Biologically Inspired Cognitive Architectures (2013) 4, 41–53

Available at www.sciencedirect.com

SciVerse ScienceDirect

journal homepage: www.elsevier.com/locate/bica

RESEARCH ARTICLE

Application of multi-dimensional scaling and artificial neural networks for biologically inspired opinion mining

Erik Cambria [a,*], Thomas Mazzocco [b], Amir Hussain [b]

[a] *Temasek Laboratories, National University of Singapore, Singapore*
[b] *Dept. of Computing Science and Mathematics, University of Stirling, United Kingdom*

Received 23 September 2012; received in revised form 18 February 2013; accepted 18 February 2013

KEYWORDS
AI;
NLP;
ANN;
Cognitive modelling;
Sentic computing

Abstract
The way people express their opinions has radically changed in the past few years thanks to the advent of online collaborative media. The distillation of knowledge from this huge amount of unstructured information can be a key factor for marketers who want to create an identity for their product or brand in the minds of their customers. These online social data, however, remain hardly accessible to computers, as they are specifically meant for human consumption. Existing approaches to opinion mining, in fact, are still far from being able to infer the cognitive and affective information associated with natural language as they mainly rely on knowledge bases that are too limited to efficiently process text at concept-level. In this context, standard clustering techniques have been previously employed on an affective common-sense knowledge base in attempt to discover how different natural language concepts are semantically and affectively related to each other and, hence, to accordingly mine on-line opinions. In this work, a novel cognitive model based on the combined use of multi-dimensional scaling and artificial neural networks is exploited for better modelling the way multi-word expressions are organised in a brain-like universe of natural language concepts. The integration of a biologically inspired paradigm with standard principal component analysis helps to better grasp the non-linearities of the resulting vector space and, hence, improve the affective common-sense reasoning capabilities of the system.

1. Introduction

Emotions are intrinsically part of our mental activity and play a key role in decision-making and cognitive communication

* Corresponding author.
E-mail addresses: cambria@nus.edu.sg (E. Cambria), tma@cs.stir.ac.uk (T. Mazzocco), ahu@cs.stir.ac.uk (A. Hussain).

2212-683X/$ - see front matter
http://dx.doi.org/10.1016/j.bica.2013.02.003

processes. They are special states of the mind, shaped by natural selection, for adjusting various aspects of human organism in a way that it can better face particular situations, e.g., anger evolved for reaction, fear evolved for protection, and affection evolved for reproduction. Therefore, emotions cannot be shelved in the development of intelligent systems: in order for a machine to be really intelligent, it has to possess the ability to recognise, understand, and express emotions. To this end, a great number of emotion categorisation models and emotion-sensitive systems has been developed in recent years for performing tasks such as affect recognition and polarity detection.

In the context of sentic computing[1] (Cambria and Hussain, 2012), in particular, graph mining techniques and multi-dimensionality reduction techniques (Cambria et al., 2012) have been employed on a knowledge base obtained by blending ConceptNet (Speer and Havasi, 2012), a directed graph representation of common-sense knowledge, with WordNet-Affect (WNA) (Strapparava and Valitutti, 2004), a linguistic resource for the lexical representation of affect. In this work, a novel cognitive model based on the combined use of principal component analysis (PCA) and artificial neural networks (ANNs) is exploited on the same knowledge base to further improve the way multi-word expressions are organised in a brain-like universe of natural language concepts. Results demonstrate noticeable enhancements in emotion recognition from natural language text with respect to previously adopted strategies and pave the way for future development of more biologically inspired approaches to the emulation of affective common-sense reasoning.

The rest of this paper is organised as follows: the next section introduces related works in the field of opinion mining; the following one illustrates how the affective common-sense knowledge base is constructed; next, a section describes the multi-dimensional scaling techniques adopted to perform reasoning on such a knowledge base; the following section presents the emotion categorisation model used for clustering affective knowledge; then, a section describes in detail the proposed cognitive architecture and how this can be exploited for brain-inspired opinion mining; finally, the last section offers some concluding remarks and future work recommendations.

2. Related work

Existing approaches to opinion mining can be grouped into three main categories, with few exceptions: keyword spotting, lexical affinity, and statistical methods. Keyword spotting is the most naïve approach and probably also the most popular because of its accessibility and economy. Text is classified into affect categories based on the presence of fairly unambiguous affect words like 'happy', 'sad', 'afraid', and 'bored'. Elliott's Affective Reasoner (Elliott, 1992), for example, watches for 198 affect keywords, e.g., 'distressed' and 'enraged', plus affect intensity modifiers, e.g., 'extremely', 'somewhat', and 'mildly', plus a handful of cue phrases, e.g., 'did that' and 'wanted to'. Other popular sources of affect words are Ortony's Affective Lexicon (Ortony et al., 1988), which groups terms into affective categories, and Wiebe's linguistic annotation scheme (Wiebe et al., 2005). The weaknesses of this approach lie in two areas: poor recognition of affect when negation is involved and reliance on surface features. About its first weakness, while the approach can correctly classify the sentence ''today was a happy day'' as being happy, it is likely to fail on a sentence like ''today wasn't a happy day at all''. About its second weakness, the approach relies on the presence of obvious affect words which are only surface features of the prose. In practice, a lot of sentences convey affect through underlying meaning rather than affect adjectives. For example, the text ''My husband just filed for divorce and he wants to take custody of my children away from me'' certainly evokes strong emotions, but uses no affect keywords, and therefore, cannot be classified using a keyword spotting approach.

Lexical affinity is slightly more sophisticated than keyword spotting as, rather than simply detecting obvious affect words; it assigns arbitrary words a probabilistic 'affinity' for a particular emotion. For example, 'accident' might be assigned a 75% probability of being indicating a negative affect, as in 'car accident' or 'hurt by accident'. These probabilities are usually trained from linguistic corpora (Rao and Ravichandran, 2009; Somasundaran et al., 2008; Stevenson et al., 2007; Wilson et al., 2005). Though often outperforming pure keyword spotting, there are two main problems with the approach. First, lexical affinity, operating solely on the word-level, can easily be tricked by sentences like ''I avoided an accident'' (negation) and ''I met my girlfriend by accident'' (other word senses). Second, lexical affinity probabilities are often biased toward text of a particular genre, dictated by the source of the linguistic corpora. This makes it difficult to develop a reusable, domain-independent model.

Statistical methods, such as latent semantic analysis (LSA) and support vector machine (SVM), have been popular for affect classification of texts and have been used by researchers on projects such as Goertzel's Webmind (Goertzel et al., 2000), Pang's movie review classifier (Pang et al., 2002), and many others (Abbasi et al., 2008; Hu and Liu, 2004; Pang and Lee, 2005; Turney and Littman, 2003; Velikovich et al., 2010). By feeding a machine learning algorithm a large training corpus of affectively annotated texts, it is possible for the systems to not only learn the affective valence of affect keywords as in the keyword spotting approach, but such a system can also take into account the valence of other arbitrary keywords (like lexical affinity), punctuation, and word co-occurrence frequencies. However, statistical methods are generally semantically weak, meaning that, with the exception of obvious affect keywords, other lexical or co-occurrence elements in a statistical model have little predictive value individually. As a result, statistical text classifiers only work with acceptable accuracy when given a sufficiently large text input. So, while these methods may be able to affectively classify user's text on the page- or paragraph-level, they do not work well on smaller text units such as sentences.

The proposed alternative approach aims to focus on emulating the human reasoning process. The motivation is to enable machines to represent knowledge and perform reasoning in many different ways so that, whenever they reach a dead end, they can switch among different points of view and find one that may work. To bridge the cognitive and affective

[1] http://sentic.net/sentics.

gap between 'word-level' natural language data and the 'concept-level' opinions and sentiments conveyed by them, intelligent cognitive systems able to learn new affective common-sense knowledge and perform reasoning on it are needed.

3. Building the affective common-sense knowledge base

The affective common-sense knowledge base developed within this research work is built upon ConceptNet, the graph representation of the Open Mind corpus, which structurally similar to WordNet (Fellbaum, 1998), but whose scope of contents is general world knowledge, in the same vein as Cyc (Lenat and Guha, 1989). Instead of insisting on formalising common-sense reasoning using mathematical logic (Mueller, 2006), 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 (NLP). The prerogative of ConceptNet, in fact, is contextual common-sense reasoning: while WordNet is optimised for lexical categorisation 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, e.g., 'satisfy hunger' and 'follow recipe', to represent knowledge around a greater range of concepts found in everyday life. Moreover WordNet's repertoire of semantic relations is extended from the triplet of synonym, *IsA* and *PartOf*, 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, i.e., it describes something that is often true but not always, e.g., EffectOf('fall off bicycle', 'get hurt'), which is something that cannot be left aside in common-sense reasoning. 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 semi-structured English sentences of the OMCS corpus, and then applies an additional set of 'relaxation' procedures, i.e., filling in and smoothing over network gaps, to optimise the connectivity of the semantic network. The last version of ConceptNet (ConceptNet 5) contains knowledge from English Wikipedia, specifically from DBPedia, which extracts knowledge from the info-boxes that appear on articles, and ReVerb, a machine-reading project extracting relational knowledge from the actual text of each article. It also includes a large amount of content from the English Wiktionary, including synonyms, antonyms, translations of concepts into hundreds of languages, and multiple labelled word senses for many English words. ConceptNet 5 contains more dictionary-style knowledge coming from WordNet and some knowledge about people's intuitive word associations coming from games with a purpose (GWAP).

In Chinese culture (and many others), the concepts of 'heart' and 'mind' used to be expressed by the same word ('xin') as it was believed that consciousness and thoughts came from the cardiac muscle. In human cognition, in fact, 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 are intrinsically part of our mental activity and play a key role in communication and decision-making processes. Emotion is a chain of events made up of feedback loops. Feelings and behaviour can affect cognition, just as cognition can influence feeling. Emotion, cognition, and action interact in feedback loops and emotion can be viewed in a structural model tied to adaptation (Plutchik, 2001). There is actually no fundamental opposition between emotion and reason. In fact, it may be argued that reason consists of basing choices on the perspectives of emotions at some later time. Reason dictates not giving in to one's impulses because doing so may cause greater suffering later (Frijda, 1988).

Reason does not necessarily imply exertion of the voluntary capacities to suppress emotion. It does not necessarily involve depriving certain aspects of reality of their emotive powers. On the contrary, our voluntary capacities allow us to draw more of reality into the sphere of emotion. They allow one's emotions to be elicited not merely by the proximal, or the perceptual, or that which directly interferes with one's actions, but by that which, in fact, touches on one's concerns, whether proximal or distal, whether occurring now or in the future, whether interfering with one's own life or that of others. Cognitive functions serve emotions and biological needs. Information from the environment is evaluated in terms of its ability to satisfy or frustrate needs. What is particularly significant is that each new cognitive experience that is biologically important is connected with an emotional reaction such as fear, pleasure, pain, disgust, or depression (Neisser, 1967). Emotions, in fact, are special states shaped by natural selection to adjust various aspects of our organism in order to make it better face particular situations, e.g., anger evolved for reaction, fear evolved for protection, and affection evolved for reproduction. For these reasons, the development of intelligent systems cannot prescind from emotions: 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. To this end, it is useful to build a knowledge base that contains not only common-sense concepts, but also the affective information associated with these. ConceptNet is a good source of common-sense knowledge but alone is not enough for sentiment analysis tasks as it specifies how concepts are semantically related to each other but often lacks connections between concepts that convey the same kind of emotion or similar polarity.

To overcome such a hurdle, WNA, a linguistic resource for the lexical representation of affective knowledge developed starting from WordNet, is used. WNA is built by assigning to a number of WordNet synsets one or more affective labels (a-labels). In particular, the affective concepts representing emotional states are identified by synsets marked with the a-label 'emotion', but there are also other a-labels for concepts representing moods, situations eliciting emotions, or emotional responses. WNA was developed in two stages. The first consisted of the identification of a first core of affective synsets. The second step consisted of the extension of the core with the relations defined in WordNet. ConceptNet and WNA are blended together by combining the matrix representations of the two knowledge bases linearly into a single matrix, in which the information between the two initial sources is shared. The first step to create the affective blend is to transform the input data so that it can all be represented in the same matrix. To do this, the lemma forms of ConceptNet concepts are aligned with the lemma forms of the words in WNA and the most common relations in the affective knowledge base are mapped into ConceptNet's set of relations, e.g., Hypernym into *IsA* and Holonym into *PartOf*. In particular, ConceptNet is first converted into a matrix by dividing each assertion into two parts: a concept and a feature, where a feature is simply the assertion with the first or the second concept left unspecified such as 'a wheel is part of' or 'is a kind of liquid'.

The entries in the resulting matrix are positive or negative numbers, depending on the reliability of the assertions, and their magnitude increases logarithmically with the confidence score. WNA, similarly, is represented as a matrix where rows are affective concepts and columns are features related to these. The result of aligning the matrix representations of ConceptNet and WNA is a new affective semantic network, in which common-sense concepts are linked to a hierarchy of affective domain labels. In such a semantic network, termed AffectNet (http://sentic.net/affectnet.zip) (Cambria and Hussain, 2012), common-sense and affective knowledge are in fact combined, not just concomitant, i.e., everyday life concepts like 'have breakfast', 'meet people', or 'watch tv' are linked to affective domain labels like 'joy', 'anger', or 'surprise'. Such knowledge base results very useful when performing tasks such as emotion recognition or polarity detection from natural language text, as opinions and sentiments are often conveyed implicitly through context and domain dependent concepts, rather than through specific affect words.

4. Multi-dimensional scaling for affect recognition

The best way to solve a problem is to already know a solution for it. But, if we have to face a problem we have never met before, we need to use our intuition. Intuition can be explained as the process of making analogies between the current problem and the ones solved in the past to find a suitable solution. Marvin Minsky attributes this property to the so called 'difference-engines' (Minsky, 1986). This particular kind of agents operates by recognising differences between the current state and the desired state, and acting to reduce each difference by invoking K-lines that turn on suitable solution methods. This kind of thinking is maybe the essence of our supreme intelligence since in everyday life no two situations are ever the same and have to perform this action continuously. To emulate such a process, AffectiveSpace (Cambria and Hussain, 2012), a novel affective common-sense knowledge visualisation and analysis system,[2] is used.

Human mind constructs intelligible meanings by continuously compressing over vital relations (Fauconnier and Turner, 2003). The compression principles aim to transform diffuse and distended conceptual structures to more focused versions so as to become more congenial for human understanding. To this end, principal component analysis (PCA) has been applied on the matrix representation of AffectNet. In particular, truncated singular value decomposition (TSVD) has been preferred to other dimensionality reduction techniques for its simplicity, relatively low computational cost, and compactness. TSVD, in fact, is particularly suitable for measuring the cross-correlations between affective common-sense concepts as it uses an orthogonal transformation to convert the set of possibly correlated common-sense features associated with each concept into a set of values of uncorrelated variables (the principal components of the SVD). By using Lanczos' method (Lanczos, 1950), moreover, the generalisation process is relatively fast (a few seconds), despite the size and the sparseness of AffectNet. As the dimensions of such a matrix grow, however, PCA might cease to be a good solution in the future. To this end, different techniques, e.g., independent component analysis (ICA), random projections, and non-negative matrix factorisation (NMF) are being investigated.

At the present time, TSVD is applied over the concept-feature matrix in order to conveniently reduce its dimensionality and capture the most important correlations. The objective of such compression is to allow many details in the blend of ConceptNet and WNA to be removed such that the blend only consists of a few essential features that represent the global picture. Applying TSVD on AffectNet, in fact, causes it to describe other features that could apply to known affective concepts by analogy: if a concept in the matrix has no value specified for a feature owned by many similar concepts, then by analogy the concept is likely to have that feature as well. In other words, concepts and features that point in similar directions and, therefore, have high dot products, are good candidates for analogies.

A pioneering work on understanding and visualising the affective information associated to natural language text was conducted by Osgood et al. (1975). Osgood used multi-dimensional scaling (MDS) to create visualisations of affective words based on similarity ratings of the words provided to subjects from different cultures. Words can be thought of as points in a multi-dimensional space and the similarity ratings represent the distances between these words. MDS projects these distances to points in a smaller dimensional space (usually two or three dimensions). Similarly, AffectiveSpace aims to grasp the semantic and affective similarity between different concepts by plotting them into a multi-dimensional vector space.

[2] http://sentic.net/affectivespace.zip.

Differently from Osgood's space, however, the building blocks of AffectiveSpace are not simply a limited set of similarity ratings between affect words, but rather millions of confidence scores related to pieces of common-sense knowledge linked to a hierarchy of affective domain labels. Rather than merely determined by a few human annotators and represented as a word–word matrix, in fact, AffectiveSpace is built upon an affective common-sense knowledge base, namely AffectNet, represented as a concept-feature matrix. After performing TSVD on such matrix, hereby termed A for the sake of conciseness, a low-rank approximation of it is obtained, that is, a new matrix $\widetilde{A} = U_k\,\Sigma_k\,V_k^T$. This approximation is based on minimising the Frobenius norm of the difference between A and $\widetilde{A}$ under the constraint $rank(\widetilde{A}) = k$. For the Eckart–Young theorem (Eckart and Young, 1936), it represents the best approximation of A in the least-square sense, in fact:

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

assuming that $\widetilde{A}$ has the form $\widetilde{A} = USV^*$, 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:

$$\min_{\widetilde{A}|rank(\widetilde{A})=k} |\Sigma - S| = \min_{s_i}\sqrt{\sum_{i=1}^{n}(\sigma_i - s_i)^2} \quad (2)$$

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

$$= \sqrt{\sum_{i=k+1}^{n}\sigma_i^2} \quad (3)$$

Therefore, $\widetilde{A}$ 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 all but the first k principal components are discarded, 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, i.e., the basis $e_0, \ldots, e_{k-1}$ of the vector space (Fig. 1). 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, concepts conveying the same emotion 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 concepts such as 'beautiful day', 'birthday party', and 'make person happy' are found very close in direction in the vector space, while concepts like 'feel guilty', 'be laid off', and 'shed tear' are found in a completely different direction (nearly opposite with respect to the centre of the space). The key to perform common-sense reasoning is to find a good trade-off for representing knowledge. Since in life two situations are never exactly the same, no representation should be too concrete, or it will not apply to new situations, but, at the same time, no representation should be too abstract, or it will suppress too many details. ConceptNet already supports different representations, in fact, it maintains different ways of conveying the same idea with redundant concepts, e.g., 'car' and 'automobile', that can be reconciled through background linguistic knowledge, if necessary. Within AffectiveSpace, this knowledge representation trade-off can be seen in the choice of the vector space dimensionality. The number k of singular values selected to build AffectiveSpace, in fact, is a measure of the trade-off between precision and efficiency in the representation of the affective common-sense knowledge base. The bigger is k, the more precisely AffectiveSpace represents AffectNet's knowledge, but generating the vector space is slower, and so is computing dot products between concepts.

The smaller is k, on the other hand, the more efficiently AffectiveSpace represents affective common-sense knowledge both in terms of vector space generation and of dot product computation. However, too few dimensions risk not to correctly represent AffectNet as concepts defined with too few features tend to be too close to each other in the vector space and, hence, not easily distinguishable and clusterable. In order to find a good k, AffectiveSpace was tested on a benchmark for affective common-sense knowledge (BACK) (Cambria and Hussain, 2012) built by applying CF-IOF (concept frequency - inverse opinion frequency) (Cambria et al., 2010) on the 5000 posts of the LiveJournal corpus.[3] CF-IOF is a technique that identifies common domain-dependent semantics 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:

$$CF - IOF_{c,d} = \frac{n_{c,d}}{\sum_k n_{k,d}} \log \sum_k \frac{n_k}{n_c} \quad (4)$$

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. Specifically, CF-IOF weighting was exploited to filter out common concepts in the LiveJournal corpus and to detect relevant mood-dependent semantics for the set of 24 emotions defined by Plutchik Plutchik (2001). The result was a benchmark of 2000 affective concepts that were screened by 21 English-speaking students who were asked to map each concept to the 24 different emotional categories, which form the Hourglass of Emotions (Cambria et al.,

[3] http://livejournal.com.

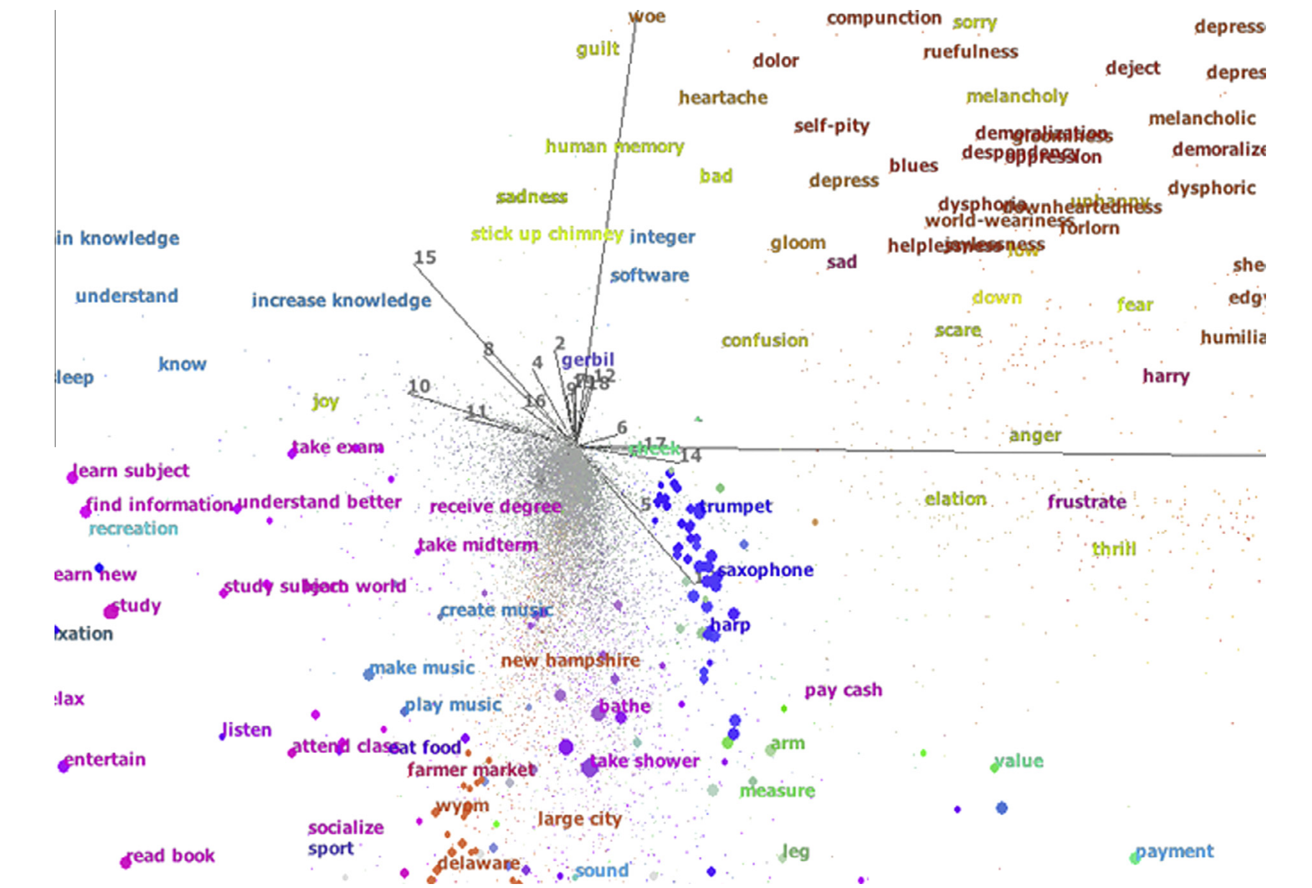

Fig. 1 A sketch of AffectiveSpace. Affectively positive concepts (in the bottom-left corner) and affectively negative concepts (in the up-right corner) are floating in the multi-dimensional vector space.

Table 1 Distribution of concepts through the Pleasantness dimension. The affective information associated with most concepts concentrates around the centre of the Hourglass, rather than its extremes.

Level	Label	Frequency (%)
$-$G(1)	Grief	14.3
$-$G(2/3)	Sadness	19.8
$-$G(1/3)	Pensiveness	11.4
0	Neutral	10.5
+G(1/3)	Serenity	20.6
+G(2/3)	Joy	18.3
+G(1)	Ecstasy	5.1

2012) (explained in next section). Results obtained were averaged (Table 1).

BACK's concepts were compared with the classification results obtained by applying the AffectiveSpace process using different values of k, from 1 to 250. The best trade-off is achieved at 100, as selecting more than 100 singular values does not improve accuracy significantly. The distribution of the values of each AffectiveSpace dimension is bell-shaped, with different centres and different degree of dispersion around them. Affective common-sense concepts, in fact, tend to be close to the origin of the vector space. In order to more uniformly distribute concept density in AffectiveSpace, an alternative strategy to represent the vector space was investigated. Such strategy consists in centring the values of the distribution of each dimension on the origin and in mapping dimensions according to a transformation $x \in \mathbb{R} \mapsto x^* \in [-1, 1]$. This transformation is often pivotal for better clustering AffectiveSpace as the vector space tends to have different grades of dispersion of data points across different dimensions, with some space regions more densely populated than others.

The switch to a different space configuration helps to distribute data more uniformly, possibly leading to an improved (or, at least, different) reasoning process. In particular, the transformation $x_{ij} \mapsto x_{ij} - \mu_i$ is first applied, being μ_i the average of all values of the ith dimension. Then a normalisation is applied, combining the previous transformation with a new one $x_{ij} \mapsto \frac{x_{ij}}{a \cdot \sigma_i}$, where σ_i is the standard deviation calculated on the ith dimension and a is a coefficient that can modify the same proportion of data that is represented within a specified interval.

Finally, in order to ensure that all components of the vectors in the defined space are within $[-1, 1]$ (i.e., that the Chebyshev distance between the origin and each vector is smaller or equal to 1), a final transformation $x_{ij} \mapsto s(x_{ij})$ is needed, where $s(x)$ is a sigmoid function. Different choices for the sigmoid function may be made, influencing how 'fast' the function approaches the unit value while the independent variable approaches infinity. Combining the proposed transformations, two possible mapping functions are expressed in the following formulae (5) and (6):

$$x^*_{ij} = tanh\left(\frac{x_{ij} - \mu_i}{a \cdot \sigma_i}\right) \quad (5)$$

$$x^*_{ij} = \frac{x_{ij} - \mu_i}{a \cdot \sigma_i + |x_{ij} - \mu_i|} \quad (6)$$

This space transformation leads to two main advantages, which could be of notable importance depending on the problem being tackled. Firstly, this different space configuration ensures that each dimension is equally important by avoiding that the information provided by dimensions with higher (i.e., more distant from the origin) averages predominates. Secondly, normalising according to the standard deviations of each dimension allows a more uniform distribution of data around the origin, leading to a full use of information potential.

5. Emotion categorisation model

In order to accordingly organise and interpret AffectiveSpace, an affective categorisation model is needed. The Hourglass of Emotions (Cambria et al., 2012), a model inspired by Plutchik's studies on human emotions (Plutchik, 2001), was selected. It reinterprets Plutchik's model by organising primary emotions around four independent but concomitant dimensions, whose different levels of activation make up 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 Minsky (2006). This way, the model can potentially synthesise the full range of emotional experiences in terms of Pleasantness, Attention, Sensitivity, and Aptitude, as the different combined values of the four affective dimensions can also model affective states we do not have a specific name for, due to the ambiguity of natural language and the elusive nature of emotions.

The main motivation for the design of the model is the concept-level inference of the cognitive and affective information associated with text. Such faceted information is needed, within sentic computing, for a feature-based sentiment analysis, where the affective common-sense knowledge associated with natural language opinions has to be objectively assessed. Therefore, the Hourglass model systematically excludes what are variously known as self-conscious or moral emotions, e.g., pride, guilt, shame, embarrassment, moral outrage, or humiliation (Lazarus, 1991; Lewis, 2000; Scherer et al., 2001; Tracy et al., 2007). Such emotions, in fact, present a blind spot for models rooted in basic emotions, because they are by definition contingent on subjective moral standards. The distinction between guilt and shame, for example, is based in the attribution of negativity to the self or to the act. So, guilt arises when believing to have done a bad thing, and shame arises when thinking to be a bad person. This matters because, in turn, these emotions have been shown to have different consequences in terms of action tendencies. Likewise, an emotion such as *schadenfreude* is essentially a form of pleasure, but it is crucially different from pride or happiness because of the object of the emotion (the misfortune of another that is not caused by the self), and the resulting action tendency (do not express). However, since the Hourglass model currently focuses on the objective inference of affective information associated with natural language opinions, appraisal-based emotions are not taken into account within the present version of the model.

The Hourglass model, in fact, is a biologically-inspired and psychologically-motivated model based on the idea that emotional states result from the selective activation/disactivation of different resources in the brain. 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. Evidence of this theory is also given by several fMRI experiments showing that there is a distinct pattern of brain activity that occurs when people are experiencing different emotions. Zeki and Romaya, for example, investigated the neural correlates of hate with an fMRI procedure (Zeki and Romaya, 2008). In their experiment, people had their brains scanned while viewing pictures of people they hated. The results showed increased activity in the medial frontal gyrus, right putamen, bilaterally in the premotor cortex, in the frontal pole, and bilaterally in the medial insula of the human brain. Also the activity of emotionally enhanced memory retention can be linked to human evolution (Cahill and McGaugh, 1995). During early development, in fact, responsive behaviour to environmental events is likely to have progressed as a process of trial-and-error.

Survival depended on behavioural patterns that were repeated or reinforced through life and death situations. Through evolution, this process of learning became genetically embedded in humans and all animal species in what is known as 'fight or flight' instinct (Bradford Cannon, 1915). The primary quantity we can measure about an emotion we feel is its strength. But, when we feel a strong emotion, it is because we feel a very specific emotion. And, conversely, we cannot feel a specific emotion like fear or amazement without that emotion being reasonably strong. For such reasons, the transition between different emotional states is modelled, within the same affective dimension, using 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 towards the unit value.

In particular, the function models how the level of activation of each affective dimension varies from the state of 'emotional void' (null value) to the state of 'heightened emotionality' (unit value). Justification for assuming that the Gaussian function (rather than a step or simple linear function) is appropriate for modelling the variation of emotion intensity is based on research into the neural and behavioural correlates of emotion, which are assumed to indicate emotional intensity in some sense. In fact, nobody genuinely knows what function subjective emotion intensity follows, because it has never been truly or directly measured (Barrett, 2006). For example, the so-called Duchenne smile (a genuine smile indicating pleasure) is characterised by smooth onset, increasing to an apex, and a smooth, relatively lengthy offset (Krumhuber and Kappas, 2005). More generally, Klaus Scherer has argued that emotion is a process characterised by non-linear relations among its component elements - especially physiological measures, which typically look Gaussian (Lewis and Granic, 2002). Emotions, in fact, are not linear (Plutchik, 2001): the stronger the emotion, the easier it is to be aware of it. Mapping this space of possible emotions leads to a hourglass shape (Fig. 2).

It is worth to note that, in the model, the state of 'emotional void' is a-dimensional, which contributes to determine the hourglass shape. Total absence of emotion, in fact, can be associated with the total absence of reasoning (or, at least, consciousness) (Csikszentmihalyi, 1991), which is not an envisaged mental state as, in human mind, there is never nothing going on.

The Hourglass of Emotions, in particular, can be exploited in the context of HCI to measure how much respectively: the user is amused by interaction modalities (Pleasantness), the user is interested in interaction contents (Attention), the user is comfortable with interaction dynamics (Sensitivity), the user is confident in interaction benefits (Aptitude). Each affective dimension, in particular, is characterised by six levels of activation (measuring the strength of an emotion), termed 'sentic levels', which represent the intensity thresholds of the expressed or perceived emotion. These levels are also labelled as a set of 24 basic emotions (Plutchik, 2001), 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 (Table 2). The dimensional form, in

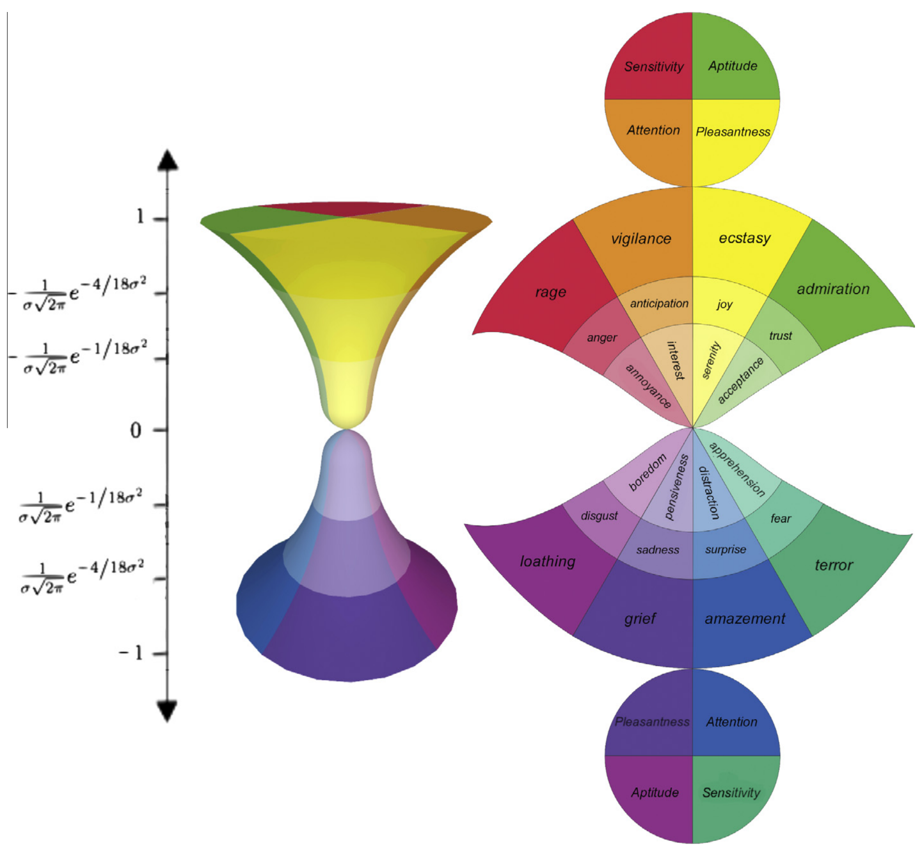

Fig. 2 The 3D model and the net of the Hourglass of Emotions. Since affective states go from strongly positive to null to strongly negative, the model assumes a hourglass shape.

Table 2 The sentic levels of the Hourglass model. Labels are organised into four affective dimensions with six different levels each, whose combined activity constitutes the 'total state' of the mind.

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

particular, is termed 'sentic vector' and it is a four-dimensional *float* vector that can potentially synthesise the full range of emotional experiences in terms of Pleasantness, Attention, Sensitivity, and Aptitude.

In the model, the vertical dimension represents the intensity of the different affective dimensions, i.e., their level of activation, while the radial dimension represents K-lines (Minsky, 1986) that can activate configurations of the mind, which can either last just a few seconds or years. The model follows the pattern used in colour theory and research in order to obtain judgements about combinations, i.e., the emotions that result when two or more fundamental emotions are combined, in the same way that red and blue make purple. Hence, 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 $\in [G(2/3), G(1/3))$, i.e., joy, a level of Aptitude $\in [G(2/3), G(1/3))$, i.e., trust, and a minor magnitude of Attention and Sensitivity, are termed 'love sentic vectors' since they specify the compound emotion of love.

More complex emotions can be synthesised by using three, or even four, sentic levels, e.g., joy + trust + anger = jealousy. Therefore, analogous to the way primary colours combine to generate different colour gradations (and even colours we do not have a name for), the primary emotions of the Hourglass model can blend to form the full spectrum of human emotional experience. Beyond emotion detection, the Hourglass model is also used for polarity detection tasks. Since polarity is strongly connected to attitudes and feelings, in fact, it is defined in terms of the four affective dimensions, according to the formula:

$$p = \sum_{i=1}^{N} \frac{Pleasantness(c_i) + |Attention(c_i)| - |Sensitivity(c_i)| + Aptitude(c_i)}{3N} \tag{7}$$

where c_i is an input concept, N the total number of concepts, and 3 the normalisation factor (as the Hourglass dimensions are defined as float $\in [-1, +1]$). In the formula, Attention is taken as absolute value since both its positive and negative intensity values correspond to positive polarity values (e.g., 'surprise' is negative in the sense of lack of Attention, but positive from a polarity point of view). Similarly, Sensitivity is taken as negative absolute value since both its positive and negative intensity values correspond to negative polarity values (e.g., 'anger' is positive in the sense of level of activation of Sensitivity, but negative in terms of polarity). The formula can be seen as one of the first attempts to show a clear connection between emotion recognition (sentiment analysis) and polarity detection (opinion mining).

6. Bio-inspired opinion mining engine

The proposed architecture extends a framework previously proposed by the authors (Mazzocco et al., 2012) and investigates if an emulation of the biological neural system, represented by two ANNs, could outperform the state-of-the-art k-medoids clustering approach (Cambria et al., 2011). Similarly to previous works (Cambria et al., 2012; Havasi et al., 2009; Cambria et al., 2010), the proposed architecture uses PCA to organise the space where concepts lie but, rather than using standard clustering techniques, e.g., k-NN or k-medoids, for reasoning on how such concepts are semantically related to each other, it exploits a human-inspired cognitive architecture paradigm (Samsonovich, 2010) to better deal with non-linearities of the resulting space and, hence, more accurately infer the semantic and affective information associated with common-sense concepts, based on the presented Hourglass model. Contrary to any clustering algorithm, the emotion recognition task is independent from both concepts' absolute and relative positions in the vector space. The eventual aim of the proposed ANNs developed in this study is to predict which class each concept belongs to (i.e., its level of affective valence in a specific dimension of the Hourglass model). Two different approaches may be adopted in order to set up an ANN: a 'discrete' neural network (DNN) and a 'continuous' neural network (CNN). DNN, in particular, is expected to return seven different real-valued outputs $y_k \in [0, 1]$ for $k = 1, 2, \ldots, 7$, each showing the degree of belonging to a specified affective level, while CNN provides a single real-valued output $y \in [-1, 1]$, corresponding to the best guess of the level of affective valence (assuming that categories are equispaced within the considered dimension).

In both cases, a further step is required in order to obtain a final classification output: for DNN the best selection strategy seems to be the choice of the class with the highest degree of belonging, while for CNN the easiest approach is to round off the output to get an integer corresponding to the class. Since the task of choosing from these two approaches is not easily solvable a priori, both approaches are adopted and compared in this study. Therefore, two multi-layer perceptron neural networks with three layers (one input layer, one hidden layer and one output layer) were set up. The input vector $\mathbf{x}_{(k)}$ is built so that $x_{(k)0} = 1$ (the 'bias node') and $[x_{(k)1}, x_{(k)2}, \ldots, x_{(k)100}] = \mathbf{a}_{(k)}$ for the kth concept of the dataset. The target output is, for DNN, a vector $\mathbf{y}_{(k)}$ having $y_{(k)i} = 1$ if $b_{(k)} = i - 4$ and $y_{(k)i} = 0$ otherwise, for each $i = 1, 2, \ldots, 7$. On the other hand, for CNN the target output is a single value $y_{(k)} = b_{(k)}$.

Let us assume that the hidden layer has H neurons; the input and the hidden layers are then linked by Eq. (8) where $w_{mn} \in \mathbf{W}$, that is the matrix of weights (of dimensions $H \times 101$) defined during the network training:

$$h_j = tanh\left(\sum_{i=0}^{100} w_{ji} x_i\right) \quad j = 1, 2, \ldots, H \tag{8}$$

The relationships between the hidden and the output layers are expressed in Eqs. (9) and (10) for DNN and CNN respectively, where $v_{mn} \in \mathbf{V}$, which is a matrix of weights (of dimensions $7 \times H$) and $\mathbf{v}$ a vector of weights (of dimension H) built during the network training. The output is computed as follows:

$$y_j = tanh\left(\sum_{i=1}^{H} v_{ji} h_i\right) \quad j = 1, 2, \ldots, 7 \tag{9}$$

$$y = tanh(\mathbf{v}^{\mathsf{T}} \cdot \mathbf{h}) \tag{10}$$

It is worth noting that different choices of the activation functions, as well as other design choices, are possible; however, a definition of the best structure of the used neural networks besides being quite difficult to state (due to, for example, the dataset dependency) is beyond the scope of this study.

Finally, further transformations of $\mathbf{y} \in \mathbb{R}^7 \mapsto b^* \in \Theta$ and $y \in \mathbb{R} \mapsto b^* \in \Theta$ are required. They are proposed in Eqs. (11) and (12):

$$b^* = \mathbf{d}^\mathrm{T} \cdot \left[-1, -\frac{1}{\sigma\sqrt{2\pi}} e^{-4}/18\sigma^2, -\frac{1}{\sigma\sqrt{2\pi}} e^{-1}/18\sigma^2, 0, \frac{1}{\sigma\sqrt{2\pi}} e^{-1}/18\sigma^2, \frac{1}{\sigma\sqrt{2\pi}} e^{-4}/18\sigma^2, +1\right] \quad (11)$$

$$b^* = round(y) \quad (12)$$

where $d_i = \delta_{i,m}$ for $i = 1, 2, \ldots, 7$, δ being the Kronecker's delta, $m = i \| y_i = \max_j y_j$ and $round(x)$ a function $\mathbb{R} \mapsto \mathbb{Z}$ rounding x up to the closest integer. The neural networks training is conducted using the gradient descent optimisation algorithm with the inclusion of a momentum term, which has been proven to improve the algorithm speed of convergence (Qian, 1999).

$$\Delta\mathbf{U}_{(\tau)} = -\nu\nabla\mathbf{E}(\mathbf{U}_{(\tau)}) + \mu\Delta\mathbf{U}_{(\tau-1)} \quad (13)$$

The matrices (or vectors) of weights (**U**) are then updated at each iteration τ according to Eq. (13), where **E** is calculated with $e_j = b_j - b_j^*$ for the second layer of weights (where $j = 1, 2, \ldots, 7$ for DNN and $j = 1$ for CNN), and $e_j = \sum_k \frac{b^*_{j(\tau)}}{b^*_{j(\tau-1)}} u_{jk}$ for the first layer of weights, respectively. In order to evaluate the designed system, the ANNs were tested on the benchmark for affective common-sense knowledge (BACK) (Cambria and Hussain, 2012). To avoid the risk of overfitting, a cross-validation approach was adopted. The networks were trained 10 times (10-fold cross-validation), each of which excluded 10% of dataset entries that are used for evaluating the performance of the system; the excluded 10% is then cycled so that, at the end of all simulations, each dataset entry has been used exactly once to test the system. In order to evaluate the accuracy of the model, the percentage of entries where $b^* = b$ ('strict accuracy') is considered. However, since the used dataset can include noise and entries may incorporate a certain degree of subjectiveness, this criterion was relaxed by considering the accuracy of entries which have $|b^* - b| \leqslant 1$ ('relaxed accuracy').

The performance of the proposed ANNs are tabulated in Table 3 where they are compared with the state-of-the-art k-medoids approach, k-nearest neighbour (k-NN), and a random classifier. A trial-and-error approach was adopted for the network parameters tuning: for CNN the best performance was obtained after three iterations (when the error stopped decreasing significantly) with $H = 10$, learning rate $\nu = 0.1$, momentum factor $\mu = 0.2$; for DNN the best set of parameters obtained were $H = 15$, $\nu = 0.05$, $\mu = 0.05$ with best performance reached after an average of 10 iterations. As it can be seen from Table 3, the proposed ANN approaches outperform the state-of-the-art k-medoids model, as well as the k-NN model and the random classifier.

Table 3 Performance comparison. Bold values represent best values for each specific category.

	Strict acc. (%)	Relaxed acc. (%)
Random	14.3	40.1
k-NN	41.9	72.3
k-medoids	43.2	74.1
DNN	**46.9**	76.5
CNN	39.7	**84.3**

Both proposed models improved the 'relaxed accuracy' (with the CNN producing a considerable 10% performance improvement) while the DNN was able to outperform the benchmark for the 'strict accuracy' case. In order to test the performance of the proposed approach in a more practical environment, the ANNs were also embedded into an opinion mining engine (Cambria and Hussain, 2012) for the inference of the cognitive and affective information associated with natural language. Such an engine consists of four main components: a pre-processing module, which performs a first skim of text; a semantic parser, whose aim is to extract concepts from the opinionated text; a target spotting module, which identifies opinion targets; an affect interpreter, for emotion recognition and polarity detection. The pre-processing module firstly interprets all the affective valence indicators usually contained in opinionated text such as special punctuation, complete upper-case words, cross-linguistic onomatopoeias, exclamation words, negations, degree adverbs and emoticons. Secondly, it converts text to lower-case and, after lemmatizing it, splits the opinion into single clauses according to grammatical conjunctions and punctuation. Then, the semantic parser deconstructs text into concepts using a lexicon based on sequences of lexemes that represent multiple-word concepts extracted from ConceptNet, WordNet and other linguistic resources. 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, e.g., the concept 'buy christmas present' in the sentence ''I bought a lot of very nice Christmas presents''. The semantic parser, additionally, provides, for each retrieved concept, the relative frequency, valence and status, that is the concept's occurrence in the text, its positive or negative connotation and the degree of intensity with which the concept is expressed.

For each clause, the module outputs a small bag of concepts (SBoC), which is later on analyzed separately by the target spotting module and the affect interpreter 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 (Havasi et al., 2010). Then, to assign the correct sense to the ambiguous concept, the sense of each concept that has the highest dot product (and thus the strongest similarity) with the clause vector has to be seeked.

Table 4 Structured output example of opinion mining engine.

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

Table 5 F-measure values relative to PatientOpinion evaluation. Bold values represent best values for each specific category.

	k-NN (%)	k-Medoids (%)	DNN (%)	CNN (%)
Clinical service	70.1	72.9	77.2	**81.8**
Communication	69.8	75.3	75.5	**78.1**
Food	79.4	79.6	**83.1**	79.7
Parking	71.0	72.5	74.0	**76.7**
Staff	76.1	76.1	**82.1**	77.9
Timeliness	72.3	73.0	77.4	**79.6**

The target spotting module aims to individuate one or more opinion targets, such as people, places, events and ideas, from the input concepts. This is done by projecting the concepts of each SBoC into the graph representation of AffectNet, in order to assign these to a specific conceptual class. The categorisation does not consist in simply labelling each concept but also in assigning a confidence score to each category label, which is directly proportional to the value of belonging to a specific conceptual cluster (number of steps in the AffectNet graph). The affect interpreter, in turn, projects the concepts of each SBoC into AffectiveSpace and feeds their coordinates to both DNN and CNN, in order to assign such concepts to a specific affective class, and hence, calculate polarity in terms of the Hourglass dimensions, as specified in formula (7).

As an example of how the opinion mining engine works, intermediate and final outputs obtained when a natural language opinion is given as input to the system can be examined. The tweet ''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!'' is selected. After the pre-processing 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 target spotting module and the affect interpreter, which detect the opinion targets and output, for each of them, the relative affective information both in a discrete way, with one or more emotional labels, and in a dimensional way, with a polarity value $\in [-1, +1]$ (as shown in Table 4). In order to evaluate the resulting opinion mining engine, a patient opinion database (Cambria and Hussain, 2012) is used, and results obtained using k-NN and k-medoids are compared with those obtained using the ANNs. The resource is a dataset obtained from PatientOpinion,[4] a social enterprise pioneering an online feedback service for users of the UK national health service to enable people to share their recent experience of local health services online. It is a manually tagged dataset of 2000 patient opinions that associates to each post a category (namely, clinical service, communication, food, parking, staff, and timeliness) and a positive or negative polarity. There are no ethical issues involved in the data used in the experimentation as tweets, blogposts, and patient opinions were all anonymised. In order to guarantee full anonymity, moreover, the text associated with tweets, blogposts, and patient opinions has never been wholly reported in the proposed tables and examples. The dataset is hereby used to test the combined detection of opinion targets and the polarity associated with these. Results show that DNN and CNN generally outperform k-medoids and k-NN. In particular, DNN achieves better accuracy for categories in which affect is usually conveyed explicitly, e.g., 'staff' and 'food', while CNN turns out to be a better choice when sentiment is expressed in a more subtle manner (Table 5).

7. Conclusions and future work

With the advent of Web 2.0, the extraction of opinions and sentiments from the huge amount of available unstructured information derived from blog, wikis, and social networks is a very arduous task. While existing approaches to opinion mining mainly work at a syntactic-level, computational techniques and tools were hereby employed to analyze text

[4] http://patientopinion.org.uk.

natural language at a semantic-level. In particular, we developed a bio-inspired opinion mining engine that, first, deconstructs natural language text into concepts, then, encodes such concepts as coordinates of a multi-dimensional vector space, and finally infers the semantic and affective information associated with them by means of two ANNs. We also demonstrated how such a human-like reasoning framework outperforms state-of-the-art clustering techniques for opinion mining.

The integration of multi-dimensional scaling and ANNs, in fact, has embedded a bio-inspired way of reasoning to carry out cognitive tasks such as emotion recognition and polarity detection. Such an ensemble model better grasps the nonlinearities of the vector space of affective common-sense knowledge and, hence, improves the performance of the opinion mining engine.

Since this study has shown promising results, further research is now planned to understand how artificial intelligence techniques can affectively analyze natural language text: structured data collection is planned to extend the BACK database, which will be made publicly available to enable comparison with other reasoning models. This extended dataset will be exploited to assess how other biologically inspired frameworks, e.g., extreme learning machines (Decherchi et al., 2013), could further improve the way multi-word expressions are organised in a brain-like universe of natural language concepts and, hence, refine the sensemaking process of the affective common-sense reasoning model.

References

Abbasi, A., Chen, H., & Salem, A. (2008). Sentiment analysis in multiple languages: Feature selection for opinion classification in web forums. *ACM Transactions on Information Systems, 26*(3), 1–34.

Barrett, L. (2006). Solving the emotion paradox: Categorization and the experience of emotion. *Personality and Social Psychology Review, 10*(1), 20–46.

Bradford Cannon, W. (1915). *Bodily changes in pain, hunger, fear and rage: An account of recent researches into the function of emotional excitement*. Appleton Century Crofts.

Cahill, L., & McGaugh, J. (1995). A novel demonstration of enhanced memory associated with emotional arousal. *Consciousness and Cognition, 4*(4), 410–421.

Cambria, E., Hussain, A., Durrani, T., Havasi, C., Eckl, C., & Munro, J. (2010). *Sentic computing for patient centered application*. Beijing.

Cambria, E., Havasi, C., & Hussain, A. (2012). *SenticNet 2: A semantic and affective resource for opinion mining and sentiment analysis*. Marco Island.

Cambria, E., Olsher, D., & Kwok, K. (2012). *Sentic activation: A two-level affective common sense reasoning framework*. Toronto.

Cambria, E., & Hussain, A. (2012). *Sentic computing: Techniques, tools, and applications*. Dordrecht, Netherlands: Springer.

Cambria, E., Livingstone, A., & Hussain, A. (2012). The hourglass of emotions. In A. Esposito, A. Vinciarelli, R. Hoffmann, & V. Muller (Eds.). *Cognitive behavioral systems* (Vol. 7403, pp. 144–157). Berlin Heidelberg: Springer.

Cambria, E., Mazzocco, T., Hussain, A., & Eckl, C. (2011). Sentic medoids: Organizing affective common sense knowledge in a multi-dimensional vector space. In D. Liu, H. Zhang, M. Polycarpou, C. Alippi, & H. He (Eds.). *Advances in neural networks* (Vol. 6677, pp. 601–610). Berlin: Springer.

Csikszentmihalyi, M. (1991). *Flow: The psychology of optimal experience*. Harper Perennial.

Decherchi, S., Gastaldo, P., Zunino, R., Cambria, E., & Redi, J. (2013). Circular-ELM for the reduced-reference assessment of perceived image quality. *Neurocomputing, 102*, 78–89.

Eckart, C., & Young, G. (1936). The approximation of one matrix by another of lower rank. *Psychometrika, 1*(3), 211–218.

Elliott, C.D. (1992). *The affective reasoner: A process model of emotions in a multi-agent system*. Unpublished doctoral dissertation, Northwestern University, Evanston.

Fauconnier, G., & Turner, M. (2003). *The way we think: Conceptual blending and the mind's hidden complexities*. Basic Books.

Fellbaum, C. (1998). *WordNet: An electronic lexical database (language, speech, and communication)*. The MIT Press.

Frijda, N. (1988). The laws of emotions. *American Psychologist, 43*(5).

Goertzel, B., Silverman, K., Hartley, C., Bugaj, S., & Ross, M. (2000). *The Baby Webmind project*. Birmingham.

Havasi, C., Speer, R., & Pustejovsky, J. (2010). Coarse word-sense disambiguation using common sense. In *AAAI CSK*. Arlington.

Havasi, C., Speer, R., Pustejovsky, J., & Lieberman, H. (2009). Digital intuition: Applying common sense using dimensionality reduction. *IEEE Intelligent Systems, 24*(4), 24–35.

Hu, M., & Liu, B. (2004). Mining and summarizing customer reviews. In *Kdd*. Seattle.

Krumhuber, E., & Kappas, A. (2005). Moving smiles: The role of dynamic components for the perception of the genuineness of smiles. *Journal of Nonverbal Behavior, 29*(1), 3–24.

Lanczos, C. (1950). An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *Journal of Research of The National Bureau of Standards, 45*(4), 255–282.

Lazarus, R. (1991). *Emotion and adaptation*. New York: Oxford University Press.

Lenat, D., & Guha, R. (1989). *Building large knowledge-based systems: Representation and inference in the cyc project*. Boston: Addison-Wesley.

Lewis, M. (2000). Self-conscious emotions: Embarrassment, pride, shame, and guilt. *Handbook of cognition and emotion* (Vol. 2, pp. 623–636). Guilford Press.

Lewis, M., & Granic, I. (2002). *Emotion, development, and self-organization: Dynamic systems approaches to emotional development*. Cambridge University Press.

Mazzocco, T., Cambria, E., Hussain, A., & Wang, Q. (2012). Sentic neural networks: A novel cognitive model for affective common sense reasoning. In H. Zhang, A. Hussain, D. Liu, & Z. Wang (Eds.). *Advances in brain inspired cognitive systems* (Vol. 7366, pp. 12–21). Heidelberg: Springer.

Minsky, M. (1986). *The society of mind*. New York: Simon and Schuster.

Minsky, M. (2006). *The emotion machine: Commonsense thinking, artificial intelligence, and the future of the human mind*. New York: Simon & Schuster.

Mueller, E. (2006). *Commonsense reasoning*. Morgan Kaufmann.

Neisser, U. (1967). *Cognitive psychology*. Appleton Century Crofts.

Ortony, A., Clore, G., & Collins, A. (1988). *The cognitive structure of emotions*. Cambridge: Cambridge University Press.

Osgood, C., May, W., & Miron, M. (1975). *Cross-cultural universals of affective meaning*. Univ. of Illinois Press.

Pang, B., & Lee, L. (2005). Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Acl* (pp. 115–124). Ann Arbor.

Pang, B., Lee, L., & Vaithyanathan, S. (2002). Thumbs up? Sentiment classification using machine learning techniques. In *EMNLP* (pp. 79–86). Philadelphia.

Plutchik, R. (2001). The nature of emotions. *American Scientist, 89*(4), 344—350.

Qian, N. (1999). On the momentum term in gradient descent learning algorithms. *Neural Networks, 12*, 145—151.

Rao, D., & Ravichandran, D. (2009). Semi-supervised polarity lexicon induction. In *Eacl* (pp. 675—682). Athens.

Samsonovich, A. (2010). *A human-inspired cognitive architecture supporting self-regulated learning in problem solving*. Atlanta.

Scherer, K., Shorr, A., & Johnstone, T. (2001). *Appraisal processes in emotion: Theory, methods, research*. Canary: Oxford University Press.

Somasundaran, S., Wiebe, J., & Ruppenhofer, J. (2008). Discourse level opinion interpretation. In *Coling*. Manchester.

Speer, R., & Havasi, C. (2012). ConceptNet 5: A large semantic network for relational knowledge. In E. Hovy, M. Johnson, & G. Hirst (Eds.), *Theory and applications of natural language processing*. Springer.

Stevenson, R., Mikels, J., & James, T. (2007). Characterization of the affective norms for english words by discrete emotional categories. *Behavior Research Methods, 39*, 1020—1024.

Strapparava, C., & Valitutti, A. (2004). WordNet-Affect: An affective extension of WordNet. In *LREC*. Lisbon.

Tracy, J., Robins, R., & Tangney, J. (2007). *The self-conscious emotions: Theory and research*. The Guilford Press.

Turney, P., & Littman, M. (2003). Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems, 21*(4), 315—346.

Velikovich, L., Goldensohn, S., Hannan, K., & McDonald, R. (2010). The viability of web-derived polarity lexicons. In *Naacl* (pp. 777—785). Los Angeles.

Wiebe, J., Wilson, T., & Cardie, C. (2005). Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation, 39*(2), 165—210.

Wilson, T., Wiebe, J., & Hoffmann, P. (2005). Recognizing contextual polarity in phrase-level sentiment analysis. In *Hlt/emnlp*. Vancouver.

Zeki, S., & Romaya, J. (2008). Neural correlates of hate. *PloS one, 3*(10), 35—56.