Onyx: Describing Emotions on the Web of Data

J. Fernando Sánchez-Rada and Carlos A. Iglesias

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

Abstract. There are several different standardised and widespread formats to represent emotions. However, there is no standard semantic model yet. This paper presents a new ontology, called Onyx, that aims to become such a standard while adding concepts from the latest Semantic Web models. In particular, the ontology focuses on the representation of Emotion Analysis results. But the model is abstract and inherits from previous standards and formats. It can thus be used as a reference representation of emotions in any future application or ontology. To prove this, we have translated resources from EmotionML representation to Onyx. We also present several ways in which developers could benefit from using this ontology instead of an ad-hoc representation. Our ultimate goal is to foster the use of semantic technologies for Emotion Analysis while following the Linked Data ideals.

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. Unfortunately, despite being shared via online services, quite often that information is just meant for human consumption and its format merely allows machines to display it. This prevents us from automatically processing these massive streams of information to aggregate, summarise or transform them to present human users with a bigger picture. In other words, data mining techniques require machine-formatted data input.

In an attempt to shorten that gap, 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

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

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 [26].

The multidisciplary field called Sentiment Analysis or Opinion Mining aims at determining the subjectivity of human opinions. This field is now evolving to determine also the human emotions. Thus, the combination of objective facts exposed as Linked Data with subjective opinions extracted using Sentiment and Emotion analysis techniques can enable a wide array of new services. Nevertheless, 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 summarises the main factors that drove the creation of Onyx; Section 3 introduces the technologies that Onyx is based upon, as well as the challenges in Emotion Analysis and creating a standard model for emotions, including an overview of the formats currently in use; Section 4 covers the Onyx ontology in detail and several use cases for this ontology; Section 5 presents the results of our evaluation of the Ontology, focusing on the coverage of current formats like EmotionML; Section 6 completes this paper with our conclusions and future work.

2 Motivation

Semantic technologies are evolving rapidly, which translates into more and more services that start benefiting from them to a certain extent. Formats like RDFa [1] or Microformats serve as a bridge between web pages for human consumption and Linked Data formatted information. Therefore, many sites like Facebook [9] have started making heavy use of those formats.

The increasingly important role of user generated content in the Web and concepts 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. In addition to the several sharing options to accompany plain text (company, location, pictures...), now users can express their feelings and emotions. The fact that one of the largest social network is including emotions in such basic and frequently used feature is a sign of the rising relevance of Emotion Analysis. Other sites are also using emotions in a more complex way, like MySmark [3]. With MySmark, users can share the emotions about anything, in a way that is similar to Twitter or Facebook. What makes it different is that it focuses on emotions and uses a model for emotions [17] which helps categorise the emotions in a more sensible and complex way.

However, there are not any ontologies to deal with emotive information and spare developers from designing their own ad-hoc implementation. Furthermore,

[3] `http://mysmark.com`

most of the times emotive information will need to be extracted from already existing sources via Emotion Analysis. Whether this information comes explicitly from the user or has been made explicit by an automated process, they represent the same thing. It is important then that a single ontology covers both the results from such analysis and native emotive information.

With all this in mind, we consider that an ontology to represent Emotion Mining results would be highly beneficial.

3 Enabling Technologies

3.1 Emotion Analysis

For the sake of simplicity, we will informally define Emotion Analysis as the extraction of emotive information from textual content. Even though this is not the only possible definition, it will help us understand the problem domain. However, bear in mind that the term Emotion Analysis and other related terms like Sentiment Analysis have been perceived and defined in several different ways throughout history [16], and that it is also possible to extract emotive information from non-textual media [6].

One of the simplest traditional uses for Emotion Analysis is in personal agents [10]. It provides information about the emotional state of a human user, which the agent can use to modify its behaviour, usually to get the user into a better mood.

When Emotion Analysis is performed performed over big amounts of data, not all of which is valuable for the purpose, it is also called Emotion Mining [28, 25]. The term bears a resemblance and is closely related to Opinion Mining [7], which as its name suggests is focused on representing and extracting human opinions in a machine readable format. The difference between these two fields is their focus: Opinion Mining helps us to attain an approximation of the objective (or otherwise generally perceived) values or attributes of entities whereas Emotion Mining provides us information about the emotional state of the subject. In other words: Opinion Mining focuses on what you think or know, Emotion Mining on how you feel, either in general or about something in particular.

An important fact about Emotions is that they change the way we communicate [16]. Moreover, they can be passed on just like any other information, in what some authors call emotional contagion [5]. 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 [15], with Emotion Mining as one of its components.

Emotion Analysis poses several technical and conceptual challenges. Natural language processing, extraction of emotions from processed language and modelling of emotions are the main ones. Besides these, it faces the same problem as any other application: exposing meaningful data that can be used in other processes or applications. With Onyx we intend to address the latter, as well

as providing tools to choose from different emotion models while keeping them linked.

3.2 Models for Emotions and Sentiment Analysis

Describing and categorising emotions is rather complex.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 [8, 18]. One of the most interesting proposals for Emotion Analysis is the one used in the emotion annotation and representation language (EARL) [13] [19], presented by the Human-Machine Interaction Network on Emotion (HUMAINE) This proposal of an XML language to represent emotions includes 48 emotions divided into 10 different categories. Plutchik takes a different approach [17], relating all the different emotions to each other in what is called the rose of emotions. It is built on top of 8 basic emotions and a set of 8 more complex emotions, which are in turn composed by two basic emotions. There are also models for affects, which include Emotions as part of them. One of them is the work done by Strapparava and Valitutti in WordNet-Affect [20]. It comprises more than 300 affects among which many are considered Emotions. What makes this categorisation 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.

Hastings et al. [12] propose the EMO ontology that reconciles the discrepancies in affective phenomena terminology. The purpose of our ontology is a different one, that is, providing a vocabulary for expressing the results of emotion analysis services as well as annotating lexical resources. Nevertheless, our proposal is compatible with EMO, since EMO can be easily linked with Onyx using the property usesEmotionModel. The situation is similar with the proposal of Lopez et al. [14], which focuses on emotions instead of affects in general.

Despite all, there does not seem to be a universally accepted model for emotions [19]. Considering this, we have decided to leave the representation of the emotion to developers, so they can freely use their emotion model of choice. Emotion Markup Language (EmotionML) [4], one of the most notable general-purpose emotion annotation and representation languages, took a similar approach. In EmotionML, developers can define their own emotion models via vocabularies. 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 [3]. We have used this description to translate its vocabularies to the Onyx format. More on this on Section 5.

In order to simplify the work of developers and to support the use of the ontology, we have translated the WordNet-Affect [20] taxonomy and published a SKOS version of it [23]. 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.

On the Sentiment Analysis side, Onyx complements and extends the Marl ontology [27], a vocabulary 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.

3.3 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 [24]. It includes a full-fledged ontology, which we will use in Onyx. 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.

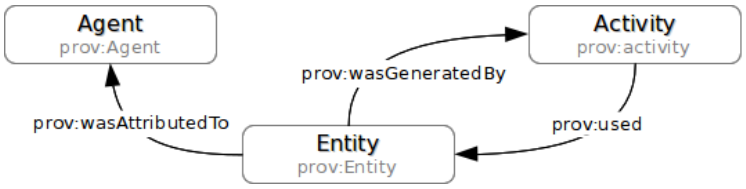

Fig. 1: Overview of the basic Provenance classes

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 Analaysis 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.

4 Onyx

The Onyx Ontology effectively combines the concepts that we have introduced in the previous sections while providing a simple model. It reflects a selection of

the basic attributes that we found are common to the main services and open source projects in Emotion Analysis. It was also designed to be very easy to adapt and extend, taking advantage of the nature of the semantic technologies.

At its core, the Onyx ontology has three main classes: EmotionAnalysis, EmotionExpresion and Emotion. In a standard Emotion Analysis, these three classes are related as follows: an EmotionAnalysis analyses a source (generally in the form of text, e.g. a status update), the result is represented as an EmotionExpression that contains one or more Emotion. Table 1, Table 2 and Table 3 contain a comprehensive list of the properties associated with each of these classes.

Property	**Description**
domain	The specific domain in which the EmotionAnalysis was carried out
algorithmConfidence	Numeric value that represents the predicted accuracy of the result, as given by the algorithm in use
extractedFrom	Text or resource that was subject to the analysis
hasEmotion	An Emotion that is shown by the EmotionExpression. An EmotionResult may contain several Emotions.

Table 1: EmotionResult class properties

Property	**Description**
source	Identifies the source of the user generated content
algorithm	Emotion analysis algorithm that was used
usesEmotionModel	Link to the emotion model used, which defines the categories, dimensions, appraisals, etc.

Table 2: EmotionAnalysis class properties

Property	**Description**
hasEmotionCategory	The type of emotion, defined by an instance of the Emotion Ontology as specified in the corresponding EmotionAnalysis
hasEmotionIntensity	Degree of intensity of the emotion
emotionText	Fragment of the EmotionExpression's source that contained emotion information
hasActionTendency	Action that is triggered by this emotion

Table 3: Emotion class properties

In addition to the classes above, there are other auxiliary classes that make it possible to use different models of emotions. Namely, these classes are: EmotionCategory, which represents an emotion category (or type), such as "sadness" or

"surprise"; ActionTendency, an action that is triggered by emotions; DimensionProperty, a class to be be inherited by Dimension properties; AppraisalProperty, the analogous to the previous one for Appraisals. Lastly, the EmotionModel class represents a model for emotions and groups a specific set of instances of the previous classes. Figure 2 shows a complete overview of all these classes, as well as all their properties.

Finally, we developed an extension to the Onyx core that includes all the categories specified by the EmotionML specification [22]. Using this extension we can translate resources expressed in EmotionML XML into Onyx for their use in the Semantic Web.

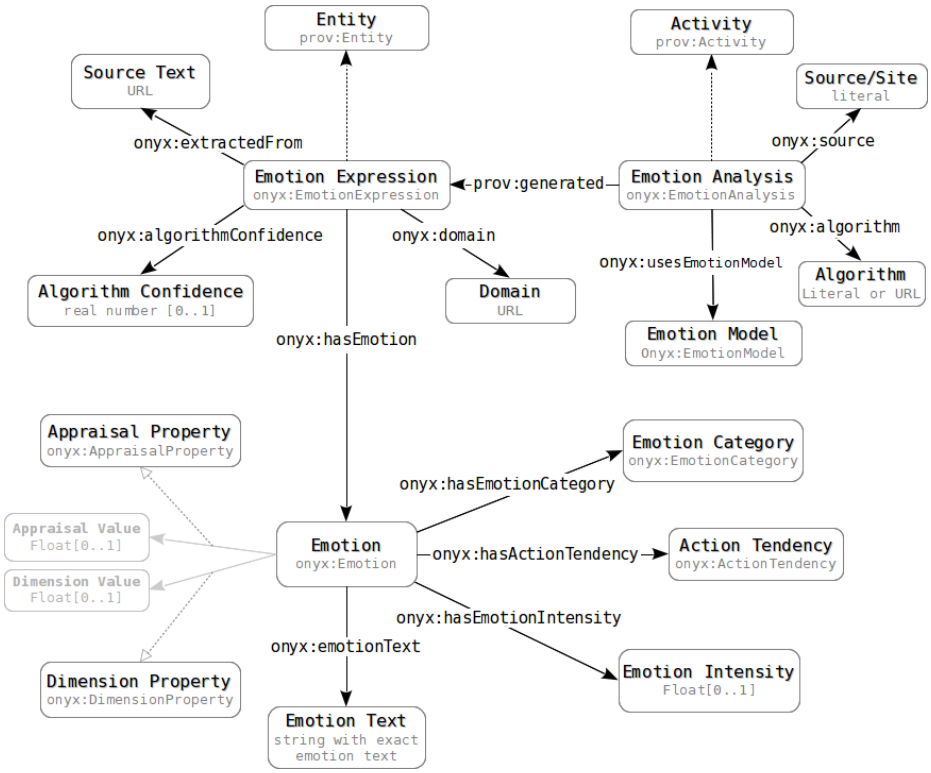

Fig. 2: Class diagram of the Onyx ontology.

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. Listing 1.1 has a comprehensive list of prefixes used, both in the n3 examples (Table 4) and the SPARQL queries (Table 5).

```
@prefix onyx: <http://gsi.dit.upm.es/ontologies/onyx/ns>.
@prefix wna: <http://gsi.dit.upm.es/ontologies/wnaffect/ns>.
@prefix prov: <http://www.w3.org/ns/prov#>.
@prefix dcterms: <http://purl.org/dc/terms>.
@prefix lemon: <http://www.monnet-project.eu/lemon#>.
@prefix wn: <http://semanticweb.cs.vu.nl/europeana/lod/purl/vocabularies/princeton
    /wn30/>.
@prefix lexinfo: <http://www.lexinfo.net/ontology/2.0/lexinfo#>.
```

Listing 1.1: "Prefixes for the examples."

5 Evaluation

Evaluating ontologies is always a difficult task, evaluation methodologies are highly debatable and there are no standards [11]. 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.

First of all, we consider that the translation of WordNet-Affect to SKOS and its use with Lemon and Onyx to represent lexicon in the context of Emotion Analysis is in its own proof of applicability. However, Onyx was created with this application in mind. To guarantee its versatility in a wider range of scenarios, we also looked for totally different applications where it could be used. Finally, we chose two different test scenarios: the translation of a well-known Emotion Analysis tool, Synesketch [21], and the translation of resources in the rather extended emotion format EmotionML.

With Synesketch, our approach has been to represent its Emotion classes to Onyx and then develop a proof-of-concept web service that performs basic Emotion Analysis using Synesketch and returns the results in valid RDF format. The Synesketch API accepts text input, and returns an object with the emotions in it. In particular, it 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 other attributes that correspond to the general emotional valence (positive, negative or neutral) and the general emotional weight. To model the big 6 emotions with Onyx, we used their equivalents in WordNet-Affect, as explained before in Section 4. After we translated the EmotionML categories to Onyx, it could have been modelled using the big-6 category from EmotionML. The Synesketch weight translated to the 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 Emotion with the PositiveEmotion, NeutralEmotion or NegativeEmotion category (as defined by WordNet-Affect) depending on the value of the valence (1, 0 and -1, respectively), and represented the weight as the intensity, just like in the other cases. We consider this a valid representation due to two facts: a) the Synesketch model uses only 6 pre-defined emotions that do not overlap with these three emotion categories we are using,

Case	N3 Representation

A personal status using the provided WordNet-Affect SKOS taxonomy	`:status1`
	`  a onyx:EmotionExpression;`
	`  onyx:hasEmotion [ onyx:hasEmotionCategory wna:anger ];`
	`  dcterms:created "2013-05-16T19:20:30+01:00"^^dcterms:W3CDTF.`

Example EmotionAnalysis.	`:customAnalysis`
	`  a onyx:EmotionAnalysis;`
	`  onyx:algorithm "SimpleAlgorithm";`
	`  onyx:usesEmotionModel wna:WNAModel.`

Example of results of an Emotion Analysis using custom emotions	`:status3`
	`  a onyx:EmotionResult;`
	`  onyx:hasEmotion [ onyx:hasEmotionCategory :blue.`
	`  onyx:hasEmotionIntensity :0.9];`
	`  prov:wasGeneratedBy :customAnalysis.`

Example of annotation of a lexical entry using Onyx and Lemon [2].	`:fifa`
	`  a lemon:Lexicalentry;`
	`  lemon:sense [`
	`  lemon:reference wn:synset-fear-noun-1;`
	`  onyx:hasEmotion`
	`  [ oyx:hasEmotionCategory wna:fear. ].`
	`  ];`
	`  lexinfo:partOfSpeech lexinfo:noun.`

Table 4: Representation with Onyx

Case	Query

Simple query of emotions of a given type.	`SELECT ?emotionResult`
	`WHERE {`
	`  ?emotionResult onyx:hasEmotion ?emo.`
	`  ?emo onyx:hasEmotionCategory wna:fear.`
	`}`

Advanced query of emotions by broader category. Select all the results that show a positive emotion.	`SELECT ?expression ?cat`
	`WHERE {`
	`  ?expression onyx:hasEmotion`
	`  [ onyx:hasEmotionCategory ?cat].`
	`  ?cat skos:broaderTransitive* wna:positive-emotion.`
	`}`

Querying the EmotionAnalysis activities for discrepancies in the results obtained using two different algorithms.	`SELECT ?source1 ?algo1 (GROUP_CONCAT(?cat1) as ?cats1)`
	`WHERE {`
	`  ?expression1 onyx:extractedFrom ?source1.`
	`  ?analysis1 prov:generated ?expression1;`
	`  onyx:algorithm ?algo1.`
	`  ?expression1 onyx:hasEmotion`
	`  [ onyx:hasEmotionCategory ?cat1 ].`
	`  FILTER EXISTS{`
	`  ?expression2 onyx:extractedFrom ?source1.`
	`  ?analysis2 prov:generated ?expression2.`
	`  ?expression2 onyx:hasEmotion`
	`  [ onyx:hasEmotionCategory ?cat2 ].`
	`  FILTER ( ?expression1 != ?expression2).`
	`  FILTER ( ?cat2 != ?cat1 ).`
	`  }`
	`}`
	`GROUP BY ?source1 ?algo1`
	`ORDER BY ?source1`

Table 5: Example SPARQL queries with Onyx

b) these three categories are very generic and probably should only be used for aggregation of emotions (like this case) or other uses that do not interfere with the rest of the categories. The final result is publicly available at our website [4].

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 [22], which has been used as namespace (emlonyx) in the translation of an EmotionML example in Table 6.

EmotionML	**Onyx**

`<emotionml xmlns="http://www.w3.org /2009/10/emotionml" xmlns:meta="http://www. example.com/metadata" category-set="http://www.w3.org/TR/ emotion-voc/xml#everyday- categories">`
`<info>`
`<classifiers:classifier classifiers:name="GMM"/>`
`</info>`
`<emotion>`
`<category name="Disgust" value="0.82"/>`
'Come, there is no use in crying like that!'
`</emotion>`
said Alice to herself rather sharply;
`<emotion>`
`<category name="Anger" value="0.57"/>`
'I advise you to leave off this minute!'
`</emotion>`
`</emotionml>`
	:Analysis1 a onyx:EmotionAnalysis;
	onyx:algorithm "GMM".
	onyx:usesEmotionModel emlonyx: everyday-categories
	:Exp1 a onyx:EmotionExpression;
	onyx:extractedFrom "Come, there is no use in crying like that! said Alice to herself rather sharply; I advice you to live off this minute!";
	onyx:hasEmotion :Emo1
	onyx:hasEmotion :Emo2
	:Emo1 a onyx:Emotion;
	onyx:hasEmotionCategory emlonyx: disgust;
	onyx:hasEmotionIntensity 0.82;
	onyx:hasEmotionText "Come, there's no use in crying like that!"
	:Emo2 a onyx:Emotion;
	onyx:hasEmotionCategory emlonyx: anger;
	onyx:hasEmotionIntensity 0.57;
	onyx:hasEmotionText "I advice you to leave off this minute!"

Table 6: Representation of EmotionML with Onyx

6 Conclusions and Future Work

With this work we have presented 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. Furthermore, it makes it possible to publish and retrieve information about emotions in a multitude of scenarios and gives a formal treatment to the Emotion Analysis process itself by integrating with the Provenance Ontology. First, we presented several simple scenarios where these features can be exploited. But we went a step further by adapting current resources and services to Onyx, which are now publicly available. This opens the door to its use by service developers that could highly benefit from Linked Data by sharing their resources and develop applications that build on different sources.

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

As part of the future plans for Onyx, it will be actively used in the EUROSENTIMENT [5] project, whose aims is to create a language resource pool for Sentiment Analysis. Together with Marl [27] and Lemon [2], 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. Finally, our research group will use the integration with EmotionML to develop intelligent personal agents that benefit from the potential of the Semantic Web.

References

[1] Ben Adida et al. "RDFa in XHTML: Syntax and processing". In: *Recommendation, W3C* (2008).

[2] In: *The Semantic Web: Research and Applications*. Ed. by Grigoris Antoniou et al. Vol. 6643. Lecture Notes in Computer Science. 2011. ISBN: 978-3-642-21033-4.

[3] Kazuyuki Ashimura et al. *EmotionML vocabularies*. May 2012. URL: http://www.w3.org/TR/2012/NOTE-emotion-voc-20120510/.

[4] Paolo Baggia et al. *Emotion Markup Language (EmotionML) 1.0*. Apr. 2013. URL: http://www.w3.org/TR/emotionml/.

[5] Sigal G. Barsade. "The Ripple Effect: Emotional Contagion and its Influence on Group Behavior". In: 47.4 (2002), pp. 644–675.

[6] Rafael A Calvo and Sidney D'Mello. "Affect detection: An interdisciplinary review of models, methods, and their applications". In: *IEEE Transactions on Affective Computing* 1.1 (2010), pp. 18–37.

[7] Kushal Dave, Steve Lawrence, and David M. Pennock. "Mining the Peanut Gallery: Opinion Extraction and Semantic Classification of Product Reviews". In: *Proceedings of the 12th International Conference on World Wide Web*. WWW '03. Budapest, Hungary, 2003, pp. 519–528. ISBN: 1-58113-680-3.

[8] Paul Ekman. "Basic emotions". In: *Handbook of cognition and emotion* 98 (1999), pp. 45–60.

[9] *Facebook Open Graph API*. June 2013. URL: http://developers.facebook.com/docs/opengraph/.

[10] Martin E Ford. *Motivating humans: Goals, emotions, and personal agency beliefs*. Sage, 1992.

[11] Asunción Gómez-Pérez. "Evaluation of ontologies". In: *International Journal of Intelligent Systems* 16.3 (). ISSN: 1098-111X.

[12] Janna Hastings et al. "Dispositions and Processes in the Emotion Ontology." In: *ICBO*. 2011.

[13] *HUMAINE Emotion Annotation and Representation Language (EARL): Proposal*. June 2006. URL: http://emotion-research.net/projects/humaine/earl/proposal#Dialects.

[5] http://eurosentiment.eu

[14] Juan Miguel López et al. "Towards an ontology for describing emotions". In: *Emerging Technologies and Information Systems for the Knowledge Society*. Springer, 2008, pp. 96–104.

[15] Alan Mislove et al. "Measurement and analysis of online social networks". In: *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*. IMC '07. San Diego, California, USA, 2007, pp. 29–42. ISBN: 978-1-59593-908-1.

[16] Bo Pang and Lillian Lee. "Opinion mining and sentiment analysis". In: *Foundations and trends in information retrieval* 2.1-2 (2008), pp. 1–135.

[17] Robert Plutchik. *Emotion: A psychoevolutionary synthesis*. Harper & Row New York, 1980.

[18] Jesse J Prinz. "Gut reactions: A perceptual theory of emotion". In: (2004).

[19] Marc Schröder, Hannes Pirker, and Myriam Lamolle. "First suggestions for an emotion annotation and representation language". In: *Proceedings of LREC*. Vol. 6. Citeseer. 2006, pp. 88–92.

[20] Carlo Strapparava and Alessandro Valitutti. "WordNet-Affect: an affective extension of WordNet". In: *Proceedings of LREC*. Vol. 4. 2004, pp. 1083–1086.

[21] *Synesketch: Free Open-Source Software for Textual Emotion Recognition and Visualization*. June 2006. URL: `http://emotion-research.net/projects/humaine/earl/proposal#Dialects`.

[22] J. Fernando Sánchez-Rada and Carlos Á. Iglesias. *EmotionML categories for Onyx*. July 2013. URL: `http://gsi.dit.upm.es/ontologies/onyx/emotionml/ns`.

[23] J. Fernando Sánchez-Rada and Carlos Á. Iglesias. *WordNet-Affect SKOS Taxonomy*. May 2013. URL: `http://gsi.dit.upm.es/ontologies/wnaffect/`.

[24] "The Open Provenance Model core specification (v1.1)". In: *Future Generation Computer Systems* 27.6 (2011), pp. 743 –756. ISSN: 0167-739X.

[25] Mike Thelwall, David Wilkinson, and Sukhvinder Uppal. "Data mining emotion in social network communication: Gender differences in MySpace". In: *Journal of the American Society for Information Science and Technology* 61.1 (2010), pp. 190–199.

[26] Giovanni Tummarello, Renaud Delbru, and Eyal Oren. "Sindice. com: Weaving the open linked data". In: *The Semantic Web*. Springer, 2007, pp. 552–565.

[27] 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, 2011.

[28] M. Yassine and H. Hajj. "A Framework for Emotion Mining from Text in Online Social Networks". In: *IEEE International Conference on Data Mining Workshops (ICDMW)*. 2010, pp. 1136–1142.