

SenticNet: A Publicly Available Semantic Resource for Opinion Mining

Erik Cambria
University of Stirling
Stirling, FK4 9LA UK
eca@cs.stir.ac.uk

Robert Speer
MIT Media Lab
Cambridge, 02139-4307 USA
rspeer@mit.edu

Catherine Havasi
MIT Media Lab
Cambridge, 02139-4307 USA
havasi@media.mit.edu

Amir Hussain
University of Stirling
Stirling, FK4 9LA UK
ahu@cs.stir.ac.uk

Abstract

Today millions of web-users express their opinions about many topics through blogs, wikis, fora, chats and social networks. For sectors such as e-commerce and e-tourism, it is very useful to automatically analyze the huge amount of social information available on the Web, but the extremely unstructured nature of these contents makes it a difficult task. SenticNet is a publicly available resource for opinion mining built exploiting AI and Semantic Web techniques. It uses dimensionality reduction to infer the polarity of common sense concepts and hence provide a public resource for mining opinions from natural language text at a semantic, rather than just syntactic, level.

Introduction

Web 2.0 gifted its users with the power of actively participating in its growth. In contrast to the read-only Web, today's read-write Web gives people the ability to interact, share, and collaborate through social networks, online communities, blogs, wikis and other online collaborative media. With the advent of services like eBay, TripAdvisor and Amazon, web-users' activity became central to most web applications, and their opinions gave birth to a collective intelligence that is often more listened to than experts' viewpoints.

However, the distillation of knowledge from this huge amount of unstructured information is a very complicated task. To this end we developed SenticNet, a publicly available semantic resource for opinion mining built using common sense reasoning techniques together with an emotion categorization model and an ontology for describing human emotions.

Opinion Mining

Opinion mining is a new discipline which has recently attracted increased attention within fields such as marketing, personal affective profiling, and financial market prediction.

Although often associated with sentiment analysis, which consists in inferring emotional states from text, opinion mining is an independent area related to natural language processing and text mining that deals with the identification of

opinions and attitudes in natural language texts. In particular, given a textual resource containing a number of opinions o about a number of topics, we would like to be able to assign each opinion a polarity $p(o) \in [-1, 1]$, representing a range from generally unfavorable to generally favorable, and to aggregate the polarities of opinions on various topics to discover the general sentiment about those topics.

Existing approaches to automatic identification and extraction of opinions from text can be grouped into three main categories: keyword spotting, in which text is classified into categories based on the presence of fairly unambiguous affect words (Ortony, Clore, and Collins 1998)(Wiebe, Wilson, and Claire 2005), lexical affinity, which assigns arbitrary words a probabilistic affinity for a particular opinion (Wilson, Wiebe, and Hoffmann 2005)(Somasundaran, Wiebe, and Ruppenhofer 2008), and statistical methods, which consist in calculating the valence of keywords, punctuation and word co-occurrence frequencies on the base of a large training corpus (Hu and Liu 2004)(Pang and Lee 2005)(Abbasi, Chen, and Salem 2008).

These approaches mainly rely on parts of text in which opinions are explicitly expressed such as positive terms (e.g. good, nice, excellent, fortunate, correct, superior, best) and negative terms (e.g. bad, nasty, poor, unfortunate, wrong, inferior, worst). But in general opinions are expressed implicitly through context and domain dependent concepts, which makes purely syntactical approaches ineffective.

Sentic Computing

Sentic Computing (Cambria et al. 2010c) is a novel opinion mining and sentiment analysis paradigm which exploits AI and Semantic Web techniques to better recognize, interpret and process opinions and sentiments in natural language text. In Sentic Computing, whose term derives from the Latin 'sentire' (the root of words such as sentiment and sensation) and 'sense' (intended as common sense), the analysis of text is not based on statistical learning models but rather on common sense reasoning tools (Cambria et al. 2009b) and domain-specific ontologies (Cambria et al. 2010a).

Differently from statistical classification, which generally requires large inputs and thus cannot appraise texts with satisfactory granularity, Sentic Computing enables the analysis of documents not only on the page or paragraph-level but even on the sentence level.

ConceptNet

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

The Open Mind Common Sense project has been collecting this kind of knowledge from volunteers on the Internet since 2000 to provide intuition to AI systems and applications. ConceptNet (Havasi, Speer, and Alonso 2007) represents the information in the Open Mind corpus as a directed graph in which the nodes are concepts and the labelled edges are assertions of common sense that interconnect them.

AffectiveSpace

AffectiveSpace (Cambria et al. 2009a) is a n -dimensional vector space built from ConceptNet and WordNet-Affect, a linguistic resource for the lexical representation of affective knowledge (Strapparava and Valitutti 2004).

After aligning the lemma forms of ConceptNet concepts with the lemma forms of the words in WordNet-Affect, we perform singular value decomposition (SVD) on the resulting matrix and use dimensionality reduction to discard those components representing relatively small variations in the data. This yields a multi-dimensional space which we call AffectiveSpace (illustrated in Fig. 1), in which different vectors represent different ways of making binary distinctions among concepts and emotions.

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 center of the space).

If we choose to discard all but the first 100 principal components, common sense concepts and emotions are represented by vectors of 100 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_{99} 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.

The Hourglass of Emotions

The Hourglass of Emotions (Fig. 2) is an affective categorization model developed starting from Plutchik’s studies on human emotions (Plutchik 2001). In the model sentiments

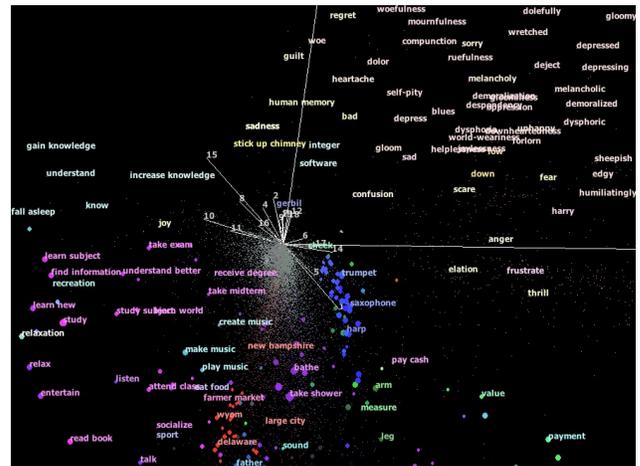


Figure 1: AffectiveSpace

are reorganized around four independent dimensions whose different levels of activation make up the total emotional state of the mind.

The Hourglass of Emotions, in fact, is based on the idea that the mind is made of different independent resources and that emotional states result from turning some set of these resources on and turning another set of them off. Each such selection changes how we think by changing our brain’s activities: the state of anger, for example, appears to select a set of resources that help us react with more speed and strength while also suppressing some other resources that usually make us act prudently.

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

1. the user is happy with the service provided
2. the user is interested in the information supplied
3. the user is comfortable with the interface
4. the user is disposed to use the application

Each of the four affective dimensions is characterized by six levels of activation, called ‘sentic levels’, which determine the intensity of the expressed/perceived emotion, for a total of 24 labels specifying ‘elementary emotions’.

The concomitance of the different affective dimensions makes possible the generation of ‘compound emotions’ such as ‘love’, which is given by the sum of positive values of Pleasantness and Aptitude, ‘aggressiveness’, given by the concomitance of Attention and Sensitivity, or ‘disappointment’, which results from the combination of negative values of Pleasantness and Attention.

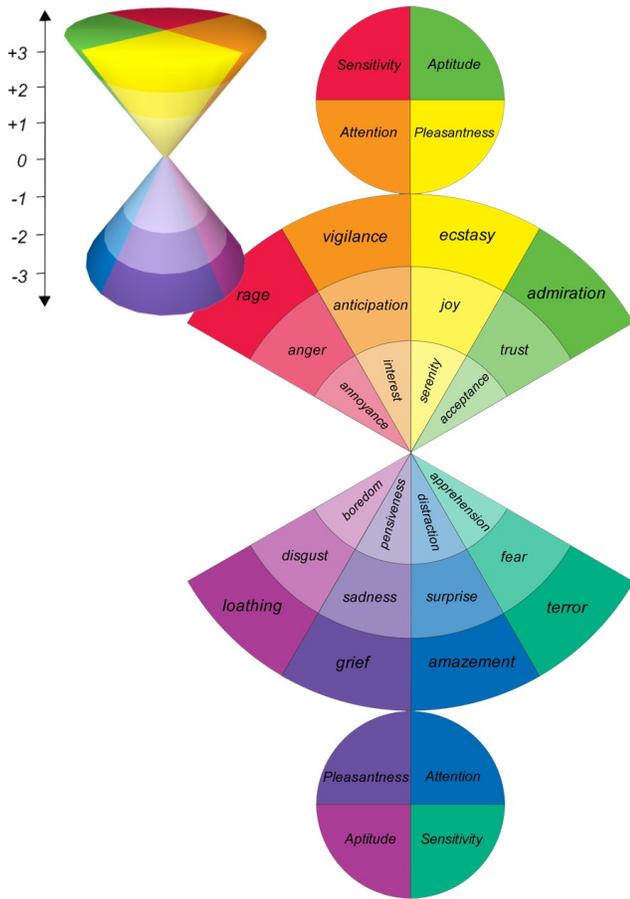


Figure 2: The Hourglass of Emotions

The Human Emotion Ontology

The Human Emotion Ontology (HEO) (Grassi 2009) is conceived as a high level ontology for human emotions that supplies the most significant concepts and properties which constitute the centrepiece for the description of every human emotion. If necessary, these high level features can be further refined using lower level concepts and properties related to more specific descriptions or linked to other more specialized ontologies.

The main purpose of HEO is thus to create a description framework that could grant at the same time enough flexibility, by allowing the use of a wide and extensible set of descriptors to represent all the main features of an emotion, and interoperability, by allowing to map concepts and properties belonging to different emotion representation models.

HEO has been developed in OWL description logic (OWL DL) to take advantage of its expressiveness and its inference power in order to map the different models used in the emotion description. OWL DL, in fact, allows a taxonomical organization of emotion categories and properties restriction to link emotion description made by category and dimension.

In HEO, for example, Ekman’s ‘joy’ archetypal emotion represents a superclass for Plutchik’s ‘ecstasy’, ‘joy’

and ‘serenity’ emotions. Using property restriction, the Plutchik’s ‘joy’ emotion can also be defined as an emotion that ‘has Pleasantness some float $\in [+1,+2]$ ’, ‘interest’ as an emotion that ‘has Attention $\in [0,+1]$ ’ and ‘love’ as an emotion that ‘has Pleasantness some float $\in [0,+3]$ and Aptitude some float $\in [0,+3]$ ’.

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

SentiWordNet

The development of SenticNet was inspired by SentiWordNet (Esuli and Sebastiani 2006), a lexical resource in which each WordNet synset is associated to three numerical scores describing how objective, positive and negative the terms contained in the synset are. Each of the three scores ranges from 0.0 to 1.0, and their sum is 1.0 for each synset. This means that a synset may have non-zero scores for all the three categories, which would indicate that the corresponding terms have, in the sense indicated by the synset, each of the three opinion-related properties only to a certain degree.

The method used to develop SentiWordNet is based on the quantitative analysis of the ‘glosses’ associated to synsets, and on the use of the resulting vector representations for semi-supervised synset classification. The three scores are derived by combining the results produced by a committee of eight ternary classifiers, all characterized by similar accuracy levels but different classification behaviour.

SentiWordNet currently represents a good resource for opinion mining however it contains a lot of noise and it mainly provides opinion polarity at syntactical level, leaving out polarity information for common sense knowledge concepts such as ‘accomplish goal’, ‘bad feeling’, ‘celebrate special occasion’, ‘lose temper’ or ‘be on cloud nine’, which are usually found in natural language text to express positive and negative viewpoints.

Building SenticNet

The aim of this work is to create a collection of commonly used ‘polarity concepts’ i.e. common sense concepts with relatively strong positive or negative polarity.

To this end, differently from SentiWordNet (which also contains null polarity terms), we discard concepts with neutral or almost neutral polarity i.e. concepts with polarity magnitude close to zero. Moreover, while SentiWordNet stores three values for each synset, in SenticNet each concept c is associated to just one value p_c , i.e. a float $\in [-1,1]$ representing its polarity, in order to avoid redundancy and more easily represent SenticNet as a semantic network.

Therefore in SenticNet concepts like ‘make good impression’, ‘look attractive’, ‘show appreciation’ or ‘good deal’ are likely to have p_c very close to 1 while concepts such as ‘being fired’, ‘leave behind’ or ‘lose control’ are likely to have $p_c \approx -1$.

Defining Concept Polarity

In Sentic Computing we define concept polarity as the algebraic sum of the Hourglass model’s sentic labels. But while positiveness and negativeness of Pleasantness and Aptitude reflect positive and negative polarity, Attention and Sensitivity are mono-polarized dimensions.

Sentic values such as distraction or surprise, in fact, represent negative activation values in terms of Attention but, in terms of polarity, they are associated to positive concepts. Viceversa positive values of Sensitivity, e.g. anger or annoyance, are generally associated with negative polarity concepts. For these reasons we define concept polarity as:

$$p_c = \frac{Plsn(c) + |Attn(c)| - |Snst(c)| + Aptt(c)}{9}$$

Concept Polarity Inference

The calculation of concept polarity is based on the assumption that relative distances between concepts in AffectiveSpace, i.e. their dot product, are directly proportional to their polarity degree difference.

This is a very logical assumption since in AffectiveSpace concepts concerning the same emotions are likely to be close to each other and usually the positiveness/negativeness of an emotion expressed in a sentence is directly proportional to the positiveness/negativeness of its polarity.

For each affective dimension of the Hourglass model we seek for concepts which are semantically correlated to positive sentic values and, at the same time, uncorrelated to negative sentic values, and viceversa. For example, to find positive polarity concepts associated to Pleasantness, we seek for concepts semantically correlated to ‘ecstasy’, ‘joy’ and ‘serenity’ and, at the same time, semantically uncorrelated to ‘pensiveness’, ‘sadness’ and ‘grief’. We then do the opposite to seek for negative polarity concepts.

This process is performed through two different techniques called ‘blending’ and ‘spectral association’.

Blending

Blending (Havasi et al. 2009) is a technique that performs inference over multiple sources of data simultaneously, taking advantage of the overlap between them. It basically combines two sparse matrices linearly into a single matrix in which the information between the two initial sources is shared.

When we perform SVD on a blended matrix, the result is that new connections are made in each source matrix taking into account information and connections present in the other matrix, originating from the information that overlaps. By this method, we can combine different sources of general knowledge, or overlay general knowledge with domain-specific knowledge, such as medical, geological or financial knowledge.

In this work we use it to combine the domain-general knowledge in ConceptNet with the affective knowledge contained in WordNet-Affect.

Spectral Association

Spectral association (Havasi, Speer, and Holmgren 2010) consists in assigning values, or activations, to key affective concepts such as ‘good’ and ‘interesting’, and applying an operation that spreads their values across the ConceptNet graph. This operation, an approximation of many steps of spreading activation, transfers the most activation to concepts that are connected to the key concepts by short paths or many different paths in common sense knowledge.

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 WordNet-Affect. In particular, we build a matrix C that relates concepts to other concepts, instead of their features, and add up the scores over all relations that relate one concept to another, disregarding direction.

Applying C to a vector containing a single concept spreads that concept’s value to its connected concepts. Applying C^2 spreads that value to concepts connected by two links (including back to the concept itself). But what we’d really like is to spread the activation through any number of links, with diminishing returns, so perhaps the operator we want is:

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

We can calculate this odd operator, e^C , because we can factor C . C is already symmetric, so instead of applying Lanczos’ method to CC^T and getting the SVD, we can apply it directly to C and get the spectral decomposition $C = V\Lambda V^T$. As before, we can raise this expression to any power and cancel everything but the power of Λ . Therefore, $e^C = Ve^\Lambda V^T$. This simple twist on the SVD lets us calculate spreading activation over the whole matrix instantly.

As with the SVD, we can truncate these matrices to k axes and therefore save space while 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 $Ve^\Lambda V^T$.

Encoding SenticNet

After retrieving polarity concepts through blending and spectral association operations, we need to reorganize them in a way that they can be represented in a unique and consistent resource. We deal with possible conflicts by discarding duplicate concepts with smaller polarity magnitude since bigger concept polarity values usually correspond to more reliability (higher dot products in the vector space).

Since concepts are usually strongly related to just one or two affective dimensions (most of compound emotions are in fact given by summing just two elementary emotions), the average magnitude is pretty low. Therefore, in order to obtain more homogeneous and intelligible polarity values, we run a normalization process over SenticNet before storing its contents in a Semantic Web aware format.

In order to represent SenticNet in a machine-accessible and machine-processable way, results are encoded in RDF triples using a XML syntax. In particular, concepts are identified using the ConceptNet Web API and statements, which have the form concept-hasPolarity-polarityValue, are encoded in RDF/XML format on the base of HEO.

Working with SenticNet

The current version of SenticNet, freely available at <http://cs.stir.ac.uk/~ecalsentics>, contains more than 5700 polarity concepts (nearly 40% of Open Mind corpus).

It is very easy to interface SenticNet with any kind of opinion mining application and, especially if used within Open Mind software (for a full correspondence of concepts), it is a very precise polarity detection tool.

We are currently using SenticNet in the field of e-health for analyzing online patient opinions in order to make a comprehensive and dynamic evaluation of the UK National Health Service (Cambria et al. 2010b). Each patient opinion is processed through a NLP module, which performs a first skim of text and determines the lemma forms of each word, a Semantic Parser, which deconstructs lemmatized text into concepts, and SenticNet, which detects the polarity (if any) of each retrieved concept. The overall opinion polarity p is given by simply averaging these concepts' polarity values:

$$p = \sum_{i=1}^N \frac{Plsn(c_i) + |Attn(c_i)| - |Snst(c_i)| + Aptt(c_i)}{9N}$$

Evaluation

As a preliminary evaluation, we compared SenticNet's with SentiWordNet's capacity of detecting opinion polarity over a collection of 2,000 patient opinions, of which 57% are labelled as negative, 32% as positive and the rest as neutral. After extracting concepts from each opinion, we looked up for relative polarity values in SentiWordNet and SenticNet and then compared these with the dataset labels to calculate statistical classifications such as precision and recall.

Results showed SenticNet to be much more accurate than SentiWordNet. The former, in particular, can identify positive opinions with much higher precision (79% against 53%) and significantly better recall rate (58% against 46%), for a total F-measure value of 67% versus 49%.

Conclusion and Future Work

In this work we developed SenticNet, a publicly available semantic resource for opinion mining built exploiting common sense reasoning techniques, such as blending and spectral activation, together with an emotion categorization model and an ontology for describing human emotions.

First tests confirmed the superiority of SenticNet with respect to currently available lexical resources for opinion mining developed using purely syntactical approaches.

We already employed SenticNet in real-world applications with remarkable results and we plan to exploit it a lot in the future for other projects in the field of opinion mining such as UI design and opinion visualization.

References

- Abbasi, A.; Chen, H.; and 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.
- Cambria, E.; Hussain, A.; Havasi, C.; and Eckl, C. 2009a. Affectivespace: Blending common sense and affective knowledge to perform emotive reasoning. In *CAEPIA09*.
- Cambria, E.; Hussain, A.; Havasi, C.; and Eckl, C. 2009b. *Common Sense Computing: From the Society of Mind to Digital Intuition and Beyond*, volume 5707 of *LNCS*. Berlin Heidelberg: Springer-Verlag. 252–259.
- Cambria, E.; Grassi, M.; Hussain, A.; and Havasi, C. 2010a. Sentic computing for social media marketing. Forthcoming.
- Cambria, E.; Hussain, A.; Havasi, C.; Eckl, C.; and Munro, J. 2010b. Towards crowd validation of the uk national health service. In *WebSci10*.
- Cambria, E.; Hussain, A.; Havasi, C.; and Eckl, C. 2010c. *Sentic Computing: Exploitation of Common Sense for the Development of Emotion-Sensitive Systems*, volume 5967 of *LNCS*. Berlin Heidelberg: Springer-Verlag. 148–156.
- Esuli, A., and Sebastiani, F. 2006. Sentiwordnet: A publicly available lexical resource for opinion mining. In *LREC06*.
- Grassi, M. 2009. *Developing HEO Human Emotion Ontology*, volume 5707 of *LNCS*. Berlin Heidelberg: Springer-Verlag. 244–251.
- 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 *RANLP07*.
- Havasi, C.; Speer, R.; and Holmgren, J. 2010. Automated color selection using semantic knowledge. In *AAAI CSK10*.
- Hu, M., and Liu, B. 2004. Mining opinion features in customer reviews. In *AAAI04*.
- Ortony, A.; Clore, G.; and Collins, A. 1998. *The Cognitive Structure of Emotions*. New York: Cambridge Press.
- Pang, B., and Lee, L. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *ACL05*.
- Plutchik, R. 2001. The nature of emotions. *American Scientist* 89(4):344–350.
- Somasundaran, S.; Wiebe, J.; and Ruppenhofer, J. 2008. Discourse level opinion interpretation. In *COLING08*.
- Strapparava, C., and Valitutti, A. 2004. Wordnet-affect: an affective extension of wordnet. In *LREC04*.
- Wiebe, J.; Wilson, T.; and Claire, C. 2005. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation* 2(3):165–210.
- Wilson, T.; Wiebe, J.; and Hoffmann, P. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *HLT-EMNLP05*.