Sentic Computing: Exploitation of Common Sense for the Development of Emotion-Sensitive Systems

Erik Cambria[1], Amir Hussain[1], Catherine Havasi[2], and Chris Eckl[3]

[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. Emotions are a fundamental component in human experience, cognition, perception, learning and communication. In this paper we explore how the use of Common Sense Computing can significantly enhance computers' emotional intelligence i.e. their capability of perceiving and expressing emotions, to allow machines to make more human-like decisions and improve the human-computer interaction.

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

1 Introduction

Today text is one of the most important modalities for affective analysis and generation because the bulk of computer user interfaces are still text-based. But the inference of emotions from text has always been a difficult task since people usually use non-verbal cues, such as facial expression, vocal inflections and body movement, to reveal, either intentionally or unintentionally, the emotional information. Previous attempts to perform this task mainly relied on statistical methods which have shown to have strong limitations.

We introduce a new paradigm, which we call Sentic Computing, in which we use a novel emotion representation and a Common Sense [1] based approach to infer affective states from short texts over the web.

2 The Importance of Emotions

In normal human cognition, thinking and feeling are mutually present: our emotions are often the product of our thoughts and our reflections are frequently the product of our sentiments. Emotions in fact are intrinsically part of our mental activity and play a key role in decision-making processes: they are special states shaped by natural selection to balance the reaction of our organism to particular situations e.g. anger evolved for reaction, fear evolved for protection and affection evolved for reproduction.

A. Esposito et al. (Eds.): COST 2102 Int. Training School 2009, LNCS 5967, pp. 148–156, 2010.

For these reasons we can't prescind from emotions in the development of intelligent systems: if we want computers to be really intelligent, not just have the veneer of intelligence, we need to give them the ability to recognize, understand and express emotions.

3 The Emotional Web Era

Differently from early web development, retroactively labeled Web 1.0, today Internet is a dynamic being in which information is no more the core – the user is now at the center of it.

This sort of Copernican revolution brought us to the Web 2.0 era, in which the first simple websites evolved to become more and more interactive, from static to dynamically generated, from handcrafted to CMS-driven, from purely informative to more and more social. The passage from a read-only to a read-write web in fact made web users more inclined to express their emotions through blogs, fora, wikis, feeds and chats.

The efforts to understand this new cultural and social phenomenon gave birth to Web Science [2], a new discipline which brings computer scientists and social scientists together across the disciplinary divide, to explore the development of the Web across different areas of everyday life and technological development.

4 Sentic Computing

Within the field of Web Science we introduce a new paradigm, which we call Sentic Computing, whose aim is to use Common Sense to better recognize, interpret and process human emotions in webposts i.e. short texts over the web such as blog posts, forum entries, RSS feeds, tweets and instant messages. The term 'sentic' derives from the Latin 'sentire', the root of words like sentiment and sensation, and it was first adopted in 1977 by Manfred Clynes [3], who discovered that when people have emotional experience, their nervous system always responds in a characteristic way which is measurable.

Sentic Computing is part of the efforts in the fields of computer science, psychology, linguistics, sociology and cognitive science, to develop a kind of computing that relates to, arises from, or influences emotions [4].

In the past, emotion extraction from text witnessed the implementation of different techniques: hand-crafted models [5], keyword spotting e.g. ANEW [6] or LIWC [7] which rely on an often-used source of affect words, lexical affinity [8] and statistical methods e.g. LSA (Latent Semantic Analysis) which has been frequently used by researchers on projects such as Webmind [9]. The problem with these kinds of methods consists in the fact that they mainly rely on parts of text in which emotional states are explicitly expressed: the so called 'attitudinal inscriptions' of the appraisal theory [10].

They are verbs of emotion e.g. to love/to hate, to frighten/to reassure, to interest/to bore, to enrage/to placate, adverbs of emotion e.g. happily/sadly and adjectives of emotion e.g. happy/sad, worried/confident, angry/pleased, keen/uninterested. But more often emotions are expressed implicitly through concepts with an affective valence such as 'playing a game', 'being laid off' or 'going on a first date'. To extract from text these latent emotional states, termed 'attitudinal tokens' within the appraisal framework, we exploit the recent developments in the field of Common Sense Computing.

4.1 A Common Sense Based Approach

Statistical affective classification using statistical learning models generally requires large inputs and thus cannot appraise texts with satisfactory granularity. Our approach allows to affectively classify webposts not only on the page or paragraph-level but even on sentence-level.

Sentic Computing develops an approach previously adopted by Hugo Liu et al. [11] in which a Common Sense knowledge base was exploited to try to extract affective information from emails using the standard notion of basic emotions provided by Ekman. We now use a much richer semantic network, ConceptNet [12] version 4, with almost 10,000 concepts and a set of 72,000+ features extracted from the Open Mind corpus, and the power of cumulative analogy provided by AnalogySpace [13], a process which reveals large-scale patterns in the data, smooths over noise, and predicts new knowledge.

4.2 A Novel Emotion Representation

Our aim is to develop emotion-sensitive systems in fields such as e-health, software agents, e-games, customer care, e-learning and e-tourism.

For this reason, instead of trying to categorize webposts into basic emotional categories, we are interested in understanding 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

Thus we adopt a new emotion representation where the user's affective states are organized around four independent dimensions: Pleasantness, Attention, Sensitivity and Aptitude. 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 emulate this process we organize the different mental resources around four affective dimensions, each of them representing an independent emotional sphere, whose different levels of activation, which we call 'sentic levels', give the total emotional state of the mind (see Table 1).

Table 1. The four affective dimensions and their sentic levels

	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

5 The Sentics Extraction Process

The Sentics Extraction Process goes through a Natural Language Processing module, which performs a first skim of the webpost, 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.

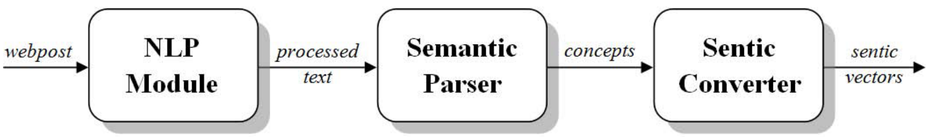

Fig. 1. The Sentics Extraction Process

5.1 The NLP Module

The module interprets all the affective valence indicators usually contained in webposts such as special punctuation e.g. suspension points to express impatience, question marks to convey doubt or exclamation marks to express irony, complete upper-case words, onomatopoeic repetitions e.g. repetition of letters to emulate a shout or a sensation of surprise and ditto syllables to reproduce a laughter or a moan, exclamation words, and emoticons.

Emoticons in particular are a very good indicator of a webpost's affective valence: they can convey information about the user's affective state either directly e.g. ':-)' = 'joy' or indirectly e.g. '|-o' = 'yawning' which implies the state of boredom. We currently use a database of 300+ smileys, either in Western and East Asian style, which covers almost the totality of the emoticons used today. The module also detects negations and degree adverbs, to let the Semantic Parser correctly weigh the expressed concepts, and finally stems the text i.e. removes function words, pronouns and inflections.

5.2 The Semantic Parser

The aim of this module is to deconstruct text into concepts. This is not an easy task since the semantic analysis of a text is far more difficult than the lexical one and it's still an open problem in many fields e.g. the semantic web, where the task is currently being tackled by using dereferenceable URIs, RDF statements and web ontologies.

To scan the processed text coming out of the NLP module, the Semantic Parser looks for matches in the Common Sense knowledge base combining the closer stemmed words: this way the parser is able to collect not just the atomic concepts such as 'fun' or 'hug' but also compound concepts e.g. 'eat spaghetti', 'play hockey', 'say hello' or 'blow out candle'.

The output of the Semantic Parser is the set of concepts retrieved in the web-post with their 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.

5.3 The Sentic Converter

This module converts the set of concepts given by the Semantic Parser into a list of four-dimensional vectors, which we call 'sentic vectors'.

The 'sentic vector' of a concept is the tuple [Pleasantness, Attention, Sensitivity, Aptitude] whose values are floating point numbers in the range (-3,+3). To build it we rely on an Affective Similarity Map containing all the relevant 'affective concepts' mapped into the four affective dimensions. By 'affective concepts' we mean the concepts which, in AnalogySpace, are neighbours to the 'sentic centroids' i.e. the concepts representing the 24 sentic levels such as joy, surprise and anger (see Table 1).

AnalogySpace is a vector space representation of Common Sense knowledge describing the analogical closure of a semantic network: by selecting the closest vectors to the 24 centroids we practically clusterize the vector space using a K-means approach. These clusters represent groups of concepts semantically related to the concepts embodying the sentic levels, while the rest of the space is labeled as 'limbo' i.e. absence of affective valence.

This way we obtain an Affective Similarity Map whose rows represent the 'affective concepts', whose columns stand for the 'sentic centroids' and whose entries are the distances, i.e. the dot products, between them. For each concept provided by the Semantic Parser, we look up in 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 (see Table 1).

Depending on the corresponding concept's status, the sentic vector's magnitude is then increased or decreased of 20% and, in case the concept has a negative valence, the vector is switched with its opposite.

6 Interpreting the Sentic Vectors

The sentic vectors represent a webpost's affective valence rating in terms of Pleasantness, Attention, Sensitivity and Aptitude: depending on our needs we can use this information to quantify user's emotive load or visualize it by plotting the vectors in the space specified by the four affective dimensions.

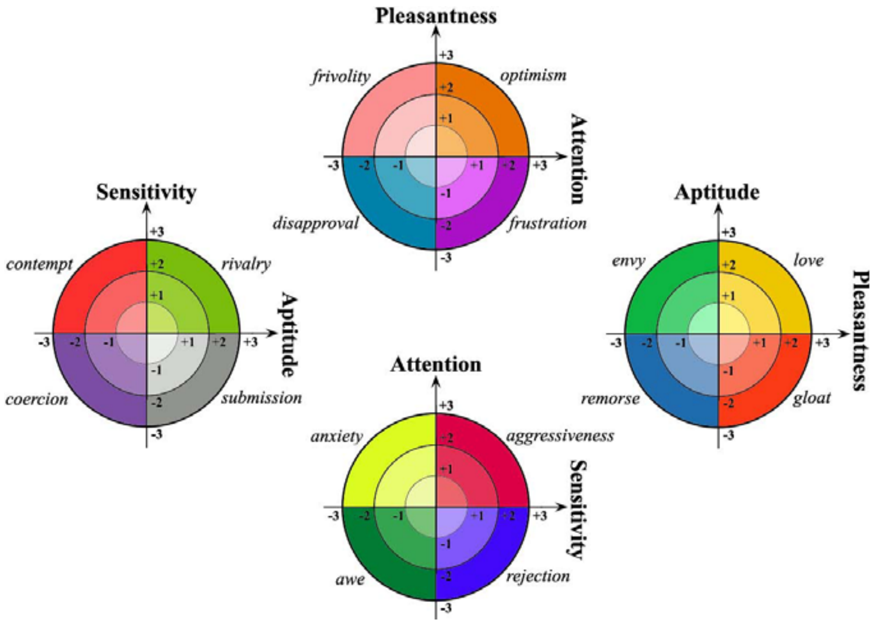

Fig. 2. Particular planes in the sentic vector space

By setting equal to zero the affective dimensions pairwise we can also individuate particular planes representing advanced emotions: the projection of the sentic vectors onto these particular planes gives the webpost's affective valence in terms of compound emotional states such as aggressiveness, which is given by the sum of positive values of Attention and Sensitivity, or remorse, which is given by the sum of negative values of Aptitude and Pleasantness (see Fig. 2).

By averaging the sentic vectors of a sentence we can then summarize its affective valence in a single vector and perform a study of the user's emotional state over time for example in a chat room or in a microblogging application.

This process can be iterated to affectively classify more complex webposts such as blog posts, RSS feeds or forum entries, and evaluate the user's changes of attitudes within a specific web application.

At a higher level we can even perform a user categorization by representing each user with a sentic vector, built with the information gathered from all his/her webposts we have at our disposal, and by plotting it in the vector space to make comparisons with other users and try to affectively classify them.

7 Evaluation

To make a first approximate evaluation of our system we considered a corpus of blogposts from the LiveJournal community, annotated with happy and sad moods. Since the indication of the mood is optional when posting on LiveJournal, the posts we used are likely to reflect the true mood of the authors, and hence form a good test set for the evaluation of the Sentics Extraction Process.

However, since the blogposts are classified just in terms of positive and negative affective valence, we had to take in consideration only the first entry of the sentic vectors i.e. the Pleasantness dimension.

We selected 500 blogposts labeled as happy and as many labeled as sad and processed them through Sentic Computing. To be able to reliably classify the posts, we considered just the absolute average values of Pleasantness superior to 5% and left the rest unlabeled. After running the Sentics Extraction Process over the one thousand blogposts, we obtained almost 70% hits with about 20% false positives and 10% false negatives.

Despite the non-specificity of the test, the results were quite encouraging and left the doors open for future evaluations in which, for example, the Sentics Extraction Process may be tested with a bunch of posts manually classified by users in terms of Pleasantness, Attention, Sensitivity and Aptitude.

To improve the process we plan to enhance the NLP Module's sensibility by making it able to intepret more affective valence indicators in webposts such as superlatives, double negatives and affectively relevant rhetorical devices e.g. climax, anaphora, epistrophe, commoratio, palilogy, aposiopesis.

We also want to refine the Semantic Parser by exploring different Python procedures for text scanning and retrieval to perform a faster and more accurate webpost semantic analysis.

We finally plan to improve the Sentic Converter by blending ConceptNet with a linguistic resource for the lexical representation of affective knowledge.

8 Future Work

General Common Sense knowledge is very useful to discover how concepts are affectively related but we need more specific information to improve the Sentics Extraction Process. To this end we plan to start soon the implementation of a new web interface for the affective Common Sense knowledge acquisition from general public: Open Mind Common Sentics.

The idea is to develop a GUI in which users are asked to give an affective interpretation to random assertions extracted from the Open Mind corpus. After selecting a topic, the user would be given an assertion and asked to type the kind of sensations it arises or simply select one of the available emoticons to intuitively classify the assertion from an emotional point of view.

Open Mind Common Sentics will be complementary to Open Mind Commons, the current interface for collecting Common Sense knowledge from users over the web. In ConceptNet concepts are linked by 24 relations such as 'IsA', 'AtLocation' and 'MadeOf' which respectively answer the questions 'What kind of thing

Fig. 3. Open Mind Common Sentics logo

is it?', 'Where would you find it?' and 'What is it made of?'. We plan to add to ConceptNet the information gathered through Open Mind Common Sentics by inserting a new relation, 'ArisesEmotion', which answer the question 'What kind of emotion it arises?'.

Recruiting new customers is often more expensive than retaining the existing ones. For this reason it's very important to have tools to measure how much your clients are happy with your product or service.

We plan to design a customer care tool which exploits Sentic Computing to evaluate users' level of satisfaction for enterprise 2.0 or e-tourism content management systems. The tool will also be embedded in a health care expert system for clinical decision support to gather patients' attitudes and thus provide better prescriptions.

Today instant messaging clients, which are increasingly used for interpersonal communication, lack the richness of face-to-face conversations. We plan the development of a MSN or Skype add-on in which the chat background or the style font and color change according to the current emotional state of the user and in which a cartoon avatar instantly changes its expression according to the last emotion detected.

A similar approach will be finally employed in fields such as software agents, e-games and e-learning for the development of embodied conversational agents: the sentic vectors will be used as inputs for a facial action coding system to better respond to the user's emotional changes.

9 Conclusions

In this paper we showed how Sentic Computing can help the development of emotion-sensitive systems in fields such as e-health, software agents, e-games, customer care, e-learning and e-tourism.

The approach hereby described is very flexible and lends itself to be used for various purposes: we opted for a four-dimensional vector representation because we were interested in four particular affective dimensions but, depending on the kind of analysis performed and on the emotion representation used, the technique can be easily extended.

The importance of Sentic Computing, anyway, consists not only in introducing a new method for affectively analyze text but also in highlighting the importance of emotions for the development of next-generation systems because, as Marvin Minsky would say, the question is not whether intelligent machines can have emotions, but whether machines can be intelligent without any emotions.

References

1. Cambria, E., Hussain, A., Havasi, C., Eckl, C.: Common Sense computing: From the society of mind to digital intuition and beyond. In: Fierrez, J., Ortega-Garcia, J., Esposito, A., Drygajlo, A., Faundez-Zanuy, M. (eds.) Biometric ID Management and Multimodal Communication. LNCS, vol. 5707, pp. 252–259. Springer, Heidelberg (2009), http://cs.stir.ac.uk/~eca/commansense
2. Hendler, J., Shadbolt, N., Hall, W., Berners-Lee, T., Weitzner, D.: Web Science: an Interdisciplinary Approach to Understanding the World Wide Web. Communications of the ACM 51(7), 60–69 (2008)
3. Clynes, M.: Sentics: The Touch of the Emotions, Doubleday, New York (1977)
4. Picard, R.: Affective Computing. MIT Press, Cambridge (1997)
5. Dyer, M.: Emotions and Their Computations: Three Computer Models. Cognition and Emotion 1(3), 323–347 (1987)
6. Bradley, M., Lang, P.: Affective Norms for English Words (ANEW): Instruction manual and affective ratings. Technical Report C-1, University of Florida (1999)
7. Pennebaker, J., Francis, M., Booth, R.: Linguistic Inquiry and Word Count (LIWC). Erlbaum Publishers, Mahwah (2001)
8. Valitutti, A., Strapparava, C., Stock, O.: Developing Affective Lexical Resources. PsychNology Journal 2(1), 61–83 (2004)
9. Goertzel, B., Silverman, K., Hartley, C., Bugaj, S., Ross, M.: The Baby Webmind Project. In: AISB 2000 (2000)
10. Martin, J., White, P.: The Language of Evaluation: Appraisal in English. Palgrave Macmillan, Basingstoke (2005)
11. Liu, H., Lieberman, H., Selker, T.: A Model of Textual Affect Sensing using Real-World Knowledge. In: IUI 2003 (2003)
12. Havasi, C., Speer, R., Alonso, J.: ConceptNet 3: a Flexible, Multilingual Semantic Network for Common Sense Knowledge. In: RANLP 2007 (2007)
13. Speer, R., Havasi, C., Lieberman, H.: Analogy Space: Reducing the Dimensionality of Common Sense Knowledge. In: AAAI 2008 (2008)
14. Plutchik, R.: The Nature of Emotions. American Scientist 89(4), 344–350 (2001)
15. Minsky, M.: The Emotion Machine. Simon and Schuster, New York (2006)