SentiSense: An easily scalable concept-based affective lexicon for sentiment analysis

Jorge Carrillo de Albornoz, Laura Plaza, Pablo Gervás

Universidad Complutense de Madrid
C/ Profesor José García Santesmases, s/n
28040 Madrid, Spain
jcalbornoz@fdi.ucm.es, lplazam@fdi.ucm.es, pgervas@sip.ucm.es

Abstract

This paper presents SentiSense, a concept-based affective lexicon. It is intended to be used in sentiment analysis-related tasks, specially in polarity and intensity classification and emotion identification. SentiSense attaches emotional meanings to concepts from the WordNet lexical database, instead of terms, thus allowing to address the word ambiguity problem using one of the many WordNet-based word sense disambiguation algorithms. SentiSense consists of 5,496 words and 2,190 synsets labeled with an emotion from a set of 14 emotional categories, which are related by an antonym relationship. SentiSense has been developed semi-automatically using several semantic relations between synsets in WordNet. SentiSense is endowed with a set of tools that allow users to visualize the lexicon and some statistics about the distribution of synsets and emotions in SentiSense, as well as to easily expand the lexicon. SentiSense is available for research purposes.

Keywords: affective lexicon, emotional lexicon, sentiment analysis

1. Introduction and Motivation

Sentiment analysis and affective computing is becoming a key area of natural language processing (NLP) which aims to discover and interpret sentiments and opinions expressed in text. The grown of this discipline is mainly due to the interest of companies to quickly understand consumers' opinions about their products and services as a means to improve their marketing mix.

Sentiment analysis involves different research tasks, such as *subjectivity detection* (Wiebe et al., 1999; Pang and Lee, 2004), *polarity classification* (Pang et al., 2002; Turney, 2002), *intensity classification* (Wilson et al., 2009; Brooke, 2009), and *emotion identification* (Chaumartin, 2007; Katz et al., 2007). Subjectivity detection aims to discover subjective or neutral terms, phrases or sentences, and it is frequently used as a previous step in polarity and intensity classification with the aim of separating subjective information from objective one. Polarity classification attempts to classify texts into positive or negative. The intensity classification (or rating inference) task goes a step further and tries to identify different degrees of positivity and negativity, e.g., *strongly-negative*, *negative*, *fair*, *positive*, and *strongly-positive*. Finally, the emotion identification task seeks to identify the specific emotion (e.g., *sadness*, *fear*, etc.) that best reflects the meaning of the text.

To accomplish these tasks, different linguistic resources have been developed. On the one hand, several lexicons have been created to help determine if a term expresses a fact or an opinion (i.e., if the term is objective or subjective), and thus to support the subjectivity detection task. Among these resources, the most outstanding are SentiWordNet (Esuli and Sebastiani, 2006) and the Subjectivity Lexicon (Wilson et al., 2005). On the other hand, the second group of affective lexicons aims at deciding if a subjective term expresses a positive or negative opinion, and even the strength of such polarity. Therefore, such lexicons are frequently used in polarity and intensity classification systems. Examples of such resources are SentiWordNet and the General Inquirer (Stone et al., 1966).

Even though positive/negative annotation is interesting for some tasks, usually a more fine-grained emotion annotation is needed. When analyzing opinions about a phone, for instance, the manufacturer is interested in distinguishing a customer who is unhappy with the battery life, from a customer who is angry and frustrated with the treatment of the customer service. In this situation, it is important to understand the emotional meaning of the elementary textual units that make up the text. To this end, a lexicon that attaches emotional meanings or categories is needed. Examples of these types of resources are the LIWC Dictionary (Pennebaker et al., 2001), the LEW list (Francisco et al., 2010), and WordNet Affect (Strapparava and Valitutti, 2004).

However, these lexicons present several handicaps. Regarding the LIWC Dictionary and the LEW list, they attach emotions to words instead of concepts, and thus do not allow us to distinguish different meanings of the same word. Concerning WordNet Affect, we find two main limitations. First, there is an issue with the granularity of representations of emotional categories. We consider the set of emotional categories in WordNet Affect to be excessively broad. Second, there is an issue of labeling ambiguity. We have detected a good number of synsets in WordNet Affect (113 out of 911) that have been labeled more than once, and with different emotional categories, making it difficult to discern which of them is more appropriate in each situation.

To overcome such limitations, we have developed the SentiSense affective lexicon.[1] SentiSense attaches emotional

[1] http://nil.fdi.ucm.es/index.php?q=node/456

meanings or categories to concepts from the WordNet lexical database, instead of terms, allowing end-user applications to correctly disambiguate the terms using one of the many WordNet-based word sense disambituation algorithms. Moreover, the emotional categories in SentiSense are well-supported by most accepted psychological theories. SentiSense can be used for both polarity and intensity classification and emotion identification.

The coverage of vocabulary is another important issue. When developing an affective lexicon, two methodologies may be followed: an automatic labeling process (e.g., SentiWordNet) or a manual labeling one (e.g., the LEW list). The automatic labeling usually generates resources with high coverage of vocabulary but low precision. In this way, SentiWordNet covers all synsets in WordNet, but precision is sometimes poor (for instance, the concept *SID-14051451-N-{cancer#1}* is only assigned a negativity score of *0.125*). Automatic techniques use a seed of manually labeled terms or concepts, which are then used to train some classifiers in order to label new terms or concepts (Esuli and Sebastiani, 2006) or used to generate rules that inferrer the emotional meaning of new terms or concepts (Strapparava and Valitutti, 2004). These rules make use of the relations between words, or the structure of the graph in the thesaurus, etc. On the other hand, the manual labeling techniques generate resources with very low coverage but very high precision, obtaining affective lexicons that are intended for specific domains. The manually generated resources are usually developed by two or more annotators that label each term or concept. The performance of these resources is considerably high for the target domain, but drops substantially when they are used in other domains.

Our goal is to build a resource that combines both high vocabulary coverage and high precision. In this way, SentiSense may be developed in a collaborative manner, so that people may easily expand the resource in order to fit the desired application domain, and these extensions may be used to enrich the core of the lexicon. To assist this process, SentiSense is endowed with a set of tools that allow users to expand and visualize the lexicon and some statistics about the distribution of emotions in SentiSense.

The paper is organized as follows. Section 2 introduces the design principles and decisions. Section 3 describes the development process. Section 4 presents the labeling and visualization tools. Finally, section 5 provides concluding remarks and future lines of work.

2. Design of SentiSense

SentiSense classifies WordNet synsets (Miller, 1995) representing emotional meanings into a set of emotional categories. The main reason for using WordNet synsets instead of terms is that words usually have multiple senses so that a word can act as subjective or objective within a sentence depending on its context, and even present a different polarity. Other reasons are the wide coverage of the English lexicon and the availability of WordNet-based resources.

The emotional categories in SentiSense are based in those proposed by Arnold (1960), Plutchik (1980), and Parrot (2001). Arnold proposed one of the first classification of emotions. He defined a list of eleven fundamental emotions (*anger*, *aversion*, *courage*, *dejection*, *desire*, *despair*, *fear*, *hate*, *hope*, *love*, and *sadness*). Plutchik considers a narrower set of eight basic emotions: *acceptance*, *anger*, *anticipation*, *disgust*, *joy*, *fear*, *sadness*, and *surprise*. Parrot presents an even more reduced set of six primary emotions: *anger*, *fear*, *joy*, *love*, *sadness*, and *surprise*.

Category	Antonym	Category	Antonym
Ambiguous	-	Hate	Love
Anger	Calmness	Hope	Despair
Calmness	-	Joy	Sadness
Despair	Hope	Like	Disgust
Disgust	Like	Love	Hate
Anticipation	Surprise	Sadness	Joy
Fear	Calmness	Surprise	Anticipation

Table 1: Emotional categories in SentiSense and antonym relation among them

We first considered the list of sixteen emotions that result from combining the three models above. This list of emotions was showed to three experts in computational linguistics. They were first asked to propose, for each emotion, its closer antonym emotion, provided that they have a clear antonym. As a result, we got the following set of 20 emotions and antonym relations among them: {*acceptance-refusal*, *anger-calmness*, *anticipation-surprise*, *aversion-desire*, *courage-cowardice*, *dejection-hope*, *despair-hope*, *disgust-like*, *fear-calmness*, *hate-love*, and *joy-sadness*}. During the labeling process, however, the annotators noted that seven of them did not appear in the annotation corpus, they were not expected to be commonly used in opinionated texts. Therefore, such emotions were removed from the lexicon. Also suggested by the annotators, we introduced an *ambiguous* category in order to label those concepts with unclear or ambiguous emotional meaning. Consequenly, SentiSense presents 14 emotional categories, which are also related by an antonym relationship (see Table 1).

SentiSense consists of 5,496 words and 2,190 synsets labeled with an emotional category. Table 2 shows the number of synsets per emotional category and part of speech. The main part of the lexicon consists of nouns and adjectives, followed by verbs and a small set of adverbs. Table 3 shows some example synsets for each emotional category.

SentiSense consists of two data files in XML. The first file, *categories.xml*, defines the emotional categories and the antonym relationship between them (see Table 4). The second file, *synsets.xml*, contains the WordNet synsets that make up the lexicon. In this file, each entry contains the WordNet synset identifier (SID), its part of speech (POS), its gloss or definition in WordNet and the emotional category assigned to it. An extract of the synset.xml file may be shown in Table 5.

3. Development of SentiSense

SentiSense has been created semi-automatically in a two-phase process, following the development methodology of WordNet Affect. First, two annotators were presented the same 500 texts (250 news headlines and 250 hotel reviews

Category	Nouns	Verbs	Adjectives	Adverbs
Ambiguous	22	15	20	3
Anger	23	11	24	6
Calmness	18	6	35	10
Despair	3	2	5	3
Disgust	222	134	254	51
Anticipation	75	35	46	9
Fear	104	43	38	14
Hate	2	4	11	1
Hope	33	14	15	6
Joy	63	30	52	19
Like	146	70	185	25
Love	30	4	28	6
Sadness	58	27	67	22
Surprise	13	9	15	3
Total	812	405	795	178

Table 2: Distribution of synsets among different emotional and parts of speech in SentiSense

Category	Examples	
Ambiguous	SID-01739623-A	“rare, uncommon”
Anger	SID-07406119-N	“upset, discompose ...”
Calmness	SID-00087725-A	“unafraid, fearless”
Despair	SID-00201948-R	“hopelessly”
Disgust	SID-01791206-V	“disgust, revolt ...”
Anticipation	SID-01787793-V	“anticipate ...”
Fear	SID-01172163-A	“fearful, frightful”
Hate	SID-01171307-A	“atrocious, abominable ...”
Hope	SID-13805740-N	“freedom”
Joy	SID-07425411-N	“exultation, jubilance ...”
Like	SID-04628402-N	“beauty”
Love	SID-01512123-A	“adorable, endearing ...”
Sadness	SID-07435041-N	“depression”
Surprise	SID-00214835-R	“amazingly, surprisingly ...”

Table 3: Emotional categories and corresponding example synsets

from a development set). For each text, the list of WordNet synsets and the glosses describing them were also shown. The annotators were asked to select, from the set of emotional categories in Table 1, the one that best described the sentiment expressed by each synset, provided that the synset conveyed affective meaning. It must be noted that the task of assigning emotional categories to WordNet synsets is quite subjective. In order to solve interjudge disagreement and ensure the reliability of our resource, only the synsets for which the two annotators had

```
<SentiSenseEmotionalCategories>
    <EmotionalCategory name=“joy’ antonym=“sadness” />
    <EmotionalCategory name=“fear’ antonym=“calmness” />
    <EmotionalCategory name=“love’ antonym=“hate” />
    <EmotionalCategory name=“hope’ antonym=“despair” />
    ...
</SentiSenseEmotionalCategories>
```

Table 4: An extract of the *categories.xml* file

```
<SentiSenseCorpus>
    <Concept synset=“SID-00152712-A” pos=“adjective”
    gloss=“lacking cordiality...” emotion=“disgust”/>
    <Concept synset=“SID-00050667-R” pos=“adverb”
    gloss=“in a joyous manner...” emotion=“joy”/>
    <Concept synset=“SID-03430539-N” pos=“noun”
    gloss=“a weapon that discharges...” emotion=“fear”/>
    <Concept synset=“SID-02571914-V” pos=“verb”
    gloss=“come upon or...” emotion=“surprise”/>
    ...
</SentiSenseCorpus>
```

Table 5: An extract of the *synset.xml* file

emitted the same judgement were included in the lexicon (1200 synsets). We decided to choose this strategy in order obtain the highest possible precision in the manual process. The most frequent emotional categories in the affective lexicon are *like* and *disgust*, which have been described by the judges as the widest emotional categories in the corpus.

In the second step, these synsets were automatically expanded using several relations in WordNet. In particular, the following relations were considered: *antonym*, *hypernymy*, *derived-from-adjective*, *entailment*, *pertains-to-noun*, *participle-of-verb*, *attribute*, and *also-see*. For each relation, we studied if it generates synsets that preserve the same emotional meaning than the original synset. We concluded that only the *derived-from-adjective*, *pertains-to-noun*, and *participle-of-verb* relations typically maintain such emotional meaning. Therefore, all synsets obtained by the application of those relations were automatically labeled with the same emotions than the original synsets and included in the lexicon. We also found that antonym synsets present antonym emotional meanings, and therefore, all synsets obtained by applying the *antonym* relation were automatically labeled with the opposite emotional categories than the original synsets using the antonym relation between emotional categories defined in SentiSense (see Table 1) and included in the lexicon. For instance, if the synset *SID-02420512-A-{superior#1}* is manually labeled with the emotional category *like*, then its antonym synset *SID-02424479-A-{inferior#2}* will be annotated with the antonym emotional category of *like*; i.e., *disgust*.

4. Tagging and Visualization Tools

In order to help with the development process, we have implemented a tagging software. It allows annotators to select the data set from which they want to collect the vocabulary to be labeled. Note that, since the lexicon is based on concepts instead of terms, selecting the vocabulary from texts rather than labeling isolated words provides a context from which to obtain the correct meaning of each word. The tagging tool may be shown in Figure 1.

The data set used in the labeling process must conform to the format shown in Table 6.

Once the data set is loaded, each text is shown along with the list of terms that are found within it. When a term is selected, this is mapped to WordNet, and all candidate synsets are shown. For the linguistic processing of the text,

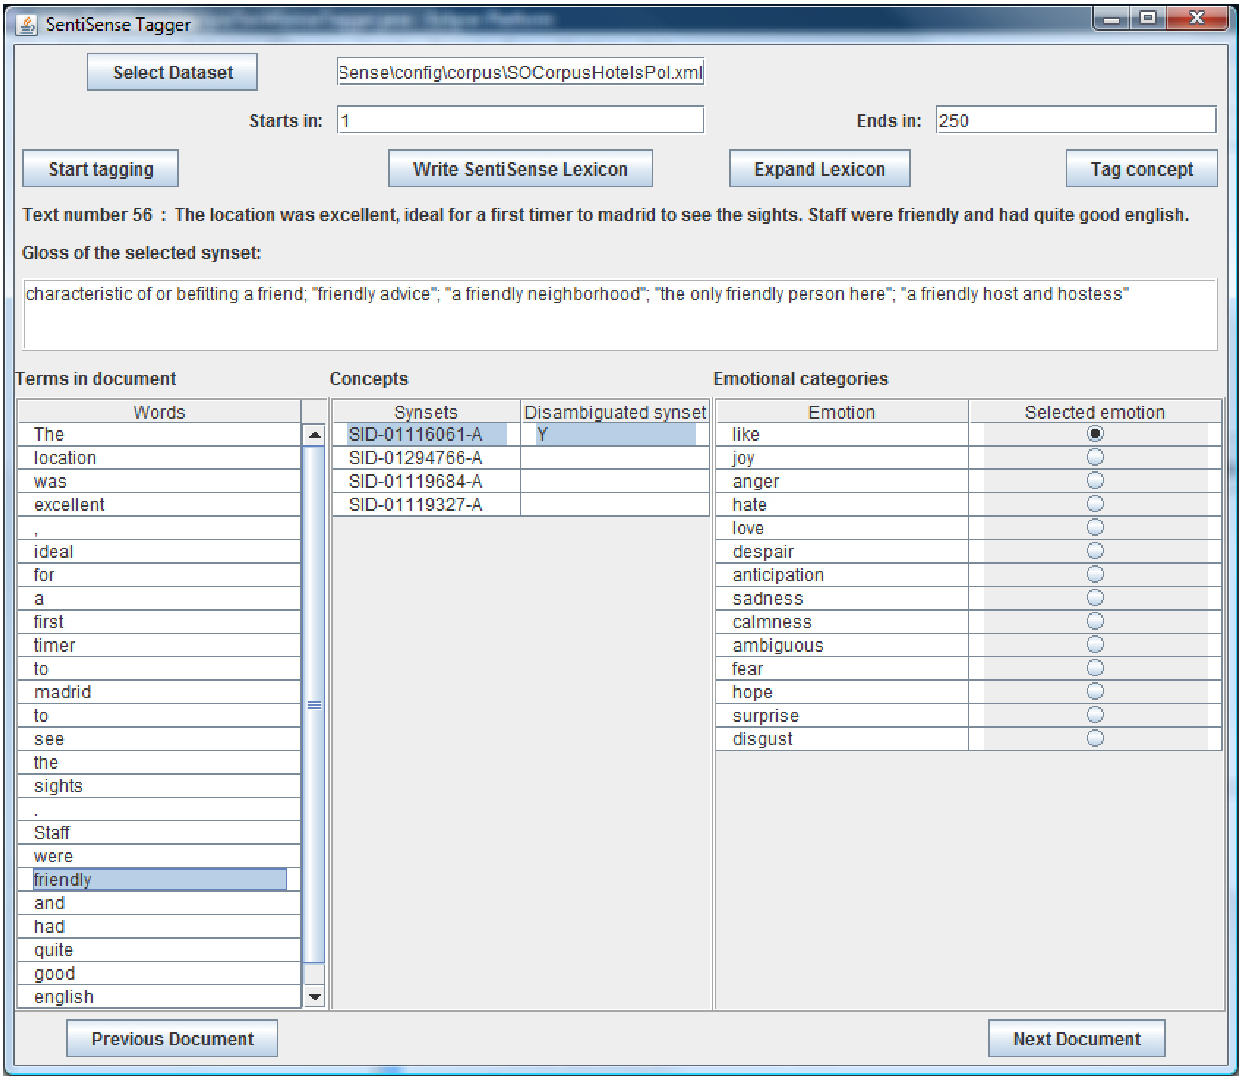

Figure 1: SentiSense tagging software

<**SentiSenseCorpus**>
<**SentiSenseDoc** id="0">
This hotel is new or recently renovated. There is a Monoprix ...
</**SentiSenseDoc**>
<**SentiSenseDoc** id="1">
Breakfast could have been better for the price paid.
</**SentiSenseDoc**>
...
</**SentiSenseCorpus**>

Table 6: An example of data set for assisting the tagging process

the GATE architecture [2] and the Stanford parser [3] are used. The Lesk disambiguation algorithm (Lesk, 1986), as implemented in the WordNet Sense-Relate package (Patwardhan et al., 2005), is executed and the correctly disambiguated synset is indicated. However, since the disambiguation algorithm may introduce some errors, the tool allows the user to manually change the wrongly disambiguated synset to the correct one. When the synset is selected, its gloss is shown and the annotator may select the emotional category that will be associated to the synset. Finally, the lexicon may be expanded automatically via WordNet relations.

SentiSense also offers a visualization tool, which can be seen in Figure 2. This tool shows, for each synset in the lexicon, its synset identifier and the words or terms that compose the synset. When a synset is selected, its emotional category is shown. Moreover, the application permits to change the emotional category associated to a given synset. The bottom right corner also displays statistics of the number of synset/word within the lexicon and their distribution among different emotional categories and parts of speech.

5. Conclusions and Future Work

In this paper, we have presented SentiSense, an affective lexicon that attaches emotional categories to WordNet synsets. We believe this lexicon can prove a useful resource for opinion mining and affective computing applications. One of its main advantages is the availability of a set of tools that allow users to easily expand the coverage of the lexicon, both manually and automatically, in order to cover the emotional vocabulary of each specific application do-

[2]http://gate.ac.uk/.
[3]http://nlp.stanford.edu/software/lex-parser.shtml

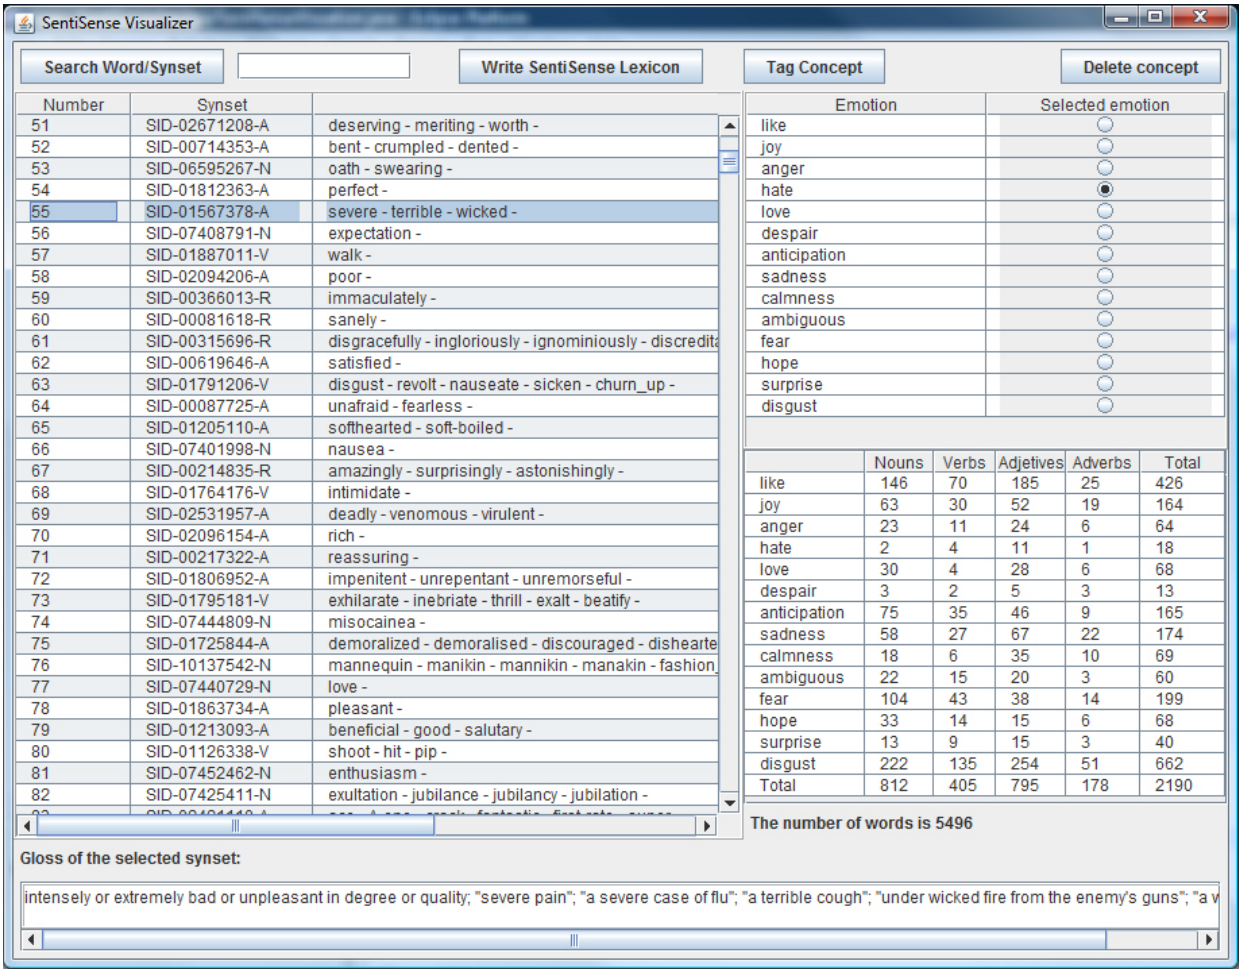

Figure 2: SentiSense visualitation tool

main. In this way, the lexicon may be extended collaboratively, so that users extensions may be used to enrich the core of the lexicon.

As future work, we plan to test new WordNet relations among synsets in order to automatically expand the number of tagged synsets only if the emotional meaning is preserved. Moreover, we will improve our tagging tool to allow users to select the specific relations they want to use to expand the lexicon, as well as to employ different WSD algorithms. We will also study the possibility of tagging not only unigrams, but also bigrams and expressions, and how to expand these emotional units with the relations among synsets. Finally, in a near future we want to compare SentiSense to other lexicons in the context of a real sentiment analysis application.

6. Acknowledgements

This research is funded by the Spanish Ministry of Science and Innovation (TIN2009-14659-C03-01), the Comunidad Autónoma de Madrid and the European Social Fund through the IV PRICIT program, and the Spanish Ministry of Education through the FPU program.

7. References

M.B. Arnold. 1960. *Emotion and Personality. Volume I: Psychological Aspects*. New York: Columbia University Press.

Julian Brooke. 2009. A semantic approach to automated text sentiment analysis. Master's thesis, Simon Fraser University.

Francois Regis Chaumartin. 2007. UPAR7: A knowledge-based system for headline sentiment tagging. In *Proceedings of the Fourth International Workshop on Semantic Evaluations*, Prague, Czech Republic.

Andrea Esuli and Fabrizio Sebastiani. 2006. SENTIWORDNET: a publicly available lexical resource for opinion mining. In *Proceedings of the 5th Conference on Language Resources and Evaluation (LREC 2006)*, pages 417–422, Genoa, Italy.

Virginia Francisco, Pablo Gervás, and Federico Peinado. 2010. Ontological reasoning for improving the treatment of emotions in text. *Knowledge and Information Systems*, 24(2):421–443.

Phil Katz, Matthew Singleton, and Richard Wicentowski. 2007. SWAT-MP:The SemEval-2007 systems for task 5 and task 14. In *Proceedings of the Fourth International Workshop on Semantic Evaluations*, Prague, Czech Republic.

Michael Lesk. 1986. Automatic sense disambiguation using machine readable dictionaries: how to tell a pine cone from an ice cream cone. In *Proceedings of the 5th Annual International Conference on Systems Documen-*

tation, pages 24–26, New York, USA.

George A. Miller. 1995. WordNet: a lexical database for English. *Communications of the ACM*, 38:39–41.

Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the 42nd Meeting of the Association for Computational Linguistics (ACL 04), Main Volume*, pages 271–278, Barcelona, Spain.

Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? sentiment classification using machine learning techniques. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2002)*, pages 79–86, Philadelphia, USA.

W.G Parrot. 2001. *Emotions in Social Psychology: Essential Readings*. Psychology Press.

Siddharth Patwardhan, Satanjeev Banerjee, and Ted Pedersen. 2005. SenseRelate::TargetWord: a generalized framework for word sense disambiguation. In *Proceedings of the ACL 2005 on Interactive poster and demonstration sessions*, pages 73–76, Stroudsburg, USA.

James Pennebaker, Martha Francis, and Roger Booth. 2001. *Linguistic Inquiry and Word Count (LIWC): LIWC2001*.

R. Plutchik, 1980. *Emotion: Theory, research, and experience: Vol. 1. Theories of emotion*, chapter A general Psychoevolutionary Theory of Emotion, pages 3–33. New York: Academic.

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

Carlo Strapparava and Alessandro Valitutti. 2004. WordNet-Affect: an Affective Extension of WordNet. In *In Proceedings of the 4th International Conference on Language Resources and Evaluation*, volume 2004, pages 1083–1086, Lisbon, Portugal.

Peter Turney. 2002. Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL 2002)*, pages 417–424, Philadelphia, USA.

Janyce M. Wiebe, Rebecca F. Bruce, and Thomas P. O'Hara. 1999. Development and use of a gold-standard data set for subjectivity classifications. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics on Computational Linguistics*, pages 246–253, Stroudsburg, USA.

Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing*, pages 347–354, Stroudsburg, USA.

Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2009. Recognizing contextual polarity: An exploration of features for phrase-level sentiment analysis. *Computational Linguistics*, 35(3):399–433.