

AffectiveSpace: Blending Common Sense and Affective Knowledge to Perform Emotive Reasoning

Erik Cambria¹, Amir Hussain¹, Catherine Havasi², and Chris Eckl³

¹ Dept. of Computing Science and Maths, University of Stirling, Scotland, UK

² MIT Media Lab, MIT, Massachusetts, USA

³ Sitekit Labs, Sitekit Solutions Ltd, Scotland, UK

`{eca,ahu}@cs.stir.ac.uk,havasi@media.mit.edu,chris.eckl@sitekit.net`
`http://cs.stir.ac.uk/~eca/sentics`

Abstract. The detection of emotions in text is a key issue for the development of intelligent systems. As demonstrated by the Turing test, a machine cannot be considered really intelligent unless it is also capable of perceiving and expressing emotions. In this work we focus on building a knowledge base which merges Common Sense and affective knowledge and use dimensionality reduction to perform emotive reasoning on it.

Key words: Common Sense Computing, AI, Semantic Networks, NLP, Analogies, Knowledge Base Management, Emotion and Affective UI

1 Introduction

Although we often tend to separate sense and sensibility, there is no such hard line in our brain between rationality and emotions. In normal human cognition, in fact, thinking and feeling are inseparable – emotions are often the product of our thoughts while our reflections are often the product of our affective states.

Therefore we need to give machines the ability to understand and express emotions, and hence allow them to make more human-like decisions and to improve the human-computer interaction.

A related paper on this work [1] explored how Sentic Computing, whose term derives from the Latin ‘sentire’ (root of words like sentiment and sensation), can significantly enhance computers’ emotional intelligence. In this paper we focus on blending a linguistic resource for the lexical representation of affective knowledge with a Common Sense knowledge base, and propose the use of principal component analysis (PCA) to reveal large-scale patterns in the data and thus predict new emotional knowledge.

2 Sentic Computing

In the past, emotion extraction from text involved the implementation of different techniques such as hand-crafted models [2], keyword spotting [3], fuzzy logic [4], lexical affinity [5] and statistical methods [6].

Unfortunately these methods turned out to be semantically weak since they mainly rely on parts of text in which emotional states are explicitly expressed i.e. verbs, adjectives and adverbs of emotions. In fact emotions are more often expressed implicitly through concepts with an affective valence such as ‘play a game’, ‘be laid off’ or ‘go on a first date’.

Sentic Computing overcomes this problem by using a Common Sense reasoning approach and a novel emotion categorization born from the idea that our mind consists of four independent emotional spheres, whose different levels of activation make up the total emotional state of the mind.

3 Common Sense Computing

Our approach for the affective categorization of text exploits recent developments in the field of Common Sense Computing [7]. In particular we rely on a semantic network, a process to reason on this knowledge base, a technique to perform categorization on it and a method to combine different datasets.

3.1 ConceptNet

When people communicate with each other, they rely on similar background knowledge e.g. 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 [8] represents the information in the Open Mind corpus as a directed graph in which the nodes are concepts and the labeled edges are assertions of Common Sense that interconnect them (Fig. 1).

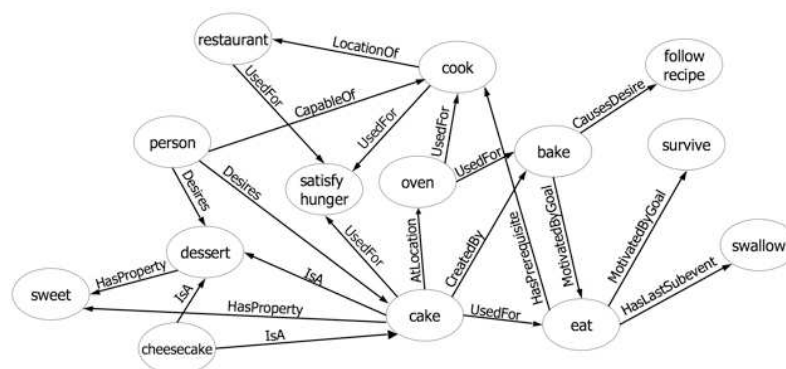


Fig. 1. A sketch of ConceptNet

3.2 AnalogySpace

AnalogySpace [9] is a way of representing a Common Sense knowledge base in a multidimensional vector space. In this process, we represent ConceptNet as a sparse matrix and use singular value decomposition (SVD) to reduce its dimensionality, capturing the most important correlations in it.

The principle of SVD is that any matrix A can be factored into an orthonormal matrix U , a diagonal matrix Σ , and an orthonormal matrix V^T , so that $A = U * \Sigma * V^T$. The singular values in Σ are ordered from largest to smallest, where the larger values correspond to the vectors in U and V that are more significant components of the initial A matrix.

When making use of the SVD results, we often discard all but the first k components, i.e. the principal components of A , which form a low-rank approximation of the original data.

This factorization allows the row space of A and the column space of A to be projected into a common space by the transformations U and V . We can think of these spaces as containing two types of objects, which we can represent as row and column vectors of A , which are related to each other by the values where they meet. After the SVD transformation, AnalogySpace represents both kinds of objects in the same space, where it can compare them to one another as k -dimensional vectors by means of dot products.

3.3 Ad-hoc Categories

The vectors that are compared in AnalogySpace do not have to correspond to existing concepts or properties in ConceptNet. In some cases, it is useful to construct a new vector in AnalogySpace, such as a synthetic concept made of a weighted sum of concepts. Such a vector can represent an ad-hoc category that is useful in an application.

The idea of ad-hoc categories [10] is similar to the idea of a ‘mini-document’ in latent semantic indexing (LSI), in which a collection of a few words represents, for example, a search query.

As an example, we can create a category of furniture from the linear combination of the concepts ‘chair’ + ‘table’ + ‘desk’ + ‘bed’ + ‘couch’. If we add up the AnalogySpace vectors for these concepts, we get a combined vector that represents the category as a whole. One simple way to use an ad-hoc category is to look for its strongest similarities to existing concepts, and thus discover more objects that seem to belong in the category.

3.4 Blending

Blending [11] 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.

4 Emotive Reasoning

The capability of perceiving and expressing emotions is a fundamental component in human experience, cognition, perception, learning and communication.

Today conventional computers lack this kind of skill. They just do what they are programmed to do without caring at all if the user is experiencing fascination or frustration. This is why nowadays we have plenty of programs that exceed the capabilities of world experts but are not able to do what even a puppy can do – understand us from an emotional point of view.

In order to better connect us, entertain us, help us work and keep us informed, computers must get to know how to recognize, understand and express emotions. Thus we can't exclude emotions in the development of intelligent systems: if we want computers to be really comprehensive and user friendly, we need to give them the capacity for emotive reasoning.

To do this we first use the blending technique to build an affective knowledge base of Common Sense, and then employ an ad-hoc categories approach, together with Sentic Computing, to perform the affective categorization.

4.1 WordNet-Affect

WordNet-Affect [12] is a linguistic resource for the lexical representation of affective knowledge, developed starting from WordNet [13].

The knowledge base 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.

WordNet-Affect 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.

4.2 Building AffectiveSpace

To build a suitable knowledge base for emotive reasoning, we apply the blending technique on ConceptNet and WordNet-Affect.

The first step to create a blend is to transform the input data so that it can all be represented in the same matrix. To do this we align the lemma forms of ConceptNet concepts with the lemma forms of the words in WordNet-Affect and

‘depression’. However similarity and analogy in AffectiveSpace do not depend on concepts’ absolute positions in the vector space but only on their positions relative to each other.

Concepts and emotions are represented by vectors of 50 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_{49} 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, have negative affective valence.

4.3 Emotion Categorization in AffectiveSpace

The aim of Sentic Computing is to develop emotion-sensitive systems for application in fields such as e-health, software agents, e-games, customer care, e-learning and e-tourism.

For this reason, in Sentic Computing, affective states are not categorized, as often happens in the field of emotion extraction, into basic emotional categories, but they are organized around four independent dimensions, Pleasantness, Attention, Sensitivity and Aptitude, to be able to understand how much:

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 keen on using the application

This model is a variant of Plutchik’s wheel of emotions [14] and constitutes an attempt to emulate Marvin Minsky’s conception of emotions.

Minsky sees the mind as made of thousands of different resources and believes that our emotional states result from turning some set of these resources on and turning another set of them off [15]. 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.

To be able to affectively analyze text according to this emotion categorization, we build an ad-hoc category for each sentic level i.e. each level of activation of the four affective dimensions (Table 1).

Each sentic category is the weighted sum of a set of concepts corresponding to the list of hypernyms and hyponyms of the sentic level, extracted from WordNet. The result of this operation is a new strong vector in AffectiveSpace which assumes all the peculiar features of the concept representing the sentic level.

So, for example, to find concepts semantically related to the sentic level ‘joy’, such as ‘birthday party’ or ‘have fun’, we first build an ad-hoc category by summing together the concepts ‘cheer’, ‘elation’, ‘amusement’, ‘exhilaration’, ‘exuberance’, ‘exultation’ and ‘joy’ itself, and then we look for concepts in AffectiveSpace that point in the same or similar directions as this vector i.e. concepts

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

Table 1. The four affective dimensions and their sentic levels

whose dot product with the ad-hoc category vector is bigger than a certain threshold (that is set by trial and error).

After building the 24 ad-hoc categories, we use k-means clustering to group together clouds of related concepts in AffectiveSpace. These affective clusters are then stored in a tensor, called Affective Similarity Map, containing the distances (the dot products) between concepts and the cluster means identified by the ad-hoc categories.

4.4 Parsing Emotions from Text

The technique we propose for gathering information from text and categorizing it according to Sentic Computing concepts, is termed the Sentic Extraction Process (Fig. 3), and it comprises three main components: a Natural Language Processing (NLP) module, which performs a first skim of the document, a Semantic Parser, whose aim is to extract concepts from the processed text, and eventually the Sentic Converter, a module for analyzing concepts' affective valence.



Fig. 3. The Sentic Extraction Process

The NLP module interprets all the affective valence indicators usually contained in text such as special punctuation, complete upper-case words, onomatopoeic repetitions, exclamation words, negations, degree adverbs and emoticons.

The Semantic Parser then deconstructs text into concepts and provides, for each of them, the relative frequency, valence and status i.e. the concept's occurrence in the text, its positive or negative connotation, and the degree of intensity with which the concept is expressed.

The Sentic Converter finally extracts, from the set of concepts so far obtained, a list of four-dimensional vectors, called ‘sentic vectors’, which contain the affective information of each concept in terms of Pleasantness, Attention, Sensitivity and Aptitude.

The conversion takes place thanks to the values stored in the Affective Similarity Map. For each concept provided by the Semantic Parser, we look up this map and, whenever a match is found, we extract the relative information, a ‘raw sentic vector’ containing 24 values, and encode it.

The codification process goes through a normalization step, the identification of the maximum affective similarity value for each affective dimension, and the addition of the corresponding sentic level value (Table 1).

Depending on the corresponding concept’s status, the sentic vector’s magnitude is then increased or decreased by 20% and, if the concept has a negative valence, the vector is replaced with its opposite.

5 Evaluation

To make a first evaluation of our system we considered a corpus of blogposts from LiveJournal (LJ), a virtual community where Internet users can keep a blog, journal or diary.

One of the interesting features of this website is that LJ bloggers, who number over 23 millions, are allowed to label their posts with a mood tag, by choosing from 130+ predefined moods or by creating custom mood themes. Since the indication of the affective status is optional, the mood-tagged posts are likely to reflect the true mood of the authors, and hence form a good test set for the evaluation of the Sentic Extraction Process.

However, since LJ mood themes do not perfectly match the sentic levels, we had to consider just a small set of moods i.e. ecstatic, cheerful, pensive, surprised, enraged, sad, angry, annoyed, scared and bored. Moreover we couldn’t consider non-affective webposts since untagged blog entries do not necessarily lack emotions.

Thus we selected approximately 5,000 webposts labeled with the above-mentioned tags and processed them through Sentic Computing, powered by AffectiveSpace. The posts’ average length was 242 words and the average number of concepts extracted per blog was 53.

After running the Sentic Extraction Process over the selected blogposts, the system showed a very high precision (73%) and significantly good recall and F-measure rates (65% and 68% respectively). The affective categorization was particularly good, in terms of precision, for positive moods, probably due to the fact that the Open Mind corpus contains more concepts expressing positive emotions than concepts related to sad affective states.

Despite the non-specificity of the test, the results were quite encouraging and left the doors open for future evaluations in which, for example, LJ moods could be fully mapped into the sentic levels or in which the Sentic Extraction Process

could be tested with a sample of posts manually classified by users in terms of Pleasantness, Attention, Sensitivity and Aptitude.

6 Future Work

We plan to explore different dimensionality reduction techniques e.g. independent component analysis (ICA) in place of PCA for building AffectiveSpace, to better find large-scale patterns in the knowledge base, smooth over noise and predict new affective information.

We also aim to improve the Sentic Extraction Process by adding more NLP functionalities and by refining the granularity of the Affective Similarity Map i.e. the way the affectively related concepts are selected in AffectiveSpace.

We then plan to embed the AffectiveSpace process in some marketable products e.g. a customer care tool for the evaluation of users' level of satisfaction in enterprise 2.0 and e-tourism applications or an affective analysis tool, to be embedded into a health care expert system, to assess patients' attitudes and thus provide better prescriptions.

Since instant messaging clients, which are more and more frequently and widely used for interpersonal communication, lack the richness of face-to-face conversations, we are also thinking about exploiting AffectiveSpace to develop a MSN or Skype add-on. Thanks to emotive reasoning the chat background or the style font and color could change according to the current emotional state of the user or a cartoon avatar could instantly change its expression according to the last emotion detected.

Finally a similar approach could be employed in the fields of software agents, e-games and e-learning for the development of embodied conversational agents: the emotions extracted through the Sentic Extraction Process could be used as inputs for a facial action coding system (FACS) to better respond to the user's emotional changes.

7 Conclusions

In this paper we explored a new method to improve machines' emotional intelligence. By blending Common Sense and affective knowledge and applying SVD on the resulting matrix, we built a vector space in which concepts can be affectively classified according to some ad-hoc categories.

AffectiveSpace's capability for emotive reasoning makes it an attractive tool for the development of emotion-sensitive systems in fields such as e-health, software agents, e-games, customer care, e-learning and e-tourism.

Next-generation intelligent systems must recognize, understand and express emotions to make human-like decisions and improve the human-computer interaction, because the question is not whether intelligent machines can have emotions, but whether machines can be intelligent without any emotions [16].

References

1. Cambria, E., Hussain, A., Havasi, C., Eckl, C.: Sentic Computing: Exploitation of Common Sense for the Development of Emotion-Sensitive Systems. Submitted to: Esposito, A. et al. (eds.) *Development of Multimodal Interfaces: Active Listening and Synchrony*. LNCS, Springer-Verlag (2009)
2. Dyer, M.: *Emotions and Their Computations: Three Computer Models*. *Cognition and Emotion*. 1, 323–347 (1987)
3. Bradley, M., Lang, P.: *Affective Norms for English Words (ANEW): Instruction manual and affective ratings*. Technical Report C-1, University of Florida (1999)
4. Subasic, P., Huettner, A.: *Affect Analysis of Text Using Fuzzy Semantic Typing*. *IEEE Transactions On Fuzzy Systems*. 9, 483–496 (2000)
5. Valitutti, A., Strapparava, C., Stock, O.: *Developing affective lexical resources*. *PsychNology Journal*. 2, 61–83 (2004)
6. Goertzel, B., Silverman, K., Hartley, C., Bugaj, S., Ross, M.: *The Baby Webmind Project*. In: AISB (2000)
7. Cambria, E., Hussain, A., Havasi, C., Eckl, C.: *Common Sense Computing: From the Society of Mind to Digital Intuition and Beyond*. In: Fierrez, J. et al. (eds.) *BioID_MultiComm 2009*. LNCS, vol. 5707, pp. 260–267. Springer-Verlag, Berlin Heidelberg (2009)
8. Havasi, C., Speer, R., Alonso, J.: *ConceptNet 3: a Flexible, Multilingual Semantic Network for Common Sense Knowledge*. In: RANLP (2007)
9. Speer, R., Havasi, C., Lieberman, H.: *AnalogySpace: Reducing the Dimensionality of Common Sense Knowledge*. In: AAAI (2008)
10. Havasi, C., Speer, R., Pustejovsky, J.: *Automatically Suggesting Semantic Structure for a Generative Lexicon Ontology*. In: *Generative Lexicon* (2009)
11. Havasi, C., Speer, R., Pustejovsky, J., Lieberman, H.: *Digital Intuition: Applying Common Sense Using Dimensionality Reduction*. *IEEE Intelligent Systems*. 24, 24–35 (2009)
12. Strapparava, C., Valitutti, A.: *WordNet-Affect: an Affective Extension of WordNet*. In: LREC (2004)
13. Fellbaum, C.: *WordNet: An Electronic Lexical Database*. MIT Press, Cambridge (1998)
14. Plutchik, R.: *The Nature of Emotions*. *American Scientist*. 89, 344 (2001)
15. Minsky, M.: *The Emotion Machine*. Simon and Schuster, New York (2006)
16. Minsky, M.: *The Society of Mind*. Simon and Schuster, New York (1986)