Onyx: Describing Emotions on the Web of Data

J. Fernando Sànchez-Rada and Carlos A. Iglesias

Intelligent Systems Group
Telematic Systems Engineering Deparment
Technical University of Madrid (UPM)
Email: jfernando@dit.upm.es
cif@dit.upm.es

Abstract. Textual emotion analysis is a new field whose aim is to detect emotions in user generated content. It complements Sentiment Analysis in the characterization of users subjective opinions and feelings. Nevertheless, there is a lack of available lexical and semantic emotion resources that could foster the development of emotion analysis services. Some of the barriers for developing such resources are the diversity of emotion theories and the absence of a vocabulary to express emotion characteristics. This article presents a semantic vocabulary, called Onyx, intended to provide support to represent emotion characteristics in lexical resources and emotion analysis services. Onyx follows the Linked Data principles as it is aligned with the Provenance Ontology. It also takes a linguistic Linked Data approach: it is aligned with the Provenance Ontology, it represents lexical resources as linked data, and has been integrated with Lemon, an increasingly popular RDF model for representing lexical entries. Furthermore, it does not prescribe any emotion model and can be linked to heterogeneous emotion models expressed as Linked Data. Onyx representations can also be published using W3C EmotionML markup, based on the proposed mapping.

Keywords: ontology, emotions, emotion analysis, sentiment analysis, semantic, semantic web, linked data, provenance, emotionml, lemon

1 Introduction

From the tech-savvy to elders, our society is exponentially moving its social and professional activity to the Internet, with its myriad of services and social networks. Facebook[1] or Twitter[2] are only two of the most successful examples, producing flooding streams of user-generated data. Unluckily, quite often that information is just meant for human consumption and is only formatted to be displayed. This prevents us from automatically processing these massive streams of information to aggregate, summarize or transform them and present human users with a bigger picture. In other words, data mining techniques require machine-formatted data input.

[1] https://facebook.com
[2] https://twitter.com

In an attempt to shorten that gap, the multidisciplinary field called Sentiment Analysis or Opinion Mining was born, which aims at determining the subjectivity of human opinions. Many tools have been created to enrich or make sense out of human generated content by applying natural language processing and adding the results as annotations or tags. Whilst this solves the issue at a small scale, for each ad-hoc solution, it raises another problem: data collected by different programs presents different and sometimes incompatible formats. Linked Data introduced a lingua franca for data representation as well as a set of tools to process and share such information. Many services embraced the Linked Data concepts and are providing tools to interconnect the previously closed silos of information [28].

The Sentiment Analysis field is now evolving to determine also human emotions. An important fact about emotions is that they change the way we communicate [20]. They can be passed on just like any other information, in what some authors call emotional contagion [7]. That is a phenomenon that is clearly visible in social networks. Most of them offer a public API that makes studying the networks and information flow relatively easy. For this very reason social network analysis is an active field [18], with Emotion Mining as one of its components.

Social networks aside, another field of application of Emotion Analysis is Affective computing. There are a variety of systems whose only human-machine communication is purely text-based. These systems are often referred to as dialog systems (e.g. Q&A systems). Such systems can use the emotive information to change their behaviour and responses [20].

On the other hand, the rise of services like microblogging will inevitably lead to services that exchange and use affective information. Some social sites are already using emotions natively, giving their users the chance to share emotions or use them in queries. Facebook, for instance, recently updated the way its users can share personal statuses.

These sites have started making heavy use [3] of formats like RDFa [4] or Microformats as a bridge between web pages for human consumption and Linked Data. This made it possible to provide a better user experience and better search results despite the big amount of information these networks contain.

Combining the objective facts already published as Linked Data with subjective opinions extracted using Sentiment and Emotion analysis techniques can enable a wide array of new services. Unfortunately, there is not yet any widely accepted Linked Data representation for emotions. This paper aims at bridging this gap with the definition of a new vocabulary, Onyx.

This paper is structured as follows: Section 2 introduces the technologies that Onyx is based upon, as well as the challenges related to Emotion Analysis and creating a standard model for emotions, including a succinct overview of the formats currently in use; Section 3 covers the Onyx ontology in detail and several use cases for this ontology; Section 4 presents the results of our evaluation of the Ontology, focusing on the coverage of current formats like EmotionML; Section 5 completes this paper with our conclusions and future work.

2 Enabling Technologies

2.1 Models for Emotions and Sentiment Analysis

To work with Emotions and reason about them, we first need to have a solid understanding and model of emotions. This, however, turns out to be a rather complex task.It is comprised of two main components: modelling (including categorisation) and representation.

There are several models for emotions, ranging from the most simplistic and ancient that come from Chinese philosophers to the most modern theories that refine and expand older models [11, 22]. The literature on the topic is vast, and it is out of the scope of this paper to reproduce it. The recent work by Cambria et al. [10] contains a comprehensive state of the art on the topic, as well as an introduction to a novel model, The Hourglass of Emotions, inspired by Plutchik's studies [21]. Plutchik's model has been extensively used [8, 9] in the area of Sentiment Analysis and Affective Computing, relating all the different emotions to each other in what is called the rose of emotions.

Other models cover affects in general, which include Emotions as part of them. One of them is the work done by Strapparava and Valitutti in WordNet-Affect [27]. It comprises more than 300 affects, many of which are considered emotions. What makes this categorization interesting is that it effectively provides a taxonomy of emotions. It both gives information about relationship between emotions and makes it possible to to decide the level of granularity of the emotions expressed.

Despite all, there does not seem to be a universally accepted model for emotions [26]. This complicates the task of representing emotions. In a discussion regarding Emotion Markup Language (EmotionML), Schroder et al. pose that *any attempt to propose a standard way of representing emotions for technological contexts seems doomed to fail* [25]. Instead they claim that the markup should offer users choice of representation, including the option to specify the affective state that is being labelled, different emotional dimensions and appraisal scales. The level of intensity completes their definition of an affect in their proposal.

EmotionML [6] is one of the most notable general-purpose emotion annotation and representation languages. It was born from the efforts made for Emotion Annotation and Representation Language (EARL) [1, 26] by Human-Machine Interaction Network on Emotion (HUMAINE) EARL originally included 48 emotions divided into 10 different categories. EmotionML offers twelve vocabularies for categories, appraisals, dimensions and action tendencies. A vocabulary is a set of possible values for any given attribute of the emotion. There is a complete description of those vocabularies and its computer-readable form available [5].

In the field of Semantic Technologies, Grassi presented Human Emotion Ontology (HEO). This ontology presents an ontology for human emotions for its use for annotating emotions in multimedia data. Another work worth mentioning is that of Hastings et al. [14] in Emotion Ontology (EMO), an ontology that tries to reconcile the discrepancies in affective phenomena terminology.

For Opinion Mining we find the Marl vocabulary [29]. Marl was designed to annotate and describe subjective opinions expressed in text. In essence, it provides the conceptual tools to annotate Opinions and results from Sentiment Analysis in an open and sensible format. However, it is focused on polarity extraction and is not capable of representing Emotions. Onyx aims to remedy this and offer a complete set of tools for any kind of Sentiment Analysis, including advanced Emotion Analysis.

Lastly, it is worth mentioning lemon, the Lexicon Model for Ontologies. As its name indicates, it is a model that supports the sharing of terminological and lexicon resources on the Semantic Web as well as their linking to the existing semantic representation provided by ontologies [16]. Onyx will be used together with lemon to annotate lexicon resources for Emotion Analysis, as will be shown in some of the examples below.

2.2 W3C's Provenance

Provenance is information about entities, activities, and people involved in producing a piece of data or thing, which can be used to form assessments about its quality, reliability or trustworthiness. The PROV Family of Documents defines a model, corresponding serializations and other supporting definitions to enable the inter-operable interchange of provenance information in heterogeneous environments such as the Web [19]. It includes a full-fledged ontology, to which Onyx is linked. The complete ontology is covered by the PROV-O Specification. However, to understand the role of Provenance in Onyx and vice versa, it is enough to understand Figure 1.

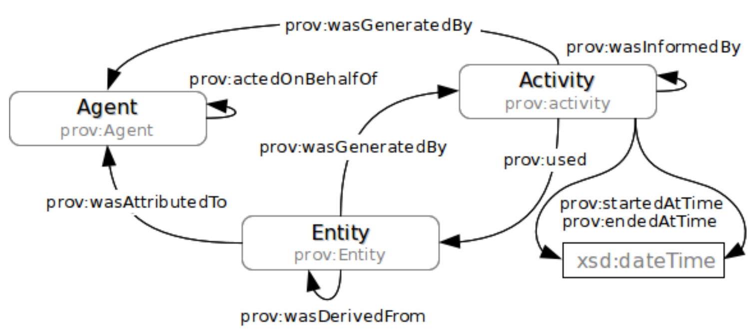

Fig. 1: Simple overview of the basic classes in the Provenance Ontology [12]

As we can see, Agents take part in Activities to transform Entities (data) into different Entities (modified data). This process can be aggregation of information, translation, adaptation, etc. In our case, this activity is an Emotion Analysis, which turns plain data into semantic emotion information.

There are many advantages to adding provenance information in Sentiment Analysis in particular as different algorithms may produce different results. By including the Provenance classes in our Emotion Mining Ontology we can not only link results with the source from which it was extracted, but also with the algorithm that produced them.

3 Onyx

Onyx is a vocabulary to represent the Emotion Analysis process and its results, as well as annotating lexical resources for Emotion Analysis. It includes all the necessary classes and properties to provide structured and meaningful Emotion Analysis results, and to connect results from different providers and applications.

At its core, the Onyx ontology has three main classes: EmotionAnalysis, EmotionSet and Emotion. In a standard Emotion Analysis, these three classes are related as follows: an EmotionAnalysis is run on a source (generally in the form of text, e.g. a status update), the result is represented as one or more EmotionSet instances that contain one or more Emotion instances.

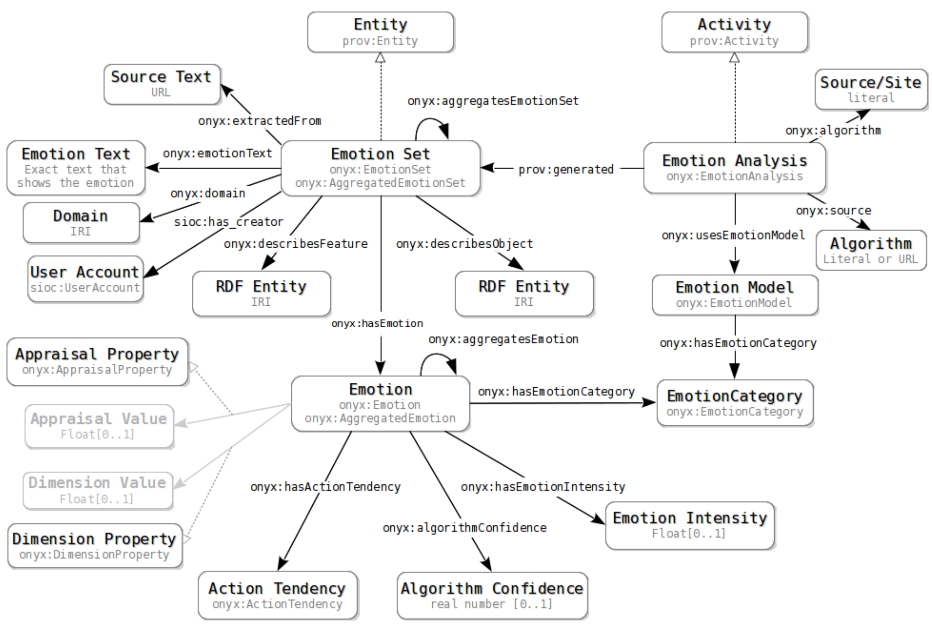

Fig. 2: Class diagram of the Onyx ontology.

The EmotionAnalysis instance contains information about: the source (e.g. dataset, website) from which the information was taken, the algorithm used, and the emotion model that was used to represent emotions. Additionally, it can make use of Provenance to specify the Agent in charge of the analysis, the resources used (e.g. dictionaries), and other useful information.

An EmotionSet contains a group of emotions found in the text or one of its parts. As such, it contains information about: the original text (extractedFrom); the exact excerpt that contains the emotion or emotions (emotionText); the person that showed the emotions (sioc:has_creator); the entity that the emotion is related to (describesObject); the concrete part of that object it refers to (describesObjectPart); the feature about that part or object that triggers the emotion (describesFeature); and, lastly, the domain detected. All this properties are straightforward, but a note should be given about the domain property. Different emotions could have different interpretations in different contexts (e.g., fear is positive when referred to a thriller, but negative when it comes to cars and safety).

When several EmotionSet instances are related, an AggregatedEmotionSet can be created that links to all of them. For instance, we could aggregate all the emotions about a Movie, or all the emotions shown by a particular user. An AggregatedEmotionSet is a subclass of EmotionSet which contains additional information about the original EmotionSet instances it aggregates.

Considering the lack of consensus on modeling and categorizing emotions, our model of emotions is very generic. In this Emotion model we include: an EmotionCategory or type of emotion (although more could be specified), through the hasEmotionCategory property (e.g. "sadness"); the emotion intensity; action tendencies (ActionTendency) related to this emotion, or actions that are triggered by the emotion; appraisals and dimensions. Appraisals and dimensions are defined as properties, whose value is a float number. On top of that generic model, we have adapted two different systems: the WordNet-Affect taxonomy, and the EmotionML vocabularies for categories, dimensions and appraisals.

WordNet-Affect [27] contains the relationships (concepts and superconcepts) of affects, among which we find emotions. We processed the list of affects and published a SKOS version of the taxonomy [24]. The taxonomy specification includes a navigable tree that contains the concepts (i.e. affect types) in it, aligned with WordNet concepts. This makes it trivial to select an affect that represents the desired emotion. Besides providing a good starting point for other ontologies, this taxonomy also serves as a base to translate between the several different ontologies in the future.

Regarding EmotionML, we have converted its vocabularies [5] the Onyx format. Using this extension we can translate EmotionML resources into Onyx for their use in the Semantic Web.

This is further developed in Section 4.

It is also possible that two separate emotions, when found simultaneously, imply a third emotion. A more complex one. For instance, "thinking of the awful things I've done makes me want to cry" might reveal sadness and disgust, which together might be interpreted as remorse. In such situation, we could add an AggregatedEmotion that represents remorse to the EmotionSet, linking it to the primary emotions with the aggregatesEmotion property.

To group all the attributes that correspond to a specific emotion model, we created the EmotionModel class. Each EmotionModel will be linked to the different categories it contains (hasEmotionCategory), the AppraisalProperty or DimensionProperty instances it introduces (through hasAppraisalProperty and hasDimensionProperty), etc.

Figure 2 shows a complete overview of all these classes, as well as all their properties.

After this introduction of the ontology, we will present several use cases for it. This should give a better understanding of the whole ontology by example. Rather than exhaustive and complex real life applications, these examples are meant as simple self-contained showcases of the capabilities of semantic Emotion Analysis using Onyx. For the sake of brevity, we will omit the prefix declaration in the examples.

Case	**N3 Representation**

An example EmotionAnalysis.	:customAnalysis
	  a onyx:EmotionAnalysis;
	  onyx:algorithm "SimpleAlgorithm";
	  onyx:usesEmotionModel wna:WNAModel.

Processing "I lost one hour today because of the strikes!!", by the user JohnDoe	:result1
	  a onyx:EmotionSet;
	  prov:wasGeneratedBy :customAnalysis;
	  sioc:has_creator [
	    sioc:UserAccount <http://blog.example.
	         com/JohnDoe>.];
	  onyx:hasEmotion [
	    onyx:hasEmotionCategory wna:anger;
	    onyx:hasEmotionIntensity :0.9];
	  onyx:emotionText "I lost one hour today
	          because of the strikes!!" ;
	  dcterms:created "2013-05-16T19:20:30+01:00"
	               ^^dcterms:W3CDTF.

Example of annotation of a lexical entry using Onyx and lemon [17].	:fifa
	  a lemon:Lexicalentry;
	  lemon:sense [
	    lemon:reference wn:synset-fear-noun-1;
	    onyx:hasEmotion
	        [oyx:hasEmotionCategory wna:fear.].
	  ];
	  lexinfo:partOfSpeech lexinfo:noun.

Table 1: Representation with Onyx

4 Evaluation

Evaluating ontologies is always a difficult task. Evaluation methodologies are highly debatable and there are no standards [13]. For the evaluation of Onyx we focused on its practical use as well as in its correctness. This means testing the adequacy of the model for existing applications as well as scenarios with several emotion models. In particular we have chosen two different test scenarios: the

Case	Query

Finding all the users that did not feel good during last New Year's Eve, and the exact emotions they felt.	`SELECT DISTINCT ?creator ?cat`
	`WHERE {`
	`  ?set onyx:hasEmotion`
	`      [ onyx:hasEmotionCategory ?cat];`
	`    dcterms:created ?date;`
	`    sioc:has_creator ?creator.`
	`  ?cat skos:broaderTransitive* wna:negative`
	`             -emotion.`
	`  FILTER( ?date >= xsd:date("2012-12-31")`
	`      ?date <= xsd:date("2013-01-01") )`
	`}`

Comparing two Emotion Mining algorithms by comparing the discrepancies in the results obtained using both.	`SELECT ?source1 ?algo1 (GROUP_CONCAT(?cat1)`
	`            as ?cats1)`
	`WHERE {`
	`  ?set1 onyx:extractedFrom ?source1.`
	`  ?analysis1 prov:generated ?set1;`
	`             onyx:algorithm ?algo1.`
	`  ?set1 onyx:hasEmotion`
	`          [ onyx:hasEmotionCategory ?cat1 ].`
	`  FILTER EXISTS{`
	`    ?set2 onyx:extractedFrom ?source1.`
	`    ?analysis2 prov:generated ?set2.`
	`    ?set2 onyx:hasEmotion`
	`          [ onyx:hasEmotionCategory ?cat2 ].`
	`    FILTER ( ?set1 != ?set2).`
	`    FILTER ( ?cat2 != ?cat1 ).`
	`  }`
	`}`
	`GROUP BY ?source1 ?algo1`
	`ORDER BY ?source1`

Table 2: Example SPARQL queries with Onyx

adaptation of a well-known Emotion Analysis tool to output Onyx, Synesketch [2,], and the translation of EmotionML resources to Onyx and vice versa.

For the EmotionML part, the evaluation process is split into two parts: transforming the EmotionML categories into a semantic format, and representing EmotionML cases with Onyx. The result of the former can be seen in [23], which has been used as namespace (emlonyx) in the translation of an EmotionML example in Table 3. The specification of EmotionML is public, including its XML schema, which eased the process of mapping it to Onyx. We have focused especially on representing EmotionML emotions in Onyx.

Synesketch is a library and application that detects emotions in English texts and can generate images that reflect those emotions. Originally written in Java, it has been unofficially ported to several programming languages (including PHP), which shows the interest of the community in this tool. The aim of the PHP port was, among others, to offer a public endpoint for emotion analysis, which later had to be taken down due to misuse. The relevance of this tool and its Open Source license were the leading factors in choosing this tool. Our approach has been to develop a proof-of-concept web service that performs Emotion Analysis using Synesketch's emotion analysis. The service can be accessed via a REST API and its results are presented in Onyx, using the RDF format.

The Synesketch library uses the big-6 emotional model, which comprises: happiness, sadness, fear, anger, disgust and surprise. Each of those emotions are present in the input text with a certain weight that ranges from 0 to 1. Additionally, it has two attributes more that correspond to the general emotional valence (positive, negative or neutral) and the general emotional weight. In other words, these attributes together show how "positive", "negative" or "neutral" the overall emotion is.

To represent the big-6 emotion category in Onyx we used EmotionML's big-6 category, which we previously mapped to Onyx. The Synesketch weight directly mapped to hasEmotionIntensity in Onyx.

However, the General Emotional valence and weight do not directly match any Onyx property or class. To solve it, we simply added an AggregatedEmotion with the PositiveEmotion, NeutralEmotion or NegativeEmotion category (as defined by WordNet-Affect) depending on the value of the valence. The general emotional weight is then the intensity of this AggregatedEmotion, just like in the other cases.

The final result is a REST service that is publicly available at our website[3].

EmotionML	**Onyx**
<emotionml	:Set1 a onyx:EmotionSet;
xmlns="http://.../emotionml"	onyx:extractedFrom "Come, there is no use
xmlns:meta="http://.../	in crying like that! said Alice to
metadata"	herself rather sharply; I advice
category-set="http://.../#	you to live off this minute!";
everyday-	onyx:hasEmotion :Emo1
categories">	onyx:hasEmotion :Emo2
<info>	:Emo1 a onyx:Emotion;
<classifiers:classifier	onyx:hasEmotionCategory
classifiers:name="GMM"/>	emlonyx:disgust;
</info>	onyx:hasEmotionIntensity 0.82;
<emotion>	onyx:hasEmotionText "Come, there's no use in
<category name="Disgust" value	crying like that!"
="0.82"/>	:Emo2 a onyx:Emotion;
'Come, there is no use in	onyx:hasEmotionCategory emlonyx:anger;
crying like that!'	onyx:hasEmotionIntensity 0.57;
</emotion>	onyx:hasEmotionText "I advice you to leave
said Alice to herself rather	off this minute!"
sharply;	:Analysis1 a onyx:EmotionAnalysis;
<emotion>	onyx:algorithm "GMM";
<category name="Anger" value="	onyx:usesEmotionModel emlonyx:everyday-
0.57"/>	categories;
'I advise you to leave off	prov:generated Set1.
this minute!'	
</emotion>	
</emotionml>	

Table 3: Representation of EmotionML with Onyx

[3] http://demos.gsi.dit.upm.es/onyxemote/

5 Conclusions and Future Work

With this work we have introduced an option to represent Emotions that takes advantage of the work conducted in the field of Semantic Web. This ontology presents characteristics that are particularly beneficial for any process of Emotion Analysis. Onyx provides a structured format for Emotion Analysis. It addresses the problem of supporting heterogeneous categories of emotions, and new categories and features can be added, using the recommended taxonomy to link them and retain compatibility.

We also presented how Onyx would be used in several scenarios. Furthermore, we adapted some of the existent resources and services to Onyx, making them publicly available.

Although this paper is focused on Emotion Analysis, emotive information can also be directly provided by users. Either given explicitly or extracted via an automated process (Emotion Analysis), the information they represent is the same. A single ontology should thus cover both scenarios. This is possible with Onyx, as we demonstrate in this paper.

We would like to note that our proposal is compatible with EMO, since EMO can be easily mapped to Onyx using the property usesEmotionModel. The situation is similar with the proposal of Lopez et al. [15], which focuses on emotions instead of affects in general. The integration with HEO will be investigated. Onyx's and HEO's Emotion classes are very similar overall, but follow different approaches in several aspects.

With all this in mind, we consider that using Onyx to represent Emotion Mining results is highly beneficial.

As part of the future plans for Onyx, it will be actively used in the EUROSENTIMENT[4] project, whose aims is to create a language resource pool for Sentiment Analysis. Together with Marl [29] and Lemon [17], they will be the standard formats for representation of lexicons and results. Therefore all the services provided in the frame of the EUROSENTIMENT project will export emotional information using Onyx. Marl has already been integrated in NIF 2.0[5] to represent opinions, and efforts will be made to integrate Onyx as well for emotions.

There is also room for experimentation emotion composition and inference using tools such as SPIN[6]. It is possible to infer complex emotions whenever other simple emotions are present, and vice versa. The same techniques could be used to work with different emotion models

Finally, our research group will use the integration with EmotionML to develop intelligent personal agents that benefit from the potential of the Semantic Web.

---

[4] `http://eurosentiment.eu`

[5] `http://persistence.uni-leipzig.org/nlp2rdf/ontologies/nif-core/nif-core.html`

[6] `http://spinrdf.org/`

6 Acknowledgements

This work was partially funded by the EUROSENTIMENT FP7 Project (Grant Agreement no: 296277)

References

1. Humaine Emotion Annotation and Representation Language (EARL): Proposal., June 2006, `http://emotion-research.net/projects/humaine/earl/proposal#Dialects`.
2. Synesketch: Free Open-Source Software for Textual Emotion Recognition and Visualization, June 2006, `http://emotion-research.net/projects/humaine/earl/proposal#Dialects`.
3. Facebook Open Graph API, June 2013, `http://developers.facebook.com/docs/opengraph/`.
4. Ben Adida, Mark Birbeck, Shane McCarron, and Steven Pemberton. RDFa in XHTML: Syntax and processing. *Recommendation, W3C*, 2008.
5. Kazuyuki Ashimura, Paolo Baggia, Felix Burkhardt, Alessandro Oltramari, Christian Peter, and Enrico Zovato. EmotionML vocabularies, May 2012, `http://www.w3.org/TR/2012/NOTE-emotion-voc-20120510/`.
6. Paolo Baggia, Felix Burkhardt, Catherine Pelachaud, Christian Peter, and Enrico Zovato. Emotion Markup Language (EmotionML) 1.0, April 2013,`http://www.w3.org/TR/emotionml/`.
7. Sigal G. Barsade. The ripple effect: Emotional contagion and its influence on group behavior. *Administrative Science Quarterly*, 47(4):644–675, 2002.
8. Damian Borth, Tao Chen, Rongrong Ji, and Shih-Fu Chang. SentiBank: large-scale ontology and classifiers for detecting sentiment and emotions in visual content. In *Proceedings of the 21st ACM international conference on Multimedia*, MM '13, pages 459–460, New York, NY, USA, 2013. ACM.
9. Erik Cambria, Catherine Havasi, and Amir Hussain. SenticNet 2: A semantic and affective resource for opinion mining and sentiment analysis. In *FLAIRS Conference*, pages 202–207, 2012.
10. Erik Cambria, Andrew Livingstone, and Amir Hussain. The hourglass of emotions. In *Cognitive Behavioural Systems*, pages 144–157. Springer, 2012.
11. Paul Ekman. Basic emotions. *Handbook of cognition and emotion*, 98:45–60, 1999.
12. Paul Groth and Luc Moreau. Prov-O W3C Recommendation, April 2013, `http://www.w3.org/TR/prov-o/`.
13. Asunción Gómez-Pérez. Evaluation of ontologies. *International Journal of Intelligent Systems*, 16(3):391–409, 2001.
14. Janna Hastings, Werner Ceusters, Barry Smith, and Kevin Mulligan. Dispositions and processes in the emotion ontology. In *ICBO*, 2011.
15. Juan Miguel López, Rosa Gil, Roberto García, Idoia Cearreta, and Nestor Garay. Towards an ontology for describing emotions. In *Emerging Technologies and Information Systems for the Knowledge Society*, pages 96–104. Springer, 2008.
16. John McCrae, Dennis Spohr, and Philipp Cimiano. Linking lexical resources and ontologies on the semantic web with lemon. In *The Semantic Web: Research and Applications*, pages 245–259. Springer, 2011.

17. John McCrae, Dennis Spohr, and Philipp Cimiano. Linking lexical resources and ontologies on the semantic web with Lemon. In Grigoris Antoniou, Marko Grobelnik, Elena Simperl, Bijan Parsia, Dimitris Plexousakis, Pieter Leenheer, and Jeff Pan, editors, *The Semantic Web: Research and Applications*, volume 6643 of *Lecture Notes in Computer Science*, pages 245–259. Springer Berlin Heidelberg, 2011.
18. Alan Mislove, Massimiliano Marcon, Krishna P. Gummadi, Peter Druschel, and Bobby Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, IMC '07, pages 29–42, New York, NY, USA, 2007. ACM.
19. Luc Moreau, Ben Clifford, Juliana Freire, Joe Futrelle, Yolanda Gil, Paul Groth, Natalia Kwasnikowska, Simon Miles, Paolo Missier, Jim Myers, Beth Plale, Yogesh Simmhan, Eric Stephan, and Jan Van den Bussche. The open provenance model core specification (v1.1). *Future Generation Computer Systems*, 27(6):743 – 756, 2011.
20. Bo Pang and Lillian Lee. Opinion mining and sentiment analysis. *Foundations and trends in information retrieval*, 2(1-2):1–135, 2008.
21. Robert Plutchik. *Emotion: A psychoevolutionary synthesis.* Harper & Row New York, 1980.
22. Jesse J Prinz. Gut reactions: A perceptual theory of emotion. 2004.
23. J. Fernando Sànchez-Rada and Carlos A. Iglesias. EmotionML categories for Onyx, July 2013, `http://gsi.dit.upm.es/ontologies/onyx/emotionml`.
24. J. Fernando Sànchez-Rada and Carlos A. Iglesias. WordNet-Affect SKOS Taxonomy, May 2013,`http://gsi.dit.upm.es/ontologies/wnaffect/`.
25. Marc Schröder, Laurence Devillers, Kostas Karpouzis, Jean-Claude Martin, Catherine Pelachaud, Christian Peter, Hannes Pirker, Björn Schuller, Jianhua Tao, and Ian Wilson. What should a generic emotion markup language be able to represent? In *Affective Computing and Intelligent Interaction*, pages 440–451. Springer, 2007.
26. Marc Schröder, Hannes Pirker, and Myriam Lamolle. First suggestions for an emotion annotation and representation language. In *Proceedings of LREC*, volume 6, pages 88–92. Citeseer, 2006.
27. Carlo Strapparava and Alessandro Valitutti. Wordnet-affect: an affective extension of wordnet. In *Proceedings of LREC*, volume 4, pages 1083–1086, 2004.
28. Giovanni Tummarello, Renaud Delbru, and Eyal Oren. Sindice. com: Weaving the open linked data. In *The Semantic Web*, pages 552–565. Springer, 2007.
29. Adam Westerski, Carlos A. Iglesias, and Fernando Tapia Rico. Linked opinions: Describing sentiments on the structured web of data. In *4th international workshop Social Data on the Web (SDoW2011)*, Bonn, Germany, October 2011.