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

Matteo Baldoni, Cristina Baroglio, Viviana Patti and Paolo Rena

Abstract Affective computing is receiving increasing attention in many sectors, ranging from advertisement to politics. This work, set in a Social Semantic Web framework, presents ArsEmotica, an application software for associating the predominant emotions to artistic resources of a social tagging platform. Our aim is to extract a rich emotional semantics (i.e. not limited to a positive or a negative reception) of tagged resources through an ontology driven approach. This is done by exploiting and combining available computational and sentiment lexicons with an ontology of emotional categories.
The information sources we rely upon are the tags by which users annotated resources, that are available through the ArsMeteo platform, and the ontology OntoEmotion, that was enriched by means of our tool with over four hundred Italian emotional words referring to the about eighty-five emotional concepts of the ontology. Tags directly referring to ontological concepts are identified, while potentially affective tags, can be annotated by using the ontology thanks to the spontaneous intervention of the user, in a pure Web 2.0 approach. Finally, the tagged artworks are related with the emerging predominant emotions.

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

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

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

Paolo Rena
Dipartimento di Informatica, Università degli Studi di Torino e-mail: paolo.rena@gmail.com

1 Introduction

In the last years, the conception of the Web evolved from a web of documents to a web of users: the Social Web. Users are more and more involved in the production of contents or in their elaboration, e.g. by publishing and organizing own materials, by posting comments for discussing newspaper articles, by participating into wikis, by rating resources. Social networks and platforms (Facebook, Flickr, Youtube, LastFM, Anobii, StumbleUpon to cite some among the best known) promote the participation of users in many ways, stimulating the expression of opinions about the contents inserted by other users, by supplying simple "Like" or "Dislike" tools, star-rating systems, tag-based annotation and navigation, and so forth. This huge amount of data is a precious information source, about perceptions, trends, and feelings, and a lot of research work is being carried on to identify ways for extracting meaningful information from these data.

One of the emerging research fields, aimed at extracting information from the data supplied by the Social Web users, is emotion-oriented computing (a.k.a. Affective Computing [12]), whose focus is to automatically recognize the users' emotions by analyzing their tagging or writing behavior. Since emotions are often related to appreciation, knowing the feelings of the users towards target issues is an important feedback that can support many decisional tasks. The recent success of Sentiment Analysis (or Opinion Mining [9]) techniques applied to business applications, and the development of tools like Twitter Sentiment are significant cases (http://twittersentiment.appspot.com/). Such techniques, however, return simple feedbacks, i.e. the appreciation of a community of users given in terms of positive and negative reception.

This work faces a more complex task: the *identification of the prevalent emotions* associated by the users of a Social Network to (broadly speaking) resources. In particular our approach to Sentiment Analysis presents two main novelties. First, the extraction of the emotional content is driven by a *taxonomy of emotional concepts*, resulting in something richer than a polarized appreciation. Second, we focus on very essential form of textual resources: *collections of tags*. In artistic domains, where resources represent artworks, movies, books, users often tag resources to supply *concise reviews*. Thus, it is possible to leave aside the complications due to text analysis (often aimed at identifying the keywords inside text) and focus on the emotional analysis.

Our goal is providing as output a set of specific emotional concepts. These concepts emerge as the most significant ones for capturing the user's emotions toward a specific resource. Each concept is enriched with a score that express the emotion's strength. Our application (*ArsEmotica*) does this by exploiting and combining computational lexicons and tools that are already available. The information sources we rely upon are the tags by which users annotated resources, that are available through the ArsMeteo platform (http://www.arsmeteo.org [1]), and an ontology of emotions, obtained by enriching the one proposed in [7] with over four hundred Italian words having an affective value.

Of course, not all tags have an affective value. In order to identify the ones having an emotional content, we integrated various Semantic Web and NLP tools. In particular, we integrated the use of multi-lingual and Italian computational lexicons [11] and affective lexicons (WordNet-Affect [16] and SentiWordNet [2]) with the aim of identifying (i) words that directly refer to emotional concepts (e.g. *paura*, fear), (ii) words that indirectly deliver an emotional content (e.g. *infinito*, infinite). However, since in the latter case, SentiWordNet does not associate to the term at issue an emotion but only distinguishes between objective/neutral terms and terms with an affective impact, we decided to involve the users of the Social Network in the definition of the tags emotional contents. Therefore, we designed a simple interface by which users can indicate for an artwork and a tag, that indirectly delivers emotional contents, which emotions it raises and score them. After this analysis, it is possible to apply algorithms for the identification of the prevalent emotions to the tagged resources. The algorithm that we have implemented is described in [7] and exploits the ontology of emotions as well as the associations of tags to emotional concepts to compute its output.

The paper begins with a brief overview of the background. Section 3 describes our proposal. Section 4 reports the case study, followed by a scenario that outlines prospected applications. Final remarks end the paper.

2 Background

The research, that we have been carrying on, strictly embraces or is in some way related to different technologies and topics, including the Social Semantic Web (folksonomies and ontologies), classical computational lexicons (MultiWordNet [11], WordNetAffect [16], WordNet 3.0, SentiWordNet 3.0 [2]), studies on emotions in social, psychological, and computational contexts, and sentiment analysis. Let us briefly overview the background.

One of the key problems in the Web is how to index resources so as to efficiently and effectively retrieve them. Ontologies and folksonomies are two ways for indexing resources: the former are to be designed by knowledge engineers, while the latter are spontaneously produced by the *tagging* activity of the members of a community. Tagging is one of the ways offered to users by the so called Web 2.0 to become actively involved into the web experience, and amounts to attaching freely chosen labels to resources. Often such labels are used to categorize the resources, but they are also used to express reception, opinions, feelings. Our choice to work on tags rather than on text is due to many reasons. In our view, in many applications, users use tags to supply concise reviews. In these cases, it is possible to leave aside the complications due to text analysis (often aimed at identifying the key-words inside text) and focus on the emotional analysis. This happens mainly for artistic domains, where resources represent artworks, movies, music, books. Nevertheless, a technique like the one that we propose can easily be extended to the analysis of the latent emotions of a text, given proper linguistic tools, because it is orthogonal

from the way in which key words are identified. Moreover, when the user's feedback is supplied in a language like Italian, for which there are few affective lexical resources that are freely available for the development of new application softwares, tags can be the only processable information source that can be used.

Whatever they represent, tags are used as meta-data on top of which it is possible to devise many algorithms for navigation and retrieval. More precisely, these algorithms exploit the folksonomy, arising from tags, using it as an open, distributed and social classification system, though with a *flat* structure. This final remark suggests one of most challenging tasks that are currently being studied in the field, which is the identification of relationships that tie the folksonomy terms with one another, and possibly capture a machine-processable semantics. There are many attempts to reconcile folksonomies with ontologies, for instance by inducing ontologies from folksonomies or by matching terms in some way. It is out of the scope of this paper to get into the details. We would just like to mention the survey by Dotsika [5], together with a previous work of the authors [3], where the association was done based on the outcomes returned by a search engine. The reason for these interests is that, indeed, the Social Web aims at developing applications that combine the ease of use, which is typical of its platforms, with the advantages deriving from a formal semantics, i.e. interoperability, data/service integration, personalization, better recommendation and retrieval performances [14].

In this context, the identification of the feelings of a community or of its single members is receiving an increasing attention, as an indicator of the appraisal of topics, people, situations, resources, trends. Hence the development of opinion Mining [9], of ontologies of emotions, like the one we started out work from [6], and of W3C markup language proposals However, there are still few applications that use the most advanced results in Semantic Web technology to deal with emotions and most of the approaches use ontologies where emotions are individual isolated units (e.g. WordAffect). Such considerations motivated our focus on the ontology of emotions in [6], an OWL ontology where emotions are structured and organized in levels, trying to integrate the results of the most recent psychological models. As we will see in the next section, OntoEmotion provided a good starting point to explore an ontology-driven approach to Sentiment Analysis, where tags (and then tagged resources) are related to emotions. Such approach, to the best of our knowledge, is original w.r.t. previous work on Sentiment Analysis and allows to extract from tags affective information which is richer than a polarized appreciation.

The study of emotions is particularly relevant in artistic domains where it is a common experience for users to share the feelings kindled by an artwork (be it a painting, a video or some music) with friends. As a consequence, a software which, starting from tags freely associated to resources, can extract a rich emotional semantics, could find many interesting applications. For instance, by creating new, emotion-driven navigations of the items shown by on-line museum portals [18, 4], or by creating apps for portable devices, on top of which creating new ways of participating to art exhibitions. Moreover, most of current portals allow users to express their appreciation on artworks by means of simple star-rating systems. Our

ontology-driven approach could provide users with new means for expressing and sharing emotions stirred by artworks.

3 ArsEmotica

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

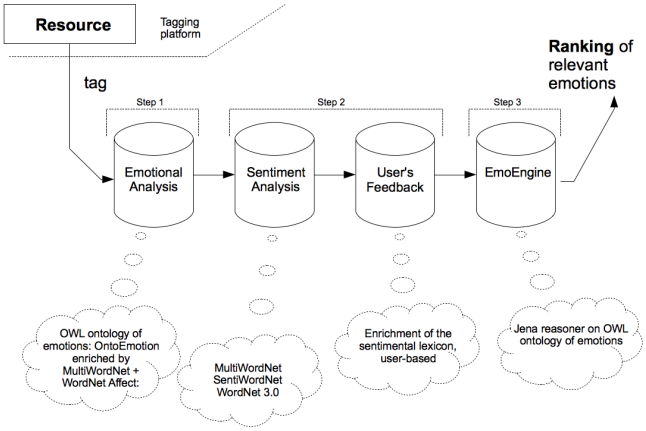

Fig. 1 ArsEmotica overall architecture.

steps that characterize the computation after the pre-processing:

Step 1: Checking tags against the ontology of emotions. This step checks whether a tag belongs to the ontology of emotions. Tags belonging to the ontology are immediately classified as "emotional".
Step 2: Checking tags with SentiWordNet. Tags that do not correspond to terms in the ontology are further analyzed by means of *SentiWordNet*, in order to distinguish *objective* tags, which do not bear an emotional meaning, from *subjective* and, therefore, affective tags. The latter will be the only ones presented to the user in order to get a feedback on which emotional concept they deliver.
Step 3: Ranking of Emotions. Based on data collected in the previous steps, the tool ranks the emotions associated by the users to the resource.

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

3.1 The ontology of emotions and the Italian emotional words

The first step checks if the tags of a given resource are "emotion-denoting" words directly referring to some emotional categories of the ontology. Our starting point was *OntoEmotion*, an emotional ontology developer at Universidad Complutense de Madrid [6] that met our requirement to have a taxonomic structure, mirroring well-founded psychological models of emotions, and that was implemented by using semantic web technologies. The ontology is written in OWL and structures emotional categories in a taxonomy that includes 87 emotional *concepts*. The basic emotions are *Sadness*, *Happiness*, *Surprise*, *Fear* and *Anger* and the taxonomic structure basically refers to the psychological model by Parrot [10], adapted to these five basic emotions, and integrated with emotions which appear in other well-established models. OntoEmotions has been conceived for categorizing emotion-denoting words. Classes corresponding to the emotional concepts were originally populated by about 250 instances, consisting in emotion-denoting words of English and Spanish. The ontology has two root concepts: *Emotion* and *Word*. *Emotion* is the root for all the emotional concepts. *Word* is the root for the emotion-denoting words, i.e. the words which each language provides for denoting emotions, and originally had two subclasses: *EnglishWord* and *SpanishWord*. Each instance of these two concepts has two parents: one is a concept from the *Emotion* hierarchy (the type of emotion denoted by the word), while the other is a concept from the *Word* hierarchy (e.g. the language of the word). For instance, the word *rage* is both an instance of the concept *Fury*, and an instance of the concept *EnglishWord*, which means that *rage* is an English word for denoting fury.

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

WordNet is a lexical database, in which nouns, verbs, adjectives and adverbs (lemmas) are organized into sets of synonyms (synsets), representing lexical concepts. After choosing the representative Italian emotional words for each concepts, such words were used as entry lemmas for querying the lexical database. The result for a word is a synset, representing the 'senses' of that word, and are labeled by MultiWordNet unique synset identifiers. Each synset was then processed by using WordNet-Affect: when a synset is annotated as representing affective information, then all the synonyms belonging to that synset are imported in the ontology as relevant Italian emotion-denoting words. This allowed us to automatically enrich the ontology with synonyms of the representative emotional words, but also to filter out synsets which do not convey affective information. Let us see an example. When we query the MultiWordNet database with the italian word *panico* (noun, representative for the emotion *Panic*), only two of the three resulting synsets are affective (WordNet senses n#10337390 and n#05591377). In particular, the third not affec-

tive synset refers to the sense of the word "panico" described by the following gloss: [coarse drought-resistant annual grass grown for grain, hay and forage in Europe and Asia and chiefly for forage and hay in United States]. Thanks to our affective filter we can exclude words belonging to that synset (Setaria_italica, pabbio_coltivato) when populating the concept *Panic* of our ontology.

The resulting ontology contains more than 450 Italian words referring to the 87 emotional categories of OntoEmotion. In order to keep trace in the ontology of the synonymy among words belonging to a same synset, we have defined the OWL object property *hasSynonym*.

ArsEmotica uses the enhanced ontology for checking if a tag describing a resource *directly* refer to some emotional category (Emotional Analysis). If yes, the tag is immediately classified as "emotional. The information collected during this phase is stored as a set of triples having the form: (t,e,s), meaning that tag t is related to emotion e with a strength value s. The range of the score s is $[0,100]$. When a tag is an instance of an emotional concept, the strength will be 100. So, for example, since the word "affanno" (breathlessness) is an instance of "anxiety", the corresponding triple will be: ("affanno", "Ansia", 100). The analysis of the tags that starts with Step 1 produces, in the end, a table that reports, for each tag and for each emotional concept, the associated score.

3.2 Sentiment analysis and User Feedback

The previous analysis identifies a set of tags as directly bearing an affective meaning. However, other tags can potentially convey affective meaning and *indirectly* refer to emotional categories of the ontology. As observed in [17], some words can be emotional for someone due to her individual story. In other cases the affective power is part of the collective imagination (e.g. words like "war"). As a consequence, it seems to be appropriate and promising to involve the community in the definition of such indirect affective meanings.

In order to minimize the effort requested to the users, before offering the tags to their judgment, we select the most promising ones by using SentiWordNet 3.0, a lexical resource for opinion mining where synsets of Princeton WordNet 3.0 are annotated according to their degree of neutrality, positiveness and negativity. Each synset s is associated the scores $Pos(s)$, $Neg(s)$ and $Obj(s)$ indicating how neutral (Obj) or affective (Pos and Neg) the terms contained in the synset are. Each score ranges in [0.0,1.0] and their sum is 1.0 for each synset. Sentiment analysis is performed by applying the steps sketched in Figure 2 to each tag that does not belong to the ontology of emotions. Since SentiWordNet was created for the English language, we needed to use MultiWordNet to align the Italian lemmas corresponding to the English ones. Moreover since SentiWordNet annotates a newer version of Princeton WordNet (3.0) with respect to the version MultiWordNet is based on (1.6), we have to query such newer lexical database. As an example, let us consider the Italian word *infinito*. MultiWordNet returns many synsets denoted by the iden-

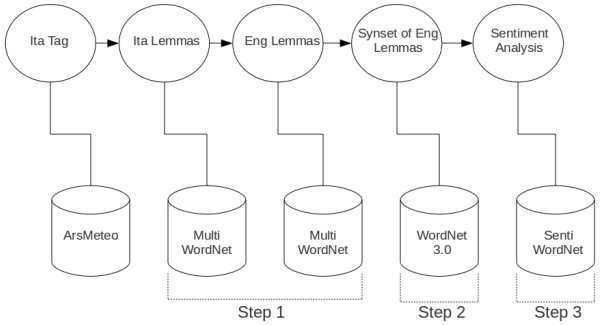

Fig. 2 Steps of the sentiment analysis on tags.

tifiers *n#04767390*, *a#00955446*, *a#00877590*. Each of them has a corresponding English lemma. In particular, *n#04767390* refers to the English word "infinitive"; *a#00955446* to "infinite"; while the last synset refers to "inexhaustible". We now use these three English words as entry lemmas for querying WordNet 3.0. The result is the list of all their possible meanings: (*infinite: 00028651*): the uninflected form of the verb; (*infinite: 01007354*): having no limits or boundaries in time or space or extent or magnitude; (*infinite: 01008745*): of verbs, having neither person nor number nor mood (as a participle or gerund or infinitive); (*infinitive: 00301951*): too numerous to be counted; (*inexhaustible: 00005718*): incapable of being entirely consumed or used up; (*inexhaustible: 01008289*): that cannot be entirely consumed or used up.

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

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

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

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

3.3 Getting the predominant emotions

Once the analysis of the tags associated to a resource is finished, during the last step ArsEmotica ranks the emotions associated by the users to the resource and computes the prevalent emotion. This is done with the help of the Jena Reasoner applied to the triples resulting from the previous analysis steps. The implemented algorithm relies on the taxonomic structure of the ontology and is inspired to the one in [7], where an analysis is performed to emotionally mark up a sentence by analyzing the words that compose it. Intuitively, the algorithm allows to select the *most specific emotion* which represents the affective information related to the artwork from the probability (score) that each of its tag has of indicating different emotions. The basic steps are: (1) processing the emotional concepts appearing in the triples (t,e,s), so as to identify also those emotions in the ontology that are related of the ones appearing in the triples. The identified emotional concepts can be organized into *layers* by following the parent-child relationship (in this phase a Jena Reasoner has been applied to the collected triples); (2) starting from the leaves and moving upward towards the root, compose and propagate the scores.

4 Case Study

The ArsEmotica prototype (developed in Java and NetBeans) was tested against a corpus of tagged multimedia artworks from the web platform *ArsMeteo* [1]. The test involved the analysis of 47 sample artworks and of 842 tags.[1] The prototype strictly implements the architecture designed in Figure 1, so it is characterized by three stages, that can be enacted by pressing in sequence the buttons *Tag to Emo*, *Tag to Senti* and *Emo Engine* (see Figure 3 for a screenshot of the interface). Given a tagged arkwork, by pressing the *Tag to Emo* button the user queries the OWL ontology of emotions; the result is a set of tags referring to emotions. For instance, by applying the emotional analysis to the artwork "Forse un giorno arriveranno al mare" by Claudio Guasti (1997) (Figure 3), three tags are identified as emotion-denoting words. In particular, the tag "desiderio" refers to emotion *Lust*, while "stupore" refers to *Amazement* and "paura" to *Fear*, *Trepidation*, etc.. The second stage works on set of tags with no correspondence in the ontology. By pressing the *Tag To Senti* button,

[1] See http://www.giorgiovaccarino.it/mostre/emozioni.html, click on images for details.

the affective potential of these tags is studied thanks to SentiWordNet. This lexicon classifies terms as subjective or objective, returning a positivity score, a negativity score, from which it is possible to compute the objectivity score. We consider as tags possibly conveying an emotional meaning all those whose objectivity score is less than the threshold 0.5. For these, we ask the user to give a classification driven by the ontology of emotions, as explained in Section 3.2. The third stage amounts to compute the predominant emotions and can be enacted by pressing the *Emo Engine* button. We are currently completing the prototype in order to allow the formatting of the result according to the standard markup language *Emotion ML* by W3C [13].

First experiments show that, by applying the emotional analysis only, i.e. just by checking whether the tags belong to the ontology of emotions (Step 1) ArsEmotica returns an emotional meaning for 14 images out of 47.

In order to evaluate the match between these outcomes and the perceptions of the community we mean to integrate in the ArsMeteo portal a feedback mechanism by which users can express their (dis)agreement with the proposed prominent emotion as computed based on the meaning of tags. The collected data will allow us to test the generality of the emerging emotion w.r.t. the feelings of the community.

Fig. 3 Emotion -denoting tags for an artwork

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

5 Conclusion and Future Work

This work presents a software solution that, by combining lexicons and libraries that are already available, allows both the population of an ontology of emotions (based on [6]) with Italian emotion-bearing terms and the extraction of the prevalent emotions from the set of tags associated to a resource. The extracted information is richer than a polarized appreciation, as instead usually done by sentiment analysis. Preliminary tests were done on a subset of the artworks from the art portal ArsMeteo.

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

The proposed solution can be refined in many ways. For what concerns the preprocessing of tags, we intend to improve the current prototype by applying stemming and word-similarity algorithms. For instance, in Italian, adjectives are declined in many ways, depending whether they refer to males or females, singular or plural. Stemming and lemmatization algorithms would help reduce the noise due to this variability. Word similarity could, instead, help to find relations among concepts that are not detected by the studied computational lexicons.

An extensive evaluation of our approach with real users of the ArsMeteo community will be part of the future work. Moreover, we intend to work on an effective way for motivating users to annotate tags having an indirect affective meaning by means of emotional concepts from our ontology. In order to face this issue, one promising direction could be to rely on the Game With A Purpose [19] paradigm and to develop a proper game in which users, as a side effect of playing, perform the task of associating emotional concepts to tag-words. This is on the line of recent approaches which face the challenge of increasing the user involvement in building the Semantic Web [15]. An alternative could be to integrate in ArsEmotica the use of automatic techniques, e.g. the one proposed in [3], for identifying the association of terms having an emotional value (that is recognized by the sentiment analysis step) with the proper ontological concepts.

Moreover, it is known that the emotional semantics may vary depending on the context. Psychological theories concerning emotions, that tie perception to context, could be integrated in Arsemotica to refine the outcome [8].

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

Acknowledgements The authors thank all the persons who supported the work with comments and contributions. In particular, Andrea Bolioli and Flavio Portis, Giorgio Vaccarino and the *Associazione Culturale ArsMeteo* for providing the data corpus, Fondazione Bruno Kessler and ISTI

CNR for supplying the lexicons, the web agency *EasyBit*, which supplied the tagging platform of Arsmeteo, the NIL group at the Complutense University of Madrid, which supplied *OntoEmotion*.

References

1. E. Acotto, M. Baldoni, C. Baroglio, V. Patti, F. Portis, and G. Vaccarino. Arsmeteo: artworks and tags floating over the planet art. In *Proc. of the 20th ACM conference on Hypertext and hypermedia*, HT '09, pages 331–332. ACM, 2009.
2. A. E. S. Baccianella and F. Sebastiani. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proc. of the 7th conference on International Language Resources and Evaluation (LREC'10)*. ELRA, May 2010.
3. M. Baldoni, C. Baroglio, A. Horváth, V. Patti, F. Portis, M. Avilia, and P. Grillo. Folksonomies meet ontologies in arsmeteo: from social descriptions of artifacts to emotional concepts. In S. Borgo and L. Lesmo, editors, *Formal Ontologies Meet Industry, FOMI 2008*, pages 132–143. IOS Press, 2008.
4. S. Chan. Tagging and searching: serendipity and museum collection databases. In *Museums and the Web 2007*, pages 87–99, 2007.
5. F. Dotsika. Uniting formal and informal descriptive power: Reconciling ontologies with folksonomies. *International Journal of Information Management*, pages 407–415, October 2009.
6. V. Francisco, P. Gervás, and F. Peinado. Ontological reasoning to configure emotional voice synthesis. In *Proc. of Web Reasoning and Rule Systems, RR 2007*, volume 4524 of *LNCS*, pages 88–102. Springer, 2007.
7. V. Francisco, F. Peinado, R. Hervás, and P. Gervás. *Semantic Web Approaches to the Extraction and Representation of Emotions in Texts*. NOVA Publishers, 2010.
8. J. R. Ogorek. Normative picture categorization: Defining affective space in response to pictorial stimuli. In *Proc. of REU'05*, 2005.
9. B. Pang and L. Lee. *Opinion Mining and Sentiment Analysis (Foundations and Trends(R) in Information Retrieval)*. Now Publishers Inc., 2008.
10. W. Parrot. *Emotions in Social Psychology*. Psychology Press, Philadelphia, 2001.
11. E. Pianta, L. Bentivogli, and C. Girardi. Multiwordnet: developing an aligned multilingual database. In *Proc. of the First International Conference on Global WordNet*, January 2002.
12. R. W. Picard. Affective computing. Technical Report 321, MIT, 1995.
13. M. Schröder, P. Baggia, F. Burkhardt, A. Oltramari, C. Pelachaud, C. Peter, and E. Zovato. Emotion Markup Language (EmotionML) 1.0. W3C Working Draft, 2010. http://www.w3.org/TR/emotionml/.
14. N. Shadbolt, T. Berners-Lee, and W. Hall. The semantic web revisited. *IEEE Intelligent Systems*, 21(3):96–101, 2006.
15. K. Siorpaes and M. Hepp. Games with a purpose for the semantic web. *IEEE Intelligent Systems*, 23:50–60, 2008.
16. C. Strapparava and A. Valitutti. WordNet-Affect: an affective extension of WordNet. In *Proc. of LREC*, volume 4, pages 1083–1086, 2004.
17. C. Strapparava, A. Valitutti, and O. Stock. The affective weight of lexicon. In *Proc. of LREC*, pages 1–83, 2006.
18. J. Trant and B. Wyman. Investigating social tagging and folksonomy in art museums with steve.museum. In *Proc. of the Collaborative Web Tagging Workshop (WWW'06)*, 2006.
19. L. von Ahn. Games with a purpose. *IEEE Computer*, 39(6):92–94, 2006.