Combining Supervised and Unsupervised Polarity Classification for non-English Reviews

José M. Perea-Ortega, Eugenio Martínez-Cámara,
María-Teresa Martín-Valdivia, and L. Alfonso Ureña-López

SINAI Research Group, Computer Science Department, University of Jaén
Escuela Politécnica Superior, Campus Las Lagunillas s/n, 23071, Jaén, Spain
{jmperea,emcamara,maite,laurena}@ujaen.es

Abstract. Two main approaches are used in order to detect the sentiment polarity from reviews. The supervised methods apply machine learning algorithms when training data are provided and the unsupervised methods are usually applied when linguistic resources are available and training data are not provided. Each one of them has its own advantages and disadvantages and for this reason we propose the use of meta-classifiers that combine both of them in order to classify the polarity of reviews. Firstly, the non-English corpus is translated to English with the aim of taking advantage of English linguistic resources. Then, it is generated two machine learning models over the two corpora (original and translated), and an unsupervised technique is only applied to the translated version. Finally, the three models are combined with a voting algorithm. Several experiments have been carried out using Spanish and Arabic corpora showing that the proposed combination approach achieves better results than those obtained by using the methods separately.

1 Introduction

Opinion Mining (OM), also known as Sentiment Analysis (SA) is a challenging task that combines data mining and Natural Language Processing (NLP) techniques in order to computationally treat subjectivity in textual documents [1]. This new area of research is becoming more and more important mainly due to the growth of social media where users continually generate contents on the web in the form of comments, opinions, emotions, etc. There are several issues related to OM like subjectivity detection, opinion extraction, irony detection and so on. However, perhaps the most widely-studied task is sentiment polarity classification. This task aims to determine which is the overall sentiment-orientation (positive or negative) of the opinions contained within a given document. The document contains subjective information such as product reviews or opinionated posts in blogs.

Although different approaches have been applied to polarity classification, the mainstream basically consists of two major methodologies. On the one hand, the Machine Learning (ML) approach (also known as the supervised approach) is based on using a collection of data to train the classifiers [2]. On the other hand,

A. Gelbukh (Ed.): CICLing 2013, Part II, LNCS 7817, pp. 63–74, 2013.

the approach based on Semantic Orientation (SO) does not need prior training, but takes into account the positive or negative orientation of words [3]. This method, also known as the unsupervised approach, makes use of lexical resources like lists of opinionated words, lexicons, dictionaries, etc. Both methodologies have their advantages and drawbacks. For example, the ML approach depends on the availability of labeled data sets (training data), which in many cases are impossible or difficult to achieve, partially due to the novelty of the task. On the other hand, the SO strategy requires a large amount of linguistic resources which generally depend on the language, and often this approach obtains lower recall because it depends on the presence of the words comprising the lexicon in the document in order to determine the orientation of opinion. In order to overcome the weaknesses of both approaches, we have performed several experiments, combining ML and SO through different strategies.

Most of the studies on polarity classification only deal with English documents, perhaps due to the lack of resources in other languages. However, people increasingly comment on their experiences, opinions, and points of views not only in English but in many other languages. Consequently, the management and study of subjectivity and SA in languages other than English is a growing need. The work presented herein is mainly motivated by the need to develop polarity detection systems in languages other than English.

According to Mihalcea, Banea and Wiebe [4], there are two main approaches in the context of multilingual SA. The first one is a Lexicon-based approach, where a target-language subjectivity classifier is generated by translating an existing lexicon into another idiom. The second one is a Corpus-based approach, where a subjectivity-annotated corpus for the target language is built through projection, training a statistical classifier on the resulting corpus. In this paper we follow this second approach and we generate an English parallel corpus by applying machine translation to the original corpus.

The aim of this study is to evaluate an approach based on the combination of supervised and unsupervised methods to improve the results obtained using these methods separately. Specifically, this study has been carried out on two different corpora of reviews in Arabic and Spanish. The main idea is to translate the original corpus into English, generating a parallel corpus. Thus, we could apply the supervised approach to the original corpus and the unsupervised one to the translated version of the original corpus, since it is more feasible to find linguistic resources for this language. languages that have few lexical resources for tackling the polarity classification problem.

The rest of the paper is organized as follows: the next section presents work related to polarity detection dealing with languages other than English and multilingual opinion mining. Section 3 presents the approach proposed in this work. Section 4 describes the different resources used in our experiments including the MC and MCE corpora and SentiWordNet. The different experiments carried out and the results obtained are expounded in Section 5. In Section 6 the obtained results are analyzed. Finally, the main conclusions and ideas for further work are expounded in Section 7.

2 Background

Most of the research papers on SA that we can find in the literature have been applied to English exclusively, although works on other languages are growing increasingly. There are some interesting papers that have studied the problem of polarity classification using non-English collections such as German, French, Chinese, Arabic or Spanish. Below, we summarize some of the most interesting related works.

Kim and Hovy [5] compared opinion expressions between an aligned corpus of emails in German and English. They developed two models: for the first one they translated German emails into English and then applied opinion-bearing words. For the second one they translated English opinion-bearing words into German and then analyzed the German emails using the German opinion-bearing words. The results showed that the first model worked slightly better than the second one. Following this work, Denecke [6] worked on German comments collected from Amazon. These reviews were translated into English using standard machine translation software. Then the translated reviews were classified as positive or negative, using three different classifiers: LingPipe, SentiWordNet with classification rule, and SentiWordNet with machine learning.

Tan and Zhang [7] were among the first researchers to study opinion mining in Chinese. They carried out a widely experimental revision using lots of different models. Zhang et al. [8] applied Chinese SA on two datasets. In the first one, euthanasia reviews were collected from different web sites, while the second dataset was about six product categories collected from Amazon (Chinese reviews). They proposed a rule-based approach including two phases: firstly, by determining each sentence's sentiment based on word dependency, and secondly, by aggregating sentences in order to predict the document sentiment. Wan [9] studied the sentiment polarity identification of Chinese product reviews using a semantic orientation. He made use of bilingual knowledge including both Chinese resources and English resources. The corpus was composed of 886 Chinese documents that were translated into English by using Google Translate and Yahoo Babel Fish. In addition, the approach used ensemble methods to combine the individual results over Chinese and English datasets. The results for the combination methods improved the performance of individual results.

Ghorbel and Jacot [10] used a corpus with movie reviews in French. They applied a supervised classification combined with SentiWordNet in order to determinate the polarity of the reviews. French is also managed in Balahur and Turchi [11], along with Spanish and German. Different machine translation systems and meta-classifiers were tested in order to demonstrate that multilingual SA using these techniques is comparable to the English performance.

In Rushdi-Saleh et al. [12] a corpus of movies reviews in Arabic annotated with polarity was presented and several experiments using machine learning techniques were performed. Subsequently, they generated the parallel EVOCA corpus (English version of OCA) by translating the OCA corpus automatically into English. The results showed that, although the results obtained with EVOCA were worse than those obtained with OCA, they are comparable to other English experiments, since the loss of precision due to the translation process is very slight, as can be seen in Rushdi-Saleh et al. [13].

Regarding opinion mining focused on Spanish, there are also some remarkable studies. For example, Banea et al. [14] proposed several approaches to cross lingual subjectivity analysis by directly applying the translations of opinion corpus in English to training an opinion classifier in Romanian and Spanish. This study showed that automatic translation is a viable alternative for the construction of resources and tools for subjectivity analysis in a new target language. Brooke et al. [15] presented several experiments dealing with Spanish and English resources. They concluded that although the ML techniques can provide a good baseline performance, it is necessary to integrate language-specific knowledge and resources in order to achieve an improvement. Finally, Cruz et al. [16] generated the MuchoCine corpus by recollecting manually Spanish movie reviews from the MuchoCine website. This corpus was generated in order to develop a sentiment polarity classifier based on semantic orientation. On the other hand, Martínez-Cámara, Martín-Valdivia and Ureña-López [17] applied the supervised approach to the MuchoCine corpus using different ML algorithms, obtaining better results than those obtained by applying the unsupervised approach proposed by Cruz et al.

One of the drawbacks for the investigation in SA over non-English texts is the lack of linguistic resources. In Steinberger et al. [18] is presented a novelty method to develop multilingual and comparable sentiment dictionaries, which consists of using two high-level gold-standard sentiment dictionaries for two languages (English and Spanish) and then translated them automatically into third languages. The third languages dictionaries are formed by the overlap of the translations, i.e. via triangulation. The obtained dictionaries are manually filtered and expanded.

3 Combination of Supervised and Unsupervised Methods

The aim of the approach proposed in this study is to improve the polarity classification of the reviews provided by a corpus whose documents are in a language other than English. The main proposal is to translate the original corpus into English and work with parallel corpora, generating several learning models by using both corpora. Furthermore, since we have a corpus translated into English, we can make use of semantic resources for opinion mining tasks such as SentiWordNet[1] in order to apply a non-supervised approach to that corpus. In this way, the models (supervised and unsupervised) generated using the parallel corpora can be combined in a meta-classifier that could apply different algorithms to establish the final polarity classification. Figure 1 illustrates this approach.

One of the advantages of our architecture is its modularity, allowing the use of different supervised algorithms for both corpora (original and translated) and even in the meta-classifier, for combining previous generated models. As can be seen in Figure 1, we apply a processing to the corpora, which usually consists

[1] `http://sentiwordnet.isti.cnr.it`

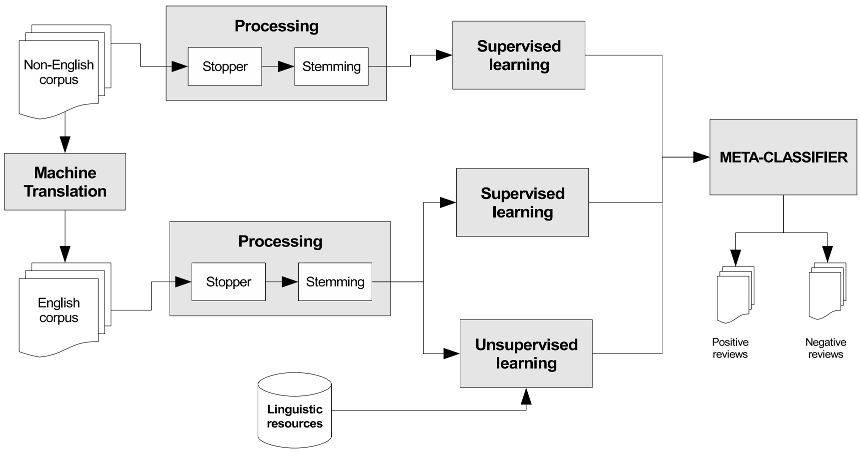

Fig. 1. Overview of the approach proposed

of a stemming process for extracting the root of each word after removing the words without semantic meaning (stopwords).

Once the corpora were processed, we generated the learning models that were used later in the meta-classifier. The supervised approach was applied to both corpora using different learning algorithms such as SVM or NB. However, the unsupervised approach was applied solely to the translated corpus because the linguistic resources, such as SentiWordNet or WordNet-Affect[2], are available in English only. Finally, the meta-classifier process combined several features from the supervised and unsupervised models previously generated, allowing to apply different combination algorithms.

The approach proposed in this paper is especially suitable when we work with non-English corpora because using the translated version of the original corpus we could apply unsupervised approaches on it, since there are very few linguistic and semantic resources for non-English corpora. In this way, we could improve the results obtained by using the supervised methods and to gain some independence from the domain.

4 Experiment Framework

In order to verify the performance of the proposed approach, we decided to apply it on two non-English corpora, specifically on the MuchoCine corpus in Spanish and the OCA corpus in Arabic. In this section we explain the main tools used in carrying out the experiments presented in this study. Then, we describe both corpora employed for the experiments.

[2] `http://wndomains.fbk.eu/wnaffect.html`

For the processing carried out to the parallel corpora we used the RapidMiner[3] tool, which allows to apply the stopper and stemming for different languages. The supervised approach was also performed using this tool, since it allows to apply the cross-validation method using different learning algorithms such as Support Vector Machines (SVM) or Naïve Bayes (NB).

Regarding the unsupervised approach, we used SentiWordNet 3.0 [19] as semantic resource. SentiWordNet (SWN) is a lexical resource for SA which assigns three sentiment scores to each *synset* of WordNet[4]: positivity, objectivity and negativity. Each of the scores ranges from 0 to 1, and their sum is equals 1. A good example is the word *beautiful*, which belongs to two *synsets* (00217728, 01800764). For the *synset* 00217728, the SWN score of *beautiful* is $(0.75, 0.25, 0)$ and for the synset *01800764* is $(0.625, 0.375, 0)$. We used nouns, adjectives, verbs and adverbs as linguistic features. In a first step, the translated corpus was processed by applying a POS tagger like TreeTagger[5]. The aim of this process was to obtain all the nouns, adjectives, verbs and adverbs of each review. The second step after tagging the translated corpus was to generate a total of 15 sub-corpora by making a combination of the four possibilities (nouns, adjectives, verbs and adverbs) in order to analyze the impact of each type of word. Finally, we calculated the SWN score for each review as the polarity score of the document. This score was obtained following the method proposed by Denecke [6] based on the calculation of a triplet of positivity, negativity and objectivity scores.

Below, we explain the main features of the both parallel corpora used for the experiments carried out in this study.

4.1 The OCA-EVOCA Corpus

The Arabic corpus called OCA (Opinion Corpus for Arabic) was generated by Rushdi-Saleh et al. [12] to be freely used for the research community related to OM[6]. It is composed of 500 film reviews that were extracted from different blogs in Arabic found on the Internet. 250 reviews were labeled as positive and the other 250 as negative. In Rushdi-Saleh et al. [12] can be found more details about the process of generation of OCA and its evaluation carried out by applying the cross-validation method.

The same authors conducted the machine translation of OCA into English, generating the parallel corpus called EVOCA (English Version of OCA), also available for research purposes[7]. This translation was carried out using the PROMT[8] tool. In Rushdi-Saleh et al. [13] can be found the evaluation performed on the EVOCA corpus also using the cross-validation method.

3 http://rapid-i.com

4 WordNet is a large lexical database of English. Nouns, verbs, adjectives and adverbs are grouped into sets of cognitive synonyms (*synsets*), each expressing a distinct concept. It is available in http://wordnet.princeton.edu

5 http://www.ims.uni-stuttgart.de/projekte/corplex/TreeTagger

6 http://sinai.ujaen.es/wiki/index.php/OCA_Corpus_(English_version)

7 http://sinai.ujaen.es/wiki/index.php/EVOCA_Corpus_(English_version)

8 http://translation2.paralink.com

4.2 The MC-MCE Corpus

The MuchoCine corpus (MC) was presented in Cruz et al. [16] and it is freely available for the research community. It is composed of 3,878 movie reviews collected from the MuchoCine website[9]. The reviews are written by web users instead of professional film critics. This increases the difficulty of the task because the sentences found in the documents may not always be grammatically correct, or they may include spelling mistakes or informal expressions. The corpus contains about 2 million words and an average of 546 words per review.

The opinions are rated on a scale from 1 to 5. One point means that the movie is very bad and 5 means very good. Films with a rating of 3 can be considered as "*neutral*", which means that the user considers the film is neither bad nor good. In our experiments we have discarded the neutral examples because the polarity classification task is binary, i.e. we have to classify the reviews as positive or negative only. Therefore, the opinions with ratings of 1 or 2 were considered as negative and those with ratings of 4 or 5 were considered as positive.

The MuchoCine English corpus (MCE) is the English version of MC. We generated MCE by applying a machine translation process using the Microsoft Translator[10] tool, formerly known as Bing Translator. Specifically we used the Java API provided for that tool. The MCE corpus is also freely available[11].

5 Experiments and Results

In this section we describe the experiments carried out and the results obtained after applying the proposed approach to the OCA-EVOCA and MC-MCE corpora. In the first subsection, the best individual results obtained for each parallel corpus are shown. Then, in the second subsection, we show the results obtained using the proposed approach.

5.1 Individual Results

According to the evaluation carried out by Rushdi-Saleh et al. [13] using supervised approaches over OCA and EVOCA, the configuration that reported the best results for the OCA corpus used SVM and TF·IDF as learning algorithm and weighting scheme, respectively, and did not apply the stemming process. The score obtained for the F1 measure was 0.9073. However, for the EVOCA corpus, the best F1 score (0.8840) was obtained by applying the stemming process and also using SVM and TF·IDF. For the unsupervised method, we carried out several experiments, as explained at the beginning of Section 4, and the configuration that reported the best F1 score used nouns and adjectives solely, obtaining a F1 score of 0.6698, which is lower than that obtained using the supervised approach, as expected.

[9] `http://www.muchocine.net`
[10] `http://www.bing.com/translator`
[11] `http://sinai.ujaen.es/wiki/index.php/MCE_Corpus_(English_version)`

Regarding the evaluation of the MC corpus, Martínez-Cámara et al. [20] followed a similar procedure based on the cross-validation method for the supervised approach. The best configuration for the MC corpus used SVM, TF·IDF, stopper and did not apply the stemming process. The best F1 score was 0.8767. For the translated version of MC (MCE), we considered the same configuration as the best one, achieving 0.8698 of F1 score. Regarding the semantic orientation approach, we carried out the same experiments as for the EVOCA corpus and the configuration that reported the best F1 score used adjectives and verbs solely, achieving a F1 value of 0.6879.

Table 1 summarizes the best individual results obtained for both corpora, showing the score obtained for the typical measures in classification tasks, such as *precision* (P), *recall* (R) and F1.

Table 1. Best results obtained for both parallel corpora individually

Corpora	Approach	Setting	P	R	F1
OCA	supervised	SVM, TF·IDF and no stemming	0.8699	0.9480	**0.9073**
EVOCA	supervised	SVM, TF·IDF and stemming	0.9007	0.8680	0.8840
	unsupervised	nouns + adjectives	0.5535	0.8480	0.6698
MC	supervised	SVM, TF·IDF and no stemming	0.8771	0.8763	**0.8767**
MCE	supervised	SVM, TF·IDF and no stemming	0.8704	0.8693	0.8698
	unsupervised	adjectives + verbs	0.5669	0.8744	0.6879

5.2 Results Obtained Using the Proposed Approach

After carrying out the individual experiments we propose the following method: if we use several classifiers for the same data then we will obtain several models that have learned different patterns from that data. In this manner it is very likely that the correct combination of the models achieves better results than those obtained by each classifier individually. Therefore we adapted the idea of the ensemble classifiers, but working with parallel corpora instead of the same corpus.

Taking into account the best results obtained individually over the OCA-EVOCA and MC-MCE corpora, we decided to combine them in order to improve the performance achieved separately. Specifically we tried *voting* as one of the most widely used combination algorithms in order to carry out the meta-classifier process that combines the three models generated from each corpora. The proposed algorithm makes use of the well-known voting system called majority rule [21]. Then we proposed two possible combinations for both corpora:

- Combination of the three models generated: the supervised approach applied to the original corpus (OCA-SVM and MC-SVM), the supervised approach applied to the translated corpus (EVOCA-SVM and MCE-SVM), and the unsupervised approach applied to the translated corpus (EVOCA-SWN and MCE-SWN).

– Combination of the supervised models: OCA-SVM + EVOCA-SVM for the OCA-EVOCA corpora, and MC-SVM + MCE-SVM for the MC-MCE corpora.

Due to the fact that the number of voters in the first combination is odd, the application of the voting system always returns a single-winner. However, in the second combination (supervised models from original and translated corpora) it is possible to obtain a draw because the predicted class for the OCA-SVM/MC-SVM voters may be different from that obtained by the EVOCA-SVM/MCE-SVM voters, respectively. In order to solve this problem we have considered two possible heuristics:

– Assign a final positive prediction only if both voters return a positive prediction (otherwise negative prediction), or
– Assign a final positive prediction if at least one of the voters returns a positive prediction (negative prediction only when both voters return a negative prediction)

Taking into account these possible combinations and heuristics, Table 2 shows the results obtained by applying the proposed approach to the OCA-EVOCA and MC-MCE corpora.

Table 2. Results obtained by applying the proposed approach

Corpora	Combination	Heuristic	P	R	F1
OCA-EVOCA	OCA-SVM + EVOCA-SVM + EVOCA-SWN	-	0.8566	0.9800	**0.9142**
	OCA-SVM + EVOCA-SVM	pos. if both voters	0.8984	0.9200	0.9091
		pos. if one voter	0.8483	0.9840	0.9111
MC-MCE	MC-SVM + MCE-SVM + MCE-SWN	-	0.8160	0.9608	0.8825
	MC-SVM + MCE-SVM	pos. if both voters	0.8551	0.8893	0.8719
		pos. if one voter	0.8003	0.9843	**0.8828**

6 Analysis of the Results

In this section we analyze the results obtained for both individual and combined experiments. Regarding the individual experiments is noteworthy the good behavior of the supervised approach versus the unsupervised one, as expected. Taking into account the translated versions of the corpora evaluated, the difference obtained for the supervised approach was around +32% and +26% regarding the unsupervised one for the EVOCA and MCE corpora, respectively. If we compare the supervised approach between the original corpus and its translation, the results obtained for the original corpus improve slightly those obtained for the translated version. For the OCA-EVOCA corpora, this improvement was around +3%, while for the MC-MCE corpora was around +0.8%. This behavior is also expected due to the noise that almost all automatic translation tools introduce during the process, although, specifically for the corpora evaluated, it is important to note the good performance of this translation process.

If we compare the results obtained by using the proposed combination approach with those obtained by using the supervised and unsupervised approaches separately, we can observe the improvement achieved by using the proposed approach. As can be seen in Table 3, for the OCA-EVOCA corpora we obtained an improvement of +0.76% regarding the supervised approach applied to the OCA corpus (OCA-SVM). On the other hand, for the MC-MCE corpora we obtained an improvement of +0.70% regarding the supervised approach applied to MC corpus (MC-SVM). This means that the proposed approach can be considered an interesting strategy for applying in polarity classification tasks when we work with parallel corpora.

Table 3. Comparison between the best results obtained by applying the proposed combination approach and those obtained by using the supervised and unsupervised approaches separately

Corpora	Approach	P	R	F1
OCA-EVOCA	OCA-SVM	0.8699	0.9480	0.9073
	EVOCA-SVM	0.9007	0.8680	0.8840
	OCA-SVM + EVOCA-SVM + EVOCA-SWN (combined)	0.8566	0.9800	**0.9142**
MC-MCE	MC-SVM	0.8771	0.8763	0.8767
	MCE-SVM	0.8704	0.8693	0.8698
	MC-SVM + MCE-SVM (combined)	0.8003	0.9843	**0.8828**

7 Conclusions and Further Work

In this paper we have presented a study about polarity classification over corpora written in different languages of English. In the proposed approach, firstly we translated the original corpus into English in order to generate its parallel corpus. Then, several experiments were carried out in order to build supervised and unsupervised models using these corpora. SentiWordNet was used as linguistic resource for the unsupervised experiments. Finally, the individual models were combined by applying a voting algorithm based on the majority rule. Although the results obtained with individual models were very promising, we have shown that the combination approach improved the performances achieved individually. In addition, this improvement was achieved in two parallel corpora so the robustness of the proposed method was evaluated in different frameworks.

For further work, we would like to test the performance using linguistic resources other than SentiWordNet, like for example WordNet-Affect or General

Inquirer. Moreover, it could be interesting to generate several lists of affective words for languages other than English. Thus, we could apply a semantic orientation approach directly to the original corpus and obtain a new model to consider in the meta-classifier architecture.

Acknowledgments. This study has been partially supported by a grant from the Fondo Europeo de Desarrollo Regional (FEDER), TEXT-COOL 2.0 project (TIN2009-13391-C04-02), ATTOS project (TIN2012-38536-C03-0) from the Spanish Government. Also, this study is partially funded by the European Commission under the Seventh (FP7 - 2007-2013) Framework Programme for Research and Technological Development through the FIRST project (FP7-287607). This publication reflects the views only of the authors, and the Commission cannot be held responsible for any use which may be made of the information contained therein.

References

1. Pang, B., Lee, L.: Opinion mining and sentiment analysis. Found.Trends Inf.Retr. 2, 1–135 (2008)
2. Pang, B., Lee, L., Vaithyanathan, S.: Thumbs up?: Sentiment classification using machine learning techniques. In: Proceedings of the ACL 2002 Conference on Empirical Methods in Natural Language Processing, EMNLP 2002, vol. 10, pp. 79–86. Association for Computational Linguistics, Stroudsburg (2002)
3. Turney, P.D.: Thumbs up or thumbs down?: semantic orientation applied to unsupervised classification of reviews. In: Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL 2002, pp. 417–424. Association for Computational Linguistics, Stroudsburg (2002)
4. Mihalcea, R., Banea, C., Wiebe, J.: Learning multilingual subjective language via cross-lingual projections. In: Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics, pp. 976–983. Association for Computational Linguistics, Prague (2007)
5. Kim, S.M., Hovy, E.: Identifying and analyzing judgment opinions. In: Proceedings of the Main Conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics, HLT-NAACL 2006, pp. 200–207. Association for Computational Linguistics, Stroudsburg (2006)
6. Denecke, K.: Using sentiwordnet for multilingual sentiment analysis. In: ICDE Workshops, pp. 507–512 (2008)
7. Tan, S., Zhang, J.: An empirical study of sentiment analysis for chinese documents. Expert Systems with Applications 34, 2622–2629 (2008)
8. Zhang, C., Zeng, D., Li, J., Wang, F.Y., Zuo, W.: Sentiment analysis of chinese documents: From sentence to document level. Journal of the American Society for Information Science and Technology 60, 2474–2487 (2009)
9. Wan, X.: Co-training for cross-lingual sentiment classification. In: Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP, ACL 2009, vol. 1, pp. 235–243. Association for Computational Linguistics, Stroudsburg (2009)
10. Ghorbel, D.J.H.: Sentiment analysis of french movie reviews. In: Proceedings of the 4th International Workshop on Distributed Agent-based Retrieval Tools (DART 2010) (2010)

11. Balahur, A., Turchi, M.: Multilingual sentiment analysis using machine translation? In: Proceedings of the 3rd Workshop in Computational Approaches to Subjectivity and Sentiment Analysis, WASSA 2012, pp. 52–60. Association for Computational Linguistics, Stroudsburg (2012)
12. Rushdi-Saleh, M., Martín-Valdivia, M.T., Ureña López, L.A., Perea-Ortega, J.M.: OCA: Opinion corpus for Arabic. Journal of the American Society for Information Science and Technology 62, 2045–2054 (2011)
13. Rushdi-Saleh, M., Martín-Valdivia, M.T., Ureña-López, L.A., Perea-Ortega, J.M.: Bilingual Experiments with an Arabic-English Corpus for Opinion Mining. In: Angelova, G., Bontcheva, K., Mitkov, R., Nicolov, N. (eds.) RANLP 2011 Organising Committee, pp. 740–745 (2011)
14. Banea, C., Mihalcea, R., Wiebe, J., Hassan, S.: Multilingual subjectivity analysis using machine translation. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP 2008, pp. 127–135. Association for Computational Linguistics, Stroudsburg (2008)
15. Brooke, J., Tofiloski, M., Taboada, M.: Cross-linguistic sentiment analysis: From english to spanish. In: International Conference RANLP, pp. 50–54 (2009)
16. Cruz, F.L., Troyano, J.A., Enriquez, F., Ortega, J.: Clasificación de documentos basada en la opinión: experimentos con un corpus de críticas de cine en español. Procesamiento del Lenguaje Natural 41, 73–80 (2008)
17. Martínez-Cámara, E., Martín-Valdivia, M.T., Ureña-López, L.A.: Opinion classification techniques applied to a spanish corpus. In: Muñoz, R., Montoyo, A., Métais, E. (eds.) NLDB 2011. LNCS, vol. 6716, pp. 169–176. Springer, Heidelberg (2011)
18. Steinberger, J., Ebrahim, M., Ehrmann, M., Hurriyetoglu, A., Kabadjov, M., Lenkova, P., Steinberger, R., Tanev, H., Vázquez, S., Zavarella, V.: Creating sentiment dictionaries via triangulation. Decision Support Systems 53, 689–694 (2012)
19. Baccianella, S., Esuli, A., Sebastiani, F.: Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In: Chair, N.C.C., Choukri, K., Maegaard, B., Mariani, J., Odijk, J., Piperidis, S., Rosner, M., Tapias, D. (eds.) Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC 2010), European Language Resources Association (ELRA), Valletta (2010)
20. Martínez-Cámara, E., Martín-Valdivia, M.T., Perea-Ortega, J.M., Ureña-López, L.A.: Opinion classification techniques applied to a Spanish corpus. Procesamiento del Lenguaje Natural 47 (2011)
21. Johnson, P.E.: Voting systems. a textbook-style overview of voting methods and their mathematical properties. Technical report, University of Kansas (2005)