Procesamiento del Lenguaje Natural, Revista nº 51, septiembre de 2013, pp 179-186 recibido 30-04-2013 revisado 16-06-2013 aceptado 21-06-2013

Improving Subjectivity Detection using Unsupervised Subjectivity Word Sense Disambiguation

Mejoras en la Detección de Subjetividad usando Desambiguación Semántica del Sentido de las Palabras

Reynier Ortega
Adrian Fonseca
Center for Pattern Recognition
and Data Mining
Santiago de Cuba, Cuba

Yoan Gutiérrez
Andrés Montoyo
DLSI,University of Alicante
Alicante, Spain

Resumen: En este trabajo se presenta un método para la detección de subjetividad a nivel de oraciones basado en la desambiguación subjetiva del sentido de las palabras. Para ello se extiende un método de desambiguación semántica basado en agrupamiento de sentidos para determinar cuándo las palabras dentro de la oración están siendo utilizadas de forma subjetiva u objetiva. En nuestra propuesta se utilizan recursos semánticos anotados con valores de polaridad y emociones para determinar cuándo un sentido de una palabra puede ser considerado subjetivo u objetivo. Se presenta un estudio experimental sobre la detección de subjetividad en oraciones, en el cual se consideran las colecciones del corpus MPQA y Movie Review Dataset, así como los recursos semánticos SentiWordNet, Micro-WNOp y WordNet-Affect. Los resultados obtenidos muestran que nuestra propuesta contribuye de manera significativa en la detección de subjetividad.
Palabras clave: detección de subjetividad, desambiguación semántica, análisis de sentimiento

Abstract: In this work, we present a sentence-level subjectivity detection method, which relies on Subjectivity Word Sense Disambiguation (SWSD). We use an unsupervised sense clustering-based method for SWSD. In our method, semantic resources tagged with emotions and sentiment polarities are used to apply subjectivity detection, intervening Word Sense Disambiguation sub-tasks. Through an experimental study, we empirically validated the proposed method over two subjectivity collections, MPQA Corpus and Movie Review Dataset, using three widely popular opinion-mining resources SentiWordNet, WordNet-Affect and Micro-WNOp. The results show that our proposal performs significantly better than our proposed baseline.
Keywords: subjectivity detection, subjective word sense disambiguation, sentiment analysis

1 Introduction

Subjectivity detection consists in identifying whether a phrase, word or sentence is used to express opinion, emotion, evaluation, speculation, etc., (Wiebe and Riloff, 2005). It besides contributes in many Natural Language Processing (NLP) tasks. For instance, Information Retrieval systems incorporate subjectivity detection to provide opinionated and factual information, separately (Pang and Lee, 2008); and Question Answering systems increase their performances when using criteria for discrimination among types of factual versus opinionated questions (Lloret et al., 2011). On the other hand, Summarization systems pretends to resume factual and subjective content differently (Murray and Carenini, 2008).

Motivated by the usability and applicability of this task, some researchers have pro-

ISSN 1135-5948

posed methods for deal with Subjective Detection Resolution (SDR). Many approaches rely on lexicons[1] of words that may be used to express subjectivity. These approaches do not make distinction between different senses of a word, so terms included in such lexicons are treated as subjective regardless of their sense. Moreover, most subjectivity lexicons are compiled as keyword lists, rather than word meanings. However, many keywords have both subjective and objective senses, depending on the context where the corresponding word appears.

Recent approaches have proposed to profit from Word Sense Disambiguation (WSD) in subjective analysis. It could be either by adding semi-automatically subjectivity tags to annotated senses in WSD corpora, or training a supervised classifier to determine whether a word is being used in a subjective sense or not, without explicitly tagging senses. The WSD uses in this area have been necessaries to know the context meaning to provide a better efficiency SDR.

In this paper, we propose using subjectivity annotated resources to solve the SDR, unlike previous approaches, which depend heavily on manual or semiautomatic annotation for training supervised classifiers. We use an unsupervised strategy consisting in a coarse-grained clustering-based WSD method that differentiates objective, subjective and highly subjective uses of every word, and classify sentences as subjective or objective. Our method is able to integrate the affective usabilities of SentiWordNet "SWN" (Baccianella, Esuli, and Sebastiani, 2010), Micro-WNOp "WNOp" (Cerini and Gandini, 2007) and WordNet-Affect"WNA" (Valitutti, 2004) to resolve Subjective Analysis Task.

The paper is organized as follows. We review related works in Section 2. Section 3 is dedicated to describing our approach, whereas Section 4 contains the descriptions and results analysis of the conducted experiments. Finally, we present in Section 5 our conclusions and further works.

2 Related Work

Methods for subjectivity detection span a wide range of viewpoints. An early work proposed by Hatzivassiloglou and Wiebe (2000) examined the effects of adjective orientation and gradability on sentence subjectivity. Its goal has been to determine whether a given sentence is subjective or not, judging from the adjectives involved in current sentence. An attempt to classify subjective and objective sentences have been introduced in (Riloff and Wiebe, 2003), which explores syntactic pattern extraction using semi-supervised learning.

Other works have focused in annotating senses with emotion labels or polarity values. For instance, a WordNet (Miller, 1995) extension has been presented by (Valitutti, 2004), where every sense is annotated with one of the six basic emotion labels "anger", "happiness", "surprise", "digust", "sadness" and "fear". Esuli and Sebastiani (2006) determine the polarity of word senses in WordNet, distinguishing among positive, negative and objective. They manually annotate a seed set of positive/negative senses and by following the internal relations in WordNet expand a small set using a supervised approach. They extend their work (Baccianella, Esuli, and Sebastiani, 2010) by applying the PageRank algorithm for ranking the WordNet senses in terms of how strongly a sense possesses a given semantic property (e.g., positive or negative).

A large number of works have applied WSD in sentiment analysis for instance, Rentoumi et al. (2008) determine the polarity by disambiguating the words and then mapping the senses to models of positive and negative polarity. To compute these models and produce the mappings of senses, they adopt a graph-based method which takes into account contextual and sub-word information. Similarly to earlier work, Martín-Wanton et al. (2010a) exploits full word sense disambiguation for determining the correct sense of a word and assigning polarity using SentiWordNet and General Inquirer (Stone et al., 1966). Martín-Wanton et al. (2010b) study the behavior of SWN, WNA, and WNOp in polarity detection.

Recently, Akkaya, Wiebe, and Mihalcea (2009) introduced Subjectivity Word Sense Disambiguation (SWSD), which consists in automatically determining which word instances in a corpus are being used in subjective senses, and which are being used in objective senses. They use a supervised system for SWSD, and exploit the SWSD output

[1] Are a stock of words used in a particular profession, subject or domain.

to improve the performance of multiple contextual opinion analysis tasks. Akkaya et al. (2010) carried out a pilot study where a subjectivity sense-tagged dataset was created for eight SENSEVAL[2] words through MTurk[3], a web-based non-expert manual annotation interface.

These works have focused in creating new datasets for subjectivity contextual analysis by using existing polarity classification resources. They all rely heavily on manual or semiautomatic annotation for training supervised classifiers.

3 *Our Proposal*

As we mentioned previously, we use an unsupervised strategy consisting in a coarse-grained clustering-based WSD method that differentiates objective, subjective and highly subjective uses of every word.

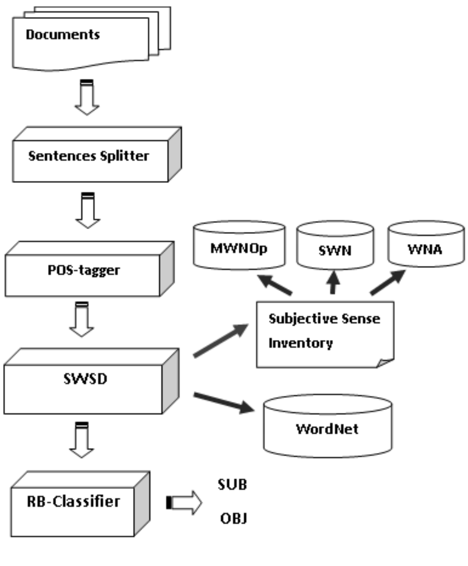

Figure 1: Overall architecture of contextual subjective classifier.

In this work, we evaluate the behavior of our proposal using different opinion mining resources. The overall architecture of our contextual subjective classifier is shown in Figure 1.

Firstly, the text is segmented into sentences, lemmatized and POS-tagged using TreeTagger tool (Schmid, 1994) being removed the stopwords. Then a Subjectivity Word Sense Disambiguation (SWSD) method is applied to content words (nouns, adjectives, verbs and adverbs). Once all content words are disambiguated (Section 3.2), we apply a rule-based classifier (Section 3.3) to decide whether the sentence is subjective or objective.

The use of SWSD in our proposal is motivated by considerations exposed in (Akkaya, Wiebe, and Mihalcea, 2009), where they explain that a same word may be used subjectively or objectively in different contexts. For example, the word *"earthquake"* is used in a subjective sense in the sentence:

"Selling the company caused an earthquake among the employees".

Whereas it is used in an objective sense in the sentence:

"An earthquake is the result of a sudden release of energy in the Earth's crust that creates seismic waves".

We adapted the unsupervised word sense disambiguation method proposed by Anaya-Sánchez, Pons-Porrata, and Berlanga-Llavori (2006) which is based on clustering as a manner of identifying related senses, for SWSD. Unlike the authors, who aim at obtaining the correct sense of a word; we use the method to determine when a word is subjective or objective relying on a subjective sense inventory. We constructed subjective senses inventories based on affective and polarity annotations in opinion mining resources.

3.1 Subjective Sense Inventories

Creating subjective sense-tagged data is a hard and expensive task. For this reason, we decided to use existent sense-level resources for fine-grained and coarse-grained subjective sense labeling. We considered three different resources for building our subjective sense inventory: SWN, WNOp and WNA. These resources have not explicit subjectivity labels; therefore we mapped polarity or affect labels to subjectivity labels.

SWN and WNOp contain positive, negative and objective scores between 0 and 1. In this case the mapping was defined in the following manner: senses whose sum of positive and negative scores is greater than or equal to 0.75 are considered to be highly subjective (HS); whereas those whose sum is lower than 0.75 and greater than or equal to 0.5 are con-

[2] http://www.senseval.org/
[3] http://mturk.amazon.com

sidered to be subjective (S). In the remaining cases, the senses are considered to be objective (O). In WNOp is important to clarify that this resource only contains 1105 WordNet senses annotate manually, the remainder WordNet senses were considered as objective.

In WNA, the senses are annotated with emotion labels. In order to match these labels with ours, we apply a similar strategy to (Balahur et al., 2009). Here, senses labeled with the following emotions: "anger", "disgust" and "surprise", are considered as highly subjectives. Others like "guilt", "sadnes" and "joy" are considered as subjectives; and the rest are considered as objectives. In Table 1, we show the distribution of the subjectivity labels assigned for each resource.

Resources	**HS**	**S**	**O**
SentiWordNet	1766	6429	107229
Micro-WNOp	216	118	115090
WordNet-Affect	110	148	115166

Table 1: Senses highly subjective (HS), subjective(S) and objective (O) distributions by resources.

For all three resources, exists a notable unbalance between the number of objective and subjective senses, which is particularly strong in the case of WNA.

Once tagged the sense with subjectivity label, these are grouped for building the coarse-grained sense. For instance, considering the following adjective, *"sad"*, using SentiWordNet, this adjective has three word senses in WordNet 2.0, from which we can obtain its lemma, part-of-speech, sense offset (id), definition and subjective label assigned.

i. sad#a#1 – *experiencing or showing sorrow or unhappiness* – (HS)

ii. sad#a#2 – *of things that make you feel sad* – (O)

iii. sad#a#3 – *bad; unfortunate* – (HS)

As we can see, the first and third senses are considered as highly subjective and second is considered as objective. These considerations were taken using the defined mapping above. For this reason sense 1 and 3 are merged in only one sense representing an highly subjective unique sense, keeping sense 2 as objective sense.

3.2 Adaptations introduced in the WSD

As we expressed, the selected disambiguation method was developed for the traditional WSD task. In this WSD method, the senses are represented as topic signatures (Lin and Hovy, 2000) built from WordNet concept repositories. The disambiguation process starts from a clustering distribution of all possible senses of the ambiguous words by applying the Extended Star clustering algorithm (Gil-García, Badía-Contelles, and Pons-Porrata, 2003). Such clustering tries to identify cohesive groups of word senses, which are assumed to represent different meanings for the set of words. Then, clusters that best match with the context are selected. If the selected clusters disambiguate all words, the process stops and the senses belonging to the selected clusters are interpreted as the disambiguating ones. Otherwise, the clustering is performed again (regarding the remaining senses) until a complete disambiguation is achieved.

Thus, it does not distinguish between highly subjective, subjective and objective senses. We propose two strategies to adapt this method for the task at hand. The first strategy is based in fine-grained WSD. It consists in applying the original WSD method (Anaya-Sánchez, Pons-Porrata, and Berlanga-Llavori, 2006) and searching the subjective sense inventory for subjectivity labels for each fine-grained sense.

The second strategy is based on coarse-grained WSD. Many authors (Chan and Ng, 2007; Navigli, Litkowski, and Hargraves, 2007) have demonstrated that WSD methods increase their performance by using coarse-grained senses. In this paper, we defined new coarse-grained senses. All highly subjective senses of a word are collapsed into a single sense, as well as all subjective senses. On the other hand, objective senses are kept separated.

For our selected WSD method, word senses are represented by means of topic signatures (Lin and Hovy, 2000). The topic signature for coarse-grained senses is the sum of the topic signature of the corresponding original fine-grained senses. To take again the example in section 3.1 referring the adjective *"sad"*, it represents an instance of coarse topic signatures, where first and third senses, were grouped in a coarse-grained sense (highly subjective HS). The signature for this new sense is obtained in following manner:

$Tsign(sad\#a\#HS) = Tsign(sad\#a\#1) + Tsign(sad\#a\#3)$
$Tsign(sad\#a\#O) = Tsign(sad\#a\#2)$

Where $Tsign(sense_i)$ compute the related topic signature with the $sense_i$.

3.3 Subjective Sentence Classifier

We use a rule-based classifier to classify sentences into subjective or objective. A voting scheme is used. Every word disambiguated as highly subjective has assigned a score of 4 and every word disambiguated as subjective has assigned a score of 2. If the sum of all scores is greater than a threshold, the sentence is classified as subjective. This method is similar to that proposed by (Riloff, Wiebe, and Wilson, 2003). Equation 1 is used to classify a new sentence:

$$RL\left(f\right) = \begin{cases} subjective \quad if \quad \sum_{i=1}^{n} Score\left(w_i\right) \geq \lambda \\ \\ objective \quad e.o.c \end{cases} \tag{1}$$

Where:

$$Score\left(w_i\right) = \begin{cases} 4.0 \quad if \quad w_i \quad is \quad high \quad subjective \\ 2.0 \quad if \quad w_i \quad subjective \\ 0.0 \quad if \quad w_i \quad objective \end{cases} \tag{2}$$

In both equations (1, 2) w_i is the sense which the word is using in the sentence f. In our proposal the threshold used was $\lambda = 4.0$. This value was estimated using empirical evaluation over a subset of the SemCor corpus for English, being it automatically annotated with OpinionFinder tool by (Carmen Banea and Hassan, 2008). Thus, we employ this rule-based classifier with the aim to obtain an unsupervised method to classify sentences in subjective or objective categories.

4 Result and Discussion

We conducted a series of experiments in order to evaluate the validity of our proposal. The aim has been focused on the impact of using different resources for constructing the sense inventories to solve the SWSD.

In our experiments, we use two collections of subjectivity detection: the manually annotated MPQA Corpus (Wilson, 2005) and the automatically annotated collection over movie domain, Movie Review Dataset (Pang and Lee, 2004).

MPQA Corpus contains news (for version 1.2 contains 11115 sentences) where opinions are spreaded at sentence level. They are annotated with sentiment polarities and its respective strength value. In order to experimenting our SWSD method, we used the approach presented by Riloff and Wiebe (2003), obtaining 8026 subjective and 3089 objective sentences respectively.

Movie Review Dataset covers the movie domain. It contains 5000 subjective sentences extracted from movie reviews collected from the Rotten Tomatoes web site, and 5000 objective sentences collected from movie plot summaries from the Internet Movie Database (IMDB). The underlying assumption is that all the snippets from the Rotten Tomatoes pages are subjective (as they come from a review site), while all sentences from IMDB are objective (as they focus on movie plot descriptions).

In order to constructing a baseline which to evaluate the effect of applying SWSD on each resource, we use the same classification scheme, but without applying word sense disambiguation.

The polarity score of the words were defined as the average of the positive and negative score sum, for all associated senses to each word. All words with a new score above 0.75 were tagged as highly subjective, the words with score in the range 0.5 and 0.75 were tagged as subjective, and the rest were tagged as objective.

As score measures we computed precision, recall and F1 for both subjective and objective classes, moreover we compute the average of the F1 over subjective and objective sentences.

The behavior of all variants are shown in Tables 2 and 3. In these tables we can observe that except for one case, the classification is improved when is used both forms of SWSD, respect to the baseline ([resource] without WSD). The improvement is higher than 25% in both collections when we use WNOp (see Tables 2 and 3). This fact confirms our prior hypothesis that taking into account the individual subjectivity levels of different senses

Strategy	Ps	Po	Rs	Ro	Fs	Fo	F1-avg
SWN with Fine WSD	0.9305	0.270	0.5677	0.7903	0.7052	0.4225	**0.5538**
SWN with Coarse WSD	**0.9322**	0.2861	0.6052	0.7823	0.7339	0.4190	**0.5765**
SWN without WSD	0.8996	0.3309	0.8483	0.3710	0.8558	0.3498	**0.6043**
WNOp with Fine WSD	0.**9305**	0.2700	0.5677	0.7903	0.7052	0.4025	0.**5538**
WNOp with Coarse WSD	0.9237	0.2533	0.5334	0.7823	0.6763	0.3826	**0.5295**
WNOp without WSD	0.9302	0.1744	0.06525	0.9758	0.1220	0.2958	0.2089
WNA with Fine WSD	1.0	0.1710	0.0196	1.0	0.0384	0.2921	**0.1653**
WNA with Coarse WSD	1.0	0.1710	0.0196	1.0	0.0384	0.2921	**0.1653**
WNA without WSD	0.9091	0.1694	0.0163	0.9919	0.0321	0.2894	0.1607

Table 2: Experimental evaluation using MPQA Corpus.

Strategy	Ps	Po	Rs	Ro	Fs	Fo	F1-avg
SWN with Fine WSD	0.6303	0.5204	0.7745	0.35	0.6950	0.4185	**0.5568**
SWN with Coarse WSD	0.6203	0.4988	0.7734	0.3226	0.6884	0.3918	**0.5401**
SWN without WSD	0.6066	0.5268	0.89064	0.1742	0.7218	0.2618	0.4918
WNOp with Fine WSD	**0.6303**	0.5204	0.7745	0.35	0.6950	0.4185	**0.5568**
WNOp with Coarse WSD	0.6267	0.5011	0.7475	0.3629	0.6818	0.4210	**0.5514**
WNOp without WSD	0.5044	0.4046	0.0643	0.9097	0.1140	0.9097	0.5118
WNA with Fine WSD	0.7436	0.4245	0.0.981	0.9516	0.1733	0.5871	**0.3802**
WNA with Coarse WSD	**0.7479**	0.4251	0.1003	0.9516	0.1769	0.5876	**0.3821**
WNA without WSD	0.6724	0.4148	0.04397	0.9694	0.0825	0.5810	0.3317

Table 3: Experimental evaluation using Movie Review Dataset.

of a word, it may helping in SDR. Surprisingly, very small differences are observed between fine-grained and coarse-grained SWSD variants. We suppose that this situation is due to the high correlation among sense clusters obtained by the WSD method and those manually defined for the coarse-grained variant.

On the other hand the results using SWN are higher than the rest, whereas are lower those obtained using WNA. We may observe that as the unbalance between subjective and objective senses is higher, less accuracy have the obtained results with this resource. In case of WNA, we should note additionally that the mapping established between affect categories and subjectivity labels does not reflect all circumstances under which words are used subjectively.

In case of WNOp, despite being significantly smaller than SWN and suffering an unbalance between objective and subjective senses, when it is used for SWSD, the obtained results are similar to the obtained when SWN is used. This fact is encouraging, as it suggests that a small resource with a high-quality of annotated data is able to perform at the same level than annotated resources much bigger. A further exploration could be required to determine if the growing of WNOp may resulting in an improving of the performance of our proposal.

5 *Conclusion and Further Works*

In this work, we have presented an unsupervised SWSD-based approach to subjectivity detection, which relies on sense-level polarity and emotion-labeled resources. We conduct an experimental study, where the behavior of our proposed method is evaluated using three widely used resources: SWN, WNOp and WNA. As a result of our experiments, we show that subjectivity detection using our unsupervised SWSD-based approach outperforms a baseline where disambiguation techniques are not used. Besides, we obtain a characterization of the method's behavior using different resources, and remarking that SWN and WNOp are the most suitable for the task.

On the other hand, in order to find out other ways to obtain semantic labels of coarse-grained, we will adapt our method to the use (Gutiérrez, Vázquez, and Montoyo, 2011) proposal, which is able to obtain relevant domains associated to the sentences, where these domains involve polarity values.

Another attractive direction for future works is determining the influence of subjectivity-annotated resources, rather than approximating a subjectivity annotation from existing polarity or affect annotations.

Acknowledgments

This research work has been partially funded by the Spanish Government through the project TEXT-MESS 2.0 (TIN2009-13391-C04), "Análisis de Tendencias Mediante Técnicas de Opinión Semántica" (TIN2012-38536-C03-03), and SAM - Dynamic Social & Media Content Syndication for 2nd Screen (FP7-611312); and by the Valencian Government through the project PROMETEO (PROMETEO/2009/199).

References

Akkaya, Cem, Alexander Conrad, Janyce Wiebe, and Rada Mihalcea. 2010. Amazon mechanical turk for subjectivity word sense disambiguation. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*, CSLDAMT 10, pages 195–203, Stroudsburg, PA, USA. Association for Computational Linguistics.

Akkaya, Cem, Janyce Wiebe, and Rada Mihalcea. 2009. Subjectivity word sense disambiguation. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1 - Volume 1*, EMNLP '09, pages 190–199, Stroudsburg, PA, USA. Association for Computational Linguistics.

Anaya-Sánchez, Henry, Aurora Pons-Porrata, and Rafael Berlanga-Llavori. 2006. Word sense disambiguation based on word sense clustering. In *Proceedings of the 10th Ibero-American Conference on AI 18th Brazilian conference on Advances in Artificial Intelligence*, IBERAMIA-SBIA '06, pages 472–481, Berlin, Heidelberg. Springer-Verlag.

Baccianella, Stefano, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In Nicoletta Calzolari (Conference Chair), Khalid Choukri, Bente Maegaard, Joseph Mariani, Jan Odijk, Stelios Piperidis, Mike Rosner, and Daniel Tapias, editors, *Proceedings of the Seventh International Conference on Language Resources and Evaluation*, LREC '10, Valletta, Malta, may.

Balahur, Alexandra, Ralf Steinberger, Erik van der Goot, Bruno Pouliquen, and Mijail Kabadjov. 2009. Opinion mining on newspaper quotations. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology - Volume 03*, WI-IAT '09, pages 523–526, Washington, DC, USA. IEEE Computer Society.

Carmen Banea, Rada Mihalcea, Janyce Wiebe and Samer Hassan. 2008. Multilingual subjectivity analysis using machine translation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Honolulu, Hawaii.

Cerini, S., Compagnoni V. Demontis A. Formentelli M. and G Gandini. 2007. Language resources and linguistic theory: Typology, second language acquisition, english linguistics (forthcoming), chapter micro-wnop: A gold standard for the evaluation of automatically compiled lexical resources for opinion mining.

Chan, Yee Seng and Hwee Tou Ng. 2007. Word sense disambiguation improves statistical machine translation. In *In 45th Annual Meeting of the Association for Computational Linguistics*, ACL '07, pages 33–40.

Esuli, Andrea and Fabrizio Sebastiani. 2006. Sentiwordnet: A publicly available lexical resource for opinion mining. In *In Proceedings of the 5th Conference on Language Resources and Evaluation*, LREC '06, pages 417–422.

Gil-García, R. J., J. M. Badía-Contelles, and A. Pons-Porrata. 2003. Extended star clustering algorithm. *Lecture Notes on Computer Sciences*, 2905:480–487.

Gutiérrez, Yoan, Sonia Vázquez, and Andrés Montoyo. 2011. Sentiment classification using semantic features extracted from wordnet-based resources. In *Proceedings of the 2nd Workshop on Computational Approaches to Subjectivity and Sentiment*

Analysis (WASSA 2.011), pages 139–145, Portland, Oregon, June. Association for Computational Linguistics.

Hatzivassiloglou, Vasileios and Janyce M. Wiebe. 2000. Effects of adjective orientation and gradability on sentence subjectivity. pages 299–305.

Lin, Chin-Yew and Eduard Hovy. 2000. The automated acquisition of topic signatures for text summarization. In *Proceedings of the 18th conference on Computational linguistics - Volume 1*, COLING '00, pages 495–501, Stroudsburg, PA, USA. Association for Computational Linguistics.

Lloret, Elena, Alexandra Balahur, Manuel Palomar, and Andrés Montoyo. 2011. Towards a unified approach for opinion question answering and summarization. In *Proceedings of the 2nd Workshop on Computational Approaches to Subjectivity and Sentiment Analysis*, WASSA '11, pages 168–174, Stroudsburg, PA, USA. Association for Computational Linguistics.

Martín-Wanton, Tamara, Aurora Pons-Porrata, Andrés Montoyo-Guijarro, and Alexandra Balahur. 2010a. Opinion polarity detection - using word sense disambiguation to determine the polarity of opinions. In *ICAART (1)*, pages 483–486.

Martín-Wanton, Tamara, Aurora Pons-Porrata, Andrés Montoyo-Guijarro, and Alexandra Balahur. 2010b. Word sense disambiguation in opinion mining: Pros and cons. In *Research in Computing Science*, pages 119 – 129.

Miller, George A. 1995. Wordnet: A lexical database for english. *Communications of the ACM*, 38:39–41.

Murray, Gabriel and Giuseppe Carenini. 2008. Summarizing spoken and written conversations. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, EMNLP '09, pages 773–782.

Navigli, Roberto, Kenneth C. Litkowski, and Orin Hargraves. 2007. Semeval-2007 task 07: coarse-grained english all-words task. In *Proceedings of the 4th International Workshop on Semantic Evaluations*, SemEval '07, pages 30–35, Stroudsburg, PA, USA. Association for Computational Linguistics.

Pang, Bo and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the ACL*.

Pang, Bo and Lillian Lee. 2008. Opinion mining and sentiment analysis.

Rentoumi, Vassiliki, George Giannakopoulos, Vangelis Karkaletsis, and George A. Vouros. 2008. Sentiment analysis of figurative language using a word sense disambiguation approach.

Riloff, Ellen and Janyce Wiebe. 2003. Learning extraction patterns for subjective expressions. In *Conference on Empirical Methods in Natural Language Processing*, EMNLP '03, pages 105–112.

Riloff, Ellen, Janyce Wiebe, and Theresa Wilson. 2003. Learning subjective nouns using extraction pattern bootstrapping. pages 25–32.

Schmid, Helmut. 1994. Probabilistic part-of-speech tagging using decision trees.

Stone, Philip J., Dexter C. Dunphy, Marshall S. Smith, and Daniel M. Ogilvie. 1966. *The General Inquirer: A Computer Approach to Content Analysis*. MIT Press, Cambridge, MA.

Valitutti, Ro. 2004. Wordnet-affect: an affective extension of wordnet. In *In Proceedings of the 4th International Conference on Language Resources and Evaluation*, pages 1083–1086.

Wiebe, Janyce and Ellen Riloff. 2005. Creating subjective and objective sentence classifiers from unannotated texts. In *In Proceedings of the 6th International Conference on Intelligent Text Processing and Computational Linguistics*, CICLing '05, Mexico City, Mexico.

Wilson, Theresa. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *In Proceedings of HLT-EMNLP*, pages 347–354.