Sentic Computing for Social Media Analysis, Representation, and Retrieval

Erik Cambria, Marco Grassi, Soujanya Poria, and Amir Hussain

Abstract As the web is rapidly evolving, web users are evolving with it. In the era of social colonisation, people are getting more and more enthusiastic about interacting, sharing and collaborating through social networks, online communities, blogs, wikis and other online collaborative media. In recent 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. These online social data, however, remain hardly accessible to computers, as they are specifically meant for human consumption. To overcome such obstacle, we need to explore more concept-level approaches that rely more on the implicit semantic texture of natural language, rather than its explicit syntactic structure. To this end, we further develop and apply sentic computing tools and techniques to the development of a novel unified framework for social media analysis, representation and retrieval. The proposed system extracts semantics from natural language text by applying graph mining and multidimensionality reduction techniques on an affective common sense knowledge base and makes use of them for inferring the cognitive and affective information associated with social media.

E. Cambria (✉)
Temasek Laboratories, National University of Singapore, Singapore, 117411, Singapore
e-mail: cambria@nus.edu.sg

M. Grassi
Department of Information Engineering, Universitá Politecnica delle Marche,
Ancona, 60131, Italy
e-mail: m.grassi@univpm.it

S. Poria
Department of Computer Science and Engineering, Jadavpur University, Kolkata, 700 032, India
e-mail: soujanya.poria@gmail.com

A. Hussain
Department of Computing Science and Mathematics, University of Stirling,
Scotland, FK9 4LA, UK
e-mail: ahu@cs.stir.ac.uk

1 Introduction

Web 2.0 has changed the ways people communicate, collaborate, express their opinions and sentiments. The distillation of knowledge from this huge amount of unstructured information is an extremely difficult task as today web contents are perfectly suitable for human consumption, but they remain hardly accessible to machines. The web, in fact, mostly owes its success to the development of search engines like Google and Yahoo, which represent the starting point for information retrieval. Such engines, which base their searches on keyword-based algorithms relying on the textual representation of the web page, are very good in retrieving texts, splitting them into parts, checking the spelling and counting their words. But when it comes to interpreting sentences and extracting useful information for users, their capabilities result still very limited.

Current attempts to perform automatic understanding of text, for example, textual entailment and machine reading, still suffer from numerous problems including inconsistencies, synonymy, polysemy, entity duplication and more, as they focus on a mere syntactical analysis of text. To bridge the cognitive and affective gap between word-level natural language data and the concept-level opinions and sentiments conveyed by them, we need intelligent user interfaces able to learn new affective common sense knowledge and to perform reasoning on it, in order to semantically and affectively analyse natural language text. In human cognition, thinking and feeling are mutually present: emotions are often the product of our thoughts as well as our reflections are often the product of our affective states. Emotions, in fact, are intrinsically part of our mental activity and play a key role in decision-making processes: they are special states, shaped by natural selection, to adjust various aspects of our organism to make it better face particular situations, for example, anger evolved for reaction, fear evolved for protection and affection evolved for reproduction [1].

For these reasons, we cannot prescind from emotions in the development of intelligent systems: if we want computers to be really intelligent, not just have the veneer of intelligence, we need to give them the ability to recognise, understand and express emotions. In this work, we further develop and apply AI tools and techniques (Sect. 2) to the development of a novel unified framework for analysing (Sect. 3), representing (Sect. 4) and retrieving (Sect. 5) social media. The developed system (Fig. 1), in particular, consists of four main modules:

1. NLP module: This module is in charge of preprocessing the input text by using the affective valence indicators that are usually contained in opinionated text such as special punctuation, complete upper-case words, onomatopoeic repetitions, exclamation words, degree adverbs and emoticons.
2. Semantic parsing module: This module deconstructs the text into concepts using a lexicon contains several n-grams extracted from different semantic resources. The input text is deconstructed into several small bags of concepts (SBoCs), which are used as inputs to the ConceptNet module and AffectiveSpace module to infer their relative cognitive and affective information, respectively.
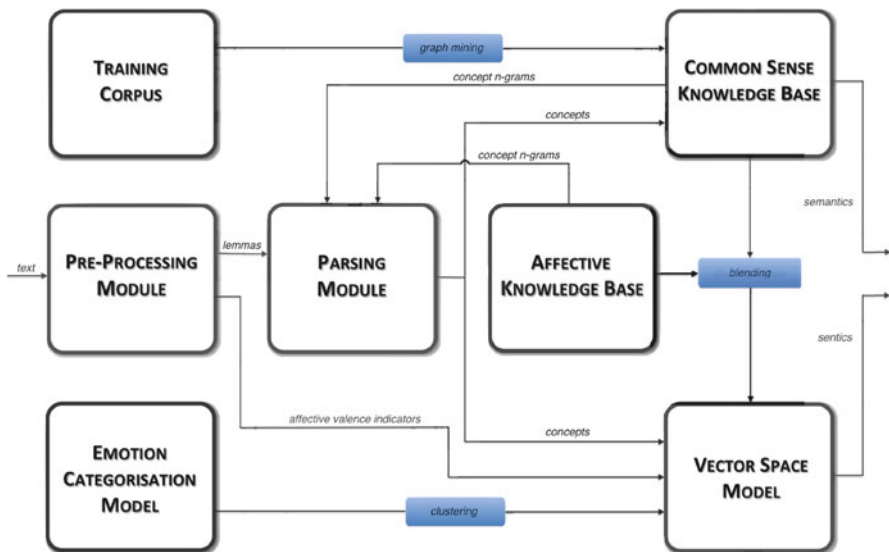
Fig. 1 System architecture. Graph mining and dimensionality reduction techniques are employed on two knowledge bases for open-domain sentiment analysis

3. ConceptNet module: This module exploits the graph representation of a common sense knowledge base to detect semantics. The concepts of each SBoC obtained from the output of the semantic parser are projected on the matrix resulting from many steps of spreading activation in order to calculate their semantic relatedness to each seed concept and, hence, their degree of belonging to each different class.
4. AffectiveSpace module: The concepts of each SBoC are projected into a vector space of affective common sense knowledge and clustered according to their coordinates in such space. This module assigns a score to each concept of the SBoC which defines the affinity of a concept belonging to a particular affective cluster.

2 Methodology

Sentic computing is a multidisciplinary approach to sentiment analysis, recently proposed by Cambria and Hussain [2], at the crossroads between affective computing and common sense computing. In the field of opinion mining, in fact, not only common sense knowledge but also emotional knowledge is important to grasp both the cognitive and affective information (termed semantics and sentics) associated with natural language opinions and sentiments.

Although scientific research in the area of emotion stretches back to the nineteenth century when Charles Darwin and William James proposed theories of

emotion that continue to influence thinking today [3, 4], the injection of affect into computer technologies is much more recent. During most of the last century, research on emotions was conducted by philosophers and psychologists, whose work was based on a small set of emotion theories that continue to underpin research in this area. The first researchers to try linking text to emotions were actually social psychologists and anthropologists who tried to find similarities on how people from different cultures communicate [5]. This research was also triggered by a dissatisfaction with the dominant cognitive view centred around humans as 'information processors' [6]. Later on, in the 1980s, researchers such as Turkle [7] began to speculate about how computers might be used to study emotions.

Systematic research programmes along this front began to emerge in the early 1990s. For example, Scherer [8] implemented a computational model of emotion as an expert system. A few years later, Picard's landmark book Affective Computing [9] prompted a wave of interest among computer scientists and engineers looking for ways to improve human–computer interfaces by coordinating emotions and cognition with task constraints and demands. Picard described three types of affective computing applications:

1. Systems that detect the emotions of the user
2. Systems that express what a human would perceive as an emotion
3. Systems that actually 'feel' an emotion

Although touching upon HCI [10] and affective modelling [11, 12], sentic computing primarily focuses on affect detection from text. Affect detection is critical because an affect-sensitive interface can never respond to users' affective states if it cannot sense their affective states. Affect detection need not be perfect but must be approximately on target. Affect detection is, however, a very challenging problem because emotions are constructs (i.e., conceptual quantities that cannot be directly measured) with fuzzy boundaries and with substantial individual difference variations in expression and experience.

To overcome such problem, sentic computing builds upon a brain-inspired and psychologically motivated affective categorisation model, proposed by Cambria et al. [13], that can potentially describe the full range of emotional experiences in terms of four independent but concomitant dimensions, whose different levels of activation make up the total emotional state of the mind. In sentic computing, whose term derives from the Latin *sentire* (root of words such as sentiment and sentience) and *sensus* (intended both as capability of feeling and as common sense), the analysis of natural language is based on affective ontologies and common sense reasoning tools, which enable the analysis of text not only at document, page or paragraph level but also at sentence and clause level.

In particular, sentic computing involves the use of AI and Semantic Web techniques, for knowledge representation and inference; mathematics, for carrying out tasks such as graph mining and multidimensionality reduction; linguistics, for discourse analysis and pragmatics; psychology, for cognitive and affective modelling; sociology, for understanding social network dynamics and social influence; and finally ethics, for understanding related issues about the nature of mind and

the creation of emotional machines. In this work, we exploit sentic computing tools and techniques to extract the semantics and sentics (i.e., the cognitive and affective information) associated with social media and, hence, bridge the gap between unstructured natural language data and structured machine-processable data. In particular, for the extraction of semantics, we use the following sentic computing tools and techniques:

1. A directed graph representation of common sense knowledge (Sect. 2.1)
2. A statistical method for the identification of common semantics (Sect. 2.2)
3. A technique that expands semantics through spreading activation (Sect. 2.3)

In turn, for the extraction of sentics, we use:

1. A language visualisation and analysis system (Sect. 2.4)
2. A novel emotion categorization model (Sect. 2.5)
3. A technique for clustering sentics (Sect. 2.6)

2.1 ConceptNet

ConceptNet [14] is a semantic resource structurally similar to WordNet, but whose scope of contents is general world knowledge, in the same vein as Cyc [15].Instead of insisting on formalising common sense reasoning using mathematical logic [16], ConceptNet uses a new approach: it represents data in the form of a semantic network and makes it available to be used in natural language processing. The prerogative of ConceptNet, in fact, is contextual common sense reasoning: while WordNet is optimised for lexical categorization and word-similarity determination, and Cyc is optimised for formalised logical reasoning, ConceptNet is optimised for making practical context-based inferences over real-world texts.

In ConceptNet, WordNet's notion of node in the semantic network is extended from purely lexical items (words and simple phrases with atomic meaning) to include higher-order compound concepts, for example, 'satisfy hunger' and 'follow recipe', to represent knowledge around a greater range of concepts found in everyday life (see Table 1). Moreover, WordNet's repertoire of semantic relations is extended from the triplet of synonym, is-a and part-of, to a repertoire of twenty semantic relations including, for example, EffectOf (causality), SubeventOf (event hierarchy), CapableOf (agent's ability), MotivationOf (affect), PropertyOf and LocationOf. ConceptNet's knowledge is also of a more informal, defeasible and practically valued nature. For example, WordNet has formal taxonomic knowledge that 'dog' is a 'canine', which is a 'carnivore', which is a 'placental mammal'; but it cannot make the practically oriented member-to-set association that 'dog' is a 'pet'. ConceptNet also contains a lot of knowledge that is defeasible, that is, it describes something that is often true but not always, for example, EffectOf ('fall off bicycle', 'get hurt'), which is something we cannot leave aside in common sense reasoning.

Table 1 Comparing WordNet and ConceptNet: while WordNet synsets contain vocabulary knowledge associated with concepts, ConceptNet assertions convey generic knowledge about what such concepts are used for

Term	WordNet hypernyms	ConceptNet assertions
Cat	Feline; felid; adult male; man; gossip; gossiper; gossipmonger; rumormonger; rumourmonger; newsmonger; woman; adult female; stimulant; stimulant drug; excitant; tracked vehicle	Cats can hunt mice; cats have whiskers; cats can eat mice; cats have fur; cats have claws; cats can eat meat; cats are cute
Dog	Canine; canid; unpleasant woman; disagreeable woman; chap; fellow; feller; lad; gent; fella; scoundrel; sausage; follow	Dogs are mammals; a dog can be a pet; a dog can guard a house; you are likely to find a dog in kennel; an activity a dog can do is run; a dog is a loyal friend; a dog has fur
Language	Communication; auditory communication; word; higher cognitive process; faculty; mental faculty; module; text; textual matter	English is a language; French is a language; language is used for communication; music is a language; a word is part of language
iPhone	N/A	An iPhone is a kind of a telephone; an iPhone is a kind of computer; an iPhone can display your position on a map; an iPhone can send and receive emails; an iPhone can display the time
Birthday gift	Present	Card is birthday gift; present is birthday gift; buying something for a loved one is for a birthday gift

Most of the facts interrelating ConceptNet's semantic network are dedicated to making rather generic connections between concepts. This type of knowledge can be brought back to Minsky's K-lines as it increases the connectivity of the semantic network and makes it more likely that concepts parsed out of a text document can be mapped into ConceptNet. ConceptNet is produced by an automatic process, which first applies a set of extraction rules to the semistructured English sentences of the Open Mind Common Sense (OMCS) corpus and then applies an additional set of 'relaxation' procedures, that is, filling in and smoothing over network gaps, to optimise the connectivity of the semantic network (Fig. 2). In ConceptNet version 2.0, a new system for weighting knowledge was implemented, which scores each binary assertion based on how many times it was uttered in the OMCS corpus and on how well it can be inferred indirectly from other facts in ConceptNet. In ConceptNet version 3.0 [17], users can also participate in the process of refining knowledge by evaluating existing statements on Open Mind Commons [18], the new interface for collecting common sense knowledge from users over the web.

By giving the user many forms of feedback and using inferences by analogy to find appropriate questions to ask, Open Mind Commons can learn well-connected structures of common sense knowledge, refine its existing knowledge and build analogies that lead to even more powerful inferences. The pieces of common
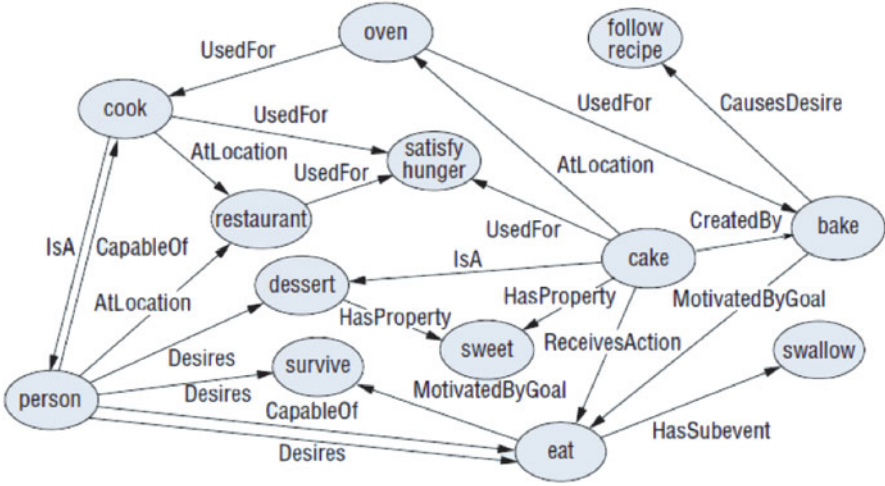
Fig. 2 ConceptNet represents the information in the Open Mind corpus as a directed graph where nodes are concepts and labelled edges are assertions of common sense that interconnect them

sense knowledge acquired through this interface are made publicly available in ConceptNet, which is released periodically both as an SQL database and through an API.

2.2 CF-IOF Weighting

CF-IOF (concept frequency – inverse opinion frequency) [19] is a technique that identifies domain-dependent semantics, using an approach similar to TF-IDF weighting, in order to evaluate how important a concept is to a set of opinions concerning the same topic. Firstly, the frequency of a concept c for a given domain d is calculated by counting the occurrences of the concept c in the set of available d-tagged opinions and dividing the result by the sum of number of occurrences of all concepts in the set of opinions concerning d. This frequency is then multiplied by the logarithm of the inverse frequency of the concept in the whole collection of opinions, that is:

$$\text{CF-IOF}_{c,d} = \frac{n_{c,d}}{\sum_k n_{k,d}} \log \sum_k \frac{n_k}{n_c}$$

where $n_{c,d}$ is the number of occurrences of concept c in the set of opinions tagged as d, n_k is the total number of concept occurrences and n_c is the number of occurrences of c in the whole set of opinions. A high weight in CF-IOF is reached by a high concept frequency in a given domain and a low frequency of the concept in the whole collection of opinions. Therefore, thanks to CF-IOF weights, it is possible to filter out common concepts and detect relevant domain-dependent semantics.

2.3 Spectral Association

Spectral association [20] is a technique that involves assigning values, or activations, to 'seed concepts' and applying an operation that spreads their values across the ConceptNet graph. This operation, an approximation of many steps of spreading activation, transfers the most activation to concepts that are connected to the key concepts by short paths or many different paths in common sense knowledge. In particular, we build a matrix C that relates concepts to other concepts, instead of their features, and add up the scores over all relations that relate one concept to another, disregarding direction. Applying C to a vector containing a single concept spreads that concept's value to its connected concepts. Applying C^2 spreads that value to concepts connected by two links (including back to the concept itself). But what we would really like is to spread the activation through any number of links, with diminishing returns, so the operator we want is:

$$1 + C + \frac{C^2}{2!} + \frac{C^3}{3!} + \ldots = e^C$$

We can calculate this odd operator, e^C, because we can factor C. C is already symmetric, so instead of applying Lanczos' method to CC^T and getting the singular value decomposition (SVD), we can apply it directly to C and get the spectral decomposition $C = V\Lambda V^T$. As before, we can raise this expression to any power and cancel everything but the power of Λ. Therefore, $e^C \approx Ve^\Lambda V^T$. This simple twist on the SVD lets us calculate spreading activation over the whole matrix instantly. As with the SVD, we can truncate these matrices to k axes and therefore save space while generalising from similar concepts. We can also rescale the matrix so that activation values have a maximum of 1 and do not tend to collect in highly connected concepts such as 'person', by normalising the truncated rows of $Ve^{\Lambda/2}$ to unit vectors, and multiplying that matrix by its transpose to get a rescaled version of $Ve^\Lambda V^T$.

2.4 AffectiveSpace

AffectiveSpace is a multidimensional vector space built by 'blending' [21] ConceptNet with WordNet-Affect (WNA) [22], a linguistic resource for the lexical representation of affective knowledge. Blending is a technique that performs inference over multiple sources of data simultaneously, taking advantage of the overlap between them. It basically combines two sparse matrices linearly into a single matrix in which the information between the two initial sources is shared. When we perform SVD on a blended matrix, the result is that new connections are made in each source matrix taking into account information and connections present in the other matrix, originating from the information that overlaps. The alignment operation operated over ConceptNet and WNA yields a new matrix, A, in which

common sense and affective knowledge coexist, that is, a matrix $14{,}301 \times 117{,}365$ whose rows are concepts (e.g., 'dog' or 'bake cake'), whose columns are either common sense or affective features (e.g., 'isA-pet' or 'hasEmotion-joy') and whose values indicate truth values of assertions. Therefore, in A, each concept is represented by a vector in the space of possible features whose values are positive for features that produce an assertion of positive valence (e.g., 'a penguin is a bird'), negative for features that produce an assertion of negative valence (e.g., 'a penguin cannot fly') and zero when nothing is known about the assertion.

The degree of similarity between two concepts, then, is the dot product between their rows in A. The value of such a dot product increases whenever two concepts are described with the same feature and decreases when they are described by features that are negations of each other. In particular, we use truncated singular value decomposition (TSVD) [23] in order to obtain a new matrix containing both hierarchical affective knowledge and common sense. The resulting matrix has the form $\tilde{A} = U_k * \Sigma_k * V_k^T$ and is a low-rank approximation of A, the original data. This approximation is based on minimising the Frobenius norm of the difference between A and $\tilde{A}$ under the constraint $\mathrm{rank}(\tilde{A}) = k$. For the Eckart–Young theorem [24], it represents the best approximation of A in the mean-square sense, in fact

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

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:

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

Therefore, $\tilde{A}$ of rank k is the best approximation of A in the Frobenius norm sense when $\sigma_i = s_i$ $(i = 1, \ldots, k)$, and the corresponding singular vectors are the same as those of A. If we choose to discard all but the first k principal components, common sense concepts and emotions are represented by vectors of k coordinates: these coordinates can be seen as describing concepts in terms of 'eigenmoods' that form the axes of AffectiveSpace, that is, the basis $e_0, \ldots, e_{k-1}$ of the vector space (Fig. 3). 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. Concepts with negative e_0 components, then, are likely to have negative affective valence. Thus, by exploiting the information sharing property of TSVD, concepts with the same affective valence are likely to have similar features – that is, tend to fall near each other in AffectiveSpace. Concept similarity does not depend on their absolute positions in the vector space, but rather on the angle they make with the origin. For example, we can find concepts such as 'beautiful day', 'birthday
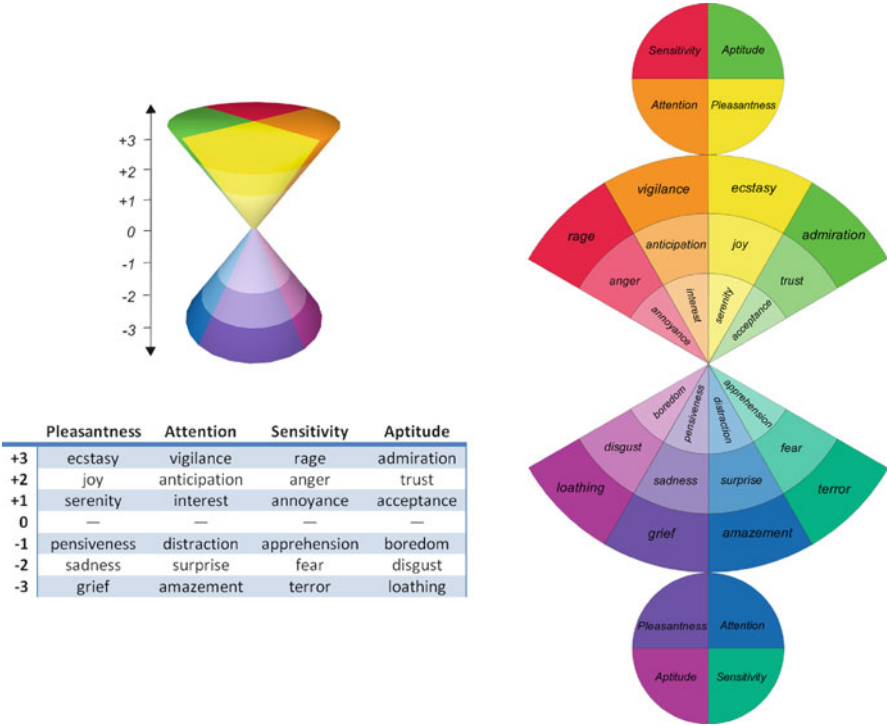
Fig. 3 Affectively positive (*bottom-left* corner) and affectively negative (*up-right* corner) common sense concepts in AffectiveSpace

party', 'laugh' and 'make person happy' very close in direction in the vector space, while concepts like 'sick', 'feel guilty', 'be laid off' and 'shed tear' are found in a completely different direction.

2.5 The Hourglass of Emotions

This model is a variant of Plutchik's emotion categorization [25] and constitutes an attempt to emulate Marvin Minsky's theories on human emotions. Minsky sees the mind as made up of thousands of different resources and believes that our emotional states result from turning one set of these resources on and turning another set of them off [1]. 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 [44]. The Hourglass of Emotions (Fig. 4) is specifically designed to recognise, understand and express emotions in the context of human–computer interaction (HCI). In the model, in fact, affective

	Pleasantness	Attention	Sensitivity	Aptitude
+3	ecstasy	vigilance	rage	admiration
+2	joy	anticipation	anger	trust
+1	serenity	interest	annoyance	acceptance
0	—	—	—	—
−1	pensiveness	distraction	apprehension	boredom
−2	sadness	surprise	fear	disgust
−3	grief	amazement	terror	loathing

Fig. 4 The 3D model and the net of the Hourglass of Emotions. Dimensional and discrete forms of the different sentic levels are summarised in the proposed emotion categorization table

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 in order to understand how much respectively:

1. The user is happy with the service provided (pleasantness).
2. The user is interested in the information supplied (attention).
3. The user is comfortable with the interface (sensitivity).
4. The user is disposed to use the application (aptitude).

Each affective dimension is characterised by six levels of activation, called 'sentic levels', which determine the intensity of the expressed/perceived emotion as an int $\in [-3,+3]$. These levels are also labelled as a set of 24 basic emotions (six for each of the affective dimensions) in a way that allows the model to specify the affective information associated with text both in a dimensional and in a discrete form. The dimensional form, in particular, is called 'sentic vector', and it is a four-dimensional *float* vector that can potentially express any human emotion in terms of pleasantness, attention, sensitivity and aptitude. Some particular sets of sentic vectors have special names as they specify well-known compound emotions. For

example, the set of sentic vectors with a level of pleasantness $\in (+1,+2]$ ('joy'), a null attention, a null sensitivity and a level of aptitude $\in (+1,+2]$ ('trust') are called 'love sentic vectors' since they specify the compound emotion of 'love'.

2.6   Sentic Medoids

Sentic medoids [26] is a clustering technique that adopts a k-medoids approach [27] to partition affective common sense concepts in AffectiveSpace into k clusters around as many centroids, trying to minimise a given cost function. Differently from the k-means algorithm [28], which does not pose constraints on centroids, k-medoids do assume that centroids must coincide with k observed points. The k-means approach finds the k centroids, where the coordinate of each centroid is the mean of the coordinates of the objects in the cluster and assigns every object to the nearest centroid. Unfortunately, k-means clustering is sensitive to the outliers, and a set of objects closest to a centroid may be empty, in which case centroids cannot be updated. For this reason, k-medoids are sometimes used, where representative objects are considered instead of centroids. In many clustering problems, in fact, one is interested in the characterisation of the clusters by means of typical objects, which represent the various structural features of objects under investigation. Because it uses the most centrally located object in a cluster, k-medoids clustering is less sensitive to outliers compared with k-means.

Among many algorithms for k-medoids clustering, partitioning around medoids (PAM) is one of the most widely used. The algorithm, proposed by Kaufman and Rousseeuw [27], first computes k representative objects, called medoids. A medoid can be defined as that object of a cluster, whose average dissimilarity to all the objects in the cluster is minimal. PAM determines a medoid for each cluster selecting the most centrally located centroid within the cluster. After selection of medoids, clusters are rearranged so that each point is grouped with the closest medoid. Compared to k-means, PAM operates on the dissimilarity matrix of the given data set. It is more robust, because it minimises a sum of dissimilarities instead of a sum of squared Euclidean distances. A particularly nice property is that PAM allows clustering with respect to any specified distance metric. In addition, the medoids are robust representations of the cluster centres, which is particularly important in the common context that many elements do not belong well to any cluster. However, PAM works inefficiently for large data sets due to its complexity.

To this end, a modified version of the algorithm recently proposed by Park and Jun [29] was used, which runs similarly to the k-means clustering algorithm. This has shown to have similar performance when compared to PAM algorithm while taking a significantly reduced computational time. In particular, we have N concepts ($N = 14,301$) encoded as points $x \in \mathbb{R}^p (p = 50)$. We want to group them into k clusters, and, in our case, we can fix $k = 24$ as we are looking for one cluster for each sentic level s of the Hourglass model. Generally, the initialization of clusters for clustering algorithms is a problematic task as the process often risks to get

stuck into local optimum points, depending on the initial choice of centroids [30]. However, we decide to use as initial centroids the concepts that are currently used as centroids for clusters, as they specify the emotional categories we want to organise AffectiveSpace into. For this reason, what is usually seen as a limitation of the algorithm can be seen as advantage for this approach, since we are not looking for the 24 centroids leading to the best 24 clusters but indeed for the 24 centroids identifying the required 24 sentic levels (i.e., the centroids should not be 'too far' from the ones currently used). In particular, as the Hourglass affective dimensions are independent but concomitant, we need to cluster AffectiveSpace four times, once for each dimension. According to the Hourglass categorization model, in fact, each concept can convey, at the same time, more than one emotion (which is why we get compound emotions), and this information can be expressed via a sentic vector specifying the concept's affective valence in terms of pleasantness, attention, sensitivity and aptitude. Therefore, given that the distance between two points in AffectiveSpace is defined as $D(a, b) = \sqrt{\sum_{i=1}^{p} (a_i - b_i)^2}$ (note that the choice of Euclidean distance is arbitrary), the used algorithm, applied for each of the four affective dimensions, can be summarised as follows:

1. Each centroid $C_n \in \mathbb{R}^{50}$ $(n = 1, 2, \ldots, k)$ is set as one of the six concepts corresponding to each s in the current affective dimension.
2. Assign each record x to a cluster Ξ so that $x_i \in \Xi_n$ if $D(x_i, C_n) \leq D(x_i, C_m)$ $m = 1, 2, \ldots, k$.
3. Find a new centroid C for each cluster Ξ so that $C_j = x_i$
   if $\sum_{x_m \in \Xi_j} D(x_i, x_m) \leq \sum_{x_m \in \Xi_j} D(x_h, x_m)$ $\forall x_h \in \Xi_j$.
4. Repeat steps 2 and 3 until no changes on centroids are observed.

Note that condition posed on steps 2 and 3 may occasionally lead to more than one solution. Should this happen, our model will randomly choose one of them. This clusterization of AffectiveSpace allows to calculate, for each common sense concept x, a four-dimensional sentic vector that defines its affective valence in terms of a degree of fitness $\mathbf{f}(x)$ where $f_a = D(x, C_j)$ $C_j | D(x, C_j) \leq D(x, C_k)$ $a = 1, 2, 3, 4$ $k = 6a\text{-}5, 6a\text{-}4, \ldots, 6a$.

3 System Architecture

In order to effectively mine and analyse opinions and sentiments, it is necessary to bridge the gap between unstructured natural language data and structured machine-processable data. To this end, an intelligent software engine has been proposed by Cambria et al. [31] that aims to extract the semantics and sentics, that is, the cognitive and affective information, associated with natural language text, in a way that the opinions and sentiments contained in it can be more easily aggregated and interpreted. The engine exploits graph mining and multidimensionality reduction techniques on ConceptNet, and it is based on the Hourglass model (Fig. 5).
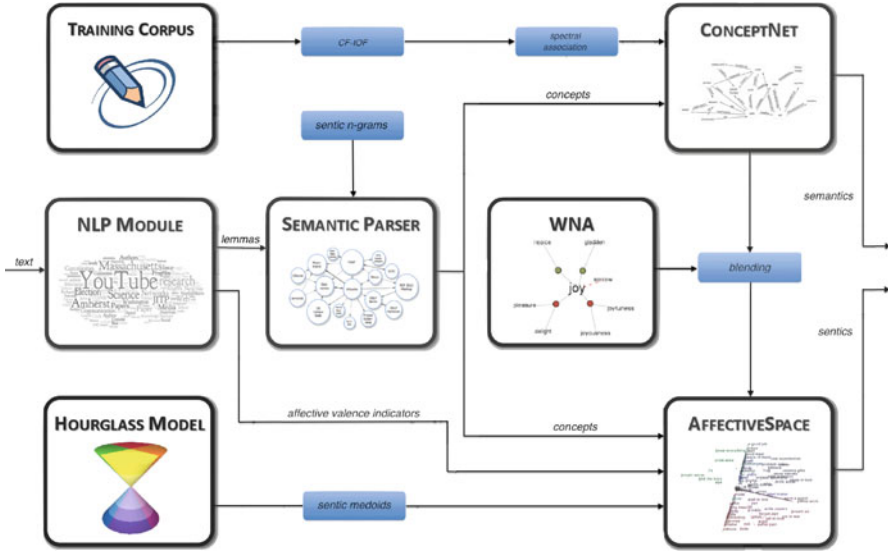
Fig. 5 Opinion mining engine *block diagram*. After performing a first skim of the input text, the engine extracts concepts from it and, hence, infers related semantics and sentics

Table 2 An overview of recent model-based affect recognition and sentiment analysis systems. Studies are divided by techniques applied, number of categories of the model adopted, corpora and knowledge base used

Study	Techniques	Model	Corpora	Knowledge base
[32]	NB, SVM	2 categories	Political articles	None
[33]	LSA, MLP, NB, KNN	3 categories	Dialogue turns	ITS interaction
[34]	Cohesion indices	4 categories	Dialogue logs	ITS interaction
[35]	VSM, NB, SVM	5 categories	ISEAR	ConceptNet
[36]	WN presence, LSA	6 categories	News stories	WNA
[37]	WN presence	6 categories	Chat logs	WNA
[38]	Winnow linear, C4.5	7 categories	Children stories	None
[39]	VSM, KNN	24 categories	LiveJournal	ConceptNet, WNA
[31]	VSM, k-means	24 categories	YouTube, LiveJournal	ConceptNet, WNA, HEO
[40]	VSM, k-means	24 categories	LiveJournal, Patient Opinion	ConceptNet, WNA
[41]	VSM, k-medoids	24 categories	Twitter, LiveJournal, Patient Opinion	ConceptNet, Probase

Several other affect recognition and sentiment analysis systems [32–38] are based on different emotion categorisation models, which generally comprise a relatively small set of categories (Table 2). The Hourglass of Emotions, in turn, allows the opinion mining engine to classify affective information both in a categorical way (according to a wider number of emotion categories) and in a

dimensional format (which facilitates comparison and aggregation). Such engine, in particular, consists of four main components: an NLP module, which performs a first skim of the opinion (Sect. 3.1); a semantic parser, whose aim is to extract concepts from the opinionated text (Sect. 3.2); the ConceptNet module, for inferring the semantics associated with the given concepts (Sect. 3.3); and the AffectiveSpace module, for the extraction of sentics (Sect. 3.4). Eventually, this section illustrates an output example of the engine, given a short natural language sentence as input (Sect. 3.5).

3.1 NLP Module

This preprocessing module firstly interprets all the affective valence indicators usually contained in opinionated text such as special punctuation, complete upper-case words, onomatopoeic repetitions, exclamation words, degree adverbs and emoticons. Secondly, the module detects negation and spreads it in a way that it can be accordingly associated to concepts during the parsing phase. Handling negation is an important concern in opinion- and sentiment-related analysis, as it can reverse the meaning of a statement.

Such task, however, is not trivial as not all appearances of explicit negation terms reverse the polarity of the enclosing sentence and that negation can often be expressed in rather subtle ways, for example, sarcasm and irony, which are quite difficult to detect. Lastly, the module converts text to lower case and, after lemmatising it, splits the opinion into single clauses according to grammatical conjunctions and punctuation.

3.2 Semantic Parser

The semantic parser deconstructs text into concepts using a lexicon based on sequences of lexemes that represent multiple-word concepts extracted from ConceptNet and WordNet. These n-grams are not used blindly as fixed word patterns but exploited as reference for the module, in order to extract multiple-word concepts from information-rich sentences. So, differently from other shallow parsers, the module can recognise complex concepts also when irregular verbs are used or when these are interspersed with adjective and adverbs, for example, the concept 'buy christmas present' in the sentence 'I bought a lot of very nice Christmas presents'. For each clause, the module outputs a small bag of concepts (SBoC), which is later on analysed separately by the ConcepNet and AffectiveSpace modules to infer the cognitive and affective information associated with the input text, respectively. In case any of the detected concepts is found more than once in the vector space (that is, any of the concepts has multiple senses), all the SBoC concepts are exploited

for a context-dependent coarse sense disambiguation. In particular, to represent the expected semantic value of the clause as a whole, the vectors corresponding to all concepts in the clause (in their ambiguous form) can be averaged together. The resulting vector does not represent a single meaning but the 'ad hoc category' of meanings that are similar to the various possible meanings of concepts in the clause [42]. Then, to assign the correct sense to the ambiguous concept, the concept sense with the highest dot product (and thus the strongest similarity) with the clause vector is searched.

3.3   ConceptNet Module

Once natural language text is deconstructed into concepts, these are given as input to both the ConceptNet and the AffectiveSpace modules. While the former exploits the graph representation of the affective common sense knowledge base to detect semantics, the latter exploits the vector space representation of ConceptNet to infer sentics. In particular, the ConceptNet module applies spectral association for assigning activation to key concepts, that is, nodes of the semantic network, which are used as seeds or centroids for classification. Such seeds can simply be the concepts corresponding to the class labels of interest plus their available synonyms and antonyms, if any.

As shown in Sect. 2.3, seeds can also be found by applying CF-IOF on a training corpus (when available), in order to perform a classification that is more relevant to the data under analysis. After seeds concepts are identified, the module spreads their values across the ConceptNet graph. This operation, an approximation of many steps of spreading activation, transfers the most activation to concepts that are connected to the seed concepts by short paths or many different paths in affective common sense knowledge. Therefore, the concepts of each SBoC provided by the semantic parser are projected on the matrix resulting from spectral association in order to calculate their semantic relatedness to each seed concept and, hence, their degree of belonging to each different class. Such classification measure is directly proportional to the degree of connectivity between the nodes representing the retrieved concepts and the seed concepts in the affective common sense knowledge graph.

3.4   AffectiveSpace Module

In the ConceptNet module, graph-mining techniques are exploited to extract semantics from the concepts retrieved by the semantic parser. Such concepts are also given as input to the AffectiveSpace module, which, in turn, exploits dimensionality reduction techniques to infer the affective information associated with them. To this

end, the concepts of each SBoC are projected into AffectiveSpace, and, according to their position in the vector space representation of affective common sense knowledge, they are assigned to an affective class defined through the sentic medoids technique.

As well as in the ConceptNet module, the categorisation does not consist in simply labelling each concept but also in assigning a confidence score to each emotional label, which is directly proportional to the degree of belonging to a specific affective cluster (dot product between the given concept and the relative sentic medoid). Such affective information can also be exploited to calculate a polarity value associated with each SBoC provided by the semantic parser as well as to detect the overall polarity associated with the opinionated text.

3.5 Output Example

As an example of how the software engine works, intermediate and final outputs obtained when a natural language opinion is given as input to the system can be examined. The following tweet was chosen: 'I think iPhone4 is the top of the heap! OK, the speaker is not the best i hv ever seen bt touchscreen really puts me on cloud 9...camera looks pretty good too!'. After the preprocessing and semantic parsing operations, the following SBoCs are obtained:

SBoC#1:
  <Concept: 'think'>
  <Concept: 'iphone4'>
  <Concept: 'top heap'>
SBoC#2:
  <Concept: 'ok'>
  <Concept: 'speaker'>
  <Concept: !'good'++>
  <Concept: 'see'>
SBoC#3:
  <Concept: 'touchscreen'>
  <Concept: 'put cloud nine'++>
SBoC#4:
  <Concept: 'camera'>
  <Concept: 'look good'−−>

These are then concurrently processed by the ConceptNet and the AffectiveSpace modules, which output the cognitive and affective information associated with each SBoC, both in a discrete way, with one or more labels, and in a dimensional way, with a polarity value $\in [-1,+1]$ (Table 3).

Table 3 Structured output example of opinion mining engine. For each clause, the engine detects the opinion target, the category it belongs to and the affective information associated with it

Opinion target	Category	Moods	Polarity
'iphone4'	'phones', 'electronics'	'ecstasy', 'interest'	+0.71
'speaker'	'electronics', 'music'	'annoyance'	−0.34
'touchscreen'	'electronics'	'ecstasy', 'anticipation'	+0.82
'camera'	'photography', 'electronics'	'acceptance'	+0.56

4   Data Model

Our framework for social media representation and analysis aims to be applicable to most of online resources (videos, images, text) coming from different sources, for example, online video sharing services, blogs and social networks.

To such purpose, it is necessary to standardise as much as possible the descriptors used in encoding the information about multimedia resources and people to which the text refer (considering that every website uses its own vocabulary) in order to make it univocally interpretable and suitable to feed other applications. To achieve this purpose, Semantic Web techniques are exploited.

The Semantic Web initiative by W3C[1] 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 RDF data model to describe such resources in univocally interpretable format, whose basic building block is an object-attribute-value triple, that is, 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', and '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 XML syntax. Each triple can also be seen as a directed graph with labelled 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.

---

[1] http://w3.org

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, that is, 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) and OWL (Ontology Web Language). RDFS can be seen as a RDF vocabulary and a primitive ontology language. It offers certain modelling 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 modelling 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). The proposed framework for opinions and affective information description aims to be applicable to most of online resources (videos, images, text) coming from different sources, for example, online video sharing services, blogs and social networks. To such purpose, it is necessary to standardise 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.

To this end, we encode the cognitive and affective information associated with multimedia resources and people using the descriptors provided by OMR (Ontology for Media Resources), FOAF (Friend of a friend ontology), HEO (Human Emotion Ontology) [43], and WNA (WordNet-Affect) [22]. 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 recognised 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 needed for describing media and people and other descriptors are added only when necessary. For example, OMR, at least in the current realisation, does not supply vocabulary for describing comments, which are analysed to extract the
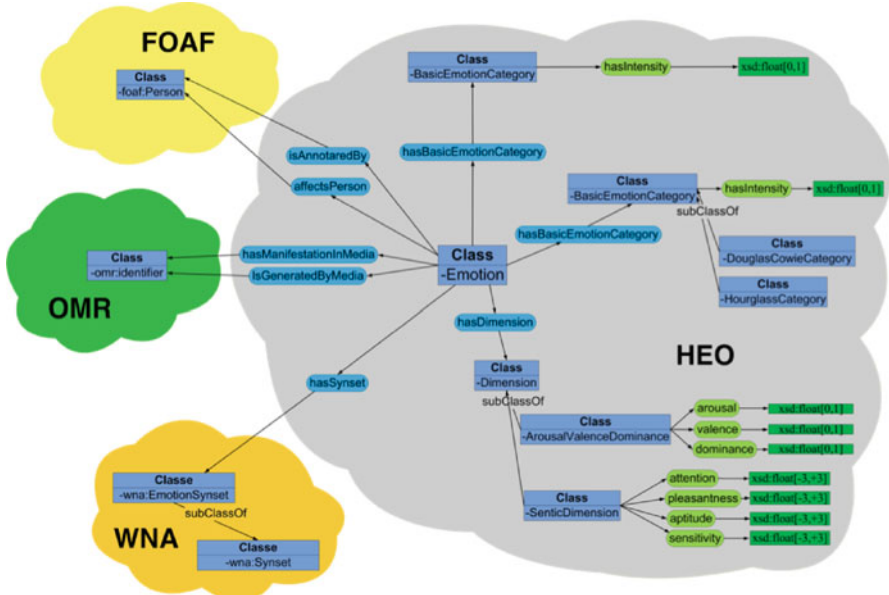
Fig. 6 Social media representation. *HEO*, *WNA*, *OMR* and *FOAF* are accordingly merged for effectively representing the cognitive and affective information associated with social media

affective information relative to media. This ontology is extended by introducing the 'Comment' class and by defining for it the 'author', 'text' and 'publicationDate' properties.

HEO is a high-level ontology for human emotions that supplies the most significant concepts and properties which constitute the centrepiece for the description of every human emotion. The main purpose of HEO is 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 emotion feature descriptors, and interoperability, by allowing to map concepts and properties belonging to different emotion representation models. 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 organisation of emotions in HEO, we exploit WNA, a linguistic resource for the lexical representation of affective knowledge, built by assigning to a number of WordNet synsets one or more affective labels (a-labels) and then by extending the core with the relations defined in WordNet. Thus, the combination of HEO with WNA, OMR and FOAF provides a complete framework to describe not only multimedia contents and the users that have created, uploaded or interacted with them but also the opinions and the affective content carried by the media and the way they are perceived by people (Fig. 6).

5 User Interface

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.

Such interdependence can be exploited, for example, to find similar patterns in customer reviews of commercial products and hence to gather useful information for marketing, sales, public relations and customer service. To visualise the cognitive and affective information associated to social media, we exploit the multifaceted 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, predetermined, 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 multidimensional information space by concurrently writing queries in a text box and progressively narrowing choices in each dimension.

For our framework, we use SIMILE Exhibit API, 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 multifaceted classification paradigm to display semantically structured data stored in a Semantic Web aware format, for example, 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 encode the cognitive and affective information associated with social media in RDF/XML, using the descriptors defined by HEO, WNA, OMR and FOAF, and store it in a Sesame triple store, a purpose-built database for the storage and retrieval of RDF meta-data. 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 (Fig. 7).

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. We export all the information contained in the triplestore into a JSON file to feed the Exhibit interface, in order to make it available for being browsed as a unique knowledge base. We choose to use Exhibit in our framework due to the ease with which it allows to create rich and interactive web pages. Social media are displayed in a dynamic gallery that can be ordered according to different parameters and the cognitive and affective information associated with them. Using faceted
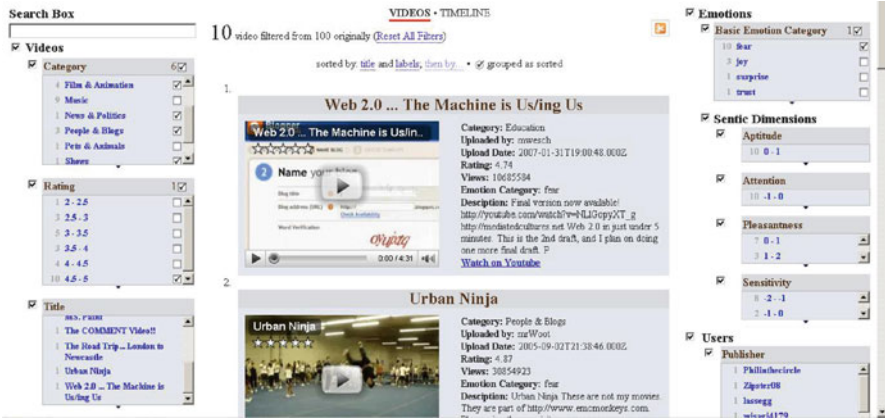
Fig. 7 Social media retrieval. The faceted classification interface allows multimodal social media retrieval according to the semantics and sentics associated with them

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, that is, by adding or removing constraints on the facet properties. 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.

6 Conclusions

With the advent of the social web, the way people express their views and opinions has dramatically changed. They can now post reviews of products at merchant sites and express their views on almost anything in Internet forums, discussion groups, and blogs. Such online word-of-mouth behaviour represents new and measurable 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. In many cases, in fact, opinions are hidden in long forum posts and blogs. It is extremely time consuming for a human reader to find relevant sources, extract related sentences with opinions, read them, summarise them and organise them into usable forms. Thus, automated opinion discovery and summarisation systems are needed. Sentiment analysis, also known as opinion mining, grows out of this need. It is a challenging NLP or text mining problem. Due to its tremendous value for practical applications, there has been an explosive growth of both research in academia and applications in the industry.

Due to many challenging research problems and a wide variety of practical applications, opinion mining has been a very active research area in recent years. All the sentiment analysis tasks, however, are very challenging. Our understanding and knowledge of the problem and its solution are still limited. The main reason is that it is a NLP task, and NLP has no easy problems. Another reason may be due to our popular ways of doing research. So far, in fact, researchers have probably relied too much on machine learning algorithms. Some of the most effective machine learning algorithms, for example, SVM and CRF, in fact, produce no human understandable results such that, although they may achieve improved accuracy, little about how and why is known, apart from some superficial knowledge gained in the manual feature engineering process.

All such approaches, moreover, rely on syntactical structure of text, which is far from the way human mind processes natural language. In this work, common sense computing techniques were further developed and applied to bridge the semantic gap between word-level natural language data and the concept-level opinions conveyed by these. In particular, the ensemble application of graph mining and multi-dimensionality reduction techniques was exploited on a common sense knowledge base to develop a novel intelligent engine for open-domain opinion mining and sentiment analysis. The proposed framework performs a clause-level semantic analysis of text, which allows the inference of both the conceptual and emotional information associated with natural language opinions and, hence, a more efficient passage from (unstructured) textual information to (structured) machine-processable data.

References

1. Minsky, M.: The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind. Simon and Schuster, New York (2006)
2. Cambria, E., Hussain, A.: Sentic Computing: Techniques, Tools, and Applications. Dordrecht, Netherlands: Springer (2012)
3. Charles, D.: The Expression of the Emotions in Man and Animals. John Murray, London (1872)
4. James, W.: What is an emotion? Mind **34**, 188–205 (1884)
5. Osgood, C., May, W., Miron, M.: Cross-cultural universals of affective meaning. University of Illinois, Urbana (1975)
6. Lutz, C., White, G.: The anthropology of emotions. Ann. Rev. Anthropol. **15**, 405–436 (1986)
7. Turkle, S.: The Second Self: Computers and the Human Spirit. Simon and Schuster, New York (1984)
8. Scherer, K.: Studying the emotion-antecedent appraisal process: an expert system approach. Cognit. Emot. **7**, 325–355 (1993)
9. Picard, R.: Affective computing. MIT, Boston (1997)
10. Cambria, E., Hupont, I., Hussain, A., Cerezo, E., Baldassarri, S.: Sentic avatar: multimodal affective conversational agent with common sense. In: Esposito, A., Hussain, A., Faundez-Zanuy, M., Martone, R., Melone, N. (eds.) Toward Autonomous, Adaptive, and Context-Aware Multimodal Interfaces: Theoretical and Practical Issues. Lecture Notes in Computer Science, vol. 6456, pp. 82–96. Springer, Berlin/Heidelberg (2011)

11. Cambria, E., Olsher, D., Kwok, K.: Sentic activation: a two-level affective common sense reasoning framework. In: Proceedings of the AAAI, Toronto, pp. 186–192 (2012)
12. Cambria, E., Olsher, D., Kwok, K.: Sentic panalogy: swapping affective common sense reasoning strategies and foci. In: Proceedings of the CogSci, Sapporo, pp. 174–179 (2012)
13. Cambria, E., Hussain, A., Havasi, C., Eckl, C.: SenticSpace: visualizing opinions and sentiments in a multi-dimensional vector space. In: Setchi, R., Jordanov, I., Howlett, R., Jain, L. (eds.) Knowledge-Based and Intelligent Information and Engineering Systems. Lecture Notes in Artificial Intelligence, vol. 6279, pp. 385–393. Springer, Berlin (2010)
14. Liu, H., Singh, P.: ConceptNet: a practical commonsense reasoning toolkit. BT Technol. J. **22**(4), 211–226 (2004)
15. Lenat, D., Guha, R.: Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project. Addison-Wesley, Boston (1989)
16. Mueller, E.: Commonsense Reasoning. Morgan Kaufmann, Amsterdam/Boston (2006)
17. Havasi, C., Speer, R., Alonso, J.: ConceptNet 3: a flexible, multilingual semantic network for common sense knowledge. In: Proceedings of the RANLP, Borovets (2007)
18. Speer, R.: Open mind commons: an inquisitive approach to learning common sense. In: Proceedings of the Workshop on Common Sense and Interactive Applications, Honolulu (2007)
19. Cambria, E., Hussain, A., Durrani, T., Havasi, C., Eckl, C., Munro, J.: Sentic computing for patient centered application. In: Proceedings of the IEEE ICSP, Beijing, pp. 1279–1282 (2010)
20. Havasi, C., Speer, R., Holmgren, J.: Automated color selection using semantic knowledge. In: Proceedings of the AAAI CSK, Arlington (2010)
21. Havasi, C., Speer, R., Pustejovsky, J., Lieberman, H.: Digital intuition: applying common sense using dimensionality reduction. IEEE Intell. Syst. **24**(4), 24–35 (2009)
22. Strapparava, C., Valitutti, A.: WordNet-affect: an affective extension of WordNet. In: Proceedings of the LREC, Lisbon (2004)
23. Wall, M., Rechtsteiner, A., Rocha, L.: Singular value decomposition and principal component analysis. In: Berrar, D., Dubitzky, W., Granzow, M. (eds.) A Practical Approach to Microarray Data Analysis, pp. 91–109. Kluwer Academic Publishers, Boston (2003)
24. Eckart, C., Young, G.: The approximation of one matrix by another of lower rank. Psychometrika **1**(3), 211–218 (1936)
25. Plutchik, R.: The nature of emotions. Am. Sci. **89**(4), 344–350 (2001)
26. Cambria, E., Mazzocco, T., Hussain, A., Eckl, C.: Sentic medoids: organizing affective common sense knowledge in a multi-dimensional vector space. In: Liu, D., Zhang, H., Polycarpou, M., Alippi, C., He, H. (eds.) Advances in Neural Networks. Lecture Notes in Computer Science, vol. 6677, pp. 601–610. Springer, Berlin (2011)
27. Kaufman, L., Rousseeuw, P.: Finding Groups in Data: An Introduction to Cluster Analysis. Wiley, New York (1990)
28. Hartigan, J., Wong, M.: Algorithm AS 136: a k-means clustering algorithm. J. R. Stat. Soc. **28**(1), 100–108 (1979)
29. Park, H., Jun, C.: A simple and fast algorithm for k-medoids clustering. Expert Syst. Appl. **36**(2), 3336–3341 (2009)
30. Duda, R., Hart, P.: Pattern Classification and Scene Analysis. Wiley, New York (1973)
31. Cambria, E., Grassi, M., Hussain, A., Havasi, C.: Sentic computing for social media marketing. Multimed. Tools Appl. **59**(2), 557–577 (2012)
32. Lin, W., Wilson, T., Wiebe, J., Hauptmann, A.: Which side are you on? Identifying perspectives at the document and sentence levels. In: Proceedings of the Conference on Natural Language Learning, New York, pp. 109–116 (2006)
33. D'Mello, S., Craig, S., Sullins, J., Graesser, A.: Predicting affective states expressed through an emote-aloud procedure from autotutor's mixed-initiative dialogue. Int. J. Artif. Intell. Educ. **16**, 3–28 (2006)
34. D'Mello, S., Dowell, N., Graesser, A.: Cohesion relationships in tutorial dialogue as predictors of affective states. In: Proceedings of the Conference Artificial Intelligence in Education, pp. 9–16. Springer, New York (2009)

35. Danisman, T., Alpkocak, A.: Feeler: emotion classification of text using vector space model. In: Proceedings of the AISB, Aberdeen (2008)
36. Strapparava, C., Mihalcea, R.: Learning to identify emotions in text. In: Proceedings of the ACM Symposium Applied Computing, pp. 1556–1560. ACM, New York (2008)
37. Ma, C., Osherenko, A., Prendinger, H., Ishizuka, M.: A chat system based on emotion estimation from text and embodied conversational messengers. In: Proceedings of the International Conference Active Media Technology, pp. 546–548. IEEE, Piscataway (2005)
38. Alm, C., Roth, D., Sproat, R.: Emotions from text: machine learning for text-based emotion prediction. In: Proceedings of the HLT/EMNLP, pp. 347–354. Association for Computing Linguistics, Morristown (2005)
39. Grassi, M., Cambria, E., Hussain, A., Piazza, F.: Sentic web: a new paradigm for managing social media affective information. Cognit. Comput. **3**(3), 480–489 (2011)
40. Cambria, E., Benson, T., Eckl, C., Hussain, A.: Sentic PROMs: application of sentic computing to the development of a novel unified framework for measuring health-care quality. Expert Syst. Appl. **39**(12), 10533–10543 (2012)
41. Cambria, E., Song, Y., Wang, H., Howard, N.: Semantic Multi-Dimensional Scaling for Open-Domain Sentiment Analysis. In press: IEEE Intelligent Systems (2013)
42. Havasi, C., Speer, R., Pustejovsky, J.: Coarse word-sense disambiguation using common sense. In: Proceedings of the AAAI CSK, Arlington (2010)
43. Grassi, M.: Developing HEO human emotions ontology. Lecture Notes in Computer Science, vol. 5707, pp. 244–251. Springer, Berlin/Heidelberg (2009)
44. Cambria, E., Livingstone, A., Hussain, A.: The Hourglass of Emotions. LNCS, vol. 7403, pp. 144–157. Springer, Heidelberg (2012)