Improving Emotional Intensity Classification using Word Sense Disambiguation

Jorge Carrillo de Albornoz[1], Laura Plaza[2], Pablo Gervás[3]
Computer Science Faculty, Universidad Complutense de Madrid,
C/ Prof. José García Santesmases, s/n. 28040, Madrid, Spain
[1]jcalbornoz@fdi.ucm.es, [2]lplazam@fdi.ucm.es, [3]pgervas@sip.ucm.es

Abstract. During the last years, sentiment analysis has become a very popular task in Natural Language Processing. Affective analysis of text is usually presented as the problem of automatically identifying a representative emotional category or scoring the text within a set of emotional dimensions. However, most existing approaches determine these categories and dimensions by matching the terms in the text with those presented in an affective lexicon, without taking into account the context in which these terms are immersed. This paper presents a method for the automatic tagging of sentences with an emotional intensity value, which makes use of the WordNet Affect lexicon and a word sense disambiguation algorithm to assign emotions to concepts rather than terms. An extensive evaluation is performed using the metrics and guidelines proposed in the SemEval 2007 Affective Text Task. Results are discussed and compared with those obtained by similar systems in the same task.

Keywords: Sentiment Analysis, Word Sense Disambiguation, Emotional Intensity, Machine Learning Techniques

1 Introduction

Sentiment analysis is becoming increasingly important in recent years. This discipline comprises very distinct research areas, such as text analysis, speech and facial expression, which have motivated a great number of systems, each one with specific properties and frequently using *ad hoc* and rarely available resources.

Focusing on text applications, emotional analysis usually relies on the identification of emotional keywords in the text [1]. These emotional keywords are either compared with those existing in an affective lexicon or used as the keys of a set of rules. Furthermore, most approaches work at the lexical level, using words or stems as emotional keywords, instead of the appropriate concepts according to their contexts. The few approaches that work at a conceptual level rarely make use of word sense disambiguation techniques in order to obtain the correct senses of the concepts. Instead, they simply obtain the first meaning or all possible meanings [2].

On the other hand, most sentiment analysis systems are focused on the identification of emotional categories or dimensions in the sentences [3, 4]. This information may be insufficient or even irrelevant in some applications. For instance,

Special issue: Natural Language Processing and its Applications.
Research in Computing Science 46, 2010, pp. 131-142
Received 26/11/09
Accepted 16/01/10
Final version 08/03/10

the emotional intensity as well as its polarity (positive vs. negative) can be of interest in public opinion analysis systems [5].

In this paper, an emotional intensity classifier capable of determining the intensity and polarity of the sentences in a text is presented. This system has been conceived to be used in an automatic camera management system for virtual environments based on the emotional analysis of the film script. In this context, the emotional intensity of the scene is an important parameter if the adequate camera position and movements need to be selected from the dialogs and actions described in the film script. The system is based on the identification of concepts in the sentences rather than terms, using a word sense disambiguation tool to obtain the correct senses for these concepts. The WordNet Affect lexicon is used to identify those concepts which are a priori candidates to denote an emotion or feeling.

The paper is organized as follows. Section 2 exposes the background and related work on sentiment analysis. Section 3 presents the method proposed for the tagging of sentences with emotional intensities. Section 4 introduces the evaluation methodology as well as the results obtained and the comparison with other similar systems. In section 5, the experimental results are discussed. Finally, section 6 provides concluding remarks and identifies pending problems and future work.

2 Background

This section presents the background of this study, as well as some recent works on sentiment analysis.

2.1 Emotion

Nowadays the most outstanding psychological theories on sentiment analysis are the emotional dimensions theory and the emotional categories theory. The first theory proposes the interpretation of human emotions throughout a set of emotional dimensions. This theory is based on James Russel studies [6], where the human emotional space is presented as a circular bipolar structure that represents the *pleasure* and the *activation*. The emotional dimensions were employed in the development of the ANEW list (*Affective Norms for English Words*) according to the SAM standard (*Self Assessment Manikin*) [7], where each English word is assigned a value, between 1 and 9, in each of the three dimensions proposed: *pleasure, arousal* and *dominance*. This work was supported by the idea that humans understand emotions as opposite poles.

On the other hand, the emotional categories theory exposes the emotions as entities, primitive units with boundaries that can be countable [8]. This idea can be traced to Rene Descartes [9], who proposes a set of primitive emotions from which the other emotions can be derived. In order to represent these categories, this theory makes use of the everyday language (i.e. *joy* and *anger* are emotional categories). The main handicap of this theory is the disagreement on the most adequate set of emotion categories. While some works argue that a small set of categories is the most suitable selection [10], others studies expose that a bigger hierarchical set of categories is necessary to enclose the rich human feeling [11]. This debate is also encouraged by the specific purpose of the studies. For instance, Ekman [12] argues that only six basic emotions (*anger, disgust, fear, joy, sadness* and *surprise*) are needed to analyze facial expressions, while Ortony et *al.* [13] present a set of 22 emotions in their OCC standard to emotion synthesis in speech.

2.2 Corpora and Affective Lexical Dictionaries

According to the different emotional theories, a wide range of resources has been developed, each of one supported by a psychological study and most of them specific for a given task. Focusing on the text sentiment analysis, any system that attempts to identify the affective meaning of a text will need, at least, two types of NLP resources: an annotated corpus to trail and test the system performance and an affective lexical dictionary that attaches affective meanings to words.

It is difficult to find affective corpora publicly available for researchers. Besides, most of them are very specific to the task and domain to which they have been designed; so that it is either impossible or surprisingly difficult to use them in a different task. An example of emotional corpus is Emotag [14]. Emotag consists of a set of sentences extracted from eight popular tales marked up by human evaluators with an emotional category and values for three emotional dimensions (*evaluation*, *activation* and *power*). A radically different corpus is the one proposed for the SemEval 2007 Affective Text task [15], where a set of 1250 sentences from news headings are manually tagged with six basic emotions (*anger, disgust, fear, joy, sadness* and *surprise*). Each emotion is scored between 0 and 100, where 0 indicates that the emotion is not present in the headline, and 100 indicates that the maximum amount of emotion is found in the headline. This corpus also includes a *valence* value in the interval [-100, 100] for each sentence, which expresses the negative (-100), neutral (0) or positive (100) intensity of the sentence.

Similarly to the corpora, the affective lexical dictionaries strongly depend on the underlying psychological theory. In relation to the emotional dimensions theory, the most popular lexicons are the ANEW word list (*Affective Norms for English Words*) [7] and the DAL dictionary (*Whissell's Dictionary of Affect Language*) [16]. The first one consists of a list of words scored within three emotional dimensions: *pleasure*, *arousal* and *dominance*, according to the SAM standard; while the second one contains 8742 words rated by people for their *activation*, *evaluation* and *imagery*. In relation to the emotional categories theory, the LIWC Dictionary (*Linguistic Inquiry and Word Count Dictionary*) [17] provides a set of 2290 words and stems, classified in one or more categories, such as *sadness*, *negative emotion* or *overall affect*.

The main limitation of these lexicons is the use of words or stems, instead of concepts, as the primitive units, without recognizing the context in which the words are used. On the contrary, the WordNet Affect database [18] provides a list of 911 WordNet synsets labeled with a hierarchical set of emotional categories. Most of the synsets labeled are representative of the meaning of the emotional categories (nouns and adjectives), while others are suitable to denote affective meanings.

2.3 Sentiment Analyzers

As already mentioned, the most accepted approach to sentiment analysis is the identification of a set of emotional keywords in the text. However, a great variety of methods have been proposed to achieve this purpose. Francisco and Gervás [14] present a system which is consistent with both theories: the emotional categories and the emotional dimensions, based on the averaged frequencies of the words found in a corpus of tales. Subasic and Huettner [19] propose a fuzzy approximation that uses an affective lexicon containing words manually annotated with two properties: *centrality* and *intensity*. These properties are, respectively, the membership degree of an emotional category and the strength of the affective meaning. A similar approach is presented in [20] where the emotional categories are represented as fuzzy hipercubes with a range of intensity between 0 and 1.

More sophisticated are the methods that aim to improve the emotional information extracted by means of semantic rules and syntactical analysis. Zhe and Boucouvalas [21] present an emotion extraction engine that uses a parser to identify auxiliary verbs, negations, subject of the sentence, etc. Wu et *al.* [3] expose a system where the annotation process is guided by emotional generation rules manually deduced from psychology. In the same line, Mostafa Al Masum et *al.* [22] present the system ASNA (*Affective Sensitive News Agent*) that uses the OCC model with a set of rules and natural language processing techniques to automatically classify news. Nicolov et *al.* [23] analyzed the effect of coreference resolution in sentiment analysis.

A third common approximation to the problem is the use of Machine Learning techniques. Devillers et *al.* [24] study the applicability of different machine learning algorithms to identify relevant emotional states in real life spoken interactions. In contrast, Seol et *al.* [25] present a hybrid system that uses emotional keywords if these are present in the text or a set of domain knowledge-based artificial neural networks if no emotional keywords are found. The neural networks are initialized with a set of rules that determine possible emotional categories states.

3 The Emotion Classifier

In this section, the method for automatically labeling sentences with an emotional intensity is presented. The problem is faced as a text classification task. The classifier aims to identify the emotional intensity of each sentence in a text, as well as if this intensity denotes a positive or negative emotional meaning. The method accomplishes the task throughout four steps. Each step is explained in detail in the following subsections, along with a working example that illustrates the algorithm. Besides, in order to clarify how the system works and what resources are used, its architecture is shown in Fig. 1.

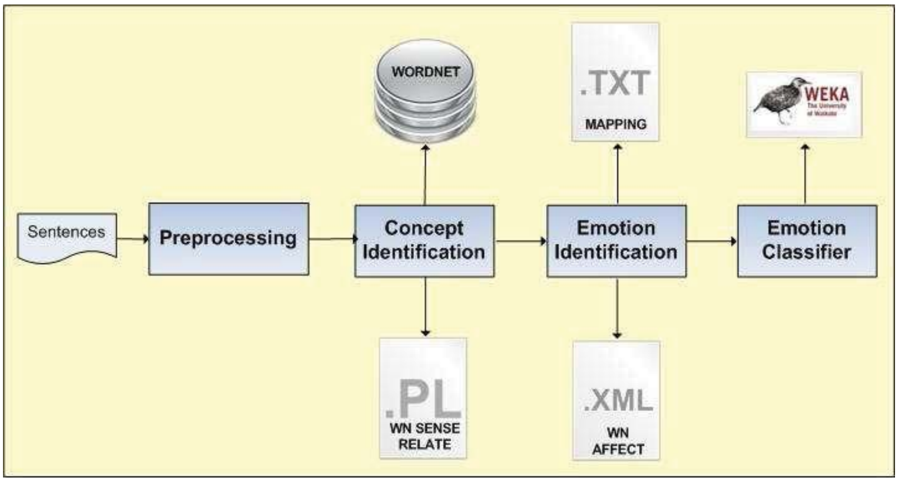

Fig. 1. Architecture of the emotional intensity tagger

3.1 Preprocessing

As most NLP systems, a preliminary preprocessing of the input text is needed. This includes splitting the text into sentences and tagging the words with their part of speech (POS). To this purpose, the *Tokenizer, Part of Speech tagger* and *Sentence splitter* modules in GATE [26] have been used. Generic and high frequency terms are removed using a stop list. Besides, as only the *nouns*, *verbs*, *adjectives* and *adverbs* can present an emotional meaning, only terms from these grammatical categories are considered.

3.2 Concept Identification

Once the text has been split into sentences and the words have been labeled with their POS, the next step is the mapping of the terms in the sentences to their appropriated concepts in the *WordNet* lexical database [27].

In order to correctly translate these terms to WordNet concepts, a word sense disambiguation tool is needed. To this aim, the implementation of the *lesk* algorithm in the *WordNet Sense Relate* Perl package was used [28]. As a result, for each word its corresponding stem and sense in WordNet are obtained. This information is used to retrieve the appropriate synset that represents the concept in WordNet. Next, the hypernyms of each concept are also retrieved.

Fig. 2 shows the concept identification process for the sentence: *Foetal mechanism helps heart failure.* In this sentence, the term *foetal* has not been correctly disambiguated by WordNet Sense Relate and its synset could not be retrieved from WordNet.

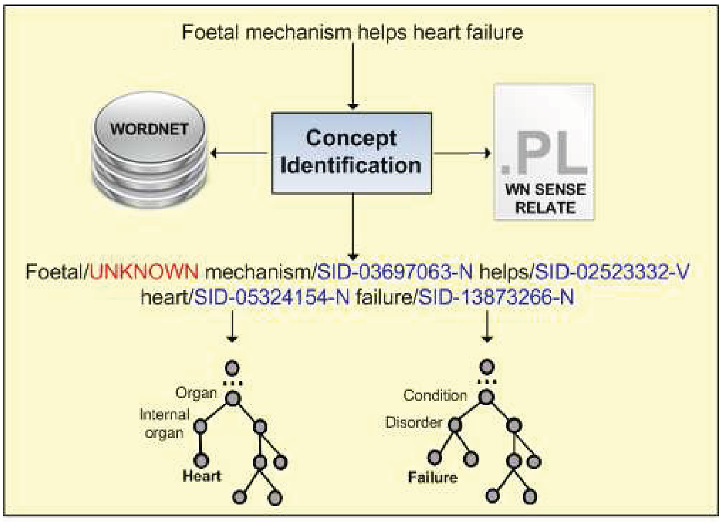

Fig.2. An example of the concept identification process

3.3 Emotion Identification

The goal of this third step is to match the previously identified WordNet concepts to their emotional categories in the WordNet Affect affective lexicon.

Focusing on the WordNet Affect emotion sub hierarchy, the first level distinguishes between *positive-emotions*, *negative-emotion*, *neutral-emotions* and *ambiguous-emotions,* while the second level encloses the emotional categories themselves. This level contains most of the basic emotions exposed in the emotional categories theories, such as *sadness, joy* and *surprise.* As the hierarchy used in WordNet Affect is considerably broader than those frequently used in sentiment analysis, the authors have identified that the second level is a good representation of human feeling and a good starting point for the attribute selection for the classifier and its evaluation. This subset contains 32 emotional categories.

Thus, once the synset of each word in the sentence has been identified, its emotional category is retrieved from WordNet Affect (if the concept appears in the lexicon). The same analysis is carried out over their hypernyms if no entry in WordNet Affect is found for the synset. To this aim, a previous mapping between synsets of WordNet 2.1 and WordNet 1.6 versions is needed, since the method and the affective lexicon works on different versions of WordNet. The use of the WordNet 2.1 version instead of the WordNet 1.6 is motivated by the fact that this version is the most updated for windows operative systems.

This process is illustrated in Fig. 3. It can be observed that only two concepts in the example sentence have been assigned an emotional meaning after this step. The emotional category for the concept *helps* is retrieved from its own synset, which is assigned the *liking* category. In contrast, the synset of the concept *failure* is not labeled in WordNet Affect, so the analysis of its hypernyms is carried out. As it first level hypernyms (*disorder*) is labeled with the *general-dislike* category, this same category is assigned to *failure.*

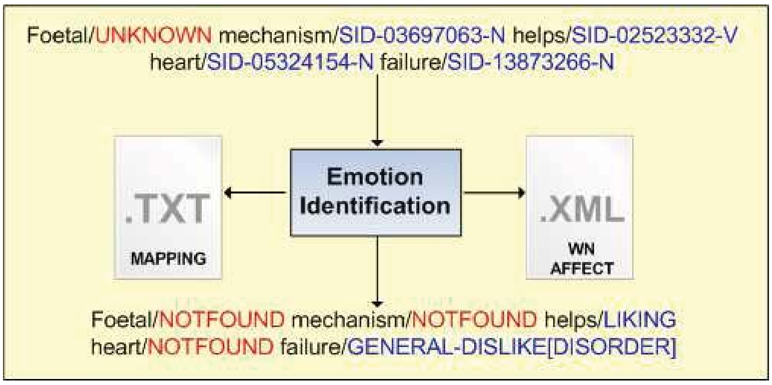

Fig. 3. An example of the emotion identification process

3.4 Emotional Intensity Classification

Up to this point, all the words in the sentence have been labeled with their emotional category (if any). Next, a vector of emotional occurrences (*VEO*) is created, which will be used as input for the *Random Forest* classifier implementation as provided by the *Weka* machine learning tool [29].

A VEO vector is an array of 32 positions, one representing each emotional category. To construct this vector, each concept is evaluated in order to determine its degree of affective meaning. If the concept has been assigned an emotional category, then the position in the VEO vector that represents that category is increases in 1. If no emotional category was retrieved for the concept, then the nearest labeled hypernym is used. As a hypernym is a generalization of the concept, a lower weight is assigned to the category position in the VEO vector, which depends on the depth of the hypernym in the hierarchy, as defined in (1). In order to avoid an excessive emotion generalization, only the first *n* levels of hypernyms are considered, where *n* has been empirically set to 3.

$$\text{VEO [i]} = \text{VEO [i]} + 1/(\text{Hyper. Depth}+1) \tag{1}$$

Fig. 4 shows the VEO vector for the example sentence. Since the emotional category *liking* is assigned to *helps*, the position of *liking* in the VEO vector is increased in one degree. On the other hand, the concept *failure* is labeled with the emotional category *general-dislike* through its first level hypernym, so its position in the VEO vector is increased in 0,5.

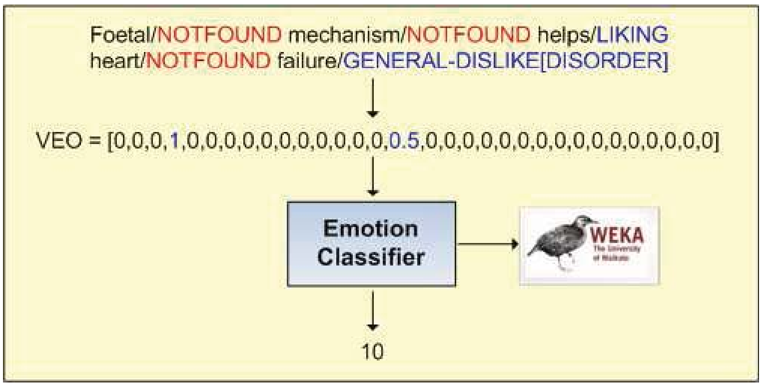

Fig. 4. An example of the emotional intensity classifier

Finally, the VEO vector is used as input for the classifier, and an intensity value between -100 and 100 is obtained as output.

4 Evaluation

In order to evaluate the performance of the method, a large–scale evaluation is accomplished following the guidelines observed in the SemEval 2007 Affective Text Task and using the corpus developed for this task.

The evaluation corpus consists of a training set of 250 sentences and a test set of 1000 sentences of news headlines. Each sentence has been manually labeled by experts with a score between -100 and 100, where -100 means a strongly negative emotional intensity, 100 means a strongly positive emotional intensity and 0 means neutral.

In this task, the intensity score assigned to each sentence has been mapped to three broader classes, -100 (between -100 and -50), 0 (between -50 and 50) and 100 (between 50 and 100). The results are evaluated using precision and recall metrics.

Due to the specific language used in the news domain most of the affective words in the headlines are not found in WordNet Affect. Thus, a set of emotional meanings has been manually added to the lexicon, such as *dead*, *kill* and *kidnap*. In particular, 188 synsets have been added: 140 nouns, 56 verbs and 22 adjectives.

To determine the best classification algorithm, all the experiments have been carried out over the different classification techniques implemented in Weka, although only the results of the best four algorithms are presented here.

The first group of experiments is directed to find out the best value for the number of hypernym levels as explained in section 3.4. For this parameter, we have considered all possible values from 0 to 5. As it can be observed in Table 1, using three levels of hypernyms produces the best results in 2 out of the 4 algorithms.

Table 1. Evaluation of the number of hypernym levels

Algorithm		**Number of hypernyms level**					
		0	*1*	*2*	*3*	*4*	*5*
Functional Trees	*Precision*	52,2	57,0	58,1	**60,4**	60,2	60,2
	Recall	62,2	62,2	62,8	**63,1**	62,9	62,8
Random Forest	*Precision*	57,5	57,0	57,4	**57,6**	57,4	57,3
	Recall	62,0	62,0	62,1	**62,5**	62,3	62,2
Naïve Bayes	*Precision*	56,3	**58,5**	58,3	58,3	58,3	58,3
	Recall	60,1	**59,8**	59,7	59,5	59,5	59,5
Multinomial Logistic	*Precision*	55,8	57,5	57,4	57,4	57,8	**58,1**
	Recall	62,0	62,6	62,9	63,1	63,1	**63,2**

The aim of the second group of experiments is to determine the best set of emotional categories for the classifier. Starting from the initial set of 32 categories, three algorithms for attribute selection implemented in Weka have been applied, which lead to three subsets of 3, 4 and 12 attributes respectively. Next, the four classifiers have been evaluated over these reduced sets of attributes. For these experiments, the number of hypernym levels was set to 3. Table 2 shows that two classifiers perform better with the subset of attributes selected by the *CFS Best First* algorithm.

Table 2. Evaluation of the set of attributes

Algorithm		**Consistency Best First**	**CFS Best First**	**Consistency Genetic**
Functional Trees	*Precision*	**57,7**	57,0	57,2
	Recall	**62,8**	62,2	62,3
Random Forest	*Precision*	63,1	**64.0**	63,5
	Recall	63,1	**63,5**	63,5
Naïve Bayes	*Precision*	55,8	**58,7**	55,8
	Recall	62,4	**61,3**	62,4
Multinomial Logistic	*Precision*	**57,7**	57,0	57,3
	Recall	**62,8**	62,2	62,5

A third group of experiments has been carried out in order to evaluate the effect of the distribution of the intensity within classes. To this aim, the emotional intensity of the sentences has been mapped to 3 balanced classes: -100 (from -100 to -35), 0 (from -35 to 35) and 100 (from 35 to 100), and to 5 balanced classes: -100 (-100 to -60), -50 (-60 to -20), 0 (-20 to 20), 50 (20 to 60) and 100 (60 to 100). For these experiments, the number of hypernym levels was set to 3, while the subset of attributes selected by the *CFS Best First* algorithm was used. Table 3 shows that the method performance decreases substantially when the number of classes is increased, while only a small reduction is reported when the 3 classes are equitably distributed with respect to the original distribution.

Therefore, the best configuration implies using the *Random Forest* algorithm along with the 4 attribute subset obtained by the *CFS Best First* algorithm and 3 hypernym levels.

Table 3. Evaluation of the intensity distribution in classes

Algorithm		3 Classes	5 Classes
Functional Trees	*Precision*	54,8	26,7
	Recall	51,0	35,5
Random Forest	*Precision*	**55,4**	36,6
	Recall	**52,1**	35,7
Naïve Bayes	*Precision*	52,1	**39,1**
	Recall	43,9	**32,9**
Multinomial Logistic	*Precision*	55,9	26,7
	Recall	50,7	35,6

Finally, Table 4 summarizes the best results of our method along with those of the systems participating in the SemEval 2007 Affective Text Task. Our method outperforms the other systems in precision, while CLaC-NB obtains the best recall.

Table 4. Comparison with SemEval 2007 systems

Systems	Precision	Recall
CLaC	61.42	9.20
UPAR7	57.54	8.78
SWAT	45.71	3.42
CLaC-NB	31.18	**66.38**
SICS	28.41	60.17
Our method	**64.00**	63.50

5 Discussion

As already mentioned, the system has obtained very promising results. When compared to other systems that take part in the SemEval task, our method obtains the best precision (64%), while providing the second highest recall (63.5%). Furthermore, both metrics are well balanced, which does not occur in the other systems.

The use of concepts instead of terms, along with the use of a word sense disambiguation algorithm allows the method to correctly map the words in the sentence to their emotional categories in the affective lexicon. Besides, the use of hypernyms has permitted to increase the number of concepts labeled with an emotion, which has significantly improved the evaluation results. This indicates that the method is strongly dependent on the number of concepts labeled in the lexicon.

Another interesting result is that the initial set of 32 emotional categories has proved to be too large. Some of these categories introduces noise and decreases the efficiency of the classification algorithms. According to the experiments, the best set consists of 4 emotional categories: *liking, negative-fear, sadness* and *general-dislike*. However, it is important to note that this subset will depend on the regarded domain.

A detailed examination of the precision and recall obtained by each intensity class has shown that the positive class encloses most of the classification errors. The reason seems to be that a good number of the positive sentences are expressed as the negation of negative emotional concepts (i.e. *Jet flips in snowstorm, none dead*). A previous process of negation detection which adapts negated sentences to a positive form through antonym relations could minimize this problem.

6 Conclusions and Future Work

In this paper, an effective approach to automatically assign an emotional intensity and polarity to a sentence is presented. The method obtains both high precision and recall, which are also well balance. The evaluation results outperform those obtained by the systems participating in the SemEval 2007 Affective Text Task.

However, several problems have been identified. First, the experimentation has shown that the amount of index concepts is low, since most concepts in the corpus cannot be retrieved from WordNet Affect. This clearly has an influence on the precision of the method, which can be improved by enriching the lexicon with the corpus specific vocabulary.

A second handicap to the concept identification is the part of speech tagging errors reported by the GATE POS tagger, which has resulted in an good number of concepts that could not be correctly disambiguated and retrieved from WordNet. Future work will include a further evaluation using the statistical Stanford parser [30].

Finally, different techniques will be studied in order to detect negated concepts, as well as their scope, since they can invert the polarity of the sentences.

Acknowledgments. This research is funded by the Spanish Ministry of Science and Innovation (TIN2009-14659-C03-01). It is also partially funded by the Comunidad Autonoma de Madrid (CAM) and the European Social Fund (ESF) through the IV PRICIT program, and by the Spanish Ministry of Science and Innovation through the FPU program.

References

1. Strapparava, C., Mihalcea, R.: Learning to identify emotions in text. In: Proceedings of the 2008 ACM symposium on Applied computing, pp. 1556--1560. Ceara, Brazil (2008)
2. Chaumartin, F.R.: UPAR7: a knowledge-based system for headline sentiment tagging. In: Proceedings of the 4th International Workshop on Semantic Evaluations, pp. 422--425. Prague, Czech Republic (2007)
3. Wu, C.H., Chuang, Z.J., Lin, Y.C.: Emotion recognition from text using semantic labels and separable mixture models. Journal of ACM Transactions on Asian Language Information Processing 5, pp. 165--183 (2006)
4. Rubin, V.L., Stanton, J.M., Liddy, E.D.: Discerning Emotions in Texts. In: Proceedings of the AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications. Stanford, C.A. (2004)
5. Mullen, T., Collier, NSentiment Analysis using Support Vector Machines with Diverse Information Sources. In: Proceedings of EMNLP, pp. 412--418. Barcelona, Spain (2004)
6. Russell, J. A.: A circumplex model of affect. Journal of Personality and Social Psychology 36, pp. 1161--1178 (1980)

7. Bradley, M.M., Lang, P.J.: Affective norms for English words (ANEW): Instruction manual and affective ratings. Technical Report C-1, The Center for Research in Psychophysiology, University of Florida (1999)
8. James, W.: What is an Emotion? Published in Mind 9, pp. 188--205 (1884)
9. Anscombe, E., Geach, P.T.: Descartes: Philosophical Writings. Ed. Nelson University Paperbacks (1972)
10. Plutchik, R.: A general Psycho Evolutionary Theory of Emotion. Emotion: Theory, Research, and Experience (1980)
11. Parrott, W.G.: Emotions in Social Psychology: Essential Readings. Psychology Press, Philadelphia (2001)
12. Ekman, P.: Are there basic emotions? Journal Psychol. Rev. 99, pp 550--553 (1992)
13. Ortony, A., Clore, G. L., Collins, A.**:** The Cognitive Structure of Emotions. Cambridge University Press (1988)
14. Francisco, V., Gervás, P.: Ontology-Supported Automated Mark Up of Affective Information in Texts. Special Issue of Language Forum on Computational Treatment of Language 34, pp. 26--36 (2008)
15. SemEval 2007 Affective Text Task Corpus, http://www.cse.unt.edu/~rada/affectivetext/
16. Whissell, C.M.: The dictionary of affect in language. In: Robert Plutchik and Henry Kellerman (Ed.), Emotion: Theory, Research, and Experience, pp. 113--131 (1989)
17. Pennebaker, J.W., Francis, M.E., Booth, R.J.: Linguistic Inquiry and Word Count (LIWC): LIWC2001. Erlbaum Publishers, Mahwah, NJ (2001)
18. Strapparava, C., Valitutti, A.: Wordnet-affect: an affective extension of WordNet. In: Proceedings of the 4th International Conference on Language Resources and Evaluation, pp. 108--1086. Lisbon, Portugal (2004)
19. Subasic, P., Huettner, A.: Affect Analysis of Text Using Fuzzy Semantic Typing. Journal of IEEE Transactions on Fuzzy Systems 9, pp 483--496 (2001)
20. Esau, N., Kleinjohann, L., Kleinjohann, B.: An Adaptable Fuzzy Emotion Model for Emotion Recognition. In: Proceeding of EUSFLAT Conference, pp. 73--78 (2005)
21. Zhe, X., Boucouvalas, A.: Text-to-Emotion Engine for Real Time Internet Communication. In: Proceedings of International Symposium on CSNDSP, pp. 164--168 (2002)
22. Mostafa Al Masum, S., Islam, M. T., Ishizuka, M.: ASNA: An Intelligent Agent for Retrieving and Classifying News on the Basis of Emotion-Affinity. In: Proceedings of the CIMCA-IAWTIC'06 (2006)
23. Nicolov, N., Salvetti, F., Ivanova, S.: Sentiment Analysis: Does Coreference Matter? In: of the Symposium on Affective Language in Human and Machine, pp. 37--40. (2008)
24. Devillers, L., Vidrascu, L., Lamel, L.: Challenges in real-life emotion annotation and machine learning based detection. Journal of Neural Netw. 18, pp 407--422 (2005)
25. Seol, Y.S., Kim, D.J., Kim, H.W.: Emotion Recognition from Text Using Knowledge-based ANN. In: Proceedings of 23rd International Technical Conference on Circuits/Systems, Computers and Communications, pp. 1569--1572 (2008)
26. General Architecture for Text Engineering, http://gate.ac.uk/
27. Miller, G.A., Beckwith, R., Fellbaum, C., Gross, D., Miller, K.J.**:** Introduction to WordNet: An On-Line Lexical Database. International Journal of Lexicography 3 (4), pp. 235--244 (1990)
28. Patwardhan, S., Banerjee, S., Pedersen, T.: SenseRelate::TargetWord - A Generalized Framework for Word Sense Disambiguation. In: Proceedings of the Twentieth National Conference on Artificial Intelligence (Intelligent Systems Demonstrations), pp. 1692–1693. Pittsburgh, PA (2005)
29. Weka Machine Learning Tool, http://www.cs.waikato.ac.nz/ml/weka/
30. Statistical Stanford Parser, http://nlp.stanford.edu/software/lex-parser.shtml