

# SENTIC COMPUTING FOR SOCIAL NETWORK ANALYSIS

# 5

F. Bisio<sup>a</sup>, L. Oneto<sup>a</sup>, E. Cambria<sup>b</sup>

*University of Genoa, Genoa, Italy<sup>a</sup> Nanyang Technological University, Singapore, Singapore<sup>b</sup>*

## 1 INTRODUCTION

Sentiment analysis or opinion mining can be defined as a particular application of data mining, which aims to aggregate and extract emotions and feelings from different types of documents [1]. The amount of data available on the web is growing exponentially; these data, however, are mainly in an unstructured format and hence are not machine processable and machine interpretable. Therefore graph-mining and natural language processing (NLP) techniques may contribute to the distillation of knowledge and opinions from the huge amount of information present in the web.

Sentiment analysis can enhance the capabilities of customer relationship management and recommendation systems by allowing one, for example, to find out which features customers are particularly interested in or to exclude from ads items that have received unfavorable feedback. Likewise, it can be used in social communication to enhance antispam systems.

Business intelligence can also benefit from sentiment analysis. Since predicting the attitude of the public toward a brand or a product has become of crucial importance for companies, an increasing amount of money has been invested in marketing strategies involving opinion and sentiment mining.

That scenario led to sentic computing [2], which tackles those crucial issues by exploiting affective commonsense reasoning (ie, the intrinsically human capacity to interpret the cognitive and affective information associated with natural language). In particular, sentic computing leverages a commonsense knowledge base built through crowdsourcing [3,4]. Common sense is useful in many different computer science applications, including data visualization [5], text recognition [6], and human-computer interaction [7]. In this context, common sense is used to bridge the semantic gap between word-level natural language data and the concept-level opinions conveyed by these [8].

To perform affective commonsense reasoning [9], a knowledge database is required for storage and extraction of the affective and emotional information associated with word and multiword expressions. Graph-mining [10] and dimensionality reduction [11] techniques have been employed on a knowledge base obtained by the blending of ConceptNet [12], a directed graph representation of commonsense knowledge, with WordNet-Affect [13], a linguistic resource for the lexical representation of feelings. Unlike WordNet-Affect, SenticNet [14] and its extended versions [15,16] exploit an ensemble of

common and commonsense knowledge to go beyond word-level opinion mining, and hence to associate semantics and sentics with a set of natural language concepts.

With the rise of the Social Web, there are now millions of humans offering their knowledge online, which means that the information is stored, searchable, and easily shared. This trend has created and maintained an ecosystem of participation, where the value is created by the aggregation of many individual user contributions. Such contributions, however, are meant for human consumption and hence are hardly accessible and processable by computers. Making sense of the huge amount of social data available on the web requires the adoption of novel approaches to natural language understanding that can give a structure to such data in a way that they can be more easily aggregated and analyzed [17,18].

Social networks are indeed a popular means of sharing data and ideas, and have witnessed an ever-increasing diffusion. The amount of data generated in 30 seconds on the Internet is about 600 GB of traffic<sup>1</sup>: this fact confirms that online information, with a particular focus on social networks, has become a source of big data. For example, considering the specific case of the Twitter community, every minute more than 320 new accounts are created and more than 98,000 tweets are posted. This makes the analysis of Twitter microblogging a topmost and significant domain for business intelligence and marketing strategies. A multiplicity of users populate this social network, sharing different types of information. The average age of users on Twitter ranges from 14 to 60 years, equally distributed among individuals of both genders. Hence among the multitude of tweets, one may want to retrieve information associated with specific relevant topics and identify the related polarity and affective characterization. Therefore one of the most interesting applications of sentiment analysis in recent years has involved social networks, so as to be able to analyze opinions on Twitter, Facebook, or other digital communities in real time [19–21]. In this context, sentic computing can be exploited for NLP tasks requiring the inference of semantic and/or affective information associated with both text (eg, for the analysis of social network interaction dynamics) [22] and online multimodal data [23].

This chapter presents an overview of the most recent and advanced technologies of sentiment analysis, with particular focus on the applications related to social networks. The main result consists of a review of the most interesting methods employed to compare, classify, and visualize messages on social media platforms. The chapter is organized as follows: [Section 2](#) provides a description of state-of-the-art sentiment analysis techniques, with particular regard to the affective characterization of the data; [Section 3](#) describes an affective characterization of emotions and sentiments; [Section 4](#) describes several applications that employ sentiment analysis methods; [Section 5](#) discusses some future trends and directions; and [Section 6](#) presents our conclusions and final remarks.

---

## 2 RELATED WORK

The Social Web has provided people with new content-sharing services that allow them to create and share personal content, ideas, and opinions, in a time-and cost-efficient way, with virtually millions of other people connected to the World Wide Web. Since this amount of information is mainly unstructured, research has so far focused on online information retrieval, aggregation, and processing.

---

<sup>1</sup><http://pennystocks.la/internet-in-real-time/>

The potential applications of concept-level sentiment analysis are countless and span interdisciplinary areas, such as stock market prediction, political forecasting, social network analysis, social stream mining, and man-machine interactions.

Today, most of the existing approaches still rely on word co-occurrence frequencies (ie, the syntactic representation of text). Therefore computational models aim to bridge the cognitive gap by emulating the way the human brain processes natural language. For instance, by leveraging semantic features that are not explicitly expressed in text, one may accomplish complex NLP tasks such as word-sense disambiguation, textual entailment, and semantic role labeling. Computational models are useful for both scientific purposes (such as exploring the nature of linguistic communication), and practical purposes (such as enabling effective human-machine communication [24]).

Most existing approaches to sentiment analysis rely on the extraction of a vector representing the most salient and important text features, which is later used for classification purposes [25,26]. Commonly used features include term frequency and the presence and the position of a token in a text. An extension of this feature is the presence of  $n$ -grams, typically bigrams and trigrams. Other methods rely on the distance between terms or on the part-of-speech information: for example, certain adjectives may be good indicators of sentiment orientation. A drawback of these approaches is the strict dependency on the domain of application considered and the related topics.

Sentiment analysis systems aim to classify entire documents by associating content with some overall positive or negative polarity or rating scores (eg, 1–5 stars) [27–31]. These approaches are typically supervised and rely on manually labeled samples [32,33]. We can distinguish between knowledge-based systems [34], based on approaches such as keyword spotting and lexical affinity, and statistics-based systems [35]. At first, the identification of emotions and polarity was performed mainly by means of knowledge-based methods; recently, sentiment analysis researchers have been increasingly using statistics-based approaches, with a special focus on supervised statistical methods.

Keyword spotting is the most straightforward, and possibly also the most popular, approach thanks to its accessibility and economy. Text is classified into affect categories on the basis of the presence of fairly unambiguous “affect words” such as “happy,” “sad,” “afraid,” and “bored.” Elliott’s Affective Reasoner [36], for example, looks for 198 affect keywords plus affect intensity modifiers and a handful of cue phrases. Other popular sources of affect words are Ortony’s Affective Lexicon [37], which groups terms into affective categories, and Wiebe’s linguistic annotation scheme [38]. The crucial issue of these approaches lies in the ineffectiveness at handling negations and in the structure based on the presence of obvious affect words.

Rather than simply detecting affect words, lexical affinity assigns each word a probabilistic “affinity” for a particular emotion. These probabilities are usually learned from linguistic corpora [39–41]. Even if this method often outperforms pure keyword spotting, it still works at word level and can be easily tricked by negations and different senses of the same word. Besides, lexical affinity probabilities are often biased by the linguistic corpora adopted, which makes it difficult to develop a reusable, domain-independent model.

Statistical-based approaches, such as Bayesian inference and support vector machines (SVMs), have been used in several projects [27,42–44]. By feeding a machine learning algorithm [45] a large training corpus of affectively annotated texts, one can learn not only the affective valence of affect keywords (as in the keyword spotting approach), but also the affective valence of other arbitrary keywords (like lexical affinity), punctuation, and word co-occurrence frequencies. Anyway, it is worth noting that statistical classifiers work well only when a sufficiently large text dataset is given as input. This is

because, with the exception of affect keywords, other lexical or co-occurrence elements possess little predictive value individually.

For example, Pang et al. [27] used a movie review dataset to compare the performance of different machine learning algorithms: in particular, they obtained 82.90% accuracy when they used a large number of textual features. Socher et al. [46] proposed a recursive neural tensor network and improved the accuracy (85%). Yu and Hatzivassiloglou [47] identified polarity at sentence level using semantic orientation of words. Melville et al. [48] developed a framework exploiting word-class association information for domain-dependent sentiment analysis. Bisio et al. [49] tackled a particular aspect of the sentiment classification problem: the ability of the framework itself to operate effectively in heterogeneous commercial domains. The approach adopts a distance-based predictive model to combine computational efficiency and modularity.

Some approaches are based on the following assumption: many short  $n$ -grams are usually neutral, while longer phrases are well distributed among positive and negative subjective sentence classes. Therefore matrix representations for long phrases and matrix multiplication to model composition can also be used to evaluate sentiment. In such models, sentence composition is modeled with use of deep neural networks such as recursive auto-associated memories. Recursive neural networks predict the sentiment class at each node in the parse tree and try to capture the negation and its scope in the entire sentence.

Several unsupervised learning approaches have also been proposed, and rely on the creation of the lexicon via the unsupervised labeling of words or phrases with their sentiment polarity or subjectivity [50]. To this aim, early work was mainly based on linguistic heuristics. For example, Hatzivassiloglou and McKeown [51] built a system based on “opposition constraints” to help label decisions in the case of polarity classification.

Other work exploited the seed words, defined as terms for which the polarity is known, and propagated them to terms that co-occur with them in general text, or in specific WordNet-defined relations. Popescu and Etzioni [52] proposed an algorithm that, starting from a global word label computed over a large collection of generic topic text, gradually tried to redefine such a label for a more and more specific corpus, until the one that is specific to the particular context in which the word occurs. Snyder and Barzilay [53] exploited the idea of utilizing discourse information to aid in the inference of relationships between product attributes.

Concept-based approaches [54–57] focus on a semantic analysis of text through the use of web ontologies or semantic networks, which allow the handling of conceptual and affective information associated with natural language opinions. By relying on large semantic knowledge bases, such approaches step away from blind use of keywords and word co-occurrence counts, and rather rely on the implicit meaning/features associated with natural language concepts. Unlike purely syntactical techniques, concept-based approaches are also able to detect sentiments that are expressed in a subtle manner (eg, through the analysis of concepts that do not explicitly convey any emotion but are implicitly linked to other concepts that do so).

More recent studies [58–61] tackled the problem of sentiment analysis in social networks: for example, by enhancing sentiment analysis of tweets by exploiting microblogging text or Twitter-specific features such as emoticons, hashtags, URLs, @ symbols, capitalization, and elongations. Tang et al. [62] developed a convolutional neural network–based approach to obtain word embeddings for the words mostly used in tweets. These word vectors were then fed to a convolutional neural network for sentiment analysis. A deep convolutional neural network for sentiment detection in short text was also

proposed by dos Santos and Gatti [63]. The approach based on sentiment-specific word embeddings [64] considers word embeddings based on a sentiment corpus: this means the inclusion of more affective clues than regular word vectors and the production of a better result.

### 3 AFFECTIVE CHARACTERIZATION

Emotions and affective information can be represented through the hourglass of emotions model [65], an affective categorization model that allows emotions to be deconstructed into independent but concomitant affective dimensions, whose different levels of activation compose the total emotional state of the mind. Such a reinterpretation is inspired by Minsky’s theory of the mind, according to which brain activity consists of different independent resources and that emotional states result from turning some set of these resources on and turning another set of them off.

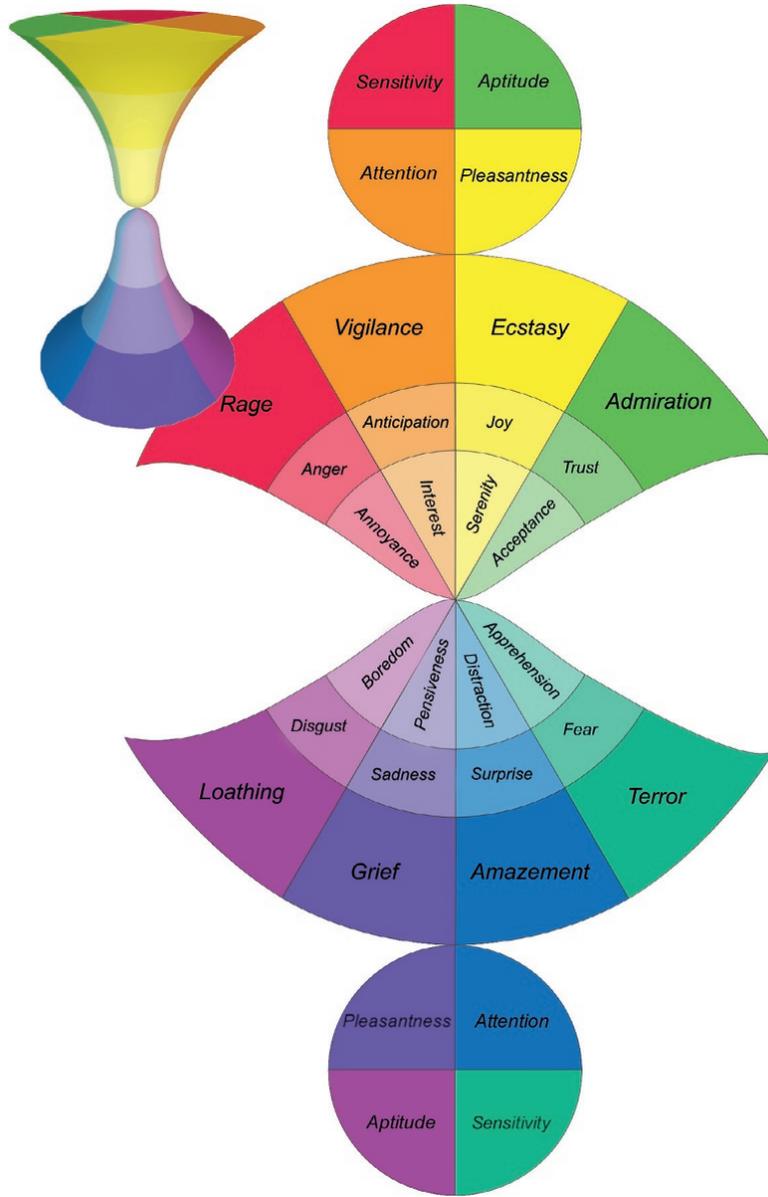
The model can potentially synthesize the full range of emotional experiences in terms of four affective dimensions—pleasantness, attention, sensitivity, and aptitude—which define the intensity of the expressed/perceived emotion in the interval  $\in [-1, +1]$ . Each affective dimension is characterized by six levels of activation, termed “sentic levels,” which are also labeled as a set of 24 basic emotions (six for each affective dimension) (see Table 5.1). In the model the vertical dimension represents the intensity of the different affective dimensions (ie, their level of activation). The model follows the pattern used in color theory and research to obtain judgments about combinations (ie, the emotions that result when two or more fundamental emotions are combined).

The transition between different emotional states is modeled, within the same affective dimension, by means of the function  $G(x) = -\frac{1}{\sigma\sqrt{2\pi}}e^{-x^2/2\sigma^2}$ , for its symmetric inverted bell curve shape that quickly rises up toward 1. In particular, the function models how the valence or intensity of an affective dimension varies according to different values of arousal or activation, ranging from 0 (emotional void) to 1 (heightened emotionality). Mapping this space of possible emotions leads to a hourglass shape (see Fig. 5.1). The justification for assuming that the Gaussian function (rather than a step or simple linear function) is appropriate for modeling the variation of emotion intensity is based on research into the neural and behavioral correlates of emotion, which are assumed to indicate emotional intensity in some sense.

It is worth noting that, in the model, the state of “emotional void” is adimensional, which contributes to the creation of the hourglass shape. Total absence of emotion can be associated with the total absence

**Table 5.1 The Sentic Levels of the Hourglass Model**

Interval	Pleasantness	Attention	Sensitivity	Aptitude
$G(1), G(2/3)$	Ecstasy	Vigilance	Rage	Admiration
$[G(2/3), G(1/3))$	Joy	Anticipation	Anger	Trust
$[G(1/3), G(0))$	Serenity	Interest	Annoyance	Acceptance
$(G(0), G(-1/3])$	Pensiveness	Distraction	Apprehension	Boredom
$(G(-1/3), G(-2/3])$	Sadness	Surprise	Fear	Disgust
$(G(-2/3), G(-1))$	Grief	Amazement	Terror	Loathing



**FIG. 5.1**  
The three-dimensional model of the hourglass of emotions.

**Table 5.2 The Emotions Generated by Pairwise Combination of the Sentic Levels of the Hourglass Model**

	Attention >0	Attention <0	Aptitude >0	Aptitude <0
Pleasantness >0	Optimism	Frivolity	Love	Gloat
Pleasantness <0	Frustration	Disapproval	Envy	Remorse
Sensitivity >0	Aggressiveness	Rejection	Rivalry	Contempt
Sensitivity <0	Anxiety	Awe	Submission	Coercion

of reasoning (or, at least, consciousness), which is not an envisaged mental state as, in the human mind, there is never nothing going on.

Complex emotions can be synthesized by use of different sentic levels, as shown in [Table 5.2](#), which represent the intensity thresholds of the expressed or perceived emotion.

## 4 APPLICATIONS

This section describes how sentic computing tools and techniques can be employed for the development of applications in the context of big social data analysis.

### 4.1 TROLL FILTERING

The democracy of the web has allowed a high degree of freedom of expression, which also gave birth to negative side effects: in the Social Web context the possibility of anonymity often results in the posting of inflammatory, extraneous, or off-topic messages in an online community, with the primary intent of provoking other users into making a desired emotional response or of otherwise disrupting normal on-topic discussion.

Such practice is usually referred to as “trolling,” and the generator of such messages is called a “troll.” Trolling appears to have spread a lot recently and it is alarming most of the biggest social networking sites since, in extreme cases such as abuse, it has led some teenagers to commit suicide. These attacks usually address not only individuals but also entire communities.

At present, users cannot do much other than to manually delete abusive messages. Current antitrolling methods mainly consist in identifying additional accounts that use the same IP address and blocking fake accounts on the basis of the name and anomalous site activity (eg, users who send lots of messages to nonfriends or whose friend requests are rejected at a high rate). Even though there is the possibility of reporting trouble, social networking websites usually cannot react instantly to these alarms.

A prior analysis of the trustworthiness of statements published on the web was presented by Rowe and Butters [66]. Their approach adopts a contextual trust value determined for the person who asserted a statement as the trustworthiness of the statement itself. Their study, however, focuses not on the problem of trolling but rather on defining a contextual accountability for the detection of web, email, and opinion spam.

The main aim of the troll filter [67] is to identify malicious content in natural language text with a certain confidence level and hence to automatically block trolls. To train the system, the concepts most commonly used by trolls are first identified by use of the concept frequency–inverted opinion frequency technique [2], and then this set is expanded through spectral association. In particular, after a set of 1000 offensive phrases extracted from Wordnik<sup>2</sup> had been analyzed, it was found that, statistically, a post is likely to be edited by a troll when its average sentic vector has a high absolute value of sensitivity (one of the dimensions of the hourglass of emotions [65]) and a very low polarity. Hence the *trollness*  $t_i$  associated with a concept  $c_i \in [0, 1]$  is

$$t_i(c_i) = \frac{s_i(c_i) + |\text{sensitivity}(c_i)| - p_i(c_i)}{3}, \quad (5.1)$$

where  $s_i \in [0, 1]$  is the semantic similarity of  $c_i$  with any of the concept frequency–inverted opinion frequency seed concepts,  $p_i \in [-1, 1]$  is the polarity associated with the concept  $c_i$ , and 3 is the normalization factor. Hence the total *trollness* of a post containing  $N$  concepts is defined as

$$t = \sum_{i=1}^N \frac{3 s_i(c_i) + 4 |Se(c_i)| - Pl(c_i) - |At(c_i)| - Ap(c_i)}{9N}, \quad (5.2)$$

where  $Se$  is the sensitivity,  $Pl$  the pleasantness,  $At$  the attention, and  $Ap$  the aptitude. This information is stored, together with the post type and content plus the sender ID and the receiver ID, in an interaction database that keeps track of all the messages and comments interchanged between users within the same social network.

Posts with a high level of *trollness* (with a threshold set, by means of a trial-and-error approach, to 60%) are labeled as troll posts and, whenever a specific user addresses more than two troll posts to the same person or community, his/her sender ID is labeled as a troll for that particular receiver ID. All the past troll posts sent to that particular receiver ID by that specific sender ID are then automatically deleted from the website. Moreover, any new post with a high level of *trollness* edited by a user labeled as a troll for that specific receiver is automatically blocked.

The evaluation of this approach [67] was performed by consideration of a set of 500 tweets manually annotated as troll and nontroll posts, most of which were fetched from Wordnik.

In Table 5.3 a comparison of three different vector space representations for commonsense concepts is presented. AnalogySpace [68] forms the analogical closure of a semantic network through dimensionality reduction. IsaCore [69] is a strongly connected core of hyponym-hypernym common knowledge. AffectNet is a semantic network in which commonsense concepts are linked to a hierarchy of affective domain labels (eg, “joy,” “amazement,” “fear,” “admiration”). To enable affective analogical reasoning on natural language concepts, AffectiveSpace [70] is obtained as the vector space representation of such a semantic network.

The results show that, by use of the troll filtering process, inflammatory and outrageous messages can be identified with good precision (82.5%) and a good recall rate (75.1%). In particular, the  $F$  score (78.6%) is significantly high compared with the corresponding  $F$  scores obtained by use of IsaCore and AnalogySpace in place of the AffectiveSpace process (Table 5.3).

<sup>2</sup><https://www.wordnik.com/>

<b>Metric</b>	<b>IsaCore (%)</b>	<b>AnalogySpace (%)</b>	<b>AffectiveSpace (%)</b>
Precision	57.1	69.1	82.5
Recall	40.0	56.6	75.1
<i>F</i> score	47.0	62.2	78.6

## 4.2 SOCIAL MEDIA MARKETING

The online review of commercial services and products allows users to share their opinions about services they have received or products they have bought, and it constitutes an 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 online advertising and positioning. Given the growing interest in social media marketing, several commercial tools have been developed recently to provide companies with the possibility to analyze the blogosphere on a large scale so they can extract information about the trend of the opinions relative to their products. Nevertheless, most of the existing tools and 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.

To this end, a novel social media marketing tool has been proposed [71] to provide marketers with a tool for the management of social media information at the semantic level, able to capture both opinion polarity and affective information.

In particular, YouTube was selected as a social media source since, with its more than 2 billion views per day, 24 hours of video uploaded every minute, and 15 minutes a day spent by the average user, it represents more than 40% of the online video market.<sup>3</sup> Specifically, the focus was on video reviews of cell phones because of the quantity and the quality of the comments usually associated with them.

The social media analysis was performed through three main steps: firstly, comments are analyzed with an opinion mining engine; secondly, the extracted information is encoded on the basis of different web ontologies; finally, the resulting knowledge base is made available for browsing through a multifaceted classification website.

To make the tool applicable to most online resources (videos, images, text), the information relative to multimedia resources and people is encoded with use of the descriptors provided by the Ontology for Media Resources (OMR)<sup>4</sup> and Friend of a Friend (FOAF) ontology<sup>5</sup> respectively.

OMR offers a core vocabulary to describe media resources on the web, introducing descriptors such as “title,” “creator,” “publisher,” “createDate,” and “rating,” and it defines semantic-preserving mappings between elements from existing formats to foster interoperability among them.

<sup>3</sup><http://www.viralblog.com/research-cases/youtube-statistics>

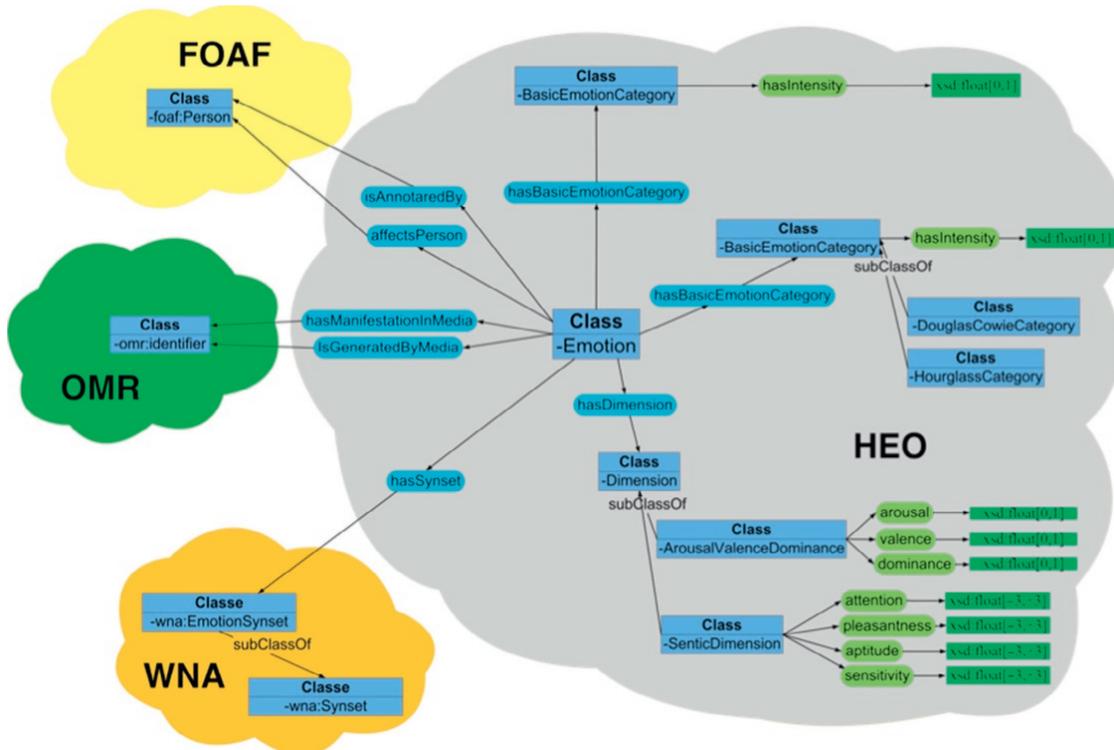
<sup>4</sup><https://www.w3.org/TR/mediaont-10>

<sup>5</sup><http://www.foaf-project.org/>

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 appear in social networks and communities.

OMR and FOAF together supply most of the vocabulary needed to describe media and people; other descriptors are added only when necessary. For example, OMR does not currently supply vocabulary for describing comments, which are hereby analyzed to extract the affective information relative to media. Hence the ontology is extended by the introduction of the “Comment” class and by defining for it the “author,” “text,” and “publicationDate” properties.

In the human emotion ontology (HEO) [72], properties to link emotions to multimedia resources and people are introduced (see Fig. 5.2). Additionally, WordNet-Affect [13] is exploited as an ontology to improve the hierarchical organization of emotions in HEO. Thus the combination of HEO and WordNet-Affect, OMR, and FOAF provides a complete framework to describe not only multimedia content and the users that have created, uploaded, or interacted with it but also the opinions and the affective content carried by the media and the way they are perceived by web users. To make this information visualizable, the multifaceted categorization paradigm is exploited (see Fig. 5.3). Faceted



**FIG. 5.2** Merging different ontologies. *FOAF*, Friend of a Friend; *HEO*, Human Emotion Ontology; *OMR*, Ontology for Media Resources; *WNA*, WordNet-Affect.

The screenshot displays a web-based interface for a social media marketing tool. At the top, there is a search box and a navigation menu with 'PHONES' and 'VIDEO REVIEWS'. The main content area shows 78 items, sorted by product name, brand, and labels. On the left, there are several filter panels: 'Mobile Phone' (listing Apple iPhone 3G 8GB, Apple iPhone 3G S 16 GB, Apple iPhone 3G S 32 GB, BlackBerry 8520, and BlackBerry 9700 Bold), 'Brand' (Apple, BlackBerry, Generic, HTC, LG), 'Supported Networks' (2G, 3G, 3G+/HSDPA, 3G/UMTS, 3GPP/HSPA), and 'MultimediaFormats' (3GP, 3GPP, AAC, AAC+). The central panel displays the details for an 'Apple iPhone 3G 8GB', including its specifications and three video review thumbnails. On the right, there are three panels for 'Affective Information': 'Satisfaction' (ranging from 1-2 to 4-5), 'Basic Emotion Category' (anticipation, disgust, fear, joy, sadness), and 'Sentic Dimensions' (Aptitude, Attention, Pleasantness).

FIG. 5.3

A screenshot of the social media marketing tool.

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, and taxonomic order. This makes it possible to perform searches combining the textual approach with the navigational one.

To evaluate the proposed system on the level of both opinion mining and sentiment analysis, its polarity detection accuracy was separately tested with a set of like/dislike-rated video reviews from YouTube and its affect recognition capabilities were evaluated with a corpus of mood-tagged blogs from LiveJournal.

To evaluate the system in terms of polarity detection accuracy, the YouTube data application programming interface was used to retrieve from the YouTube database the ratings relative to 220 video reviews of cell phones. On YouTube, users can express their opinions about videos either by adding comments or by simply rating them using a like/dislike button. The YouTube data application programming interface makes this kind of information available by providing, for each video, the number of raters and the average rating (ie, the sum of likes and dislikes divided by the number of raters).

This takes a value in the range [1, 5] and indicates if a video is generally considered bad (values in the range [1, 3]) or good (values in the range [3, 5]). This information was compared with the polarity values previously extracted by use of sentic computing on the comments relative to each of the 220 videos. True positives were identified as videos with both an average rating  $r \in [3, 5]$  and a polarity  $p \in [0, 1]$  (for positively rated videos), or videos with both an average rating  $r \in [1, 3]$  and a polarity  $p \in [-1, 0]$  (for negatively rated videos). The evaluation showed that, by use of the system to perform polarity detection, negatively and positively rated videos (37.7% and 62.3% of the total respectively) can be identified with a precision of 97.1% and recall of 86.3% ( $F$  score of 91.3%).

**Table 5.4 Evaluation Results of the Sentics Extraction Process**

Mood	Precision (%)	Recall (%)	F score (%)
Ecstatic	73.1	61.3	66.6
Happy	89.2	76.5	82.3
Pensive	69.6	52.9	60.1
Surprised	81.2	65.8	72.6
Enraged	68.9	51.6	59.0
Sad	81.8	68.4	74.5
Angry	81.4	53.3	64.4
Annoyed	77.3	58.7	66.7
Scared	82.6	63.5	71.8
Bored	70.3	55.1	61.7

Since no mood-labeled dataset regarding commercial products is currently available, the LiveJournal database was used to test the system's affect recognition capabilities. For this test, a reduced set of 10 moods was considered: ecstatic, happy, pensive, surprised, enraged, sad, angry, annoyed, scared, and bored. After relevant data and metadata for 5000 posts had been retrieved and stored, the sentics extraction process was conducted on each of these and the outputs were compared with the relative LiveJournal mood tags to compute recall and precision rates as evaluation metrics.

On average, each post contained around 140 words and, from this about 4 affective valence indicators and 60 sentic vectors were extracted. According to this information, mood labels were assigned to each post and compared with the corresponding LiveJournal mood tags, resulting in very good accuracy for each of the 10 selected moods (Table 5.4). Among these, "happy" and "sad" posts were identified with particularly high precision (89.2% and 81.8% respectively) and good recall rates (76.5% and 68.4%). The *F* scores obtained were also good (82.3% and 74.5% respectively), especially when compared with the corresponding *F* scores for a standard keyword spotting system based on a set of 500 affect words (65.7% and 58.6%).

### 4.3 A MODEL FOR SENTIMENT CLASSIFICATION IN TWITTER

This section describes a Twitter sentiment analysis system able to classify a tweet as positive or negative on the basis of its overall tweet-level polarity. The need for such a tool comes from the following consideration: supervised learning classifiers often misclassify tweets containing conjunctions such as "but" and conditionals such as "if" because of their special linguistic characteristics. Moreover, tweets often contain misspelled words, slang, URLs, elongations, repeated punctuation, emoticons, abbreviations, and hashtags. To counter these challenges, a system that enhances supervised learning for polarity classification by leveraging linguistic rules and sentic computing resources is proposed [61]. The general scheme of the system is presented in Fig. 5.4.

First all tweets are preprocessed to change all the @<username> references to @USER and all the URLs to http://URL.com. Then tweets are tokenized and a part-of-speech tag is assigned to each token. Then all the emoticons are identified. Since people often repeat certain punctuation to

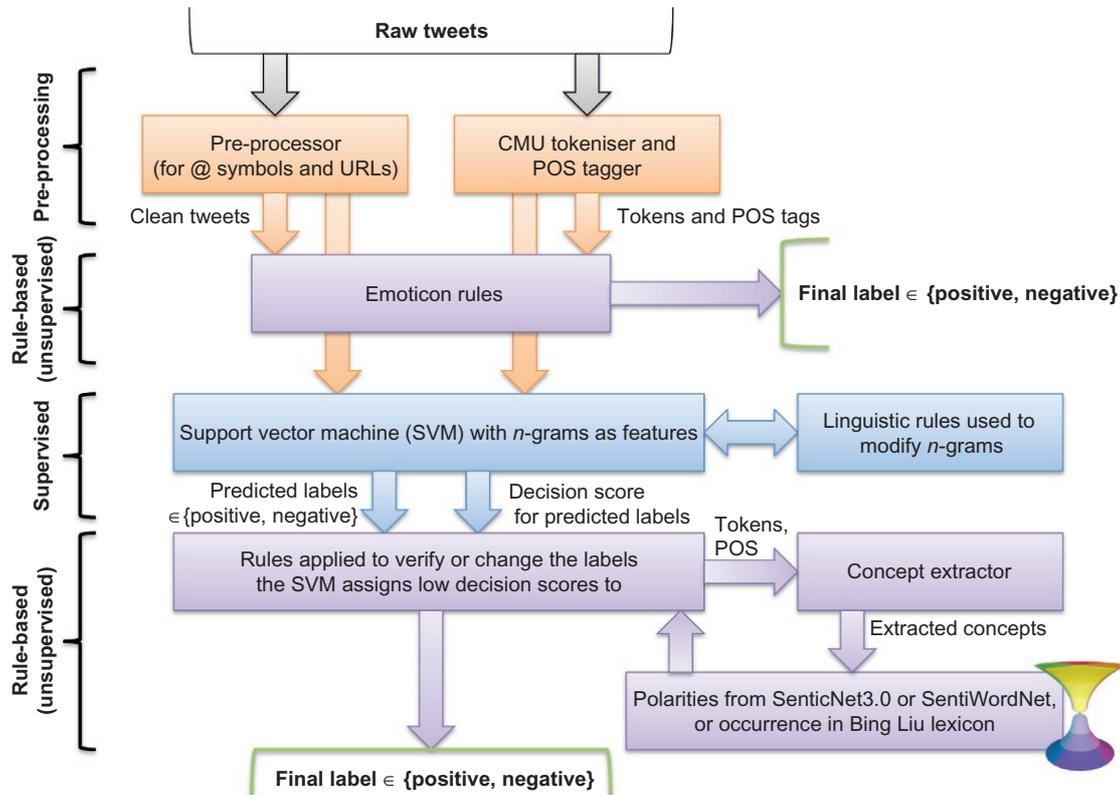


FIG. 5.4

Flowchart of the Twitter sentiment analysis system. *CMU*, Carnegie Mellon University; *POS*, part of speech.

emphasize emoticons, all repeated characters from every emoticon string are removed to obtain the bag of emoticons present in the tweet.

The number of positive and negative emoticons in the tweet are counted and the following rules are applied:

- If a tweet contains one or more positive emoticons and no negative emoticons, it is labeled as *positive*.
- If a tweet contains one or more negative emoticons and no positive emoticons, it is labeled as *negative*.
- If neither of the two rules above can be applied, the tweet is labeled as *unknown*.

If these emoticon-based rules label a tweet as *positive* or *negative*, this is considered the final label outputted by the system. Otherwise all tweets labeled as *unknown* are passed into a supervised learning classifier. In this case, each tweet is represented as a feature vector of case-sensitive  $n$ -grams (unigrams, bigrams, and trigrams). These  $n$ -grams are frequencies of sequences of one, two, or three contiguous tokens in a tweet. Besides, negation is taken into consideration. Moreover, all tweets containing the

conjunction “but” and the conditionals “if,” “unless,” “until,” and “in case” are considered, and specific linguistic rules are formulated to enable removal of irrelevant or oppositely oriented  $n$ -grams from the tweet’s feature vector.

Finally, a SVM classifier is trained so as to obtain the tweet’s label. For tweets with an absolute decision score or confidence below 0.5, the class labels assigned by the SVM are discarded and an unsupervised classifier is employed. This unsupervised classification process works as follows:

1. Single-word and multiword concepts are extracted from the tweets so as to fetch their polarities from SenticNet [14].
2. If a single-word concept is not found in SenticNet, it is queried in SentiWordNet [73], and if it is not found in SentiWordNet, it is searched for in the list of positive and negative words from the Bing Liu lexicon [74].
3. On the basis of the number of positive and negative concepts, and the most polar value occurring in the tweet, the following rules are applied:
  - If the number of positive concepts is greater than the number of negative concepts and the most polar value occurring in the tweet is greater than or equal to 0.6, the tweet is labeled as positive.
  - If the number of negative concepts is greater than the number of positive concepts and the most polar value occurring in the tweet is less than or equal to  $-0.6$ , the tweet is labeled as negative.
  - If neither of the two rules stated above can be applied, the tweet is labeled as unknown by the rule-based classifier, and the SVM’s low confidence prediction is taken as the final output of the system.

An SVM classifier [75] was trained on 1.6 million positive and negative tweets provided by Go et al. [76]. After a standard 10 cross-validation procedure to set the model parameters [77,78], the method was evaluated on two publicly available datasets: SemEval 2013 [79] and SemEval 2014 [80]. Tables 5.5 and 5.6 show the results obtained on these two datasets. In these tables, each row shows the

Method	Positive			Negative			Average		
	P	R	F	P	R	F	P	R	F
$n$ -grams	90.48	82.67	86.40	61.98	76.45	68.46	76.23	79.56	77.43
$n$ -grams and emoticon rules	90.62	83.36	86.84	62.99	76.65	69.15	76.80	80.00	78.00
Modified $n$ -grams	89.95	84.05	86.90	63.33	74.59	68.50	76.64	79.32	77.70
Modified $n$ -grams and emoticon rules	90.10	84.73	87.33	64.41	74.79	69.22	77.26	79.76	78.27
Modified $n$ -grams, emoticon rules, and word-level unsupervised rules	91.40	86.79	89.04	68.55	77.89	72.92	79.97	82.34	80.98
Modified $n$ -grams, emoticon rules, and concept-level unsupervised rules	92.42	86.56	89.40	68.96	80.79	74.41	80.69	83.68	81.90

*F, F score; P, precision; R, recall.*

**Table 5.6 Results Obtained on 3584 Positive/Negative Tweets From the SemEval 2014 Dataset**

Method	Positive			Negative			Average		
	P	R	F	P	R	F	P	R	F
<i>n</i> -grams	89.92	81.90	85.72	61.20	75.66	67.67	75.56	78.78	76.69
<i>n</i> -grams and emoticon rules	89.74	83.05	86.27	62.50	74.85	68.11	76.12	78.95	77.19
Modified <i>n</i> -grams	89.39	82.90	86.02	62.00	73.93	67.44	75.69	78.41	76.73
Modified <i>n</i> -grams and emoticon rules	89.25	83.97	86.53	63.29	73.22	67.89	76.27	78.60	77.21
Modified <i>n</i> -grams, emoticon rules, and word-level unsupervised rules	90.22	86.24	88.19	67.37	75.25	71.09	78.80	80.75	79.64
Modified <i>n</i> -grams, emoticon rules, and concept-level unsupervised rules	90.41	86.20	88.25	67.45	75.76	71.37	78.93	80.98	79.81

*F*, *F* score; *P*, precision; *R*, recall.

precision, recall, and *F* score for the positive and negative classes, followed by the average positive and negative precision, recall, and *F* score. In the tables we indicate as modified *n*-grams the method that takes into account the linguistic rules employed on conjunctions and conditionals.

When comparing the standard *n*-grams model with the *n*-grams and emoticon rules model, one may notice that emoticon rules increase the average *F* by 0.57 and 0.50 in the 2013 and 2014 datasets respectively. A comparison between the modified *n*-grams model and the modified *n*-grams and emoticon rules model also shows that emoticon rules increase the average *F* by 0.57 and 0.48 in the two datasets respectively. This confirms that emoticon rules may significantly improve the classification. Moreover, modification of *n*-grams with use of linguistic rules for conjunctions and conditionals increases the average *F* by 0.27 and 0.04 in the two datasets respectively. Even if the increase is not very significant for the 2014 dataset, modified *n*-grams are still useful since a typical Twitter corpus usually contains a very small percentage of tweets with such conjunctions and conditionals.

A further comparison involves the analysis of the unsupervised rules: the results obtained by use of a bag-of-concepts model are compared with the ones obtained by use of a bag-of-words (or single-word concepts only) model. The average *F* for the bag-of-concepts model is 0.92 greater than for the bag-of-words model for the 2013 dataset, and 0.17 greater than the bag-of-words model for the 2014 dataset, thus confirming the effectiveness of the concept-level sentiment features. Overall, the final sentiment analysis system achieves an average *F* score that is greater by 4.47 and 3.12 than for the standard *n*-grams model.

## 5 FUTURE TRENDS AND DIRECTIONS

The research work carried out so far has formed a solid base for the development of a variety of emotion-sensitive systems and applications in the fields of opinion mining and sentiment analysis. This chapter has presented an overview of approaches to the analysis of opinions and sentiments in

social networks, which go beyond keyword-based methods by using commonsense reasoning tools and affective ontologies. The techniques developed, however, are still far from perfect as the common and commonsense knowledge bases need to be further expanded and the reasoning tools built on top of them adjusted accordingly.

To overcome these limitations, current research work is focusing on expanding the affective knowledge by employing external resources (eg, Cyc, Freebase, and Yago). New graph-mining and multidimensionality reduction techniques are also being explored to perform reasoning on the commonsense knowledge bases. Moreover, new classification techniques are being tested, together with the ensemble application of dimensionality reduction and neural networks techniques for emulation of fast affective learning. Therefore, looking ahead, we believe the combined use of different knowledge bases and affective commonsense reasoning techniques will, eventually, pave the way for the development of more bio-inspired approaches to the design of intelligent systems capable of handling knowledge, making analogies, learning from experience, and perceiving and expressing affect.

---

## 6 CONCLUSION

With the advent of the Social Web, the way people express their views and opinions has dramatically changed. Social networks now represent huge sources of information with many practical applications. However, finding opinion sources and monitoring them can be a formidable task because there are a large number of diverse sources and each source may also have a huge volume of opinionated text. Thus automated opinion discovery and summarization systems are needed. Because of their tremendous value for practical applications, there has been an explosive growth of sentiment analysis techniques in both research in academia and applications in industry. However, most of the existing approaches still rely on syntactical structure of text, which is far from the way the human mind processes natural language.

In this chapter we have presented an overview of the most recent state-of-the-art sentiment-based techniques applied to social networks applications and tools. To assess the capability of sentiment analysis to tackle real-world NLP tasks, we considered several applications in different domains. Specifically, sentiment analysis tools and techniques were used for the design of Social Web applications (ie, a troll filtering system, a social media marketing tool, and a Twitter sentiment analysis system).

All these applications demonstrate how sentiment analysis can be a useful resource for the analysis of affective information in social platforms, relying not on domain-dependent keywords but rather on commonsense knowledge bases that allow one to extrapolate the cognitive and affective information associated with natural language text.

---

## REFERENCES

- [1] E. Cambria, *Affective computing and sentiment analysis*, *IEEE Intell. Syst.* 31 (2) (2016) 102–107.
- [2] E. Cambria, A. Hussain, *Sentic Computing: A Common-Sense-Based Framework for Concept-Level Sentiment Analysis*, Springer, Cham, Switzerland, 2015.

- [3] E. Cambria, Y. Xia, A. Hussain, Affective common sense knowledge acquisition for sentiment analysis, in: LREC, Istanbul, 2012, pp. 3580–3585.
- [4] E. Cambria, D. Rajagopal, K. Kwok, J. Sepulveda, GECKA: game engine for commonsense knowledge acquisition, in: FLAIRS, 2015, pp. 282–287.
- [5] E. Cambria, A. Hussain, C. Havasi, C. Eckl, SenticSpace: visualizing opinions and sentiments in a multi-dimensional vector space, in: R. Setchi, I. Jordanov, R. Howlett, L. Jain (Eds.), Knowledge-Based and Intelligent Information and Engineering Systems, Lecture Notes in Artificial Intelligence, vol. 6279, Springer, Berlin, 2010, pp. 385–393.
- [6] Q. Wang, E. Cambria, C. Liu, A. Hussain, Common sense knowledge for handwritten Chinese recognition, *Cogn. Comput.* 5 (2) (2013) 234–242.
- [7] S. Poria, E. Cambria, N. Howard, G.-B. Huang, A. Hussain, Fusing audio, visual and textual clues for sentiment analysis from multimodal content, *Neurocomputing* 174 (2016) 50–59.
- [8] E. Cambria, S. Poria, F. Bisio, R. Bajpai, I. Chaturvedi, The CLSA model: a novel framework for concept-level sentiment analysis, in: LNCS, vol. 9042, Springer, Berlin, 2015, pp. 3–22.
- [9] F. Bisio, P. Gastaldo, R. Zunino, E. Cambria, A learning scheme based on similarity functions for affective common-sense reasoning, in: IJCNN, 2015, pp. 2476–2481.
- [10] E. Cambria, D. Olsher, K. Kwok, Sentic activation: a two-level affective common sense reasoning framework, in: AAI, Toronto, 2012, pp. 186–192.
- [11] E. Cambria, J. Fu, F. Bisio, S. Poria, AffectiveSpace 2: enabling affective intuition for concept-level sentiment analysis, in: AAI, Austin, 2015, pp. 508–514.
- [12] R. Speer, C. Havasi, ConceptNet 5: a large semantic network for relational knowledge, in: Theory and Applications of Natural Language Processing, 2012.
- [13] C. Strapparava, A. Valitutti, WordNet-Affect: an Affective Extension of WordNet, in: Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC 2004), Lisbon, May 2004, pp. 1083–1086.
- [14] E. Cambria, D. Olsher, D. Rajagopal, SenticNet 3: a common and common-sense knowledge base for cognition-driven sentiment analysis, in: AAI, Quebec City, 2014, pp. 1515–1521.
- [15] S. Poria, A. Gelbukh, E. Cambria, D. Das, S. Bandyopadhyay, Enriching SenticNet polarity scores through semi-supervised fuzzy clustering, in: IEEE ICDM, Brussels, 2012, pp. 709–716.
- [16] S. Poria, A. Gelbukh, E. Cambria, P. Yang, A. Hussain, T. Durrani, Merging SenticNet and WordNet-Affect emotion lists for sentiment analysis, in: 2012 IEEE 11th International Conference on Signal Processing (ICSP), vol. 2, IEEE, 2012, pp. 1251–1255.
- [17] E. Cambria, H. Wang, B. White, Guest editorial: big social data analysis, *Knowl. Based Syst.* 69 (2014) 1–2.
- [18] S. Poria, E. Cambria, A. Gelbukh, F. Bisio, A. Hussain, Sentiment data flow analysis by means of dynamic linguistic patterns, *IEEE Comput. Intell. Mag.* 10 (4) (2015) 26–36.
- [19] S. Chinthala, R. Mande, S. Manne, S. Vemuri, Sentiment analysis on twitter streaming data, in: Annual Convention of the Computer Society of India, 2015.
- [20] S. Rosenthal, P. Nakov, S. Kiritchenko, S.M. Mohammad, A. Ritter, V. Stoyanov, Semeval-2015 task 10: sentiment analysis in twitter, in: International Workshop on Semantic Evaluation, 2015.
- [21] A. Ortigosa, J.M. Martín, R.M. Carro, Sentiment analysis in Facebook and its application to e-learning, *Comput. Hum. Behav.* 31 (2014) 527–541.
- [22] P. Chandra, E. Cambria, A. Hussain, Clustering social networks using interaction semantics and sentics, in: J. Wang, G. Yen, M. Polycarpou (Eds.), Advances in Neural Networks, Lecture Notes in Computer Science, vol. 7367, Springer, Heidelberg, 2012, pp. 379–385.
- [23] S. Poria, E. Cambria, A. Hussain, G.-B. Huang, Towards an intelligent framework for multimodal affective data analysis, *Neural Netw.* 63 (2015) 104–116.
- [24] J.L. Reyes-Ortiz, L. Oneto, A. Samá, X. Parra, D. Anguita, Transition-aware human activity recognition using smartphones, *Neurocomputing* 171 (2016) 754–767.

- [25] E. Fersini, E. Messina, F. Pozzi, Expressive signals in social media languages to improve polarity detection, *Inform. Process. Manage.* 52 (1) (2016) 20–35.
- [26] D. Tang, B. Qin, F. Wei, L. Dong, T. Liu, M. Zhou, A joint segmentation and classification framework for sentence level sentiment classification, *IEEE/ACM Trans. Audio Speech Lang. Process.* 23 (11) (2015) 1750–1761.
- [27] B. Pang, L. Lee, S. Vaithyanathan, Thumbs up? Sentiment classification using machine learning techniques, in: *Conference on Empirical Methods on Natural Language Processing*, 2002.
- [28] B. Pang, L. Lee, Seeing stars: exploiting class relationships for sentiment categorization with respect to rating scales, in: *Annual Meeting on Association for Computational Linguistics*, 2005.
- [29] K.L. Devi, P. Subathra, P. Kumar, Tweet sentiment classification using an ensemble of machine learning supervised classifiers employing statistical feature selection methods, in: *International Conference on Fuzzy and Neuro Computing*, 2015.
- [30] G. Wang, J. Sun, J. Ma, K. Xu, J. Gu, Sentiment classification: the contribution of ensemble learning, *Decis. Support Syst.* 57 (2014) 77–93.
- [31] E. Fersini, E. Messina, F. Pozzi, Sentiment analysis: Bayesian ensemble learning, *Decis. Support Syst.* 68 (2014) 26–38.
- [32] A. Cernian, V. Sgarciu, B. Martin, Sentiment analysis from product reviews using SentiWordNet as lexical resource, in: *International Conference on Electronics, Computers and Artificial Intelligence*, 2015.
- [33] D. Tang, B. Qin, T. Liu, Y. Yang, User modeling with neural network for review rating prediction, in: *International Joint Conferences on Artificial Intelligence*, 2015.
- [34] E. Cambria, B. Schuller, B. Liu, H. Wang, C. Havasi, Knowledge-based approaches to concept-level sentiment analysis, *IEEE Intell. Syst.* 28 (2) (2013) 12–14.
- [35] E. Cambria, B. Schuller, B. Liu, H. Wang, C. Havasi, Statistical approaches to concept-level sentiment analysis, *IEEE Intell. Syst.* 28 (3) (2013) 6–9.
- [36] C.D. Elliott, *The Affective Reasoner: A Process Model of Emotions in a Multi-Agent System*, Northwestern University, Evanston, IL, 1992.
- [37] A. Ortony, G. Clore, A. Collins, *The Cognitive Structure of Emotions*, Cambridge University Press, Cambridge, UK, 1988.
- [38] J. Wiebe, T. Wilson, C. Cardie, Annotating expressions of opinions and emotions in language, *Lang. Resour. Eval.* 39 (2) (2005) 165–210.
- [39] R. Stevenson, J. Mikels, T. James, Characterization of the affective norms for English words by discrete emotional categories, *Behav. Res. Methods* 39 (2007) 1020–1024.
- [40] S. Somasundaran, J. Wiebe, J. Ruppenhofer, Discourse level opinion interpretation, in: *International Conference on Computational Linguistics*, 2008.
- [41] D. Rao, D. Ravichandran, Semi-supervised polarity lexicon induction, in: *Conference of the European Chapter of the Association for Computational Linguistics*, 2009.
- [42] A. Weichselbraun, S. Gindl, A. Scharl, Extracting and grounding context-aware sentiment lexicons, *IEEE Intell. Syst.* 28 (2) (2013) 39–46.
- [43] L. García-Moya, H. Anaya-Sanchez, R. Berlanga-Llavori, A language model approach for retrieving product features and opinions from customer reviews, *IEEE Intell. Syst.* 28 (3) (2013) 19–27.
- [44] G. Di Fabbriozio, A. Aker, R. Gaizauskas, Summarizing on-line product and service reviews using aspect rating distributions and language modeling, *IEEE Intell. Syst.* 28 (3) (2013) 28–37.
- [45] M. Vahdat, L. Oneto, D. Anguita, M. Funk, M. Rauterberg, Can machine learning explain human learning? *Neurocomputing* 192 (2016) 14–28.
- [46] R. Socher, A. Perelygin, J.Y. Wu, J. Chuang, C.D. Manning, A.Y. Ng, C. Potts, Recursive deep models for semantic compositionality over a sentiment treebank, in: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2013.

- [47] H. Yu, V. Hatzivassiloglou, Towards answering opinion questions: separating facts from opinions and identifying the polarity of opinion sentences, in: *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*, 2003.
- [48] P. Melville, W. Gryc, R.D. Lawrence, Sentiment analysis of blogs by combining lexical knowledge with text classification, in: *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2009.
- [49] F. Bisio, P. Gastaldo, C. Peretti, R. Zunino, E. Cambria, Data intensive review mining for sentiment classification across heterogeneous domains, in: *Advances in Social Networks Analysis and Mining (ASONAM)*, IEEE, 2013, pp. 1061–1067.
- [50] B. Pang, L. Lee, Opinion mining and sentiment analysis, *Found. Trends Inform. Retrieval*. 2 (1–2) (2008) 1–135.
- [51] V. Hatzivassiloglou, K.R. McKeown, Predicting the semantic orientation of adjectives, in: *Annual Meeting of the Association for Computational Linguistics and Eighth Conference of the European Chapter of the Association for Computational Linguistics*, 1997.
- [52] A.-M. Popescu, O. Etzioni, Extracting product features and opinions from reviews, in: *Natural Language Processing and Text Mining*, 2007.
- [53] B. Snyder, R. Barzilay, Multiple aspect ranking using the good grief algorithm, in: *Conference of the North American Chapter of the Association for Computational Linguistics—Human Language Technologies*, 2007.
- [54] A. Tsai, R. Tsai, J. Hsu, Building a concept-level sentiment dictionary based on commonsense knowledge, *IEEE Intell. Syst.* 28 (2) (2013) 22–30.
- [55] S. Poria, A. Gelbukh, A. Hussain, N. Howard, D. Das, S. Bandyopadhyay, Enhanced SenticNet with affective labels for concept-based opinion mining, *IEEE Intell. Syst.* 28 (2) (2013) 31–38.
- [56] C. Hung, H.-K. Lin, Using objective words in SentiWordNet to improve sentiment classification for word of mouth, *IEEE Intell. Syst.* 28 (2) (2013) 47–54.
- [57] C. Bosco, V. Patti, A. Bolioli, Developing corpora for sentiment analysis and opinion mining: a survey and the Senti-TUT case study, *IEEE Intell. Syst.* 28 (2) (2013) 55–63.
- [58] E. Kouloumpis, T. Wilson, J. Moore, Twitter sentiment analysis: the good the bad and the omg!, in: *International AAAI Conference on Weblogs and Social Media*, 2011.
- [59] D. Davidov, O. Tsur, A. Rappoport, Enhanced sentiment learning using twitter hashtags and smileys, in: *Proceedings of the 23rd International Conference on Computational Linguistics: Posters*, 2010.
- [60] S.M. Mohammad, S. Kiritchenko, X. Zhu, NRC-Canada: building the state-of-the-art in sentiment analysis of tweets, in: *Joint Conference on Lexical and Computational Semantics*, 2013.
- [61] P. Chikersal, S. Poria, E. Cambria, SeNTU: sentiment analysis of tweets by combining a rule-based classifier with supervised learning, in: *International Workshop on Semantic Evaluation (SemEval 2015)*, 2015.
- [62] D. Tang, F. Wei, B. Qin, T. Liu, M. Zhou, Coooolll: a deep learning system for twitter sentiment classification, in: *International Workshop on Semantic Evaluation*, 2014.
- [63] C.N. dos Santos, M. Gatti, Deep convolutional neural networks for sentiment analysis of short texts, in: *International Conference on Computational Linguistics*, 2014.
- [64] D. Tang, F. Wei, N. Yang, M. Zhou, T. Liu, B. Qin, Learning sentiment-specific word embedding for twitter sentiment classification, in: *Annual Meeting of the Association for Computational Linguistics*, 2014.
- [65] E. Cambria, A. Livingstone, A. Hussain, The hourglass of emotions, in: A. Esposito, A. Vinciarelli, R. Hoffmann, V. Muller (Eds.), *Cognitive Behavioral Systems*, Lecture Notes in Computer Science, vol. 7403, Springer, Berlin, Heidelberg, 2012, pp. 144–157.
- [66] M. Rowe, J. Butters, Assessing trust: contextual accountability, in: *European Semantic Web Conference*, 2009.

- [67] E. Cambria, P. Chandra, A. Sharma, A. Hussain, Do not feel the trolls, in: *International Semantic Web Conference*, 2010.
- [68] R. Speer, C. Havasi, H. Lieberman, AnalogySpace: reducing the dimensionality of common sense knowledge, in: *AAAI*, 2008.
- [69] E. Cambria, Y. Song, H. Wang, N. Howard, Semantic multi-dimensional scaling for open-domain sentiment analysis, *IEEE Intell. Syst.* 29 (2) (2014) 44–51.
- [70] E. Cambria, P. Gastaldo, F. Bisio, R. Zunino, An ELM-based model for affective analogical reasoning, *Neurocomputing* 149 (2015) 443–455.
- [71] E. Cambria, M. Grassi, A. Hussain, C. Havasi, Sentic computing for social media marketing, *Multimedia Tools Appl.* 59 (2) (2012) 557–577.
- [72] M. Grassi, Developing HEO human emotions ontology, in: J. Fierrez, J. Ortega-Garcia, A. Esposito, A. Drygajlo, M. Faundez-Zanuy (Eds.), *Biometric ID Management and Multimodal Communication*, Lecture Notes in Computer Science, vol. 5707, Springer, Berlin, Heidelberg, 2009 pp. 244–251.
- [73] A. Esuli, F. Sebastiani, SentiWordNet: a publicly available lexical resource for opinion mining, in: *International Conference on Language Resources and Evaluation*, 2006.
- [74] B. Liu, M. Hu, J. Cheng, Opinion observer: analyzing and comparing opinions on the web, in: *WWW*, 2005.
- [75] L. Oneto, S. Ridella, D. Anguita, Tikhonov, Ivanov and Morozov regularization for support vector machine learning, *Mach. Learn.* 103 (1) (2016) 103–136.
- [76] A. Go, R. Bhayani, L. Huang, Twitter sentiment classification using distant supervision, in: *CS224N Project Report*, Stanford 1 (2009) 12.
- [77] D. Anguita, A. Ghio, L. Oneto, S. Ridella, In-sample and out-of-sample model selection and error estimation for support vector machines, *IEEE Trans. Neural Netw. Learn. Syst.* 23 (9) (2012) 1390–1406.
- [78] L. Oneto, A. Ghio, S. Ridella, D. Anguita, Fully empirical and data-dependent stability-based bounds, *IEEE Trans. Cybernet.* 45 (9) (2015) 1913–1926.
- [79] P. Nakov, Z. Kozareva, A. Ritter, S. Rosenthal, V. Stoyanov, T. Wilson, Semeval-2013 task 2: sentiment analysis in twitter, in: *Annual Conference of the North American Chapter of the Association for Computational Linguistics*, 2013.
- [80] S. Rosenthal, P. Nakov, A. Ritter, V. Stoyanov, Semeval-2014 task 9: sentiment analysis in twitter, in: *Semantic Evaluation Exercises*, (2014).