

Sentic Computing for social media marketing

Erik Cambria · Marco Grassi ·
Amir Hussain · Catherine Havasi

Published online: 19 May 2011
© Springer Science+Business Media, LLC 2011

Abstract In a world in which millions of people express their opinions about commercial products in blogs, wikis, fora, chats and social networks, the distillation of knowledge from this huge amount of unstructured information can be a key factor for marketers who want to create an image or identity in the minds of their customers for their product, brand or organization. Opinion mining for product positioning, in fact, is getting a more and more popular research field but the extraction of useful information from social media is not a simple task. In this work we merge AI and Semantic Web techniques to extract, encode and represent this unstructured information. In particular, we use Sentic Computing, a multi-disciplinary approach to opinion mining and sentiment analysis, to semantically and affectively analyze text and encode results in a semantic aware format according to different web ontologies. Eventually we represent this information as an interconnected knowledge base which is browsable through a multi-faceted classification website.

Keywords AI · Semantic Web · Knowledge base management · NLP · Opinion mining and sentiment analysis

E. Cambria (✉) · A. Hussain
Department of Computing Science and Mathematics,
University of Stirling, Stirling, FK94LA, UK
e-mail: eca@cs.stir.ac.uk
URL: <http://cs.stir.ac.uk/~eca>

A. Hussain
e-mail: ahu@cs.stir.ac.uk

M. Grassi
Department of Biomedical, Electronic and Telecommunication Engineering,
Università Politecnica delle Marche, Ancona, 60131, Italy
e-mail: m.grassi@univpm.it

C. Havasi
MIT Media Laboratory, Massachusetts Institute of Technology,
Boston, MA 02139-4307, USA
e-mail: havasi@media.mit.edu

1 Introduction

The passage from a read-only to a read-write Web made users more enthusiastic about interacting, sharing and collaborating through social networks, on-line communities, blogs, wikis and other on-line collaborative media. In the last years this collective intelligence has spread to many different areas in the Web, with particular focus on fields related to our everyday life such as commerce, tourism, education and health. The on-line review of commercial services and products, in particular, is an action that users usually perform with pleasure, to share their opinions about services they have received or products they have just bought, and it constitutes immeasurable value for other potential buyers. This trend opened new doors to enterprises that want to reinforce their brand and product presence in the market by investing in on-line advertising and positioning i.e. in social media marketing (SMM).

What mainly makes SMM work is the buzz mechanism. A buzz replicates a message through user-to-user contact, rather than purchasing some advertising or promoting a press release. The message does not have to necessarily deal with the product. Many successful viral campaigns, in fact, have spread thanks to a compelling message, with the company logo included incidentally. At the heart of buzz is an understanding that the natural, spontaneous networks that comprise the social universe are the most effective means of reaching people in a meaningful way. The power of marketing lies, therefore, not in pushing information to the masses but in effectively tapping those individuals who wield influence over others. The marketers who are winning are the ones using consumers and culture to their advantage, crafting messages with consumers rather than throwing messages at them [20].

In confirmation of the growing interest in SMM, several commercial tools have been recently developed by Umbria (<http://umbrialistens.com>), Cymfony (<http://www.cymfony.com>), Nielsen (<http://nielsen-online.com>) and Evolve24 (<http://evolve24.com>, among many others) to provide companies with a way to analyze the blogosphere on a large scale in order to extract information about the trend of the opinions relative to their products. Nevertheless most of the existing tools and the research efforts are limited to a polarity evaluation or a mood classification according to a very limited set of emotions. In addition, such methods mainly rely on parts of text in which emotional states are explicitly expressed and hence they are unable to capture opinions and sentiments that are expressed implicitly.

The main aim of this research work is to provide marketers with a novel SMM tool for the management of social media information at semantic level, able to capture both opinion polarity and affective information associated with user-generated contents. A polarity value associated to an opinion, in fact, sometimes can be restrictive. Enriching automatic analysis of social media with affective labels such as ‘joy’ or ‘disgust’ can help marketers to have a clearer idea of what their customers think about their products. In particular, we selected YouTube (YT, <http://youtube.com>) as a social media source since, with its over 2 billion views per day, 24 h of video uploaded every minute and 15 min a day spent by the average user, it represents more than 40% of on-line video market (<http://viralblog.com/research/youtube-statistics>). Specifically, we focused on video reviews of mobile phones because of the quantity and the quality of the comments usually associated with them. The social media analysis is performed through three main steps: firstly comments are analyzed using Sentic Computing [4], a multi-disciplinary approach to opinion mining and sentiment

analysis (Section 3), secondly the extracted information is encoded on the base of different web ontologies (Section 4), and eventually the resulting knowledge base is made available for browsing through a multi-faceted classification website (Section 5).

2 Opinion mining and sentiment analysis

Opinion mining and sentiment analysis are new disciplines that have recently raised more and more interest, especially in fields such as marketing, personal affective profiling and financial market prediction. Although often discussed as a unique discipline in the sense of extracting people's attitudes from text, opinion mining focuses on opinion polarity detection while sentiment analysis on emotion inference. Given a generic resource (textual, visual, audio or multi-modal) containing a set of opinions O about a set of topics T with different polarity $p \in [-1, 1]$, we define opinion mining as the process that aims to extract, for each $t \in T$, the subset of opinions $o \subset O$ concerning t and determine p . Equally, we define sentiment analysis as the process aiming to infer the affective information (generally an emotional label) associated with each $o \subset O$.

In this section, we briefly present some of the most relevant approaches to automatic identification and extraction of opinions and sentiments from text. A complete review of the state of the art goes beyond the purpose of this paper and can be found in [17]. Most of the existing systems focus on polarity classification of opinions (typically positive, negative, neutral) in a target document [18, 23, 25]. Other approaches deal with extracting the most relevant text fragments containing subjective opinions using machine learning approaches [1] as well as feature selection [16]. Some works, eventually, focus on extraction of moods from informal text resources such as blogs [15, 22].

More generally, existing approaches to opinion mining and sentiment analysis can be grouped into three main categories: keyword spotting, in which text is classified into categories based on the presence of fairly unambiguous words [7, 27], lexical affinity, which assigns arbitrary words a probabilistic affinity for a particular concept [21, 28], and statistical methods, which calculate the valence of keywords, punctuation and word co-occurrence frequencies on the base of a large training corpus [9, 13]. The problem with these approaches is that they mainly rely on parts of text in which opinions are explicitly expressed such as positive terms (e.g. good, nice, excellent, fortunate, correct, superior, best) and negative terms (e.g. bad, nasty, poor, unfortunate, wrong, inferior, worst). In general, in fact, opinions are expressed implicitly through context and domain dependent concepts, which make keyword-based approaches ineffective.

3 Sentic Computing

Sentic Computing is a multi-disciplinary approach to opinion mining and sentiment analysis that exploits both computer and social sciences to better recognize, interpret and process opinions and sentiments over the Web. Sentic Computing, whose term

derives from the Latin *sentire* (root of words such as sentiment and sentience) and *sensus* (intended as common sense), overcomes some of the limits of existing tools by revealing the implicit content hidden in texts. To this end, Sentic Computing uses a common sense reasoning [3] approach coupled with a novel emotion categorization model [5], which enable the analysis of documents not only at page or paragraph level but also at sentence level. In particular, Sentic Computing involves the use of AI and Semantic Web techniques, for knowledge representation and inference; mathematics, for tasks such as graph mining and multi-dimensionality reduction; linguistics, for discourse analysis and pragmatics; psychology, for cognitive and affective modeling; sociology, for understanding social network dynamics and social influence; finally ethics, for related issues about the nature of mind.

3.1 Common sense

When people communicate with each other, they rely on similar background knowledge e.g. the way objects relate to each other in the world, people’s goals in their daily lives, the emotional content of events or situations. This ‘taken for granted’ information is what we call common sense—obvious things people normally know and usually leave unstated.

The Open Mind Common Sense project has been collecting this kind of knowledge from volunteers on the Internet since 2000 to provide intuition to AI systems and applications. ConceptNet [11] represents the information in the Open Mind corpus as a directed graph in which the nodes are concepts and the labeled edges are assertions of common sense that interconnect them (Fig. 1).

3.2 The hourglass of emotions

This model is a variant of Plutchik’s emotion categorization [19] and constitutes an attempt to emulate Marvin Minsky’s conception of emotions. Minsky sees the mind as made of thousands of different resources and believes that our emotional states result from turning some set of these resources on and turning another set of them off

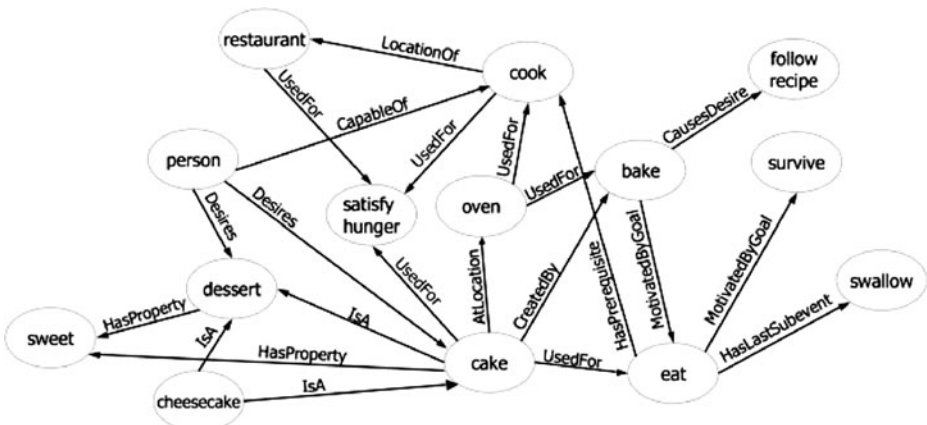


Fig. 1 A sketch of ConceptNet

[14]. Each such selection changes how we think by changing our brain’s activities: the state of anger, for example, appears to select a set of resources that help us react with more speed and strength while also suppressing some other resources that usually make us act prudently.

The Hourglass of Emotions (Fig. 2) is specifically designed to recognize, understand and express emotions in the context of human-computer interaction (HCI). In the model, in fact, affective states are not classified, as often happens in the field of emotion analysis, into basic emotional categories, but rather into four concomitant but independent dimensions—Pleasantness, Attention, Sensitivity and Aptitude—in order to understand how much respectively:

1. the user is amused by interaction modalities (Pleasantness)
2. the user is interested in interaction contents (Attention)
3. the user is comfortable with interaction dynamics (Sensitivity)
4. the user is confident in interaction benefits (Aptitude)

Each affective dimension is characterized by six levels of activation, called ‘sentic levels’, which determine the intensity of the expressed/perceived emotion. The concomitance of the different affective dimensions makes possible the generation of compound emotions such as ‘love’, which is given by the sum of ‘joy’ and ‘trust’, or ‘aggressiveness’, given by the concomitance of ‘anger’ and ‘anticipation’.




Fig. 2 The hourglass of emotions

3.3 The sentsics extraction process

The process for the inference of opinions and sentiments from text (illustrated in Fig. 3) numbers three main components: a NLP module, which performs a first skim of the document, a Semantic Parser, whose aim is to extract concepts from the lemmatized text, and AffectiveSpace [2], a multi-dimensional vector space clustered according to the Hourglass model.

The NLP module interprets all the affective valence indicators usually contained in text such as special punctuation, complete upper-case words, onomatopoeic repetitions, exclamation words, negations, degree adverbs and emoticons, and eventually lemmatizes text.

The Semantic Parser then deconstructs text into concepts and provides, for each of them, the relative frequency, valence and status i.e. the concept's occurrence in the text, its positive or negative connotation, and the degree of intensity with which the concept is expressed.

AffectiveSpace eventually calculates, for each retrieved concept, the resulting 'sentic vector' i.e. the relative affective information in terms of Pleasantness, Attention, Sensitivity and Aptitude, which is ultimately exploited to infer moods and polarity.

AffectiveSpace, the core of the Sentsics Extraction Process, is built by applying the blending technique [12] over ConceptNet and WordNet-Affect (WNA) [24], a linguistic resource for the lexical representation of sentiments (see Section 4.3 for more details). This alignment operation yields a new dataset in which common sense and affective knowledge coexist i.e. a matrix, A , whose rows are concepts (e.g. 'dog' or 'bake cake'), whose columns are either common sense and affective features (e.g. 'isA-pet' or 'hasEmotion-joy'), and whose values indicate truth values of assertions.

After performing truncated singular value decomposition (TSVD) [26] on this data, we obtain a new matrix $\tilde{A} = U_k * \Sigma_k * V_k^T$, which forms a low-rank approximation of the original data. This approximation is based on minimizing the Frobenius

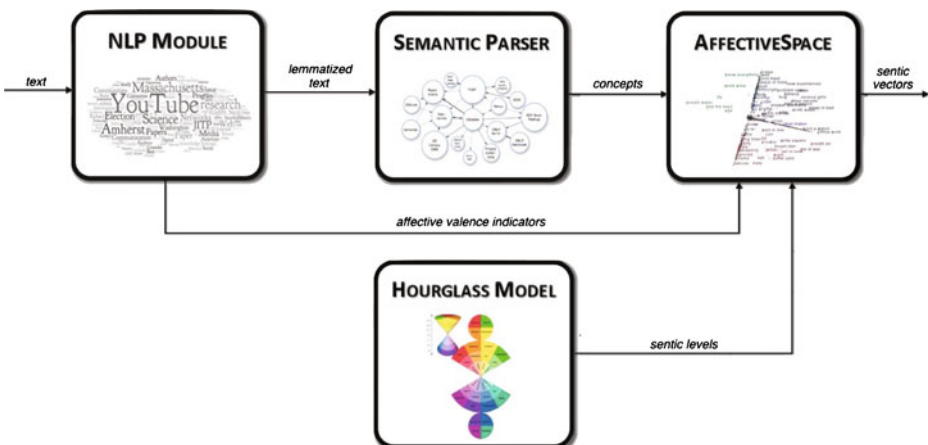


Fig. 3 Sentsics extraction process

norm of the difference between A and \tilde{A} under the constraint $rank(\tilde{A}) = k$. For the Eckart–Young theorem [6] it represents the best approximation of A in the mean-square sense, in fact:

$$\min_{\tilde{A}|rank(\tilde{A})=k} \|A - \tilde{A}\| = \min_{\tilde{A}|rank(\tilde{A})=k} \|\Sigma - U^* \tilde{A} V\| = \min_{\tilde{A}|rank(\tilde{A})=k} \|\Sigma - S\| \quad (1)$$

assuming that \tilde{A} has the form $\tilde{A} = USV^*$, where S is diagonal.

From the rank constraint, i.e. S has k non-zero diagonal entries, the minimum of the above statement is obtained as follows:

$$\min_{s_i} \sqrt{\sum_{i=1}^n (\sigma_i - s_i)^2} = \min_{s_i} \sqrt{\sum_{i=1}^k (\sigma_i - s_i)^2 + \sum_{i=k+1}^n \sigma_i^2} = \sqrt{\sum_{i=k+1}^n \sigma_i^2} \quad (2)$$

Therefore, \tilde{A} of rank k is the best approximation of A in the Frobenius norm sense when $\sigma_i = s_i$ ($i = 1, \dots, k$) and the corresponding singular vectors are same as those of A . We use a trial and error approach to set $k = 50$ and hence obtain AffectiveSpace (Fig. 4), a 50-dimensional space in which different vectors represent different ways of making binary distinctions among concepts and emotions.

In AffectiveSpace common sense and affective knowledge are in fact combined, not just concomitant, i.e. everyday life concepts like ‘have breakfast’, ‘meet people’ or ‘watch tv’ are linked to a hierarchy of affective domain labels.

By exploiting the information sharing property of TSVD, concepts with the same affective valence are likely to have similar features i.e. concepts concerning the same




Fig. 4 A sketch of AffectiveSpace

opinion tend to fall near each other in the vector space. Concepts and emotions are represented by vectors of 50 coordinates: these coordinates can be seen as describing concepts in terms of ‘eigenmoods’ that form the axes of AffectiveSpace i.e. the basis e_0, \dots, e_{49} of the vector space. For example, the most significant eigenmood, e_0 , represents concepts with positive affective valence. That is, the larger a concept’s component in the e_0 direction is, the more affectively positive it is likely to be. Consequently concepts with negative e_0 components have negative affective valence.

After organizing AffectiveSpace according to the Hourglass model (using a k-means clustering approach), each retrieved concept is projected into the vector space and hence its affective valence is calculated, for each sentic dimension, according to the position the concept occupies in the space. This information is encoded as four dimensional vectors (sentic vectors) which are eventually averaged in order to infer moods and polarity. The latter, in particular, is calculated according to the formula:

$$p = \sum_{i=1}^N \frac{Pleasantness(o_i) + |Attention(o_i)| - |Sensitivity(o_i)| + Aptitude(o_i)}{9N} \quad (3)$$

where N is the total number of retrieved concepts and 9 is the normalization factor (the maximum value of the numerator in fact is given by the sentic vector $[3, \pm 3, 0, 3]$ and the minimum by $[-3, 0, \pm 3, -3]$). In the formula, Attention and Sensitivity are taken in absolute value since, from the point of view of polarity rather than affection, all of their sentic values represent positive and negative values respectively (e.g. ‘anger’ is positive in the sense of level of activation of Sensitivity but negative in terms of polarity and ‘surprise’ is negative in the sense of lack of Attention but positive from a polarity point of view).

As an example of Sentic Extraction Process output, we can consider a short video comment such as “This phone looks pretty good! I’ll definitely consider this option!!!”. In this case the inferred moods would be ‘joy’ and ‘interest’ and the resulting polarity would be positive, as shown below:

<Concept: ‘phone’>
 <Concept: ‘look good’++>
 <Concept: ‘consider option’++>
 Sentic: [1.037, 0.748, 0.0, 0.0]
 Moods: joy and interest
 Polarity: 0.19

4 Describing media, people and their opinions

Web resources, and social media resources in particular, represent a peculiar kind of data that is characterized for a deeply interconnected nature. Web itself is in fact based on links that bound together different data and information, and community contributed multimedia resources characterize themselves for the collaborative way in which they are created and maintained.

An effective description of such resources needs therefore to capture and manage such interconnected nature, allowing to encode information not only about the resource itself but also about the linked resources into an interconnected knowledge

base. Encoding information relative to a market product to analyze its market trends represents a situation in which this approach is particularly suitable and useful. In this case, in fact, it is necessary not only to encode the information relative to product features but also the information about the producer, the consumers and their opinions. To achieve this purpose we exploit Semantic Web techniques.

4.1 The semantic Web

The recent advent of the Web has made available to its users a previously unbelievable amount of information and it has changed the way information is disseminated and retrieved, how business is conducted and how people communicate with each other.

As the Web grows at a constantly increasing rate, the volume of available information multiplies exponentially and the knowledge contained in the Web becomes more and more complete and vast. On the other hand, paradoxically, it become more and more difficult to capture the required information among a mass of irrelevant or mildly irrelevant information. Furthermore the demanded information is often spread over different sources and a highly time consuming activity is required to collect it, performing several queries and merging manually the partial information. This is due to the fact that most of web-content is perfectly suitable for human consumption, but it remains completely inaccessible to machines.

Today information retrieval is mainly performed using search engines like Google or Yahoo, that base their searches on keyword-based algorithms relying on the textual representation of web-pages. Such engines are very good in retrieving texts, splitting them into parts, checking the spelling, counting their words. But when it comes to interpreting sentences and capture their semantic, their capabilities result still very limited. The Semantic Web initiative by W3C (<http://w3.org>) tackles this problem through an appropriate representation of information in the web-page, able to univocally identify resources and encode the meaning of their description.

In particular, the Semantic Web uses uniform resource identifiers (URIs) to univocally identify entities available on the Web as documents or images but not as concepts or properties, and the resource description framework (RDF) (<http://w3.org/TR/REC-rdf-syntax>) data model to describe such resources in an univocally interpretable format, whose basic building block is an object-attribute-value triple i.e. a statement. Resources may be authors, books, publishers, places, people, hotels, rooms, search queries, etc., while properties describe relations between resources such as 'writtenBy', 'age', 'title'. Statements assert the properties of resources and their values can be either resources or literals (strings).

To provide machine-accessible and machine-processable representations, it is usual to encode RDF triples using a XML syntax. Each triple can also be seen as a directed graph with labeled nodes and arcs, where the arcs are directed from the resource (the subject of the statement) to the value (the object of the statement). Each statement describes the graph node or connects it to other nodes, linking together multiple data from different sources without pre-existing schema. It is according to this representation that indeed the Semantic Web in its whole can be envisioned as a Giant Global Graph of Linked Data.

RDF, however, does not make assumptions about any particular application domain, nor does it defines the semantics of any domain. For this purpose it is

necessary to introduce ontologies. Ontologies basically deal with knowledge representation and can be defined as formal explicit descriptions of concepts in a domain of discourse (named classes or concepts), properties of each concept describing various features and attributes of the concept (roles or properties), and restrictions on property (role restrictions). Ontologies make possible the sharing of common understanding about the structure of information among people or software agents. In addition, ontologies make possible reasoning i.e. it is possible, starting from the data and the additional information expressed in the form of ontology, to infer new relationships between data. Different languages have been developed for the design of ontologies, among the most popular there are RDFS (RDF Schema) (<http://w3.org/TR/rdf-schema>) and OWL (Ontology Web Language) (<http://w3.org/TR/owl-features>).

RDFS can be seen as a RDF vocabulary and a primitive ontology language. It offers certain modeling primitives with fixed meaning. Key concepts of RDF are class, subclass relations, property, sub-property relations, and domain and range restrictions. OWL is a language more specifically conceived for ontologies creation. It builds upon RDF and RDFS and a XML-based RDF syntax is used. Instances are defined using RDF descriptions and most RDFS modeling primitives are used. Moreover OWL introduces a number of features that are missing in RDFS such as local scope of property, disjointness of classes, Boolean combination of classes (like union, intersection and complement), cardinality restriction and special characteristics of properties (like transitive, unique or inverse).

4.2 The human emotion ontology

It is undisputed that emotions play a fundamental role in our lives. They occur in every relationship we care about—in the workplace, in our friendships, in dealings with family members, and in our most intimate relationships—and the most important moments of our life are linked indissolubly to the emotions we felt.

In recent years, the study of emotions has clearly revealed how emotions are fundamental to human experience, influencing cognition, perception, and everyday tasks such as learning, communication and decision-making. As a consequence, the study of emotions has attracted a growing attention from different research fields, ranging from advanced signal processing to psychology, from AI to linguistics. This has led to great advances in the understanding of emotions and in the multiplication of the models and classification techniques.

As a result, a standardization of the knowledge about emotions is getting more and more important but at the same time more difficult. Within the scientific community, in fact, the debate over human emotions is still open and there is still not a common agreement about which features are the most relevant in the definition of an emotion and which are the basic emotions and their names. Furthermore, dealing with the emotion extraction from media, the kind of descriptors and techniques used strictly rely on the considered medium. It is evident, for example, that the characteristics extracted from an audio sample rely basically on speech prosody, while features extracted from a video sample focus on facial expressions.

For this reason, there exists an unavoidable and intrinsic heterogeneity in the emotion descriptors that makes impossible to define a standard and unique set of such descriptors that could grant at the same time flexibility and interoperability. To

standardize the main sets of existing human emotion descriptors into a computational ontology, we developed the human emotion ontology (HEO) [10]. HEO is conceived as a high level ontology for human emotions that supplies the most significant concepts and properties which constitute the centerpiece for the description of every human emotion (Fig. 5). If necessary, these high level features can be further refined using lower level concepts and properties related to more specific descriptions or linked to other more specialized ontologies. HEO has been developed on the base of the activities carried on inside the Humaine Project (<http://emotion-research.net>), the COST Action 2102 (<http://cost2102.cs.stir.ac.uk>) and in particular the W3C Emotion Markup Language Incubator Group, which has recently developed EmotionML (<http://w3.org/TR/2009/WD-emotionml-20091029>), a markup language for the description and annotation of emotions.

The main purpose of HEO is thus to create a description framework that could grant at the same time enough flexibility, by allowing the use of a wide and extensible set of descriptors to represent all the main features of an emotion, and interoperability, by allowing to map concepts and properties belonging to different emotion representation models. In HEO an emotion can be described both in a discrete way, by using the ‘hasCategory’ property, and in a dimensional way, by using the ‘hasDimension’ property. HEO introduces two main disjoint classes for describing

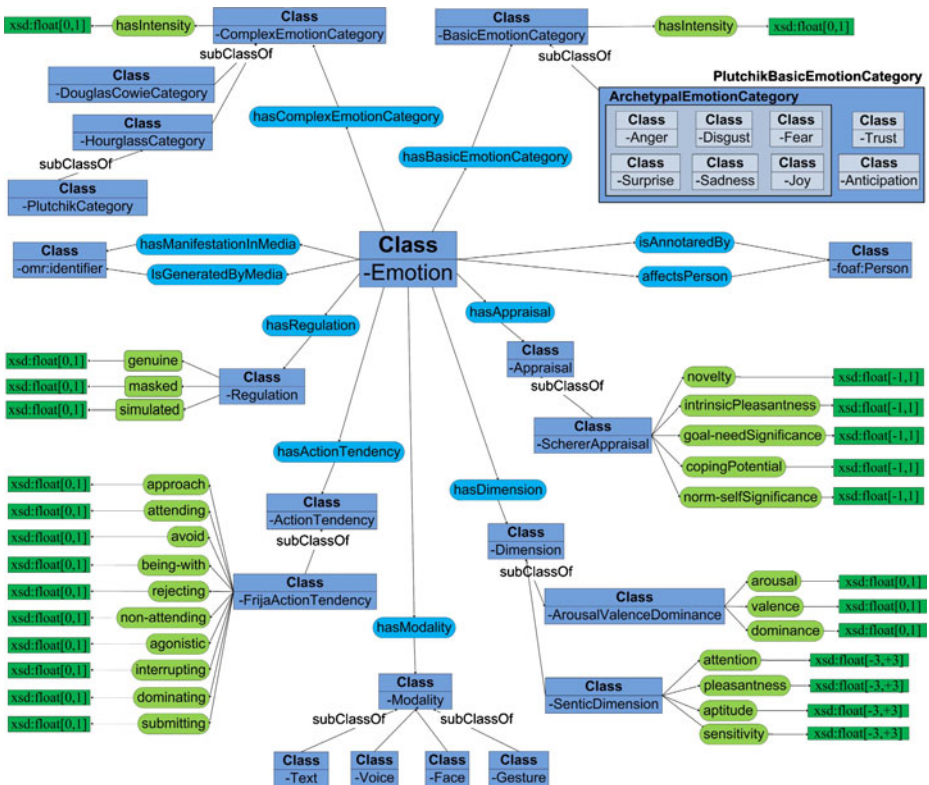


Fig. 5 The human emotion ontology schema

emotions by category: ‘BasicEmotionCategory’ and ‘ComplexEmotionCategory’. The reason for this distinction is twofold. Firstly, there are some emotion models that define complex emotions as a combination of two or more basic emotions. Secondly, the same descriptor is often used as a basic emotion or as complex emotion, according to different models, but the represented concept is different.

Different models can be used both for expressing the basic emotions e.g. Ekman’s 6 archetypal emotions (‘anger’, ‘disgust’, ‘fear’, ‘joy’, ‘sadness’, ‘surprise’) or Plutchik’s 8 basic emotions (‘acceptance’, ‘anger’, ‘anticipation’, ‘disgust’, ‘joy’, ‘fear’, ‘sadness’, ‘surprise’), and for the complex emotions using wider emotion sets e.g. Douglas and Cowie’s 48 descriptors or the 24 sentic levels of the Hourglass model. Instead, to describe emotions by dimension, HEO uses the ‘hasDimension’ property which includes the arousal-valence-dominance model and the sentic level values ($\text{float} \in [-3, +3]$).

HEO has been developed in OWL description logic (OWL DL) to take advantage of its expressiveness and its inference power in order to map the different models used in the emotion description. OWL DL, in fact, allows a taxonomical organization of emotion categories and properties restriction to link emotion description made by category and dimension. In HEO, for example, Ekman’s ‘joy’ archetypal emotion represents a superclass for Plutchik’s ‘ecstasy’, ‘joy’ and ‘serenity’ emotions. Using property restriction, the basic Plutchik’s ‘joy’ emotion can also be defined as an emotion that ‘has Pleasantness some float $\in [+1, +2]$ ’, ‘interest’ as an emotion that ‘has Attention some float $\in [0, +1]$ ’ and ‘love’ as an emotion that ‘has Pleasantness some float $\in [0, +3]$ and Aptitude some float $\in [0, +3]$ ’.

In this way, for example, querying a database that support OWL DL inference for basic emotions of type ‘joy’ will return not only the emotions expressly encoded as Ekman archetypal emotions of type ‘joy’, but also the emotions encoded as Plutchik basic emotion of type ‘joy’ and the emotions that have Pleasantness values in $[+1, +2]$.

4.3 A standardized framework for representing social media

Our framework for opinions and affective information description aims to be applicable to most of on-line resources (videos, images, text) coming from different sources e.g. on-line video sharing services, blogs and social networks. To such purpose it is necessary to standardize as much as possible the descriptors used in encoding the information about multimedia resources and people to which the opinions refer (considering that every website uses its own vocabulary) in order to make it univocally interpretable and suitable to feed other applications. For this reason we encode the information relative to multimedia resources and people using respectively the descriptors provided by OMR (Ontology for Media Resources) and FOAF (Friend of a Friend Ontology).

OMR represents an important effort to help circumventing the current proliferation of audio/video meta-data formats, currently carried on by the W3C Media Annotations Working Group. It offers a core vocabulary to describe media resources on the Web, introducing descriptors such as ‘title’, ‘creator’, ‘publisher’, ‘createDate’ and ‘rating’. It defines semantic-preserving mappings between elements from existing formats. This ontology is supposed to foster the interoperability among various kinds of meta-data formats currently used to describe media resources on the Web.

FOAF represents a recognized standard in describing people, providing information such as their names, birthdays, pictures, blogs, and especially other people they know, which makes it particularly suitable for representing data that appears on social networks and communities.

OMR and FOAF together supply most of the vocabulary we need for describing media and people and we add other descriptors only when necessary. For example OMR, at least in the current realization, does not supply vocabulary for describing comments, which we analyze to extract the affective information relative to media.

We extend this ontology introducing the ‘Comment’ class, and define for it the ‘author’, ‘text’ and ‘publicationDate’ properties. In HEO we introduce properties to link emotions to multimedia resources and people. In particular we have defined the ‘hasManifestationInMedia’ and ‘isGeneratedByMedia’, to describe emotions that respectively occur and are generated in media, and the property ‘affectPerson’ to connect emotions to people.

Moreover, to improve the hierarchical organization of emotions in HEO, we exploit WNA, the linguistic resource for the lexical representation of affective knowledge we use to build AffectiveSpace. WNA is built by assigning to a number of WordNet [8] synsets one or more affective labels (a-labels) and then by extending the core with the relations defined in WordNet. In particular, the affective concepts representing emotional states are identified by synsets marked with the a-label ‘EMOTION’, but there are also other a-labels for concepts representing moods, situations eliciting emotions or emotional responses. Thus, the combination of HEO with WNA, OMR and FOAF (Fig. 6) provides a complete framework to describe not only multimedia contents and the users that have created, uploaded or interacted

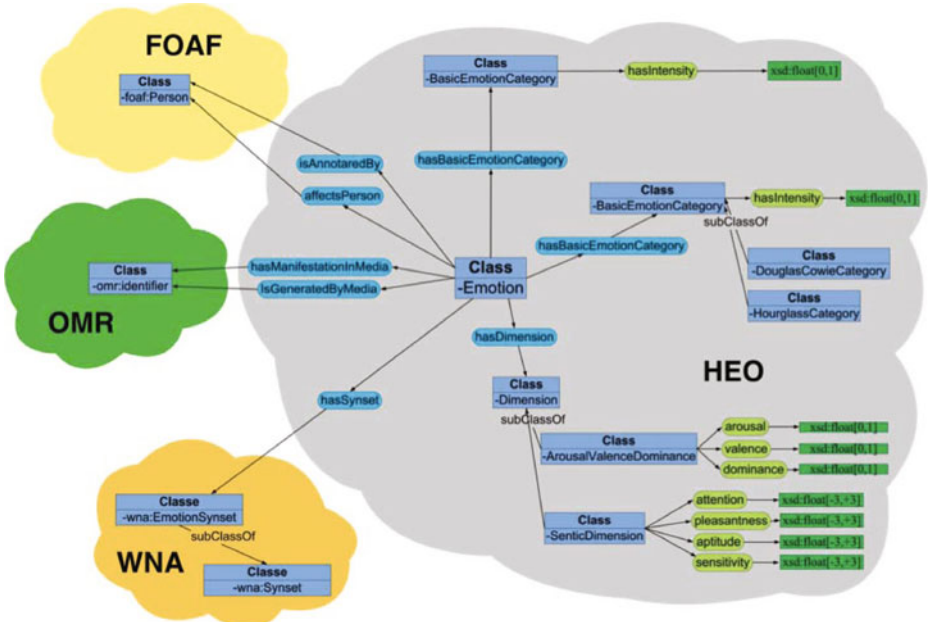


Fig. 6 Merging different ontologies to represent social media affective information

with them, but also the opinions and the affective content carried by the media and the way they are perceived by people.

5 Displaying product positioning

As remarked above, due to the way they are created and maintained, community-contributed multimedia resources are very different from standard web-data. One fundamental aspect is constituted by the collaborative way in which such data is created, uploaded and annotated. A deep interconnection emerges in the nature of these data and meta-data, allowing for example to associate videos of completely different genre, but uploaded by the same user, or different users, even living in opposite sides of the world, who have appreciated the same pictures. In the context of SMM, this interdependence can be exploited to find similar patterns in customer reviews of commercial products and hence to gather useful information for marketing, sales, public relations and customer service.

On-line reviews of electronics products, in particular, usually offer substantial and reliable information about the perceived quality of the products because of the size of electronics on-line market and the type of customers related to it. To visualize this information we exploit the multi-faceted categorization paradigm. Faceted classification allows the assignment of multiple categories to an object, enabling the classifications to be ordered in multiple ways, rather than in a single, pre-determined, taxonomic order. This makes possible to perform searches combining the textual approach with the navigational one. Faceted search, in fact, enables users to navigate a multi-dimensional information space by concurrently writing queries in a text box and progressively narrowing choices in each dimension.

For our application we use SIMILE Exhibit API (<http://simile.mit.edu/wiki/Exhibit>), a set of Javascript files that allows to easily create rich interactive web-pages including maps, timelines and galleries, with very detailed client-side filtering. Exhibit pages use the multi-faceted classification paradigm to display semantically structured data stored in a Semantic Web aware format e.g. RDF or JavaScript object notation (JSON). One of the most relevant aspects of Exhibit is that, once the page is loaded, the web-browser also loads the entire data set in a lightweight database and performs all the computations (sorting, filtering, etc.) locally on the client-side, providing high performances.

We selected one of the most prolific types of electronic products in terms of data reviews available on the Web i.e. mobile phones. In particular, we considered a set of 220 models ranked as the most popular according to Kelkoo (<http://kelkoo.com>), a shopping site featuring on-line shopping guides and user reviews, from which we parsed all the available information about the handsets such as model, brand, input type, screen resolution, camera type, standby time and weight. We encoded this information in RDF and stored it in a Sesame (<http://openrdf.org>) triple-store, a purpose-built database for the storage and retrieval of RDF meta-data. We then exploited YT DATA API to retrieve from YT database the most relevant video reviews for each mobile phone and their relative meta-data such as duration, rating, upload date and name, gender and country of the uploaders. For each video we also extracted the relative comments and processed them using Sentic Computing to extract the related opinions.

We then encoded the extracted opinions in RDF/XML, using the descriptors defined by HEO, WNA, OMR and FOAF, and inserted them into the triple-store. Sesame can be embedded in applications and used to conduct a wide range of inferences on the information stored, based on RDFS and OWL type relations between data. In addition, it can also be used in a standalone server mode, much like a traditional database with multiple applications connecting to it. In this way all the knowledge stored inside Sesame can be queried and the results can also be retrieved in a semantic aware format and used for other applications.

For our demo (available at <http://semedia.dibet.univpm.it/senticweb/smm>), we exported all the information contained in the triple-store into a JSON file to feed our Exhibit application, in order to make it available for being browsed as a unique knowledge base. We chose to use Exhibit in our demo due to the ease with which it allows to create rich and interactive web-pages. Mobile phones are displayed in a dynamic gallery, that can be ordered according to different parameters like model, price and rating, showing technical information jointly with their video reviews and the opinions extracted from the relative comments (Fig. 7). Using faceted menus it is possible to explore such information both using the search box, to perform keyword-based queries, and filtering the results using the faceted menus i.e. by adding or removing constraints on the facet properties.

In this way it becomes very easy and intuitive to search for mobile phones of interest: users can specify the technical features required using the faceted menus and compare different phones that match such requirements consulting the video reviews and the opinions extracted from the relative comments. In addition it is possible to explore in detail the comments of each video review through a specific Exhibit page in which comments are organized in a timeline and highlighted in different colors, according to the value of their polarity. Moreover faceted menus allow to filter the comments according to the reviewers' information e.g. age, gender and nationality.

The screenshot shows a web application interface for browsing mobile phones. On the left, there is a 'Search Box' and a list of facets for filtering: 'Mobile Phone', 'Product Name' (listing models like Apple iPhone 3G 8GB), 'Brand' (listing Apple, BlackBerry, HTC, LG), 'Supported Networks' (listing 2G, 3G, 3G+, 3GUMTS, 3GPP/HSDPA), and 'Multimedia Formats' (listing 3GP, 3GPP, AAC, AAC+). The main area displays '78 items' sorted by 'productName, brand, and label; then by ...' and 'grouped as sorted'. The selected item is an 'Apple iPhone 3G 8GB', showing its product name, brand, supported networks, display type, camera type, still image resolution, multimedia formats, messaging protocols, height, width, depth, weight, and rating. Below the specifications are three video review thumbnails. On the right, there are two panels: 'Affective Information' showing satisfaction levels (1-2, 2-3, 3-4, 4-5) and 'Basic Emotion Category' (anticipation, disgust, fear, joy, sadness). Below that is 'Sentiment Dimensions' with 'Attitude' (7-3, 5-2, 8-1, 6-1) and 'Attention' (7-3, 7-2, 8-1, 7-1) and 'Pleasantness' (8-3, 4-2).

Fig. 7 A snapshot of the faceted classification interface

Using such a tool a marketer can easily get an insight about the trend of a product, e.g. at the end of an advertising campaign, observing how the number of reviews and the relative satisfaction evolve in time and also monitoring this trend for different campaign targets.

6 Evaluation

In order to evaluate our system both on the level of opinion mining and sentiment analysis, we separately tested its polarity detection accuracy with a set of like/dislike-rated video reviews from YT (Section 6.1) and evaluated its affect recognition capabilities with a corpus of mood-tagged blogs from LiveJournal (LJ, <http://livejournal.com>) (Section 6.2), a virtual community of more than 23 million users who keep a blog, journal or diary.

6.1 Evaluating opinion mining accuracy

In order to evaluate the system in terms of polarity detection accuracy, we exploited YT DATA API to retrieve from YT database the ratings relative to the 220 video reviews previously selected for displaying in the faceted classification interface. On YT, in fact, users can express their opinions about videos either by adding comments or by simply rating them using a like/dislike button. YT DATA API makes this kind of information available by providing, for each video, number of raters and average rating i.e. sum of likes and dislikes divided by number of raters. This information is expressed as a float $\in [1, 5]$ and indicates if a video is generally considered as bad (float $\in [1, 3]$) or good (float $\in [3, 5]$). We compared this information with the polarity values previously extracted by employing Sentic Computing on the comments relative to each of the 220 videos. We considered true positives those videos with both an average rating $\in [3, 5]$ and a polarity $\in [0, 1]$ (for positively rated videos) and those with both an average rating $\in [1, 3]$ and a polarity $\in [-1, 0]$ (for negatively rated videos). The evaluation showed that, by using Sentic Computing to perform polarity detection, negatively and positively rated videos (37% and 63% of the total respectively) can be identified with precision of 97% and recall of 86% (91% F-measure).

6.2 Evaluating sentiment analysis accuracy

Since no mood-labeled dataset about commercial products is currently available, we used a corpus blogs from LJ to test our system's affect recognition capabilities. One of the interesting features of this website is that LJ bloggers are allowed to label their posts with a mood tag, by choosing from more than 130 predefined moods or by creating custom mood themes. Since the indication of the affective status is optional, the mood-tagged posts are likely to reflect the true mood of the authors and, hence, form a good test-set for Sentic Computing. However, since LJ mood themes do not perfectly match sentic levels, we had to consider a reduced set of 10 moods i.e. 'ecstatic', 'happy', 'pensive', 'surprised', 'enraged', 'sad', 'angry', 'annoyed', 'scared' and 'bored'. Moreover we could not consider non-affective web-posts since untagged blog entries do not necessarily lack emotions.

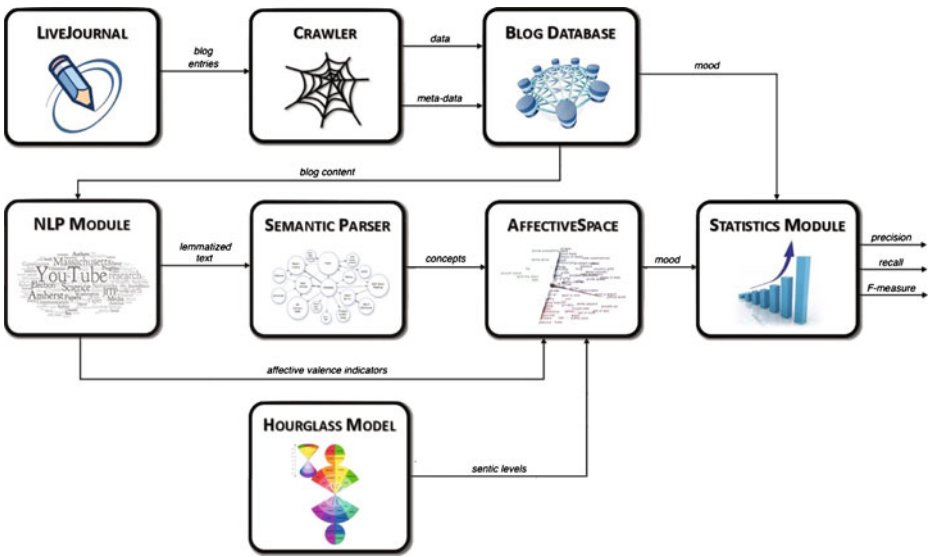


Fig. 8 Sentic Computing evaluation process

All LJ accounts have Atom, RSS and other data feeds which show recent public entries, friend relationships and interests. Unfortunately there is no possibility to get mood-tagged blog-posts via data feeds so we had to design our own crawler. After retrieving and storing relevant data and meta-data for a total of 10,000 posts, we ran the Sentic Extraction Process on each of these and compared the output with the relative mood-tags, in order to calculate statistical classifications such as precision and recall (Fig. 8). On average, each post contained around 140 words and, from it, about four affective valence indicators and 60 sentic vectors were extracted.

According to this information, we assigned mood-labels to each post and compared these with the corresponding LJ mood-tags, obtaining very good accuracy for each of the ten selected moods (as shown in Table 1). Among these, ‘happy’ and ‘sad’ posts were identified with particularly high precision (89 and 81%, respectively) and decorous recall rates (76 and 68%). The F-measure values obtained, hence, were significantly good (82 and 74%, respectively), especially if compared to the

Table 1 Precision, recall and F-measure rates of the ten selected moods

	Precision (%)	Recall (%)	F-measure (%)
Ecstatic	73	61	66
Happy	89	76	82
Pensive	69	52	59
Surprised	81	65	72
Enraged	68	51	58
Sad	81	68	74
Angry	81	53	64
Annoyed	77	58	66
Scared	82	63	71
Bored	70	55	62

corresponding F-measure rates of the baseline methods (53 and 51% for keyword spotting, 63 and 58% for lexical affinity, 69 and 62% for statistical methods).

7 Conclusion and future efforts

As the Web plays a more and more significant role in people's social lives, it contains more and more information concerning their feelings and opinions but most of this informative content slips from all the searches performed every day.

In this paper we showed how AI and Semantic Web techniques can be used in conjunction to manage affective information associated to community-created data and meta-data. In particular, we used Sentic Computing, a new paradigm for the affective analysis of natural language text, to semantically analyze opinions and we exploited different web ontologies to encode the results in a semantic aware format. We eventually made these data accessible as a whole knowledge base allowing advanced features selection, both in querying and in displaying the relevant information.

The output of this research work is an intelligent SMM tool addressed to both marketers and end-users for the management of social media information relative to products, brands or organizations. We plan to improve our system by making most of its functionalities available as web-services in order to enable users to dynamically analyze, encode and display data in a way that they can create their own video review mash-ups in real-time and hence concurrently analyze various social media information in a unique interface. Linking data coming from different sources and enriching it with relative affective information, in fact, represent key factors for the development of next-generation intelligent web-applications—especially in the field of SMM—since no application can be really intelligent unless it is aware of user's opinions and feelings.

References

1. Alm C, Roth D, Sproat R (2005) Emotions from text: machine learning for text-based emotion prediction. In: HLT/EMNLP, Vancouver
2. Cambria E, Hussain A, Havasi C, Eckl C (2009) AffectiveSpace: blending common sense and affective knowledge to perform emotive reasoning. In: WOMSA at CAEPIA, Seville
3. Cambria E, Hussain A, Havasi C, Eckl C (2009) Common sense computing: from the society of mind to digital intuition and beyond. In: LNCS, vol 5707. Springer, Berlin, pp 252–259
4. Cambria E, Hussain A, Havasi C, Eckl C (2010) Sentic computing: exploitation of common sense for the development of emotion-sensitive systems. In: LNCS, vol 5967. Springer, Berlin, pp 148–156 (2010)
5. Cambria E, Hussain A, Havasi C, Eckl C (2010) SenticSpace: visualizing opinions and sentiments in a multi-dimensional vector space. In: LNAI, vol 6279. Springer, Berlin, pp 385–393
6. Eckart C, Young G (1936) The approximation of one matrix by another of lower rank. *Psychometrika* 1(3):211–218
7. Elliott C (1992) The affective reasoner: a process model of emotions in a multi-agent system. Technical Report no. 32, The Institute for the Learning Sciences
8. Fellbaum C (1998) WordNet: an electronic Lexical database. MIT Press, Cambridge
9. Goertzel B, Silverman K, Hartley C, Bugaj S, Ross M (2000) The baby webmind project. In: AISB, Birmingham
10. Grassi M (2009) Developing HEO human emotion ontology. In: LNCS, vol 5707. Springer, Berlin, pp 244–251


11. Havasi C, Speer R, Alonso J (2007) ConceptNet 3: a flexible, multilingual semantic network for common sense knowledge. In: RANLP, Borovets
12. Havasi C, Speer R, Pustejovsky J, Lieberman H (2009) Digital intuition: applying common sense using dimensionality reduction. *IEEE Intell Syst* 24(4):24–35
13. Hu M, Liu B (2004) Mining opinion features in customer reviews. In: AAAI, San Jose
14. Minsky M (2006) *The emotion machine*. Simon and Schuster, New York
15. Nicolov N, Salvetti F, Liberman M, Martin J (eds) (2006) AAAI symposium on computational approaches to analysing weblogs. AAAI, Menlo Park
16. Pang B, Lee L (2004) A sentimental education: sentiment analysis using subjectivity summarization based on minimum cuts. In: ACL, Barcelona
17. Pang B, Lee L (2008) Opinion mining and sentiment analysis. In: *Foundations and trends in information retrieval*, vol 2, pp 1–135
18. Pang B, Lee L, Vaithyanathan S (2002) Thumbs up? Sentiment classification using machine learning techniques. In: EMNLP, Philadelphia
19. Plutchik R (2001) The nature of emotions. *Am Sci* 89(4):344–350
20. Salzman M, Matathia I, O'Reilly A (2003) *Buzz: harness the power of influence and create demand*. Wiley, New York
21. Somasundaran S, Wiebe J, Ruppenhofer J (2008) Discourse level opinion interpretation. In: COLING, Manchester
22. Sood S, Vasserman L (2009) Esse: exploring mood on the Web. In: ICWSM, Seattle
23. Sood S, Owsley S, Hammond K, Birnbaum L (2007) Reasoning through search: a novel approach to sentiment classification. In: WWW, Banff
24. Strapparava C, Valitutti A (2004) WordNet-Affect: an affective extension of WordNet. In: LREC, Lisbon
25. Turney P (2002) Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. In: ACL, Philadelphia
26. Wall M, Rechtsteiner A, Rocha L (2003) Singular value decomposition and principal component analysis. In: Berrar D et al (eds) *A practical approach to microarray data analysis*. Kluwer, Norwell, pp 91–109
27. Wiebe J, Wilson T, Claire C (2005) Annotating expressions of opinions and emotions in language. *Lang Resour Eval* 2(3) 39:165–210
28. Wilson T, Wiebe J, Hoffmann P (2005) Recognizing contextual polarity in phrase-level sentiment analysis. In: HLT-EMNLP, Vancouver



Erik Cambria is an EPSRC researcher on a project born from the collaboration between the University of Stirling (UK), the MIT Media Laboratory (USA) and Sitekit Labs (UK). His interests include AI, Semantic Web, NLP, cognitive and affective modeling, opinion mining and sentiment analysis. Erik is also a Web Science lecturer at the University of the Highlands and Islands (UK) and a research affiliate at HP Labs India (India), Chinese Academy of Sciences (China) and Microsoft Research Asia (China).



Marco Grassi obtained his MSc in Electronic Engineering and his PhD degree in Electronic and Computing Engineering at the Polytechnic University of Marche (Italy), in 2007 and 2011 respectively. Since then he is a researcher fellow at Department of Biomedical, Electronic and Telecommunication Engineering. His research interests include Semantic Web, knowledge modelling, opinion mining and sentiment analysis. Marco is also the working group coordinator of the COST Action 2102 metadata encoding task.



Amir Hussain is a reader in the Department of Computing Science and Mathematics at the University of Stirling (UK). From 2000 until 2008, he held lecturer and senior lecturer positions within the same university. Amir is editor-in-chief of the (Springer) Journal of Cognitive Computation, and serves on the editorial board of four other international journals. He is IEEE chapter chair of the UK and RI Industry Applications Society, a senior member of the IEEE, and a fellow of the UK Higher Education Academy.



Catherine Havasi is a co-founder of the Open Mind Common Sense project at the MIT Media Laboratory (USA) where she works as a post-doctoral associate. She received her PhD in computer science at Brandeis University (USA) with James Pustejovsky, and her interests include common sense reasoning, dimensionality reduction, machine learning, language acquisition, cognitive modeling and intelligent user interfaces.