

Ontology-Supported Automated Mark Up of Affective Information in Texts

Virginia Francisco and Pablo Gervás

Departamento de Ingeniería del Software e Inteligencia Artificial
Universidad Complutense de Madrid, Spain
virginia@fdi.ucm.es, pgervas@sip.ucm.es

Abstract. This paper presents the application of an ontology of emotions to an existing approach for the automated mark up of affective information in texts. The emotional ontology has three main applications in this system: to select the most specific emotion which represents the affective information of the sentences from the probability that each word in the sentence has of indicating different emotions, to establish what emotion should be assigned to a sentence given the set of emotional assignments suggested by a group of evaluators, and to determine if the emotion assigned to the sentence during automated mark up is correct. The enhanced markup system has been tested and the results show improvements with respect to the previous version.

1 Introduction

Human languages have produced extremely powerful labels for emotional states, for example, English provides at least 107 emotion-denoting adjectives and German at least 235. To consider each of these categories as individual labels, not related with any other category, produces an uncontrolled proliferation of labels which multiplies the complexity of tasks which involve the use of these labels. If the emotional categories were not individual isolated units but units related with each other this might simplify tasks such as comparing two different emotional labels or deciding which is the emotion that better represents the generalization of two different emotions. A taxonomy of emotional categories where emotional labels are structured and organized in levels, from the most general to the most specific might provide a very useful tool in the treatment of emotional categories.

If we have two different emotional labels and we want to compare them, the different granularity of the labels could be an important aspect to determine if the two emotions are completely different, both are synonyms used to describe the same emotion, or both are different specifications of the same general emotion. For example, if we want to compare *gladness* and *joy* and we consider both emotional categories as individual units we will determine these two labels as indicating different emotions, but really *gladness* and *joy* are two emotion-denoting words provided by English to denote the same abstract emotion: *Happiness*. Another example, *sulking* and *annoyance* are specifications of the same abstract emotion *Anger* which differ only in the degree of arousal.

Another important application of an ontology of emotional labels would be the generalization of emotions, that is, we have two (or more) different emotional labels and we want to find a more general emotional category. In this case we will need a tool for identifying relations between emotional categories. For example, if we want an emotional label which generalizes the emotional labels *fright* and *dread* our the emotional ontology might suggest *Fear* as a possible generalization.

In this paper we present the application of an existing ontology of emotional categories to an approach for automated mark up of texts with emotional information. This approach marked up text with an extensive set of emotional categories so the ontology can be used to improve the treatment of emotions at three different levels: to select the most specific emotion which represents the affective information of the sentences from the probability that each word in the sentence has of indicating different emotions, to establish what emotion should be assigned to a sentence given the set of emotional assignments suggested by a group of evaluators, and to determine if the emotion assigned to the sentence by mark up application is correct.

2 Previous work

This section provides a brief outline of EmoTag, the existing approach for marking up text with emotions used in this paper, and the ontology of emotions.

2.1 EmoTag

EmoTag is a system that mark up texts with emotional labels. These labels are selected from a group of 92 emotional labels and EmoTag considers each of these categories as individual concepts not related with any other category. This uncontrolled proliferation of labels multiplies the complexity involved in tasks such as deciding which is the correct emotion for a sentence, marking the sentences automatically with an emotion, or deciding if a particular assignment of emotion to text is correct. If the emotional categories were not individual units but units related with each other this might improve significantly the results obtained by EmoTag.

EmoTag [1] relies on a dictionary of word-to-emotion assignments. This is obtained from a corpus of human evaluated texts by applying language analysis techniques such as stemming, POS tagging ¹ or dependency analysis [2]. Similar techniques are later applied to assign emotions to sentences from the assignments for the words that compose them.

Representation of Emotions There are different methods for representing emotions in research [3]: emotional categories, descriptions based on psychology, descriptions based on evaluation, circumflex models and emotional dimensions. EmoTag marks texts with *emotional categories* which is a method based in the use of emotion-denoting words, or category labels such as *happy, sad, angry ...*

¹ <http://www.english.bham.ac.uk/staff/omason/software/qtag.html>

Corpus Annotation . Each of the texts which forms part of the corpus may be marked by more than one person because assignment of emotions is a subjective task so the “subjective extremes” must be avoided. Once all the texts of the corpus are marked up by each of the evaluators we have to obtain the average emotion for each of the sentences from the emotions selected by the evaluators. In order to get the reference value EmoTag selects the emotion most often assigned to the sentence by the evaluators, not taking into account the different granularity of emotional concepts or the relations between them.

Construction of the Dictionary Based on the tales marked up by human evaluators EmoTag obtains a data base of words (List of Emotional Words or LEW) which includes their relation with emotional dimensions and categories. In order to obtain the LEW list the first step is to obtain the relevant words (words whose part-of-speech is not included in a POS tag stop list). For each of the relevant words EmoTag obtains its stem and associated emotion. Next, the system computes the complement of the emotional content of the words under the scope of negations, and it inserts the final values into the LEW list. Finally, it carries out a process of normalization and expansion.

A Method for Automated Mark Up of Emotions EmoTag classifies sentences into emotions based on the relation between words and different emotions. This process is carried out in the following way:

- Perform sentence detection and tokenization and obtain: the words affected by negations, the stem and the part-of-speech of each of the words.
- Obtain the emotional value associated to each word by looking it up in the affective dictionary (LEW list).
- Process the words under the scope of negations.
- Obtain the final value of the sentence based on the emotions associated to the words which compose it. Once all the words of the sentences have been evaluated, the probability of conveying each emotion for the different words are added up and the emotion which has a higher probability is assigned to the sentence. This process does not consider the different granularity of emotional concepts. If we have several generalizations of the same emotional concept, each of its probabilities is considered individually.

A sample part of a marked tale is given below:

```
...
<neutral>The knight threw the spear. </neutral>
<sad>It killed the fierce lioness. </sad>
...
<happy>The knight resurrected the pretty blonde princess. </happy>
<delight>She returned to the strong castle. </delight>
<happy>The knight and the princess lived happy ever afterwards. </happy>
```

2.2 Ontology of Emotions

There is a general agreement that there are some emotions that are more basic than others [3]. Scherer [4] suggests that an emotion A is more fundamental than other emotion B if the set of evaluation components of the emotion A is a subset of the evaluation components of the emotion B. Based on this premise emotional categories can be structured an ontology by organizing the labels in levels, from the most general to the most specific. The application of such ontology might provide the means of improving significantly the results of the tasks in which different emotional categories take part, as well as a method for identifying relations between them.

The ontology of emotional categories is explained in detail in [5]. This ontology is written in OWL [6] and uses Pellet² as reasoner. Emotional categories are structured as a taxonomy that covers from basic emotions to the most specific emotional categories. The ontology has two root concepts:

- *Emotion*: This is the root for all the emotional concepts which are used to refer to emotions. Each of the emotional concepts are subclasses of the root concept Emotion. Some examples of these subclasses are: *Happiness*, *Sadness*, *Fear*, *Envy* . . .
- *Word*: This is the root for the emotion-denoting words, the specific words which each language provides for denoting emotions. The ontology is currently available for two different languages: English and Spanish. In order to classify the words into their corresponding language the root concept Word has two subclasses: *EnglishWord* and *SpanishWord*.

Figure 1 shows a fragment of the ontology. In this fragment it can be seen how the words are related both to one emotional concept and to one word concept, for example the word *cheerfulness* is an instance of the emotional concept *Happiness* at the same time it is an instance of the word concept *EnglishWord*, which means that *cheerfulness* is an English word for denoting the emotion *Happiness*.

In handling words, the system may need to identify synonyms for particular words. Two instances of the Word-concept can be considered to be synonyms if they are also instances of the same single Emotion-concept from the parallel Emotion subhierarchy. For example, in the figure above, we can find that the words *amazement* and *astonishment* may be considered synonyms because they are both instances of the Emotion-concept *Amazement*.

To summarize we can conclude that the emotional ontology represents the emotional categories as instances of a tree structure of emotional concepts. Each emotional word is instance of two concepts: an emotional concept which represents the emotion denoting by the emotional word and a word-of-a-particular-language concept which determines the language to which the word belongs. From a given emotion-denoting word by means of our ontology we obtain the direct emotional concept associated to it as well as the more general emotional concept related to the direct emotional concept. We can get, too, the synonyms

² <http://pellet.owldl.com/>

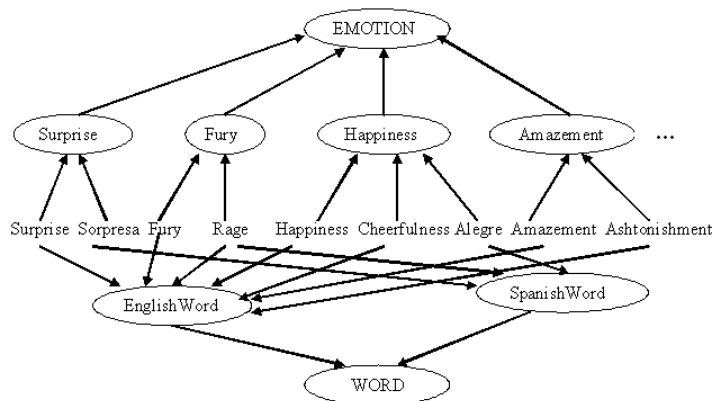


Fig. 1. Fragment of the emotional ontology

for an emotional word and the corresponding word in other language. For example, given the emotional word *grief*, we have as direct emotional concept, *Grief*, as general emotional concepts, *Distress*, *Sadness* and *Emotion*, as Spanish translation, *Agonía*, and finally as synonyms, *agony*, *anguish* and *sorrow*.

3 Application of the Emotional Ontology in EmoTag

The emotional ontology has three main fields of application in our approach for automated marked up of Emotions (EmoTag). It can help us to:

1. Select the appropriate emotion for a sentence based on the emotions associated with each word, and the probability of that association. In order to select the emotion assigned to a sentence, EmoTag considers the emotions as independent units, so the emotion with a higher resulting probability is the emotion assigned to the sentence. This is a good first solution but it can be improved by distinguishing the different granularity of emotional concepts and the relation between emotions derived from this taxonomy. For example, if we have that emotion A is a generalization of emotion B and C it could be an interesting approach to consider the probability of emotion A as the addition of the probability of emotion B and the probability of emotion C.
2. Obtain the reference value for the emotion of each sentence from the emotions selected by evaluators. In the original version of EmoTag the reference value (value most often assigned to a sentence by the evaluators) did not take into account the emotions as related units as in the previous case. A similar use of the ontology could improve the agreement between evaluators if we look not only for the choice of a specific value but for the choice of a similar value (similarity in the ontology).

- Evaluate the results. In order to determine how well a text is marked up, EmoTag considers that each sentence is successfully tagged if the emotion assigned by the tagger matched the reference value mentioned in the previous point. This way a sentence only can be correctly marked (1) or incorrectly marked(0), but there are other intermediate values that it may be interesting to consider. If the sentence is marked with an emotion A which is a generalization of the emotion B used by the tagger, it would be wrong to consider this sentence as incorrectly marked (0).

3.1 Using the Ontology to Identify Most Frequently Assigned Emotions

In corpus used by EmoTag a *reference value* for the emotion of each sentence is obtained by choosing the emotion most often assigned to that sentence by the human evaluators. This value is used to compare with the results generated by our tagger. This is a good first solution but it can be improved, because of the different granularity of emotional concepts. In the case of the assignment of emotions by evaluators to a sentence shown in Table 1, the emotion *sadness* is taken as the reference emotion for evaluation, but it is not the best reference value. If we consult the emotional ontology we can see that *agony*, *anguish*, *grief* and *sorrow* are synonyms which are related to the emotional concept *grief*, so the best reference value for that sentence is *grief* instead of *sadness*.

Eval 1	Eval 2	Eval 3	Eval 4	Eval 5	Eval 6
agony	anguish	grief	sorrow	sadness	sadness

Table 1. Example of set of emotion assignments for different evaluators.

In this new version of EmoTag we consider emotions as instances of the ontology and we select the most specific emotion supported by at least half of the evaluators. More precisely we carry out the following process:

- If at least half of the evaluators agrees on the assignment of one emotion, this emotion is taken as the reference value for the sentence as in the previous version.
- Otherwise, we group the emotions in levels (according to the level of the emotional concept they refer to). We obtain the related concepts for the emotion with the lower level. Once these new concepts are added to their corresponding levels, if we have any emotion supported by at least half of the evaluators we take it as the reference value. In the case of two emotions supported by most of the evaluators we get the emotion with a lower level.
- Otherwise we repeat the previous step for each of the levels in ascending order.

Figure 2 shows an example of this process. In this example we have a sentence marked up by six evaluators. There is no emotion supported by at least half of the evaluators so we follow the process explained above. First, we group the emotions in levels shown in the second table, then we obtain the related concepts for *Grief* (*Distress* and *Sadness*), *Helpless* (*Powerless* and *Sadness*) and *Remorse* (*Regret* and *Sadness*) which are the emotions with the lower level. Third, we add these new concepts to the previous ones, in this case we add *Sadness* which is supported by 4 evaluators (Eval 1, Eval 3, Eval 4 and Eval 5), *Distress* which is supported by 2 evaluators (Eval2, Eval3), *Powerless* supported by Eval4 and *Regret* supported by Eval 5. Finally, we take *Sadness* as the reference value because it is supported by 4 evaluators which is more than half of the evaluators.

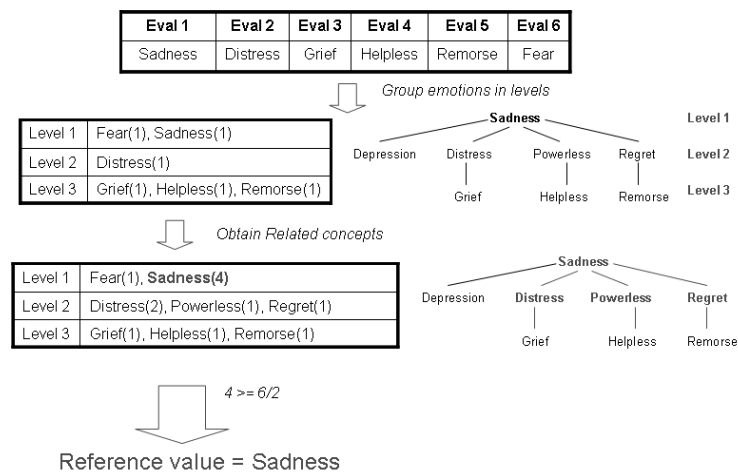


Fig. 2. Example of the reference value obtained from the assignment of emotions by the evaluators.

3.2 Ontology-Supported Automated Mark Up of Emotions

As mentioned in Section 2.1, in order to mark up a sentence with an emotion EmoTag splits it into words and it assigns to each one the probability of carrying the different emotions. Based on these probabilities of the words EmoTag obtains the final emotion of the sentence as follows. Once all the words of the sentences have been evaluated, it adds up the probability of each emotion of the different words and assigns to the sentence the emotion which has a higher probability. In this case we have a similar problem as in the previous section: if we have several generalizations of the same emotional concept, each of its probabilities are considered individually. In this case it will be better to consider their probabilities together under the more general concept. An example of this is the following case,

suppose we have the emotions and their corresponding probabilities showing in the Table 2:

Anger	Indignation	Sulking	Displeasure	Happiness	Happiness
0.2	0.1	0.05	0.1	0.1	0.2

Table 2. Example of probability assignment to a set of emotions.

In this example EmoTag considers that the emotion with the higher probability is *happiness* (0.3), but with the emotional ontology we can see that *indignation*, *sulking* and *displeasure* are generalizations of the emotional concept *anger*. With this information we can determine that the emotion with the higher probability is *anger* (0.45) which seems to be a better result.

The process we follow in the new version of EmoTag using our ontology is the following: once all the words of the sentences have been evaluated EmoTag adds up the probability of each emotion of the different words, and it carries out the following process for each of the possible emotions:

- Process all the emotions in order to obtain the related emotional concepts, the parents of the emotion in the ontology.
- The related emotional concepts are added to the previous ones with the probability associated to the more specific concept.
- Emotions are grouped by their corresponding level in the ontology.
- The more general emotion (with the lower level in the ontology) with the higher probability is assigned to the sentence.

Figure 3 shows an example of this process. In this example we have a sentence with three emotional words, each word has the probabilities shown in the first table. The first step is to get the level of each emotion and put all the emotions in the second table with their probabilities. Second we obtain for each emotion its related concepts (*Indignation - Anger*, *Sulking - Anger*, *Displeasure - Anger* and *Amazement - Surprise*) which are added to the third table with the probability of its child ($Anger = 30\%$ -previous probability- + 20% -*Indignation* probability- + 15% -*Sulking* probability- + 30% -*Displeasure* probability = 95% , $Surprise = 65\%$ -previous probability- + 5% -*Amazement* probability- = 70%). Finally *Anger* which is the more general emotion with the higher probability is assigned to the sentences.

3.3 Ontology-Supported Evaluation of Mark Up

In order to determine how well a text is marked up, EmoTag considers that each sentence is successfully tagged if the emotion assigned by the tagger matched the reference value obtained in the previous section. This way a sentence can only be correctly marked (1) or incorrectly marked (0), but with the emotional ontology we can determine how well a sentence is marked up with a finer granularity.

Emotion	Anger	Fear	Surprise	Sadness	Indignation	Sulking	Displeasure	Amazement
Word 1	30%	40%	0%	30%	0%	0%	0%	0%
Word 2	0%	20%	50%	10%	20%	0%	0%	0%
Word 3	0%	20%	15%	15%	0%	15%	30%	5%
Total	30%	80%	65%	45%	20%	15%	30%	5%

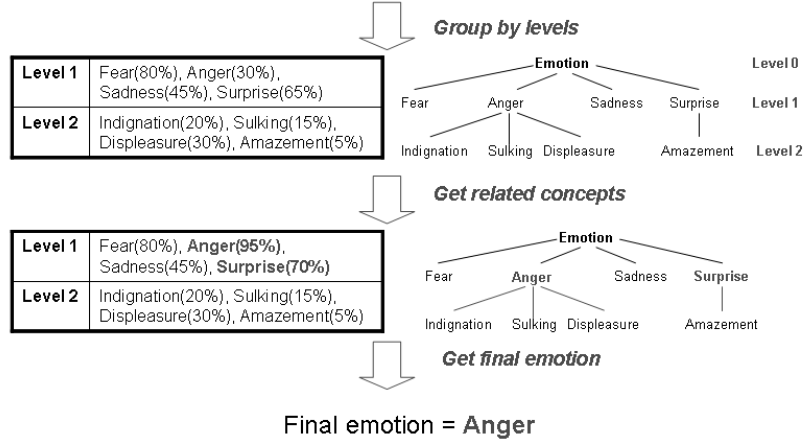


Fig. 3. Example of the assignment of the final emotion to a sentence based on the emotions of the words which compod it.

In the new version of EmoTag we assign to each sentence a score, which ranges from 0.0 to 1.0, in order to determine how well the sentence is marked up. The score is obtained, based on the number of levels in common between the two emotions and the level of the more specific emotion, as follows:

$$Correct = \frac{number_common_levels}{level_specific_emotion}$$

For example, in the case of a sentence marked by the evaluators as *excitement* and by EmoTag as *enthusiasm*, as can be seen in Figure 4, we have that the emotional category *excitement* is related to the emotional concept *Excitement* in the emotional ontology, this is the direct emotional concept. *Excitement* (level 3) is related to *Enthusiasm* (level 2) and this is related to *Happiness* (level 1). On the other hand, we have *enthusiasm* which is related to the emotional concept *Enthusiasm* in the ontology, its direct emotional concept, and *Enthusiasm* (level 2) is related to *Happiness* (level 1). To summarize, *enthusiasm* has level 2, *excitement* has level 3 and *enthusiasm* and *excitement* have 2 emotional concepts in common: *Enthusiasm* and *Happiness*, so the result is: $correct = 2/3 = 0.67$.

4 Evaluation of Ontology-Enhanced version of EmoTag

In order to evaluate our work we carried out the same tests as in EmoTag. In these tests four texts are going to take part. Each of our four texts will be marked

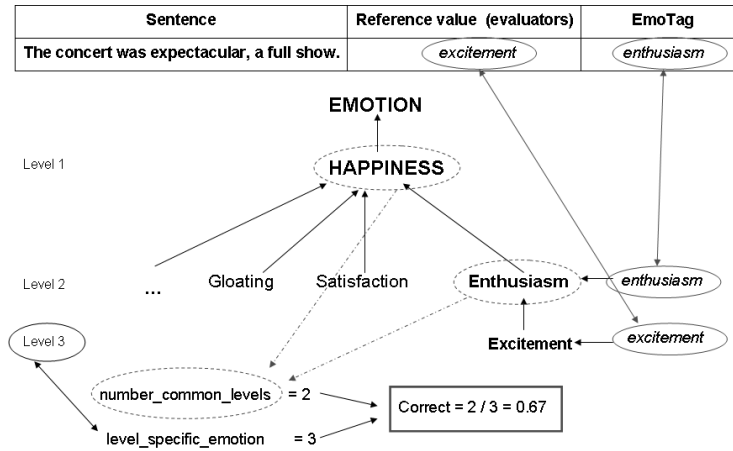


Fig. 4. Example of the assignment to each sentence a score.

with the emotional categories by different evaluators. Then by marking up every emotion with our process, we will obtain two measures: results obtained from the human evaluators (or evaluator's results), and results obtained by the tagger (tagger's results).

- Evaluator's results: A reference value for the emotion of each sentence is obtained by choosing the emotion most often assigned to that sentence by the human evaluators, using the emotional ontology as explained above.
- Tagger's results: The reference value obtained in the evaluator's results is used to compare with the results generated by our tagger. The graph in Figure 5 shows the percentages of success obtained for each text and the percentage of sentences incorrectly annotated which correspond to a sentence in which the majority of the evaluators did not agree.

4.1 Using the Ontology to Identify Most Frequently Assigned Emotions

The percentage of sentences on which the majority of the human evaluators - half of their number plus one - agrees on the assignment of an emotion is higher than in the previous version, around 70% (as opposed to, around 45%). This is an important improvement of the new version.

4.2 Ontology-Supported Automated Mark Up of Emotions

In Figure 6 we compare the results obtained for the new version with the ones obtained for the previous version. We can conclude that the emotional ontology increase the percentage of sentences correctly marked around a 15%.

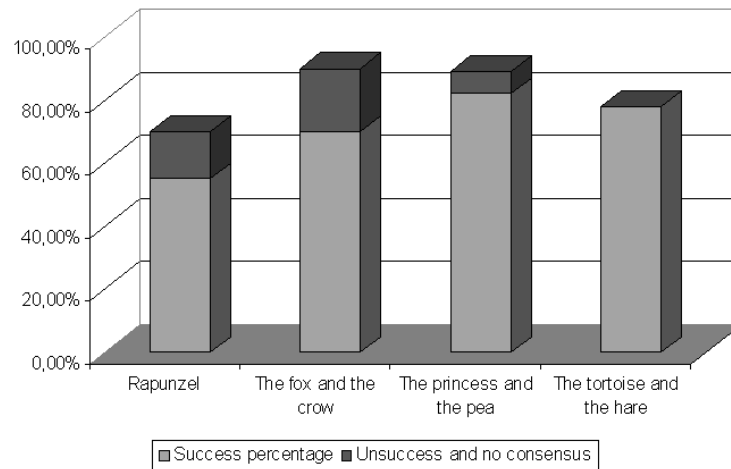


Fig. 5. Percentage of success in automated tagging with new version of EmoTag, with the emotional ontology.

4.3 Ontology-Supported Evaluation of Mark Up

Only the sentences with a correct level of 1 have been taken into account in the evaluation as correct sentences, because the percentage of sentences with a lower value is so low that can be disregarded. We can conclude that the use of an emotional ontology for the comparison of the tagger value with the reference value does not improved significantly the results.

5 Conclusions

The use of a ontology of emotions is a good solution for improving the results obtained for EmoTag. We have employed the ontology in three different tasks: marking up of sentences with emotions, obtaining the reference value of each sentence as the emotion assign to it by the majority of the human evaluators and determining if a sentence is correctly marked up. After analysis of the results obtained with the new version of EmoTag we can conclude that the ontology improves the percentage of sentences correctly marked by 15% and it improves the percentage of reference value supported by the majority of evaluators by 25%. The only change in the new version which has not improved results is the application of the ontology in the determination of the correctness of a sentence, because the percentage of sentences with a value between 0 and 1 is so low that it can be disregarded.

An interesting improvement of our tests might be a comparison with other marking tools but at present other existing mark up tools tag texts with emotional dimensions or with a reduced set of 5 or 6 emotional categories. The




Fig. 6. Relative improvement in mark up success of the ontology supported version with respect of the previous one.

emotional ontology is useful only in the case that an extensive set of emotional categories take part in the mark up process as in EmoTag.

As future work we will apply the emotional ontology in order to make EmoTag more flexible and more adaptable to the different environments of application which can benefit from our marker. Using the emotional ontology EmoTag will mark up text with basic emotional categories or with more specific emotional categories.

References

1. V. Francisco, P.G.: Automated Mark Up of Affective Information in English Texts. In: International Conference on Text, Speech and Dialogue (TSD06), Brno, Czech Republic, Springer, LNCS series (2006)
2. Lin, D.: Dependency-based Evaluation of MINIPAR. In: Proc. of Workshop on the Evaluation of Parsing Systems, Granada, Spain (May 1998)
3. Cowie, R., Cornelius, R.: Describing the Emotional States that are Expressed in Speech. In: Speech Communication Special Issue on Speech and Emotion. (2003)
4. Scherer, K.: On the Nature and Function of Emotion: A Component Process Approach. Scherer and K.R. and Ekman P and editors, Erlbaum, Hillsdale, NJ (1984)
5. Francisco, V., Gervás, P., Peinado, F.: Ontological Reasoning to Configure Emotional Voice Synthesis. In: First International Conference on Web Reasoning and Rule Systems (RR 2007), Innsbruck, Austria, Springer, LNCS series (2007)
6. Bechhofer, S., Harmelen, F.V., Hendler, J., I.Horrocks, McGuinness, D., Patel-Schneider, P., Stein, A.: OWL Web Ontology Language Reference. Volume 12/2. W3C (2004)