

Ontological Reasoning to Configure Emotional Voice Synthesis

Virginia Francisco, Pablo Gervás, and Federico Peinado

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

Abstract. The adequate representation of emotions in affective computing is an important problem and the starting point of studies related to emotions. There are different approaches for representing emotions, selecting one of this existing methods depends on the purpose of the application. Another problem related to emotions is the amount of different emotional concepts which makes it very difficult to find the most specific emotion to be expressed in each situation. This paper presents a system that reasons with an ontology of emotions implemented with semantic web technologies. Each emotional concept is defined in terms of a range of values along the three-dimensional space of emotional dimensions. The capabilities for automated classification and establishing taxonomical relations between concepts are used to provide a bridge between an unrestricted input and a restricted set of concepts for which particular rules are provided. The rules applied at the end of the process provide configuration parameters for a system for emotional voice synthesis.

1 Introduction

An important challenge in addressing issues of affective computing is having an adequate representation of emotions. Existing approaches vary between identifying a set of basic categories - with a name tag assigned to each one of them - to designing a multi-dimensional space in terms of primitive elements - or emotional dimensions - such that any particular emotion can be defined in terms of a tuple of values along the different dimensions. 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. For assigning emotional values to given utterances, on the other hand, human evaluators find it much easier to provide numbers along given dimensions. If one were to operate computationally with a representation of emotions expressed in more than one format, one is faced with the task of being able to convert from one to another. This task is reasonably easy when converting from emotional categories to emotional dimensions: it would suffice to assign a particular tuple of values for the emotional dimensions of each emotional category. When trying

to convert 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.

A separate problem arises from the fact that there is a large number of emotional categories, and the differences and similarities between them are not clear cut. In some cases, it is reasonable to assume that certain emotional categories may be subsumed by others. For example, the emotion *anger* subsumes the emotions *sulking*, *displeasure* and *annoyance* which may be seen as different types of *anger*. This suggests that a taxonomy of emotional categories as a hierarchy might be useful in finding correspondence between more specific emotional categories and more general 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. 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. Thanks to a taxonomical reasoning system an implicit hierarchy for the concepts represented in the ontology can be inferred automatically.

This paper describes the development of such a system, together with its application as an interface between a text input marked up in terms of emotional dimensions and a set of rules for configuring an emotionally-enabled voice synthesizer. By reasoning over the 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 that satisfies the condition that specific rules are available for generating an appropriate voice synthesis configuration for expressing the intended emotional impression. Section 2 provides a basic introduction to the representation of emotions. Section 3 summarises the Semantic Web technologies employed in this approach. Section 4 describes how input texts are tagged with information describing their emotional content in terms of emotional dimensions. Section 5 gives an overview of the ontology of emotions we have developed. Section 6 describes the operation of the emotional synthesizer. Section 7 provides an example of the complete process for a particular input. Finally, section 8 discusses the technological issues that have arisen, and section 9 summarises our conclusions and future work.

2 State of the Art: Representation of Emotions

This section provides a brief review of the different methods used in the study of emotions in order to classify them. Interested readers can find more detail in the work of Randolph Cornelius [1] and Marc Schröder [2].

Emotions are not an easy reaction, there are a lot of factors that contribute to them. For Izard [3] a good definition of Emotion must take into account: conscious feeling of emotion, process which takes place in the nervous system and in the brain and expressive models of emotion. Emotions take place when something unexpected happens and the so-called “emotional effects” begin to take control.

Many of the terms used to describe emotions and their effects are difficult to tell apart from one another, as they are usually not well defined. This is due to the fact that the abstract concepts and the feelings associate with such concepts are very difficult to express with words. For this reason, there are a lot of methods for describing the characteristics of emotions.

There are different methods in order to represent emotions: *emotional categories* - based on the use of emotion-denoting words -, *descriptions based on psychology* [4] and *evaluation* [1], *circumflex models* - emotional concepts are represented by means of a circular structure [5], so that two emotional categories close in the circle are conceptually similar - and *emotional dimensions* which represent the essential aspects of emotional concepts.

In the following subsections we describe in detail the two methods which are employed in our work: *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 labels of varying degrees of expressiveness for the description of emotional states. There are significant differences between languages in terms of the granularity with which these labels describe particular areas of emotional experience. Even within a given language, some areas of emotional experience have a higher density of labels than others. This diversity presents an additional difficulty. A lot of methods have been proposed in order to reduce the number of labels used to identify emotions. Some of them are listed below:

- Basic emotions: There is a general agreement that there are some emotions that are more basic than others. The number of basic emotions generally is small (in early studies 10, in more recent ones between 10 and 20), so it is possible to characterize each emotional category in terms of its intrinsic properties [1].
- Super ordinate emotional categories: Some emotional categories have been proposed as more fundamental than others on the grounds that they include the others. Scherer [6] and Ortony suggest that an emotion A is more fundamental than other emotion B if the set of evaluation components of the emotion A are a subset of the evaluation components of the emotion B.
- Essential everyday emotion terms: A pragmatic approach is to ask for the emotion terms that play an important role in everyday life. The approach is

exemplified by the work of Cowie [7], who proposed a Basic English Emotion Vocabulary. Starting from lists of emotional terms from the literature, subjects were asked to select a subset which appropriately represents the emotions relevant in everyday life. A subset of 16 emotion terms emerged.

Emotional Dimensions. Emotional dimensions represent the essential aspects of emotional concepts. There are two basic dimensions: *evaluation* and *activation*, occasionally these two dimensions are completed with a third dimension: *power*. *Evaluation* represents how positive or negative is an emotion. For example in a scale for the evaluation dimensions at one extreme we have emotions such as happy, satisfied, hopeful ... the other end of the scale is for emotions such as unhappy, unsatisfied, despaired ... *Activation* represents an active / passive scale for emotions, at one extreme of the activation are emotions such as excited, aroused ... At the other end of this scale are emotions such as calm, relaxed ... The last dimension, *power*, represent the control which exerts the emotion, at one end of the scale we have emotions characterized as completely controlled, such as care for, submissive ... At the opposite end of this scale we have emotions such as dominant, autonomous ... For all dimensions, if the emotion is completely neutral with respect to the emotional dimensions it should be assigned to the middle point of the scale.

This method is very useful because it provides a way of measuring the similarity between emotional states. Another important property of that method is that shifting the representational weight away from the actual labels employed allows for a relative arbitrariness when naming the different dimensions.

3 State of Art: 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 living nowadays.

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

Ontology Web Language. Semantic Web relies heavily on ontologies. Concretely, ontologies based on Description Logics paradigm include definitions of concepts –OWL classes–, roles –OWL properties– and individuals. The most common language to formalize Semantic Web ontologies is OWL (Ontology Web Language [8]), a proposal of the W3C. The goal of this standard is to formalize the semantics that was created *ad hoc* in old frame systems and semantic networks. OWL has three increasingly-expressive sublanguages: OWL Lite, OWL DL, and OWL Full.

OWL Full is powerful for representing complex statements but not useful for reasoning with them due to their computational properties.

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, roles, 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 polihierarchical automatic classification.

Frameworks and APIs. The first thing a novice Semantic Web application developer is searching for is an all-in-one framework or a versatile application programming interface. Java is probably the most important general-purpose language for developing Semantic Web applications, and it is also the language in which the original voice synthesizer was made, so the choice was obvious. But there are at least two very promising Java frameworks available. One of them is Sesame [9], an open source RDF framework with support for RDF Schema inferencing and querying. The other one is Jena [10], 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 (recently added 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.

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 a client building a wrapper for Jena has been the easiest way of working.

DLModel [11] is a very straightforward open source API for accessing a Description Logic model instantiated in an external ontology and knowledge base. Although it has an abstract DL interface (called DLModel), it can act as a wrapper on top of Jena (called JenaModel), offering simple methods to access concepts, roles and individuals of the knowledge base of our Java application .

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 [12] is a hypermedia-based OWL ontology browser and editor written in Java. Is is open souce and it tries to simplify the ontology development using an interface similtar 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 [13], specially the Protégé-OWL version, focuses on editing OWL ontologies. It is a powerful Java open source tool with a user-friendly interface that let 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, the extensibility of its architecture (plug-and-play environment) software allows rapid prototyping and application development, just what we were looking for. But this choice was not an easy decision.

Reasoner. Two different reasoners were considered for this project: Pellet [14] and Racer Pro [15].

Pellet is an open source DL reasoner completely implemented in Java. It deals not only with taxonomical reasoning but also with datatype reasoning, which is very important for our project. Pellet is the default reasoner integrated with SWOOP.

Compared to Racer Pro, a well-know commercial system for OWL/RDF which claims to be the most efficient and robust DL reasoner available, 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. There are educational licenses for Racer Pro, but we have chosen Pellet as a tool for our prototype.

4 Tales Marked Up with Emotions

As a starting point of our approach we have some texts marked up with emotions. In these texts every emotional unit is marked up with the three emotional dimensions (activation, evaluation and power). We are currently using as emotional units the sentences of the text. This implies that every sentence has a value for each of the three dimensions. The emotions associated to each of the sentences try to rate how the listener will feel while listening each sentence as it is read out aloud by the synthesizer.

Texts are marked up with emotions by means of EmoTag [16] a tool for automated mark up of texts with emotional labels. The approach considers the representation of emotions as emotional dimensions. A corpus of example texts previously annotated by human evaluators was mined for an initial assignment of emotional features to words. This results in a List of Emotional Words (LEW) which becomes a useful resource for later automated mark up. EmoTag employs for the assignment of emotional features a combination of the LEW resource, the ANEW word list [17]¹, and WordNet [18] for knowledge-based expansion of words not occurring in either.

A sample part of a marked tale by EmoTag is given in Table 1.

¹ The ANEW word list is a set of normative emotional ratings for a large number of words in the English language. Words are rated in terms of evaluation, activation and power.

Table 1. Fragment of a Marked Up Tale

```

...
<emotion act=9 eval=7 pow=5>"How well you are looking today: how glossy your feathers; how
bright your eye."</emotion>
<emotion act=9 eval=7 pow=5>"I feel sure your voice must surpass that of other birds, just
as your figure does;</emotion>
<emotion act=9 eval=7 pow=5>let me hear but one song from you that I may greet you as the
Queen of Birds."</emotion>
<emotion act=3 eval=9 pow=1>The Crow lifted up her head and began to caw her best, but the
moment she opened her mouth the piece of cheese fell to the ground, only to be snapped up
by Master Fox.</emotion>
...

```

5 Emotional Ontology

We have developed an ontology for all the emotional categories. They are structured in a taxonomy that covers from the basic emotions to the most specific emotional categories. Each of the emotional categories are related with the three emotional dimensions by means of data ranges.

5.1 Structure

Our 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: *Anger*, *Annoyance*, *Displeasure*, *Sad*, *Happy*, *Surprise*, *Fright*, *Horror* ...
- Word: This is the root for the emotion-denoting words, the specific words which 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 there are emotion-denoting words, which are all the words used for denoting *Anger*, *Annoyance*, *Displeasure*, *Terror* ... 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).

It is important to note here that, because the ontology is intended to operate over input in the form of language utterances, the ontology must include the means for representing words. Therefore it includes the specific concept of Word. All actual words handled by the system must be instances of this concept or one of its subclasses. Specific subhierarchies are added to group together all words in a given language.

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 *unhappiness* is an instance of the emotional concept *Sadness*

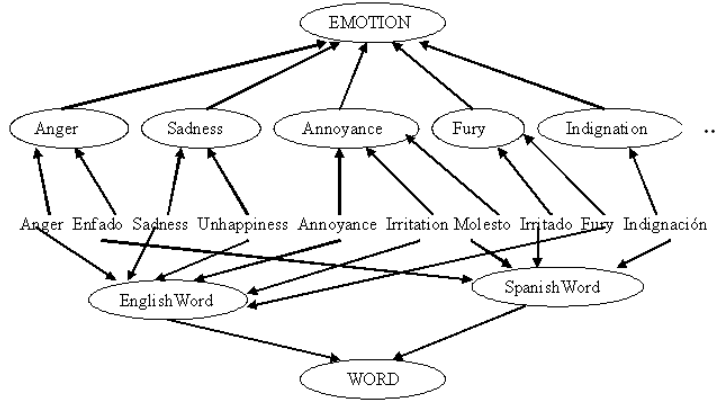


Fig. 1. Fragment of the emotional ontology

at the same time it is an instance of the word concept *EnglishWord*, which means that *unhappiness* is an English word for denoting the emotion sadness.

Another valid way of representing these relations might be to create a new property called “language” to connect each word with 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”.

In handling words, the system may need to identify synonyms for a particular words, that is, other words which may be used to refer to the same concept. Given the semantics we have chosen for our ontology, 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 *annoyance* and *irritation* are synonyms because they are both instances of the Emotion-concept *Annoyance*.

5.2 Datatype Properties

Once we have a hierarchy of emotions, relations between the emotion-denoting words and their language and the concept they represent, we want to link the emotional concepts with the three emotional dimensions. Numeric data can be represented in OWL using datatype properties. To achieve this we have declared three datatype properties: *hasEvaluation*, *hasActivation* and *hasPower*. Each of the emotional concepts is defined by specifying appropriate data ranges for these properties as described in the following section.

5.3 Data Range

We have defined each of the emotional concepts through the emotional dimensions defined as datatype properties. Each emotional concept 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, we have the *Anger* emotional concept, we can describe the region of the space associated to it in the following way: $7 \leq \text{hasActivation} \leq 10$, $0 \leq \text{hasEvaluation} \leq 3$, $3 \leq \text{hasPower} \leq 5$.

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

Table 2. Fragment of the OWL Ontology

```

<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:Restriction>
<owl:Restriction>
  <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>

```

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.

5.4 Automatic Classification of Emotions Using Datatype Properties

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 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, we have the following emotional concepts: *Anger*, *Annoyance*, *Fury* and *Indignation*. Anger is one of the basic emotions and Annoyance, Indignation and Fury are different forms of anger that differ from one another in their intensity of arousal. We define the four concepts as subclasses of the root concept Emotion, and we define the following ranges for the three datatype properties:

- Anger: $7 \leq \text{hasActivation} \leq 10; 0 \leq \text{hasActivation} \leq 3; 3 \leq \text{hasPower} \leq 5$
- Annoyance: $7 \leq \text{hasActivation} < 8; 0 \leq \text{hasActivation} \leq 3; 3 \leq \text{hasPower} \leq 5$
- Indignation: $8 \leq \text{hasActivation} < 9; 0 \leq \text{hasActivation} \leq 3; 3 \leq \text{hasPower} \leq 5$
- Fury: $9 \leq \text{hasActivation} \leq 10; 0 \leq \text{hasActivation} \leq 3; 3 \leq \text{hasPower} \leq 5$

Just by loading the ontology in DLModel, the reasoner automatically classifies the concepts *Annoyance*, *Indignation* and *Fury* as subclasses of the emotional concept *Angry* which is automatically identified as more general.

6 Emotional Synthesizer

EmoSpeech [19] is a system capable of modulating the voice quality of a synthesizer while reading out 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*. The aspects of the voice that act as personality identifiers are: volume, rate, pitch baseline and pitch range. EmoTag uses a group of rules which relates 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 [2], based on the analysis of emotional material generated by actors. The optimal values were obtained through the systematic variation of the parameters during the synthesis. Table 3 summarizes the rules of the synthesizer for the basic emotions.

Table 3. Configuration Parameters for Emotional Voice Synthesis

	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%

7 Example of the Entire Process

The complete process from text input to voice output is described in this section. We have a text as the input of our system. EmoTag marks up this text with the emotional dimensions (*activation, evaluation and power*). Each sentence of the marked up text is related to a point in the three-dimensional space of emotions. This point is the input to our ontology of emotions, which by means of the datatype properties and the dataRange restrictions, automatically classifies this point under a given emotional concept. Once we have identified the specific emotional concept

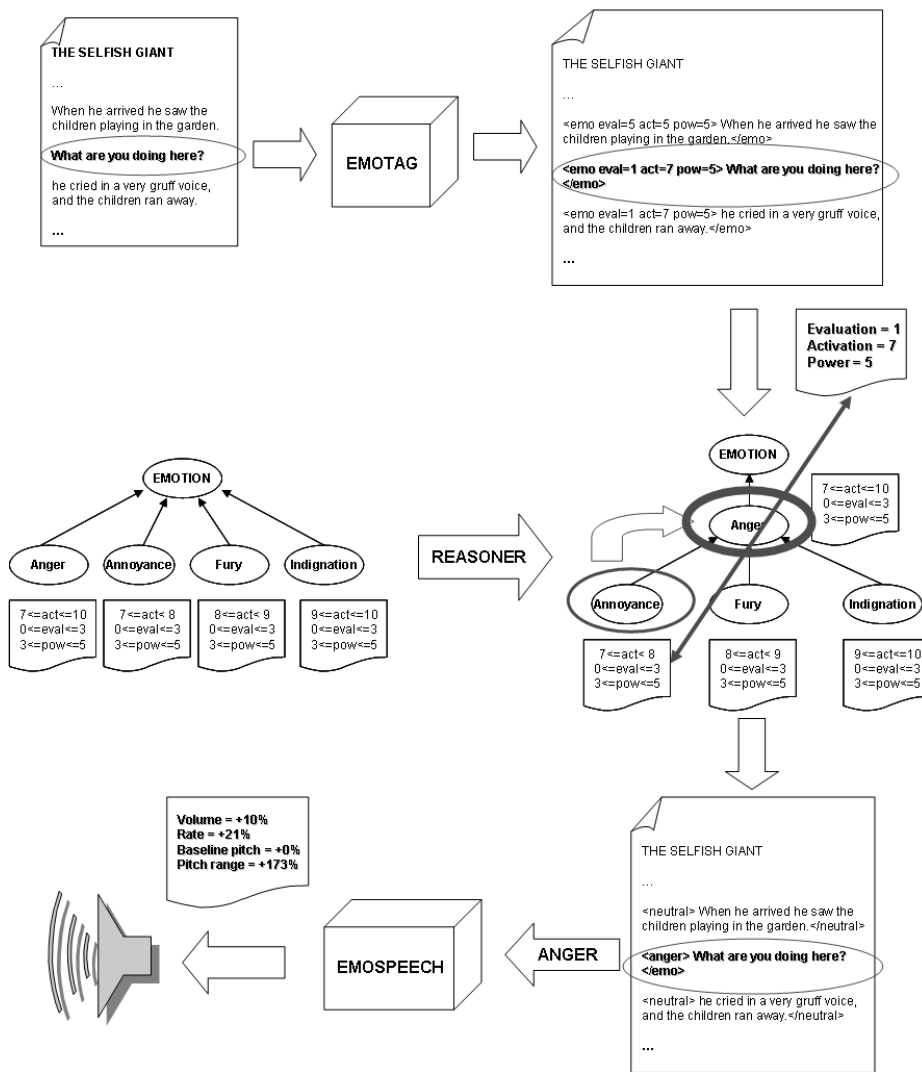


Fig. 2. Example of the entire process

to which the input point is related, by means of DLModel we recursively obtain its ancestors until we locate the one which corresponds to one of the five basic emotions (*anger*, *happiness*, *sadness*, *fear* and *surprise*). Using the particular configuration of parameters for that particular basic emotion, the synthesizer reads out aloud the text with the emotion assigned by EmoTag to the sentences.

In Figure 2 we can see how this process works for a concrete example.

In the example, we have a sentence of input text which EmoTag marks up with the following values: activation = 7, evaluation = 1 and power = 5. This point is classified by means of the 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 emotions. Once EmoSpeech has the suitable rules for the emotional meaning of the sentence, the synthesizer reads out aloud the sentence in an angry way.

8 Discussion

Because the ontology is being used only as interface between the emotional mark up application and the voice synthesizer, its effect on the quality of speech output is limited.² 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, restricted to the five basic emotions. On the other hand, emotional dimensions provide a 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 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 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 paper, 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 to it.

² The quality and emotional precision of the resulting voice has been discussed elsewhere. Details can be found in [19].

Regarding the technologies that have been applied in this proposal, 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. It seems that in the next version of OWL (1.1) support for datatypes is going to be improved, because they are useful for many applications. But the current version of OWL just supports some standard XML Schema datatypes and not 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 makes 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 the new OWL 1.1 will offer those features, but for now other shortcuts must be used in order to reason with restrictions on datatype properties. Protégé 3.2.1 now has a proprietary solution to represent user-defined datatypes, which allows the creation of restrictions with interesting datatype properties and even visualization of the limits of a numeric interval and things like that in the GUI. However, DIG does not allow that kind of information to travel to a DL reasoner. Pellet 1.4, by itself, can deal with user-defined datatype restrictions, and now the last version supports the inline syntax proposed by OWL 1.1³ 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 the “Pellet-Java” default configuration. We hope that some of these shortcomings might be solved in later versions of the technologies.

9 Conclusions

An emotional ontology based on description logics has been implemented using semantic web technologies. Each emotional concept is defined in terms of a range of values along the three-dimensional space of emotional dimensions, that allows the system to make inferences concerning the location of new concepts with respect to the taxonomy. This constitutes a valid solution to the problem of finding a relationship between an arbitrary point in a space of emotional dimensions and the set of basic emotional categories usually identified with specific names. The importance of being able to identify such relationships is strengthened by the fact that configuration of synthesizer parameters for artificially producing emotional voice tends to be established in terms of basic emotional categories.

The ontology described in this paper has demonstrated its usefulness as part of a complex process of converting unmarked input text to emotional voice, resolving the problems that originated at the interface between the emotional tagging in terms of emotional dimensions and the synthesis of emotional voice in terms of basic emotional categories. In this process, both the capability for automatic classification provided by the reasoner, and the hierarchical structure provided by the ontology played important roles.

³ http://owl1_1.cs.manchester.ac.uk/owl_specification.html#4.3

Although reasoning support for datatype properties in OWL DL is still not standard, technologies are available that let us experiment with these features and allow us to develop affective computing applications like the emotional voice synthesizer described in the paper. OWL, Jena, DLModel, Protégé and Pellet are the choices we made before developing this new iteration of the software.

Still more improvements are needed in editors as Protégé to be compatible with reasoners as Pellet. Testing SWOOP is going to be one of our next steps in order to facilitate the acquisition of knowledge for the emotional knowledge base.

Acknowledgements

This research is funded by the Spanish Ministry of Education and Science (TIN2006-14433-C02-01 project) and a joint research group grant (UCM-CAM-910494) from the Universidad Complutense de Madrid and the Dirección General de Universidades e Investigación of the Comunidad Autónoma de Madrid.

References

1. Cowie, R., Cornelius, R.: Describing the emotional states that are expressed in speech. In: *Speech Communication Special Issue on Speech and Emotion* (2003)
2. Schroder, M.: Dimensional emotion representation as a basis for speech synthesis with non-extreme emotions. In: *Proc. Workshop on Affective Dialogue Systems, Kloster Irsee, Germany* (2004)
3. Izard, C.: *The face of emotion*. Appleton-Century-Crofts, New York (1971)
4. Alter, K., Rank, E., Kotz, S., Toepel, U., Besson, M., Schirmer, A., Friederici, A.: Accentuation and emotions - two different systems? In: *Proceedings of the ISCA Workshop on Speech and Emotion*, pp.138–142 Northern Ireland (2000)
5. Russell, J.: A circumflex model of affect. *Journal of Personality and Social Psychology* 39, 1161–1178 (1980)
6. Scherer, K.R.: On the nature and function of emotion: A component process approach. *Approaches to emotion*, pp. 293–317 (1984)
7. Cowie, R., Douglas-Cowie, E., Romano, A.: Changing emotional tone in dialogue and its prosodic correlates. In: *Proc ESCA International Workshop on Dialogue and prosody, Veldhoven, The Netherlands* (1999)
8. Bechhofer, S., van Harmelen, F., Hender, J., Horrocks, I., McGuinness, D.L., Patel-Schneider, P.F., Stein, A.: OWL web ontology language reference. <http://www.w3.org/TR/2004/REC-owl-ref-20040210/>
9. Aduna, NLnet-Foundation: Sesame. <http://www.openrdf.org>
10. Hewlett-Packard: Jena: A semantic web framework for java. <http://jena.sourceforge.net/>
11. Peinado, F.: Dlmodel, a tool for dealing with description logics. <http://federicopeinado.com/projects/dlmodel/>
12. Mindswap: Swoop: A hypermedia-based featherweight OWL ontology editor. <http://www.mindswap.org/2004/SW00P/>
13. Crubézy, M., Dameron, O., Fergerson, R.W., Knublauch, H., Musen, M.A., Noy, N.F., Rubin, D., Tu, S.W., Vendetti, J.: Protégé project. <http://protege.stanford.edu/>

14. Mindswap: Pellet OWL reasoner. <http://pellet.owldl.com/>
15. Racer-Systems: Racerpro. <http://www.racer-systems.com/>
16. Francisco, V., Gervás, P.: Exploring the compositionality of emotions in text: Word emotions, sentence emotions and automated tagging. In: Proceedings of the AAAI-06 Workshop on Computational Aesthetics: AI Approaches to Beauty and Happiness, Boston (2006)
17. Bradley, M., Lang, P.: Affective norms for English words (ANEW): Stimuli, instruction manual and affective ratings. technical report c-1. Technical report, The Center for Research in Psychophysiology, University of Florida (1999)
18. Miller, G.: Wordnet: a lexical database for english. Communications of the ACM 38, 39–41 (1995)
19. Francisco, V., Hervás, R., Gervás, P.: Análisis y síntesis de expresión emocional en cuentos leídos en voz alta. In: Sociedad Española para el Procesamiento del Lenguaje Natural, Procesamiento de Lenguaje Natural, Granada, Spain (2005)