Emotional Responses to Artworks in Online Collections

Federico Bertola and Viviana Patti

Università degli Studi di Torino
Dipartimento di Informatica
c.so Svizzera 185, I-10149, Torino, Italy
bertola@celi.it, patti@di.unito.it

Abstract. Artworks have a strong emotional impact on the visitors of an exhibition. Many museums and associations opened their collections for access on the web and have studied the potential of social tagging. User data collected by tagging in art social platforms are a precious information source about emotional responses to artworks and a feedback about the way in which users perceives collections. In this paper, we present our new achievements on this topic within the ArsEmotica framework. Our focus is on eliciting sharable emotional meanings from online collection visitors tags, by interactively involving the users of the virtual communities in the process of capturing the latent emotions behind the tags. We rely on methods and tools from a set of disciplines ranging from Semantic and Social Web to NLP. Such disciplines provide the building blocks for creating a semantic social space, where artworks can be dynamically organized according to a new ontology of emotions inspired by the well-known Plutchik's model of human emotions. The final aim is to involve users in the creation of such emotional space and, then, to offer them an emotion-driven access to the artworks.

Keywords: ontologies, social tagging, affective computing, online collections

1 Introduction

One of the emerging research fields, targeted at extracting information from the data supplied by the Social Web users, is emotion-oriented computing (a.k.a. affective computing), whose focus is to automatically recognize the users' emotions by analyzing their tagging or writing behavior, their conversations or their facial expressions [16]. In particular, the rise of social media has fueled interest in Sentiment Analysis and Opinion Mining: since emotions are often related to appreciation, knowing the feelings of the users towards target issues is important feedback that can support many decisional tasks and has found interesting applications in the business world. Affective computing is receiving increasing attention in many sectors. Its application to the Planet Art, however, is quite at its beginning [9, 5, 6]. In the last years, many museums and cultural heritage

institutions opened their collections and archives for access on the web (see e.g. the American Guggenheim, Metropolitan Museum of Art, and San Francisco Museum of Modern Art, the European Rijksmuseum or, more recently, the Google Art project[1]), and the potential of social tagging, in this context, has been explored [19, 11, 5]. User data collected by art social platforms are a precious information source about trends and emotions. Collecting such emotional responses to artworks and collections can be important not only for artists and curators, but also for policymakers in the cultural heritage sector, that need advanced e-participation tools for being supported in their work, both at the decision-making stage, and in the ex-post evaluation of the impact of their policies (e.g. What is the sentiment of citizens about a publicly funded exhibition?).

In this paper we study how to apply Sentiment Analysis to the Planet Art, by exploiting, as information source, tags intended as textual traces that the visitors leave for commenting artworks on social platforms. In particular, we present our new achievements within the ArsEmotica framework [5, 6], where the focus is on methods to extract, from the floating individual meanings and reactions of the visitors of a collection, a *shared emotional semantics*. In previous work [5, 6], we have presented, and evaluated by user study, application software that analyzes tagged artworks from a social tagging platform and provides as output, for each artwork, a set of related emotions. Such emotions emerge as the most significant ones for capturing the users emotional responses toward that resource.

ArsEmotica creates a semantic social space where artworks can be dynamically organized according to an ontology of emotions. The sentiment analysis approach is, indeed, *ontology-driven*. Given a tagged resource, the correlation between tags and emotions is computed by referring to an ontology of emotional categories: from unstructured texts to concepts of an ontology. This is done by exploiting and combining Semantic Web tools and lexical resources. Intuitively, first, tags directly referring to ontological concepts are identified; then, users can give emotional feedback on potentially affective tags, by using the ontology; the final emotional output for the given resource is calculated based on identified emotional concepts, by exploiting automated reasoning on ontology relationships. Notice that there is a strong accent on the social dimension of the process: the key aim is to identify the emotions which better capture the affective meaning, that visitors *collectively* give to the artworks.

In the following, we will describe a new version of the ArsEmotica prototype (Arsemotica 2.0), and we will report about the most significant advancements. Our efforts were mainly devoted to improve the application framework, by devising proper graphical representations of the computed emotional outcomes, in order to effectively involve the user in the process of extracting the latent emotional semantics from tags. In other words, the goal was to setup a platform were the outcomes of our ontology-driven sentiment analysis could be *presented* in a graphical, intuitive way to users, and possibly *enriched* by a further user emotional feedback. To this aim we have designed a new ontology of emotions, which refers to the Robert Plutchik's circumplex model [15]. The ontology pro-

[1] http://www.googleartproject.com/.

vided a guidance in the design a new interface for ArsEmotica, where emotional responses to artworks are represented by means of a graphical representation inspired to the Plutchik's *emotion wheel*, a bi-dimensional representation of the model. Summarizing, the strong points of the new version are: a) the implementation in OWL of a new ontology of emotions inspired by the Plutchik's psychoevolutionary theory of emotion, and its use in the computation of the emotional outcomes; b) the design of a new interactive user interface inspired by the model underlying the ontology of emotions; c) an extended and generalized architecture for the ArsEmotica application; d) the possibility to delivery the output of the emotional analysis in a machine-understandable format, compliant to emerging standards. The paper is organized as follows. Section 2 contains a brief overview of related work. Section 3 presents ArsEmotica 2.0, with a special focus on the new ontology of emotions. Section 4 presents the new ArsEmotica's interactive user interface. Final remarks end the paper.

2 Related Work

A high interest in monitoring the sentiment of the visitors in environments like museums and exhibitions is recently raised among art practitioners, curators and cultural heritage stakeholders, as witnessed for instance by projects like *e-motion*[2], where the experience of museum-goers are analyzed by means of visitor tracking and biometric measurements technologies. Moreover, many curators, cultural organizations, and museums explored the ways social technologies and principles of Web 2.0 can be applied in museums or exhibitions to foster social awareness and reflections about artworks. In this spirit, many cultural heritage institutions opened their collections for access on the web. Someone went one step beyond [17] by investigating the important role that artworks can play as "social objects", i.e. as a basis for an object-centered sociality, and in the foundation of a new generation of social applications, where the key aim is to stimulate participation and encourage visitors to share their experiences, both in case of virtual and physical collections [9]. Artworks, are a good example of social objects: they are objects that connect people, for instance by fostering conversation. They raise personal feelings and personal degrees of understanding and of acceptance. People can attach to them stories that are taken from their own experience. Recent works such as the visitor's guide in [10] or storytelling-based mobile museum guides [12, 7] can be viewed as contributions in this direction, especially [7] and [10], which pose an accent on social aspects, moving from the consideration that most people visit museums in *groups* [2].

The kind of work we are doing with ArsEmotica [5, 6] can be settled in this picture, but the emotional dimension moves to the foreground. ArsEmotica analyses tags and extracts the prevalent emotions associated to the artworks. It 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

[2] Mapping Museum Experience: `http://www.mapping-museum-experience.com/en`

to share their experiences about the artworks as social objects. The outcome of the collective experience can provide a means for integrating virtual and physical collections. A recent experiment in this direction is the project "A map of emotions #palazzomadama" for the Robert Wilson's video-portraits exhibition at Palazzo Madama Museum (Turin, Italy). Here visitors were invited to tag artwork's photos, taken during their visit of the physical collection, with the emotions evoked (e.g. #joy, #fear...), by relying on the use of Instagram mobile app running on their smart phones. The Followgram web application was, then, used as sharing platform for collection's photos and affective tags, with the final aim to offer new ways to visit Wilson's collection to online visitors.

Other relevant contributions on art and emotions can be found in a wide range of disciplines, ranging from art and aesthetics to psychology. In particular, recent investigations on psychology of art, which recognizes a specific role to emotions in the aesthetic experience, provide a cognitive grounding to the ArsEmotica perspective and aims. The role of the senses in the arts, both for what concerns creation and reception, has been studied in [3] by taking a multidisciplinary perspective, where computational approaches to emotion recognition can play a key role. Specifically, on this line, recently an interesting experiment has been conducted [20], aimed at analyzing abstract paintings and investigating why a specific painting is perceived as emotional, by combining the use of machine learning algorithms for automatic recognition and eye tracking technology. The approach has been tested on a dataset of abstract paintings (MART museum, Rovereto, Italy), which have been classified as being positive and negative. These results give us the opportunity to observe that, when dealing with the Art World, since affective information evoked by artworks is usually richer than a polarized appreciation, reducing emotional evaluations to positive (or negative) classifications is an over-simplification. On this respect, in ArsEmotica the reference to a rich cognitive-based ontological model of emotions allows us to classify artworks according to a congruous emotional space.

3 Arsemotica 2.0: a New Prototype

In this section, we describe the new version of ArsEmotica, the application software that we developed for testing our ideas. We will briefly recall the main characteristics of the application, by focussing on the original aspects and advancements with respect to previous work [5, 6]. In particular, we will present here: a new extended and generalized architecture (Sec. 3.1); a new OWL ontology, which refers to a state-of-the-art cognitive model of emotions (Sec. 3.2), and inspires the new interactive user interface discussed in Section 4.

3.1 Architecture

Figure 1 reports the four main steps that characterize the computation in ArsEmotica. They are briefly described in the following. The application can be interfaced with any resource sharing and tagging system which provides the data

to be processed, i.e. digital artworks and their tags. Social tagging platform for art collections as steve.museum or armeteo.org, with active communities that frequently visit and comment online the collections, would be ideal data sources.

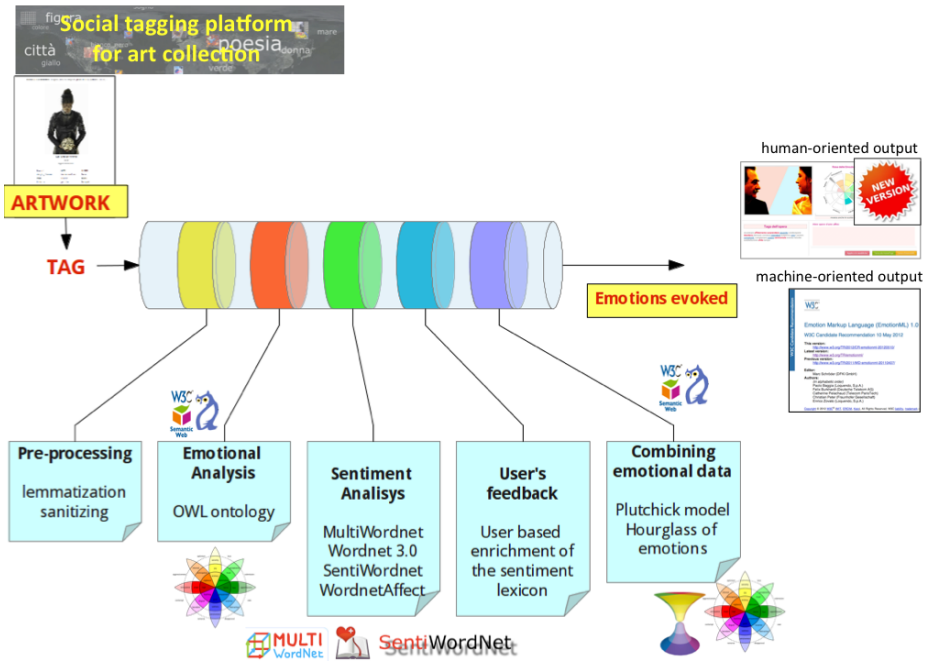

Fig. 1. ArsEmotica overall architecture.

- **Phase 1. Pre-processing: Lemmatization and string sanitizing.** In this step tags associated to a given artworks are filtered so as to eliminate flaws like spelling mistakes, badly accented characters, and so forth. Then, tags are converted into lemmas by applying a lemmatization algorithm, which builds upon *Morph-It!*, a corpus-based morphological resource for the Italian language.
- **Phase 2. Checking tags against the ontology of emotions.** This step checks whether a tag belongs to the ontology of emotions. In other words, it checks if the tags of a given resource are "emotion-denoting" words directly referring to some emotional categories of the ontology. Tags belonging to the ontology are immediately classified as "emotional".
- **Phase 3. Checking tags with SentiWordNet.** Tags that do not correspond to terms in the ontology are further analyzed by means of *SentiWordNet* [4], 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. The feedback is collected thanks to the interactive user interface described in Sec. 4, which has been designed in tune with the new ontological model of emotion presented below.

– **Phase 4. Combining emotional data and producing a set of emotions as output.** Based on data collected in the previous steps, the so-called emotional engine of the tool offers as output a set of emotions associated to the resource. We have implemented a new algorithm for accomplishing this task, where emotions collected in the previous steps are not simply ranked as in [5] but compared and combined. The algorithm implemented compare collected emotions, by exploiting ontological reasoning on the taxonomic structure of the new ontology of emotions. Moreover, it combines them by referring to the Hourglass Model [8], a reinterpretation of the Plutchik's model, where primary emotions are further organized around four independent but concomitant dimensions (*Pleasantness*, *Attention*, *Sensitivity* and *Aptitude*), whose different levels of activation can give birth to very wide space of different emotions. Shortly, in this model different emotions (basic or compound), result from different combinations of activation levels for the four dimensions. Dimensions are characterized by six levels of activation, which determine the intensity of the expressed/perceived emotion as a float $\in [-1, +1]$. This allows to classify affective information both in a categorical way (according to a number of emotion categories) and in a dimensional format (which facilitates comparison and aggregation), and provided us a powerful inspiration in implementing a new algorithm for combining emotional data in a final output.
  As an advancement w.r.t the previous version, the resulting output can be produced in different modalities. For what concerns human users, emotions evoked by artworks are visualized by a sort of *emotion wheel*, graphically inspired to the color wheel used by Plutchik for offering a bi-dimensional representation of his circumplex model of emotions [14] (Sec. 4). Moreover, in order to foster real interoperability and integration between ArsEmotica and other semantic applications, the system encodes the output in a machine-readable format, by using W3C standards such as RDF and EmotionML, which is currently an emerging standard for emotion annotation.

In the following, we introduce the ontology, which drives the emotional analysis and plays a crucial role in all core steps of the ArsEmotica computation.

3.2 A New Ontology of Emotions based on a Cognitive Model

We have designed for ArsEmotica 2.0 a new ontology of emotional categories based on Plutchik's circumplex model [15, 14], a well-founded psychological model of emotions. The ontology is written in OWL, structures emotional categories in a taxonomy, which includes 32 emotional concepts. Due to its role within the ArsEmotica architecture, the ontology has been conceived for categorizing

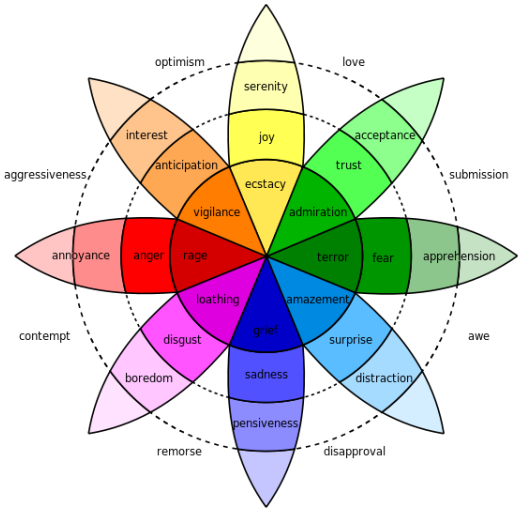

Fig. 2. Plutchik's circumplex model [15].

emotion-denoting words, as the one used in the previous version of the application. It, then, includes two root concepts: *Emotion* and *Word*.

Class Emotion For what concerns the class *Emotion*, the design of the emotional categories taxonomic structure, of the disjunction axioms and of the object and data properties mirrors the main features of Plutchik's circumplex model, (see the two-dimensional representation in Fig 2). In particular the *Emotion*'s hierarchy includes all the 32 emotional categories presented as distinguished labels in the model. Such model is represented as a *wheel of emotions*, which encodes the following elements and concepts:

- **Basic or primary emotions**: *joy*, *trust*, *fear*, *surprise*, *sadness*, *disgust*, *anger*, *anticipation* (i.e. *expectancy*); in the color wheel this is represented by differently colored sectors.
- **Opposites**: basic emotions can be conceptualized in terms of polar opposites: *joy* versus *sadness*, *anger* versus *fear*, *trust* versus *disgust*, *surprise* versus *anticipation*.
- **Intensity**: each emotion can exist in varying degrees of intensity; in the wheel this is represented by the vertical dimension.
- **Similarity**: emotions vary in their degree of similarity to one another; in the wheel this is represented by the radial dimension.
- **Complex emotions**: beside basic emotions, there are complex emotions, that are a mixtures of the primary emotions, just as some colors are primary, and others made by mixing the primary colors; in the model in Fig 2 emotions in the blank spaces are compositions of basic emotions called *primary dyads*.

Class Word For what concerns the class *Word*, it is the root for the emotion-denoting words, i.e. those words which each language provides for denoting emotions. Since we actually applied our application to use cases where tagging involved Italian communities, we actually defined and populated the subclass *ItalianWord*. However, the ontology is already designed to be extended with further subclasses of *Word*, for representing emotion-denoting words in different languages. Intuitively, each instance of the *Word* and *Emotion* concepts has two parents: one is a concept from the *Emotion* hierarchy (the emotion denoted by the word, e.g. *rage*), while the other is a concept from the *Word* hierarchy (e.g. Italian, the language the word belongs to). For instance, the following code excerpt corresponds to the description of the Italian affective word *rabbia*: it is both an instance of the concept *Rage* (an intense anger), and an instance of the concept *ItalianWord*, i.e. *rabbia* is an Italian word for denoting rage:

```
<rdf:Description rdf:about="http://www.arsemotica.unito.it/ontology/emotions.owl#rabbia">
  <rdf:type rdf:resource="http://www.w3.org/2002/07/owl#NamedIndividual"/>
  <rdf:type rdf:resource="http://www.arsemotica.unito.it/ontology/emotions.owl#ItalianWord"/>
  <rdf:type rdf:resource="http://www.arsemotica.unito.it/ontology/emotions.owl#Rage"/>
</rdf:Description>
```

Ontology Population We semi-automatically populated the ontology with Italian words by following the same methodology described in [5] for populating OntoEmotion. Shortly, we relied on the multilingual lexical database MultiWordNet [13] 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. We have manually chosen an initial set of representative Italian emotional words for each concept. At the beginning we have chosen a set of nouns. Then, by making use of morpho-semantic relations between nouns and verbs or between nouns and adjectives specified in Italian dictionaries, we have expanded the initial set including new related verbs and adjectives[3]. 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 [18]: 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. Currently, the resulting ontology contains about 500 Italian words referring to the 32 emotional categories of the ontology.

[3] Unfortunately, this process has been carried on mainly manually: available lexical resources do not provide the information on such kind of relations for Italian.

4 An Interactive User Interface for Collecting and Presenting Emotional Responses

Our choice to design an ad hoc ontology for ArsEmotica was driven, on the one hand, by the need to find a interpretative graphical representations for presenting the outcomes of the automatic elaboration of the artworks' tags. On the other hand, we were looking for a simple interface to foster the users to annotate tags having an indirect affective meaning, by means of emotional concepts from the ontology. On this perspective, we have found the Plutchik's model very attractive for three main reasons:

- The reference to a graphical wheel is very intuitive and offers a *spacial representation of emotions* and their different relations (similarities, intensities, polar oppositions). Such kind of representation allows to convey to the user a rich information on the emotional model, without referring to tree-like visualization of the ontology hierarchy. For instance, looking at the wheel, it is natural to catch a similarity relation among *joy* and *trust*, while *joy* and *sadness* are polar opposites, and *love* emotion is composed of *joy* and *trust*.
- The use of *colors* for denoting different emotions provides a very intuitive communication code. Different color nuances for different emotions, transmit naturally the idea that primary emotions can blend to form a variety of compound emotions, analogous to the way colors combine to generate different color graduations. This aspect can play an important role in the development of an user interface for a cultural heritage application, and can be exploited for proposing to the user intuitive metaphors to browse the emotional space. In particular, it inspired us in the design of an interface to easy the task of emotionally classifying tags.
- The number of emotional categories distinguished in the wheel is *limited*. This aspect facilitates the user that is involved in an emotional evaluation.

Let us now present the interactive user interface that we have developed, and its sample application on a tagged artwork from the ArsMeteo online collection [1] (http://www.arsmeteo.org). The sequence of interactions offered to the user follows the flux of computation sketched in Fig 1.

An Emotional Wheel for Presenting Emotions Evoked by the Artworks (Phase 2). After the user selects an artwork from the collection, the application applies the emotional analysis on the artwork tags. The result of this computation, i.e. the *evoked emotions*, is presented to the user by a graphical representation called "La rosa delle emozioni", which strongly recalls the Plutchik's color wheel. For instance, by applying the emotional analysis to the artwork "Dove la Raffinata Ragazza Bionda guarda il Grosso Toro Malmorto" by Filippo Valente (Fig 3), the four red colored tags are identified as emotional according to the emotional ontology: 'orrore', 'infamia', 'cattiveria', 'tristezza'; the presence of emotional responses related to *sadness* and a strong disgust

Fig. 3. Showing the results of the automatic emotional analysis

(*loathing*) is highlighted by coloring the sectors of the emotion wheel corresponding to those emotions. Internal sectors of the ArsEmotica's wheel are intended as representing light intensity of emotions, while the external ones as representing high intensity. Underlined blue colored tags denotes tags that have been recognized by the sentiment analysis stage as possibly conveying some affective meaning. Then, they appear as active links for the user's emotional feedback: see e.g. 'sangue', 'sconfiggere', and so on.

User's Emotional Feedback (Phase 3). The Arsemotica's wheel offers an effective way for fostering users to annotate tags having an indirect affective meaning by means of emotional concepts from the ontology. After the user selects a tag to evaluate, the application activates a pop-up window, where an uncolored emotional wheel is shown. Users can express the emotional evaluation in terms of basic emotions with different intensities, and color the wheel accordingly, by clicking on one of the 24 sectors of the wheel; otherwise they can select compound emotions, by selecting the wedge-shaped triangles inserted between the basic emotions. In the example in Figure 4 the user associated to the tag 'sangue' (blood) the emotions *fear* and *disgust* (with high intensity, which corresponds to *loathing*). Notice that the tag evaluation is contextual to the vision of the artwork, which indeed remains visible in the background.

Final Emotional Evaluation of the Artwork (Phase 4). After combining the emotional data of phases 2 and 3, the resulting emotional evaluation is again presented to the user by using the ArsEmotica's wheel. If requested, the application can format the result in the standard markup language EmotionML.

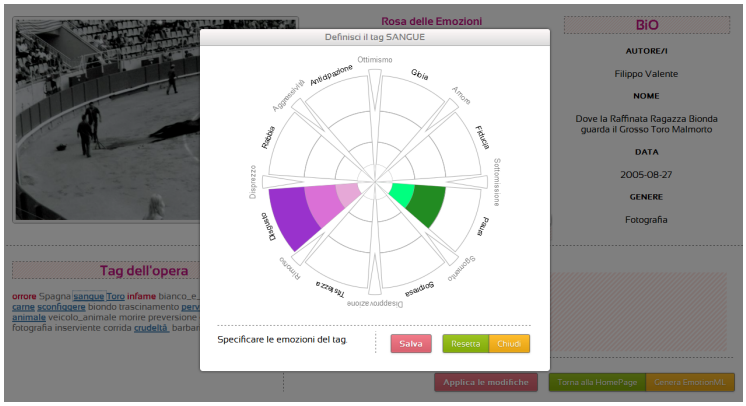

Fig. 4. Interaction with the user: collecting the tag-mediated emotional feedback

5 Conclusion and Future Work

In this paper we have presented our new achievements within the ArsEmotica framework. Our efforts were devoted to improve the application framework, by devising proper interpretative graphical representations to present the outcomes of the ArsEmotica analysis to the users. We described our new interface, that presents emotional responses to artworks by means of an emotion wheel inspired to the Plutchik's one. In general, designing engaging interfaces that allow an appropriate granularity of expression is not a trivial task. Therefore, the next step will be to evaluate the new prototype, and to carry on a user test. We plan to use tagged artworks from the ArsMeteo tagging platform as dataset and to involve users of the ArsMeteo community, which in the past have already actively participated to a user study on the first version of our prototype [6], and manifested great interest on the topic.

For what concerns the possible uses of ArsEmotica, we think that it can be exploited as a co-creation instrument for museums and virtual galleries. Moreover, the capability of extracting prevalent emotions can foster the development of emotion-aware search engines, emotional tag clouds or interactive map of emotions, which could enable new ways of exploring art collections.

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 ACM HT '09*, ACM:331–332, 2009.
2. Liliana Ardissono, Tsvi Kuflik, and Daniela Petrelli. Personalization in cultural heritage: the road travelled and the one ahead. *User Model. User-Adapt. Interact.*, 22(1-2):73–99, 2012.
3. F. Bacci and D. Melcher. *Art and the Senses.* OUP Oxford, 2011.

4. Andrea Esuli Stefano Baccianella and Fabrizio Sebastiani. SentiWordNet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proc. of LREC'10*. ELRA, May 2010.
5. Matteo Baldoni, Cristina Baroglio, Viviana Patti, and Paolo Rena. From tags to emotions: Ontology-driven sentiment analysis in the social semantic web. *Intelligenza Artificiale*, 6(1):41–54, 2012.
6. Matteo Baldoni, Cristina Baroglio, Viviana Patti, and Claudio Schifanella. Sentiment analysis in the planet art: A case study in the social semantic web. In Cristian Lai, Giovanni Semeraro, and Eloisa Vargiu, editors, *New Challenges in Distributed Information Filtering and Retrieval*, volume 439 of *Studies in Computational Intelligence*, pages 131–149. Springer, 2013.
7. Charles Callaway, Oliviero Stock, Elyon Dekoven, Kinneret Noy, Yael Citron, and Yael Dobrin. Mobile drama in an instrumented museum: inducing group conversation via coordinated narratives. *New Review of Hypermedia and Multimedia*, 18(1-2):37–61, 2012.
8. Erik Cambria, Andrew Livingstone, and Amir Hussain. The hourglass of emotions. In Anna Esposito, Antonietta Maria Esposito, Alessandro Vinciarelli, Rüdiger Hoffmann, and Vincent C. Müller, editors, *COST 2102 Training School, Revised Selected Papers*, volume 7403 of *Lecture Notes in Computer Science*. Springer, 2012.
9. Gunho Chae, S. Joon Park, Robert Stein, Jungwha Kim, and Susan Wiedenbeck. Exploring affective computing for enhancing the museum experience with online collections. In *Proc. of the Museums and the Web*, 2012.
10. Tsvi Kuflik, Oliviero Stock, Massimo Zancanaro, Ariel Gorfinkel, Sadek Jbara, Shahar Kats, Julia Sheidin, and Nadav Kashtan. A visitor's guide in an active museum: Presentations, communications, and reflection. *J. Comput. Cult. Herit.*, 3(3):1–25, February 2011.
11. Yi-Ling Lin and Lora Aroyo. Interactive curating of user tags for audiovisual archives. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*, AVI '12, pages 685–688, New York, NY, USA, 2012. ACM.
12. Vincenzo Lombardo and Rossana Damiano. Storytelling on mobile devices for cultural heritage. *New Review of Hypermedia and Multimedia*, 18(1-2):11–35, 2012.
13. Emanuele Pianta, Luisa Bentivogli, and C. Girardi. Multiwordnet: developing an aligned multilingual database. In *Proc. of Int. Conf. on Global WordNet*, 2002.
14. Robert Plutchik. The circumplex as a general model of the structure of emotions and personality. In R. Plutchik and H. R. Conte, editors, *Circumplex models of personality and emotions*, pages 17–47. American Psychological Association, 1997.
15. Robert Plutchik. The Nature of Emotions. *American Scientist*, 89(4), 2001.
16. Marc Schroeder, Hannes Pirker, Myriam Lamolle, Felix Burkhardt, Christian Peter, and Enrico Zovato. Representing emotions and related states in technological systems. In Roddy Cowie, Catherine Pelachaud, and Paolo Petta, editors, *Emotion-Oriented Systems*, Cognitive Technologies, pages 369–387. Springer, 2011.
17. Nina Simon. *Participatory Museum*. Museum 2.0, 2010.
18. C. Strapparava and A. Valitutti. WordNet-Affect: an affective extension of WordNet. In *Proc. of LREC'04*, volume 4, pages 1083–1086, 2004.
19. 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.
20. Victoria Yanulevskaya, Jasper Uijlings, Elia Bruni, Andreza Sartori, Elisa Zamboni, Francesca Bacci, David Melcher, and Nicu Sebe. In the eye of the beholder: employing statistical analysis and eye tracking for analyzing abstract paintings. In *Proc. of the ACM Int. Conf. on Multimedia*, MM '12, pages 349–358. ACM, 2012.