity of a word refers to its strength in a classification, typically in a good *vs* bad sense, for example in movie reviews. This paper describes a technique to effectively compute the polarity information for *Adjectives*. Carrying on from this, we propose to introduce a new kind of link in *WordNet* and associate a polarity score with each Adjective in the *WordNet* database. We show the inter-dependence of subjectivity and polarity of a word. We demonstrate the need for incorporating such information in *WordNet*, by showing its use in the classification of sentences as *subjective* and *objective*.

1 Introduction

By the *Polarity* of a word, we mean the extent to which it contributes in determining the sentiment or tone of a phrase or a sentence in which it is contained. For example, *excellent* and *outstanding* are words with a strong positive polarity. *Poor* and *bad* are examples of words with a high negative polarity. In the middle, we have words like *about* and *often* which are neutral, *i.e.*, lacking polarity.

A closely associated notion is that of subjectivity. Subjective phrases are those that are used to describe emotions or opinions rather than being statements of fact. The words that are frequently used in such phrases are referred to as subjective words. Words like *great*, *stunning* or *gory* are examples of subjective words. On the other hand, *general* and *other* are examples of objective words.

Recently the field of Sentiment and Subjectivity Analysis has received much attention [Turney, 2001b; Pang and Lee, 2004; Pang et al., 2002]. Typically, and intuitively too, adjectives have been found to be strong indicators of the sentimental content of a document. So it is not very surprising that many of the approaches presented for these tasks tend to rely on the notion of the strength of an adjective, and on a more general note, the sentiment contained in a phrase, in various ways. This is especially highlighted in the works of Turney [Turney, 2001b], Vegnaduzzo [Vegnaduzzo, 2004] and Baroni *et. al.* [Baroni and Vegnaduzzo, 2004].

With the increasing use of these techniques, the need has been felt for the incorporation of polarity and subjectivity information in important lexical databases like *WordNet* [Baroni and Vegnaduzzo, 2004]. This would enable applications to access polarity information in a much faster and convenient way.

A very significant use of polarity values is that clusters of *isopolar* adjectives can be further used to create clusters

of *isopolar* sentences. In applications aiming to capture the overall sentiment of a document, this is a very important step because strongly polar sentences have a far greater impact on the tone of the document than the mildly polar ones. Thus extracting a cluster of strongly polar sentences can enable a Sentiment Analysis application to evaluate the sentiment of a document by looking at a very small portion of the whole document. For such applications, just the *synonymy* information is far from sufficient because two sentences can use completely different adjectives to express the same sentiment. In such cases, *isopolarity* is a much better criterion for grouping together sentences using word level information than mere *synonymy*.

In this paper we present a technique to effectively extract polarity related information from the web by generating search engine queries and using the number of search results obtained in an Information Theory based measure for computing word polarity. We then propose to enhance the *WordNet* link structure by adding *isopolarity* links to connect adjectives with a similar degree of polarity. We discuss how useful it is to incorporate the actual polarity scores for every adjective in the *WordNet* database. We also demonstrate how the notions of polarity and subjectivity are closely associated at the word level. Then we go ahead and use this observation to determine the degree of subjectivity of complete sentences using the polarity scores of the adjectives contained in them.

The organization of the rest of the paper is as follows: section 2 is a brief review of the previous work. We describe our approach in section 3. Section 4 is on experimental setup. The experiments and results are presented in section 5. We conclude the paper in section 6.

2 Previous Work

Polarity and Subjectivity of words have been studied previously under two categories of work. In the first category, the aim was to determine the effective metrics of polarity and subjective content of words ([Vegnaduzzo, 2004; Baroni and Vegnaduzzo, 2004]). In the second, such study forms part of larger study on sentiment analysis [Turney, 2001b]. Our work lies in the first category, but we also use the information extracted to determine the subjectivity of sentences.

Initially, the task of assigning polarity scores was mainly performed by linguistic experts. However, such an approach suffers from experts' biases and can typically be done only to the extent of partitioning the words into a few categories rather than assigning ratings to every word on a continuous scale. A landmark theory on assigning polarity scores to

adjectives in various kinds of classification tasks is that of Charles Osgood [C.E. Osgood and Tannenbaum, 1957]. A technique for effective determination of adjective polarities using this theory and the *WordNet* Synonym graph was presented in [Jaap Kamps and de Rijke, 2004]. However, due to the structure of this *WordNet* graph, they could determine these ratings for only about 25% adjectives in the *WordNet* database. Moreover, this technique fails to generalize for other Parts-of-Speech.

The work closest to ours is perhaps the set of studies conducted by Turney on the use of PMI-IR technique for the mining of polarity scores for adjectives from the web [Turney, 2001b], [Turney, 2001a]. Our paper largely builds on this and then demonstrates how to use this information to augment *WordNet*.

3 Our Method

3.1 Determining Polarity scores of Adjectives

We used a Mutual Information based measure for computing the polarity scores of adjectives. The problem of data sparsity for computation of such a statistical measure is always a great bother. However, [Turney, 2001b] introduced the concept of using the Web for determination of these scores using queries made to a Search Engine.

The PointWise Mutual Information between two words is defined as

$$PMI(w_1, w_2) = \log_2 \frac{P(w_1, w_2)}{P(w_1)P(w_2)} \quad (1)$$

This is nothing but the ratio of the actual probability two words being seen together to the probability of their being seen together if their occurrence was independent of each other. Hence, thought in an intuitive way, it reflects the degree of association between the two words.

Now let us consider two words of strongly opposite polarity. In our case we took them to be *excellent* and *poor*. We define the function Polarity Score for a word w , $PS(w)$ as

$$PS(w) = PMI(w, excellent) - PMI(w, poor) \quad (2)$$

Now as evident from common sense and pointed out in [Baroni and Vegnaduzzo, 2004], words with similar polarity will have a higher rate of co-occurrence. This arises due to the fact that when we appreciate or criticize anything, we seldom use just one adjective to express our emotion. For example, take this arbitrary sentence picked up from a movie review: "This film is serious, real, and most of all BELIEVABLE. This is both why it's amazing and fails to elicit a "strong" emotional attachment. Almost too serious." A whole lot of strongly polar words with similar polarities like *serious*, *real*, *strong* and *amazing* can be seen in a very close vicinity.

This is the main intuition behind the above formula for polarity scores. It is worth noting here that the [Justeson and Katz, 1991] describe a phenomenon of *antonym co-occurrence* that is contrary to our hypothesis. This phe-

nomenon says that antonyms tend to occur in a close proximity in a sentence or a group of sentences with a probability greater than that of random co-occurrence. However, the empirical success of PMI approach and results of other works relying on a hypothesis similar to ours like [R. Mandala and Hozumi, 1998] provide sufficient reason to believe in this assumption.

Consider any positively polar word. As per the above statements, its rate of co-occurrence and hence PMI with the word *excellent* will be higher than that with *poor*. Thus a higher value of $PS(w)$ reflects a higher positive polarity. Note that the subtraction of PMI values with respect to two oppositely polar anchor words is done to ensure that common words that tend to co-occur frequently with both the anchor words get a low rating. Taking just PMI scores with respect to just one word will assign arbitrarily high scores to such words which is not desirable.

As an example, consider the adjective *usual*. We know that *usual* hasn't got any strongly polar orientation. However, the PMI scores of *usual* with the anchor words *excellent* and *poor* are -31.45 and -30.83 respectively. Looking at either of these large negative values in isolation would give us a wrong idea regarding the polarity of *usual*. For example, a large negative PMI score wrt *excellent* would indicate a strong negative polarity. However, taking the difference results in the Polarity Score of -0.614 which is a value that indicates no great orientation on either side. Thus taking the difference of PMI values is essential to obtain a rational metric for adjective polarity.

To determine the PMI values for adjectives, we used the method of gathering statistics through search engine queries as described by [Turney, 2001b]. In this method $PMI(w_1, w_2)$ is measured as:

$$PMI(w_1, w_2) = \log_2 \frac{Hits(w_1.AND.w_2)}{Hits(w_1)Hits(w_2)} \quad (3)$$

So our effective formula for $PS(w)$ reduces to:

$$PS(w) = \log_2 \frac{Hits(w.AND.excellent)Hits(poor)}{Hits(excellent)Hits(w.AND.poor)} \quad (4)$$

Here the AND operator is to ask the Search Engine to look for simultaneous occurrences of both w_1 and w_2 on the *same page*. This implies that our effective window to check co-occurrence is a complete web page. Although, [Turney, 2001b] had used the NEAR operator in his work, which is known to give better results than the AND operator [Baroni and Vegnaduzzo, 2004], AltaVista has stopped supporting the NEAR operator now. So we had to use the AND operator. These values were retrieved for each of the 21436 adjectives in *WordNet's* database with 2 queries made per word. For example, to find out the Polarity Score for the word *ideal*, the queries "*ideal*" AND "*excellent*" and "*ideal*" AND "*poor*" were issued. The number of hits for these queries along with the number of hits for the anchor words were used to compute Polarity Scores using equation 4.

It is important to point out here that although many other statistical measures like Latent Semantic Analysis have also been proposed for this task, PMI has been shown to outperform them [Turney, 2001a]. Its ease of calculation is of course an added advantage.

We experimented with some other measures ourselves and tried to take the effect of $Hits(excellent, poor)$, i.e., the number of co-occurrences of the two anchor words into account. The functions that we tried are:

$$f_1(w) = \frac{\frac{Hits(w.AND.excellent)}{Hits(excellent)} - \frac{Hits(w.AND.poor)}{Hits(poor)}}{Hits(excellent, poor)} \quad (5)$$

and

$$f_2(w) = N \frac{\frac{Hits(w.AND.excellent)}{Hits(excellent)Hits(w)} - \frac{Hits(w.AND.poor)}{Hits(poor)Hits(w)}}{Hits(excellent, poor)} \quad (6)$$

However, the results with PMI scores were much better than these measures.

3.2 Incorporation of Polarity Information in WordNet

Once we had the polarity scores for the adjectives, we ran a k -means clustering algorithm with the Calinski and Harabasz’s stopping rule [Calinski and Harabasz, 1974] to select the optimum number of clusters. In this rule, we use a Variance Ratio Criterion (VRC) defined as:

$$VRC = \frac{\frac{BGSS}{k-1}}{\frac{WGSS}{n-k}} \quad (7)$$

where,

$WGSS$ is the Within-Cluster Sum of Squared

Distances about the Centroids,

$BGSS$ is the total between cluster Sum of

Squared Distances,

k is the number of clusters, and

n is the number of data points.

These clusters contain words having similar polarities. Hence we propose to link these words with a new kind of link called the *isopolarity* link in *WordNet*. We also propose to store with each word, its polarity score calculated as in the previous section.

Here, we would like to point out the adjective polarity scores are not exactly context and domain independent. Different domains have differences in the common vocabularies used and accordingly different adjectives are used for varying degrees of polarity. However, the nature of polarity, i.e., positive or negative remains the same. Even the degrees of polarity do not usually show drastic changes across domains. Our approach averages out the strength across various domains in the sense that when we query the web, we are likely to find documents from all domains, and hence, the effect from no single domain would be in-genuinely pronounced. However, if domain specific information is desired for some application, then the set of weights will need to be altered accordingly for that application. An easy way for this would be to use a large corpus from that particular domain rather than the web for PMI evaluation.

3.3 Detection of Subjective Content of Adjectives

For this task, we employed the technique described in [Baroni and Vegnaduzzo, 2004]. We took the list of 35 seeds that [Baroni and Vegnaduzzo, 2004] and [Vegnaduzzo, 2004] had used earlier and randomly picked 10 adjectives out of it. We then calculated the PMI score for each adjective in *WordNet*’s database that had a familiarity count greater than 3.

We calculated these PMI scores with each of the 10 seed words. The formula used for PMI computation was a little different in this case. We used the following expression:

$$PMI(w1, w2) = \log_2 N \frac{Hits(w1.AND.w2)}{Hits(w1)Hits(w2)} \quad (8)$$

The extra factor of N in the numerator was taken as 1 billion, the approximate number of documents indexed by AltaVista. This expression comes from taking the Maximum Likelihood Estimators for the probabilities. Although the constant factor is not of great importance, since it is the relative scores that matter, the values with $N = 1$ were too small to be worked with.

Then we took the *first quartile* of these scores (which had given the best results in Baroni’s study [Baroni and Vegnaduzzo, 2004]) and recorded it as the subjectivity rating of the word. These scores were used to observe the correspondence between polarity and subjectivity ratings for adjectives.

3.4 Using the Adjective Polarity values to compute Sentence Subjectivity

A Support Vector Machine(henceforth SVM) based classifier was used. We used bag-of-words features. The top 1000 adjectives and 5000 non-adjective words were chosen as features on the basis of their occurrence in the dataset. With the non-adjectival features, binary values were used in the feature vectors.

For adjectival features, we tried two different approaches. One was to use the polarity score of an adjective if the given adjective was found in the document, 0 otherwise. In the second case we used binary values just like other features. The five-fold cross validation accuracies were recorded for both the cases.

The basic idea behind using polarity score in the feature vector for adjectival features was that the value in the feature vector corresponding to a feature represents the strength or weight of the feature. It’s like using a scaling factor along a particular axis in the vector space model. So when we use polarity scores, we are in effect, telling the algorithm how much importance it should associate with every single occurrence of that adjective. This is definitely closer to the human way of analyzing sentences and polarities than just having binary values in feature vectors.

4 Evaluation

4.1 Experimental Setup

Figure 1 contains a block diagram of the experimental setup.

We used the following softwares and tools:

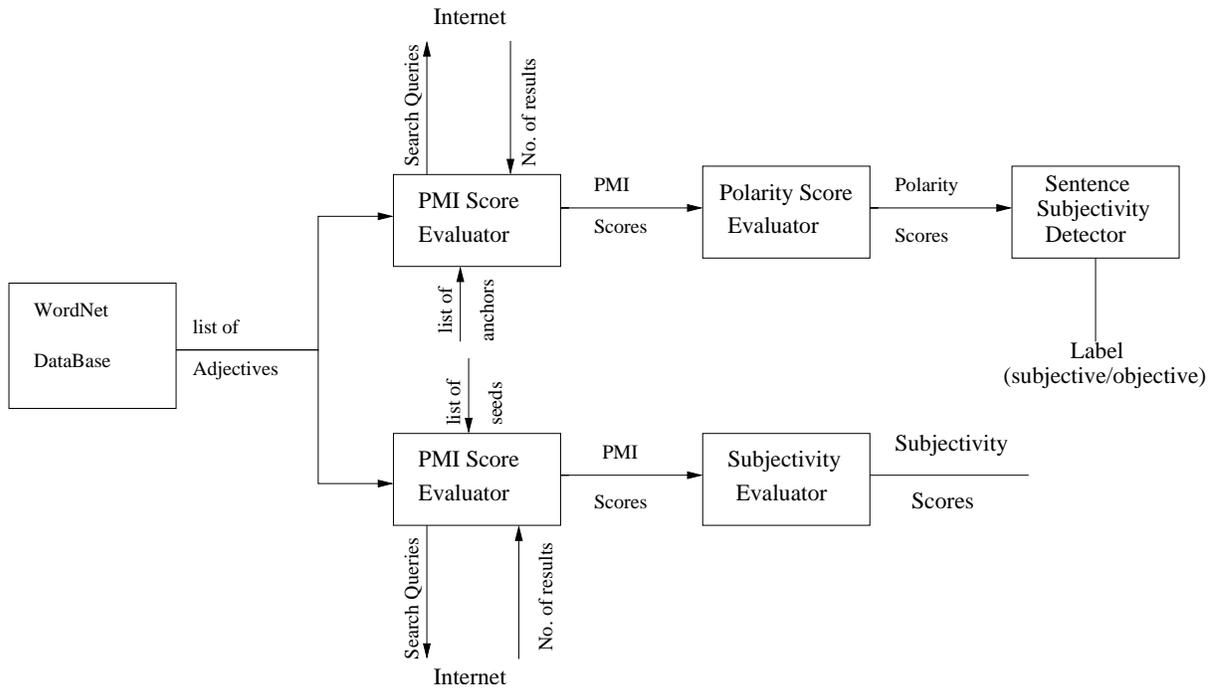


Figure 1: Flow Diagram

wget : This linux utility was used to retrieve webpages corresponding to search queries generated for words.

WordNet::QueryData : This package¹ was used to query *WordNet* for information like list of all adjectives in *WordNet* and word familiarity counts.

Stanford Log-Linear Model Tagger v1.0 : This POS Tagger² was used to tag sentences from the subjectivity corpus.

libsvm-2.71 : An implementation of SVM [Chang and Lin, 2001].

cluster-1.29 : A C Clustering Library by M.J.L. de Hoon *et al.* [M.J.L.de Hoon and S.Miyano, 2004] used to obtain *isopolar* groups.

The set of seeds for determining the subjectivity score for adjectives was derived from the list of 35 seeds used by [Baroni and Vegnaduzzo, 2004] and [Vegnaduzzo, 2004].

The sentence corpus used for Subjective *vs* Objective classification was the subjectivity corpus introduced by Bo Pang *et al.* in ACL 2004 [Pang and Lee, 2004].³ This corpus contains 5000 plot summaries and 5000 review snippets that were collected from <http://www.rottentomatoes.com>.

5 Experimental Results

First of all, we retrieved a list of all the adjectives in *WordNet's* lexical database. With each of the adjectives we

¹<http://people.csail.mit.edu/~jrennie/WordNet>

²<http://www-nlp.stanford.edu/software/tagger.shtml>

³Corpus available at www.cs.cornell.edu/people/pabo/movie-review-data (review corpus version 2.0)

performed queries to measure their PMI scores with respect to *excellent* and *poor*. Thus a total of 42874 queries was performed in this step.

When the Polarity scores obtained through this process were analyzed, we realized that the words with negative scores had greatly outnumbered those with positive scores. Also the words with very large scores on either side were mostly obscure words. On performing k-means clustering on this, we obtained the optimum number of clusters as 3. But the words were largely cluttered up into two clusters, both with their centroids on the negative side. We further tried an Expectation Maximization clustering on these scores. Here cross-validation yielded 7 clusters as the optimum value. However, words were still mainly cluttered in two clusters, both with negative mean and together accounting for over 90% adjectives in the *WordNet's* database.

We realized that the scores for the obscure words were not very good indicators of their polarities as their co-occurrence with one of the anchor word was usually too low leading to an abnormally high score on either end of the spectrum. These incorrect readings presented difficulties for the clustering algo downstream by adding a huge number of noisy points.

To take care of the above problem, we extracted the adjectives in the *WordNet's* database which had a familiarity count greater than 3. This gave 745 words. A cursory glance was enough to realize that all the obscure noise words had been filtered by this step.

We went ahead to extract the polarity scores for these words. With the noise words removed, the scores were much easier to analyze and now a clear correlation between polarities and scores could be seen. We went ahead to apply

the k-means clustering on this set. Of these there were two large clusters with 493 and 236 words. The third cluster was a small one with just 16 words. Except the anchor word *poor* none of the words in the cluster had a high degree of polarity but had large negative ratings.

However, the more important clusters are the other two. The 436 word cluster has majority of the words having a positive polarity. (Note that here the polarity refers to the actual polarity of words as we know it to be from our knowledge of english, not on the basis of polarity scores. On the basis of scores, this cluster has all the words with positive scores. But some positively polar words end up with negative scores which is a classification error that'll be always present in any unsupervised learning method).

The 236 word cluster contains most of the words with negative polarity. The cluster id and polarity scores for some of the common adjectives have been listed in Table 1.

Table 1: Polarity Scores and Clusters for a few Common Words

	Word	Polarity Score	Cluster
1.	comparative	-3.00	0
2.	poor	-2.91	0
3.	legislative	-2.81	0
4.	democratic	-2.51	0
5.	political	-1.77	0
6.	slow	-0.74	1
7.	passive	-0.72	1
8.	late	-0.71	1
9.	monotonous	-0.67	1
10.	highest	-0.15	1
11.	accurate	-0.14	1
12.	efficient	-0.13	1
13.	good	-0.12	1
14.	outstanding	0.56	1
15.	unique	0.57	1
16.	spectacular	0.59	1
17.	sporting	1.34	1
18.	raw	-1.58	2
19.	crude	-1.36	2
20.	cruel	-1.35	2
21.	furious	-0.99	2
22.	fierce	-0.99	2
23.	bad	-0.98	2
24.	dangerous	-0.96	2
25.	scattered	-0.79	2
26.	uneven	-0.78	2
27.	lesser	-0.78	2
28.	lacking	-0.77	2
29.	irregular	-0.76	2
30.	lost	-0.76	2

In the table we can clearly see how the words having their polarity scores very close, have a close association among their polarities. For example we can clearly spot the high similarity in the polarities of the words in the following word groups: (slow, passive, late); (highest, accurate, effi-

cient, good); (outstanding, unique, spectacular). Other such groupings can also be seen.

Although, it maybe pointed out that the words in a cluster vary over a wide range of polarities, but this is due to the wide range of polarity scores contained in clusters. If smaller clusters are desired, all one has to do is to increase the number of clusters in the clustering step. We kept the number of clusters to 3 as it was the optimum number predicted by the Calinski and Harbasz's stopping rule and also it provided the basic separation between positively and negatively polar words.

The other metrics described earlier that we used for measuring word polarity didn't give very good results. The values in both cases did not have a good correlation with word polarity and the results for clustering were also bad in these cases.

After this we took up the task of analyzing the correlation between word polarity and subjectivity scores. The subjectivity scores were evaluated using the technique described in [Baroni and Vegnaduzzo, 2004]. The details are as given in section 3.2. The results were encouraging. Few of the strongly polar words at the top of subjectivity list are *frightfull*, *clumsy*, *enlightened* and *sympathetic*. Of the top 25 subjective words, only 8 had the polarity scores lying in the range of -0.5 to 0.5 . This number was 19 in the first 50 words. 36 such words were present in the top 100 subjective words. Please note that the values -0.5 and 0.5 were handpicked and better analysis can be done if the optimum threshold for polarity determination is learned using supervised learning.

The final task we took up was to see the effect of polarity scores on determination of subjectivity of sentences using the approach described in section 3.3. The accuracy using binary values in the feature vector for both adjectival and non-adjectival features was found out to be 55.93%. On using polarity scores instead of a binary value for adjectival features, the accuracy rose to 61.1%.

6 Conclusions and Future Work

The incorporation of polarity information in *WordNet* is definitely needed taking into account the increasing number of works that rely on such information. Our approach shows that this is definitely feasible using automated techniques. Small misclassifications do occur and human intervention would be needed at times, if a very high precision is desired. However, our work largely suffices to provide the basic infrastructure to do such a task.

One may point out here that a drawback of our approach was the need to cut the number of adjectives down to 745 for eliminating noise words. We would like to mention, that as far as the individual scores are concerned, this approach can fetch them even for the words that we have eliminated. The real trouble is caused at the clustering step. So a possible option is to keep the individual scores for these words in *WordNet* but to omit the *isopolarity* links for them. One more thing that we would like to point out is that the value of familiarity count greater than 3 was chosen randomly and

the scope of clustering can be increased. Real problem is due to the words that have familiarity counts of 0. And also, as we have been saying from the beginning, the whole aim of this work is to make the polarity data available for practical applications that heavily rely on such information. So when we discuss the omission of *isopolarity* links for words with a low familiarity count, we are not really affecting the users of the polarity information in the sense that such words are not very common. The strength of our approach is that the common and frequent words will always be reliably linked with a fairly high degree of accuracy.

Here, we would also like to contrast our results against those of [Jaap Kamps and de Rijke, 2004], which had tried to show the association between *WordNet* graph and polarity scores of adjectives. We pointed out earlier that their approach is not able to find the polarity scores for all words. A casual observer might comment that this drawback is there in our method also. However, there is a very important difference. Their approach relied on the *WordNet* synonym link structure. So if a word is not connected to the anchor word through synonymy links, then its score cannot be evaluated. This leads to some very common adjectives like *excellent* being left out using *good* and *bad* as the two anchor words. Our approach can always calculate the polarity scores for any word and the reliability of clustering increases for common words. So the case of important words that might be frequently needed by applications being left out in our approach would hardly ever arise. Another important distinction is that our technique is not specific to adjectives or even unigrams for that matter.

The dependence between subjectivity and polarity that is reflected in our work provides an interesting line of work in future. The gain in accuracy of classification on using polarity scores is also heartening and better ways of incorporating this information in the classification technique is bound to improve the results.

Future work: A very interesting line for future work would be to try and obtain the polarity information separately for different senses of a word. This would require using the facts that there should be a very high correlation between polarity scores of the words in a *Synset*. *WordNet* glosses for the *synsets* can also be used to determine these scores.

Also other techniques to obtain the *isopolar* groups from the polarity scores can be experimented with. A very big drawback is the withdrawal of the NEAR operator of AltaVista. Looking into other techniques and sources to obtain this information can also be done.

Another thing we have pointed out earlier is that polarity scores are not evenly distributed about 0. The scores vary from -3.0 to 1.34 . A good measure here might be to try and apply supervised learning techniques to determine optimum threshold value. This task can be done using a small list of adjectives tagged with polarities.

The dependence between subjectivity and polarity can also be used in strengthening the data for one if the data for the other is available.

Acknowledgments

We would like to thank Marco Baroni for providing us with the tagged list of adjectives and set of seeds used in his work.

References

- M. Baroni and S. Vegnaduzzo. 2004. Identifying subjective adjectives through web-based mutual information. *Proceedings of KONVENS 2004, Vienna: ĀŪGAI*. 17–24.
- T. Calinski and J. Harabasz. 1974. A dendrite method for cluster analysis. *Communications in Statistics*.
- G.J.Succi C.E. Osgood and P.H. Tannenbaum. 1957. *The Measurement of Meaning*. University of Illinois Press, Urbana IL.
- Chih-Chung Chang and Chih-Jen Lin, 2001. *LIBSVM: a library for support vector machines*.
- Maarten Marx Jaap Kamps, Robert J. Mokken and Maarten de Rijke. 2004. Using wordnet to measure semantic orientation of adjectives. In *Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC 2004), volume IV, pages 1115–1118*.
- John S. Justeson and Slava M. Katz. 1991. Co-occurrences of antonymous adjectives and their contexts. *Comput. Linguist.*, 17(1):1–19.
- J. Nolan M.J.L.de Hoon, S. Imoto and S.Miyano. 2004. Open source clustering software. *Bioinformatics* 20(9): 1453-1454.
- Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *ACL*, pages 271–278.
- Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? sentiment classification using machine learning techniques. volume cs.CL/0205070.
- R. Takenobu R. Mandala and T. Hozumi. 1998. The use of wordnet in information retrieval. *Proc. of Coling-ACL*.
- Peter D. Turney. 2001a. Mining the web for synonyms: Pmi-ir versus lsa on toefl. pages 491–502.
- Peter D. Turney. 2001b. Thumbs up or thumbs down?: semantic orientation applied to unsupervised classification of reviews. In *ACL '02: Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, pages 417–424, Morristown, NJ, USA. Association for Computational Linguistics.
- S. Vegnaduzzo. 2004. Acquisition of subjective adjectives with limited resources. *AAAI Spring Symposium Technical Report: Exploring Affect and Attitude in Text*.