

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Ontology-based affective models to organize artworks in the social semantic web

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1535222> since 2017-06-17T23:11:33Z

Published version:

DOI:doi:10.1016/j.ipm.2015.10.003

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Ontology-based Affective Models to Organize Artworks in the Social Semantic Web

Federico Bertola, Viviana Patti*

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

Abstract

In this paper, we focus on applying sentiment analysis to resources from online art collections, by exploiting, as information source, tags intended as textual traces that visitors leave to comment artworks on social platforms. We present a framework where methods and tools from a set of disciplines, ranging from Semantic and Social Web to Natural Language Processing, provide us the building blocks for creating a semantic social space to organize artworks according to an ontology of emotions. The ontology is inspired by the Plutchik's circumplex model, a well-founded psychological model of human emotions. Users can be involved in the creation of the emotional space, through a graphical interactive interface. The development of such semantic space enables new ways of accessing and exploring art collections.

The affective categorization model and the emotion detection output are encoded into W3C ontology languages. This gives us the twofold advantage to enable tractable reasoning on detected emotions and related artworks, and to foster the interoperability and integration of tools developed in the Semantic Web and Linked Data community. The proposal has been evaluated against a real-word case study, a dataset of tagged multimedia artworks from the ArsMeteo Italian online collection, and validated through a user study.

Keywords: ontologies, emotion visualization, sentiment analysis, social tagging, semantic web, linked open data

1. Introduction

The development of the Web and the advent of social media has brought about new paradigms of interactions that foster first-person engagement and crowdsourcing content creation. The subjective and expressive dimensions move

*Corresponding author. Tel: +39 011 6706804
Email addresses: bertola@celi.it (Federico Bertola), patti@di.unito.it (Viviana Patti)

to the foreground, opening the way to the emergence of an affective component within a dynamic corpus of digitized contents, which advocates new techniques for automatic processing and retrieval of the available affective information. Therefore, recently a high interest raised among researchers in developing approaches and tools for sentiment analysis and emotion detection, aimed at automatically analyzing and processing of the affective information conveyed by social media (Schroeder et al., 2011; Cambria et al., 2013). In addition, the need to support users in accessing and exploring the outcomes of the emotion detection and sentiment analysis algorithms has fueled interest on research of solutions that address the sentiment summarization and visualization problem.

Organization and manipulation of social media contents, for categorization, browsing, or visualization purposes, often requires to encompass a semantic model of their affective qualities or of their reception by the users. In particular, we claim that ontologies and cognitive models of emotions can play a key role to bring advancements in this area (Cambria & Hussain, 2012). These can be defined compliant with emerging semantic web standards (which enable automated reasoning and semantic metadata processing) and integrated into traditional sentiment analysis and emotion detection techniques, with the final aim to enhance organization and access to contents.

In this paper, we address the above issues within the context of emotional aspects of Cultural Heritage. Artworks have a strong emotional impact on the visitors of an exhibition. In the last years, many cultural heritage institutions opened their collections to web access (think for instance of the Google Art project¹). User data collected by art social platforms is a precious information source about trends and emotions. However, although sentiment analysis and emotion detection of user data collected by social platforms are receiving increasing attention in many sectors, and, in spite of the fact that a high interest in monitoring the sentiment of the visitors is raised among art practitioners, curators and cultural heritage stakeholders, application of such techniques to the Cultural Heritage and Art domain is quite at its beginning (Chae et al., 2012; Baldoni et al., 2012; Bertola & Patti, 2013).

Our proposal is to elicit latent emotions behind user tags in order to recognize the emotional impact of artworks on people, relying and properly extending the ArsEmotica framework we started to develop in the very last years. The challenge is to study how to classify the sentiments of artworks in online collections by exploiting, on the one hand, tags from social media as information source, and on the other hand a structured knowledge of elicited affective information, such as affective categorization models expressed by ontologies.

The sentiment analysis of the social activity of the community (tagging the online resources) will provide an input for the “emotional engine” we are going to develop, and a basis for an emotion-driven access and browsing of the artworks. Detected emotions are meant to be the ones which better capture the affective meaning that visitors, *collectively*, give to the artworks.

¹<http://www.googleartproject.com/>

Our approach to the sentiment analysis task is *ontology-driven*. To give a short description, given a tagged resource, the relation between tags and emotions is computed by referring to an ontology of emotions and relying on the combined use of Semantic Web technologies, NLP and lexical resources. We have developed an ontology of emotional categories based on Plutchik’s circumplex model (Plutchik, 2001), a well-founded psychological model of human emotions, which has been recently exploited also in emotion analysis from a computational linguistics perspective (Suttles & Ide, 2013; Mohammad & Turney, 2013). The ontology of emotions provides a good taxonomy for classifying artworks and it is so generic that it might also be used for analyzing emotions in running text. Moreover, it inspired an interactive user interface for visualizing and summarizing the results of our emotion detection algorithm. Detected emotional responses to artworks are represented by means of a graphical representation inspired by the Plutchik’s *emotion wheel*. Such representation allows us to convey to the user in a simple way a rich information on the underlying affective model (e.g. relationships among the emotions, such as similarities, intensities, oppositions), without referring to tree-like visualization of the ontology hierarchy. This is an original trait of our proposal. Most ontology-driven information retrieval systems cannot use ontologies this way.

Our final aim is to create a semantic social space where artworks can be dynamically organized, and then accessed according to the ontology of emotions. In this scenario, emotions can serve also as a sort of controlled vocabulary for retrieving artworks.

The proposed approach has been evaluated against a real-word case study, by relying on a dataset of tagged multimedia artworks belonging to a variety of artistic forms including poems, videos, pictures and music from the ArsMeteo Italian art portal (Acotto et al., 2009). Artworks from the ArsMeteo collection, enriched with the semantic metadata about the detected emotions, can be accessed by a SPARQL endpoint. SPARQL queries can be used to explore the collection and to extract information about relationships among emotions, artworks and their genre and author (e.g., “Give me artworks belonging to the genre *Photography* classified as joyful and created by artists living in the Salento’s Italian region”). Moreover, the connection of the ArsMeteo dataset to the linked open data cloud can foster interesting and unexpected possibilities of reusing the data.

The effectiveness of the ArsEmotica application framework for an emotional tagging task has been evaluated by means of a user study, by interviewing via survey a group of human subjects about their experience in interacting with a first release of the ArsEmotica prototype.

Contributions. First, we defined an OWL ontology, which refers to an affective model of emotions well-grounded in psychology. Second, we tuned on this ontology the ArsEmotica framework, which provides an emotional engine for detecting emotions elicited by artworks via social media analysis. Third, we explored the potential of our framework to enable ontology-based, emotion-driven access to the data of an online collection. Semantic representation of knowledge about artworks allows us to enhance access to an online collection both along

the human users dimension and along the machines one, by obtaining two main results: i) enrichment of user’s experience through the proposal of an interactive user interface for emotion-driven access to the artworks, which gives to the user a flavor about the emotional classification of the artwork, in the context of the exploited ontological knowledge model, and the possibility to personally enrich this classification; ii) computer access to the online collection dataset enriched with emotion metadata, by means of Linked Open Data technologies. As a proof of concept we applied the proposed framework to a real-world case study: the dataset of tagged artworks from the ArsMeteo online collection arismeteo.org. In this context, we have developed a unified semantic data model of *artworks*, *artists* and *emotions*, and a demo SPARQL endpoint to query and access such data by combining traditional dimensions for retrieving artworks (i.e. genre, author, etc.) with the emotional axis provided by the ArsEmotica annotation. Finally, the work has been evaluated and validated through a survey with human subjects, which were asked to interact with the ArsEmotica prototype and to answer specific questions, aimed at testing the effectiveness of the application framework for an emotional tagging task from different point of views.

Organization. The paper is organized as follows. The next section contains a brief overview of the ArsEmotica framework. Section 3 focuses on the ontology-based affective model. Section 4 presents the ArsEmotica’s interactive user interface and the potential of semantic processing of the output in a linked data and semantic web perspective. In Section 5 our case-study on the arismeteo.org dataset is discussed. Section 6 presents the user study with human subjects and its results. Section 7 contains a brief overview of related work. Final remarks end the paper.

2. The ArsEmotica Framework

In this section, we describe the characteristics and the main components of ArsEmotica 2.0, the application software that we developed for testing our ideas. Details about a previous version of the application can be found in (Baldoni et al., 2012, 2013).

ArsEmotica is meant as a sort of “emotional engine”, which can be interfaced with any resource sharing and tagging system which provides the data to be processed, i.e., digital artworks and their tags. Figure 1 reports the four main steps that characterize the computation in ArsEmotica. They are briefly described in the following.

1. **Pre-processing: Lemmatization and String sanitizing.** In this step tags associated with a given artwork 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 (Zanchetta & Baroni, 2005).
2. **Checking tags against the ontology of emotions.** This step checks whether a tag belongs to the ontology of emotions, where we encoded



Figure 1: ArsEmotica overall architecture.

knowledge about affective information. 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”.

3. **Checking tags with SentiWordNet.** Tags that have not been recognized as being part of the ontology are processed by *SentiWordNet* (Esuli et al., 2010), in order to select the ones that, even if they do not refer directly to emotions of our ontology, have a significant sentiment load for some of their senses in SentiWordNet. In our view such tags potentially convey affective meaning, and indirectly could refer to emotional categories of the ontology, therefore they are good candidates to be offered to the user for an emotional annotation. For what concerns potentially affective tags, users can give a feedback on which emotional concept they deliver, by choosing from the emotional categories of the ontology one or more emotions to associate to them. So, for instance, if a word like ‘infinito’ (infinite), that has a significant sentiment load for some of its senses, appears as tag for an artwork inspected by ArsEmotica, it will be proposed to the user for a feedback (see (Baldoni et al., 2013) for details on how we calculate the sentiment load for a word, relying on SentiWord-

Net scores). Feedbacks are collected through the interactive user interface described in Sec. 4.1, which has been designed in tune with the ontological model of emotion presented below.

4. **Combining emotional data.** Based on data collected in the previous steps, the tool computes a set of emotions associated to the resource.

We have implemented a combination algorithm, that, on demand, elaborates on this results and allows one to compare and combine emotions collected in the previous steps. This can be useful, in order to have a sort of synthesis which takes into account similarities among the detected emotions encoded in the affective model. We exploited such synthesis to offer human users with a more compact graphical “vision” of the output (see Sec 4.1), but possible uses of such further computation can also be envisioned.

The algorithm compares collected emotions, by exploiting ontological reasoning on the taxonomic structure of the ontology of emotions. Moreover, it combines them by referring to the Hourglass Model (Cambria et al., 2012), a reinterpretation of the Plutchik’s model.

The resulting output can be produced in different modalities, that have been sketched in the right side of Fig 1, and can be divided into two categories: on the one hand we can provide a human-oriented output, summarized and visualized by a graphical user interface, on the other hand we encode the artwork-emotion association in a machine-processable output. Such modalities and their potential will be better described in Section 4.

3. The Ontology-based Affective Model

In this section we describe the ontology, which plays a key role in all steps of the ArsEmotica computation. It is an ontology of emotional categories based on Plutchik’s circumplex model (Plutchik, 2001, 1997), a well-founded psychological model of emotions, and includes also concepts from the Hourglass model in (Cambria et al., 2012). The ontology is written in OWL2. It is available on demand for academic purposes.

3.1. Classes, Hierarchy and Properties

The ontology 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 emotion-denoting words, as the one used in the previous version of the application. It 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 mirror the main features of Plutchik’s circumplex model, (see Fig 2, left side). Such model can be represented as a *wheel of emotions* and encodes the following elements and concepts:




Figure 2: Plutchik's circumplex model: Image from Wikimedia Commons.

- **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 radial direction.
- **Similarity:** emotions vary in their degree of similarity to one another; in the wheel this is represented by the angular direction.
- **Complex emotions:** complex emotions are a mixtures of the primary emotions; in the model in Fig. 2 emotions in the blank spaces are compositions of basic emotions called *primary dyads*.

Emotion is the root for all the emotional concepts. The *Emotion*'s hierarchy includes all the 32 emotional categories presented as distinguished labels in the

model. In particular, the *Emotion* class has two sub-classes: *BasicEmotion* and *ComplexEmotion*. *BasicEmotion* and *ComplexEmotion* are disjoint classes.

Basic emotions of the Plutchik’s model (*Disgust, Trust, Sadness, Joy, Anticipation, Surprise, Anger* and *Fear*) are direct subclasses of *BasicEmotion*. Each of them is specialized again into two subclasses representing the same emotion with weaker or the stronger intensity (e.g. the basic emotion *Joy* has *Ecstasy* and *Serenity* as subclasses). Therefore, we have 24 emotional concepts subsumed by the *BasicEmotion* concept.

Instead, the class *ComplexEmotion* has 8 subclasses, corresponding to the primary dyads² in the Plutchik’s model. Other relations proposed in the Plutchik’s model have been expressed in the ontology by means of the following *object properties*, where *Arch* is the set of the basic emotions and *Comp* the set of the complex emotions:

- **hasOpposite**: ($r : Arch \longrightarrow Arch$), encodes the notion of *polar opposites*;
- **hasSibling**: ($r : Arch \longrightarrow Arch$) encodes the notion of *similarity*;
- **isComposedOf**: ($r : Comp \longrightarrow Arch$) encodes the notion of *composition of basic emotions*.

Moreover, further properties has been introduced in order to link the ontology to the Hourglass Model of Emotions proposed in (Cambria et al., 2012). To give a short description, in this model different emotions (basic or compound), result from different combinations of activation levels for four dimensions (*Pleasantness, Attention, Sensitivity* and *Aptitude*). 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 the algorithm for combining emotional data for human access purposes. In this context, let us mention that the following *data type property* **hasScore**:($r : Arch \longrightarrow \mathbb{R}$) was introduced, to link each emotion with an intensity value mapped into the Hourglass model.

Class Word *Word* is the root for the emotion-denoting words, i.e. those words which each language provides for denoting emotions, in line with related and previous work (Francisco et al., 2010; Baldoni et al., 2012). Since we currently applied our application to use cases where tagging involved Italian communities, we defined and populated the subclass *ItalianWord*³. Intuitively, each instance of the *Word* and *Emotion* concepts, e.g. *felicità* has two parents: one is a concept from the *Emotion* hierarchy (the emotion denoted by the word, e.g.

²We are planning to extend the set of such subclasses in order to represent secondary and tertiary dyads.

³The ontology is already designed to be extended with further subclasses of *Word*, for representing emotion-denoting words in different languages.

Joy), 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>
```

3.2. Individuals and Ontology Population

We populated in a semi-automatic manner the ontology with Italian words following the methodology described in (Baldoni et al., 2012) to populate OntoEmotion, the ontology used in the previous version of the ArsEmotica prototype. We relied on the multilingual lexical database MultiWordNet (Pianta et al., 2002) and its affective domain WordNet-Affect (Strapparava & Valitutti, 2004)⁴, a well-known lexical resource that contains information about the emotions that the words convey, that was developed starting from WordNet. WordNet (Fellbaum, 1998) is a lexical database, in which nouns, verbs, adjectives and adverbs (lemmas) are organized into sets of synonyms (synsets), representing lexical concepts. The WordNet-Affect resource was developed through the selection and labeling of the synsets representing affective concepts. The population process is summarized below. For the sake of clarity, in Fig. 3 a graphical example describing the population pipeline for the emotional concept ‘Terror’ is reported.

The process started by manually selecting a set of representative Italian emotional words, at least one word for each concept. This initial set included less than 90 words classified under 32 emotional concepts, but they were only *nouns*. In order to expand with *adjectives* the set of Italian words that are representative of emotional concepts, we included and classified according to the ontology⁵ the list of 32 emotion terms in (Galati et al., 2008): *addolorato, allegro, angosciato, annoiato, ansioso, arrabbiato, contento, depresso, disgustato, disperato, divertito, entusiasta, euforico, felice, gioioso, imbarazzato, impaurito, indignato, infelice, irritato, malinconico, meravigliato, preoccupato, risentito, sbalordito, scontento, sconvolto, sereno, sorpreso, spaventato, stupito, triste*.

In a second phase we automatically expanded the set of individuals (emotion denoting words) belonging to the emotional concepts by exploiting MultiWordNet and the WordNet-Affect. All manually classified words and adjectives were used as *entry lemmas* for querying the lexical database. The result for each

⁴<http://wndomains.fbk.eu/wnaffect.html>

⁵This process has been carried on manually, by relying on morpho-semantic relations between nouns already classified in the ontology and adjectives specified in the Treccani dictionary (<http://www.treccani.it>).

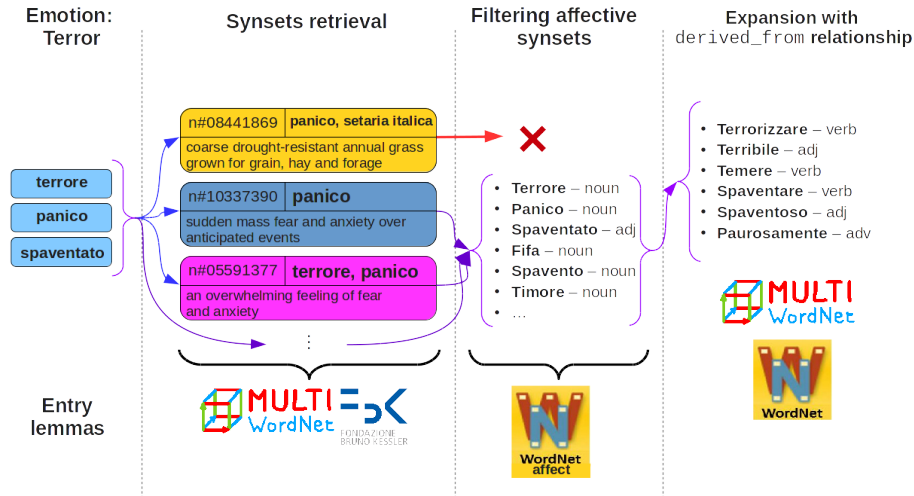


Figure 3: Ontology population process: emotion denoting words for the concept ‘Terror’.

word was a synset, representing the “senses” of that word, labeled by MultiWordNet unique synset identifiers. Each synset was then processed by using WordNet-Affect (Strapparava & Valitutti, 2004): 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 for the same concept of the entry lemmas. In other words, we automatically enriched the ontology with synonyms of the representative emotional words, but also filtered out synsets which do not convey affective information. For instance in Fig. 3, when we query the MultiWordNet database with the Italian word ‘panico’, only two thirds of the resulting synsets are affective (WordNet senses `n#10337390` and `n#05591377`). In particular, the not affective synset refers to the sense of the word ‘panico’ as one of the oldest cultivated cereal grain. Thanks to the affective filter provided by WordNet-Affect we can exclude words belonging to that synset (`Setaria_italica`, `pabbio_coltivato`) when populating the concept `Terror` of our ontology.

As a final step, we further expanded the set of emotion denoting words with further *adjectives*, *verbs* and *adverbs*, by exploiting the WordNet relation `derived-from`, for which we can assume that the affective meaning is preserved. Therefore, all synsets obtained by an application of the `derived-from` relation (and not yet classified in our ontology) were included as individuals of the proper emotional concept (e.g. *terrorizzare*, *terribile*, etc. in the example described in Fig. 3). At the end of the process a human expert checked the identified terms.

The resulting ontology contains about 700 Italian words referring to the 32 emotional categories of the ontology.

4. Processing the Outcomes: Visualization, Access and Exploitation

Semantic representation of the emotionally enriched knowledge about the artworks allows us to enhance access to cultural heritage objects, both along the human users dimension (Sec. 4.1) and along the automatic processing and retrieval one (Sec. 4.2).

4.1. A Graphical User Interface

We have developed an interactive user interface for visualizing and summarizing the results of the emotion detection process on the ArseMeteo dataset. The interface aims at giving to the human user a flavor about the emotional classification of the artwork in the context of the exploited ontological knowledge model. Essential requirements for a graphical interface in our setting were to propose to the user intuitive metaphors, to explore the emotional space and to ease the task of classifying emotionally those tags that potentially have affective meanings, by means of emotional concepts from the ontology.

The Plutchik’s model exploited in ArsEmotica has a natural visual representation as a colored wheel. It provided the inspiration for a visual presentation of the emotion ontology (see Fig. 9), having three main interesting features:

- The reference to a graphical wheel is very intuitive and offers a *spacial representation of emotions* and their different relations (similarities, intensities, polar oppositions). The representation allows to convey to the user rich information on the underlying 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 an 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 ease the task of classifying emotionally tags related to artworks.
- The number of emotional categories distinguished in the wheel is *limited*. This aspect facilitates the user exploring the emotional features of the artwork.

A sample application of the interface on an artwork from the ArseMeteo collection will be commented in 5.3.

4.2. Machine-processable Semantic Affective Information and Linked Open Data

Given an artwork, the detected emotions are also provided as semantic metadata, in a RDF machine-readable format, by relying on W3C standards, and enabling SPARQL querying⁶. Such emotional metadata includes also frequencies of emotion labels extracted from the tags of the artworks, which can be possibly processed by other applications to calculate the strength of the emotions.

As we will see in the concrete context of a real-world case study in the next section, the availability of such semantic metadata supports easy integration of affective information into a semantic data model for artworks (where emotional and standard metadata on cultural heritage objects can be combined) and, therefore, the development of ontology-driven semantic access to art collections in a Linked Open Data perspective. In this context emotions can serve as a controlled vocabulary for retrieving artworks. Moreover, inferences on semantic and interlinked representations can return new interesting relationships among artworks, emotions and their authors.

5. Case study: the arsmeteo.org Dataset

We evaluated our approach against a real-world case study, by relying on a dataset of tagged multimedia artworks from the ArsMeteo Italian art portal (<http://www.arsmeteo.org> (Acotto et al., 2009)). Social tagging platforms for art collections, having active communities that visit and comment online the artworks, are ideal data sources for applying the ArsEmotica analysis. The ArsMeteo collection fits such characteristics.

ArsMeteo is an art portal for sharing artworks and their emerging, connective meanings. Its development is led by a non-profit cultural organization called *Associazione Culturale ArsMeteo* (AMA), based in Turin, Italy. On-line since June 2007, 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. Updated statistics on artists, artworks and tags collected can be found on the portal's homepage. For what concerns the community aspects, ArsMeteo counts over 380.000 visitors. It started in a national context: most of the users are Italian contemporary artists. 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. Both artists and registered visitors may express their own reception of the artworks by tagging. Since the beginning, the community has been very active in tagging artworks, as it will be shown by the analysis in the next section.

⁶<http://www.w3.org/TR/sparql11-overview/>

Artists	Artworks	Tags	MT
282	9171	36970	11.5

Table 1: The *ArsM* dataset: Number of artworks, Number of artists, Number of Tags, Average total number of tags per artwork (MT).




Figure 4: Number of tags per artwork

5.1. The *ArsMeteo* Dataset

For our experiments we rely on the *ArsM* dataset, which includes data and comments about artworks uploaded and posted to the platform from the *ArsMeteo* community from its launch until December 2010⁷. *ArsM* includes data about a significant set of tagged artworks from the `arsmeteo.org` web portal. As summarized in Table 1, it consists of 9,171 artworks, uploaded by 282 artists. The total number of different tags in the dataset is 36,970 and the average number of tags per artwork is 11.5.

The frequency distribution of artworks in *ArsM* by number of tags is shown in Figure 4. All artworks have at least one tag. It is rare to have less than 5 tags per artwork. About 65% of the artworks have 10 or more tags. This measure

⁷All artworks and comments uploaded and posted after December 2010 are *not* included in our dataset



Figure 5: Preview page for an ArsMeteo’s artwork (left) and biographical sketch for the author, Marzia Migliora (right).

highlights an important aspect of the dataset: having many tags per artwork gives us, potentially, a richer source for the emotional analysis.

The dataset includes also structured metadata released by the authors themselves, according to a metadata schema which includes information about *author’s name*, *genre classification* according to a list of available labels, e.g. photography, music, performance. Metadata are stored in a relational database, managed by the AMA. Almost all such metadata already appears on the portal (see screenshot in Fig. 5, left). Furthermore, unstructured information about artists can be retrieved in biographical sketches (Fig. 5, right). Biographical information often includes interesting geographical data, such as the place of birth (e.g. Alessandria for Marzia Migliora), or also the place where the artist lives and works, which is possibly different (e.g. Turin, in our example). Such information has been processed in order to extract, when present, geographical data about *place of birth* and *place of living* of the author. Related additional features have been included in the ArsEmotica semantic data model that will be introduced below, together with the features aimed at modeling the original items of the ArsMeteo metadata scheme.

5.2. Emotional Analysis on ArsMeteo: Statistics

Emotions belonging to the ontology are detected in about 20 percent of our dataset⁸. Notice that the tagging activity, monitored in ArsMeteo since 2007,

⁸ According to the ArsEmotica emotional analysis, 1705 out of the 9171 artworks in the dataset bear an emotional meaning encoded in the ontology

was not performed with the aim of later applying some kind of emotion detection, but as a form of spontaneous annotation produced by the members of the community, not limited to the emotion sphere. Emotions detected could be more if we applied the analysis to a dataset of tags issued by users with the explicit intention of providing a comment on the emotions aroused by the artworks. Moreover, the fact that our system, at the current stage, processes only one-term tags and not multiword expressions surely is a limit to the recognition of emotions in tags associated to the artworks.

Let us denote with *AffectiveArsM* the set of artworks classified according to some emotions of our ontology after the emotional analysis performed by ArsEmotica.




Figure 6: Frequency distribution of the number of emotions detected in the whole *AffectiveArsM* dataset per artwork.

In ArsMeteo, artworks usually have many tags, expressing a variety of meanings, thus supporting the emergence of different emotional potentials. When this happens, the analysis performed by ArsEmotica provides multiple emotional classifications. Fig. 6 shows the distribution of the number of emotions detected in the whole *AffectiveArsM* dataset per artwork. About 40% of the artworks received multiple classification, i.e., ArsEmotica detected more than one emotion associated to the artwork. Such results seem to be consistent with the idea – confirmed by the psychological findings in the research area related to aesthetic emotions (Silvia, 2005) – that people can have different emotions in response to the same artworks. Moreover, authors of the ArsMeteo collection are mainly contemporary artists, who sometimes intentionally produced works that fall under the category of *controversial art*.

For what concerns the emotion distribution in *AffectiveArsM*, when we consider basic emotions in their varying degree of intensities, the most common emotions were the ones belonging to *sadness* (457 artworks) and *joy* family (405

artworks), followed by *anticipation*, *fear*, *disgust* and *surprise*. *Anger* was rarer, and *trust* was almost absent (see Fig. 7); when we consider complex emotions, results are summarized in Fig. 8: *love* is very common (424 artworks), *optimism* and *awe* are rare, and the other complex emotions are almost absent.




Figure 7: Distribution of emotional labels in ArsMeteo: basic emotions.

The emotion distributions in our dataset summarized in Figures 7 and 8 show that users experience a wide variety of both positive and negative emotions, especially when we look at the basic emotions of the Plutchik’s model. This is in tune with recent studies on aesthetic emotions. Indeed, while in the past, psychological aesthetics, for the most part, was concerned with people’s feelings of pleasure in response to art, recent studies have emphasized the need to cope with a wider spectrum of emotions, including positive emotions like enjoyment and interest, but also negative emotions like disgust, sadness or anger, being the range of aesthetic feelings much wider than liking, preference, and pleasure (Silvia, 2009). In particular let us notice that beside happiness, sadness, fear, and love, all emotions that we would expect to detect in an artistic domain, we find represented in a significant extent also *knowledge emotions* (i.e., emotions associated with thinking and comprehending, like interest –a type of anticipation– and surprise) or *hostile emotions* (i.e. disgust and anger), which have been recently recognized as important in the art domain, since they possibly have a role in understanding cases of artistic censorship and repression occurring frequently in the history of art (Silvia & Brown, 2007). On this perspective, our first experiments seem to show the Plutchik’s model to be adequate to detect different emotions that literature on the psychology of art and emotions identified as relevant in this domain. The encoding of further complex emotions in our ontology, such as pride and shame (belonging to the secondary and tertiary dyads, respectively, in the Plutchik model) could give further interesting results on this line (see (Silvia, 2009)).



Figure 8: Distribution of emotional labels in ArsMeteo: complex emotions (primary dyads only).

5.3. Access to Humans: Visualizing Emotions Detected in ArsMeteo Artworks

Let us now present the interactive user interface that we have developed by its sample application on a tagged artwork from the ArsMeteo collection. The sequence of interactions offered to the user follows the flux of computation sketched in Fig. 1.

5.3.1. 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 9), the four *red colored tags* are identified as emotional since they belong to the emotional ontology: ‘orrore’, ‘infame’, ‘cattiveria’, ‘tristezza’; the presence of emotional responses related to *sadness* and to a strong disgust (*loathing*) is highlighted in the interface 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.

For instance, in the example in Fig.9 all sectors in the disgust’s slice of the wheel are colored with nuances of purple increasingly deeper; in the sadness’s slice only the two most internal sectors are colored; the remaining slices are not colored at all. This means that the system detected in the community tags the



Figure 9: Graphical User Interface inspired to the Plutchik's wheel; showing the results of the automatic emotional analysis

presence of *disgust* with the highest degree of intensity coded in our model, and the presence of *sadness* with a medium intensity. For what concerns the presence of other emotions in our affective model, nothing has been detected at this stage. However, as explained in Section 2, some of the remaining tags could indirectly refer to emotional categories of the ontology and the user can be involved in the definition of such relationships. For this purpose, the system selects a set of further tags (*underlined blue colored tags* in the interface of Fig.9), which have a significant sentiment load according to SentiWordNet. As we assume that such tags could refer to emotional categories of the ontology and could then be good candidates for a user feedback, they appear in the interface as active links for the user's emotional feedback: see e.g. 'sangue' (blood), 'sconfiggere' (to defeat), and so on.

5.3.2. User's Emotional Feedback (Phase 3).

The opportunities offered by our interface to manipulate the emotion wheel, enriching with the feedback of the users the emotional output as automatically computed by the system, are described in the following. Users can annotate, with emotional concepts from the ontology, tags having a significant sentiment load by using an uncolored, neutral, emotion wheel. Since such tags are not in the ontology, they did not contribute to the automatic emotional analysis of the previous stage. The user can select one of the tags represented as active links in order to start an emotional annotation of the tag. Such an annotation will be processed as a further contribution which enriches (and has to be combined with) the outcome of the previous stage. After the tag's selection is done, a pop-up

window appears, showing the uncolored emotional wheel, where all sectors are white to show that the tag has not yet been associated with any emotion. Notice that the tag evaluation is contextual to the experience of the artwork, which indeed remains visible in the background. Users can express their contribution by using the wheel. Each empty sector refers to one of the 8 basic emotions with different intensities and wedge-shaped triangles inserted between the slices refer to complex emotions encoded in our model (primary dyads in the Plutchik’s model). A user who wants to associate a tag with a new emotion, must color the wheel accordingly, by clicking on one of the 24 sectors of the wheel for basic emotions, or on the wedge-shaped triangles for compound emotions. Multiple associations are allowed, as in case of the example in Figure 10, where the user associated to the tag ‘sangue’ (blood) two emotions: *fear* and *disgust* with high intensity (which corresponds to *loathing*). The user, if interested, can repeat this process several times, when more than one tag related to the artwork is offered for the user’s feedback. For instance, in Figure 10 after considering ‘sangue’, the user could decide to consider also ‘perversione’ (perversion) and to start a new process of emotional feedback. Notice that such possibility of iterating the emotional feedback several times for different tags was highly appreciated by the participants of the user study in Section 6.



Figure 10: Interaction with the user: collecting the tag-mediated emotional feedback

5.3.3. 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. The initial outcome is eventually enriched with new further emotional meanings for the artwork resulting from the user’s emotional feedback phase. If requested, the application can format the final output also in the standard markup lan-

guage EmotionML ⁹.

5.4. Access to Machines: ArsMeteo in the LOD

ArsEmotica provides a semantic tagging to ArsMeteo’s artworks. The affective information is encoded in the semantic web ontology of emotions described above and can be used for classifying artworks according to emotional categories. Moreover, as mentioned before, the ArsMeteo dataset includes “standard” collection structured metadata referring to artworks and unstructured information about artists who created them. In order to achieve the goal of enhancing the retrieval of information about the collection, our effort was to represent both emotions and standard metadata (knowledge about authors and on the features of the artifacts) in a unified semantic web representation (Unified ArsEmotica Ontology, UAEO henceforth). Written in OWL2, the UAEO incorporates, in a unifying model, multiple ontologies which describe different aspects of the relationships between *media objects* (e.g. the ArsMeteo artworks), *persons* (e.g. the authors), and *emotions*.

In UAEO the emotions represented in our ontology have been also linked, via `owl:sameAs`, to the corresponding emotions in DBpedia¹⁰. Furthermore, UAEO incorporates an *ontology of artifacts*, derived from the alignment of a domain ontology obtained from the DB of the ArsMeteo on line portal, with the *OMR* (Ontology for Media Resources)¹¹. Such an alignment allows us to express in a semantic unified framework the standard ontological features used for describing the artworks (e.g. the Media Resource type, the format, the creator of an artwork, and so forth). The OMR ontology offers a core vocabulary to describe media resources on the Web, introducing descriptors such as `title`, `creator`, `publisher` and `createDate`.

Finally, in order to enable the representation of the potential networks among artists, the *FOAF model*¹² has been incorporated into UAEO and the information related to the geographical area where the artists operate has been inserted. As in the case of the emotions, the locations have been connected, in a Linked Data perspective, with the corresponding DBpedia URIs. Notice that this potentially enables an interesting linkage to the GeoNames Ontology¹³. The importance to link our data to DBpedia, in fact, is not only that it includes Wikipedia data, but also that it incorporates links to other datasets on the Web, e.g., to GeoNames. By providing those extra links (in terms of RDF triples) applications may exploit the extra knowledge from other datasets for further investigations.

The above described representation enables the possibility to query automatically the ArsMeteo dataset by reasoning on semantic information about

⁹<http://www.w3.org/TR/emotionml/>

¹⁰<http://dbpedia.org/>

¹¹<http://www.w3.org/TR/mediaont-10/>

¹²Friend of a friend ontology: <http://xmlns.com/foaf/spec/>

¹³<http://www.geonames.org/ontology/documentation.html>

emotions, persons and artifacts encoded in the UAEO ontology. Recently, standards and tools for implementing the Semantic Web and the Web of Linked Data have evolved to a state of maturity, that allows us to create applications combining independently developed data sources.

Our starting point has been the development of a demo SPARQL endpoint, accessible at <http://arsemotica.di.unito.it/arsemotica>, relying on Hermit as a reasoner¹⁴, where answers to queries are returned in RDF format.

Access via SPARQL queries and connection of the ArsMeteo dataset to the linked open data cloud will foster interesting and unexpected possibility of reusing the data. In order to test the potential of our semantic representation, we encoded in SPARQL a number of predefined queries such us:

- “Give me the emotions stirred by ArsEmotica’s artworks created by artists operating in the Turin area”
- “Give me the artworks classified as sad and belonging to the *Music* genre”
- “Give me artworks classified as joyful created by artists living in the Salento’s Italian region”.
- “Given a specific author, count the occurrences of different emotions evoked by his/her artworks”

SPARQL encoding for a set of similar queries can be found in our demo SPARQL endpoint. Further reasoning is enabled by adding SWRL rules to the ontology¹⁵. For instance, if we add to the ontology a rule expressing that if two artworks are tagged with opposite emotions, then the artworks are categorized as displaying opposite emotions, a query for finding, given a specific emotion (e.g. joy), all artworks displaying an opposite emotion (e.g. sadness) can be encoded. As we will discuss in the conclusions, the relationships in our model, such as opposition, intensity and similarities could be precious aspects to exploit in web applications enabling emotion-driven browsing of the artworks in a collection.

6. User study

The purpose of the user study was mainly to evaluate the effectiveness of the ArsEmotica application framework for an emotional tagging task, both for what concerns what we call the *affective intelligence* of the prototype, and for what concerns its interface described in Section 5.3, where the hypothesis to test is that the user experience in interacting with the prototype may benefit from a *visualization interface* inspired by the Plutchik’s wheel of emotions.

¹⁴<http://hermit-reasoner.com/>.

¹⁵<http://www.w3.org/Submission/SWRL/>

6.1. Experimental Design, subjects and questionnaire

A set $\mathcal{A}=\{a_1, \dots, a_{16}\}$ of 16 artworks from the ArsMeteo corpus were randomly selected among those which were the most voted by the community (see Table A.3 in the Appendix), explicitly including also artworks with multiple emotional classifications, possibly meaning that they are capable of dividing the community in its perceptions. Artworks in \mathcal{A} were shown as icons in the ArsEmotica prototype homepage, ready to be selected from users in order to start an interaction session along the sequence of interactions described in Section 5.3.

We measured user opinions by means of a questionnaire¹⁶. Oral comments were elicited through a *thinking aloud* technique, since at this stage we were also interested in collecting free feedback on our approach and possibly requirements for the next version of the prototype. The overall experiment was conducted under the guidance of a test moderator.

We selected a group of 20 subjects, 25–70 years old, 10 females and 10 males, among students and colleagues at the Computer Science Department (University of Turin), artists from the ArsMeteo community, linguists from the CELI company¹⁷, according to an availability sampling strategy¹⁸. All subjects were frequent Internet users, familiar with social media.

After filling in the questionnaire with some demographic data such as *age*, *gender*, *qualification*, *ArsMeteo membership* (see Figure B.12), each subject was asked to select an artwork from the prototype homepage. After each interface condition, users compiled the corresponding page of the questionnaire and commented on their experience *thinking aloud*. Optionally, they could select a second artwork and play again with the system by answering the questionnaire. Finally, they were asked to answer a general question about their experience with the ArsEmotica prototype.

The questionnaire included questions devoted to deal with three aspects:

- a. *Affective intelligence*: the effectiveness of the system and its interface in detecting and visualizing the emotions evoked by the artworks according to the Plutchik’s model of emotions;
- b. *Feedback*: the effectiveness of the system and its interface in collecting emotional feedback about the artworks from users;
- c. *Satisfaction*: overall satisfaction w.r.t the experience of interacting with the prototype.

¹⁶The questionnaire was designed using the LimeSurvey tool: <https://www.limesurvey.org>.

¹⁷CELI is an Italian language information technology company, specialized in semantic analysis, multilingual text processing, and knowledge management: <http://www.celi.it/>

¹⁸Notice that, even though non-random samples are not statistically representative, they are often used in psychology researches, as well as in usability testing, especially in early evaluation phases. Motivations in (Singleton et al., 1993; Rubin et al., 2008).

Questions were answered by subjects expressing scores in a 5-point Likert scale, where they are asked to indicate the extent to which they agree with a specific statement. Subjects may respond with 5 = Strongly agree, 4 = Agree, 3 = Neither agree nor disagree (No opinion), 2 = Disagree, 1 = Strongly disagree.

In particular, to evaluate *affective intelligence* of ArsEmotica, three statements in the questionnaire were proposed to subjects, after they observed the emotional outcome provided by ArsEmotica on the selected artwork:

Q1: "The system was successful at conveying the feelings of the ArsMeteo community about the artwork";

Q2: "The graphical interface (wheel of emotions) showing the emotional outcome of the system was appropriate";

Q3: "The system was successful at conveying my emotions about the artwork".

Notice that while in *Q1* the subject is asked to give an evaluation on the emotional outcome of the system based on the information source provided by the ArsMeteo community (the tags), in *Q3* she/he is invited to answer taking into account only her/his personal feelings about the artwork. Collecting such distinct opinions can be interesting in order to reason about correspondence between the ArsEmotica emotional tagging (which is based on the social tagging of the community) and personal feelings of the individual users interacting with our system. Responses to this group of questions were collected for each artwork selected for evaluation.

Furthermore, to evaluate the aspects related to the user *feedback*, user that manifested interest in using the emotion wheel for tagging the artwork with further emotions (Fig. 10) were invited to consider the following statement in the questionnaire after accomplishing the task:

Q4: "Using the graphical interface (wheel of emotions) to associate my emotions to the artwork/tag was intuitive".

Finally, the overall satisfaction from using ArsEmotica was evaluated using statement:

Q5: "I am satisfied with the experience of interacting with the ArsEmotica prototype".

6.2. Analysis of results

Figure 11 shows the distribution of responses (percentages that agree, disagree, etc.) by means of a coordinated set of diverging stacked bar charts (one for each question), by providing an overview on survey responses to the questionnaire on interacting with the ArsEmotica prototype. The percentages of respondents who agree with the statement are shown to the right of the zero line; the percentages who disagree are shown to the left. The percentages for respondents who neither agree nor disagree are split down the middle and are shown in gray color.



Figure 11: Overview on survey responses to the questionnaire on interacting with the ArsEmotica prototype. A total of 20 subjects participated to the test and to the survey. Each person answered one of five levels of agreement or disagreement with the corresponding statement

For the sake of completeness, in the Appendix (Figure B.15) we can also see, for each of our five questions, a diverging stacked bar chart, where responses are organized according to different demographic categories.

Table 2 reports some results of a first analysis of our responses by using descriptive statistics. They were computed on the set of the 20 collected filled questionnaires. All subject but one performed the test on two artworks (with the corresponding responses reported within the same questionnaire), so on the group of questions concerning *affective intelligence* (Q1,Q2,Q3), we collected 39 responses, based on the 39 interactions of the 20 subjects. Columns correspond to questions. For each question, we reported the following values: (i) the *mean* score, which provides the average score value; (ii) the *mode*, showing the most frequent score value; (iii) the standard deviation; (iv) the relative standard deviation. The last two values are reported for understanding whether the mean score can be considered a significant value, i.e. whether it can be perceived as a representative evaluation for a certain aspect of our prototype. Actually, the mean values that we obtained for the various questions range from 3.79 to 4.65.

Table 2: Results of the questionnaires about the ArsEmotica prototype

<i>Questions</i>	<i>Affective intelligence</i>			<i>Feedback</i>	<i>Satisfaction</i>
	Q1	Q2	Q3	Q4	Q5
Mean	4.05	4.31	3.79	4.45	4.65
Mode	4	4	4	5	5
Standard Deviation	0.92	0.73	0.95	0.76	0.49
Relative Standard Deviation	0.23	0.17	0.25	0.17	0.11

All these values can be considered as representative since the relative standard deviation is, in all cases, (much) lower than 0.5.

Affective intelligence. The first question *Q1* concerns the concordance to the emotional analysis, taking into account the social tags as the expressions of the sentiment of the ArsMeteo community. The *mean* and the *mode* for this questions is 4. This suggests that users mostly judge the system effective in detecting emotions from artwork’s tags. For the question *Q2*, which instead focusses on the effectiveness of the system’s interface in visualizing the emotions evoked by the artworks according to the Plutchik’s model of emotions, we got high values as well: the mode is 4. This suggests that mostly subjects enjoyed the interface design inspired by the wheel of emotions. For what concerns *Q3*, which expresses the individual concordance to the emotional outcome, the mean is 3.79. The fact that this result is slightly inferior to the others is, actually, not surprising (many expressions of Art can be controversial for various reasons) and interesting, since it confirms the importance of giving the individual user the possibility to interact with the system in order to express their personal emotional interpretation. Also in this case, however, the mode is 4.

Visualization interface and Feedback. Values for question *Q4*, were very high (the mean is 4.45 and the mode is 5). As confirmed also by the oral free comments collected during the thinking aloud session, subjects mostly enjoyed the phase of the interaction where they were requested to give their emotional feedback to be added and combined with the system’s one. Moreover, statistics show that they evaluated very positively the possibility to enter such feedback by adding colors in an uncolored emotional wheel (see Section 5.3.2).

Overall Satisfaction. Survey results about *Q5* are also very encouraging. *Q5* was asked at the end of the session, and it was meant as measuring the overall satisfaction w.r.t the experience of interacting with ArsEmotica. For this questions we got the highest values: the mean is 4.65 and the mode is 5. We can take it as a signal that a real application based on the ArsEmotica functionalities could be interesting for social web users.

Analysis of the comments from thinking aloud session. Many subjects underlined positive aspects of the possibility to express their emotional evaluation about tags related to artworks and to contribute in this way with their own emotions to the affective annotation of the artwork. They found this possibility

helpful, because (1) commenting through the tags/words is considered precious in expressing feelings about the artwork; (2) this modality offered by the application invites the user to go beyond the immediate aesthetic experience (I like/dislike the artwork); (3) the application, by suggesting the subject to reflect on the emotions evoked by the artworks based on related tags/concepts, drives her/him towards a more conceptual evaluation.

However, many users complained that they could not express their evaluation on *all* available tags, but only on the ones filtered by the systems as the ones that most probably convey some affective meaning. This is for us a useful feedback for improving the next version of the prototype: since ArsEmotica can be intended as a sort of game about the emotions evoked by artworks, it can be better to relax our constraints and to allow users to annotate with their own emotions *any tag* associated to the artwork.

Some subjects complained that the applications did not process tags specified as multiword expressions. Obviously, this is an important issue to address in future work. Other subjects noticed that the emotional evaluation of ArsEmotica do not consider the presence of irony in an artwork (e.g. the irony presence in the artwork “Octopussy in love”, id = 1904), but this can be an element to consider in an emotional evaluation. Since the presence of irony in an artwork is often explicitly marked by the ArsMeteo community (using tags as *ironia* (irony) or *umorismo* (humor)), it could be interesting in the future to try to automatically process this information and add this dimension to the overall evaluation of the artwork.

Some subjects suggested an interesting direction of development: providing the possibility of visualizing the individual emotional annotations (also anonymized) of other users on the same artwork, in order to get an idea about the distance between their emotional assessment and that of others. This is for sure a good suggestion in order to evolve ArsEmotica in the direction of a real *social* application.

The user opinions regarding the graphical aspect of the wheel of emotions were mainly positive about the ‘concept’, but few users commented that it was not easy at the beginning to understand how the intensity function worked.

Finally, ArsMeteo artists commented in a positive way the possibility to develop a real web application offering the ArsEmotica functionalities in order to make available to the community an emotional-driven access to their artworks.

7. Related Work

Relevant contributions for the issues addressed in this paper can be found in a wide range of disciplines. A structured overview is presented hereafter.

7.1. Museums in the Social and Semantic Web

Recently, many cultural heritage institutions opened their collections and archives for access on the web. The list includes several major art museums, such as the Guggenheim Museum, the Metropolitan Museum of Art, the San

Francisco Museum of Modern Art, and the European Rijksmuseum. Initiatives like the Google Art project¹⁹ can be also considered in this perspective.

In this context, on the one hand there have been notable efforts to publish museum data to the Linked Open Data cloud (LOD): let us mention the European project (Haslhofer & Isaac, 2011), the CLAROS initiative (Kurtz et al., 2009), and other valuable efforts of single institutions such as the Amsterdam Museum (de Boer et al., 2013) and the Smithsonian American Art Museum (Szekely et al., 2013). On the other hand, there has been much interest in exploring the ways social technologies and principles of the Social Web can be applied in enhancing the exploration of cultural heritage objects, both on site, and on-line. According to the new concept of *Participatory Museum* (Simon, 2010), artworks can play an important role as “social objects”, i.e., as a basis for an object-centered sociality. Artworks, indeed, are objects that connect people, for instance by fostering conversation. They raise personal feelings and personal degrees of understanding. Thus, an interesting challenge in this domain is in developing 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 (Chae et al., 2012; Trant & Wyman, 2006; Lin & Aroyo, 2012; Baldoni et al., 2012). Recent works such as the visitor’s guide in (Kuffik et al., 2011) or storytelling-based mobile museum guides (Lombardo & Damiano, 2012; Callaway et al., 2012) can be viewed as contributions in this direction.

The kind of work we are presenting here is related to both of the above mentioned research efforts. The main novelty is that the *emotional dimension* moves to the foreground, providing a new key to access and retrieve the data.

7.2. Art, Sentiment and Emotions

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*²⁰, where the experience of museum-goers are analyzed by means of visitor tracking and biometric measurements technologies.

ArsEmotica analyses tags and extracts 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. This can be a key element for encouraging visitors of collections to share their experiences about the artworks. 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

¹⁹<http://www.googleartproject.com/>

²⁰Mapping Museum Experience: <http://www.mapping-museum-experience.com/en>

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 explore the Wilson's collection to online visitors.

Other relevant contributions on art and emotions can be found in a wide set of disciplines, ranging from art and aesthetics to psychology. In particular, recent investigations on psychology of art, which recognizes a specific role of 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 (Bacci & Melcher, 2011) by taking a multi-disciplinary perspective, where computational approaches to emotion recognition can play a key role. Specifically, on this line, recently an interesting experiment has been conducted (Yanulevskaya et al., 2012), 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.

7.3. Ontologies of emotions and affective resources

Recently, many researchers are devoting efforts in developing ontology of emotions in the Semantic Web context (Cambria & Hussain, 2012; Grassi, 2009; Francisco et al., 2010; Sánchez-Rada & Iglesias, 2015), and some of them addressed the issue from a foundational point of view. In particular, the Human Emotion Ontology (HEO) developed in OWL (Grassi, 2009), was introduced with the explicit aim to standardize the knowledge about emotions and to support very broad semantic interoperability among affective computing applications. It will be interesting to study how to link the ArsEmotica's ontology of emotions with HEO, which could play for us the role of "upper ontology" for emotions, by providing an ontological definition of the general concept of emotion. In fact, in our ontology the root concept of the *Emotion* hierarchy is treated as primitive.

In (Sánchez-Rada & Iglesias, 2015) a semantic vocabulary, called Onyx, is proposed, with the main aim to describe emotion analysis processes and results using semantic web and linked data technologies²¹. Onyx includes classes and properties to provide structured and meaningful emotion analysis results, and to connect results from different providers and applications. We plan to explore

²¹<https://www.gsi.dit.upm.es/ontologies/onyx/>

the possibility to use this vocabulary in order to semantically annotate the key ingredients and the results of the emotion analysis provided by ArsEmotica.

The Hourglass Model, which is integrated in our affective model in order to combine detected and collected emotions in ArsEmotica, allows us to design a fluid and continuous emotional space, where artworks (but also possibly user’s tag) can be positioned. Moreover, it is well-known that the Hourglass of Emotions is a model specifically designed for concept-level sentiment analysis. Hence, it would be interesting to study the use in ArsEmotica of concept-level resources (e.g. SenticNet (Cambria et al., 2014)) instead of word-level ones.

8. Conclusions and Future Work

In this paper we presented an application framework where semantic technologies, linked data and natural language processing techniques have been exploited to investigate the emotional aspects of cultural heritage artifacts, based on user sentiment detected in art social platforms. The aim is twofold: to recognize and visualize the emotional impact of artworks on people; to enable an emotion-driven organization, access and retrieval of artworks and related data in on-line collection.

We have described the OWL ontology of emotions developed for the ArsEmotica prototype. The ontology has been conceived for categorizing emotion-denoting words and has been semi-automatically populated with Italian terms from WN-Affect. In order to make the framework applicable to other languages, such as English, and eventually to multilingual descriptions of artworks, we are currently working to an integration of the ArsEmotica ontology with the LEXicon Model for ONtologies (LEMON) (McCrae et al., 2011). Such an integration will allow us to classify Princeton WordNet English synsets representing affective concepts under our emotional categories, and, then, to exploit the available multilingual lexical databases aligned to WordNet to express the link to affective lexical entries (emotion-denoting words) in the language we are interested to deal with (e.g. MultiWordNet for Italian)²². The use of a multilingual encyclopedic dictionary such as BabelNet²³ will be also useful to deal with the multilinguality issue.

Moreover, we plan to develop the ontology and extend the coverage of our emotion lexicon, also in order to cope with multi-word expressions, that may not explicitly convey emotions, but are related to concepts that do, as it has been confirmed also by the results of our user study. To deal with such an issue, it will be convenient to rely on resources like EmoSenticNet (Poria et al., 2014a) and to adopt a concept-based approach, like the one described in (Poria et al., 2014b).

²²A list of wordnets in several languages, all linked to the Princeton WordNet of English, has been collected in the context of the Open Multilingual WordNet initiative: <http://compling.hss.ntu.edu.sg/omw/>

²³<http://babelnet.org/>

The affective model refers to a state-of-the-art cognitive model of emotions and inspired an interactive user interface for visualizing and summarizing the results of the emotion detection algorithm. The current ArsEmotica interface provides our users with the possibility to access the outcomes of the emotional analysis. On this line, the next step is to study innovative strategies to *browse* the artworks, by relying on their semantic organization in the ArsEmotica emotional space. The aim is to provide users with the possibility to explore the resources by exploiting the various dimensions suggested by the ontological model. Possible user preferences to deal with could be: “show me sadder artworks” (intensity relation); “show me something emotionally completely different” (polar opposites); “show me artworks conveying similar emotions” (similarity relation). Another interesting direction we are going to investigate is the possibility to use the ArsEmotica framework in order to define innovative recommendation strategies based on affective information on contents (Tkalcic et al., 2013). Notice that, in the current interface, we are interested to offer a summary of the emotions evoked by artworks, detecting their simple presence by analysing tags, without introducing a measure about their relative strength (e.g. tags related to emotion *A* are more frequent than tags related to emotion *B*). However, in other contexts it would be interesting to use frequencies of emotion labels to give to the users a measure about different strengths of the emotions. This information, could be exploited to rank evoked emotions in case of multiple classification, and to recommend users to visit artworks emotionally similar w.r.t. the prevalent emotional category.

Another interesting issue to address is investigating how the emotional user experiences in relation to a given artwork can vary over time. The positioning of an artwork in the emotional space created by the emotional ontology is not static in time, but *dynamic*. Users of a tagging platform can insert new tags (from which new emotions can be inferred), new users can insert their emotional feedbacks about tags, or we envision the possibility that the same user can have different responses to the same artwork at different point of time. Accordingly, the emotion-driven browsing experience provided can be different from time to time. We can envisage an enhanced application framework where, initially, the position of the artworks in the emotional space will be mainly determined by the interaction of artists and curators: they will be the first ones to add meanings to the artworks, and then to give an input to the “emotional engine”. Later, when the application starts to collect the new meanings expressed by the visitors, artworks will start to float in the emotional space. New artworks and meanings can be added anytime, the emotional relations will continuously change, by reflecting the evolution of the community and its latent perception of a sort of emotional “zeitgeist”. Dealing with these dynamic aspects is out of the scope of the current proposal, but it could be an interesting line of research to follow in future work.

ArsEmotica has been tested on social data and artworks from the real-world online collection arismetoo.org. The ArsMeteo dataset, enriched with the metadata about the emotions detected by ArsEmotica, is available in RDF format and can be accessed via a demo SPARQL endpoint. A novel unified semantic

data model has been defined, where artworks belonging to the collection were semantically described by referring to emotional categories of an ontology of emotions. The framework allows us to model relations among artworks, persons and emotions, by combining the ArsEmotica ontology of emotions with available ontologies, such as FOAF and OMR. Moreover, where possible and relevant, our data were linked to external repositories of the LOD, such as DBpedia.

A user study has been proposed with the main aim to test the effectiveness of the framework for an emotional tagging tasks, with a special focus, on the one hand, on the underlying affective model, on the other hand, on the *visualization interface* inspired by the Plutchik's wheel of emotions. Results were very positive on both aspects, and comments of subjects during the thinking aloud session will be very helpful in designing future developments of the framework. In particular, some subjects suggested to develop the social level of the application, since they are interested in comparing their emotional responses with the maybe different emotional responses of other users (friends or simply visitors) interested in the same artwork. Moreover, it emerged that users want to have more freedom in choosing a tag for the feedback, as they expressed the desire to evaluate tags that were not selected as having a significant sentiment load according to SentiWordNet. As observed in (Strapparava et al., 2006), in some cases the affective power of a word is part of the collective imagination (think for instance of words like 'war'), but some words can be emotional for someone due to her individual story. Therefore, it is maybe better to widen the filter of words to be offered for the user feedback, since the current one was perceived as restrictive.

For what concerns the *prospected applications* of the ArsEmotica framework, it could be exploited, along the direction traced in (Simon, 2010), as a co-creation instrument for *museums*, *virtual galleries*, and other activities falling under the general umbrella of *creative industries*. In this sector, the demand is growing for new user experiences, and therefore for applications, including smart-phone apps, where the key aim is to stimulate user-community interaction and encourage visitors to share their experiences. Such applications can have a cultural flavor but can also be more intrinsically related to leisure, and should help transforming classical art-fruition experiences into innovative, more immersive experiences, with a greater impact on visitors.

Art and emotions are naturally related (Silvia, 2005) and, as also our user study confirmed, the possibility to speculate about the artistic arousal of emotions and to reason on artworks, authors and evoked emotions within a social dimension, where personal feelings about an artworks can be compared with what is felt by others, seems to be attractive from a visitor's perspective. Moreover, the possibility to collect emotional responses to artworks and collections can be important for artists and curators, in order to get a feedback about their creations, 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?). Another interesting application field is given by a growing number of virtual galleries, such as <http://www.saatchiart.com>, that raised in the last

years with a business perspective, having the main aim to sell online artworks (visual arts, especially). Also, in this case, detecting emotions in keywords and other information associated to the artworks could be useful in order to offer new emotion-driven searching functionalities to the website customer, to be possibly combined with traditional searching criteria based on genres (e.g. paintings, photography or sculpture) which encompass stylistic characteristics. Recently, a similar approach has been already successfully applied to digital music, with the proposal of music streaming services that instead of finding music by genre or musical-relation, propose a selection tailored to the user's mood and feelings, e.g. Stereomood, <http://www.stereomood.com>.

Finally, the possibility to exploit the geographic information about the place where artists work, together with the information about the emotion detected in their artworks by social tagging, encourages reflections on new applications exploiting the relationship between art and places (on the line of the ArtMaps project on the Tate's collection ²⁴), where the collaborative creation of dynamic art maps to visualize emotions stirred by artworks in different geographical area could be envisioned.

Acknowledgements

The authors thank all the persons who supported the work with comments and contributions. In particular, Flavio Portis, Giorgio Vaccarino and the *Associazione Culturale ArsMeteo* for providing the data corpus; Antonio Lieto for the knowledge engineering support; Andrea Bolioli and *CELI s.r.l.* for the fruitful discussions on lexical and sentiment resources, Fondazione Bruno Kessler and ISTI CNR for supplying the lexicons; subjects of the user study for their time and helpful comments. This research is partially supported by MIUR, Italian local research grant for young researchers 2013: "Ontology-based Affective Models to Organize and Access Artworks in the Social Semantic Web".

References

- Acotto, E., Baldoni, M., Baroglio, C., Patti, V., Portis, F., & Vaccarino, G. (2009). Arsmeteo: artworks and tags floating over the planet art. *In Proc. of ACM HT '09, ACM*, 331–332.
- Bacci, F., & Melcher, D. (2011). *Art and the Senses*. OUP Oxford. URL: <http://books.google.it/books?id=61Hj9o3mPBUC>.
- Baldoni, M., Baroglio, C., Patti, V., & Rena, P. (2012). From tags to emotions: Ontology-driven sentiment analysis in the social semantic web. *Intelligenza Artificiale*, 6, 41–54.

²⁴<http://artmaps.tate.org.uk/>

- Baldoni, M., Baroglio, C., Patti, V., & Schifanella, C. (2013). Sentiment analysis in the planet art: A case study in the social semantic web. In C. Lai, G. Semeraro, & E. Vargiu (Eds.), *New Challenges in Distributed Information Filtering and Retrieval* (pp. 131–149). Springer volume 439 of *Studies in Computational Intelligence*.
- Bertola, F., & Patti, V. (2013). Organizing artworks in an ontology-based semantic affective space. In C. Battaglino, C. Bosco, E. Cambria, R. Damiano, V. Patti, & P. Rosso (Eds.), *ESSEM@AI*IA* (pp. 119–130). CEUR-WS.org volume 1096 of *CEUR Workshop Proceedings*.
- de Boer, V., Wielemaker, J., van Gent, J., Oosterbroek, M., Hildebrand, M., Isaac, A., van Ossenbruggen, J., & Schreiber, G. (2013). Amsterdam museum linked open data. *Semantic Web*, 4, 237–243.
- Callaway, C., Stock, O., Dekoven, E., Noy, K., Citron, Y., & Dobrin, Y. (2012). Mobile drama in an instrumented museum: inducing group conversation via coordinated narratives. *New Review of Hypermedia and Multimedia*, 18, 37–61.
- Cambria, E., & Hussain, A. (2012). *Sentic Computing: Techniques, Tools, and Applications*. SpringerBriefs in Cognitive Computation Series. Springer-Verlag GmbH.
- Cambria, E., Livingstone, A., & Hussain, A. (2012). The hourglass of emotions. In A. Esposito, A. M. Esposito, A. Vinciarelli, R. Hoffmann, & V. C. Müller (Eds.), *COST 2102 Training School, Revised Selected Papers*. Springer volume 7403 of *Lecture Notes in Computer Science*.
- Cambria, E., Olsher, D., & Rajagopal, D. (2014). Senticnet 3: A common and common-sense knowledge base for cognition-driven sentiment analysis. In *AAAI Conference on Artificial Intelligence*.
- Cambria, E., Schuller, B., Xia, Y., & Havasi, C. (2013). New avenues in opinion mining and sentiment analysis. *IEEE Intelligent Systems*, 28, 15–21.
- Chae, G., Park, S. J., Stein, R., Kim, J., & Wiedenbeck, S. (2012). Exploring affective computing for enhancing the museum experience with online collections. In *Proc. of the Museums and the Web*.
- Esuli, A., Baccianella, S., & Sebastiani, F. (2010). SentiWordNet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proc. of LREC'10*. ELRA. URL: http://www.lrec-conf.org/proceedings/lrec2010/pdf/769_Paper.pdf.
- Fellbaum, C. (Ed.) (1998). *WordNet: An Electronic Lexical Database*. MIT Press.

- Francisco, V., Gervas, P., & Peinado, F. (2010). Ontological reasoning for improving the treatment of emotions in text. *Knowledge and Information Systems, 25*, 421–443.
- Galati, D., Sini, B., Tinti, C., & Testa, S. (2008). The lexicon of emotion in the neo-latin languages. *Social Science Information, 47*, 205–220.
- Grassi, M. (2009). Developing the human emotions ontology. In *Proc. of the 2009 joint COST 2101 and 2102 international conference on Biometric ID management and multimodal communication* (pp. 244–251). Springer-Verlag.
- Haslhofer, B., & Isaac, A. (2011). data.europeana.eu - the europeana linked open data pilot. In *Proc. of International Conference on Dublin Core and Metadata Applications*.
- Kuflik, T., Stock, O., Zancanaro, M., Gorfinkel, A., Jbara, S., Kats, S., Sheidin, J., & Kashtan, N. (2011). A visitor's guide in an active museum: Presentations, communications, and reflection. *J. Comput. Cult. Herit., 3*, 1–25. URL: <http://doi.acm.org/10.1145/1921614.1921618>. doi:10.1145/1921614.1921618.
- Kurtz, D., Parker, G., Shotton, D. M., Klyne, G., Schroff, F., Zisserman, A., & Wilks, Y. (2009). Claros - bringing classical art to a global public. In *eScience* (pp. 20–27). IEEE Computer Society. URL: <http://dblp.uni-trier.de/db/conf/eScience/eScience2009.html#KurtzPSKSZW09>.
- Lin, Y., & Aroyo, L. (2012). Interactive curating of user tags for audiovisual archives. In *Proceedings of the International Working Conference on Advanced Visual Interfaces AVI '12* (pp. 685–688). New York, NY, USA: ACM. URL: <http://doi.acm.org/10.1145/2254556.2254685>. doi:10.1145/2254556.2254685.
- Lombardo, V., & Damiano, R. (2012). Storytelling on mobile devices for cultural heritage. *New Review of Hypermedia and Multimedia, 18*, 11–35. doi:10.1080/13614568.2012.617846.
- McCrae, J., Spohr, D., & Cimiano, P. (2011). Linking lexical resources and ontologies on the semantic web with lemon. In *The Semantic Web: Research and Applications* (pp. 245–259). Springer.
- Mohammad, S., & Turney, P. D. (2013). Crowdsourcing a word-emotion association lexicon. *Computational Intelligence, 29*, 436–465.
- Pianta, E., Bentivogli, L., & Girardi, C. (2002). Multiwordnet: developing an aligned multilingual database. In *Proc. of Int. Conf. on Global WordNet*.
- Plutchik, R. (1997). The circumplex as a general model of the structure of emotions and personality. In R. Plutchik, & H. R. Conte (Eds.), *Circumplex models of personality and emotions* (pp. 17–47). American Psychological Association.

- Plutchik, R. (2001). The Nature of Emotions. *American Scientist*, 89.
- Poria, S., Cambria, E., Gelbukh, A., Hussain, A., & Huang, G. (2014a). EmoSenticSpace: A novel framework for affective common-sense reasoning. *Knowl.-Based Syst.*, 69, 108–123.
- Poria, S., Cambria, E., Winterstein, G., & Huang, G. (2014b). Sentic patterns: Dependency-based rules for concept-level sentiment analysis. *Knowl. Based Syst.*, 69, 45–63.
- Rubin, J., Chisnell, D., & Spool, J. (2008). *Handbook of Usability Testing: How to Plan, Design, and Conduct Effective Tests*. Indianapolis, IN: Wiley.
- Sánchez-Rada, J. F., & Iglesias, C. A. (2015). Onyx: A Linked Data approach to emotion representation. *Information Processing and Management*, . In Press.
- Schroeder, M., Pirker, H., Lamolle, M., Burkhardt, F., Peter, C., & Zovato, E. (2011). Representing emotions and related states in technological systems. In R. Cowie, C. Pelachaud, & P. Petta (Eds.), *Emotion-Oriented Systems Cognitive Technologies* (pp. 369–387). Springer.
- Silvia, P. J. (2005). Emotional Responses to Art: From Collation and Arousal to Cognition and Emotion. *Review of General Psychology*, 9, 342–357.
- Silvia, P. J. (2009). Looking past pleasure: Anger, confusion, disgust, pride, surprise, and other unusual aesthetic emotions. *Psychology of Aesthetics, Creativity, and the Arts*, 3, 48–51.
- Silvia, P. J., & Brown, E. M. (2007). Anger, disgust, and the negative aesthetic emotions: Expanding an appraisal model of aesthetic experience. *Psychology of Aesthetics, Creativity, and the Arts*, 1, 100–106.
- Simon, N. (2010). *Participatory Museum*. Museum 2.0. URL: <http://www.participatorymuseum.org/>.
- Singleton, R., Straits, B., & Straits, M. (1993). *Approaches to social research*. Oxford University Press.
- Strapparava, C., & Valitutti, A. (2004). WordNet-Affect: an affective extension of WordNet. In *Proc. of LREC'04* (pp. 1083–1086). volume 4.
- Strapparava, C., Valitutti, A., & Stock, O. (2006). The affective weight of lexicon. In *Proceedings of the 5th International Conference on Language Resources and Evaluation (LREC-2006)* (pp. 423–426).
- Suttles, J., & Ide, N. (2013). Distant supervision for emotion classification with discrete binary values. In *Computational Linguistics and Intelligent Text Processing - 14th International Conference, CICLing 2013* (pp. 121–136). Springer volume 7817 of LNCS.

- Szekely, P. A., Knoblock, C. A., Yang, F., Zhu, X., Fink, E. E., Allen, R., & Goodlander, G. (2013). Connecting the smithsonian american art museum to the linked data cloud. In P. Cimiano, O. Corcho, V. Presutti, L. Hollink, & S. Rudolph (Eds.), *Proc. of ESWC 2013* (pp. 593–607). Springer volume 7882 of *Lecture Notes in Computer Science*.
- Tkalcic, M., Odic, A., Kosir, A., & Tasic, J. F. (2013). Affective labeling in a content-based recommender system for images. *IEEE Transactions on Multimedia*, *15*, 391–400.
- Trant, J., & Wyman, B. (2006). Investigating social tagging and folksonomy in art museums with steve.museum. In *Proc. of the Collaborative Web Tagging Workshop (WWW'06)*.
- Yanulevskaya, V., Uijlings, J., Bruni, E., Sartori, A., Zamboni, E., Bacci, F., Melcher, D., & Sebe, N. (2012). 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* (pp. 349–358). ACM. doi:10.1145/2393347.2393399.
- Zanchetta, E., & Baroni, M. (2005). Morph-it! a free corpus-based morphological resource for the Italian language. *Corpus Linguistics 2005*, *1*.


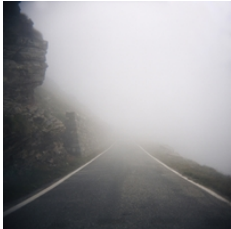

Appendix A. Survey: artworks in \mathcal{A}

We show artworks in \mathcal{A} exploited in the user study in Sec. 6. For each artwork we show: thumbs, title, author, id in the ArsMeteo DB, tags from the ArsMeteo community.

Table A.3: Artworks selected for our user study.

Artwork, title, artist and ID	tags
 <p data-bbox="378 909 607 999"><i>Che autoritratto!</i> by <i>S. Zito</i> - id: 1614</p>	<p data-bbox="638 642 1242 814">#tags: 27 - holy_smoke, oro, giovinezza, icona, we_shall_overcome, fumo, comandante, puro_humo, somos, sigaro, sornioneria, maschio, zito_mito, chi, un_bel_di_vedremo, mito_zito, un_fil_di_fumo, uomo, disincanto, che, bello_e_impossibile, rivoluzionario, occhi, guevara, querida_presentia, barbudo, di_vedremo</p>
 <p data-bbox="378 1339 607 1430"><i>Cannibalismo sociale</i> by <i>R. Palma</i> - id: 1795</p>	<p data-bbox="638 1041 1242 1381">#tags: 43 - mar_terra_cielo, rigetto, raffaele_palma, homo_homini_lupus, uomo_mangia_uomo, cannibalismo, lettera_c, alimento, evoluzione, paura, mare, pesce, grande_piccolo, scatoletta, pericolo, acquatico, fauci, cosi_è_la_vita, china_pastelli_carta, acqua, ironia, vita_morte, natura, destino, segno_zodiacale, la_vita_e_un_grande_ristorante, segno_zodiaco_pesci, alfameteo, caus, raccapricciante, sopraffazione, pressione_sociale, pesci_ascendente_pesci, pesce_cani, pescioska, catena_alimentare, sardinoske, economia_spaziale, morta_eppur_natura, pesce_mangia_pesce, inscatolati, sardine, metafora</p>
 <p data-bbox="378 1675 607 1766"><i>Cartolina di Natale</i> by <i>M. Vitale</i> - id: 1892</p>	<p data-bbox="638 1472 1242 1703"># of tags: 36 - natale_operai, dignità, neve, albero, gelata, alberello, lacrima_di_natale, silenzio, dolore, essenziale, lutto_festivo, guerra, concettuale, mortale, festività, abete_piangente, specchio, espressione, volto, donna, abbasso_il_natale, rughe, depressione, natale, irak, sangue_sudore_e_lacrime, tristezza, esaustivo, grande_zot, segnal_etico, lacrima, stilla, maccaturu, abete, smorfia, auguri</p>

 <p><i>Una farfalla</i> by L. Alessandri - id:1957</p>	<p># of tags: 25 - ingenuità, ferita, lorenzo_alessandri, surfanta, farfalla, la_bella_e_la_bestia, mostro_collina, sguardo_triste, surreale, fantasia, pittore_surrealista, surrealista, surrealismo, mutanti, dopobomba, sofferenza, degradazione, deforme, italian_surrealist, albero, pascalosi, ominide, lettera_j, meraviglia, alessandri</p>
 <p><i>Insomnia Metafisica</i> by M. Vitale - id: 800</p>	<p># of tags: 34 - rotazione, caffettiera, dialogo, meditazione, incanto, parlo, poesia, insonnia, luna, caffè, profumo, donna, bagliori, slancio, preghiera, magia, riflessione, danza, sospensione, dervishi, scacchiera, estasi, fuoco, illuminazione, silenzio, sufi, aroma, luce, passo, quarto_di_luna, caffeina, gigante, metafisica, intrappolato</p>
 <p><i>Let Down</i> by T. Cerasuolo - id: 1293</p>	<p># of tags: 11 - forza_gravità, capelli, lasciarsi_andare, resistenza, caduta, precipizio, nuvola, paura, sedile, discesa, precipitare</p>

 <p><i>Atterraggio a Leumann</i> by L. Alessandri - id: 1902</p>	<p># of tags: 57 - devastazione, apocalisse, scritte_sul_muro, biplano, aereo, operai, danza_macabra, jeronimus, tramontana, demoniaco, felix_the_cat, il_barone_azzurro, tegole, cielo_plumbeo, tempesta, azzurro, sigla, cielo, surrealismo.e.fantasia, oltre_la_fantasia, alberi, strada, frastornati, frastuono, villaggio_leuman, panico, ombrello, trasportati_dal_vento, surfanta_shop, ciminiera, fumo, case, partenza, arrivo, bovindo, balconi, satanasso, pilota, emergenza, pittore_surrealista, scheletri, lorenzo_alessandri, vento_da_morire, non_morti, morti, capolavoro, opera_somma, bufera, surfanta, fuggire, s.o.s, tetti_scoperchiati, leggeri_come_piume, olio_su_masonite, vento, leumann, atterraggio</p>
 <p><i>Nebbia</i> by M. Coretto id: 1243</p>	<p># of tags: 10 - norway, scomparsa, lettera_a, muro, grigio, silent_hill, ignoto, strada, altrove, nebbia</p>
 <p><i>Octopussy in love</i> by R. Palma - id: 1904</p>	<p># of tags: 43 - texture, china_pastello_carta, pink_humor, penose_allusioni, pesce, abbaglio, necessità_amore, pòrpo, abbraccio, tenerezze_sommerse, amore, fluidità, piovra, ventosa, innamoramento, umorismo, caus, lui_lei, tenerezza, mare, acqua, segno_zodiacale, coccole, tentacolare, stura_sentimento, avvinghiarsi, magari_potessi, appiccicati, affetto_adesivo, la_polpa_del_polpo, tentacoli, segno_zodiaco_pesci, cuoricini, occhi_allusivi, mollusco, bolle_amorose, dolce_metà, no_sex, affinità_elettive, poetico, amare_ogni_cosa, un_mare_di_bene, mimetismo</p>

 <p><i>Rinoceronte</i> by L. Alessandri - id: 1964</p>	<p># of tags: 28 - medicina_cinese, prestanza, momentum, surfanta, illegale, polverone, biocorno, thrust, superpersonico, pellaccia, pascalosi, italian_surrealist, afrodisiaco, surrealista, pittore_surrealista, furore_blu, animale, surrealismo, blu, corsa, lorenzo_alessandri, furia, carica, determinazione, rinoceronte, fantasia, sogno, energia_animale</p>
 <p><i>Un colpo di vento</i> by L. Alessandri - id: 1896</p>	<p># of tags: 35 - notte_stellata, colpo_di_vento, stelle, lettera_i, lorenzo_alessandri, vento, surrealismo, precarietà, melancholia, acedia, lego, noia, cielo_stellato, surfanta, tarocco_arcano_torre, melanconia, demone_meridiano, accidia, surrealista, fantasia, crollo, meditabondo, violino, resaca, notte, pittore_surrealista, acidità, luna, torre, falce_di_luna, simbolo, infondato, il_pensatore, hang_over, gueule_de_bois</p>
 <p><i>Reading in the car</i> by D. Posdocimo - id: 783</p>	<p># of tags: 44 - ombre, mollezze, onda, incantesimo, libro, rosso, leggere, cosa_buona_e_giusta, parole_parole_parole, bianco, sol_y_sombra, eleganza, magia, drap_rouge, luce, seta, caratteri, tao, mano, relax, femmina, illuminazione, raggio, scrittura, carattere, tempo, spazio, benessere, equilibrio, concentrazione, assorta, incanto, viaggio, rojo_que_te_quiero_rojo, piacere, lettura, sospensione, rosso_nero, parole, supponiamo_un_amore, appagamento, completezza, armonia, leggersi</p>
 <p><i>Tango</i> by G. Vaccarino - id:431</p>	<p># of tags: 21 - cerimoniale, aria, lampo, sfida, sguardo, condivisione, suspense, affidamento, complanare, coppia, netville, sorpresa, intesa, deriva, carmen, cerniera, incontro, evento, soddisfazione, complicità, desiderio</p>

 <p><i>Dove la Raffinata Ragazza Bionda guarda il Grosso Toro Malmorto</i> by F. Valente id:512</p>	<p># of tags: 24 - fotografia, crudeltà, barbarie, carne, toro, tristezza, arena, trascinamento, perversione, animale, preversione, corrida, sangue, biondo, orrore, morte, bianco_e_nero, la_bionda, spagna, sconfitta, cattiveria, veicolo_animale, inservienti, infame</p>
 <p><i>Il tunnel dei Radiohead</i> by M. Comoretto - id: 599</p>	<p># of tags: 21 - edificio_sotterraneo, brivido, melania_comoretto, infinito_limitato, risucchio, angoscia, fuga, infinito, divenire, radiohead, piranesi, underground, grigio, abisso, organico, esofago, tunnel, galleria, labirinto, escher, sotterraneo</p>
 <p><i>Le vespe sentono sete ma nessuno sa come</i> by F. Valente - id: 470</p>	<p># of tags: 20 - ronzio, paura, dolce, tensione, minuta_perfezione, vita_di, vespa, brividi, senza_sudorifere, bere, fotografia, attività, zzzzzzz, goccia, bianco_e_nero, sete, macro, fobia, lavoro, acqua</p>

Appendix B. Survey: questions and results

This section includes the main screenshots of the survey designed for the ArsEmotica user study. In Figure B.15 a set of diverging stacked bar charts gives an overview about the answers to the questionnaire (see Section 6), one chart for each question, with details about the distribution of the responses in different demographic categories.

Arsemotica

The purpose of the user study is to evaluate some aspects of the Arsemotica prototype. In particular: a) the effectiveness of the system and its interface in detecting and visualizing the emotions evoked by the artworks according to the Plutchik's model of emotions; b) the effectiveness of the system and its interface in collecting emotional feedback about the artworks from users; c) overall satisfaction w.r.t the experience of interacting with the prototype.

0%
100%

Pre-test: demographic data

Insert date
01.01.2015
Format: dd.mm.yyyy

Insert user id

What is your age?
Choose one of the following answers

- 21-30
- 31-40
- 41-50
- 51-60
- 61-70
- 71 or older

What is your gender?
Choose one of the following answers

- female
- male

What is your qualification? (mark the highest degree/status received/reached)
Choose one of the following answers

- PhD
- PhD student
- graduate
- Other:

Are you an artist of the arsmeteo.org community?

- Yes
- No

Figure B.12: Collection of the demographic data.

Arsemotica

The purpose of the user study is to evaluate some aspects of the Arsemotica prototype. In particular: a) the effectiveness of the system and its interface in detecting and visualizing the emotions evoked by artworks according to the Plutchik's model of emotions; b) the effectiveness of the system and its interface in collecting emotional feedback about the artworks from users; c) overall satisfaction w.r.t the experience of interacting with the prototype.

0% 100%

Affective intelligence

▪ Select the artwork you wish to explore with Arsemotica.

Dove la Raffaella Ragazza Bionda guarda il Grosso Toro Malmorto

▪ How much do you agree with the following statements?

	Strongly agree	Agree	Neither agree nor disagree	Disagree	Strongly disagree
The system was successful at conveying the feelings of the ArseMoto community about the artwork	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The graphical interface ('wheel of emotions') showing the emotional outcome of the system was appropriate	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The system was successful at conveying my emotions about the artwork	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

▪ Would you like to add an emotional evaluation to potentially affective tags related to the artwork?

Yes No

▪ How much do you agree with the following statements?

	Strongly agree	Agree	Neither agree nor disagree	Disagree	Strongly disagree
Using the graphical interface ('wheel of emotions') to associate my emotions to the artwork/tag was intuitive	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

▪ How many tag-mediated emotional feedbacks have you expressed?

2

Only numbers may be entered in this field

Resume later
Next >>
Exit and clear survey

Figure B.13: Questions Q1, Q2, Q3 (affective intelligence) and Q4 (feedback) on a selected artwork.

Arsemotica

The purpose of the user study is to evaluate some aspects of the Arsemotica prototype. In particular: a) the effectiveness of the system and its interface in detecting and visualizing the emotions evoked by artworks according to the Plutchik's model of emotions; b) the effectiveness of the system and its interface in collecting emotional feedback about the artworks from users; c) overall satisfaction w.r.t the experience of interacting with the prototype.

0% 100%

Overall satisfaction

▪ How much do you agree with the following statement?

	Strongly agree	Agree	Neither agree nor disagree	Disagree	Strongly disagree
I am satisfied with the experience of interacting with the Arsemotica prototype	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Resume later
Submit
Exit and clear survey

Figure B.14: Question Q5 (overall satisfaction about the interaction experience).




Figure B.15: Overview of the results on the ArsEmotica survey: for each of our five questions, a diverging stacked bar chart is shown, where responses are organized according to different demographic categories.