Onyx: A Linked Data Approach to Emotion Representation

J. Fernando Sánchez-Rada, Carlos A. Iglesias

Abstract

Extracting opinions and emotions from text is becoming increasingly important, especially since the advent of micro-blogging and social networking. Opinion mining is particularly popular and now gathers many public services, datasets and lexical resources. Unfortunately, there are few available lexical and semantic resources for emotion recognition that could foster the development of new emotion aware services and applications. The diversity of theories of emotion and the absence of a common vocabulary are two of the main barriers to the development of such resources. This situation motivated the creation of Onyx, a semantic vocabulary of emotions with a focus on lexical resources and emotion analysis services. It follows a linguistic Linked Data approach, it is aligned with the Provenance Ontology, and it has been integrated with the Lexicon Model for Ontologies (lemon), a popular RDF model for representing lexical entries. This approach also means a new and interesting way to work with different theories of emotion. As part of this work, Onyx has been aligned with EmotionML and WordNet-Affect.

Keywords: emotion analysis, ontology, provenance, Linked Data, EmotionML, lexical resources

1. Introduction

With the rise of social media, more and more users are sharing their opinions and emotions online [1]. The increasing volume of information and number of users are drawing the attention of researchers and companies alike, which seek not only academical results but also profitable applications such as brand monitoring. As a result, many tools and services have been created to enrich or make sense out of human generated content. Unfortunately, they are isolated data silos or tools that use very different annotation schemata. Even worse, the scarce available resources are also suffering from the heterogeneity of formats and models of emotion, making it hard to combine different resources.

Linked Data can change this situation, with its lingua franca for data representation as well as a set of tools to process and share such information. Plenty

of services have already embraced the Linked Data concepts and are providing tools to interconnect the previously closed silos of information [2]. In fact, the Linked Data approach has proven useful for fields like Opinion Mining. Some schemata offer semantic representation of opinions [3], allowing richer processing and interoperability.

Emotions have a crucial role in our lives, and even change the way we communicate [1]. They can be passed on just like any other kind of information, in what some authors call emotional contagion [4]. That is a phenomenon that is clearly visible in social networks [5]. Public APIs make it relatively easy to study the social networks and their information flow. For this very reason Social Network Analysis is an active field [6], with Emotion Mining as one of its components. On the other hand, the growing popularity of services like micro-blogging will inevitably lead to services that exchange and use emotion in their interactions. Some social sites are already using emotions natively, giving their users the chance to share emotions or use them in queries. A noteworthy example is Facebook, which recently updated the way its users can share personal statuses.

Nevertheless, the impact of emotion analysis goes well beyond social networks. For instance, there are a variety of systems whose only human-machine communication is purely text-based. Such systems can use the emotive information to change their behavior and responses [1].

Furthermore, there are many sources that can be used for sentiment analysis beyond pure text, including video and audio. Multimodal analysis, or making use of several of these sources, is an active field that gathers experts from different disciplines. A unified schema and appropriate tooling would open up new possibilities in this field.

Lastly, combining subjective information from emotion analysis with facts already published as Linked Data could enable a wide array of new services. This would require a widely accepted Linked Data representation for emotions, which does not exist yet. In this paper we present Onyx, a new vocabulary that aims to bridge this gap and allow for interoperable tools and resources. We also provide a set of example applications, additional vocabularies to use existing models, and multilingual resources that use Onyx to annotate emotion.

The rest of 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, including some vocabularies or models of emotions, and several examples. Section 4 presents the results of our evaluation of the Ontology, focusing on the coverage of current formats like Emotion Markup Language (EmotionML). Finally, Section 5 completes this paper with conclusions and future work.

2. Enabling technologies

2.1. Models for emotions and emotion 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: modeling (including categorization) and representation.

Confusingly, the terms opinion, sentiment, emotion, feeling and affect are commonly used interchangeably. Throughout this article we follow the terminology by Cambria et al. [7] where opinion mining and sentiment analysis are focused on polarity detection and emotion recognition, respectively.¹

There are also 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 [9, 10]. 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. [11] 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 [12]. Plutchik’s model is a model of categories that has been extensively used [13, 14, 15] in the area of emotion analysis and affective computing, relating all the different emotions to each other in what is called the wheel of emotions.

All the existing motions are mainly divided in two groups: discrete and dimensional models. In discrete models, emotions belong to one of a predefined set of categories, which varies from model to model. In dimensional models, an emotion is represented by the value in different axes or dimensions. A third category, mixed models, merges both views.

Other models are more general and model affects, including emotions as a subset. One of them is the work done by Strapparava and Valitutti in WordNet-Affect [16], an affective lexicon on top of WordNet. WordNet-Affect comprises more than 300 affective labels linked by concept-superconcept relationships, 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 relationships between emotions and makes it possible to decide the level of granularity of the emotions expressed. Section 3.2.1 discusses how we formalized this taxonomy using SKOS, and converted that taxonomy into an Onyx vocabulary.

Despite all, there does not seem to be a universally accepted model for emotions [17]. This complicates the task of representing emotions. In a discussion regarding EmotionML, Schroder et al. pose that *given the fact that even emotion theorists have very diverse definitions of what an emotion is, and that very different representations have been proposed in different research strand, any attempt to propose a standard way of representing emotions for technological contexts seems doomed to fail* [18]. Instead they claim that the markup should offer users choice of representation, including the option to specify the affective state that is being labeled, different emotional dimensions and appraisal scales. The level of intensity completes their definition of an affect in their proposal.

EmotionML [19] 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) [20, 17] 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 their

¹A more detailed terminology discussion can be found in Munezero et al. [8].

computer-readable form available [21].

In the field of Semantic Technologies, Grassi introduced Human Emotion Ontology (HEO), an ontology for human emotions meant for annotating emotions in multimedia data. We discuss some differences between Onyx and HEO in Section 5. Another work worth mentioning is that of Hastings et al. [22] in Emotion Ontology (EMO), an ontology that tries to reconcile the discrepancies in affective phenomena terminology. It is, however, too general to be used in the context of emotion analysis: it provides a qualitative notion of emotions, when a quantitative one would be needed.

For Opinion Mining we find the Marl vocabulary [3]. 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 Opinion Mining in an open and sensible format.

2.2. Linguistic Linked Open Data (LLOD)

LLOD [23, 24] is an initiative that promotes the use of linked data technologies for modeling, publishing and interlinking linguistic resources. The main benefit of using linked data principles to model linguistic resources is that it provides a graph-based model that allows representing different kinds of linguistic resources (such as lexical-semantic resources, linguistic annotations or corpora) in a uniform way, thus supporting querying across resources. The creation of an LLOD cloud is a cooperative task that is been managed by several communities, such as Open Linguistics Working Group (OWLIG) of the Open Knowledge Foundation and Ontology-Lexica Community Group (OntoLex) of W3C. As a result of this activity, an initial LLOD is currently available as shown in Figure 1, where several types of resources have been identified: lexical-semantic resources (e.g. machine readable dictionaries, semantic networks, semantic knowledge bases, ontologies and terminologies), annotated corpora and linguistic annotations.

With regards to modeling lexical-semantic resources, *Lexicon Model for Ontologies (lemon)* [25] proposes a framework for modeling and publishing lexicon and machine-readable dictionaries as linked data. *lemon* provides a bridge between the most influential lexical-semantic resources, WordNet [26] and DBPedia [27]. With regards to annotated corpora, there are two initiatives [28], POWLA [29] and NLP Interchange Format (NIF) [30] that enable to link lexical-semantic resources to corpora. Finally, Ontologies of Linguistic Annotation (OLiA) [31] are a repository of annotation terminology for various linguistic phenomena that can be used in combination with POWLA, NIF or *lemon*. OLiA ontologies allow to represent linguistic annotations in corpora, grammatical specifications in dictionaries, and their respective meaning within the LLOD cloud in an operable way.

The main benefits of modeling linguistic resources as linked data include [24] interoperability and integration of linguistic resources, unambiguous identification of elements of linguistic description, unambiguous links between different resources, possibility to annotate and query across distributed resources and availability of mature technological infrastructure.

2.3. 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

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 emotion analysis in particular, as different algorithms may produce different results.

3. Onyx ontology and vocabularies

This section gives a comprehensive view of the ontology and is structured in three parts. Subsection 3.1 presents the ontology in full. Subsection 3.2 shows how different vocabularies are represented with the ontology, using three known models of emotion. Lastly, Subsection 3.3 exemplifies how to annotate data using the ontology and vocabularies from the previous section.

3.1. Onyx ontology

Onyx is a vocabulary that models emotions and the emotion analysis process itself. It can be used represent the results of an emotion analysis service or the lexical resources involved (e.g. corpora and lexicons). This vocabulary can connect results from different providers and applications, even when different models of emotions are used.

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

The model of emotions in Onyx is very generic, which reflects the lack of consensus on modeling and categorizing emotions. An advantage of this approach is that the representation and psychological models are decoupled.

The *EmotionAnalysis* instance contains information about the source (e.g. dataset) from which the information was taken, the algorithm used to process it, and the emotion model followed (e.g. Plutchik's categories). 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 in one of its parts. As such, it contains information about: the original text (*extracted-From*); 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 these 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. *AggregatedEmotionSet* is a subclass of *EmotionSet* that contains additional information about the original *EmotionSet* instances it aggregates. For instance, we could aggregate all the emotions related to a particular movie, or all the emotions shown by a particular user, and still be able to trace back to the original individual emotions.

Figure 3: Class diagram of the Onyx ontology.

Onyx’s *Emotion* model includes: *EmotionCategory* which is a specific category of emotion (e.g. “sadness”, although more than one could be specified), linked through the *hasEmotionCategory* property; the emotion intensity via *hasEmotionIntensity*; action tendencies related to this emotion, or actions that are triggered by the emotion; appraisals and dimensions. Lastly, specific appraisals, dimensions and action tendencies can be defined by sub-classing *Appraisal*, *Dimension* and *ActionTendency*, whose value should be a float number.

On top of that generic model we have included two different models: the WordNet-Affect taxonomy, and the EmotionML vocabularies for categories, dimensions and appraisals, which are detailed in Section 3.2.

Although emotional models and categories differ in how they classify or quantify emotions, they describe different aspects of the same complex phenomenon emotion [34]. Hence, there are equivalence relationships between different categories or emotions in different models. To state such equivalence between emotion categories in Onyx one can use the properties defined in SKOS² such as *skos:exactMatch* or *skos:closeMatch*. This approach falls short when dealing with dimensional emotional theories or complex category theories. Since dimensional models are widely used in practice, Section 4.4 covers how to deal with this issue in detail.

Within a single model, it is also possible that two separate emotions, when found simultaneously, imply a third one. For instance, “thinking of the awful things I’ve done makes me want to cry” might reveal sadness and disgust,

²<http://www.w3.org/TR/skos-reference/#mapping>

EmotionAnalysis	
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.
EmotionSet	
Property	Description
domain	The specific domain in which the EmotionAnalysis was carried out
algorithmConfidence	Numeric value that represents the predicted accuracy of the result
extractedFrom	Text or resource that was subject to the analysis
hasEmotion	An Emotion that is shown by the EmotionSet. An EmotionExpression may contain several Emotions.
Emotion	
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 EmotionSet’s source that contained emotion information

Table 1: Main properties in the ontology. The specification contains the full description of all properties: <http://www.gsi.dit.upm.es/ontologies/onyx>

which together might be interpreted as remorse. Some representations would refer to remorse as a complex emotion. Onyx purposely does not include the notion of complex emotions. It follows the same approach as EmotionML in this respect, as HUMAINE EARL included this distinction between simple and complex emotions, but it was not included in the EmotionML specification. This simplifies the ontology and avoids discussion about the definition of complex emotions, since there are several possible definitions of a complex emotion, and different levels of emotions (e.g. the Hourglass of Emotions model). One possible way to deal with such situation is to add an *AggregatedEmotion* that represents remorse to the *EmotionSet*, linking it to the primary emotions with the *aggregatesEmotion* property.

Table 1, contains a comprehensive list of the properties associated with each of these classes. Figure 3 shows a complete overview of all these classes and their properties.

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 *Appraisal* or *Dimension* instances it introduces (through *hasAppraisal* and *hasDimension*), etc.

Having a formal representation of the categories and dimensions proves very useful when dealing with heterogeneous datasets in emotion analysis. In addition to being necessary to interpret the results, this information can be used to filter out results and for automation.

3.2. Vocabularies

An *EmotionModel*, or at least an *EmotionCategory*, has to be defined in order to make a valid annotation. Annotators can define their own ad-hoc models and categories, but the Linked Data approach dictates that vocabularies and entities should be reused when appropriate. Hence, we offer several *EmotionModel* vocabularies that can be used with Onyx.

As of this writing, we have modeled the quite extensive WordNet-Affect taxonomy as an *EmotionModel*, to be used as the reference for categorical representation. We also ported the main vocabularies defined for EmotionML [21], and created a model based on the The Hourglass of Emotions [11]. A list of vocabularies with a detailed explanation is publicly available³. Onyx’s Github repository⁴ contains the tools used to generate all these models.

3.2.1. WordNet-Affect

WordNet-Affect [16] contains a subset of synsets suitable to represent affective concepts. Each synset is given one or more affective labels (a-labels) or categories. These labels are linked via concept/superconcept relationships. We processed the list of labels in WordNet-Affect 1.1 and generated a SKOS taxonomy. This taxonomy and its specification are available on our website⁵. In this specification we included a navigable tree with all the affects and their relationships. This tree makes it trivial to select an affect that represents the desired emotion. Figure 4 shows part of the tree, with some nodes collapsed.

The full taxonomy contains 305 affects, 291 of which are related to emotions. The RDF version of the taxonomy also includes an *EmotionModel* that contains these 291 affects as *EmotionCategory* entities.

Besides providing a good starting point for other ontologies, the resulting taxonomy also serves as reference for mapping between several different ontologies in the future.

3.2.2. EmotionML

EmotionML does not include any emotion vocabulary in itself. However, the Multimodal Interaction Working Group released a series of vocabularies that cover the most frequent models of emotions [21]. Users can either define their own vocabularies or reuse one of the existing ones.

We have developed a tool that generates an *EmotionModel* model from a vocabulary definition, including all its dimension, category, appraisal or action tendency entries. Using this tool, we have processed the vocabularies released by the Multimodal Interaction Working Group.

EmotionML has four types of vocabularies, according to the type of characteristic of the emotion phenomenon they represent: emotion categories, emotion

³<http://www.gsi.dit.upm.es/ontologies/onyx/vocabularies>

⁴<https://github.com/gsi-upm/onyx>

⁵<http://gsi.dit.upm.es/ontologies/wnaffect/>

Figure 4: Small part of the full WordNet-Affect labels taxonomy. Blue nodes are contracted.

Listing 1: Excerpt of the models from EmotionML in Onyx

```

emoml:big6 a onyx:EmotionModel ;
  onyx:hasEmotionCategory emoml:big6_anger,
                           emoml:big6_disgust,
                           emoml:big6_fear,
                           emoml:big6_happiness,
                           emoml:big6_sadness,
                           emoml:big6_surprise .

emoml:ema a onyx:EmotionModel ;
  onyx:hasAppraisal emoml:ema_adaptability,
                    emoml:ema_agency,
                    emoml:ema_blame,
  ...

emoml:frijda a onyx:EmotionModel ;
  onyx:hasActionTendency emoml:frijda_agonistic,
                          emoml:frijda_approach,
                          emoml:frijda_attending,
  ...
  onyx:hasEmotionCategory emoml:frijda_anger,
                           emoml:frijda_arrogance,
                           emoml:frijda_desire,

emoml:pad a onyx:EmotionModel ;
  onyx:hasDimension emoml:pad_arousal,
                    emoml:pad_dominance,
                    emoml:pad_pleasure .

```

dimensions, appraisals and action tendencies. If an emotion model addresses several of these characteristics, there will be an independent vocabulary for each. For instance, Frijda’s model defines action tendencies and categories, which results in the *frijda-categories* and *frijda-action-tendencies vocabularies*. In Onyx, instead of following this approach, we opted for adding all characteristics in the same model. This results in cleaner URIs, and helps represent the emotion model as a whole.

With these vocabularies it is possible to translate EmotionML resources into Onyx for their use in the Semantic Web. Table 2 contains an example of how EmotionML resources can be translated to Onyx.

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

Table 2: Example translation of an EmotionML resource to Onyx. emoml: [<http://www.gsi.dit.upm.es/ontologies/onyx/vocabularies/emotionml/ns#>](http://www.gsi.dit.upm.es/ontologies/onyx/vocabularies/emotionml/ns#)

3.2.3. The Hourglass of Emotions

The Hourglass of Emotions [11] is an interesting example of mixed models, including both dimensions and categories. Based on the paper by Cambria et al. we have created a basic model in Onyx, which includes the four dimensions and

24 first-level emotions, and 32 second-level emotions. Using the generated Onyx vocabulary, we can perform simple experiments with this interesting model. Nonetheless, a complete representation would include the relationships between the different categories and the dimensions, or restrictions.

Listing 2: Excerpt from the definition of the Hourglass of Emotions model in Onyx

```
hg:HourglassModel a onyx:EmotionModel ;
  onyx:hasDimension hg:Pleasantness ,
    hg:Attention ,
    hg:Sensitivity ,
    hg:Aptitude ;
  onyx:hasEmotionCategory hg:ecstasy ,
    hg:vigilance ,
    hg:rage ,
    ...
    hg:coercion .

hg:Pleasantness a onyx:EmotionDimension .
hg:Attention a onyx:EmotionDimension .
hg:Sensitivity a onyx:EmotionDimension .
hg:Aptitude a onyx:EmotionDimension .

hg:ecstasy a onyx:EmotionCategory .
hg:vigilance a onyx:EmotionCategory .
hg:rage a onyx:EmotionCategory .
...
hg:coercion a onyx:EmotionCategory .
```

3.3. Examples

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

4. Applications

This section examines how Onyx has been applied in several use cases, such as modeling and publication of emotion lexicons (Section 4.1), annotation of emotion in corpora (Sect. 4.2), enabling interoperability of emotion services (Sect. 4.3) or mapping and composition heterogeneous emotion representations (Section 4.4).

4.1. Emotion annotation of lexical entries

Annotated lexical resources are the bases for analysis of emotions in text. These resources currently available can be classified into opinion and emotion lexicons.

Opinion lexicons supply a polarity value for a given lexical entry. Some examples of opinion lexicons are SentiWordNet [36], SenticNet 1.0 [37], Multi-Perspective Question Answering (MPQA) [38] or GermanPolarityClues [39].

Case	Turtle Representation
An example Emotion-Analysis.	<pre> :customAnalysis a onyx:EmotionAnalysis; onyx:algorithm "SimpleAlgorithm"; onyx:usesEmotionModel wna:WNAModel. </pre>
Processing “I lost one hour today because of the strikes!”, by the user JohnDoe	<pre> :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. </pre>
Example of annotation of a lexical entry using Onyx and <i>lemon</i> [35].	<pre> :fifa a lemon:Lexicalentry; lemon:sense [lemon:reference wn:synset-fear-noun-1; onyx:hasEmotion [onyx:hasEmotionCategory wna:fear.];]; lexinfo:partOfSpeech lexinfo:noun. </pre>

Table 3: Representation with Onyx

Case	Query
Finding all the users that did not feel good during last New Year’s Eve, and the exact emotions they felt.	<pre> 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")) } </pre>
Comparing two Emotion Mining algorithms by comparing the discrepancies in the results obtained using both.	<pre> 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 </pre>

Table 4: Example SPARQL queries with Onyx

Emotion lexicons supply an emotion description for a given lexical entry. Some examples of emotion lexicons are WordNet-Affect [16], the Affective Norms for English Words (ANEW) [40], EmoLex [41] or SenticNet 3.0 [42].

In this section we present the application of Onyx for annotating lexical entries with emotion representations. This application has been carried out in the context of the Eurosentiment R&D project where a LLOD approach has been followed [43]. In particular, lexical entries described with the *lemon* model have been extended with sentiment and emotion features using Marl and Onyx, respectively, as illustrated in Listing 3.

Listing 3: Lexicon expressed in *lemon* extended with sentiment and emotion features.

```
le:literature_en a lemon:Lexicon;
  lemon:language "en";
  lemon:entry lee:book, lee:terrifying.

lee:sense/terrifying_3 a lemon:Sense;
  lemon:reference dbp:Horror_fiction;
  lemon:reference "00111760";
  lexinfo:partOfSpeech lexinfo:adjective;
  lemon:context lee:sense/book;
  marl:hasPolarity marl:Positive;
  marl:polarityValue 0.7;
  onyx:hasEmotionSet [
    onyx:hasEmotion [
      onyx:hasEmotionCategory wna:fear;
      onyx:hasEmotionIntensity 1.
    ] ;
  ] .

lee:book a lemon:LexicalEntry;
  lemon:canonicalForm [ lemon:writtenRep "book"@en ];
  lemon:sense [ lemon:reference wn:synset-book-noun-1;
    lemon:reference dbp:Book. ];
  lexinfo:partOfSpeech lexinfo:noun.

lee:terrifying a lemon:LexicalEntry;
  lemon:canonicalForm [ lemon:writtenRep "terrifying"@en ];
  lemon:sense lee:sense/terrifying_3;
  lexinfo:partOfSpeech lexinfo:adjective.
```

In this example, we illustrate how a lexical entry (terrifying) has a positive polarity value (0.7) and a emotion category (fear) when referring to the word book. In other words, a terrifying book is linked to feeling fear, but it usually represents a positive quality. In contrast, terrifying is a negative quality when referred to news. For a more in-depth explanation of the format, see [44, 43].

The main benefits of this approach are:

- Lexical entries are aligned with WordNet [26]. This enables interoperability with other affect lexicons, such as WordNet-Affect, SentiWordNet or GermanPolarityClues are also aligned with WordNet. Moreover, WordNetDomains has been used to define the domain of the lexical entry.
- Lexical entries are aligned with Linked Data entities in datasets such as DBPedia. None of the reviewed affect lexicons provide this feature yet, except for SenticNet 3.0.
- Most of the affect lexicons assign prior polarities or emotions to lexical entries, with the exception of SenticNet, that uses multi-word expressions (i.e. small room). The use of multi-word expressions is also possible in

Lexicons		
Language	Domains	#Entities
German	General	107417
English	Hotel,Electronics	8660
Spanish	Hotel,Electronics	1041
Catalan	Hotel,Electronics	1358
Portuguese	Hotel,Electronics	1387
French	Hotel,Electronics	651

Table 5: Summary of the lexicons in the LRP

Corpora		
Language	Domains	#Entities
English	Hotel,Electronics	22373
Spanish	Hotel,Electronics	18191
Catalan	Hotel,Electronics	4707
Portuguese	Hotel,Electronics	6244
French	Electronics	22841

Table 6: Summary of the corpora in the LRP

lemon (e.g. *Siamese_cat*), which provides a formalism for describing its decomposition into their component words.

- Onyx is extensible and the same formalism can be used for different emotion descriptions and it is not tied to a particular emotion representation, as the other affect lexicons. Following the EmotionML model, Onyx uses pluggable ontologies (vocabularies) that complement the central ontology.

This formalism has been used to generate the Eurosentiment dataset as detailed in [44, 45]. This dataset is composed of fourteen domain-specific opinion and emotion lexicons covering six languages (German, English, Spanish, Catalan, Portuguese and French) and two domain (Hotels and Electronics) as shown in Table 5.

4.2. Emotion annotations in corpora

One common use case for affective technologies [19] is the annotation of material involving emotionality, such as texts, videos or speech recordings.

Onyx has been used for emotion annotation of corpora as described in [44, 46, 45]⁶. An overview of this lexical resources is provided in Table 6.

These corpora were used as gold standard for the evaluation of the services in the Eurosentiment [45] pool. The evaluation material is publicly available⁷ and can be used as an example of evaluating semantic emotion analysis services using semantic resources in Onyx.

⁶A tool for translating other formats to Eurosentiment is available at <http://eurosentiment.readthedocs.org/en/latest/corpusconverter.html>

⁷<https://github.com/EuroSentiment/evaluation>

4.3. Emotion service specification

Defining a common service API and format is important for interoperability between services and to boost the ecosystem of emotion analysis services. This section shows how this can be achieved through a combination of NIF, originally intended for NLP services, and Onyx for annotation of emotion.

NLP Interchange Format (NIF) 2.0 [30] defines a semantic format and API for improving interoperability among natural language processing services. The classes to represent linguistic data are defined in the NIF Core Ontology. All ontology classes are derived from the main class *nif:String* which represents strings of Unicode characters. One important subclass of *nif:String* is the *nif:Context* class. It represents the whole string of the text and is used to calculate the indices of the substrings. There are other classes (*nif:Word*, *nif:Sentence*, *nif:Phrase*) for representing partitions of a text. NIF individuals are identified by URIs following a *nif:URIScheme* which restricts URI's syntax. NIF can be extended via vocabularies modules. It uses Marl for sentiment annotations and Onyx have been proposed as a NIF vocabulary for emotions.

In addition, NIF defines an input and output format for REST web services in the NIF 2.0 public API specification. This specification defines a set of parameters that should be supported by NIF compliant services.

Listing 4 shows the output of a service call with the input parameter value “My iPad is an awesome device”⁸

The main benefit of using NIF for emotion services is that NIF compliant services can be easily combined as shown in the NIF combinator [47]⁹, where well known NLP tools are combined thanks to the use of NIF wrappers.

To demonstrate the creation of an emotion analysis service using this specification, we developed a proof-of-concept web service on top of the Synesketch [48] library.

Synesketch is a library and application that detects emotions in English texts and generates 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. 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 six categories of emotion proposed by Paul Ekman [9]. This model is included among the vocabularies of EmotionML under the name big6, and can be represented in Onyx as shown in Section 3.2.2. Each emotion is 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. These two attributes together show how *positive*, *negative* or *neutral* the overall emotion is.

⁸More details about the service output of sentiment and emotion services based on NIF, Marl and Onyx can be found at <http://eurosentiment.readthedocs.org/en/latest/format/servicesformat.html>

⁹The NIF combinator is available at <http://demo.nlp2rdf.aksw.org/>

Listing 4: Service output expressed in NIF using the vocabulary Onyx.

```
{
  "@context": [ "http://example.com/context.jsonld" ],
  "analysis": [{
    "@id": "http://example.com/analyse",
    "@type": [ "onyx:EmotionAnalysis" ],
    "dc:language": "en",
    "onyx:maxEmotionIntensity": 1.0,
    "onyx:minEmotionIntensity": 0.0
    "prov:wasAssociatedWith": "http://www.gsi.dit.upm.es/"
  }],
  "entries": [{
    "@id": "http://example.com/analyse?input=My%20ipad%20is%20an%20awesome%20device",
    "dc:language": "es",
    "opinions": [],
    "emotions": [{
      "onyx:aboutObject": "http://dbpedia.org/page/IPad"
      "prov:generatedBy": "http://example.com/analyse",
      "onyx:hasEmotion": [{
        "onyx:hasEmotionCategory": "wna:liking",
        "onyx:hasEmotionIntensity": 0.7
      }, {
        "onyx:hasEmotionCategory": "wna:approval",
        "onyx:hasEmotionIntensity": 0.1 }]]],
    "nif:isString": "My ipad is an awesome device",
    "prov:generatedBy": "http://example.com/analyse",
    "strings": [{
      "@id": "http://example.com/analyse?input=My%20ipad%20is%20an%20awesome%20device#char=3,6",
      "itsrdf:taIdentRef": "http://dbpedia.org/page/IPad",
      "nif:anchorOf": "ipad" }]]]
}
```

The Synesketech weight directly maps 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 add an *AggregatedEmotion* to the *EmotionSet* 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 website¹⁰.

Several other implementations have been developed in the context of the Eurosentiment project, mainly as wrappers of already existing resources. The majority of them uses the WordNet-Affect categories, although there are some that required the specification of ad-hoc categories.

4.4. Emotion mapping and rules

Some of the potential benefits of the use of ontologies [49] are: mapping between different emotion representations, the definition of the relationship between concepts in an emotion description and emotion composition. Each of these topics is a subject of study in its own right. However, we want to illustrate how Onyx's semantic approach could be used in this direction. In particular, we will focus on SPARQL Inference Notation (SPIN)¹¹ rules.

¹⁰<http://demos.gsi.dit.upm.es/onyxremote/>

¹¹<http://spinrdf.org/spin.html>

SPIN is a powerful tool to create rules and logical constraints to any entity using standard SPARQL queries. It does this by using a set of RDF classes and properties defined in its specification. For our purposes, the simplest use of SPIN is to attach a SPARQL query (rule from now on) to a certain class, so that it is performed on entity creation, update or deletion. What is interesting about SPIN is that it has mechanisms to generalize this procedure, and to create templates from these rules so that they can be applied to several classes. An Open Source Java API¹² is also publicly available which can be used to test the examples in this section.

Exploring the full potential of SPIN is out of the scope of this paper. However, we will cover two ways in which Onyx and SPIN can be used together to provide more flexibility than any non-semantic approach could. For the sake of clarity, we will only include the relevant rules that should be used in each case, rather than the complete SPIN RDF excerpt.

The first example (Listing 5) shows the rule that should be added to the *EmotionSet* class in order to infer a complex Plutchik emotion from two basic ones. In particular, it annotates an emotion with the Optimism category if it is already annotated with Anticipation and Joy. The query also shows how to add new entities and specifying their URIs using a random identifier and the base URI.

Listing 5: Automatic composition of emotions. Anticipation and Joy result in Optimism.

```
INSERT {
  ?this onyx:hasEmotion ?emotion .
  ?emotion a onyx:Emotion .
  ?emotion a onyx:AggregatedEmotion .
  ?emotion onyx:hasEmotionCategory plutchik:Optimism .
}
WHERE {
  ?this onyx:hasEmotion _:0 .
  ?this onyx:hasEmotion _:1 .
  _:0 onyx:hasEmotionCategory plutchik:Joy .
  _:1 onyx:hasEmotionCategory plutchik:Anticipation .
  BIND (IRI(CONCAT("http://example.org/id/#Emotion", SUBSTR
    (str(RAND()), 3, 16))) AS ?emotion) .
}
```

The second example (Listing 6) assigns an emotional category based on dimensional values. It takes as an example the Hourglass of Emotions [11] model where second-level emotions can be expressed as a combination of two sentic levels. In this case, a positive level of attention and pleasantness results in optimism. Its dimensions and the main emotional categories have been represented with Onyx (Section 3.2.3).

Finally, by using these two rules together and a simple mapping of both hourglass' and Plutchik's Optimism categories, it is possible to find optimistic results among entries annotated using Plutchik's basic categories and entries annotated using the dimensions from the Hourglass of Emotions.

¹²<http://topbraid.org/spin/api/>

Listing 6: Applying categories based on the dimensions of an emotion.

```
INSERT {  
  ?this onyx:hasEmotionCategory hg:Optimism .  
}  
WHERE {  
  ?this hg:Pleasantness ?p .  
  ?this hg:Attention ?a .  
  FILTER ( ?p > 0 && ?a > 0 )
```

5. Conclusions and future work

With this work we have introduced the Onyx ontology, which can represent emotions taking advantage of the work conducted in the Web of Data (Linked Data) and in emotion research community (EmotionML). The ontology is extendable through pluggable vocabularies that enable its adaptation to different emotion models and application domains.

This paper discusses several applications and use cases, such as the emotion annotation of lexical entries and corpora, the specification of inter-operable emotion analysis services as well as the definition of emotion mappings. To this end, Onyx has been linked with vocabularies such as *lemon*, NIF and the Provenance Ontology. One remarkable example is the role of Onyx in the R&D European project Eurosentiment¹³, whose aim is the creation of a language resource pool for Sentiment Analysis. As a result, a set of lexical resources are publicly available annotated with Onyx.

A key factor for the adoption of proposals such as Onyx is its use and extension by the emotion research community. To this end, the W3C Community Group *Linked Data Models for Emotion and Sentiment Analysis*¹⁴ has been set up with participants from industry and academia. We hope that it will contribute to the evolution of Onyx, shaping it to address the issues arising from feedback of the community. The next step for Onyx is to go beyond the applications in this paper (Section 4) and be leveraged in state-of-the-art sentiment analysis applications. As a start point, it could be used in any of the several sentiment analysis challenges available. In particular, it is a perfect fit for semantic sentiment analysis challenges such as SemEval¹⁵.

For the most part, Onyx relies on the main elements introduced by EmotionML. That seems to be the case with HEO as well, which explains the similarities between both ontologies. Given these similarities, we will discuss some of the differences between both ontologies that justify the use of Onyx.

First and foremost, Onyx intends to be a model as generic as possible. It has been integrates with other ontologies, When in doubt about a property or concept, we chose to leave it outside the ontology and link to other vocabularies (e.g. OpenAnnotation[50], NIF [47]). It also integrates perfectly with the provenance ontology, adding great value to the ontology. Although HEO also provides a generic model of emotion, it is tailored to a specific use case and includes concepts that would now be better expressed with other ontologies.

¹³<http://eurosentiment.eu>

¹⁴<http://www.w3.org/community/sentiment/>

¹⁵SemEval's sentiment analysis track: [http://alt.qcri.org/semeval2015/index.php?id=](http://alt.qcri.org/semeval2015/index.php?id=tasks)
tasks

For instance, Open Annotation or Prov-O for annotation. The fact that Onyx follows an approach analogous to Marl also facilitates the integration of Opinion and Emotion, which is crucial for sentiment analysis.

There are other practical differences between the two, such as action tendencies and appraisal being properties in Onyx. This reduces the number of nodes necessary for annotation, and makes annotations more convenient.

These differences prove that HEO and Onyx are actually very different. Not only on the ontological level, but also in approach. We believe Onyx was a missing piece in the Linked Data puzzle, it integrates with other ontologies and covers aspects that no other ontology for emotions did.

Moreover, the presented applications show the applicability and usefulness of the ontology. Using the concepts shown in Section 4.4, it would be possible to combine resources that use different emotion models. For instance, an application could leverage the power of WordNet-Affect [16], EmoLex [15] (Plutchik's categories) and DepecheMood [51] (Ekman's categories).

As future work, we will study how emotion synthesis and emotional embodied conversational agents can be applied in e-learning. In particular, our aim is to explore how conversational agents can benefit from integration of semantic, emotion and user models for improving user engagement. There are two interesting aspects in this sense. The first one is using a behavioral model based on emotions to interact with students on a deeper level. For this, we will build on our experience with conversational agents and make use of Onyx annotations to reason about emotions. The second one is trying to gain a better understanding of students. This analysis will build on the concepts in Section 4.4 to use various models that characterize different aspects of emotions.

6. Acknowledgements

This work was partially funded by the EUROSENTIMENT FP7 Project (Grant Agreement no: 296277). The authors want to thank the reviewers for their thorough and helpful feedback, which improved the quality of this paper. Lastly, we thank Paul Buitelaar and Gabriela Vulcu as well for their contribution to the representation of lexical information using *lemon*.

References

- [1] B. Pang, L. Lee, Opinion mining and sentiment analysis, Foundations and trends in information retrieval 2 (1-2) (2008) 1–135.
- [2] G. Tummarello, R. Delbru, E. Oren, Sindice.com: Weaving the open linked data, in: The Semantic Web, Springer, 552–565, 2007.
- [3] A. Westerski, C. A. Iglesias, F. Tapia, Linked Opinions: Describing Sentiments on the Structured Web of Data, in: Proceedings of the 4th International Workshop Social Data on the Web, vol. Vol-830, 13–23, URL onlineatCEUR-WS.org/Vol-830, 2011.
- [4] S. G. Barsade, The Ripple Effect: Emotional Contagion and its Influence on Group Behavior, Administrative Science Quarterly 47 (4) (2002) 644–675.

- [5] A. D. Kramer, J. E. Guillory, J. T. Hancock, Experimental evidence of massive-scale emotional contagion through social networks, *Proceedings of the National Academy of Sciences* (2014) 201320040 URL <http://www.pnas.org/content/early/2014/05/29/1320040111.short>.
- [6] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, B. Bhattacharjee, Measurement and analysis of online social networks, in: *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement, IMC '07*, ACM, New York, NY, USA, ISBN 978-1-59593-908-1, 29–42, 2007.
- [7] E. Cambria, B. Schuller, Y. Xia, C. Havasi, New Avenues in Opinion Mining and Sentiment Analysis, *Intelligent Systems, IEEE* 28 (2) (2013) 15–21.
- [8] M. Munezero, C. Montero, E. Sutinen, J. Pajunen, Are They Different? Affect, Feeling, Emotion, Sentiment, and Opinion Detection in Text, *Affective Computing, IEEE Transactions on* 5 (2) (2014) 101–111.
- [9] P. Ekman, Basic emotions, *Handbook of cognition and emotion* 98 (1999) 45–60.
- [10] J. J. Prinz, *Gut reactions: A perceptual theory of emotion*, Oxford University Press, 2004.
- [11] E. Cambria, A. Livingstone, A. Hussain, The hourglass of emotions, in: *Cognitive Behavioural Systems*, Springer, 144–157, 2012.
- [12] R. Plutchik, *Emotion: A psychoevolutionary synthesis*, Harper & Row New York, 1980.
- [13] D. Borth, T. Chen, R. Ji, S.-F. 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*, ACM, New York, NY, USA, ISBN 978-1-4503-2404-5, 459–460, URL <http://doi.acm.org/10.1145/2502081.2502268>, 2013.
- [14] E. Cambria, C. Havasi, A. Hussain, SenticNet 2: A Semantic and Affective Resource for Opinion Mining and Sentiment Analysis., in: *FLAIRS Conference*, 202–207, 2012.
- [15] S. M. Mohammad, P. D. Turney, Crowdsourcing a Word-Emotion Association Lexicon 29 (3) (2013) 436–465.
- [16] C. Strapparava, A. Valitutti, WordNet-Affect: an affective extension of WordNet, in: *Proceedings of LREC*, vol. 4, 1083–1086, 2004.
- [17] M. Schröder, H. Pirker, M. Lamolle, First suggestions for an emotion annotation and representation language, in: *Proceedings of LREC*, vol. 6, 88–92, 2006.
- [18] M. Schröder, L. Devillers, K. Karpouzis, J.-C. Martin, C. Pelachaud, C. Peter, H. Pirker, B. Schuller, J. Tao, I. Wilson, What should a generic emotion markup language be able to represent?, in: *Affective Computing and Intelligent Interaction*, Springer, 440–451, 2007.

- [19] P. Baggia, C. Pelachaud, C. Peter, E. Zovato, Emotion Markup Language (EmotionML) 1.0 W3C Recommendation, Tech. Rep., W3C, URL <http://www.w3.org/TR/emotionml/>, 2014.
- [20] H. N. of Excellence, HUMAINE Emotion Annotation and Representation Language (EARL): Proposal., Tech. Rep., HUMAINE Network of Excellence, URL <http://emotion-research.net/projects/humaine/earl/proposal#Dialects>, 2006.
- [21] K. Ashimura, P. Baggia, F. Burkhardt, A. Oltramari, C. Peter, E. Zovato, EmotionML vocabularies, Tech. Rep., W3C, URL <http://www.w3.org/TR/2012/NOTE-emotion-voc-20120510/>, 2012.
- [22] J. Hastings, W. Ceusters, B. Smith, K. Mulligan, Dispositions and Processes in the Emotion Ontology, in: Proc. ICBO, 71–78, 2011.
- [23] C. Chiarcos, S. Hellmann, S. Nordhoff, Towards a linguistic linked open data cloud: The Open Linguistics Working Group, TAL 52 (3) (2011) 245–275.
- [24] C. Chiarcos, J. McCrae, P. Cimiano, C. Fellbaum, Towards open data for linguistics: Linguistic linked data, in: New Trends of Research in Ontologies and Lexical Resources, Springer, 7–25, 2013.
- [25] P. Buitelaar, P. Cimiano, J. McCrae, E. Montiel-Ponsoda, T. Declerck, Ontology lexicalisation: The lemon perspective, in: Workshop at 9th International Conference on Terminology and Artificial Intelligence (TIA 2011), 33–36, 2011.
- [26] C. Fellbaum, WordNet, Wiley Online Library, 1998.
- [27] C. Bizer, J. Lehmann, G. Kobilarov, S. Auer, C. Becker, R. Cyganiak, S. Hellmann, DBpedia-A crystallization point for the Web of Data, Web Semantics: science, services and agents on the world wide web 7 (3) (2009) 154–165.
- [28] C. Chiarcos, S. Hellmann, S. Nordhoff, Linking linguistic resources: Examples from the open linguistics working group, in: Linked Data in Linguistics, Springer, 201–216, 2012.
- [29] C. Chiarcos, POWLA: Modeling linguistic corpora in OWL/DL, in: The Semantic Web: Research and Applications, Springer, 225–239, 2012.
- [30] S. Hellmann, Integrating Natural Language Processing (NLP) and Language Resources using Linked Data, Ph.D. thesis, Universität Leipzig, 2013.
- [31] C. Chiarcos, Ontologies of Linguistic Annotation: Survey and perspectives., in: LREC, 303–310, 2012.
- [32] L. Moreau, B. Clifford, J. Freire, J. Futrelle, Y. Gil, P. Groth, N. Kwasnikowska, S. Miles, P. Missier, J. Myers, B. Plale, Y. Simmhan, E. Stephan, J. Van den Bussche, The Open Provenance Model core specification (v1.1), Future Generation Computer Systems 27 (6) (2011) 743 – 756, ISSN 0167-739X, URL <http://www.sciencedirect.com/science/article/pii/S0167739X10001275>.

- [33] P. Groth, L. Moreau, Prov-O W3C Recommendation, Tech. Rep., W3C, URL <http://www.w3.org/TR/prov-o/>, 2013.
- [34] M. Schröder, The SEMAINE API: Towards a Standards-Based Framework for Building Emotion-Oriented Systems, *Advances in Human-Computer Interaction* 2010.
- [35] J. McCrae, D. Spohr, P. Cimiano, Linking Lexical Resources and Ontologies on the Semantic Web with Lemon, in: G. Antoniou, M. Grobelnik, E. Simperl, B. Parsia, D. Plexousakis, P. Leenheer, J. Pan (Eds.), *The Semantic Web: Research and Applications*, vol. 6643 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, ISBN 978-3-642-21033-4, 245–259, 2011.
- [36] A. Esuli, F. Sebastiani, Sentiwordnet: A publicly available lexical resource for opinion mining, in: *Proceedings of LREC*, vol. 6, 417–422, 2006.
- [37] E. Cambria, R. Speer, C. Havasi, A. Hussain, SenticNet: A Publicly Available Semantic Resource for Opinion Mining., in: *AAAI Fall Symposium: Commonsense Knowledge*, vol. 10, 02, 2010.
- [38] J. Wiebe, T. Wilson, C. Cardie, Annotating expressions of opinions and emotions in language, *Language resources and evaluation* 39 (2-3) (2005) 165–210.
- [39] U. Waltinger, GermanPolarityClues: A Lexical Resource for German Sentiment Analysis., in: *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC)*, 1638–1642, 2010.
- [40] M. M. Bradley, P. J. Lang, Affective norms for English words (ANEW): Instruction manual and affective ratings, Tech. Rep., Citeseer, 1999.
- [41] S. Mohammad, P. D. Turney, Crowdsourcing a Word-Emotion Association Lexicon, *Computation Intelligence* 29 (3) (2013) 436–465.
- [42] E. Cambria, D. Olsher, D. Rajagopal, SenticNet 3: a common and common-sense knowledge base for cognition-driven sentiment analysis, in: *Twenty-eighth AAAI conference on artificial intelligence*, 1515–1521, 2014.
- [43] P. Buitelaar, M. Arcan, C. A. Iglesias, J. F. Sánchez-Rada, C. Strapparava, Linguistic Linked Data for Sentiment Analysis, in: *2nd Workshop on Linked Data in Linguistics (LDL-2013): Representing and linking lexicons, terminologies and other language data*, Pisa, Italy, 1–8, 2013.
- [44] V. Gabriela, B. Paul, N. Sapna, P. Bianca, A. Mihael, C. Barry, S. J. Fernando, I. C. A., Generating Linked-Data based Domain-Specific Sentiment Lexicons from Legacy Language and Semantic Resources, in: *Proceedings of the 5th International Workshop on Emotion, Social Signals, Sentiment and Linked Open Data (ES3LOD)*, co-located with LREC 2014, Reykjavik, Iceland, 6–9, 2014.
- [45] J. F. Sánchez-Rada, G. Vulcu, C. A. Iglesias, P. Buitelaar, EUROSENTIMENT: Linked Data Sentiment Analysis, in: J. v. O. Matthew Horridge, Marco Rospocher (Ed.), *Proceedings of the 13th International Semantic*

Web Conference ISWC 2014 Posters and Demonstrations Track, 145–148, URL http://ceur-ws.org/Vol-1272/paper_116.pdf, 2014.

- [46] V. Gabriela, B. Paul, N. Sapna, P. Bianca, A. Mihael, C. Barry, S. J. Fernando, I. C. A., Linked-Data based Domain-Specific Sentiment Lexicons, in: Proceedings of 3rd Workshop on Linked Data in Linguistics: Multilingual Knowledge Resources and Natural Language Processing, co-located with LREC 2014, 77–81, 2014.
- [47] S. Hellmann, J. Lehmann, S. Auer, M. Nitzschke, NIF Combinator: Combining NLP Tool Output, in: Proceedings of the 18th International Conference on Knowledge Engineering and Knowledge Management, EKAW’12, Springer-Verlag, Berlin, Heidelberg, 446–449, 2012.
- [48] U. Krcadinac, P. Pasquier, J. Jovanovic, V. Devedzic, Synesketch: An Open Source Library for Sentence-based Emotion Recognition, IEEE Transactions on Affective Computing 4 (3) (2013) 312–325.
- [49] M. Schröder, The SEMAINE API: A Component Integration Framework for a Naturally Interacting and Emotionally Competent Embodied Conversational Agent, Ph.D. thesis, Universität des Saarlandes, 2011.
- [50] P. Ciccarese, S. Soiland-Reyes, T. Clark, Web Annotation as a First-Class Object, Internet Computing, IEEE 17 (6) (2013) 71–75, ISSN 1089-7801.
- [51] J. Staiano, M. Guerini, DepecheMood: a Lexicon for Emotion Analysis from Crowd-Annotated News, arXiv preprint arXiv:1405.1605 URL <http://arxiv.org/abs/1405.1605>, 00002.