Chapter 1

Semantic Web Approaches to the Extraction and Representation of Emotions in Texts

Virginia Francisco Federico Peinado Raquel Hervás
Pablo Gervás
Departamento de Ingeniería del Software e Inteligencia Artificial
Facultad de Informática
Universidad Complutense de Madrid
virginia@fdi.ucm.es, email@federicopeinado.com, raquelhb@fdi.ucm.es,
pgervas@sip.ucm.es
Tel.: (+34) 91 394 7599
Fax: (+34) 91 394 7529

Keywords: Ontologies, Affective Computing, Automatic Text Annotation, Natural Language Processing, Knowledge Representation.

[PLEASE REVIEW THE CHAPTER FOR ENGLISH STYLE ISSUES]

Abstract

This chapter presents the role of ontologies in the field of Affective Computing. For illustrating the main concepts, we analyse a particular emotional ontology, designed and implemented using Semantic Web technologies, and study its use for the representation, recognition and generation of emotions that are present in textual documents.

The emotional ontology developed has been used in three different applications, which will be explained in this chapter, related to the automated markup of text with emotional content. The first one is selecting the most specific emotion that represents the general affective content of a sentence according to the emotion that each word represents. The second one is establishing what emotion should be assigned to a sentence given the set of emotional assignments suggested by a group of annotators. The third one is determining whether the emotion assigned to the sentence in automatic emotion markup is correct. These ontology-enhanced markup applications have been tested and the results show improvements with respect to the previous version created without Semantic Web technology.

There are various approaches for representing emotions. The most used representations are *emotional categories*, where fixed words are used to refer to emotions, and *emotional dimensions*, where emotions are defined using numerical values to express different aspects of emotion. To address the problem of translating dimensions into categories, a common problem in this type of applications, each of the categories of our emotional ontology is related to three different dimensions (i.e. evaluation, power and control) by means of data ranges. By applying automated classification in terms of the emotional ontology, the system can interpret unrestricted inputs in terms of a restricted set of concepts for which particular rules are provided. An example of application of this use of the ontology is presented in which the rules applied for document markup with emotional dimensions provide configuration parameters for an emotional voice synthesizer.

1. Introduction

Affective Computing is defined as the subfield of Computer Science that deals with representation and processing of emotions. This field is gaining importance as human-computer interaction systems become more sophisticated and a wide spectrum of promising applications appear in areas such as virtual reality, robotics and intelligent interfaces (Tao and Tan, 2005).

Affective Computing has always been a field strongly dependent on knowledge and semantics. Whether the goal is to organize emotions using different representations, or to understand or transmit certain emotions in a text, it is necessary to have semantic information to relate meanings of emotions and represent these meanings using several knowledge representation paradigms. In this chapter we study how an ontology can be useful for describing emotional information and we present some applications of this ontology for the automatic extraction and mark up of emotions in a text.

1.1. Ontologies for Emotion Representation

Existing approaches for representing emotions in the Affective Computing field vary from identifying a set of *emotional categories* (with a name tag assigned to each one of them) to designing a multi-dimensional space in terms of *emotional dimensions* (primitive elements) such that any particular emotion can be defined in terms of a tuple of numerical values along the different dimensions. For different purposes, one approach may be better suited than the other. Emotional categories and emotional dimensions are explained in detail next in Section 2.1.1.

One important problem to face with emotional categories representation arises from the fact that there is a large number of emotional categories as 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, and the differences and similarities between them are not clear cut. 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 equivalent, different or similar 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 indicative of different emotions, but in fact *gladness* and *joy* are two emotion-denoting words provided by English to denote the same abstract emotion: *happiness*. Again, if we compare *sulking* and *annoyance*, apparently two different emotions, we found that both are specifications of the same abstract emotion *anger* which differs only in their arousal degree. To consider each of these categories as individual labels (not related to any other category) produces an uncontrolled proliferation of labels which multiplies the complexity of tasks that involve the use of these emotions. If the emotional categories were not individual isolated units but units related to 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.

A separate problem arises when operating computationally with emotions expressed in more than one representation. In this case we would face the task of being able to find equivalences from one to another. This task is reasonably easy when changing from emotional categories to emotional dimensions: it would suffice to assign a particular tuple of numerical values for the emotional dimensions of each emotional category. When trying to find an equivalence from emotional values expressed in terms of emotional dimensions to a representation in terms of emotional categories this is not so simple. The problem lies in the fact that, given the subjectivity associated with emotional perception, the particular values assigned to a given impression by one person usually deviate slightly from what a different person would have assigned. This suggests that the process of converting from emotional dimensions to emotional categories should be carried out in a manner that allows a certain tolerance, so that a region of space in the universe of emotional dimensions is assigned to each emotional category, rather than just a single point in the universe. An ontology of emotional categories in which each of the categories is related to the emotional dimensions by means of data ranges might be a useful tool for the interpretation of emotions expressed in terms of emotional dimensions as emotional categories.

In this context, the development of an ontology of emotional categories based on description logics, where each element is defined in terms of a range of values along the space of emotional dimensions, provides a simple and elegant solution to the problems presented above. Thanks to a taxonomical reasoning system an implicit hierarchy for the concepts represented in the ontology can be inferred automatically. The ability to carry out automatic classification of concepts simplifies the addition of new concepts - possibly expressed only in terms of their values along the axes of emotional dimensions - without having to worry explicitly about where in the ontology they should be placed.

1.2. Ontologies for Emotion Recognition and Generation in Texts

Two very promising areas in the field of Affective Computing are the recognition and generation of emotions expressed in natural language. Getting corpora marked up with emotions is essential to most attempts to make human-computer interaction respond to the emotional aspects that are present in any human activity. In particular, the identification of emotions within a text is useful for studying how people relate textual information with emotions. Once these emotion-text relations are extracted, it is possible to train automatic systems that assign emotional contents to raw texts. In order to fulfill these tasks, semantic information about the relations between different emotions is required.

The use of an ontology of emotional categories could improve the performance of systems that deal with the task of automated mark up of texts with emotional information. This kind of ontology could be used to improve the treatment of emotions at three different levels within the process of automatic mark up of texts with emotions: to establish what emotion should be considered for a sentence given the set of emotional assignments suggested by a group of annotators, 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, and to determine if the emotion assigned to the sentence by a mark up application is correct.

The first step towards the automatic treatment of emotions is to understand how people annotate texts with emotional information. For this task, a corpus of texts annotated by different evaluators must be used. As emotions are subjective feelings that different people experience in different ways, some kind of process must be used to find which emotion would be considered as the most appropriate for the annotators as a whole. However, this problem is a very hard one. Even when it is possible that different annotators have experienced more or less the same emotion when reading a sentence, it is probable that they will use different words or categories to refer to it. For example, if a sentence can be considered as transmitting *happiness*, two different evaluators might use *gladness* and *joy* respectively for expressing it. It is therefore required to have some kind of semantic information about how different emotions are related to be able to identify situations like this. An ontology of emotional categories as the one we are proposing in this chapter is very useful for dealing with this kind of problems.

Once a corpora of texts marked up with emotional information is available, the mark up system can be trained with it by finding correspondences between pieces of texts and emotions. When performing the mark up of a new text using the learned information, an ontology of emotional information can be extremely useful for performing inferences over the available information. For example, we can find a sentence with a very high probability of being *happy*, but with low probabilities of being *angry*, *outrageous*, *sulking* and *displeasuring*. If the sum of the probabilities of these four emotions is higher than the probability of *happiness*, we can consider that the sentence is more *angry* than *happy*. This kind of reasoning would only be possible if an ontology of emotions representing this kind of relations is available.

Finally, an emotional ontology would also be very helpful when determining the performance of the automated mark up system. The kind of reasonings that have been explained before for the creation of the corpora and the automatic annotation of texts can be also quite useful for determining if the solution found by the mark up system is more or less similar than the emotion chosen by the initial annotators.

In addition, by reasoning over the emotion ontology, insertion of new instances of emotional concepts into the ontology results in their automatic classification under the corresponding branch of the hierarchy. The system can then trace the ascendants in the ontology of the corresponding value, until a more general concept is found, in applications where different emotional specification degrees are required. This could be useful for some applications that require the use of general emotions instead of more specific ones. An example would be a voice synthesizer that relies on rules based on a reduced set of emotions for generating an appropriate voice synthesis configuration for expressing the intended emotional impression.

The rest of the chapter is structured as follows. Section 2. presents a brief review of computational representation of emotions, Semantic Web technologies that have been employed, existing ontologies of emotions and a review of the system chosen as domain of application. Section 3. presents the emotional ontology created. Sections 4. and 5. present the application of the ontology in the mark up of texts with emotional categories and in the interpretation of emotional dimensions as emotional categories, respectively. Finally, in Section 6. we conclude by discussing how the emotional ontology is a useful resource for marking up text with emotions and for interpreting emotional dimensions as emotional categories. This last section presents, as well, some lines of future work that should be faced in order to apply and improve the obtained results.

2. Related Work

To understand the work presented here, four basic topics must be covered: computational representation of emotions, Semantic Web technologies that have been employed in the project, existing ontologies of emotions and a review of the system that has been chosen as example for the use of our ontology. These topics are covered in the following subsections.

2.1. Computational Representation of Emotions

There are plenty of emotional models proposed in the academic literature, but there seems to be agreement about the fact that emotions are subjective experiences each person feels in a very different way. Following this idea, Parrott (2001) defines emotions as a response to events that are important for us, which are governed by specific laws and that emerge and manifest themselves according to what the operating mechanism of these laws dictates. In the context of this chapter the *emotional states* defined by Cowie and Cornelius (2003) are the type of emotional content considered.

2.1.1. Classification of Emotions

In order to study emotional states we must decide how they are going to be represented. There are different ways to represent emotional states, but two of them are the most commonly accepted: the use of *emotional categories* and *emotional dimensions*.

Emotional Categories: The most common method for describing emotions is the use of emotional words or affective labels. Different languages provide assorted words of varying degrees of expressiveness for the description of emotional states. That is why several approaches have been proposed to reduce the number of words used to identify emotions, for example with the use of *basic emotions*, *super ordinate emotional* *categories* or *essential everyday emotion terms*. The use of *basic emotions* considers that there are some emotions that are more known and understandable for everybody than others (Cowie and Cornelius, 2003). In the *super ordinate emotional categories* approach some emotional categories are proposed as more fundamental, arguing that they subsume the others (Scherer, 1984). Finally, the *essential everyday emotion terms* approach focuses on emotional words that play an important role in everyday life (Cowie, Douglas-Cowie, and Romano, 1999).

Emotional Dimensions: Emotional Dimensions are measures that try to model the essential aspects of emotions numerically. Although there are different dimensional models with different dimensions and numerical scales (Fontaine et al., 2007), most of them agree on three basic dimensions called *evaluation*, *activation* and *power* (Osgood, Suci, and Tannenbaum, 1957). *Evaluation* represents how positive or negative an emotion is. At one extreme we have emotions such as *happiness*, *satisfaction* and *hope* while at the other we find emotions such as *unhappiness*, *unsatisfaction* and *desperation*. *Activation* represents an activity versus passivity scale of emotions, with emotions such as *excitation* at one extreme, and *arousal* and at the other emotions such as *calm* and *relaxation*. *Power* represents the sense of control which the emotion exerts on the subject. At one end of the scale we have emotions characterized as completely controlled, such as *fear* and *submission* and at the other end we find emotions such as *dominance* and *contempt*.

For different purposes, one approach is better suited than the other. For instance, when attempting to synthesize voice utterances that reflect emotion to some extent, it is easier to identify the parameters for voice production associated with conveying a particular emotion category. For assigning emotional values to given utterances, on the other hand, human evaluators find it much easier to provide numbers along given dimensions.

2.1.2. Structure of Emotions

Psychologists have been searching for an intelligent way to structure the existing emotional repertoire. Several methods have been proposed, each with its own advantages and disadvantages.

Methods based on emotional dimensions try to capture the similarities and differences between emotions. Some researchers propose a two-dimensional space, exclusively considering evaluation and activation. This is called the *circumflex model* where the points that correspond to all possible emotions form a circle (Russell, 1980). Viewing the multitude of emotions as points in a two-dimensional space can be useful in understanding the most generic emotions but not the most specific ones. This model ignores the variety of emotional states, and does not capture the slight differences found beyond the most generic sensations.

As an alternative to dimensional spaces some researches have used cluster analysis (Storm and Storm, 1987; Shaver et al., 1987; Parrott, 2001). These approaches group emotions into clusters, with the number of clusters dependent on each specific approach. Shaver et al. (1987) proposes the use of five clusters called *affection*, *happiness*, *sadness*, *anger* and *fear*. Parrott (2001) presents a more detailed list of emotions categorized into a short tree structure. This structure has three levels for primary, secondary and tertiary emotions. As primary emotions, Parrot presents *love*, *joy*, *surprise*, *anger*, *sadness* and *fear*. Secondary emotions specialize primary emotions, e.g. *love* has *affection*, *lust* and *longing* as secondary emotions. Finally, tertiary emotions specialize secondary emotions, e.g. *lust* is a secondary emotion with *arousal*, *desire*, *passion* and *infatuation* as tertiary emotions. Instead of grouping emotions using their global similarity, other researchers preferred to group emotions based on other criteria such as the components of their appraisals (Scherer, 1984) and the events that give rise to them (Ortony, Clore, and Collins, 1988).

To summarize, there are different ways for structuring emotions and each approach could be useful for a different purpose. Any approach that tries to be useful for a variety of applications should take advantage of these different representations of emotions.

2.2. Semantic Web Technologies

The Semantic Web is being developed with the intention of providing a global framework for describing data, its properties and relationships in a standard fashion. Many developers and researchers on knowledge systems are taking the approach of using Semantic Web technologies in order to obtain more interoperability and reusability with existing software and to take advantage of the strong trend of development that these technologies are experiencing nowadays.

In this section we review the tools used in our project explaining what were the technological choices and the different criteria behind them.

2.2.1. Ontology Web Language

The Semantic Web relies heavily on ontologies. The most common language to formalize Semantic Web ontologies is OWL (Ontology Web Language [1]), a recommendation of the W3C. The goal of this standard is to formalize the semantics that were created *ad hoc* in old frame systems and semantic networks. OWL has three increasingly-expressive sublanguages: OWL Lite, OWL DL, and OWL Full.

OWL Lite is the simplest subset of OWL, specially designed to provide a quick migration path for other taxonomical structures.

OWL DL is the subset of OWL designed for applications that need the maximum expressiveness without losing computational completeness and decidability. It is based on Description Logics, a particular fragment of first order logic, in which concepts –OWL classes–, roles –OWL properties–, individuals and axioms that relate them (using universal and existential restrictions, negation, etc.) are defined. These entailments may be based on a single document or multiple distributed documents that we combine using the import OWL mechanisms. The OWL DL reasoning capabilities relies on the good computational properties of DLs. OWL DL has support for polyhierarchical automatic classification.

OWL Full ignores some significant restrictions of OWL DL, becoming a more powerful language for representing complex statements, but less useful for reasoning with them due to their computational properties.

Nowadays, the newest version of OWL, OWL 2 has three different sublanguages: OWL 2 EL (oriented to maximize the expressivity), OWL 2 QL (oriented to query answering) and

[1] OWL 1: `http://www.w3.org/TR/owl-ref/` and OWL 2: `http://www.w3.org/TR/owl2-overview/`

OWL 2 RL (oriented to rule systems). To understand the ideas of this chapter it is not necessary to update the terminology, so we still refer to the well-known OWL 1 sublanguages.

2.2.2. Frameworks and APIs

An all-in-one framework or a versatile application programming interface would be a very desirable tool for any novice Semantic Web application developer. Java is probably the most important general-purpose language for developing Semantic Web applications, and it is also the language in which the original system we used for our experiments was made, so the choice was obvious. But there are at least two very promising Java frameworks available. One of them is Sesame [2], an open source RDF framework with support for RDF Schema inferencing and querying. The other one is Jena [3], another open source framework with a programmatic environment for RDF, RDFS, OWL, SPARQL and its own rule-based inference engine.

Sesame has a local and remote access API, several query languages (included SPARQL) and it is more oriented to offer flexible and fast connections with storage systems.

Jena has also RDF and OWL APIs, tools to deal with RDF/XML, N3 and N-Triples formats, an SPARQL query engine and also some persistent storage functionality. It is important to consider that Jena is a useful tool for exploring the strong relation between SPARQL queries and OWL-DL ontologies (Jing, Jeong, and Baik, 2008).

For our purposes performance issues can be ignored and only inference support for Description Logics is taken into account. The architecture of Sesame is probably easier to extend than the architecture of Jena, but from the point of view of the client building a wrapper for the functionality of the underlying framework, Jena is the most intuitive and usable API. Because of that we chose Jena.

The framework that has been selected to carry out the development work in this chapter is a short extension of DLModel's [4] functionalities. DLModel is a very straightforward open source API for accessing a Description Logic model instantiated in a set of ontologies, a set of DL-Safe rules (Motik, Sattler, and Studer, 2005) and a knowledge base. Although it has an abstract DL interface, it can be viewed as a wrapper on top of Jena that offers simpler methods to access concepts, roles, attributes, annotations and individuals of the knowledge base from any Java application. This management includes DL reasoning and DL-Safe rule reasoning as implicit services that will benefit the programmer acting as an "intelligent black box" that knows how to use Semantic Web low-level technologies.

2.2.3. Ontology Editor

Another important tool is the Integrated Development Environments (IDE) used to edit the ontology and the knowledge base. During our review of the state-of-art we found two interesting editors able to perform this task: SWOOP and Protégé.

SWOOP [5] is a hypermedia-based OWL ontology browser and editor written in Java. It is open source and it tries to simplify ontology development using an interface similar

[2] http://www.openrdf.org,currentlyrelaunchedasOpenJena:http://openjena.org/
[3] http://jena.sourceforge.net/
[4] http://federicopeinado.com/projects/dlmodel/
[5] http://www.mindswap.org/2004/SWOOP/

to a web browser. It includes some advanced features as ontology partitioning, debugging and different kinds of visualization, so it makes ontologies more scalable, maintainable and easy to use.

Protégé [6], specially the Protégé-OWL version, focuses on editing OWL ontologies. It is a powerful Java open source tool with a user-friendly interface that lets you edit and visualize ontologies in a very easy way. It can be seen as a framework for developing Semantic Web applications itself. The number of plugins (including some plugins for knowledge acquisition), the stability of the last version, and the extensibility of its architecture (plug-and-play environment) allow rapid prototyping and application development, just what we were looking for.

Ontology building is still more of a craft than an engineering task and that is the reason why tools as Protégé-OWL, which is fundamentally open to new collaborative features gaining importance within the Semantic Web paradigm, are a better choice for us.

2.2.4. Reasoner

Two different reasoners were considered for this project: Pellet (Mindswap, 2006) and Racer Pro [7].

Pellet is an open source DL reasoner completely implemented in Java. It deals not only with taxonomical reasoning but also with datatype reasoning, DL-Safe rules implemented in SWRL and other features that were considered very important for our project. Pellet is the default reasoner integrated with SWOOP.

RacerPro is an OWL reasoner and inference server for the Semantic Web. It is a well-known system for OWL/RDF which claims to be the most efficient and robust DL reasoner available. Nowadays it is a commercial product, but some educational licenses are available.

Compared to Racer Pro, Pellet may have drawbacks, but ignoring again the problem of performance, Pellet is certainly one of the most feature-rich OWL reasoners. It is also supported by a strong development team and community, which is important if you are looking for different approaches and uses of the same tool, allowing the programmer to see what is happening in the internal code when something goes wrong or does not act as it is supposed to.

The language that the Jena implementation of DLModel used to communicate with any DL reasoner is DIG 1.1 (Bechhofer, Moller, and Crowther, 2003), but Pellet is used as the default reasoner, being integrated as a Java library in the API.

2.3. Existing Emotional Ontologies

Representing the emotional states of a user or simulating emotional states in a machine requires a suitable representation formalism. Although there are several markup languages containing elements of emotion annotation, few of them have considered the latest psychological theories about emotion or have been designed for generality of use in a broad range of applications. And, with respect to technology, it is surprising that there are few

[6] `http://protege.stanford.edu/`
[7] `http://www.racer-systems.com/`

applications that uses the most advanced results in Semantic Web technology (i.e. representation formats as RDF or OWL, reasoning tools as Protégé or Pellet, etc.) to approach the emotional phenomena.

WordNet Affect was developed by Strapparava and Valitutti (2004) semi-automatically. In WordNet Affect each word of WordNet (Miller, 1995) has an affective label. The labels include: semantic labels based on psychological and social theories, labels for the valence (positive or negative), for the activation (active or passive), etc. It can be said that WordNet Affect is an ontology that related the synsets of WordNet to affective labels. The affective labels presented in WordNet Affect do not conform a taxonomy, that is, in WordNet Affect emotions are individual isolated units not related to each other.

In (Obrenovic et al., 2005) the authors consider the great variety of theoretical models of emotions and implementation technologies which can be used in the design of emotion-based systems. They found researchers usually made practical choices of models and develop ad-hoc solutions. To improve this situations, the authors introduce a generic approach to modeling emotional cues, using an ontology of Emotional Cues. The concepts in this ontology are grouped into three global modules representing three layers of emotions' processing that enables a more flexible and implementation-independent description of emotional cues at different levels of abstraction. The ontology is developed using a proprietary XML-based formalism, and there is also a light comparison with the UML formalism. The authors seems not to have considered the benefits of using richer semantic representations. In this case we have an ontology that relates emotional clues to emotions, for example, a change of average tone (10% to 50%) is related in this ontology to *Joy in speech*. As in the previous ontology emotions are individual isolated units.

Mathieu (2005) introduces a semantic lexicon in the field of feelings and emotions, that further develops in a representation formalism in the form of an ontology. They use Protégé, in its classic frames-based modality, as a tool for annotating emotions in texts semi-automatically, as well as establishing links for navigating the text. Although the most popular tool of Semantic Web is used (i.e. Protégé) the approach is just a plugin for emotion-based annotation and does not take advantages of the semantic expressiveness of RDF or OWL. This ontology classifies 950 French words (600 are verbs and 350 are nouns) into 38 semantic classes according to their meanings (*Amusement*, *Emotion*, *Passion*, *Indignation*, etc.). This ontology only has one level of especification for each semantic class.

The Emotion Incubator Group [8] is one of those exceptions that have started to consider SemWeb technology for modeling emotions. They have identified a set of use cases for an Emotion Markup Language (Schröder et al., 2007), compiling a set of requirements that arose from those use cases in order to obtain a standard specification compatible with existing markup language. After reviewing different taxonomies of affective states (emotions, moods, interpersonal stances, preferences/attitudes, affect dispositions, etc.) five different components of emotions are identified. The first one is the subjective component, i.e. feelings, what is most strongly associated with the term emotion in folk psychology. The second one is the cognitive component, the role of conscious and unconscious mental processes that are concerned with the generation of emotion. The third one is the physiological component, i.e. neural connections and their role in the emotional phenomena. The fourth one is

[8] http://www.w3.org/2005/Incubator/emotion/charter-20071129

the behavioral component , how emotions associate to a small set of action tendencies in a human subject. Finally the fifth one is the expressive component, i.e. facial and body gestures, voice modulation, etc. Authors also review the previous annotation markup languages used to model emotions in text and accepted at recommendations or working drafts by the W3C (i.e. SMIL Synchronized Multimedia Integration Language , SSML Speech Synthesis Markup Language, EMMA Extensible MultiModal Annotation markup language, PLS Pronunciation Lexicon Specification, InkML Ink Markup Language, EARL Emotion Annotation and Representation Language and VHML Virtual Human Markup Language) and connect these languages with the five components of emotions mentioned before. The conclusions of this research are the same as many of us perceive in the field: many of the requirements of the emotion markup language cannot be found in any of the markups considered, specially the core part: emotion-specific elements. They mention that many formalisms only allow the use of emotional categories or emotional dimensions but not both. Attitudes, moods or affect dispositions are usually ignored. On the other hand timing, modality and metadata are other non emotion-specific elements that are very useful and usually well-covered in the current languages. In this ontology emotions are related to a feeling, a cognitive component, a psychological component, a behavioral component and a expressive component. Once again emotions in this ontology are not structured or related to each other.

To summarize, existing ontologies try to relate emotions to synsets, emotional clues, behavior, etc. without first structuring the emotions themselves. In this chapter we are going to present an ontology of emotional categories where emotions are structured and organized in levels which might provide a very useful step to future work relating this ontology of emotions to synsets, emotional clues or behavior.

2.4. Automatic Extraction and Mark Up of Emotions in Text

In this section first a brief outline of the most important approaches for emotional mark up of texts are presented. Then, we explain more in detail the system selected as domain for the use of our emotional ontology, EmoTag.

2.4.1. Existing Approaches for the Emotional Marked Up of Texts

There are various approaches that use emotional categories (Section 2.1.1.) to classify emotions such as the systems created by Zhe and Boucouvalas (2002), Liu, Lieberman, and Selker (2003), Alm and Sproat (2005), Sugimoto and Yoneyama (2006) or Strapparava and Mihalcea (2008).

Zhe and Boucouvalas (2002) have developed an emotions extraction engine which can analyse sentences given by the user. This system is included in the human-machine communication domain so it only takes into account the emotions referring to the speaker. The system analyses the sentences, detects the emotion (*happy*, *sad*, *fear*, *surprise*, *anger* or *disgust*) and shows the suitable facial expression. In order to obtain the emotion the words that compound the sentence are looked up in a dictionary of 16.400 words. In order to test the system a questionnaire was presented to 50 people, the same questions were answered by everyone and 450 sentences in total were returned.

Liu, Lieberman, and Selker (2003) created an approach based on large-scale real-world knowledge. This system marks up with the emotions defined by the OCC Model. Facts with emotional relevance are extracted from the OMCS (Open Mind Common Sense) Corpus. On this basis, a ‘common sense affect model’ is constructed. The model consists of a set of component models which compete with and complement each other. To construct the models, emotion keywords are propagated in three passes over the corpus. Emotion values initially are 1, then with each propagation are reduced by some factor *d*. To classify text, it is first segmented into clauses, then “linguistically processed”, and finally evaluated by 2-stage process using the models. To test their approach they incorporated their affect sensing engine into a email browser adding it the use of emotional faces. A study of 20 users was conducted in order to see what was preferred by the users: a browser with neutral faces, a browser with randomized emotional faces or a browser with the faces generated by Liu’s system.

Sugimoto and Yoneyama (2006) have a system that marks up text with emotions for Japanese in the narrative domain. The emotions used by this system are: *joy*, *sorrow*, *anger*, *surprise* and *neutral*. It decides the emotion of the sentence from the emotion of the names, adjectives and verbs which compose the sentence and from the grammatical structure of the sentence. Japanese generally has three sentence types: adjective sentence (S+V+Adjective Complement), noun sentence (S+V+Noun Complement) or verb sentence (S+V, S+V+O). The rules in order to determine the emotion of the sentence are different depending on the type of the sentence. In the adjective and noun sentences the emotion with highest weight is the emotion assigned to the adjective or to the noun. In the verbal sentences the emotion is determined by the combination of the emotions assigned to the subject and the verb.

Aman and Szpakowicz (2007) mark up blog posts with 6 basic emotions (*happiness*, *sadness*, *anger*, *disgust*, *surprise* and *fear*) and emotion intensity (*high*, *medium* and *low*). A corpus of 10,000 sentences blog posts was collected from the Web and it was marked up by 4 evaluators. The words in the blog posts are looking up in The General Inquirer and WordNet Affect in order to extract tags as *EMOT* (emotion), *Pos/Pstv* (positive), *happiness*, *fear*, etc. associated to each word. Then, Naive Bayes and Support Vector Machines (SVM) techniques are used to automatically mark up of blog posts.

Strapparava and Mihalcea (2008) mark up texts with 6 basic emotions (*anger*, *disgust*, *fear*, *joy*, *sadness* and *surprise*). This work uses Laten Semantic Analysis (LSA) and Naive Bayes classifiers to automatically mark up text with basic emotions using the corpus of SEMEVAL 2007 (a corpus of 250 headlines annotated by 6 evaluators) and a collection of blog posts annotated with moods that were mapped to the six emotions.

The systems presented mark up text with a reduced number of emotions. In this cases the different granularity of the labels is not an important aspect to take into account. So we have selected as domain of application of our emotional ontology a system that can mark up text not only with a reduced number of emotions but with labels of different granularity (from basic emotions to the more specific ones), EmoTag, which is explained in detail in the next section.

2.4.2. EmoTag

EmoTag (Francisco and Gervás, 2006) is a system that marks up texts with emotional categories which are selected from a group of 92 emotional labels, and the system considers each of these categories as individual concepts not related to any other category. This uncontrolled proliferation of labels multiplies the complexity involved in tasks such as deciding which is the best reference emotion for a sentence from the emotions selected by some human annotators, 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 to each other this might improve significantly the results obtained by EmoTag.

EmoTag relies on a dictionary of word-emotion assignments. This is obtained from a corpus of human evaluated texts by applying language analysis techniques such as stemming, POS tagging [9] or dependency analysis (Lin, 1998). Similar techniques are later applied to the assignment of emotions to sentences from the emotions associated to the words that compose the sentences.

Let us review the different stages required for the use of EmoTag, from the annotation of a textual corpus by human evaluators, to the automatic annotation of emotions in a raw text using the information extracted from the corpus.

The corpus used to train EmoTag consisted in 18 fairy tales with different lengths, written in English, making a total of 16.816 words and 1.389 sentences. The objective was to cover different styles by having tales from different authors and time periods. The sentence is the common unit of linguistic communication as it contains a complete thought-package. Therefore, it seemed reasonable to assign a different emotional content to each sentence, so the sentence was considered as emotional unit of the corpus. Each sentence of the tales had an emotion assigned to it.

Identification and assignment of emotions to a sentence is a subjective task, so each text from the corpus had to be annotated by several annotators in order to reduce possible annotator bias. The intention was to have the maximum number of stories annotated by a significant number of people. A reference value for each sentence has to be obtained based on the emotions assigned to the sentence by the annotators. In order to get this reference value, EmoTag selects the emotion most often assigned to the sentence, not taking into account the different granularity of emotional concepts or the relations between them. If the assigned emotions showed too much variability between annotators, the sentence was discarded and not included in the corpus.

An affective dictionary is required by EmoTag to know which emotions are transmitted by the words it is going to found in text. Using the tales marked up by human annotators, EmoTag obtains a database of words (List of Emotional Words or LEW) which includes the relation of each word in the texts with different emotional categories.

In order to obtain this LEW list the first step is to decide which are the relevant words in the texts. It was considered that the words that must be taken into account where those ones whose part-of-speech was somehow relevant. For example, it is very probable that nouns and verbs are influencing the emotion conveyed by a sentences, while prepositions or pronouns are not. For each of the relevant words EmoTag obtains its stem and associated

[9] http://www.english.bham.ac.uk/staff/omason/software/qtag.html

emotion in the annotation. 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, a process of normalization and expansion is performed.

EmoTag classifies sentences into emotions by taking into account the relations words-emotion in the affective dictionary. This process is carried out in the following way:

- Sentence detection and tokenization are performed. With this information EmoTag obtains the words affected by negations using dependency analysis, the stem, and the part-of-speech of each of the words.
- The emotional value associated to each word is obtained by looking in the affective dictionary (LEW list).
- The words under the scope of the negations are correctly processed.
- EmoTag obtains the emotional 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 probabilities of 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 are considered individually.

A sample part of a tale marked by EmoTag 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>`

3. OntoEmotion. The Emotional Ontology

For supporting the semantic necessity of emotional systems, we have developed an ontology of emotional categories called OntoEmotion as a useful resource for the management of emotional content. By using this ontology we can identify relations between different levels of specification of related emotions when the emotional content is represented as emotional categories. We took emotional categories (i.e. emotion-denoting words such as *happiness*, *sadness* and *fear*) as first class citizens of our ontology. In particular, we used the following basic emotions: *sadness*, *happiness*, *surprise*, *fear* and *anger*. The ontology was developed by trying to integrate all the cluster structures described in Section 2.1.2., while taking into account the emotional structures found in existing literature. The ontology is written in OWL (Bechhofer et al., 2004) and uses Pellet as reasoner. Emotional categories are structured in a taxonomy that covers from basic emotions to the most specific emotional categories. In addition, the ontology relates each of the emotional categories to the three emotional dimensions by means of data ranges. In this way, this ontology is also appropriate when it is required to change between different representations of emotions.

Concepts in the ontology represent language-independent emotions corresponding to common experiences in life. The hypothesis is that we all have the same abstract conception of *Happiness*, for instance, while different words can be used to refer to it. The instances in the ontology represent the words provided by specific languages (e.g. English) for referring to emotions. So, a concept can have multiple instances as a language can give us multiple words to refer to the same emotion.

3.1. Structure

Our ontology has two root concepts:

- **Emotion**: Each emotional concept is a subclass of this root. Emotions are structured in a taxonomy, with the number of levels under each basic emotion depending on the level of available specification for it. For example, *Sadness* has two sublevels of specification. The second of this levels indicates different types of *Sadness*: *Despair*, *Disappointment*, *Grief* or *Shame*. Some of these emotions are specialized again in the third level. For example, *Shame* is divided into *Regret* and *Guilt*. On the other hand, *Surprise* only has one sublevel with two emotional concepts: *Amazement* and *Intrigue*.

- **Word**: This is the root for the emotion-denoting words, the specific words that each language provides for denoting emotions. Our 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*.

As instances of the *EnglishWord* and *SpanishWord* subclasses we can find emotion-denoting words, which are all the words used for denoting *Anger*, *Annoyance*, *Displeasure*, *Terror*, etc. Each of these instances has two parents: a concept from the **Emotion** hierarchy (which indicates the type of abstract emotion denoted by the word) and a concept from the **Word** hierarchy (which indicates the language of the word).

Figure 1 shows a fragment of the ontology. It shows emotional concepts as *Anger*, *Fear* and *Surprise*. Under those emotional concepts there are emotional words such as *annoyance*, *unhappiness* and *terror*. In this fragment it can be seen how the emotional categories are related both to one emotional concept and to one word concept. For example, the word *unhappiness* is an instance of the emotional concept *Sadness* at the same time it is an instance of the word concept *EnglishWord*, which means that *unhappiness* is a possible English word for denoting the sadness emotion.

Another valid way of representing these relations might be to create a new property called “language” to connect each word to an instance of the language it belongs. We have chosen the in-built “type” relation because individuals with many different types are considered natural in OWL DL, and it is easier to retrieve every word of a specific type than “every word that has a relation with a specific individual”.

According to the semantics we chose for our ontology, all the instances of the same emotional concept are synonyms. For example, the words *annoyance* and *irritation* are

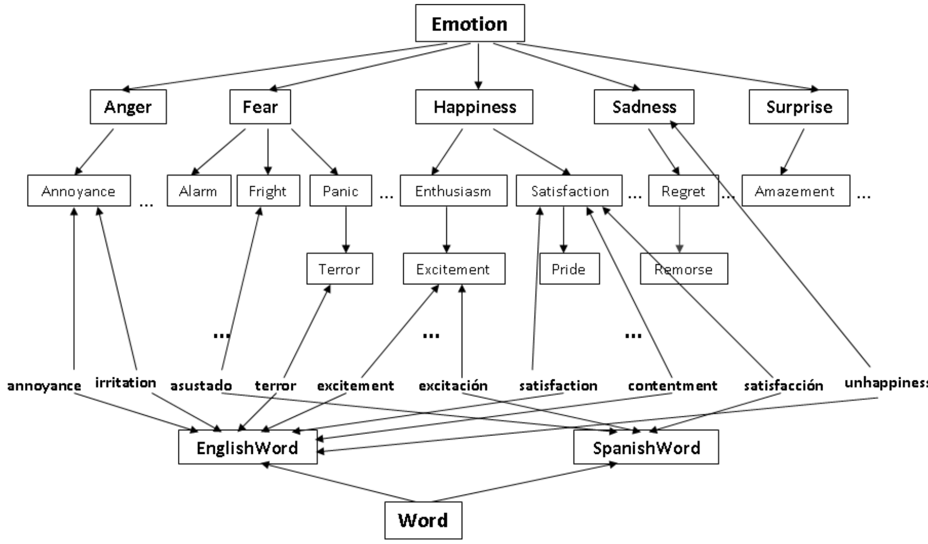

Figure 1. Partial view of the emotional ontology

considered synonyms because both are instances of the emotional concept *Annoyance*. This is specially relevant considering the importance of this type of semantic relations in text processing, compared to the classic bag of words approach (Wang et al., 2008).

To summarize, we can conclude that our emotional ontology represents the emotional categories as instances of a tree structure of emotional concepts. Each emotional word is an instance of two concepts: an emotional concept which represents the emotion denoted 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 can obtain the emotion concept directly associated with an emotional word in the ontology, i.e. its parent, as well as other more general emotion concepts related to that word, according to the conceptual hierarchy. Finally we can also obtain the synonyms of an emotion word and getting the siblings of a particular instance. For example, given the emotional word *excitement*, we have as direct emotional concept *Excitement*, as general emotional concepts *Enthusiasm*, *Happiness* and *Emotion* and as a Spanish synonym *excitación*.

3.2. Emotional Categories and Dimensions Equivalence by Means of Datatype Properties

Once we have a hierarchy of emotions and relations between the emotion-denoting words, their language and the concept they represent, we want to link the emotional concepts in the ontology with the three main emotional dimensions. Numeric data can be represented in OWL using datatype properties. To achieve this we have declared three datatype properties:

- *hasEvaluation*: Represents the data range for the evaluation dimension.
- *hasActivation*: Represents the data range for the activation dimension.

Emotion	hasActivation	hasEvaluation	hasPower
Anger	7<=x<=10	0<=y<=3	3<=z<=5
Displeasure	7<=x<7.5	0<=y<=3	3<=z<=5
Annoyance	7.5<=x<8	0<=y<=3	3<=z<=5
Sulking	8<=x<8.5	0<=y<=3	3<=z<=5
Indignation	8.5<=x<9	0<=y<=3	3<=z<=5
Fury	9<=x<=10	0<=y<=3	3<=z<=5

Table 1. Datatype properties assigns to the different forms of anger presented in our ontology.

- *hasPower*: Represents the data range for the power dimension.

Each of the emotional concepts is defined by specifying appropriate data ranges for these properties. Each emotional concept in the ontology takes up a region in the three-dimensional space of emotional dimensions. In order to describe this with the datatype properties we have to define our own datatype restrictions, because we are using specific intervals between numbers of type float. This can be done using data range definitions. For example, for the *Anger* emotional concept, the region of the space associated with it can be described in the following way: 7<=hasActivation<=10, 0<=hasEvaluation<=3, 3<=hasPower<=5.

To establish the dimensional intervals for each emotional categories we have employed the existing literature that explains the diferences between specifications of a same abstract emotion in arousal, evaluation or activation. For example, (Ekman, 2004) explains that *displeasure*, *annoyance*, *sulking*, *indignation* and *fury* are different forms of the basic emotion anger that differ from one another in their intensity of arousal, *displeasure* is the form of anger with the lowest intensity of arousal whereas *fury* is the one with the highest intensity of arousal. Taking into account that the intensity of arousal in the basic emotion *anger* ranges from 7 to 10 we have defined the ranges in Table 1 fo the three datatype properties.

The fragment of the OWL file which correspond to the data range for the *hasActivation* property is shown in Table 2.

In this way, by means of the data ranges on the datatype properties, the link between the abstract emotional concepts and the three-dimensional space of emotional dimensions is established.

3.3. Automatic Classification of Emotional Dimensions as Emotional Categories

A requirement to be taken into account when representing emotions using numerical data is to have some reasoning device capable of processing such data in an appropriate way. Pellet is able to classify concepts with restrictions formed by combinations of user-defined datatypes.

Once we have defined the emotional concepts by means of the emotional dimensions, Pellet automatically classifies the concepts into a hierarchy of emotional concepts. This

```
<owl:Restriction>
 <owl:allValuesFrom>
  <owl:DataRange>
   <owl:onDataRange rdf:resource=‘‘http://www.w3.org/2001/XMLSchema#float’’/>
   <owl:minInclusive rdf:datatype=‘‘http://www.w3.org/2001/XMLSchema#float’’>
   7.0</owl:minInclusive>
  </owl:DataRange
 </owl:allValuesFrom>
 <owl:onProperty>
  <owl:FunctionalProperty rdf:about=‘‘#hasActivation’’/>
 </owl:onProperty>
</owl:Restriction>
<owl:Restriction>
 <owl:onProperty>
  <owl:FunctionalProperty  rdf:about=‘‘#hasActivation’’/>
 </owl:onProperty>
 <owl:allValuesFrom>
  <owl:DataRange>
   <owl:onDataRange rdf:resource=‘‘http://www.w3.org/2001/XMLSchema#float’’/>
   <owl:maxInclusive rdf:datatype=‘‘http://www.w3.org/2001/XMLSchema#float’’>
   10.0</owl:maxInclusive>
  </owl:DataRange>
 </owl:allValuesFrom>
</owl:Restriction>
```

Table 2. Fragment of the OWL Ontology file which corresponds to the data range for the *hasActivation* property.

means that Pellet obtains a hierarchy of emotions in which the most basic concepts are at the top of the hierarchy and the concepts which are more specific appear as descendants of the more general ones.

Datatype properties transform the classification of the emotional concepts into a relatively simple task. It is not necessary for the designer of the ontology to know which concepts are more specific than others because it is the reasoner that carries out the task automatically. For example, if we define the ranges in Table 1 for the three datatype properties, just by loading the ontology in DLModel, the reasoner automatically classifies the concepts *Displeasure*, *Annoyance*, *Sulking*, *Indignation* and *Fury* as subclasses of the emotional concept *Anger* which is automatically identified as more general. In Figure 2 we can see how the reasoner classifies the concepts in the correct way.

4. Using the Emotional Ontology for the Automatic Treatment of Emotions in Texts

The use of an ontology like OntoEmotion could improve the performance of systems that deal with the task of automated mark up of texts with emotional information. This kind of ontology could be used to improve the treatment of emotions at three different levels within the automated marked up of emotions in texts. As a practical example, we are going to study how the ontology improves the performance of EmoTag at these three different levels:

1. For obtaining the reference value for each sentence from the emotions selected by the

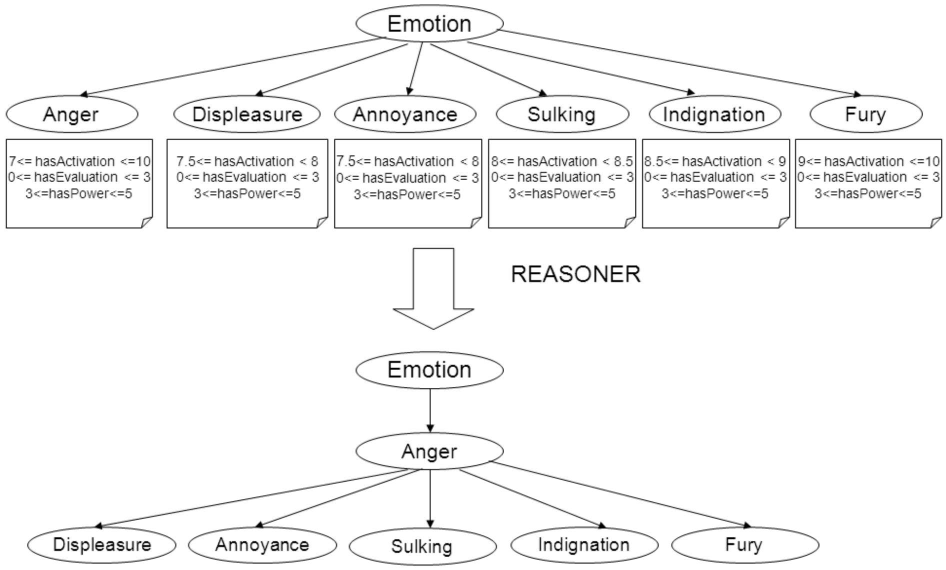

Figure 2. Example of the automatic classification of emotions using datatype properties

annotators. In the original version of EmoTag the reference value (value most often assigned to a sentence by the annotators) did not take into account the emotions as related units. If we look not only for the choice of a specific value but for the choice of a similar value the reference value could be more accurate.

2. For selecting the appropriate emotion for a sentence based on the emotions associated to each word and the probability of that association. In order to select the emotion to be assigned to a sentence, the original EmoTag considered the emotions as independent units again, so the emotion with a higher resulting probability was the emotion assigned to the sentence. This was a good first solution but it can be improved by distinguishing the different granularity of emotional concepts and the relation between emotions derived from the ontology. 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.

3. For evaluating the resulting mark up. In order to determine how well a text had been marked up, EmoTag considered that each sentence was successfully tagged if the emotion assigned by the tagger matched exactly the reference value considered for the sentence. In this way, a sentence only could be correctly marked or incorrectly marked. However, 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 annotators, it would be wrong to consider this sentence as completely incorrectly marked.

We will now see how the ontology was used to enhance the performing of EmoTag for each of the problems explained above, and how the results obtained have improved by means of the emotional ontology. In order to evaluate the effect of the use OntoEmotion in EmoTag, we have carried out different tests for studying the improvements obtained.

4.1. Using the Ontology to Identify the Reference Value for Each Sentence

In the corpus used by EmoTag, a *reference value* for the emotion of each sentence was obtained by choosing the emotion most often assigned to that sentence by human annotators. This was a good first solution but it can be improved. Because of the different granularity of emotional concepts, in the case of an assignment of emotions by annotators to a sentence as shown in Table 3, the emotion *sadness* would be taken as the reference emotion as it is the most repeated one. However, *sadness* is not the best reference value in this case. If we look in 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 would be *grief* instead of *sadness*.

Annot. 1	**Annot. 2**	**Annot. 3**	**Annot. 4**	**Annot. 5**	**Annot. 6**
agony	anguish	grief	sorrow	sadness	sadness

Table 3. Example of the assigment of emotions by annotators to a sentence.

In this ontology-enhanced 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 annotators. The detailed process undertaken to accomplish this task is the following:

1. If at least half of the annotators are in agreement on the assignment of an emotion to the sentence, we take this emotion as reference value for the sentence.

2. Otherwise, we group emotions by levels of emotional concepts from the ontology. We obtain all the ancestors of the concepts in the lowest level. If any emotion is then supported by at least half of the annotators, it is taken as reference value. If there are two emotions that are supported by at least half of the annotators, we take the emotion with the lowest level.

3. We repeat the previous step for each level in ascending order until an emotion supported by most annotators is found.

4. Finally, if there is no emotion supported by at least half of the annotators, the sentence is left out from the corpus.

This process is exemplified in Figure 3, which shows an example of how to obtain the reference value for a sentence annotated by 6 annotators. In the first table we present the assignments made to the sentence. The first step is to group the emotions. We obtain the level of each emotion by means of the emotional ontology. The result is shown in the second table. From the second table it can be seen that no emotion is supported by at least half of the annotators, so we get the concepts related to the emotions with the lowest level (*Grief*, *Helplessness* and *Remorse*). We insert all their related concepts in the table (in this case *Distress*, *Powerless*, *Regret* and *Sadness* three times). The result can be seen in the third table. Based on these results we see that *Sadness* is supported by four annotators, i.e. more than half, so this is the emotion taken as reference value for this sentence.

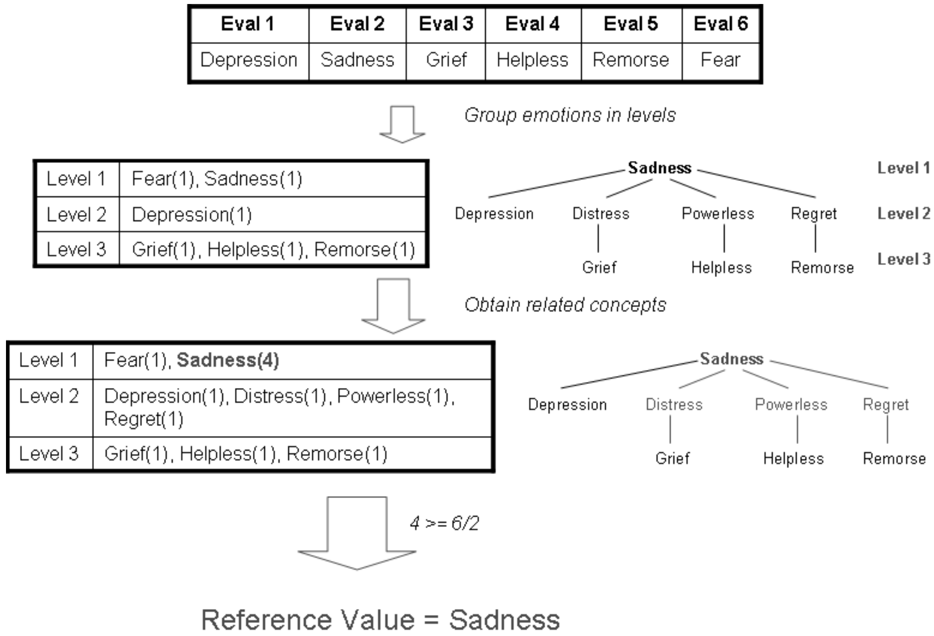

Figure 3. Example of the reference value obtained from the assigment of emotions by the annotators.

4.1.1. Comparison of Results

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

We have also empirically evaluated the category reduction performed when the annotators did not agree on the emotion assigned to a sentence. In order to evaluate our reference values we asked annotators whether the re-annotation of the tales that were assigned to them was acceptable. This evaluation was performed by the eight annotators who annotated the higher number of tales in the original experiment. We presented them pairs <*previous annotation*, *reference value*> along with some sentences that they annotated with the *previous annotation* and which had then been re-annotated by the method above with the *reference value*. For each pair they had to specify if they agreed with the new annotation or not. Each pair was tested by at least four annotators. The annotators agreed with the new reference value given in 90% of the pairs presented. Therefore it can be considered that the method used in order to get the reference value when the annotators did not agree was appropriate. Taking into account this evaluation it can be considered too that the ontology used is a valid tool for structuring and relating emotions. The pairs that were identified by the annotators as non equivalent were <*Admiration*, *Happiness*>, <*Excitement*, *Happiness*>, <*Gratification*, *Happiness*>, <*Hope*, *Happiness*>, <*Powerlessness*, *Fear*>, <*Suffering*, *Sadness*> and <*Torment*, *Sadness*>.

4.2. Ontology-Supported Automated Mark Up of Emotions

As mentioned in Section 2.4.2., in order to mark up a sentence with an emotion EmoTag splits it into words and assigns to each of them the probability of having the different emotions. Based on these probabilities for the words, EmoTag obtains the final emotion of the sentence as follows: once all the words of the sentence have been evaluated, it adds up the probability of each emotion of the different words and assigns to the sentence the emotion with 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 is considered individually. In this case it will be better to consider their probabilities together under the more general concept. As an example, suppose we have the emotions and their corresponding probabilities shown in Table 4.

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

Table 4. Example of emotions and their corresponding probabilities

In this example EmoTag would consider that the emotion with the highest probability is *happiness* (0.3=0.1+0.2), but using 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 highest probability is *anger* (0.45=0.2+0.1+0.05+0.1) which seems to be a better result.

The process followed in the ontology-enhanced version of EmoTag 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:

- It processes 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 (lower level in the ontology) with the higher probability is assigned to the sentence.

Figure 4 shows an example of this process. In this example we have a sentence with three emotional words, each of them with 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 summed up probabilities. Then, 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 their probability being the sum of the ones of their children:

- *Anger* = 30% (previous probability) + 20% (*Indignation* probability) + 15% (*Sulking* probability) + 30% (*Displeasure* probability) = 95%

- *Surprise* = 65% (previous probability) + 5% (*Amazement* probability) = 70%).

Anger, the more general emotion with the higher probability, is assigned to the sentence.

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%

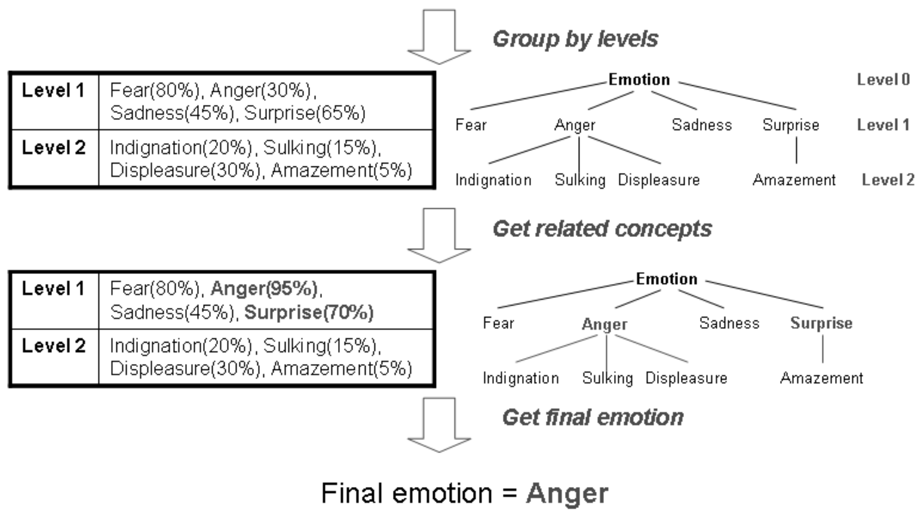

Figure 4. Example of the assignment of the final emotion to a sentence based on the emotions of the words which compond it.

4.2.1. Comparison of Results

In order to evaluate the effect of the use of the ontology in EmoTag we have carried out similar tests as the ones done for the previous version of EmoTag. In this case we have used four tales. Each of these four texts have been marked up with using emotional categories by different annotators. Then we have obtained the mark up by EmoTag. As a result, we have two different measures: the emotions assigned by the human annotators (or annotator's results), and emotions assigned by EmoTag (tagger's results). These two sets of results are explained below:

- Annotator's results: A reference emotion for each sentence is obtained by choosing the emotion most often assigned to that sentence by the human annotators, and using the emotional ontology as explained in Section 4.1..

- Tagger's results: The reference value obtained in the annotator's results is used to be compared with the results generated by the ontology-enhanced EmoTag. 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 annotators did not agree.

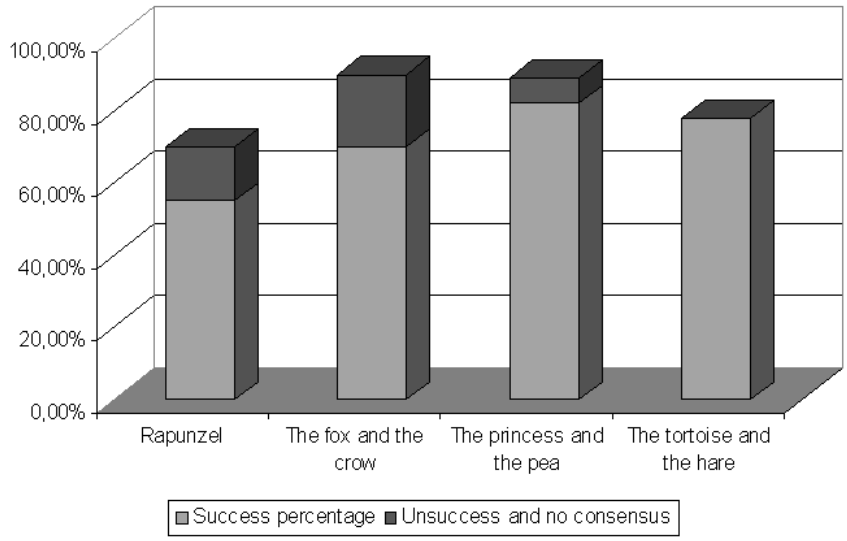

Figure 5. Percentage of success in automated tagging with ontology-enhanced version of EmoTag.

In addition, Figure 6 compares the results obtained for the ontology-enhanced version with the ones obtained for the previous version. We can see that the emotional ontology increases the percentage of sentences correctly marked around a 15%.

4.3. Ontology-Supported Evaluation of the Mark Up

In order to determine how well a text had been marked up, EmoTag considered that each sentence was successfully tagged if the emotion assigned by the tagger matched the reference value obtained from the annotators. In this way, a sentence can only be correctly marked (1) or incorrectly marked (0). However, by using the emotional ontology we can determine how well a sentence is marked up with a finer granularity.

In the ontology-enhanced 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. This score is based on the number of levels in common between the two emotions being compared and the level of the most specific emotion (see Equation 1).

$$Correctness = number\_common\_levels/level\_specific\_emotion \qquad (1)$$

For example, in the case of a sentence marked by the evaluators as *excitement* and by EmoTag as *enthusiasm* (see Figure 7), we have that the emotional category *excitement* is related to the emotional concept *Excitement* in the emotional ontology, *Excitement* (level 3) is related to *Enthusiasm* (level 2) and this is related to *Happiness* (level 1). On the other side, we have *enthusiasm* which is related to the emotional concept *Enthusiasm* in the ontology, and *Enthusiasm* (level 2) is related to *Happiness* (level 1). To summarize, *enthusiasm* has level 2, *excitement* has level 3 and *enthusiasm* and *excitement* has 2 emotional concepts in common: *Enthusiasm* and *Happiness*, so the correctness is 2/3 = 0.67.

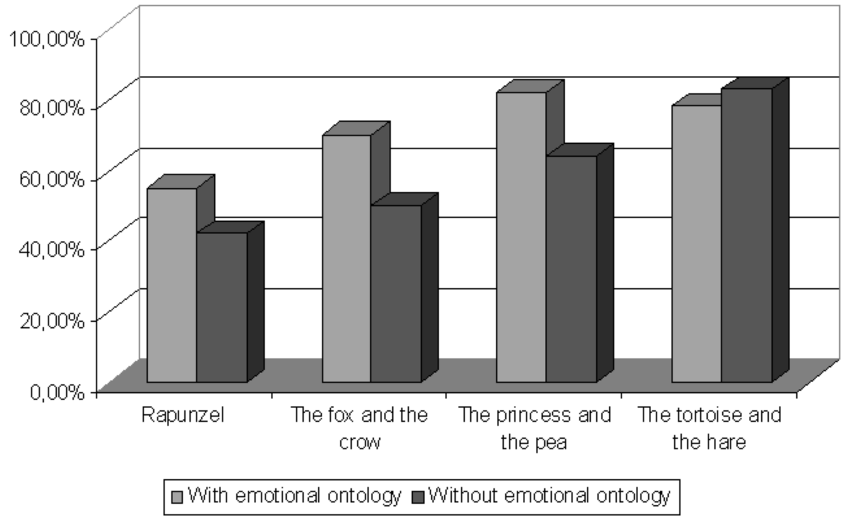

Figure 6. Relative improvement in mark up success of the ontology-enhanced version with respect to the previous one.

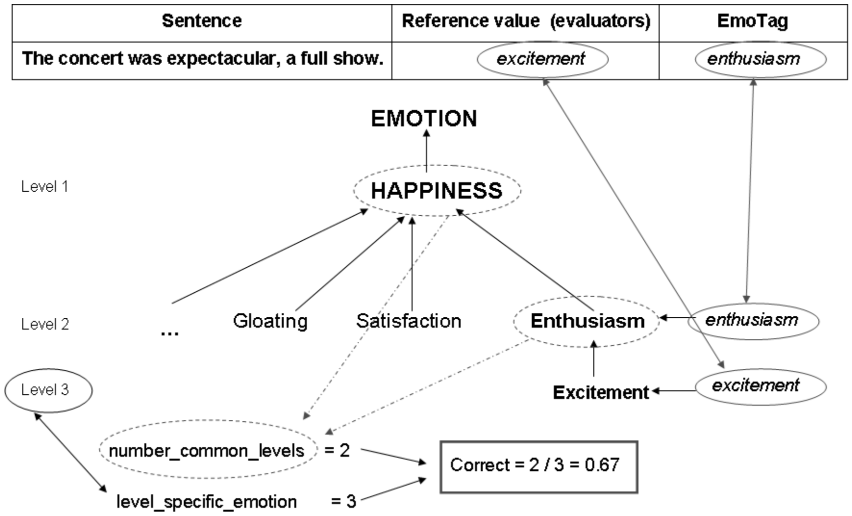

Figure 7. Example of correctness of the emotion assignment to a sentence.

4.3.1. Comparison of Results

Only the sentences with a correct level of 1 has been taken into account in the evaluation as correct sentences, because the percentage of sentences with a lower value (a value higher than 0 and lower than 1, that is sentences mark up with an emotion not equal to the reference value but similar) is so low that it can be disregarded. We can conclude that the use of a emotional ontology for the comparision of the tagger value with the reference value do not improved significantly the results.

Emotion	Volume	Rate	Pitch Baseline	Pitch Range
Anger	+10%	+21%	+0%	+173%
Surprise	+10%	+0%	+25%	+82%
Happiness	+10%	+29%	+35%	+27%
Sadness	-10%	-8%	-10%	-36%
Fear	+10%	+12,5%	+75%	+118%

Table 5. Configuration Parameters for Emotional Voice Synthesis

5. Using the Emotional Ontology to Interpret Emotional Dimensions as Emotional Categories

To address the problem of translating dimensions into categories, in our emotional ontology each of the categories is related to three different dimensions (i.e. evaluation, power and control) by means of data ranges. By applying automated classification in terms of the emotional ontology, the system can interpret unrestricted inputs in terms of a restricted set of concepts for which particular rules are provided. An example of application for this use of the ontology is presented in which the rules applied for document markup with emotional dimensions provide configuration parameters for an emotional voice synthesizer.

EmoSpeech (Francisco, Hervás, and Gervás, 2005) is a system capable of modulating the voice quality of a synthesizer while reading aloud children's tales, so that the voice conveys at least part of the emotions expressed by the corresponding text. This is achieved by controlling those parameters in the synthesizer that have been identified as having more relevance in the expression of emotions in human voice. EmoSpeech operates with five basic emotions: *anger*, *happiness*, *sadness*, *fear* and *surprise*, which correspond with the basic emotions of the emotional ontology. The aspects of the voice that act as personality identifiers are: volume, rate, pitch baseline and pitch range. EmoSpeech uses a group of rules which relate the five basic emotions to the specific changes on voice parameters involved in the communication of emotion in human voice utterances. The values of these parameters for every emotion were obtained by refining an original proposal by Schröder (Schröder, 2004), based on the analysis of emotional material generated by actors. The optimal values were obtained through the systematic variation of the parameters during synthesis. Table 5 summarizes the rules of the synthesizer for the basic emotions. Using the particular configuration of parameters for that particular basic emotion, the synthesizer reads out aloud the text with the emotion assigned to the sentences.

EmoSpeech needs as input a text mark up with basic emotional categories. If EmoSpeech receives as input a text mark up with emotional dimensions it must face the problem of translating emotional dimensions into basic emotional categories. To deal with this problem EmoSpeech needs a tool that relates emotional categories to emotional dimensions. In OntoEmotion each of the categories is related to 3 different dimensions (i.e. evaluation, activation and power) by means of data ranges. OntoEmotion allows transformation of a text marked up with emotional dimensions into a text marked up with basic emotional categories by means of the reasoner Pellet that classifies a point in the three-dimensional space of emotional dimensions into a concept of the ontology (emotional category). Once we have the related emotional category we can easy obtain the basic emotional concept related to it by means of DLModel.

In Figure 8 we can see how this process works for a specific example. In this example, we have a sentence of an input text marked up with the following values for each emotional dimension: *evaluation* = 1, *activation* = 7 and *power* = 5. These values represent a point in the dimensional space which is classified by means of the emotional ontology under the *annoyance* emotional concept. We ask DLModel for the parents of *annoyance* and the *anger* emotional concept is returned. EmoSpeech then receives the sentence of the input text and the emotion *anger* as the one associated to the sentence, so it selects the rules corresponding to this basic emotion. Once EmoSpeech has the suitable rules for the emotional meaning of the sentence, the synthesizer reads aloud the sentence in an angry way.

The option of using a description logic ontology - and the associated abilities to carry out instance recognition and automatic classification - as an interface to achieve this conversion as proposed in this chapter, presents two distinct advantages:

- It provides a method for the automatic association of any point in the three dimensional space to whatever is the closest available configuration of the speech synthesizer, based on information that is defined at the conceptual level - even if it relies on an underlying level of geometrical representation.

- Any subsequent refinement of the set of configurations available for the synthesizer - for instance, if the existing configurations are refined into a larger set of options by fine tuning them to better represent more specific emotions -, it would be enough to associate the new configurations to the corresponding concepts, and to refine the search algorithm to stop at the first ancestor that has some configuration data associated with it.

6. Conclusion

The use of an ontology of emotions is a good solution to improve the results obtained by a system that marks up text with emotional categories automatically. We have apply the ontology to a system that marks up texts with emotions (EmoTag) in three different tasks: obtaining the reference value of each sentence from the emotions selected by the human annotators, marking up of sentences with emotions and determining if a sentence has been correctly marked up. After analysis of the results obtained with the ontology-enhanced version of EmoTag we can conclude that the ontology improves the percentage of reference values supported by the majority of evaluators by 25% and the percentage of sentences correctly marked by 15%. The only change in the ontology-enhanced version which has not proven results is the application of the ontology to the determination of the correctness of a sentence, because the percentage of sentences with a value between 0 and 1 is so low that can be disregarded.

The emotional ontology provides a solution to the problem of translating emotional dimensions into emotional categories. OntoEmotion has been applied for connecting text marked up in terms of emotional dimensions to a synthesizer that is in principle only capable of processing material marked up in terms of basic emotional categories. The effect

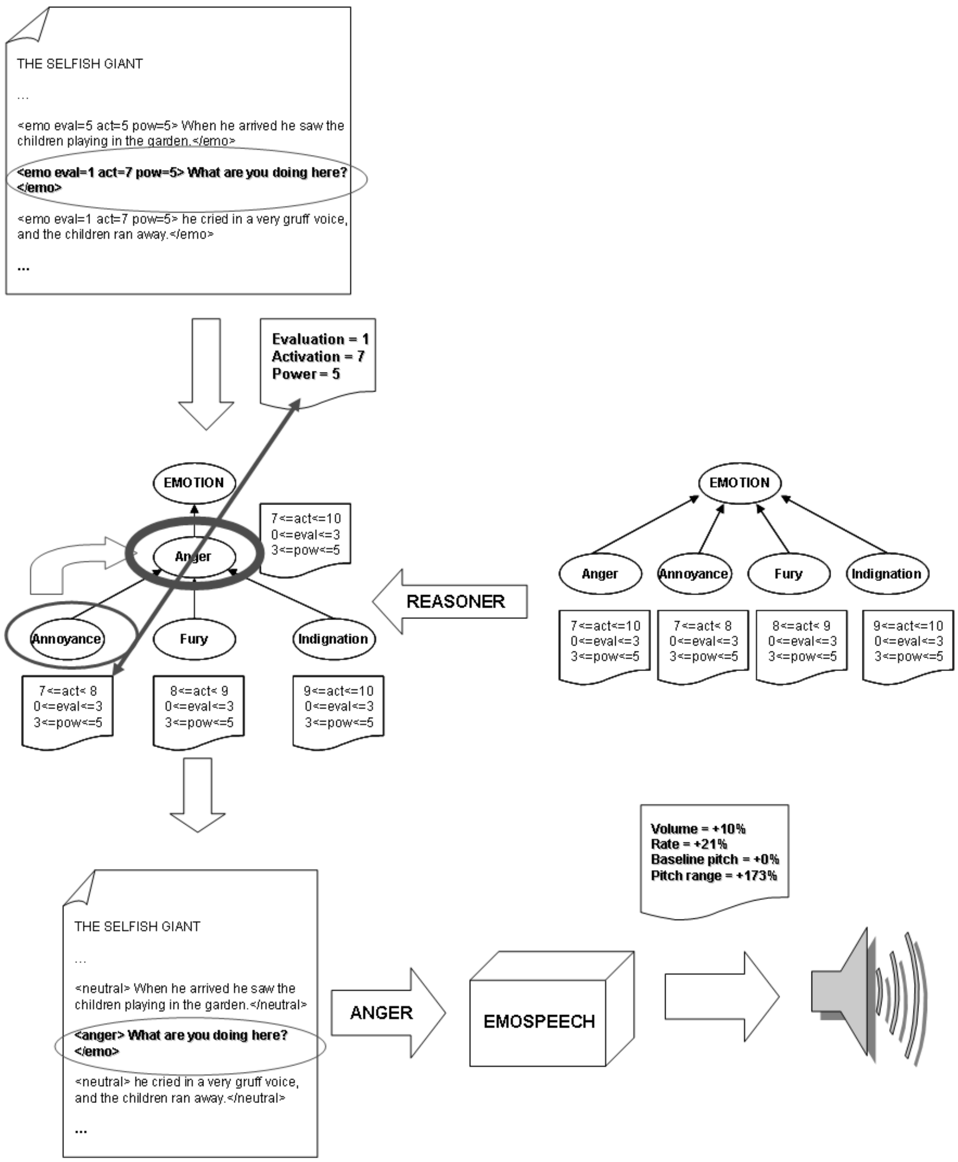

Figure 8. Example of the configuration a voice synthesizer based on a emotional mark up with emotional dimensions

of the emotional ontology on the quality of speech output is limited[10], because the emotion ontology is being used only as interface between the emotional mark up application and the voice synthesizer. For inputs originally tagged with emotional categories, the addition of the ontology has little impact. Nevertheless, emotional categories as a method of representing emotions provide only very limited granularity. On the other hand, emotional dimensions provide much more flexible means of representing emotions, with greater expressive power. The main obstacle in switching from one representation to another lies in the fact that there

[10]The quality and emotional precision of the resulting voice has been discussed elsewhere. Details can be found in (Francisco, Hervás, and Gervás, 2005).

is no easy way of converting from emotional dimensions to voice synthesizer configurations. At best, the three dimensional space of emotional dimensions could be partitioned into restricted volumes of space, and a particular configuration of the synthesizer assigned to each volume.

The next step is to relate the concepts that appear in a text with emotions. We have an ontology of emotions and we are developing an ontology of concepts which represents the concepts which take part in a tale. In order to get the semantic emotion of a tale we have to link these two ontologies. This way, we will obtain the emotions related to the concepts that take part in the tale. An example of the relation between concepts (characters and actions) and emotions can be seen in (Francisco, Hervás, and Gervás, 2006). In order to identify the links between a given sentence and the concepts in the domain we can apply text analysis techniques. We have already used techniques such as dependency analysis to identify key concepts of a sentence in order to build conceptual cases from a text corpus (Francisco, Hervás, and Gervás, 2007). This is important, because depending on the semantic content the final emotion could differ from the lexical emotion obtained at the moment by EmoTag. For example, the action "to die" is a sadness action, and the emotional dictionary used by EmoTag is marked up this way, but if the subject of this action is "the witch" the action "to die" turn into a happy action.

Regarding the technologies that have been presented in this chapter, some of these are not generally accepted as standard. Datatypes (and "reasoning" with numbers and strings) are not part of the essence of Description Logics. OWL DL considers datatypes properties disjoint with every object property. In the sublanguage profiles of OWL 2 (Motik et al., 2009) support for datatypes has been improved, because it has been showed useful for many applications. The version of OWL that we have used for this work only supports some standard XML Schema datatypes and it lacks a standard solution for representing user-defined datatypes. DIG 1.1, being a standard designed for the communication with DL reasoners, does not accept restrictions over datatype properties. This obstacle made it impossible for us to send an ontology that includes such restrictions directly from Protégé to Pellet for its automatic classification. DIG 2.0, with support for OWL 2 offers those features, but for now other shortcuts must be used in order to reason with restrictions on datatype properties. Old versions of Protégé had a proprietary solution to represent user-defined datatypes, which allowed the creation of restrictions with interesting datatype properties and even visualization of the limits of a numeric interval in the GUI. However, DIG 1.1 does not allow that kind of information to travel to a DL reasoner. Pellet, by itself, can deal with user-defined datatype restrictions, and now the last version supports the in-line syntax of OWL 2. So because we are using Protégé as the editor for our ontology and knowledge base, we have to edit the files manually to add those restrictions before loading everything in DLModel using a Pellet-Java default configuration. We hope that some of these shortcomings might be solved when updating to the later versions of these technologies. Although reasoning support for datatype properties in OWL DL is still not standard, there are available technologies that let us experiment with these features and allow us to develop Affective Computing applications like the ones described in this chapter. OWL, Jena, DLModel, Protégé and Pellet are the common choices for developing new iterations of the software. Still more improvements are needed in editors such as Protégé to take advantage of all the possibilities of reasoners such as Pellet and its latest achievements.

References

Alm, C. Ovesdotter and R. Sproat. 2005. Emotional sequencing and development in fairy tales. In *Proc. of Affective Computing and Intelligent Interaction (ACII'05)*, volume 3784 of *Lecture Notes in Computer Science*, pages 668–674. Springer.

Aman, S. and S. Szpakowicz. 2007. Identifying expressions of emotion in text. In *Proc. of Internation Conference on Text, Speech and Dialogue (TSD'07)*, volume 4629 of *Lecture Notes in Computer Science*, pages 196–205. Springer.

Bechhofer, S., F. Van Harmelen, J. Hendler, I.Horrocks, D. McGuinness, P. Patel-Schneider, and A. Stein. 2004. *OWL Web Ontology Language Reference*, volume 12/2. W3C.

Bechhofer, Sean, Ralf Moller, and Peter Crowther. 2003. The DIG description logic interface. In *Proc. of International Workshop on Description Logics (DL'03)*.

Cowie, R. and R.R. Cornelius. 2003. Describing the emotional states that are expressed in speech. *Speech Communication*, 40(1-2):5–32.

Cowie, R., E. Douglas-Cowie, and A. Romano. 1999. Changing emotional tone in dialogue and its prosodic correlates. In *Proc. of International Workshop on Dialogue and Prosody (ESCA'99)*, pages 41–46.

Ekman, P. 2004. *Emotions Revealed : Recognizing Faces and Feelings to Improve Communication and Emotional Life*. Owl Books.

Fontaine, J.R., K.R. Scherer, E.B. Roesch, and P. Ellsworth. 2007. The world of emotion is not two-dimensional. *Psychological Science*, 18:1050–1057.

Francisco, V. and P. Gervás. 2006. Automated Mark Up of Affective Information in English Texts. In *International Conference on Text, Speech and Dialogue (TSD'06)*, Lecture Notes in Computer Science, pages 375–382. Springer.

Francisco, V., R. Hervás, and P. Gervás. 2005. Análisis y síntesis de expresión emocional en cuentos leídos en voz alta. *Sociedad Española para el Procesamiento del Lenguaje Natural, Procesamiento de Lenguaje Natural*, 35:293–300.

Francisco, V., R. Hervás, and P. Gervás. 2006. Two different approaches to automated mark up of emotions in text. In *Proc. of International Conference on Innovative Techniques and Applications of Artificial Intelligence (AI'06)*, pages 101–114.

Francisco, V., R. Hervás, and P. Gervás. 2007. Dependency analysis and cbr to bridge the generation gap in template-based nlg. In *Proc. of International Conference on Intelligent Text Processing and Computational Linguistics (CICLing'07)*, Lecture Notes in Computer Science, pages 432–443. Springer.

Jing, Y., D. Jeong, and D-K. Baik. 2008. SPARQL graph pattern rewriting for OWL-DL inference queries. *Knowledge and Information Systems*, Online First.

Lin, D. 1998. Dependency-based Evaluation of MINIPAR. In *Proc. of Workshop on the Evaluation of Parsing Systems at LREC'98*.

Liu, H., H. Lieberman, and T. Selker. 2003. A model of textual affect sensing using real-world knowledge. In *Proc. of International Conference on Intelligent User Interfaces IUI'03*, pages 125–132.

Mathieu, Y. Y. 2005. Annotation of emotions and feelings in texts. In *Affective Computing and Intelligent Interaction (ACII'05)*, volume 3784 of *Lecture Notes in Computer Science*, pages 350–357. Springer.

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

Mindswap. 2006. Pellet owl reasoner. `http://pellet.owldl.com/`.

Motik, B., B. Cuenca-Grau, I. Horrocks, Z. Wu, A. Fokoue, and C. Lutz. 2009. OWL 2 Web Ontology Language Profiles. `http://www.w3.org/TR/owl2-profiles/`.

Motik, Boris, Ulrike Sattler, and Rudi Studer. 2005. Query Answering for OWL-DL with Rules. *Rules systems. Web Semantics: Science, Services and Agents on the World Wide Web*, 3(1):41–60.

Obrenovic, Z., N. Garay, J. M. López, I. Fajardo, and I. Cearreta. 2005. An ontology for description of emotional cues. In *Proc. of Affective Computing and Intelligent Interaction (ACII'05)*, volume 3784 of *Lecture Notes in Computer Science*, pages 505–512. Springer.

Ortony, A., G.L. Clore, and A. Collins. 1988. *The cognitive structure of emotions*. New York: Cambridge University Press.

Osgood, C. E., G. Suci, and P. Tannenbaum. 1957. *The measurement of meaning*. University of Illinois Press.

Parrott, W.G. 2001. *Emotions in Social Psychology: Essential Readings*. Philadelphia: PA: Psychology Press.

Russell, J.A. 1980. A circumflex model of affect. *Journal of Personality and Social Psychology*, 39:1161–1178.

Scherer, K. R. 1984. On the nature and function of emotion: A component process approach. In K.R. Scherer and P. Ekman, editors, *Approaches to emotion*. Erlbaum, Hillsdale, NJ, pages 293–317.

Schröder, M., L. Devillers, K. Karpouzis, J.C. Martin, C. Pelachaud, C. Peter, H. Pirker, B. Schuller, J. Tao, and I. Wilson. 2007. What should a generic emotion markup language be able to represent? In *Proc. of International conference on Affective Computing and Intelligent Interaction (ACII '07)*, volume 4738 of *Lecture Notes in Computer Science*, pages 440–451. Springer.

Schröder, M. 2004. Dimensional Emotion Representation as a Basis for Speech Synthesis with Non-extreme Emotions. In *Proc. of Affective Dialogue Systems: Tutorial and Research Workshop (ADS'04)*, volume 3068 of *Lecture Notes in Computer Science*, pages 209–220. Springer.

Shaver, P., J. Schwartz, D. Kirson, and C. O'Connor. 1987. Emotion knowledge: Further exploration of a prototype approach. *Journal of Personality and Social Psychology*, 52(6):1061–1086.

Storm, C. and T. Storm. 1987. A taxonomic study of the vocabulary of emotions. *Journal of Personality and Social Psychology*, 53.

Strapparava, C. and R. Mihalcea. 2008. Learning to identify emotions in text. In *Proceedings of the ACM symposium on Applied computing*, pages 1556–1560.

Strapparava, C. and A. Valitutti. 2004. Wordnet-affect: an affective extension of wordnet. In *Proc. of International Conference on Language Resources and Evaluation (LREC'04)*, pages 1083–1086.

Sugimoto, F. and M. Yoneyama. 2006. A method for classifying emotion of text based on emotional dictionaries for emotional reading. In *Proc. of International Multi-Conference Artificial Intelligence and Applications (IASTED'06)*, pages 91–96.

Tao, J. and T. Tan. 2005. Affective computing: A review. In *Proc. of Affective Computing and Intelligent Interaction (ACII'05)*, volume 3784 of *Lecture Notes in Computer Science*, pages 981–995. Springer.

Wang, P., J. Hu, H-J. Zeng, and Z. Chen. 2008. Using wikipedia knowledge to improve text classification. *Knowledge and Information Systems*, Online First.

Zhe, X. and A. Boucouvalas. 2002. Text-to-emotion engine for real time internet communication. In *Proc. of International Symposium on Communication Systems, Networks and DSPs*, pages 164–168.