

# Sentiment Analysis in the Planet Art: a Case Study in the Social Semantic Web

Matteo Baldoni, Cristina Baroglio, Viviana Patti and Claudio Schifanella

**Abstract** Affective computing is receiving increasing attention in many sectors, ranging from advertisement to politics. Its application to the Planet Art, however, is quite at its beginning, especially for what concerns the visual arts. This work, set in a Social Semantic Web framework, explores the possibility of extracting rich, emotional semantic information from the tags freely associated to digitalized visual artworks, identifying the prevalent emotions that are captured by the tags. This is done by means of ArsEmotica, an application software that we developed and that combines an ontology of emotional concepts with available computational and sentiment lexicons. Besides having made possible the enrichment of the ontology with over four-hundred Italian terms, ArsEmotica is able to analyse the emotional semantics of a tagged artwork by working at different levels: not only it can compute a semantic value, captured by tags that can be directly associated to emotional concepts, but it can also compute the semantic value of tags that can be ascribed to emotional concepts only indirectly. The results of a user study, aimed at validating the outcomes of ArsEmotica, are reported and commented. They were obtained by involving the users of the same community which tagged the artworks. It is important to observe that the tagging activity was not performed with the aim of later applying some kind of Sentiment Analysis, but in a pure Web 2.0 approach, i.e. as a form of spontaneous annotation produced by the members of the community on one another's artworks.

---

Matteo Baldoni

Dipartimento di Informatica, Università degli Studi di Torino e-mail: baldoni@di.unito.it

Cristina Baroglio

Dipartimento di Informatica, Università degli Studi di Torino e-mail: baroglio@di.unito.it

Viviana Patti

Dipartimento di Informatica, Università degli Studi di Torino e-mail: patti@di.unito.it

Claudio Schifanella

Dipartimento di Informatica, Università degli Studi di Torino e-mail: schi@di.unito.it

## 1 Introduction

Seeing artworks exposed at a vernissage or a gallery has a strong emotional impact on the visitors of the exhibition. Artistic arousal of emotions has been a traditional matter of speculation in aesthetics from Aristotle onwards. Even though the matter intrinsically stimulates challenging issues, the most recent studies in analytic philosophy seem to converge on considering emotions in the aesthetic experience as an essential dimension of cognition, in other words as a means of discerning what properties an artwork has and expresses [13]. For what concerns the Art World, recently, emotions and sentiment monitoring have been the core of some important project, conceived from both *pioneering artists*, see, for instance, the ongoing series on the mechanics of emotions by the contemporary new-media artist Maurice Benayoun<sup>1</sup>, and from *cultural institutions* and *museums*. Think, for example, to the Swiss national five-years long research project “e-motion: Mapping Museum Experience” [29], where the psycho-geographical effect of the museum on the visitor were investigated and visitors’ art reception as well as their psychological reaction were monitored by means of techniques for visitor tracking and of biometric measurements.

Moreover, it is widely recognized that technology has the power to transform visitors’ experiences at museums. In particular, the Social Web and Web 2.0 technologies can play an important role in promoting the visitors participation, with a focus on networking individual visitor experiences [23]. Many museums and artistic associations opened their collections for access on the web and have studied the potential of social tagging [28, 6, 4]. According to the Social Web principles, users can be involved in the production of contents or in their elaboration, e.g. by publishing and organizing own materials, by posting comments for discussing published contents, by participating into wikis, by rating resources. Social networks and platforms promote the participation of users in many ways, by 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, emotions and feelings*.

One of the emerging research fields, targeted at extracting information from the data that is supplied by the Social Web users, is *emotion-oriented computing* (a.k.a. *Affective Computing* [20]), whose focus is to automatically recognize the users’ emotions by analyzing their tagging or writing behavior. In particular the rise of social media has fueled interest in *Sentiment Analysis* (and *Opinion Mining*) [17, 8]: 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 and has found interesting applications in the business world.

What we propose in this paper is to study how to fruitfully apply Sentiment Analysis to the Planet Art, by exploiting, as an information source, the tags that the visitors leave for commenting artworks. Tags can be left both in the case the

---

<sup>1</sup> <http://www.benayoun.com/>

collection is accessible only through the web and in the case an exhibition supplies to its visitors appropriate tools for playing with the (physical, in this case) artworks.

In order to elicit emotional meanings we developed *ArsEmotica*, a prototype application, that extracts a shared emotional semantic from the floating individual perceptions and reactions of the visitors of a collection. *ArsEmotica* creates a semantic social space where artworks can be dynamically organized according to an ontology of emotions based on the one in [12], which includes the following base emotions: Anger, Fear, Happiness, Sadness and Surprise. It can be interfaced with a resource sharing and tagging system, which provides the data that is to be processed by the so called *emotional engine*. The emotional engine task is to automatically identify the emotions which better capture the affective meaning, that visitors collectively (though indirectly, through the tags as we will see) gave to the artworks.

We validated *ArsEmotica* by means of a *user study* which involved the community of Italian users of the *ArsMeteo* art portal (<http://www.arsmeteo.org> [1]). *ArsMeteo* is a Web 2.0 platform (on-line since 2006), that enables the collection and the presentation of digital (or digitalized) artworks and performances (including poems, videos, pictures and musical compositions), together with their tags. The user study focussed on a corpus of about 40 tagged artworks and involved about one hundred users (about 35% of the whole community), the same community which tagged the artworks. Notice that the tagging activity, monitored in *ArsMeteo* since 2006, was not performed with the aim of later applying some kind of Sentiment Analysis, but in a pure Web 2.0 approach, as a form of spontaneous annotation produced by the members of the community on one another's works. Moreover, in *ArsMeteo*, artworks usually have many tags, expressing a variety of meanings, thus supporting the emergence of different emotional potentials. This is consistent with the idea that art can emotionally affect people in different ways. When this happens, the analysis performed by *ArsEmotica* is capable of providing multiple emotional classifications.

The ontology of emotions that we used is based on the one in [12], and was constructed by enriching, with the help of *ArsEmotica*, the latter by adding over four hundred Italian words bearing an affective value. The correlation between tags and emotional concepts was computed by following the approach in [5], exploiting and combining Semantic Web tools and lexical resources. In particular, we relied on multi-lingual and Italian computational lexicons [19] and affective lexicons (WordNet-Affect [25] and SentiWordNet [2]).

Our approach to Sentiment Analysis, applied to the art domain, presents some novelties w.r.t. the literature. First, by relying on an *ontology*, we can produce results that are much richer than the classical polarized appreciation; this is necessary when dealing with the Art World, where reducing emotional evaluations to positive (negative or neutral) classifications is an inadmissible simplification. Second, we focus on a very essential form of textual resource: *collections of tags*. This choice is motivated by the fact that tags are one of the most common ways for annotating resources, and that by concentrating 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.

The paper begins with a brief overview of the background with a special focus on the art application domain and on visitor's sentiment. Section 3 describes the ArsMeteo tagging platform. Section 4 briefly presents the current version of the ArsEmotica prototype application. Section 5 reports and discusses the results of the user study. Final remarks end the paper.

## 2 Art, sentiment and social media

Over the last five years, social media, by promoting user participation and peer-to-peer interaction, have transformed the ways that web users interact with content and with each other on the web. This affected and influenced also the "Planet Art". Many curators, cultural organizations, institutions, and museums started to explore the ways social technologies and principles of Web 2.0 can be applied in museums or exhibitions, with the main aim of encouraging an active discourse among visitors, and of fostering social awareness and reflections about collections of artworks. Experiments, like the one described in [7], proved that people value connections with other persons, both for their own sake and as a way for supporting a meaningful experience.

Along this line, some important museums (among which the American museums: Guggenheim, Metropolitan Museum of Art, and San Francisco Museum of Modern Art) opened their collections for access through the web [28, 6], and studied the potential of social tagging in the development of better interfaces, by supporting the *steve.museum* research project (2006-2008) [28]. Tagging, due to its highly subjective nature, is perceived by museums as a valuable feedback that reveals the way in which their public perceives collections, exhibitions, artworks.

The curator Nina Simon in the recent book "The Participatory Museum" [23] goes one step beyond by investigating the important role that artworks can play as "social objects", i.e. as a basis for an object-centered sociality [21], and in the foundation of a new generation of social applications. The key aim is to stimulate participation and encourage visitors to share their experiences with one another. Artworks, indeed, are a good example of "social objects" because they are objects that connect people, for instance by fostering conversation. They do so because they raise personal feelings and personal degrees of understanding and of acceptance. As a consequence, people can attach to them stories that are taken from their own experience. Artworks occupy a space, either physical or a piece of display, and this allows people to create a contact with them. They can be provocative.

The kind of work we are doing with ArsEmotica, on the collection provided by the ArsMeteo portal, finds its setting this picture but the emotional dimension moves to the foreground. In the aesthetic experience, emotions are an important element of the cognition. The work of art is apprehended through the feelings as well through the senses [13]. ArsEmotica enhances the social value of the artworks by exploring the possibility of capturing their emotional value for a group of people, namely the persons who tagged the digitalized reproductions of the artworks in the collection.

The tagging platform can be seen as a mediator artifact, allowing the ArsMeteo users to interact with the social objects. ArsEmotica, the tool that we developed for analysing the tags and for extracting the prevalent emotions associated to the artworks, can be seen as an interpretative tool which produces new information by combining the individual contributions of the users.

Having the emotional dimension in the foreground can be a key element for encouraging visitors of collections to share their experiences about the artworks as social objects. The outcome of the collective experience can be presented to the society of users into several forms. For instance, in the case of digitalized collections, by allowing the construction of personalized emotion-driven visits, while in the case of physical exhibitions by supplying an emotional map of the rooms. Whatever the tool used to display the outcomes of the emotional analysis, this will be the result of a co-creation process which involved all of the users who tagged the artworks.

This way of interacting with artworks is, to the best of our knowledge, original. In fact, if it is possible to find studies on mood-driven music selection [10] (see also applications like StereoMood and SourceTone), little has been done for visual artworks, like images, paintings, and photographs. Social platforms like Flickr consider artworks as social objects, placing them in the middle of the user's activity, but the way in which users interact with photographs has nothing to do with the enhancement of the emotional dimension: users can leave comments to a picture or annotations directly inside the picture, and they can use tags, but this huge amount of information is currently not used to convey an emotional meaning. In this respect, one partly exceptional application is the flash website Moodstream [15], where the user chooses a mood and the system provides appropriate images and a soundtrack from a corpus of pre-classified data.

A high interest in monitoring the sentiment of the visitors in environments like museums and exhibitions is witnessed by projects like *e-motion* [29]<sup>2</sup>, where the psycho-geographical effect of the museum on the museum visitors is under investigation by means of methods for visitor tracking and biometric measurements. Our current work with ArsEmotica attacks a sibling issue, i.e. to monitor the visitor sentiment by relying on semantic *social web technologies* (ontologies, folksonomies, tagging), and proposes an approach based on ontology-driven sentiment analysis. To the best of our knowledge, the proposal to apply Sentiment Analysis to the Planet Art is original.

Before describing our approach, let us briefly overview the background concerning the techniques and topics involved in the ArsEmotica proposal. 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 social web 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, especially in art portals like ArsMeteo, or in artistic domains, where resources represent artworks, movies, music, books, they are also

---

<sup>2</sup> Mapping Museum Experience: <http://www.mapping-museum-experience.com/en>

used to express reception, opinions, feelings, and users seem to use tags to supply concise reviews.

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 some surveys by [9, 27], together with a previous work of the authors [4], 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 [22].

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 and Sentiment Analysis [17, 8], of ontologies of emotions, like the one we started out work from [11], and of W3C markup language proposals [3]. However, there are still few applications that use the most advanced results in Semantic Web technology to deal with emotions [14] and most of the approaches use ontologies where emotions are individual, isolated units (e.g. WordAffect). Such considerations motivated our focus on the ontology of emotions in [11], an OWL ontology where emotions are structured and organized in levels, trying to integrate the results of the most recent psychological models. 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 is new w.r.t. previous work on Sentiment Analysis and allows to extract from tags affective information which is richer than a polarized appreciation.

### 3 Arsmeteo

*ArsMeteo* (<http://www.arsmeteo.org>, see Figure 1 for a snapshot) is an art portal for sharing artworks and their emerging, connective meanings. It is inspired by an idea of the artist *Giorgio Vaccarino* and opens with a planetary vision of the Earth flown over by evolving clouds of words and images. The main aim of the project is to create a space for a community of artists and art lovers, where artists can find a dynamic, interactive and fertile background for *artistic experimentation* and *cooperative artistic creation*. Without any intermediation, artists can interact on the

same artwork, simultaneously: thus, different ages, backgrounds, languages, ways of life can merge and generate a new level of composite culture. The web platform combines social tagging and tag-based browsing technology with functionalities for collecting, accessing and presenting works of art together with their meanings. It enables the collection of digital (or digitalized) artworks and performances, belonging to a variety of artistic forms including poems, videos, pictures and musical compositions. Meanings are given by the tagging activity of the community. All contents are accessible as “digital commons”.



**Fig. 1** The ArsMeteo portal entrance.

Artists and visitors may express their own reception of the artworks by tagging them or by clicking on the plus and minus symbols next to tags to change tag weights, which will affect the ranking of search results. Currently, the portal gathers a collection of about 9936 artifacts produced by 302 artists, it receives about 300 hits per day, and it has collected almost 37000 tags (an average of 13 tags per artwork). The tagging activity of the community provides a basis for browsing the artworks in traditional ways as well as according to different navigation metaphors, respectively called *Art Time Machine*, *Argonaut* and *Serendipity*:

- *The Art Time Machine*: the whole ArsMeteo archive, including forum posts and events, can be accessed by year;
- *Serendipity*: users can find resources *without* performing an explicit and systematic search. When a user chooses a graphical artifact, the starting point of the search, the system puts a preview of such resource at the center of the search page. Then, the artwork is encircled by the previews of 24 graphical resources randomly chosen among all those which are tag-related with the initial choice. The user can then continue to browse by clicking on one of the surrounding resources: the serendipity game restarts by putting the new resource at the center of the page;

- *Argonaut* The idea is to visualize the evocative power of (poetic) sentences and text. The user starts by writing a sentence on the proper text area. Artworks tagged by words occurring in the sentence drop down from the top of the page.

This dynamic environment naturally engenders constructive contexts, where users can invent new forms of trans-individual artistic actions and reflections. In the last year we already observed the emerging of *artistic group actions*. New artifacts have been created by collecting, recycling and recontextualizing fragments of artworks from different users and from different media or instances. Different users synchronized their creative potential and generated new artworks and languages building upon (or reinventing) contents: new poems were created as rivers of words by the simultaneous and interactive tagging activity of many users; a new visual alphabet (*Alfameteo*) emerged by reinterpreting some artworks as alphabetic letters; in “The table of Niépce” series, authors started to play, again by using digital technologies, an artistic game of collective painting, which was invented in 1979.

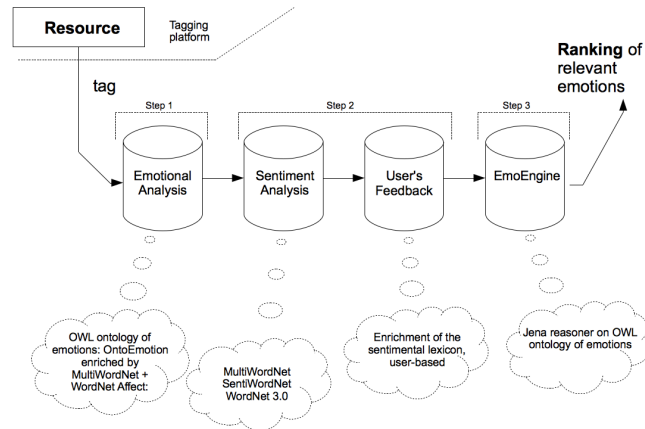
For what concerns the *community aspects*, Arsmeteo started in a national context: most of the users are Italian, contemporary artists. Since the community is growing very fast, in order to give it the time to get acquainted with the new contents and give feedbacks by tagging (keep the descriptivity index high), we implemented the *SlowMeteo policy*: artists cannot upload more than 3 new resources per day. The risk of too many uploads is that interesting artworks can be unfairly neglected. For many of the ArsMeteo authors the portal was a first appealing opportunity for accessing and exploiting the new social potential of web-based technologies. Some of them entered in this new world thanks to the help of other more technologically skillful users, who play the role of *digital curators*.

## 4 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 2 reports the three main steps that characterize the computation after the pre-processing:

- **Checking tags against the ontology of emotions.** This step checks whether a tag belongs to the ontology of emotions. Tags belonging to the ontology are immediately classified as “emotional”.
- **Checking tags with SentiWordNet.** 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. The latter will be the only ones presented to the user in order to get a feedback on which emotional concept they deliver.
- **Ranking of Emotions.** Based on data collected in the previous steps, the tool ranks the emotions associated by the users to the resource.





**Fig. 2** ArsEmotica overall architecture.

The following sections explain how the extraction of an emotional semantics is performed. We get into deeper details only for the first step of the algorithm, because the experimental results that we report and comment are aimed at validating the outcomes of this phase. Notice that for allowing the application even of the first step of the algorithm, we had to enrich the reference ontology of emotions with a suitable Italian vocabulary. The way we did this is also explained below.

#### 4.1 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. Our starting point was *OntoEmotion*, an emotional ontology developed at Universidad Complutense de Madrid [11], 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 and structures emotional categories in a taxonomy, which includes 87 emotional concepts.

The basic emotions are *Sadness*, *Happiness*, *Surprise*, *Fear* and *Anger* and the taxonomic structure basically refers to the psychological model by Parrot [18], adapted to these five basic emotions, and integrated with emotions which appear in other well-established models.

*OntoEmotions* has been conceived for categorizing emotion-denoting words. Classes corresponding to the emotional concepts were originally populated by about 250 instances, consisting of emotion-denoting words both from English and from 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 emotion-denoting words, i.e. those words which each language provides for denoting emotions. Originally it had two subclasses: *EnglishWord* and *SpanishWord*. Each instance of these two latter concepts has two parents: one is a concept from the *Emotion* hierarchy (the class of emotion denoted by the word), while the other is a concept from the *Word* hierarchy (e.g. the language the word belongs to). 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 the 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 emotions 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 synset, representing the “senses” of that word, which are labeled by MultiWordNet 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 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. The resulting ontology contains more than 450 Italian words referring to the 87 emotional categories of OntoEmotion. In order to keep track in the ontology of the synonymy among words belonging to a same synset, we have defined the OWL object property *hasSynonym*.

ArsEmotica uses the enhanced ontology for checking if a tag describing a resource *directly* refers to some emotional category (Emotional Analysis). If yes, 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  $s$ . In general, as we will see in the remainder of the section, the range of the score  $s$  is  $[0, 100]$  but when a tag is an instance of an emotional concept, its strength will be 100. So, for example, since the word “af-

fanno” (breathlessness) is an instance of “anxiety”, the corresponding triple will be: (“affanno”, “anxiety”, 100). Triples are store in a data base table.

## 4.2 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. We only sketch, here, the way in which we identify the tags which indirectly bear an affective meaning. The interested reader can find details in [5].

As observed in [26], 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 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  $[0.0,1.0]$  and their sum is 1.0 for each synset. Since SentiWordNet was created for the English 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.

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 some strong sentimental (positive or negative) meaning. Different senses of the same term can have different opinion-related properties and different scores. When for no sense of a given term has a significant sentimental score, we conclude that it is mainly descriptive and usually does not evoke emotions. Therefore, we ask the evaluation of the community only for those terms having at least one meaning with a relevant sentimental score. This was done to have no false negative.

For all tags resulting potentially affective, users will be free to associate to the word one or more emotions from the emotional categories of the ontology with a strength value, which intuitively represents the user’s measure of the semantic affinity of the term with the chosen emotional category. Again a set of triples  $(t, e, s)$  will be collected during this phase and stored in the data base.

### 4.3 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. Also this step is described in greater details in [5].

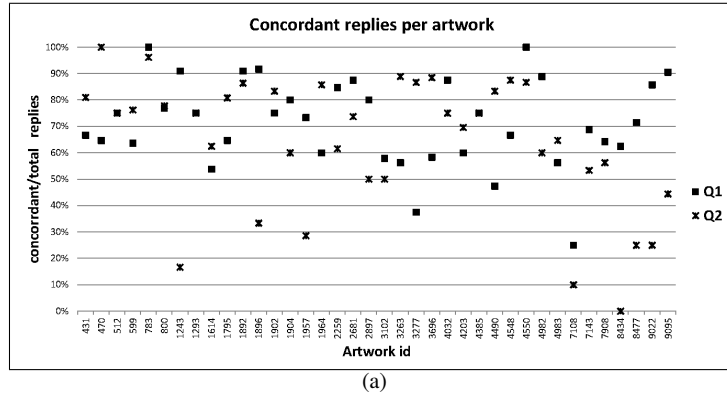
the ranking is performed with the help of the Jena Reasoner, applied to the triples collected during the previous analysis steps. The implemented algorithm relies on the taxonomic structure of the ontology and is inspired to the one in [12], where an analysis is performed to emotionally mark up a sentence by analyzing the words that compose it. Intuitively, the algorithm allows the selection of the *most specific emotion* which represents the affective information related to the artwork from the probability (score) that each of its tag has of indicating different emotions. The basic steps are: (1) processing the emotional concepts appearing in the triples  $(t, e, s)$ , so as to identify also those emotions in the ontology that are related of the ones appearing in the triples. The identified emotional concepts can be organized into *layers* by following the parent-child relationship (in this phase a Jena Reasoner has been applied to the collected triples); (2) starting from the leaves and moving upward towards the root, compose and propagate the scores.

## 5 User Study and Evaluation

In order to test the effectiveness of the proposed system, 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 selected a set  $A = \{a_1, \dots, a_{38}\}$  of 38 artworks<sup>3</sup> from the ArsMeteo corpus, randomly selected among those which are the most voted by the community, explicitly including 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  $a_i \in A$ , the set  $E_{a_i} = \{e_{a_i}^1, \dots, e_{a_i}^n\}$  of associated emotions, by analysing the tags attached to images by the ArsMeteo 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 (see Section 4).

Users were asked to answer a questionnaire composed of ten questions, belonging to two categories. In the first kind of question ( $Q1$ ), starting from a randomly selected artwork  $a_i \in A$ , users were asked to choose the most representative emotion out of a set of two elements: one of the emotions that ArsEmotica identified as associated to  $a_i$  and one that resulted as not being associated to  $a_i$ . Formally, the user had to choose between  $\{e_1, e_2\}$ , where  $e_1 \in E_{a_i}$  and  $e_2 \in (E - E_{a_i})$ . In the second kind of question ( $Q2$ ), given a randomly selected emotion  $e_j \in E$ , users were asked

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



Replies per artwork		
	Average concordant	Variance
Q1	71.43%	2.76%
Q2	64.04%	6.44%

(b)


**Fig. 3** (a) The graph shows on the X-axis all artwork identifiers, on the Y-axis percentages of replies that are concordant to ArsEmotica outcomes; (b) The table summarizes the results achieved, by reporting the averages and variances of percentages for both kinds of questions.

Overall Replies			
	#concordant	#total	%concordant
Q1	352	495	71.11%
Q2	347	495	70.10%

**Table 1** The table reports the total number of concordant replies, the total number of answers (both without dividing them per artwork) and the percentages of the former w.r.t. the latter.

to choose the most representative artwork between a set of two elements  $\{a_1, a_2\}$ , that were selected in such a way that  $e_j \in E_{a_1}$  and  $e_j \notin E_{a_2}$ .

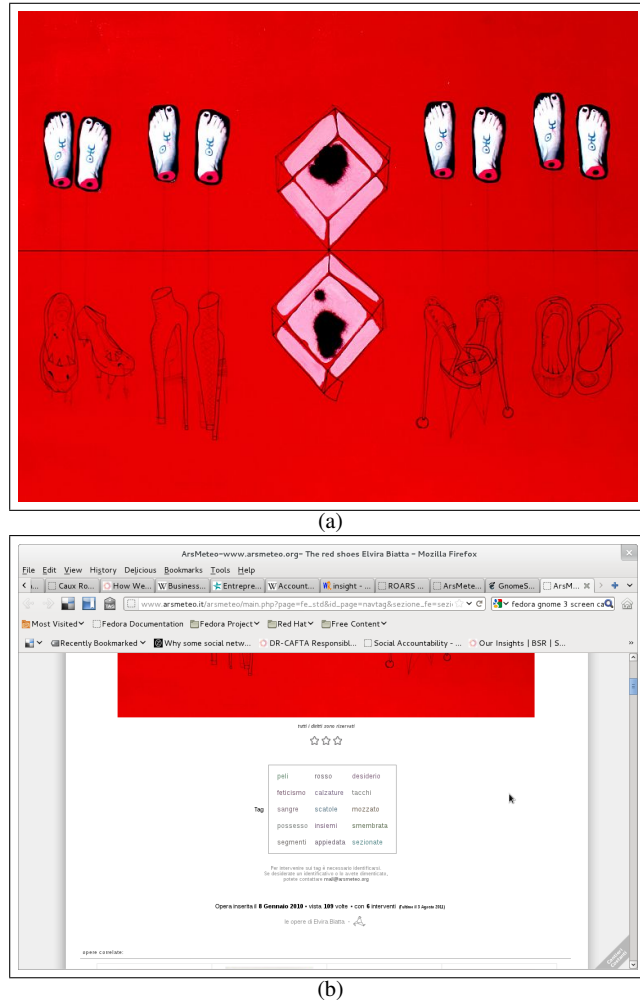
Among all, 99 ArsMeteo users replied to our call (about the 35%): each of them replied five randomly generated quiz both for question  $Q1$  and for question  $Q2$ . For each type of question, we first analyzed the users' replies by measuring, for each artwork in  $A$ , the quantity of associations  $(a_i, e_j), a_i \in A, e_j \in E$  between the artwork and emotions that are concordant to the output of ArsEmotica system w.r.t to whole set of replies. As can be seen in Figure 3(a), the majority of the artworks in the corpus shows a correlation (corresponding to a high percentage of concordant replies) between the emotions chosen by the users and those extracted by the ArsEmotica system. This behaviour can be observed both for questions  $Q1$  and  $Q2$ , with values up to 100% for artworks 783, 4550, and 470. Figure 3(b) summarizes the results into a table, reporting the average value and the variance of all percentages of concordant replies obtained for each artwork. For what concerns question  $Q1$ , there is an average agreement with ArsEmotica up to about 71%, with a variance of 2.76%; the average for question  $Q2$  decreases to 64% with a higher variance (over 6%), show-



**Fig. 4** The charts reports the behavior of the concordant replies w.r.t. (a) the total number of answers and (b) the number of emotions associated by ArsEmotica, for each artworks and kind of question.

ing in any case a correlation between the user replies and the output of ArsEmotica. Table 1 reports the results without grouping the answers artwork per artwork. In this case, while for question  $Q1$  the overall percentage of concordant replies is similar to what previously showed, this percentage increases to 70.11% for question  $Q2$ : this improvement is due to the fact that for this kind of question, poor results (for concordant replies) depend more on artworks, for which we gathered a low number of answers – indeed, due to random choice not all artworks were presented to users with the same frequency. This kind of artworks increase their weights when the overall percentage of concordant replies is computed by starting from artworks grouping. Figure 4(a) shows more details by positioning each artwork on a plane where the  $X$ -axis represents the number of replies it received and the  $Y$ -axis reports the percentage of concordant answers. A low number of answers necessarily turns out into a greater variance. Finally, in Figure 4(b) we analysed the relationship between the number of concordant replies and the number of emotions associated to an artwork by ArsEmotica: as one can see, in general, even though there is no sharp correlation, it is possible to notice that when the number of emotions increases, the number of artworks with low values of concordant replies decreases.

Some considerations it is interesting to draw, also for explaining variances, concern the sharpness of attribution of an artwork to one precise emotional concept, starting from tags. In general, one would think that when a tool like ArsEmotica identified one single emotion out of tags associated to an artwork, there will be not much surprise coming out from the users' evaluation. In some cases, however, it may happen (as for image 8434, see Figure 5) that an image which does not bear an emotional value is associated to some emotion because one of its tags has a meaning that can be related to an emotion. The fact that this correlation is almost accidental emerges from the users' evaluation: indeed, no user felt the emotion that ArsEmotica concluded image 8434 to convey, see Figure 3(a), even though in the repository there is a tag attached to it by a person, which is an emotion: desiderio (i.e. desire). This is an extreme case and, as one can see from the graph, there are not many other similar cases. However, they affect the results. Interestingly, by do-



**Fig. 5** (a) Artwork 8434: The Red Shoes, by Elvira Biatta; (b) The tags associated to artwork 8434.

ing so, they give insight on the fact that in order to distinguish between accidental (or weak) emotional correlations and sharp but effective emotional correlations, it may be useful to introduce in the ArsEmotica algorithm a notion of strength of the emotion-bearing tags. We mean to explore this kind of enhancement in future work.

Another general observation is that it is quite unlikely for an artwork to be emotion-bearing for everybody and, when it is, to be related by everybody to a same single emotion. It is much more frequent the case in which there is a sharp classification towards a little set of emotions. This is coherent with the fact that tags are produced by a community, each of whose users relates to the artwork in a totally personal way, based on his/her personal background and story. However, as the

results witness, it is possible to develop tools, like ArsEmotica, which are able to capture the main trends.

## 6 Conclusion and Future Work

With this work we have shown that it is possible to extract rich, emotional semantic information from the tags associated to a visual resource. The effort needed to obtain this outcome involves the combination of approaches, that were developed in different, though neighbouring, scientific disciplines, as well as of tools and technologies (e.g. lexicons and libraries) that are already available. The tool that we developed, ArsEmotica, succeeded in populating an ontology of emotions (based on [11]) with Italian emotion-bearing terms. Moreover, the first experiments proved that it is capable of extracting the prevalent emotions that are hidden inside the tags, associated to a digitalized visual artwork, a kind of response that is much more informative than the general polarized appreciation, returned by standard sentiment analysis tools. The analysis also gave interesting insights on images to which a single emotion seems to be associated, and images which seem to divide the visitors in groups with different perceptions.

We think that the proposed approach may be particularly suitable to those application domains where tags can be interpreted as *concise reviews* (e.g. artworks, books, movies). 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 *steve.museum*) 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. 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 proposed solution can be refined in many ways. For what concerns the pre-processing of tags, we intend to improve the current prototype by applying stemming 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 algorithms 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.

Moreover, the experiments that we carried on so far, only concern the first step of the ArsEmotica algorithm. Further extensive validation is to be performed also for the part described in Section 4.2. In this concern, the delicate point is the identification of a simple and effective way for motivating the users of ArsMeteo to annotate tags having an indirect affective meaning by means of emotional concepts, taken from the ontology. In order to face this issue, one promising direction could be to rely on the Game With A Purpose [30] 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 kind of solution would be along the line of



recent approaches, which face the challenge of increasing the user involvement in building the Semantic Web [24]. An alternative could be to integrate in ArsEmotica the use of automatic techniques, e.g. the one proposed in [4], for identifying the association of terms having an emotional value (that is recognized by the sentiment analysis step) with the proper ontological concepts.

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 [16].

For what concerns the possible uses of ArsEmotica, we think that, by devising proper interpretative graphical representations for presenting the outcomes of the elaboration of the tags associated to artworks, it would be a precious co-creation instrument for museums and virtual galleries, along the direction traced in [23]. In this sense, it would help transforming classical art-fruition experiences into innovative, more immersive experiences, with a greater impact on visitors. Finally, the capability of extracting prevalent emotions foster the development of emotion-aware search engines and of emotional tag clouds. This would open the way to a plethora of applications, including smart-phone apps, not only with a cultural flavor (along the lines of the application in the previous section) but also more intrinsically related to leisure.

**Acknowledgements** The authors thank all the persons who supported the work with comments and contributions. In particular, Paolo Rena, for the generous work on the ArsEmotica software development, Andrea Bolioli and *CELI s.r.l.* for the fruitful discussions on lexical and sentiment resources, 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 NIL group at the Complutense University of Madrid, which supplied *OntoEmotion*.

## References

1. E. Acotto, M. Baldoni, C. Baroglio, V. Patti, F. Portis, and G. 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. A. E. S. Baccianella and F. 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.
3. P. Baggia, F. Burkhardt, A. Oltramari, C. Pelachaud, C. Peter, and E. Zovato. Emotion markup language (emotionml) 1.0. <http://www.w3.org/TR/emotionml/i>.
4. M. Baldoni, C. Baroglio, A. Horváth, V. Patti, F. Portis, M. Avilia, and P. 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.
5. M. Baldoni, C. Baroglio, V. Patti, and P. Rena. From Tags to Emotions: Ontology-driven Sentiment Analysis in the Social Semantic Web. *Intelligenza Artificiale: the International Journal of the AI\*IA*, 2012. To appear.
6. S. Chan. Tagging and searching: serendipity and museum collection databases. In *Museums and the Web 2007*, pages 87–99, 2007.

7. D. Cosley, J. Lewenstein, A. Herman, J. Holloway, J. Baxter, S. Nomura, K. Boehner, and G. Gay. ArtLinks: fostering social awareness and reflection in museums. In *CHI '08: Proc. of the 26th annual SIGCHI conference on Human factors in computing systems*, pages 403–412, New York, NY, USA, 2008. ACM.
8. R. Delmonte and V. Pallotta. Opinion mining and sentiment analysis need text understanding. In V. Pallotta, A. Soro, and E. Vargiu, editors, *Advances in Distributed Agent-Based Retrieval Tools*, volume 361 of *Studies in Computational Intelligence*, pages 81–95. Springer Berlin / Heidelberg, 2011.
9. F. Dotsika. Uniting formal and informal descriptive power: Reconciling ontologies with folksonomies. *International Journal of Information Management*, pages 407–415, October 2009.
10. P. Dunker, S. Nowak, A. Begau, and C. Lanz. Content-based mood classification for photos and music: a generic multi-modal classification framework and evaluation approach. In *Proceedings of the 1st ACM international conference on Multimedia information retrieval*, MIR '08, pages 97–104, New York, NY, USA, 2008. ACM.
11. V. Francisco, P. Gervás, and F. 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.
12. V. Francisco, F. Peinado, R. Hervás, and P. Gervás. *Semantic Web Approaches to the Extraction and Representation of Emotions in Texts*. NOVA Publishers, 2010.
13. N. Goodman. *Languages of art: an approach to a theory of symbols*. Hackett, 1976.
14. M. Grassi. Developing heo human emotions ontology. In J. Fierrez-Aguilar, J. Ortega-Garcia, A. Esposito, A. Drygajlo, and M. Faúndez-Zanuy, editors, *COST 2101/2102 Conference*, volume 5707 of *Lecture Notes in Computer Science*, pages 244–251. Springer, 2009.
15. G. Images. Moodstream. 2010. <http://moodstream.gettyimages.com/>.
16. J. R. Ogorek. Normative picture categorization: Defining affective space in response to pictorial stimuli. In *Proc. of REU'05*, 2005.
17. B. Pang and L. Lee. *Opinion Mining and Sentiment Analysis (Foundations and Trends(R) in Information Retrieval)*. Now Publishers Inc., 2008.
18. W. Parrot. *Emotions in Social Psychology*. Psychology Press, Philadelphia, 2001.
19. E. Pianta, L. Bentivogli, and C. Girardi. Multiwordnet: developing an aligned multilingual database. In *Proc. of the First International Conference on Global WordNet*, January 2002.
20. R. W. Picard. Affective computing. Technical Report 321, MIT, 1995.
21. T. Schatzki, K. Knorr-Cetina, and E. Savigny. *The practice turn in contemporary theory*. Routledge, 2001.
22. N. Shadbolt, T. Berners-Lee, and W. Hall. The semantic web revisited. *IEEE Intelligent Systems*, 21(3):96–101, 2006.
23. N. Simon. *Participatory Museum*. Museum 2.0, 2010.
24. K. Siorpaes and M. Hepp. Games with a purpose for the semantic web. *IEEE Intelligent Systems*, 23:50–60, 2008.
25. C. Strapparava and A. Valitutti. WordNet-Affect: an affective extension of WordNet. In *Proc. of LREC*, volume 4, pages 1083–1086, 2004.
26. C. Strapparava, A. Valitutti, and O. Stock. The affective weight of lexicon. In *Proc. of LREC*, pages 1–83, 2006.
27. J. Trant. Studying Social Tagging and Folksonomy: A Review and Framework. *Journal of Digital Information*, 10(1), 2009.
28. J. Trant and B. Wyman. Investigating social tagging and folksonomy in art museums with steve.museum. In *Proc. of the Collaborative Web Tagging Workshop (WWW'06)*, 2006.
29. W. Tschacher, S. Greenwood, V. Kirchberg, S. Wintzerith, K. van den Berg, and M. Tróndle. Physiological correlates of aesthetic perception in a museum. *Psychology of Aesthetics, Creativity, and the Arts*, 2011.
30. L. von Ahn. Games with a purpose. *IEEE Computer*, 39(6):92–94, 2006.