

From Tags to Emotions: Ontology-driven Sentiment Analysis in the Social Semantic Web

Matteo Baldoni, Cristina Baroglio, Viviana Patti* and Paolo Rena

Dipartimento di Informatica, Università degli Studi di Torino, c.so Svizzera 185, I-10149 Torino, Italy

E-mail: {baldoni,baroglio,patti}@di.unito.it, paolo.rena@gmail.com

Abstract. Affective computing is receiving increasing attention in many sectors, ranging from advertisement to politics. This work, set in a Social Semantic Web framework, presents ArsEmotica, an application software for associating the predominant emotions to artistic resources of a social tagging platform. Our aim is to extract a rich emotional semantics (i.e. not limited to a positive or a negative reception) of tagged resources through an ontology driven approach. This is done by exploiting and combining available computational and sentiment lexicons with an ontology of emotional categories.

The information sources we rely upon are the tags by which users annotated resources, that are available through the ArsMeteo platform, and the ontology OntoEmotion, that was enriched by means of our tool with over four hundred Italian emotional words referring to the about eighty-five emotional concepts of the ontology. Tags directly referring to ontological concepts are identified, while potentially affective tags can be annotated by using the ontology, thanks to the spontaneous intervention of the users, in a pure Web 2.0 approach. Finally, the tagged artworks are related with the emerging predominant emotions. A user study involving the ArsMeteo community was conducted in order to evaluate the ArsEmotica outcome, for what concerns the emotions automatically associated by the system to the artworks.

Keywords: Semantic web, ontologies, sentiment analysis, social tagging

1. Introduction

In the last years, the conception of the Web evolved from a web of documents to a web of users: the Social Web. Users are more and more involved in the production of contents or in their elaboration, e.g. by publishing and organizing own materials, by posting comments for discussing newspaper articles, by participating into wikis, by rating resources. Social networks and platforms (Facebook, Flickr, Youtube, LastFM, Anobii, StumbleUpon to cite some among the best known) promote the participation of users in many ways, stimulating the expression of opinions about the contents inserted by other users, by supplying simple “Like” or “Dislike” tools, star-rating systems, tag-based annotation and navigation, and so forth. This

huge amount of data is a precious information source, about perceptions, trends, and feelings, and a lot of research work is being carried on to identify ways for extracting meaningful information from these data.

One of the emerging research fields, aimed at extracting information from the data supplied by the Social Web users, is emotion-oriented computing (a.k.a. Affective Computing [16,5]), whose focus is to automatically recognize the users’ emotions by analyzing their tagging or writing behavior. Since emotions are often related to appreciation, knowing the feelings of the users towards target issues is an important feedback that can support many decisional tasks. The recent success of Sentiment Analysis and Opinion Mining [13,6] techniques applied to business applications, and the development of tools like Twitter Sentiment are significant cases (<http://twittersentiment.appspot.com/>). Such techniques, however, return simple feedbacks,

* Corresponding author. E-mail: patti@di.unito.it.

i.e. the appreciation of a community of users given in terms of positive and negative reception.

This work faces a more complex task: the *identification of the prevalent emotions* associated by the users of a Social Network to (broadly speaking) resources. In particular we have developed an application software, called ArsEmotica, that analyzes tagged artworks from a social tagging platform and provides as output a set of related emotional concepts. Such concepts emerge as the most significant ones for capturing the user's emotions toward a specific resource. Each concept is enriched with a score that expresses the emotion's strength. Specifically, the social platform we worked upon is the Italian art portal ArsMeteo (<http://www.arsmeteo.org> [1]).

Our approach to Sentiment Analysis presents two main novelties. First, the extraction of the emotional content is driven by an *ontology of emotional concepts*, where emotions are structured in a taxonomy, resulting in something richer than a polarized appreciation. In particular, the ontology of emotions that is used by ArsEmotica is based on the one proposed in [10], which was semi-automatically enriched with over four hundred Italian words having an affective value. Second, we focus on very essential form of textual resources: *collections of tags*. In artistic domains, where resources represent artworks, movies, books, users often tag resources to supply *concise reviews*. Thus, by working on tags it is possible to leave aside the complications due to text analysis (often aimed at identifying the keywords inside text) and focus on the emotional analysis.

In order to identify tags bearing an emotional content, we rely on Semantic Web technology and linguistic resources. In particular, we integrated the use of multi-lingual and Italian computational lexicons [15], affective lexicons (WordNet-Affect [20] and SentiWordNet [8]) with the aim of selecting: (i) words that directly refer to emotional concepts (e.g. *paura*, fear); (ii) words that indirectly deliver an emotional content (e.g. *infinito*, infinite).

Given a tagged artwork, first, ArsEmotica selects those tags referring to the emotional concepts that are included in the ontology. Intuitively, such tags belong to the Italian affective lexicon embedded in the ontology, and are considered as directly bearing the affective meaning expressed by the emotional concepts they refer to. Then, ArsEmotica applies traditional Sentiment Analysis (by means of SentiWordNet) to the set of tags which are not recognized into the ontology, with the aim of measuring the affective potential of

such tags. However, SentiWordNet does not associate emotions to tags but it only distinguishes between objective/neutral senses and subjective senses of terms. Thus, for determining the emotional content of tags resulting as potentially affective, our system involves the community of the Social Network and asks the user to give an emotional annotation driven by the ontology. Users can annotate subjective tags with emotions in a pure Web 2.0 approach by specifying scores: such scores represent the user's measure for the semantic affinity of the term with one or more emotions, chosen from the ontology.

After collecting all the emotions related to the artworks by the tag analysis, ArsEmotica ranks the resulting emotions in order to identify the prevalent ones for the tagged resource. The algorithm that we have implemented is inspired by the one in [10]. It reasons on the ontology hierarchy and exploits the ontology of emotions as well as the associations of tags to emotional concepts to compute its output.

The paper begins with a brief overview of the background. Section 3 describes the architecture of ArsEmotica and the three main steps which lead to extract the emotional semantics from tagged resources. Section 4 reports a case study in the *Art* domain (including a first evaluation of part of the analysis performed by the system), followed by a scenario that outlines prospected applications (Section 5). Final remarks end the paper.

2. Background

The research, that we have been carrying on, strictly embraces or is in some way related to different technologies and topics, including the Social Semantic Web (folksonomies and ontologies), classical computational lexicons (MultiWordNet [15], WordNetAffect [20], WordNet 3.0, SentiWordNet 3.0 [8]), studies on emotions in social, psychological, and computational contexts, and sentiment analysis. Let us briefly overview the background.

One of the key problems in the Web is how to index resources so as to efficiently and effectively retrieve them. Ontologies and folksonomies are two ways for indexing resources: the former are to be designed by knowledge engineers, while the latter are spontaneously produced by the *tagging* activity of the members of a community. Tagging is one of the ways offered to users by the so called Web 2.0 to become actively involved into the web experience, and amounts

to attaching freely chosen labels to resources. Often such labels are used to categorize the resources, but they are also used to express reception, opinions, feelings. Our choice to work on tags has the advantage of allowing us to focus on the emotional analysis, avoiding those complications that are due to text analysis. Nevertheless, a technique like the one that we propose can easily be extended to the analysis of the latent emotions of a text, given proper linguistic tools, because it is orthogonal from the way in which key words are identified. Moreover, when the user's feedback is supplied in a language like Italian, for which there are few affective lexical resources that are freely available for the development of new application softwares, tags can be the only processable information source that can be used.

Whatever they represent, tags are used as meta-data on top of which it is possible to devise many algorithms for navigation and retrieval. More precisely, these algorithms exploit the folksonomy, arising from tags, using it as an open, distributed and social classification system, though with a *flat* structure. This final remark suggests one of most challenging tasks that are currently being studied in the field, which is the identification of relationships that tie the folksonomy terms with one another, and possibly capture a machine-processable semantics. There are many attempts to reconcile folksonomies with ontologies, for instance by inducing ontologies from folksonomies or by matching terms in some way. It is out of the scope of this paper to get into the details. We would just like to mention the survey by Dotsika [7], together with a previous work of the authors [2], where the association was done based on the outcomes returned by a search engine. The reason for these interests is that, indeed, the Social Web aims at developing applications that combine the ease of use, which is typical of its platforms, with the advantages deriving from a formal semantics, i.e. interoperability, data/service integration, personalization, better recommendation and retrieval performances [18].

In this context, the identification of the feelings of a community or of its single members is receiving an increasing attention, as an indicator of the appraisal of topics, people, situations, resources, trends. Hence the development of Opinion Mining [13], of ontologies of emotions, like the one we started out work from [9], and of W3C markup language proposals like Emotion ML [17]. However, there are still few applications that use the most advanced results in Semantic Web technology to deal with emotions and most of the ap-

proaches use ontologies where emotions are individual isolated units (e.g. WordAffect). Such considerations motivated our focus on the ontology of emotions in [9], an OWL ontology where emotions are structured and organized in levels, trying to integrate the results of the most recent psychological models. As we will see in the next section, OntoEmotion provided a good starting point to explore an ontology-driven approach to Sentiment Analysis, where tags (and then tagged resources) are related to emotions. Such approach, to the best of our knowledge, is original w.r.t. previous work on Sentiment Analysis and allows to extract from tags affective information which is richer than a polarized appreciation.

The study of emotions is particularly relevant in artistic domains where it is a common experience for users to share the feelings kindled by an artwork (be it a painting, a video or some music) with friends. As a consequence, a software which, starting from tags freely associated to resources, can extract a rich emotional semantics, could find many interesting applications. For instance, by creating new, emotion-driven navigations of the items shown by on-line museum portals [22,4], or by creating apps for portable devices, on top of which creating new ways of participating to art exhibitions. Moreover, most of current portals allow users to express their appreciation on artworks by means of simple star-rating systems. Our ontology-driven approach could provide users with new means for expressing and sharing emotions stirred by artworks.

3. ArsEmotica

This section describes the architecture of *ArsEmotica*, the application software that we developed. The analysis steps that we are about to describe rely on a pre-processing phase in which tags are filtered so as to eliminate flaws like spelling mistakes, badly accented characters, and so forth. Figure 1 reports the three main steps that characterize the computation after the pre-processing:

Step 1: Checking tags against the ontology of emotions. This step checks whether a tag belongs to the ontology, that is whether it directly refers to an emotional category included in the emotional ontology hierarchy. Tags belonging to the ontology are immediately classified as "affective" and the information on the link between the tag and the emotional category at issue is stored.




Fig. 1. ArsEmotica overall architecture.

Step 2: Sentiment analysis and user's feedback

Tags that do not correspond to terms in the ontology are further analyzed by means of *SentiWordNet*, in order to distinguish *objective* tags, which do not bear an emotional meaning, from *subjective* and, therefore, affective tags. When at least one meaning of a word has a relevant sentiment score (positive or negative), the tag is presented to the user for getting a feedback on which emotional concept it delivers. Users can select one or more emotions from the ontology and relate them to the tag at issue with a strength value (a score), which intuitively represents the user's measure for the semantic affinity of the term with the chosen emotion.

Step 3: Ranking of Emotions. Based on data collected in the previous steps, the tool ranks the emotions associated by the users to the resource.

The following sections explain in details how the extraction of an emotional semantics is performed.

3.1. The ontology of emotions and the Italian emotional words

The first step checks if the tags of a given resource are "emotion-denoting" words directly referring to some emotional categories of the ontology of emotions. In the following we describe the ontology of emotions adopted for our purpose and the methodology used for populating the ontology with Italian

emotion-denoting words. Our starting point was *OntoEmotion*, an emotional ontology developed at Universidad Complutense de Madrid [9], that met our requirement to have a taxonomic structure, mirroring well-founded psychological models of emotions, and that was implemented by using semantic web technologies. The ontology is written in OWL [11] and structures emotional categories in a taxonomy, which covers basic emotions as well as the most specific emotional categories, including 87 emotional *concepts*. The basic emotions are *Sadness*, *Happiness*, *Surprise*, *Fear* and *Anger* (Level 1 of the taxonomy). The taxonomic structure basically refers to the psychological model by Parrot [14], adapted to these five basic emotions, and integrated with emotions which appear in other well-established models.

In particular, under each basic emotion we have a number of levels which may vary and depends on the level of available specification for it. For instance, *Anger* is specialized in *Agitation*, *Annoyance*, *Hate*, *Fury*, *Displeasure*, *Sulking*, etc.), while *Surprise* has only one sublevel with two emotional concepts: *Amazement* and *Intrigue*.

OntoEmotion has been conceived for categorizing emotion-denoting words. Classes corresponding to the emotional concepts were originally populated by about 250 instances, consisting in emotion-denoting words of English and Spanish. The ontology has two root concepts: *Emotion* and *Word*. *Emotion* is the root for all the emotional concepts. *Word* is the root for the

Levels	Emotions
L1	Anger, Fear, Happiness, Sadness, Surprise
L2	Agitation, Annoyance, Arrogant, Displeasure, Envy, Fury, Hate, Hostility, Indignation, Rancour, Sulking, Alarm, Anxiety, Apprehension, Consternation, Courage, Distress, Dread, Fright, Horror, Panic, Paranoia, Phobia, Trepidation, Worry, Admiration, Affection, Ecstasy, Enthusiasm, Euphoria, Fascination, Glee, Gloating, Gratification, Hope, Jubilation, Pleasure, Relief, Satisfaction, Solidarity, Sympathy, Boredom, Contempt, Depression, Desolation, Disappointment, Discouragement, Gloom, Humiliation, Hurt, Indecisiveness, Nostalgia, Powerlessness, Regret, Rejection, Amazement, Intrigue
L3	Frustration, Grief, Terror, Care_for, Compassion, Liking, Excitement, Obsession, Pride, Apathy, Impatience, Despair, Confusion, Helplessness, Remorse, Disgust
L4	Attraction, Love, Lust, Shame
L5	Adoring, Passion, Arousal, Desire
L6	Infatuation

Table 1

Emotional concepts organized by levels according to the OntoEmotion hierarchy [9]. The colors denote subtrees in the ontology, having as root one of the top level emotions.

emotion-denoting words, i.e. the words that each language provides for denoting emotions, and originally had two subclasses: *EnglishWord* and *SpanishWord*. Each instance of these two concepts has two parents: one is a concept from the *Emotion* hierarchy (the type of emotion denoted by the word), while the other is a concept from the *Word* hierarchy (e.g. the language of the word). For instance, the word *rage* is both an instance of the concept *Fury*, and an instance of the concept *EnglishWord*, which means that *rage* is an English word for denoting fury.

Since the tags used in our case study are mainly Italian words, we enhanced the ontology by adding a new subclass *ItalianWord* to the root concept *Word* and semi-automatically populated the ontology. The approach we applied relies on the use of the multilingual lexical database MultiWordNet, in which the Italian WordNet is strictly aligned with Princeton WordNet 1.6., and its affective domain WordNet-Affect, a well-known lexical resource that contains information about the emotion that the words convey. A human expert checked the identified terms.

WordNet is a lexical database, in which nouns, verbs, adjectives and adverbs (lemmas) are organized into sets of synonyms (synsets), representing lexical concepts. After choosing the representative Italian emotional words for each concept, such words were used as entry lemmas for querying the lexical database. The result for a word is a set of synsets, representing the 'senses' of that word, and labeled by MultiWord-

Net unique synset identifiers. Each synset was then processed by using WordNet-Affect: when a synset is annotated as representing affective information, then all the synonyms belonging to that synset are imported in the ontology as relevant Italian emotion-denoting words. This allowed us to automatically enrich the ontology with synonyms of the representative emotional words, but also to filter out synsets which do not convey affective information. Let us see an example. When we query the MultiWordNet database with the Italian word *panico* (noun, representative for the emotion *Panic*), only two out of the three resulting synsets are affective (WordNet senses *n#10337390* and *n#05591377*). In particular, the third not affective synset refers to the sense of the word "panico" described by the following gloss: "*coarse drought-resistant annual grass grown for grain, hay and forage in Europe and Asia and chiefly for forage and hay in United States*". Thanks to our affective filter we can exclude words belonging to that synset (*Setaria_italica*, *pabbio_coltivato*) when populating the concept *Panic* of our ontology.

After the not affective synsets were filtered out, we automatically populated the ontology by classifying the Italian words, belonging to the selected synsets, under the proper emotional concepts. A human expert checked and polished the results, to avoid wrong effects of synset intersections, which are very natural in such a restricted domain. Intuitively, the idea behind this post-processing is that the Italian words, we

have chosen as representative lemmas for an emotional category, must univocally refer to that category, while others, which may belong to synsets of different emotional lemmas, can have a multiple emotional membership.

The resulting ontology contains about 450 Italian words referring to the 87 emotional categories of On-toEmotion. In order to keep trace in the ontology of the synonymy relations among words belonging to the same synset, we have defined the OWL object property *hasSynonym*. For instance, the following code excerpt corresponds to the description of the Italian word *stupore*. The term *stupore* is an instance of *ItalianWord* and of *Amazement*, and has all the listed synonyms, in some of its affective senses:

```
<ItalianWord rdf:about="#stupore_">
  <rdf:type
    rdf:resource="&emotions_v9;Amazement" />
  <rdf:type rdf:resource="&owl;Thing" />
  <hasSynonym rdf:resource="#attonimento_" />
  <hasSynonym rdf:resource="#meraviglia_" />
  <hasSynonym rdf:resource="#sbalordimento_" />
  <hasSynonym rdf:resource="#sbigottimento_" />
  <hasSynonym rdf:resource="#sconcerto_" />
  <hasSynonym rdf:resource="#stupefazione_" />
</ItalianWord>
```

ArsEmotica uses the enhanced ontology for checking if a tag describing a resource *directly* refers to some emotional category (Emotional Analysis in Fig 1). In case it does, the tag is immediately classified as “emotional”. The information collected during this phase is stored as a set of triples having the form: (t, e, s) , meaning that tag t is related to emotion e with a strength value s . The range of the score s is $[0, 100]$. Such triples are stored in a data base table, of the kind reported in Table 2. When a tag is an instance of an emotional concept, the strength will be 100. So, for example, since the word “spavento” is an instance of “Amazement”, the corresponding triple will be: (“spavento”, “Amazement”, 100).

3.2. Sentiment analysis and User Feedback

The previous analysis identifies a set of tags as directly bearing an affective meaning. However, other tags can potentially convey affective meaning and *indirectly* refer to emotional categories of the ontology. As observed in [21], some words can be emotional for someone due to her individual story. In other cases the affective power is part of the collective imagination (e.g. words like “war”). As a consequence, it seems to

be appropriate and promising to involve the community in the definition of such indirect affective meanings.

In order to minimize the effort requested to the users, before offering the tags to their judgment, we select the most promising ones by using SentiWordNet 3.0, a lexical resource for opinion mining where synsets of Princeton WordNet 3.0 are annotated according to their degree of neutrality, positiveness and negativity. Each synset s is associated the scores $Pos(s)$, $Neg(s)$ and $Obj(s)$ indicating how neutral (Obj) or affective (Pos and Neg) the terms contained in the synset are. Each score ranges in $[0,1]$ and their sum is 1 for each synset. Sentiment analysis is performed by applying the steps sketched in Figure 2 to each tag that does not belong to the ontology of emotions. Since SentiWordNet was created for the English




Fig. 2. Steps of the sentiment analysis on tags.

language, we needed to use MultiWordNet to align the Italian lemmas corresponding to the English ones. Moreover since SentiWordNet annotates a newer version of Princeton WordNet (3.0) with respect to the version MultiWordNet is based on (1.6), we have to query such newer lexical database.

As an example, let us consider the Italian word *infinito*. MultiWordNet returns many synsets denoted by the identifiers $n\#04767390$, $a\#00955446$, $a\#00877590$. Each of them has a corresponding English lemma. In particular, $n\#04767390$ refers to the English word “infinitive”; $a\#00955446$ to “infinite”; while the last synset refers to “inexhaustible”. We now use these three English words as entry lemmas for querying WordNet 3.0. The result is the list of all their possible meanings: (*infinite*: 00028651): the uninflected form of the verb; (*infinite*: 01007354): having no limits or boundaries in time or space or extent or magnitude; (*infinite*: 01008745): of verbs, having neither person nor number nor mood (as a participle or gerund

or infinitive); (*infinitive: 00301951*): too numerous to be counted; (*inexhaustible: 00005718*): incapable of being entirely consumed or used up; (*inexhaustible: 01008289*): that cannot be entirely consumed or used up.

Since we cannot know which sense the user had in mind, we keep them all and query SentiWordNet for getting the sentimental analysis of the synsets:

- (*infinite: 00028651*), Pos(s): 0, Neg(s): 0;
- (*infinite: 01007354*), Pos(s): 0.125, Neg(s): 0.5;
- (*infinite: 01008745*), Pos(s): 0, Neg(s): 0;
- (*infinitive: 00301951*), Pos(s): 0, Neg(s): 0;
- (*inexhaustible: 00005718*), Pos(s): 0.25, Neg(s): 0.375;
- (*inexhaustible: 01008289*), Pos(s): 0.375, Neg(s): 0.25;

The objectivity of a word in a given sense is simply measured as $1 - (Pos(s) + Neg(s))$. The value 1 indicates that the term is objective, while 0 means that the term conveys a strong sentimental (positive or negative) meaning. Different senses of the same term can have different opinion-related properties and different scores. When no sense of a given term has a significant sentimental score, we conclude that it is mainly descriptive and that it does not evoke emotions. Therefore, we ask the community to evaluate only those terms, having at least one meaning with a relevant sentimental score. This was done to have no false negative. So, for instance, since the word *infinito* has a relevant sentimental score for some of its senses, it will be proposed to the users for evaluation.

For all tags, resulting potentially affective, like *infinito*, users will be free to choose one or more emotions from the emotional categories of the ontology, specifying the related strength value. Again a set of triples (t, e, s) will be collected during this phase. The triples are stored in a data base table, having the form reported in Table 2. Notice that the scores can now be less than 100.

	Emo ₁	Emo ₂	...	Emo _n
Tag₁	s ₁₁	s ₁₂	...	s _{1n}
Tag₂	s ₂₁	s ₂₂	...	s _{2n}
...
Tag_n	s _{n1}	s _{n2}	...	s _{nn}

Table 2

Tabular representation of (t, e, s) .

3.3. Getting the predominant emotions

Once the analysis of the tags associated to a resource is finished, during the last step ArsEmotica ranks the emotions associated by the users to the resource and computes the prevalent emotion. This is done with the help of the Jena Reasoner applied to the triples resulting from the previous analysis steps. The implemented algorithm relies on the taxonomic structure of the ontology and is inspired to the one in [10], where an analysis is performed to emotionally mark up a sentence by analyzing the words that compose it. Intuitively, the algorithm allows to select the *most relevant emotion*, which represents the affective information related to the artwork, from the scores that relate each of its tags to the various emotions. The basic steps are:

- processing the emotional concepts appearing in the triples (t, e, s) , so as to identify also those emotions in the ontology that are related to the ones appearing in the triples. The identified emotional concepts can be organized into *layers* by following the parent-child relationship as shown in Table 1. In this phase a Jena Reasoner has been applied to the collected triples;
- starting from the leaves and moving upward towards the root, compose and propagate the scores.

As an example, let us suppose to start from the triples obtained for the artwork “Bianca e il suo contrario” shown in Fig. 3:

(*felicita'*, *Happiness*, 100)
 (*tristezza*, *Sadness*, 100)
 (*male*, *Hurt*, 50)
 (*male*, *Distress*, 50)
 (*bene*, *Love*, 50)
 (*bene*, *Affection*, 50)

All the above tags were identified as conveying an emotional value. Those, whose score is less than 100, are related to different emotional concepts (e.g. *bene* and *male*), therefore, their scores are distributed among such classes. For instance, the emotional value of the term *bene* is equally distributed between the emotional concepts *Love* and *Affection*.

The identified emotional concepts are divided into their corresponding layers (see Table 3) and each of them is associated the normalized scores. The algorithm rebuilds the emotional tree relating the emotions that appear in the triples, by using information about the parent-child relationship encoded in OntoEmotion. Then, the scores of the emotional concepts are com-




Fig. 3. Emotion-denoting tags for an artwork

Layer 1	Happiness (25%)	Sadness (25%)
Layer 2	Affection (12.50%) Hurt (12.50%)	Distress (12.50%)
Layer 4	Love (12.50%)	

Table 3

Emotions extracted from “Bianca e il suo contrario” divided in layers.

posed with those of the respective parents, obtaining the final score (see Table 4). In particular, according

Layer 1	Happiness (50%)	Sadness (37.50%)
Layer 2	Affection (25%) Hurt (12.50%)	Distress (12.50%)
Layer 4	Love (12.50%)	

Table 4

Final table that allows identifying the prevalent emotions.

to the OntoEmotion hierarchy [9], *Hurt* contributes to *Sadness* because the former emotional concept is child of the latter, while *Love* contributes to *Affection* that, in turn, contributes to *Happiness*. *Distress* does not contribute to any emotional concept in the first layer (it is child of the basic emotion *Fear*, which is not present). As a result, the prevalent emotion is *Happiness* although *Sadness* has quite high score too.

4. Case Study

The ArsEmotica prototype (developed in Java and NetBeans) was tested against a corpus of tagged

multimedia artworks from the *ArsMeteo* art portal (<http://www.arsmeteo.org> [1]), on-line since 2006. First tests involved the analysis of 72 artworks and of over 1000 tags, reflecting the spontaneous annotation of the resources produced by the ArsMeteo community (mainly living Italian artists and art lovers) in the last years¹. In particular, 38 images out of the 72 sample artworks bear an emotional meaning according to ArsEmotica and they were selected in order to build an image corpus, where to evaluate the match between the ArsEmotica outcomes and the perceptions of real users of the ArsMeteo community (see Section 4.1).

The prototype implements the architecture designed in Figure 1, so it is characterized by three stages, that can be enacted by pressing in sequence the buttons *Tag to Emo*, *Tag to Senti* and *Emo Engine* (see Figure 3 for a screenshot of the interface). Given a tagged artwork, by pressing the *Tag to Emo* button the user queries the OWL ontology of emotions; the result is a set of tags referring to emotions. For instance, by applying the emotional analysis to the artwork “Bianca e il suo contrario” by Marzia Migliora (2007, Figure 3), four tags are identified as emotion-denoting words. In particular, the tag “tristezza” refers to emotion *Sadness*, “felicità” refers to *Happiness*, “bene” refers to *Love* and *Affection* (which are both concepts related to *Happiness*), while, finally “male” refers to *Distress* and *Hurt*.

The second stage works on sets of tags which do not belong to the ontology. By pressing the *Tag To*

¹See <http://www.giorgiovacarino.it/mostre/emozioni.html>, click on images for details.




Fig. 4. User's feedback

Senti button, the affective potential of these tags is studied by querying SentiWordNet. This lexicon classifies terms as subjective or objective, returning a positivity score and a negativity score, from which it is possible to compute the objectivity score. We consider as tags possibly conveying an emotional meaning all those whose objectivity score is less than the threshold 0.5. For these, we ask the user to give an annotation driven by the ontology of emotions, as explained in Section 3.2. For instance, in our running example, during the second stage, subjective tags, like “lutto” (mourning) or “inquietante” (disturbing) are selected and offered to the user’s evaluation, while objective and mainly descriptive tags, like “bianca” (white) and “nera” (black) are not presented. Figure 4 shows the input form that allows the user to supply feedback about the selected tags. In the example the tag “lutto” has been associated by a user to *Compassion* with a score 80%.

The third stage amounts to compute the predominant emotions and can be enacted by pressing the *Emo Engine* button. We are currently completing the prototype in order to allow the formatting of the result according to the standard markup language *Emotion ML* by W3C [17].

The application is written in Java. Among the classes, particularly relevant to mention are:

- *DB_ArsMeteo* implements the methods needed to query the ArsMeteo data base and to collect the resource/tag data;
- *GetOntoWords* implements the methods needed to extract the emotional concepts from the ontology and the corresponding English terms;
- *EmotionTBuild* implements the methods needed to retrieve from MultiWordNet the *synset identifiers* of the terms, which correspond to the ontological concepts. Such synsets are used by *WNAffect* (see below) to check the membership of the synset to WordNet-Affect. If the answer is positive, MultiWordNet is queried again to find the Italian words of these specific synsets. These processing steps produce the ontology *emotionT.owl*;
- *EmotionT* allows to navigate through the emotional ontology and retrieve information;
- *WNAffect* checks whether a synset belongs to the WordNetAffect domain;
- *SentiWordNet* performs the sentimental analysis. It is used on the tags which do not belong to the ontology. For each of them, the corresponding English term identifier is looked for in MultiWordNet. From this, we look for the corresponding reference code in WordNet 3.0. In this way, it is possible to retrieve the related positivity and negativity scores and compute the objectivity value;

- *ReasonerJena* is used to analyze emotions. For instance, given an emotion and an ontology, it allows to gather all the relatives of the concept node of the ontology corresponding to that emotion;
- *ObjWordEmo*, implements the algorithm for identifying the prevalent emotions.

4.1. Evaluation and User Study

For testing the effectiveness of the proposed system, we started a series of experiments aimed at validating the ArsEmotica results. So far, we focussed on the validation of the emotional analysis performed by the system (Step 1). This is actually the core of ArsEmotica because it is the part of analysis that is performed in a totally automatic way, without involving users. Table 5 reports as an example some of the outcomes. It reports the name of the artwork, a thumbnail of the image, the identified affective tags and the related emotions².

For validating the first step, we conducted a user study by involving the ArsMeteo community (285 users), i.e. the same community of artists who in part produced and in part tagged the artworks on which experiments were conducted. First of all, we randomly selected a set of 38 artworks³ from the ArsMeteo corpus, among those which are the most voted by the community. The set includes also artworks which provoke contradictory reception, e.g. artworks with multiple classifications, meaning that they are capable of dividing the community in its perceptions. Then, we used ArsEmotica to compute, for each artwork a_i , the set of associated emotions, by analysing the tags attached to images by the community. The number of emotions associated to each artwork in the corpus ranges from 1 to 9, with an average value of 2.9; the corpus generated 44 different emotions out of the set E of 87 emotional concepts, contained in the ontology.

Users were asked to answer a questionnaire composed of ten questions, belonging to two categories. In the first kind of question, users were asked to associate to a randomly chosen artwork the emotion that better expresses its emotional value by selecting between one of the emotions that ArsEmotica identified as associated to the artwork (some examples are shown in Table 5), and one that was not. In the second kind of question, given a randomly selected emotion, and two

images, only one of which resulted as actually conveying the selected emotion –according to ArsEmotica–, users were asked to indicate which of the two was more related to that emotion.

On the whole, 35% of the users answered to the questionnaire. The collected results show a clear correlation between the emotions chosen by the users and those extracted by the ArsEmotica system. Experimental results are described and analysed in details in [3].

5. Prospected application

The system we have developed can be used as a basis for designing interesting application softwares. In this section we make an example, framed in an artistic domain. Seeing artworks exposed at a fair or gallery has a strong emotional impact on the visitors. A recent trend is to present the exposed artworks in multimedia environments which allow visitors to interact to some extent with the artworks themselves. In this context one could envisage the use of portable devices to enable visitors, but also artists, to express their own reception of the exposed works, by annotating them with tags. To this aim, it is possible to use well-assessed technologies, like semacodes (two-dimensional barcodes, created on the DataMatrix standard). Each exposed piece would be given a semacode label encoding a URI. Since a semacode tag can easily be read by a semacode handler, installed on a portable device that is equipped with a built-in camera, it provides an easy means to access and display information about the artwork of interest, and allow the tagging activity (see [2]).

The tagging activity of the community provides the basis for an emotion-driven browsing of the artworks. To this aim, it is possible to apply a solution based on the ArsEmotica software. The emotional analysis of tags could, then, be used as a basis to create new relationships, not only among apparently unrelated tags but also among artists and among artworks and, in particular, for offering the users the “navigation” of artworks in a common emotional space. This can be done in the physical reality of the exhibition, for instance by suggesting different emotion-related visiting paths, and in the virtual reality offered by the exhibition web site by representing in a graphical way (e.g. by exploiting a coloring code) the mood of the rooms in which the exhibition takes place.

ArsEmotica exploits such relations. For example, Agostino could start the search by querying for art-

²For sake of simplicity, in the current prototype we did not process composite tags.

³The whole set of artworks can be seen at <http://di.unito.it/arsemocorpus>.


Artwork	Title and artist	tag [Emotions]
	<i>Bianca e il suo contrario</i> by M. Migliora	“tristezza” [Sadness (L1, 50%)] “felicità” [Happiness (L1, 37.50%)] “male” [Distress (L2, 12.50%), Hurt (L2, 12.50%)] “bene” [Love (L4, 12.50%), Affection (L2, 25%)]
	<i>Forse un giorno arriveranno al mare</i> by C. Guasti	“desiderio” [Desire (L5, 33.33%)] “stupore” [Amazement (L2, 33.33%)] “paura” [Fear (L1, 33.33%)]
	<i>Da lontano il suono</i> by S. Minniti	“ossessione” [Obsession (L3, 33.33%)] “angoscia” [Apprehension (L2, 8.25%), Trepidation (L2, 8.25%), Anxiety (L2, 8.25%), Worry (L2, 8.25%)] “nostalgia” [Nostalgia (L2, 33.33%)]
	<i>Angry Germ</i> by A. Caligaris Cappio	“ira” [Anger (L1, 50%)] “furioso” [Fury (L2, 50%)]

Table 5

Some examples of artworks and the emotions they raise, identified by the tag analysis.

works related to the emotional concept *Happiness*. As a result, he will get a tag cloud, consisting of all the tags of the folksonomy resulting to have a high correlation with happiness, or with the Spanish happiness-denoting word. Then, as in normal tag-based navigation, Agostino could choose a tag and retrieve information about the artworks described by that tag and about their location. Intuitively, the idea is that tags in the cloud link Agostino to artworks having an emotional relation with happiness. Artworks are usually tagged with many words, which express a variety of meanings and thus support the emergence of different emotional potentials. This is consistent with the idea that art can emotionally affect people in different ways. In this case, Arsemotica would provide Agostino multiple accesses to the artworks, driven by different emotional concepts. Notice that our emotional engine can be in principle be interfaced, by developing a simple API, with any Web 2.0 platform (i.e. Flickr, Youtube or Arsmeteo) that shows the standard functionalities of a social resource sharing system, i.e. collecting, and presenting, browsing and accessing digital resources (artworks in our case) together with their tags.

To complete the vision, when later the visitor will take another look at the artworks in the exhibition web site, he will experience a new virtual visit and discover

emotional relations among the exhibition’s contents, that can be captured only by analyzing the individual feelings expressed by the visitors. The browsing experience will be different from time to time. Initially, just after the phase when artworks are uploaded, the position of the artworks in the emotional space will be mainly determined by the interpretations of the involved artists and curators. They will be the first ones to add meanings to the artworks, and then to give an input to Arsemotica. Later when the system will start to collect the new meanings expressed by the visitors, artworks will start to float on the emotional space in unpredictable ways. New artworks and meanings can be added anytime, so the emotional relations will continuously change, by reflecting the evolution of the community and its latent perception of a sort of impermanent emotional ‘zeitgeist’.

6. Conclusion and Future Work

This work presents a software solution that, by combining lexicons and libraries that are already available, allows both the population of an ontology of emotions (based on [9]) with Italian emotion-bearing terms and the extraction of the prevalent emotions from the set of

tags associated to a resource. The extracted information is richer than a polarized appreciation, as instead usually done by sentiment analysis. A user study was conducted on a subset of the artworks from the art portal ArsMeteo; it involved real users of the ArsMeteo community, who also contributed in the tagging activity, and concerned the validation of ArsEmotica emotional analysis (which focusses on the identification of tags with a direct emotional meaning). Further experiments on tags can indirectly bear an emotional value require the identification of an effective way for motivating users to annotate such tags by means of emotional concepts from the ontology. In order to face this issue, one promising direction could be to rely on the Game With A Purpose [23] paradigm and to develop a proper game in which users, as a side effect of playing, perform the task of associating emotional concepts to tag-words. This is on the line of recent approaches which face the challenge of increasing the user involvement in building the Semantic Web [19]. An alternative could be to integrate in ArsEmotica the use of automatic techniques, e.g. the one proposed in [2], for identifying the association of terms having an emotional value (that is recognized by the sentiment analysis step) with the proper ontological concepts.

The set of tuples that we collected by means of our tool are stored in a data base. They actually form an interesting corpus of data. In fact, by applying statistical analysis techniques or also data mining techniques, it would be possible to measure the feelings of the community towards the content of the artworks and to monitor if and how this changes along time. Emotional reception, in fact, is affected by events that occur in a lifetime. Suppose, for instance, that a painting represents the twin towers. Surely the emotions associated to the painting would have changed after “September, 11”.

The proposed approach is particularly suitable to application domains where tags can be interpreted as *concise reviews* (e.g. artworks, books, movies). Given appropriate pre-processing tools capable to extract the relevant words from a text, its use could be extended also capturing the latent emotions behind textual comments.

The current prototype can be refined in many ways. For what concerns the pre-processing of tags, we intend to improve the current prototype by applying lemmatization and word-similarity algorithms. For instance, in Italian, adjectives are declined in many ways, depending whether they refer to males or females, singular or plural. Stemming and lemmatization algo-

ritms would help reduce the noise due to this variability. Word similarity could, instead, help to find relations among concepts that are not detected by the studied computational lexicons.

In the current prototype, for sake of simplicity, we did not process composite tags, which are present in the ArsMeteo corpus (e.g. *ultima_spiaggia*, see Fig. 3). Such composite tags can be phrases or multi-word expressions that, taken together as a collocation, can convey an emotional meaning. In order to deal with emotion detection in such cases, it will be necessary to rely not only on ontologies and lexical resources but also on a deep linguistic text analysis [6].

Moreover, it is known that the emotional semantics may vary depending on the context. Psychological theories concerning emotions, that tie perception to context, could be integrated in ArsEmotica to refine the outcome [12].

Finally, future possible uses include the development of emotion-aware search engines and of emotional tag clouds. This would open the way to a plethora of applications, including iOS and Android apps, not only with a cultural flavor (along the lines of the application in the previous section) but also more intrinsically related to leisure.

7. Acknowledgements

The authors thank all the persons who supported the work with comments and contributions. In particular, Andrea Bolioli and *CELI s.r.l.*, Flavio Portis, Giorgio Vaccarino and the *Associazione Culturale ArsMeteo* for providing the data corpus, Fondazione Bruno Kessler and ISTI CNR for supplying the lexicons, the web agency *EasyBit*, which supplied the tagging platform of Arsmeteo, the NIL group at the Complutense University of Madrid, which supplied *OntoEmotion*.

References

- [1] Edoardo Acotto, Matteo Baldoni, Cristina Baroglio, Viviana Patti, Flavio Portis, and Giorgio Vaccarino. Arsmeteo: artworks and tags floating over the planet art. In *Proc. of the 20th ACM conference on Hypertext and hypermedia*, HT '09, pages 331–332. ACM, 2009.
- [2] Matteo Baldoni, Cristina Baroglio, András Horváth, Viviana Patti, Flavio Portis, Maurizio Avilia, and Pierluigi Grillo. Folksonomies meet ontologies in arsmeteo: from social descriptions of artifacts to emotional concepts. In S. Borgo and L. Lesmo, editors, *Formal Ontologies Meet Industry, FOMI 2008*, pages 132–143. IOS Press, 2008.

- [3] Matteo Baldoni, Cristina Baroglio, Viviana Patti, and Claudio Schifanella. Sentiment Analysis in the Planet Art: a Case Study in the Social Semantic Web. Technical Report 1/2012, Dipartimento di Informatica, Università degli Studi di Torino, 2012. <http://www.di.unito.it/argo/papers/RT0112.pdf>.
- [4] Sebastian Chan. Tagging and searching: serendipity and museum collection databases. In *Museums and the Web 2007*, pages 87–99, 2007.
- [5] Roddy Cowie, Catherine Pelachaud, and Paolo Petta, editors. *Emotion-Oriented Systems*. Cognitive Technologies. Springer Berlin Heidelberg, 2011.
- [6] Rodolfo Delmonte and Vincenzo Pallotta. Opinion Mining and Sentiment Analysis Need Text Understanding. In Vincenzo Pallotta, Alessandro Soro, and Eloisa Vargiu, editors, *Advances in Distributed Agent-Based Retrieval Tools*, volume 361 of *Studies in Computational Intelligence*, pages 81–95. Springer Berlin / Heidelberg, 2011.
- [7] Fefie Dotsika. Uniting formal and informal descriptive power: Reconciling ontologies with folksonomies. *International Journal of Information Management*, pages 407–415, October 2009.
- [8] Andrea Esuli, Stefano Baccianella, and Fabrizio Sebastiani. SentiWordNet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining. In *Proc. of the 7th conference on International Language Resources and Evaluation (LREC'10)*. ELRA, May 2010.
- [9] Virginia Francisco, Pablo Gervás, and Federico Peinado. Ontological reasoning to configure emotional voice synthesis. In *Proc. of Web Reasoning and Rule Systems, RR 2007*, volume 4524 of *LNCS*, pages 88–102. Springer, 2007.
- [10] Virginia Francisco, Federico Peinado, Raquel Hervás, and Pablo Gervás. *Semantic Web Approaches to the Extraction and Representation of Emotions in Texts*, chapter 4, pages 127–168. NOVA Publishers, 2010.
- [11] Deborah L. McGuinness and Frank van Harmelen. OWL Web Ontology Language Overview. W3C Recommendation, 2004. <http://www.w3.org/TR/2004/REC-owl-features-20040210/>.
- [12] John Richard Ogorek. Normative picture categorization: Defining affective space in response to pictorial stimuli. In *Proc. of REU'05*, 2005.
- [13] Bo Pang and Lillian Lee. *Opinion Mining and Sentiment Analysis (Foundations and Trends(R) in Information Retrieval)*. Now Publishers Inc., 2008.
- [14] W. Parrot. *Emotions in Social Psychology*. Psychology Press, Philadelphia, 2001.
- [15] Emanuele Pianta, Luisa Bentivogli, and Christian Girardi. Multiwordnet: developing an aligned multilingual database. In *Proc. of the First International Conference on Global WordNet*, January 2002.
- [16] Rosalind W. Picard. Affective computing. Technical Report 321, MIT, 1995.
- [17] Marc Schröder, Paolo Baggia, Felix Burkhardt, Alessandro Oltramari, Catherine Pelachaud, Christian Peter, and Enrico Zovato. Emotion Markup Language (EmotionML) 1.0. W3C Working Draft, 2010. <http://www.w3.org/TR/emotionml/>.
- [18] Nigel Shadbolt, Tim Berners-Lee, and Wendy Hall. The semantic web revisited. *IEEE Intelligent Systems*, 21(3):96–101, 2006.
- [19] Katharina Siorpaes and Martin Hepp. Games with a purpose for the semantic web. *IEEE Intelligent Systems*, 23:50–60, 2008.
- [20] Carlo Strapparava and Alessandro Valitutti. WordNet-Affect: an affective extension of WordNet. In *Proc. of LREC*, volume 4, pages 1083–1086, 2004.
- [21] Carlo Strapparava, Alessandro Valitutti, and Oliviero Stock. The affective weight of lexicon. In *Proc. of LREC*, pages 1–83, 2006.
- [22] Jennifer Trant and Bruce Wyman. Investigating social tagging and folksonomy in art museums with *steve.museum*. In *Proc. of the Collaborative Web Tagging Workshop (WWW'06)*, 2006.
- [23] Luis von Ahn. Games with a purpose. *IEEE Computer*, 39(6):92–94, 2006.