

Sentic Activation: A Two-Level Affective Common Sense Reasoning Framework

Erik Cambria

National University of Singapore
Singapore, 117411 Singapore
cambria@nus.edu.sg

Daniel Olsher

National University of Singapore
Singapore, 117411 Singapore
olsher@nus.edu.sg

Kenneth Kwok

National University of Singapore
Singapore, 117411 Singapore
kenkwok@nus.edu.sg

Abstract

An important difference between traditional AI systems and human intelligence is our ability to harness common sense knowledge gleaned from a lifetime of learning and experiences to inform our decision making and behavior. This allows humans to adapt easily to novel situations where AI fails catastrophically for lack of situation-specific rules and generalization capabilities. Common sense knowledge also provides the background knowledge for humans to successfully operate in social situations where such knowledge is typically assumed. In order for machines to exploit common sense knowledge in reasoning as humans do, moreover, we need to endow them with human-like reasoning strategies. In this work, we propose a two-level affective reasoning framework that concurrently employs multi-dimensionality reduction and graph mining techniques to mimic the integration of conscious and unconscious reasoning, and exploit it for sentiment analysis.

Introduction

Current thinking in cognitive psychology suggests that humans process information at a minimum of two distinct levels. There is extensive evidence for the existence of (at least) two processing systems within the human brain, one that involves fast, parallel, unconscious processing, and one that involves slow, serial, more conscious processing (Kirkpatrick and Epstein 1992; Chaiken and Trope 1999; Smith and DeCoster 2000; Epstein 2003; Kahneman 2011). Dual-process models of automatic and controlled social cognition have been proposed in nearly every domain of social psychology.

Evidence from neurosciences supports this separation, with identifiably different brain regions involved in each of the two systems (Lieberman 2007). Such systems, which we term U-level (unconscious) and C-level (conscious), can operate simultaneously or sequentially, and are most effective in different contexts. The former, in particular, works intuitively, effortlessly, globally, and emotionally. The latter, in turn, works logically, systematically, effortfully, and rationally.

Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In this work, we propose to emulate such dual-process model through a novel two-level affective common sense reasoning framework, termed sentic activation, that concurrently exploits multi-dimensionality reduction and graph mining techniques for the natural language processing (NLP) task of sentiment analysis.

The structure of the paper is as follows: the first section introduces the field of sentiment analysis and explains why common sense reasoning is particularly useful for it; the second section explains in detail the multi-dimensionality reduction techniques adopted to perform unconscious affective reasoning; the third section illustrates the graph mining techniques employed to perform reasoning at conscious level; next, the development of a sentiment analysis engine and its evaluation are presented; the last section, eventually, comprises concluding remarks and future directions.

Sentiment Analysis

Sentiment analysis is a branch of the broad field of text data mining (Hearst 1997) and refers generally to the process of extracting interesting and non-trivial patterns or knowledge from unstructured text documents. It can be viewed as an extension of data mining or knowledge discovery from (structured) databases (Fayyad, Piatetsky, and Smyth 1996; Simoudis 1996).

As the most natural form of storing information is text, sentiment analysis is believed to have a commercial potential higher than that of data mining. Sentiment analysis, however, is also a much more complex task as it involves dealing with text data that are inherently unstructured and fuzzy. It is a multi-disciplinary research area that involves the adoption of techniques in fields such as text analysis, information retrieval and extraction, auto-categorization, machine learning, clustering, and visualization.

Most of the existing approaches to opinion mining and sentiment analysis rely on the extraction of a vector representing the most salient and important text features, which is later used for classification purposes. Some of the most commonly used features are term frequency (Wu et al. 2008) and presence (Pang, Lee, and Vaithyanathan 2002). The latter, in particular, is a binary-valued feature vectors in which the entries merely indicate whether a term occurs (value 1) or not (value 0) formed a more effective basis for review polarity classification.

This is indicative of an interesting difference between typical topic-based text categorization and polarity classification. While a topic is more likely to be emphasized by frequent occurrences of certain keywords, overall sentiment may not usually be highlighted through repeated use of the same terms. Differently from topics, in fact, sentiments can often be expressed in a more subtle manner, making it difficult to be identified by specific keywords, especially when considering multiple domains.

Humans readers do not face such difficulty as they can infer the cognitive and affective information associated with natural language text through their common sense knowledge, that is, obvious or widely accepted things that people normally know about the world but which are usually left unstated in discourse, e.g., that things fall downwards (and not upwards) and people smile when they are happy.

An important feature of common sense reasoning, in fact, is the sensitivity to nuanced readings of language. A sentence can be read differently depending on nuances in opinionated text and such nuanced reading can lead to markedly different reasoning trajectories. The first step in human cognitive and affective information processing, in fact, is in an appraisal of the current situation (Scherer, Shorr, and Johnstone 2001). In order to accordingly infer the cognitive and affective information associated with natural language text, next-generation sentiment analysis methods need to go beyond a mere word-level analysis and use common sense reasoning to better grasp the conceptual rules that govern sentiment and the clues that can convey these concepts from realization to verbalization in the human mind.

To this end, we propose an ensemble application of multi-dimensionality reduction and graph mining techniques on AffectNet, an affective common sense knowledge base built upon WordNet-Affect (WNA) (Strapparava and Valitutti 2004), a linguistic resource for the lexical representation of affect, and ConceptNet (Havasi, Speer, and Alonso 2007), a semantic network of common sense knowledge.

Unconscious Affective Reasoning

In recent years, neuroscience has contributed a lot to the study of emotions through the development of novel methods for studying emotional processes and their neural correlates. In particular, new methods used in affective neuroscience, e.g., fMRI, lesion studies, genetics, electrophysiology, paved the way towards the understanding of the neural circuitry that underlies emotional experience and of the manner in which emotional states influence health and life outcomes.

A key contribution in the last two decades has been to provide evidence against the notion that emotions are sub-cortical and limbic, whereas cognition is cortical. This notion was reinforcing the flawed Cartesian dichotomy between thoughts and feelings (Damasio 2003). There is now ample evidence that the neural substrates of cognition and emotion overlap substantially (Daggleish, Dunn, and Mobbs 2009). Cognitive processes, such as memory encoding and retrieval, causal reasoning, deliberation, goal appraisal, and planning, operate continually throughout the experience of emotion.

This evidence points to the importance of considering the affective components of any human-computer interaction (Calvo and D’Mello 2010). Affective neuroscience, in particular, has provided evidence that elements of emotional learning can occur without awareness (Ohman and Soares 1998) and elements of emotional behavior do not require explicit processing (Calvo and Nummenmaa 2007). Affective information processing, in fact, mainly takes place at unconscious level (U-level) (Epstein 2003). Reasoning, at this level, relies on experience and intuition, which allow considering issues intuitively and effortlessly. Hence, rather than reflecting upon various considerations in sequence, the U-level forms a global impression of the different issues. In addition, rather than applying logical rules or symbolic codes (e.g., words or numbers), the U-level considers vivid representations of objects or events.

Such representations are laden with the emotions, details, features, and sensations that correspond to the objects or events. Such human capability of summarizing the huge amount of inputs and outputs of previous situations to find useful patterns that might work at the present time is hereby implemented by means of AffectiveSpace, the vector space representation of AffectNet (Cambria and Hussain 2012).

Affective Compression

In cognitive science, the term ‘compression’ refers to transforming diffuse and distended conceptual structures that are less congenial to human understanding so that they become better suited to our human-scale ways of thinking. Compression is hereby implemented by representing affective common sense knowledge in a way that it is neither too concrete nor too abstract with respect to the detail granularity needed for performing a particular task.

To this end, we first generate a matrix representation of AffectNet by applying blending (Havasi et al. 2009), a technique that performs inference over multiple sources of data simultaneously, taking advantage of the overlap between them. In particular, the alignment of ConceptNet and WNA yields A , a matrix in which common sense and affective knowledge coexist, i.e., a matrix $14,301 \times 117,365$ whose rows are concepts (e.g., ‘dog’ or ‘bake cake’), whose columns are either common sense and affective features (e.g., ‘isA-pet’ or ‘hasEmotion-joy’), and whose values indicate truth values of assertions. 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 by the same features and decreases when they are described by features that are negations of each other. In particular, we use truncated singular value decomposition (SVD) (Wall, Rechtsteiner, and Rocha 2003) 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 minimizing the Frobenius norm of the difference between A and \tilde{A} under the constraint $\text{rank}(\tilde{A}) = k$. For the Eckart-Young theorem (Eckart and Young 1936) it represents the best approximation of A in the least-square sense, in fact:

$$\begin{aligned} \min_{\tilde{A}|\text{rank}(\tilde{A})=k} |A - \tilde{A}| &= \min_{\tilde{A}|\text{rank}(\tilde{A})=k} |\Sigma - U^* \tilde{A} V| \\ &= \min_{\tilde{A}|\text{rank}(\tilde{A})=k} |\Sigma - S| \end{aligned}$$

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:

$$\begin{aligned} \min_{\tilde{A}|\text{rank}(\tilde{A})=k} |\Sigma - S| &= \min_{s_i} \sqrt{\sum_{i=1}^n (\sigma_i - s_i)^2} = \\ &= \min_{s_i} \sqrt{\sum_{i=1}^k (\sigma_i - s_i)^2 + \sum_{i=k+1}^n \sigma_i^2} = \sqrt{\sum_{i=k+1}^n \sigma_i^2} \end{aligned}$$

Therefore, \tilde{A} of rank k is the best approximation of A in the Frobenius norm sense when $\sigma_i = s_i$ ($i = 1, \dots, k$) and the corresponding singular vectors are 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, i.e., the basis e_0, \dots, e_{k-1} of the vector space. For example, the most significant eigenmood, e_0 , represents concepts with positive affective valence. That is, the larger a concept’s component in the e_0 direction is, the more affectively positive it is likely to be. Concepts with negative e_0 components, then, are likely to have negative affective valence.

Thus, by exploiting the information sharing property of truncated SVD, 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, we can find concepts such as ‘beautiful day’, ‘birthday 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 (nearly opposite with respect to the centre of the space).

By reducing the dimensionality of the matrix representation of A , AffectiveSpace compresses the feature space of affective common sense knowledge into one that allows to better gain global insight and human-scale understanding. Compression, in particular, is achieved by balancing the number of singular values discarded when synthesizing the vector space, in a way that the affective common sense

knowledge representation is neither too concrete nor too abstract with respect to the detail granularity needed for inferring the cognitive and affective information associated with opinionated text.

Affective Clustering

The vector space representation of affective common sense knowledge is clustered using sentic medoids (Cambria et al. 2011), a technique that adopts a k -medoids approach (Kaufman and Rousseeuw 1990) to partition the given observations into k clusters around as many centroids, trying to minimize a given cost function. Differently from the k -means algorithm (Hartigan and Wong 1979), which does not pose constraints on centroids, k -medoids do assume that centroids must coincide with k observed points.

The most commonly used algorithm for finding the k medoids is the partitioning around medoids (PAM) algorithm, which 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. Since k -medoids clustering is a NP-hard problem (Garey and Johnson 1979), different approaches based on alternative optimization algorithms have been developed, though taking risk of being trapped around local minima. We use a modified version of the algorithm recently proposed by Park and Jun (Park and Jun 2009), 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 of the Hourglass of Emotions (Cambria, Livingstone, and Hussain 2012), a novel biologically-inspired and psychologically-motivated emotion categorization model, based on Plutchik’s studies on human emotions (Plutchik 2001), that can potentially describe any human emotion in terms of four independent but concomitant dimensions, whose different levels of activation make up the total emotional state of the mind (Fig. 1).

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 (Duda and Hart 1973). However, we decide to use as initial centroids the sentic levels of the Hourglass, hence, 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. 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 $D(a, b) = \sqrt{\sum_{i=1}^p (a_i - b_i)^2}$, the used algorithm, applied for each of the four affective dimensions, can be summarized as follows:

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

After such clustering process, concepts that are semantically and affectively related to the input data can be intuitively retrieved by analogy and unconsciously crop out to the C-level.

Conscious Affective Reasoning

U-level and C-level are two conceptual systems that operate by different rules of inference. While the former operates emotionally and intuitively, the latter relies on logic and rationality. In particular, the C-level analyzes issues with effort, logic, and deliberation rather than relying on intuition. Hence, while at U-level the vector space representation of AffectNet is exploited to intuitively guess semantic and affective relations between concepts, at C-level associations between concepts are made according to the actual connections between different nodes in the graph representation of affective common sense knowledge.

Memory, in fact, is not a ‘thing’ that is stored somewhere in a mental warehouse and can be pulled out and brought to the fore. Rather, it is a potential for reactivation of a set of concepts that together constitute a particular meaning. Associative memory involves the activation of networks of association – thoughts, feelings, wishes, fears, and perceptions – that are connected, so that activation of one node in the network leads to activation of the others (Westen 2002).

Sentic activation aims to implement such process through the ensemble application of multi-dimensionality reduction and graph mining techniques. Specifically, the semantically and affectively related concepts retrieved by means of AffectiveSpace (at U-level) are fed into AffectNet (at C-level), in order to crawl it according to how such seed concepts are interconnected to each other and to other concepts in the semantic network.

Spreading Affect

Spreading activation theory uses an associative relevancy measure over declarative memory by an exponential decay of activation with distance in the network structure (Anderson and Pirolli 1984). In particular, we exploit spectral association (Havasi, Speer, and Holmgren 2010), a technique that assigns values, or activations, to seed concepts and spreads their values across the AffectNet graph.




Figure 1: The net of the Hourglass of Emotions

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. These related concepts are likely to have similar affective values.

This can be seen as an alternate way of assigning affective values to all concepts, which simplifies the process by not relying on an outside resource such as WNA. In particular, we build a matrix A 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 A to a vector containing a single concept spreads that concept’s value to its connected concepts. Applying A^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:

$$1 + A + \frac{A^2}{2!} + \frac{A^3}{3!} + \dots = e^A \quad (1)$$

We can calculate this odd operator, e^A , because we can factor A . A is already symmetric, so instead of applying Lanczos’ method to AA^T and getting the SVD, we can apply it directly to A and get the spectral decomposition $A = V\Lambda V^T$. As before, we can raise this expression to any power and cancel everything but the power of Λ . Therefore, $e^A = Ve^{\Lambda}V^T$. 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, hence, save space while generalizing 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 normalizing the truncated rows of $Ve^{\Lambda/2}$ to unit vectors, and multiplying that matrix by its transpose to get a rescaled version of $Ve^{\Lambda}V^T$.

Spectral association can spread not only positive, but also negative activation values. Hence, unconscious reasoning at U-level is exploited not only to retrieve concepts that are most semantically and affectively related, but also concepts that are most likely to be unrelated with the input data (lowest dot product). While the former are exploited to spread cognitive and affective energy across the AffectNet graph, the latter are used to contain such activation in a way that potentially unrelated concepts (and their twins) do not get triggered.

Evaluation

The adopted brain-inspired ensemble application of dimensionality reduction and graph mining techniques (to which we hereby refer as unconscious and conscious reasoning, respectively) allows sentic activation to more efficiently infer the cognitive and affective information associated with natural language text. In fact, we tested sentic activation on a benchmark for affective common sense knowledge (BACK) built by applying CF-IOF (concept frequency - inverse opinion frequency), a technique similar to TF-IDF, on a 5,000 blogpost database extracted from LiveJournal¹, a virtual community of users who keep a blog, journal, or diary.

An interesting feature of this website is that bloggers are allowed to label their posts with both a category and a mood tag, by choosing from predefined categories and mood themes or by creating new ones. Since the indication of mood tags is optional, posts are likely to reflect the true mood of the authors, which is not always true for category tags. After a manual evaluation of 200 posts, the category tags turned out to be very noisy (53% accuracy).

The mood tags, however, showed a good enough reliability (78% accuracy) so we used them to test the engine’s affect recognition performances. In order to have full correspondence between LiveJournal mood labels and the activation levels of the Hourglass model, a number of English-speaking students has been asked to map each of the 130 mood labels into the 24 emotional labels of the Hourglass model. A set of 80 moods (those with higher confidence level) were selected for inclusion in the blogpost database. CF-IOF 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 .

¹<http://livejournal.com>

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%

Table 1: Distribution of Pleasantness sentic levels.

This frequency is then multiplied by the logarithm of the inverse frequency of the concept in the whole collection of opinions, that is:

$$CF-IOF_{c,d} = \frac{n_{c,d}}{\sum_k n_{k,d}} \log \sum_k \frac{n_k}{n_c} \quad (2)$$

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, we exploited CF-IOF weighting to filter out common concepts in the LiveJournal corpus and detect relevant mood-dependent semantics for each of the Hourglass sentic levels.

The result was a benchmark of 2000 affective concepts that were screened by 21 English-speaking students who were asked to evaluate the level b associated to each concept $b \in \Theta = \{\theta \in \mathbb{Z} \mid -1 \leq \theta \leq 1\}$ (each integer corresponding to a level of the Hourglass model) for each of the four affective dimensions. Results obtained were averaged (Table 1). BACK’s concepts were compared with the classification results obtained by applying the AffectiveSpace process (U-level), spectral association (C-level), and sentic activation (U&C-level). Results showed that sentic activation achieves +13.9% and +8.2% accuracy than the AffectiveSpace process and spectral association, respectively.

Brain-Inspired Sentiment Analysis

In order to test sentic activation also within a real-world problem, we developed a brain-inspired software engine for sentiment analysis. This software 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 sentiment targets, and 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 AffectNet. 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 recognize complex concepts also 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 target spotting module aims to individuate one or more sentiment targets, such as people, places, events, and ideas, from the input concepts. This is done by projecting the retrieved concepts into both the graph and the vector space representation of AffectNet, in order to assign these to a specific conceptual class. The categorization does not consist in simply labeling 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 graph and dot product in the vector space). The affect interpreter, similarly, projects the retrieved concepts into the vector space representation of AffectNet, in order to assign these to a specific affective class, and therefore calculates polarity in terms of the Hourglass dimensions.

As an example of how the software engine works, we can examine intermediate and final outputs obtained when a natural language opinion is given as input to the system. We choose 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!”. After the pre-processing and semantic parsing operations, we obtain the following small bag of concepts (SBoCs):

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 2).

System Comparison

In order to evaluate the different facets of the engine from different perspectives, we used a PatientOpinion² dataset,

²<http://patientopinion.org.uk>

Target	Category	Moods	Polarity
‘iphone4’	‘phones’, ‘electronics’	‘ecstasy’, ‘interest’	+0.71
‘speaker’	‘electronics’, ‘music’	‘annoyance’	-0.34
‘touch-screen’	‘electronics’	‘ecstasy’, ‘anticipation’	+0.82
‘camera’	‘photography’, ‘electronics’	‘acceptance’	+0.56

Table 2: Structured output example of the engine

and compared results obtained using AffectiveSpace (U-level), spectral association (C-level), and sentic activation (U&C-level). The resource is a dataset obtained from PatientOpinion, a social enterprise pioneering an on-line feedback service for users of the UK national health service to enable people to share their recent experience of local health services on-line. It is a manually tagged dataset of 2,000 patient opinions that associates to each post a category (namely, clinical service, communication, food, parking, staff, and timeliness) and a positive or negative polarity. We used it to test the detection of opinion targets and the polarity associated with these (F-measure values are reported in Table 3).

Category	U-Level	C-Level	U&C-Level
clinical service	62.4%	71.5%	82.4%
communication	62.3%	59.8%	79.1%
food	70.7%	69.6%	81.7%
parking	56.3%	53.7%	77.2%
staff	58.5%	49.2%	73.8%
timeliness	69.2%	61.9%	80.6%

Table 3: F-measures relative to PatientOpinion evaluation.

Conclusions

Current thinking in psychology ascribes the involvement of multiple concurrent levels of processing, much of which is below consciousness, to human reasoning. In this paper, we proposed a brain-inspired computational model for conscious and unconscious affective common sense reasoning. In particular, we explored the ensemble application of multi-dimensionality reduction and graph mining techniques on an affective common sense knowledge base to go beyond mere word-level sentiment analysis.

In the future, novel multi-dimensionality reduction and graph mining techniques will be explored in order to dynamically configure AffectNet and, hence, model the switch between different reasoning strategies and between the foci around which such strategies are developed (Cambria, Olsner, and Kwok 2012).

References

Anderson, J., and Pirolli, P. 1984. Spread of activation. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 10:791–799.

- Calvo, R., and D’Mello, S. 2010. Affect detection: An interdisciplinary review of models, methods, and their applications. *IEEE Transactions on Affective Computing* 1(1):18–37.
- Calvo, M., and Nummenmaa, L. 2007. Processing of unattended emotional visual scenes. *Journal of Experimental Psychology: General* 136:347–369.
- Cambria, E., and Hussain, A. 2012. *Sentic Computing: Techniques, Tools, and Applications*. Berlin Heidelberg: Springer.
- Cambria, E.; Mazzocco, T.; Hussain, A.; and Eckl, C. 2011. Sentic medoids: Organizing affective common sense knowledge in a multi-dimensional vector space. In *Advances in Neural Networks*, volume 6677 of *Lecture Notes in Computer Science*, 601–610. Berlin Heidelberg: Springer-Verlag.
- Cambria, E.; Livingstone, A.; and Hussain, A. 2012. The hourglass of emotions. In *Cognitive Behavioral Systems*, *Lecture Notes in Computer Science*. Berlin Heidelberg: Springer.
- Cambria, E.; Olsher, D.; and Kwok, K. 2012. Sentic analogy: Swapping affective common sense reasoning strategies and foci. In *CogSci*.
- Chaiken, S., and Trope, Y. 1999. *Dual-Process Theories in Social Psychology*. New York: Guilford.
- Dalgleish, T.; Dunn, B.; and Mobbs, D. 2009. Affective neuroscience: Past, present, and future. *Emotion Review* 355–368.
- Damasio, A. 2003. *Looking for Spinoza: Joy, Sorrow, and the Feeling Brain*. Harcourt, Inc.
- Duda, R., and Hart, P. 1973. *Pattern Classification and Scene Analysis*. New York: John Wiley & Sons Inc.
- Eckart, C., and Young, G. 1936. The approximation of one matrix by another of lower rank. *Psychometrika* 1(3):211–218.
- Epstein, S. 2003. Cognitive-experiential self-theory of personality. In Millon, T., and Lerner, M., eds., *Comprehensive Handbook of Psychology*, volume 5. Hoboken, NJ: Wiley & Sons. 159–184.
- Fayyad, U.; Piatetsky, G.; and Smyth, P. 1996. From data mining to knowledge discovery: An overview. In *Advances in Knowledge Discovery and Data Mining*. Cambridge: MIT Press. 1–36.
- Garey, M., and Johnson, D. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman.
- Hartigan, J., and Wong, M. 1979. Algorithm AS 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society* 28(1):100–108.
- Havasi, C.; Speer, R.; Pustejovsky, J.; and Lieberman, H. 2009. Digital intuition: Applying common sense using dimensionality reduction. *IEEE Intelligent Systems* 24(4):24–35.
- Havasi, C.; Speer, R.; and Alonso, J. 2007. ConceptNet 3: A flexible, multilingual semantic network for common sense knowledge. In *RANLP*.
- Havasi, C.; Speer, R.; and Holmgren, J. 2010. Automated color selection using semantic knowledge. In *AAAI CSK*.
- Hearst, M. 1997. Text data mining: Issues, techniques, and the relationship to information access. In *UW/MS Workshop on Data Mining*.
- Kahneman, D. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.
- Kaufman, L., and Rousseeuw, P. 1990. *Finding Groups in Data: An Introduction to Cluster Analysis*. Wiley-Interscience.
- Kirkpatrick, L., and Epstein, S. 1992. Cognitive experiential self-theory and subjective probability: Further evidence for two conceptual systems. *Journal of Personality and Social Psychology* 63:534–544.
- Lieberman, M. 2007. Social cognitive neuroscience: A review of core processes. *Annual Review of Psychology* 58:259–89.
- Ohman, A., and Soares, J. 1998. Emotional conditioning to masked stimuli: Expectancies for aversive outcomes following nonre-cognized fear-relevant stimuli. *Journal of Experimental Psychology: General* 127(1):69–82.
- Pang, B.; Lee, L.; and Vaithyanathan, S. 2002. Thumbs up? Sentiment classification using machine learning techniques. In *EMNLP*, 79–86.
- Park, H., and Jun, C. 2009. A simple and fast algorithm for k-medoids clustering. *Expert Systems with Applications* 36(2):3336–3341.
- Plutchik, R. 2001. The nature of emotions. *American Scientist* 89(4):344–350.
- Scherer, K.; Shorr, A.; and Johnstone, T. 2001. *Appraisal Processes in Emotion: Theory, Methods, Research*. Canary: Oxford University Press.
- Simoudis, E. 1996. Reality check for data mining. *IEEE Expert* 11(5).
- Smith, E., and DeCoster, J. 2000. Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychological Review* 4(2):108–131.
- Strapparava, C., and Valitutti, A. 2004. WordNet-Affect: An affective extension of WordNet. In *LREC*.
- Wall, M.; Rechtsteiner, A.; and Rocha, L. 2003. Singular value decomposition and principal component analysis. In Berrar, D.; Dubitzky, W.; and Granzow, M., eds., *A Practical Approach to Microarray Data Analysis*. Springer. 91–109.
- Westen, D. 2002. Implications of developments in cognitive neuroscience for psychoanalytic psychotherapy. *Harv Rev Psychiatry* 10(6):369–73.
- Wu, H.; Luk, R.; Wong, K.; and Kwok, K. 2008. Interpreting TF-IDF term weights as making relevance decisions. *ACM Transactions on Information Systems* 26(3).