

Detecting Implicit Expressions of Sentiment in Text Based on Commonsense Knowledge

Alexandra Balahur, Jesús M. Hermida, Andrés Montoyo

Department of Software and Computing Systems

University of Alicante

Apartado de correos 99, E-03080 Alicante, Spain

{abalahur, jhermida, montoyo}@dlsi.ua.es

Abstract

Sentiment analysis is one of the recent, highly dynamic fields in Natural Language Processing. Most existing approaches are based on word-level analysis of texts and are able to detect only explicit expressions of sentiment. In this paper, we present an approach towards automatically detecting emotions (as underlying components of sentiment) from contexts in which no clues of sentiment appear, based on commonsense knowledge. The resource we built towards this aim – EmotiNet - is a knowledge base of concepts with associated affective value. Preliminary evaluations show that this approach is appropriate for the task of implicit emotion detection, thus improving the performance of sentiment detection and classification in text.

1 Introduction

Research in affect has a long established tradition in many sciences - linguistics, psychology, socio-psychology, cognitive science, pragmatics, marketing or communication science. Recently, many closely related subtasks were developed also in the field of Natural Language Processing (NLP), such as emotion detection, subjectivity analysis, opinion mining to sentiment analysis, attitude and

appraisal analysis or review mining (Pang and Lee, 2008).

Among these tasks, sentiment analysis aims at detecting the expressions of sentiment in text and subsequently classify them, according to their polarity (semantic orientation) among different categories (usually, among positive and negative). The problem is defined by Pang and Lee (2008) as “the binary classification task of labeling an opinionated document as expressing either an overall positive or an overall negative.” (Pang and Lee, 2008)

According to the Webster dictionary (<http://www.merriam-webster.com/>), sentiment suggests a settled opinion reflective of one’s feelings, where the term feeling is defined as the conscious subjective experience of emotion. (Van den Bos, 2006), “a single component of emotion, denoting the subjective experience process” (Scherer, 2005). Most of the research performed in the field of sentiment analysis has aimed at detecting explicit expressions of sentiment (i.e. situations where specific words or word combinations are found in texts). Nevertheless, the expression of emotion is most of the times not achieved through the use of emotion-bearing words (Pennebaker et al., 2003), but indirectly, by presenting situations that based on commonsense knowledge can be interpreted in an affective manner (Balahur and Montoyo, 2008; Balahur and Steinberger, 2009).

In this paper, we present a method to build a commonsense knowledge base (EmotiNet) representing situations that trigger emotions. We demonstrate that by using this resource, we are

able to detect emotion from textual contexts in which no explicit mention of affect is present.

2 State of the Art

In Artificial Intelligence (AI), the term affective computing was first introduced by Picard (1995). Previous approaches to spot affect in text include the use of models simulating human reactions according to their needs and desires (Dyer, 1987), fuzzy logic (Subasic and Huettner, 2000), lexical affinity based on similarity of contexts – the basis for the construction of WordNet Affect (Strapparava and Valitutti, 2004) or SentiWordNet (Esuli and Sebastiani, 2005), detection of affective keywords (Riloff et al., 2003) and machine learning using term frequency (Pang et al., 2002; Riloff and Wiebe, 2003), or term discrimination (Danisman and Alpkocak, 2008). Other proposed methods include the creation of syntactic patterns and rules for cause-effect modeling (Mei Lee et al., 2009). Significantly different proposals for emotion detection in text are given in the work by (Liu et al, 2003) and the recently proposed framework of sentic computing (Cambria et al., 2009), whose scope is to model affective reaction based on commonsense knowledge. For a survey on the affect models and their affective computing applications, see (Calvo and D’Mello, 2010).

3 Motivation and Contribution

The tasks of emotion detection and sentiment analysis have been approached by a large volume of research in NLP . Nevertheless, most of this research has concentrated on developing methods for detecting only explicit mentions of sentiment in text. Therefore, sentences such as “I’m going to a party”, which express an underlying emotion, cannot be classified by most of the existing approaches. A method to overcome this issue is proposed in by *sentic* computing (Cambria et al., 2009) and by (Liu et al, 2003), whose main idea is acquiring knowledge on the emotional effect of different concepts. In this manner, the system would know that “going to a party” is something that produces “joy”. However, more complex contexts, such as “I’m going to a party, although I should study for my exam.”, where the emotion expressed is most probably “guilt”, cannot be

correctly detected and classified by present systems.

In the light of these considerations, our contribution relies in proposing and implementing a framework for modeling affect based on the appraisal theories, which can support the automatic processing of texts to extract:

- The components of the situation presented (which we denote by “action chains”) and their relation (temporal, causal etc.)
- The elements on which the appraisal is done in each action of the chain (agent, action, object);
- The appraisal criteria that can automatically be determined from the text (modifiers of the action, actor, object in each action chain);

4 Modeling Affective Reaction Using Commonsense Knowledge

Our main idea is that emotion can be expressed in text by presenting a sequence of actions (situations in which different concepts appear), which, based on commonsense knowledge and previous experiences, trigger an emotional reaction. This idea is linked to the Appraisal Theories, which claim that emotions are elicited and differentiated on the basis of the subjective evaluation of the personal significance of a situation, object or event (De Rivera, 1977; Frijda, 1986; Johnson-Laird and Oatley, 1989 – among others). Viewed in a simpler manner, a situation is presented as a chain of actions, each with an actor and an object; the appraisal depends on the temporal and causal relationship between them, on the characteristics of the actors involved in the action and on the object of the action.

Given this insight, the general idea behind our approach is to model situations as chains of actions and their corresponding emotional effect using an ontological representation. According to the definition provided by Studer et al. (1998), an ontology captures knowledge shared by a community that can be easily sharable with other communities. These two characteristics are especially relevant if we want the recall of our approach to be increased. Knowledge managed in our approach has to be shared by a large community and it also needs to be fed by heterogeneous sources of common knowledge to

avoid uncertainties. However, specific assertions can be introduced to account for the specificities of individuals or contexts. In this manner, we can model the interaction of different events in the context in which they take place.

5 Building a Knowledge Base for Detecting Implicit Expressions of Emotion

In order to build a resource that is capable of capturing emotional reaction to real-world situations in which commonsense knowledge plays a significant role in the affective interpretation, we aim at representing chains of actions and their corresponding emotional labels from several situations in such a way that we will be able to extract general patterns of appraisal. Our approach defines an action chain as a sequence of action links, or simply actions that trigger an emotion on an actor. Each specific action link can be described with a tuple (actor, action type, patient, emotional reaction).

In order to manage and store action chains, the approach we propose defines a new knowledge base, called EmotiNet, which aims to be a resource for detecting emotions in text, and a (semi)automatic, iterative process to build it, which is based on existing knowledge from different sources. This process extracts the action chains from a set of documents and adds them to the KB. Specifically, EmotiNet was built by following the next steps:

1. The design of an ontology, which contains the definitions of the main concepts of the domain.
2. The extension and population of this ontology using the situations stored in the ISEAR International Survey of Emotional Antecedents and Reactions (ISEAR, <http://www.unige.ch/fapse/emotion/databanks/isear.html>) – (Scherer and Wallbott, 1997) database.
3. The expansion of the ontology using existing commonsense knowledge bases – ConceptNet (Liu and Singh, 2004) and other resources – VerbOcean (Chklovski and Pantel, 2004).

5.1 Design of the Ontology

As mentioned before, the process of building the core of the EmotiNet knowledge base (KB) of action chains started with the design of the core ontology, whose design process was specifically divided in three stages:

1. Establishing the scope and purpose of the ontology. The EmotiNet ontology needs to capture and manage knowledge from three domains: kinship membership, emotions (and their relations) and actions (characteristics and relations between them).

2. Reusing knowledge from existing ontologies. In a second stage, we searched for other ontologies on the Web containing concepts related to the knowledge cores we specified. At the end of the process, we located two ontologies that are reused in our ontological representation: the ReiAction ontology (www.cs.umbc.edu/~lkagal1/rei/ontologies/ReiAction.owl), which represents actions between entities in a general manner, and the family ontology (www.dlsi.ua.es/~jesusmhc/emotinet/family.owl), which contains knowledge about family members and the relations between them.




Figure 1. Partial RDF graph of the Emotion Ontology.

3. Creating the final knowledge core from the ontologies imported. This third stage involved the design of the last remaining core, i.e. emotion, and the combination of the different knowledge sources into a single ontology: EmotiNet. In order to describe the emotions and the way they relate and compose, we employ Robert Plutchik’s wheel of emotion (Plutchik, 2001) and Parrot’s tree-

structured list of emotions (Parrot, 2001). These models contain an explicit modeling of the relations between the different emotions. At the end of the design process, the knowledge core included different types of relations between emotions and a collection of specific instances of emotion (e.g. anger, joy). In the last step, these three cores were combined using new classes and relations between the existing members of these ontologies (Fig. 2).

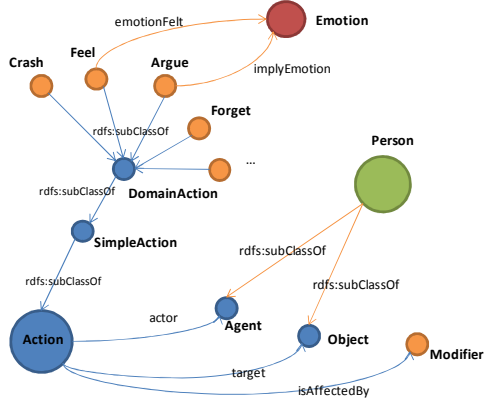


Figure 2. Main concepts of EmotiNet.

5.2 Extension and Population of the Ontology

In order to have a homogenous starting base, we selected from the 7667 examples in the ISEAR database only the 1081 cases that contained descriptions of situations between family members. Subsequently, the examples were POS-tagged using *TreeTagger*. Within each emotion class, we then computed the similarity of the examples with one another, using the implementation of the Lesk distance in Ted Pedersen’s Similarity Package. This score was used to split the examples in each emotion class into six clusters using the Simple K-Means implementation in Weka. The idea behind this approach, confirmed by the output of the clusters, was to group examples that are similar, in vocabulary and structure. From this collection, we manually selected a subset of 175 documents with 25 expressions related to each of the emotions: anger, disgust, guilt, fear, sadness, joy and shame.

The next step was to extract the actions chains described in each of the examples. For this, we employed *Semrol*, the semantic role labeling (SRL) system introduced by Moreda et al. (2007). For the

core of knowledge in the EmotiNet KB, we need 100% accurate information. Therefore, we manually extract the agent, the verb and the patient (the surface object of the verb) from the output of *Semrol*. For example, if we use the input sentence “I’m going to a family party because my mother obliges me to”, the system extracts two triples with the main actors of the sentences: (I, go, family party) and (mother, oblige, me), related by the causal adverb “because”.

Further on, we resolve the anaphoric expressions automatically, using a heuristic selection of the family member mentioned in the text that is closest to the anaphoric reference and whose properties (gender, number) are compatible with the ones of the reference. The replacement of the references to the speaker, e.g. ‘I’, ‘me’, ‘myself’, is resolved by taking into consideration the entities mentioned in the sentence. In case of ambiguity, we choose the youngest, female member. Following the last example, the subject of the action would be assigned to the daughter of the family and the triples would be updated: (daughter, go, family_party) and (mother, oblige, daughter). Finally, the action links (triplets) are grouped and sorted in action chains. This process of sorting is determined by the adverbial expressions that appear within the sentence, which actually specify the position of each action on a temporal line (e.g. “although” “because”, “when”). We defined pattern rules according to which the actions introduced by these modifiers happen prior to or after the current context.

Using our combined emotion model as a reference, we manually assigned one of the seven most basic emotions, i.e. anger, fear, disgust, shame, sadness, joy or guilt, or the neutral value to all the action links obtained, thus generating 4-tuples (subject, action, object, emotion), e.g. (daughter, go, family party, neutral) or (mother, oblige, daughter, disgust).

Once we carried out these processes on the chosen documents, we obtained 175 action chains (ordered lists of tuples). In order to be included in the EmotiNet knowledge base, all their action links needed to be mapped to existing concepts or instances within the KB. When these did not exist, they were added to it. We would like to highlight that in EmotiNet, each tuple (actor, action, patient, emotion) extracted has its own representation as an instance of the subclasses of *Action*. Each in-stance

of *Action* is related to an instance of the class *Feel*, which represents the emotion felt in this action. Subsequently, these instances (action links) were grouped in sequences of actions (class *Sequence*) ended by an instance of the class *Feel*, which determine the final emotion felt by the main actor(s) of the chain.

In our example, we created two new classes *Go* and *Oblige* (subclasses of *DomainAction*) and two new instances of them: instance *act1* (“Go”, “daughter”, “family_party”, “Neutral”); and instance *act2* (“Oblige”, “mother”, “daughter”, “Angry”). The last action link already existed within EmotiNet from another chain so we reused it: instance *act3* (“Feel”, “daughter”, “anger”). The next step consisted in grouping these instances into sequences by means of instances of the class *Sequence*, which is a subclass of *Action* that can establish the temporal order between two actions (which one occurred first). Fig. 3 shows an example of a RDF graph with the action chain of our example. We used Jena (<http://jena.sourceforge.net/>) and MySQL for the management and storage of EmotiNet on a database.




Figure 3. RDF graph of an action chain.

5.3 Ontology Expansion

In order to extend the coverage of the resource, we expanded the ontology with the actions and relations from VerbOcean. This process is essential for EmotiNet, since it adds new types of action and relations between actions, which might not have been analyzed before, thus reducing the degree of dependency between the resource and the initial set of examples. In particular, 299 new actions were automatically included as subclasses of

DomainAction, which were directly related to any of the actions of our ontology through three new relations: *can-result-in*, *happens-before* and *similar*.

6 Experiments and Evaluation

The evaluation of our approach consists in testing if by employing the model we built and the knowledge contained in the core of EmotiNet (which we denote by “knowledge sets”), we are able to detect the emotion expressed in new examples pertaining to the categories in ISEAR. Therefore, we use a test set (marked with B) that contains 895 examples (ISEAR phrases corresponding to the seven emotions modeled, from which core examples were removed).

In order to assess the system performance on the two test sets, we followed the same process we used for building the core of EmotiNet, with the exception that the manual modeling of examples into tuples was replaced with the automatic extraction of (actor, verb, patient) triples from the output given by Semrol. Subsequently, we eliminated the stopwords in the phrases contained in these three roles and performed a simple coreference resolution. Next, we ordered the actions presented in the phrase, using the adverbs that connect the sentences, through the use of patterns (temporal, causal etc.). The resulted action chains for each of the examples in the two test sets will be used in carrying different experiments:

(1). In the first approach, for each of the situations in the test sets (represented now as action chains), we search the EmotiNet KB to encounter the sequences in which these actions in the chains are involved and their corresponding subjects. As a result of the search process, we obtain the emotion label corresponding to the new situation and the subject of the emotion based on a weighting function. This function takes into consideration the number of actions and the position in which they appear in the sequence contained in EmotiNet. The issue in this first approach is that many of the examples cannot be classified, as the knowledge they contain is not present in the ontology.

(2). A subsequent approach aimed at surpassing the issues raised by the missing knowledge in EmotiNet. In a first approximation, we aimed at introducing extra knowledge from VerbOcean, by adding the verbs that were similar to the ones in

the core examples (represented in VerbOcean through the “similar” relation). Subsequently, each of the actions in the examples to be classified that was not already contained in EmotiNet, was sought in VerbOcean. In case one of the similar actions was already contained in the KB, the actions were considered equivalent. Further on, each action was associated with an emotion, using the ConceptNet relations and concepts (HasSubevent, Causes, ConceptuallyRelatedTo, HasPrerequisite). Finally, new examples were matched against chains of actions containing the same emotions, in the same order. While more complete than the first approximation, this approach was also affected by lack of knowledge about the emotional content of actions. To overcome this issue, we proposed two heuristics:

(2a) In the first one, actions on which no affect information was available, were sought in within the examples already introduced in the EmotiNet and were assigned the most frequent class of emotion labeling them. The corresponding results are marked with A2a and B2a, respectively.

(2b) In the second approximation, we used the most frequent emotion associated to the known links of a chain, whose individual emotions were obtained from ConceptNet. In this case, the core of action chains is not involved in the process. The corresponding results are marked with A2b and B2b.

We performed the steps described on test set B. The results are shown in Table 1 (results on classified examples) and Table 2 (results on all examples).

Emotion	Correct			Total			Accuracy		
	B1	B2a	B2b	B1	B2a	B2b	B1	B2a	B2b
disgust	16	16	21	44	42	40	36.3 6	38.0 9	52.5 0
shame	25	25	26	70	78	73	35.7 1	32.0 5	35.6 2
anger	31	47	57	10 5	11 5	121	29.5 2	40.8 6	47.1 1
fear	35	34	37	58	65	60	60.3 4	52.3 0	61.6 7
sadness	46	45	41	11 1	12 3	125	41.4 4	36.5 8	32.8 0
joy	13	16	18	25	29	35	52	55.1 7	51.4 3
guilt	59	68	64	15 8	16 5	171	37.3 4	41.2 1	37.4 3
Total	22 5	251	264	57 1	61 7	625	39.4 0	40.6 8	42.2 4

Table 1. Results of the emotion detection using EmotiNet on classified examples in test set B

Emotion	Correct			Total	Recall		
	B1	B2a	B2b	B1	B1	B2a	B2b
Disgust	16	16	21	59	27.11	27.11	35.59
Shame	25	25	26	91	27.47	27.47	28.57
Anger	31	47	57	145	21.37	32.41	39.31
Fear	35	34	37	85	60.34	52.30	61.67
Sadness	46	45	41	267	17.22	16.85	15.36
Joy	13	16	18	50	26	32	36.00
Guilt	59	68	64	198	29.79	34.34	32.32
Total	225	251	264	895	25.13	28.04	29.50
Baseline	126	126	126	895	14.07	14.07	14.07

Table 2. Results of the emotion detection using EmotiNet on all test examples in test set B

7 Discussion and conclusions

From the results in Table 1 and 2, we can conclude that the approach is valid and represents a method that is appropriate for the detection of emotions from contexts where no affect-related words are present. Nonetheless, much remains to be done to fully exploit the capabilities of EmotiNet. We showed that the approach has a high degree of flexibility, i.e. new information can be easily introduced from existing common-sense knowledge bases, such as ConceptNet, mainly due to its internal structure and degree of granularity.

The error analysis we performed shed some light on the causes of error of the system. The first finding is that extracting only the action, verb and patient semantic roles is not sufficient. There are other roles, such as the modifiers, which change the overall emotion in the text. Therefore, such modifiers should be included as attributes of the concepts identified in the roles. A further source of errors was that lack of knowledge on specific actions. Thus, the results of our approach can be practically limited by the structure, expressivity and degree of granularity of the imported resources. Therefore, to obtain the final, extended version of EmotiNet we should analyze the interactions between the core and the imported resources and among these re-sources as well.

Finally, other errors were produced by NLP processes and propagated at various steps of the processing chain (e.g. SRL, coreference resolution). Some of these errors cannot be eliminated; however, others can be partially solved by using alternative NLP tools.

Future work aims at extending the model by adding affective properties to the concepts

included, so that more of the appraisal criteria can be introduced in the model, testing new methods to assign affective value to the concepts and adding new knowledge from sources such as CYC.

Acknowledgments

This paper has been supported by the Spanish Ministry of Science and Innovation (grant no. TIN2009-13391-C04-01), by the Spanish Ministry of Education under the FPU Program (AP2007-03076), and by the Valencian Ministry of Education (grant no. PROMETEO/2009/119 and ACOMP/ 2010/288).

References

- A. Balahur and A. Montoyo. 2008. Applying a Culture Dependent Emotion Triggers Database for Text Valence and Emotion Classification, proceedings of the AISB 2008 Convention "Communication, Interaction and Social Intelligence".
- A. Balahur and R. Steinberger. 2009. Rethinking Opinion Mining in Newspaper Articles: from Theory to Practice and Back, proceedings of the first workshop on Opinion Mining and Sentiment Analysis (WOMSA 2009).
- A. Esuli and F. Sebastiani. 2005. Determining the semantic orientation of terms through gloss analysis", proceedings of CIKM 2005.
- B. Pang and L. Lee. 2008. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, Vol 2, Nr. 1-2, 2008.
- B. Pang, L. Lee and S. Vaithyanathan. 2002. Thumbs up? Sentiment classification using machine learning techniques, proceedings of EMNLP-02.
- C. Strapparava and R. Mihalcea. 2007. Semeval 2007 task 14: Affective text, proceedings of ACL 2007.
- E. Cambria, A. Hussain, C. Havasi and C. Eckl. 2009. Affective Space: Blending Common Sense and Affective Knowledge to Perform Emotive Reasoning, proceedings of the 1st Workshop on Opinion Mining and Sentiment Analysis (WOMSA).
- E. Riloff and J. Wiebe. 2003. Learning extraction patterns for subjective expressions, proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing.
- E. Riloff, J. Wiebe and T. Wilson. 2003. Learning subjective nouns using extraction pattern bootstrapping. In Proceedings of the Conference on Natural Language Learning (CoNLL) 2003, pp.25-32, Edmonton, Canada.
- G. Van den Bos. 2006. *APA Dictionary of Psychology*. Washington, DC: American Psychological Association.
- H. Liu and P. Singh. 2004. ConceptNet: A Practical Commonsense Reasoning Toolkit, *BT Technology Journal*, Volume 22, Kluwer Academic Publishers.
- H. Liu, H. Lieberman and T. Selker. 2003. A Model of Textual Affect Sensing Using Real-World Knowledge, proceedings of IUI 2003.
- J. De Rivera. 1977. A structural theory of the emotions, *Psychological Issues*, 10 (4), Monograph 40.
- J. W. Pennebaker, M. R. Mehl and K. Niederhoffer. 2003. Psychological aspects of natural language use: Our words, our selves, *Annual Review of Psychology* 54, 547-577.
- K. Scherer and H. Wallbott. 1997. The ISEAR Questionnaire and Codebook, Geneva Emotion Research Group.
- K. Scherer, K. 2005. What are emotions? and how can they be measured? *Social Science Information*, 3(44), 695-729.
- M. Dyer. 1987. Emotions and their computations: three computer models, *Cognition and Emotion*, 1, 323-347.
- N. Frijda. 1986. *The emotions*, Cambridge University Press.
- P. Moreda, B. Navarro and M. Palomar. 2007. Corpus-based semantic role approach in information retrieval, *Data Knowl. Eng. (DKE)* 61(3):467-483.
- P. N. Johnson-Laird and K. Oatley. 1989. The language of emotions: An analysis of a semantic field, *Cognition and Emotion*, 3, 81-123.
- P. Subasic and A. Huettner. 2000. Affect Analysis of text using fuzzy semantic typing, *IEEE Transactions on Fuzzy System*, 9, 483-496.
- R. A. Calvo and S. D'Mello. 2010. Affect Detection: An Interdisciplinary Review of Models, Methods and Their Applications, *IEEE Transactions on Affective Computing*, Vol. 1, No. 1, Jan.-Jun.
- R. Picard. 1995. *Affective computing*, Technical report, MIT Media Laboratory.
- R. Plutchik. 2001. The Nature of Emotions. *American Scientist*. 89, 344.
- R. Studer, R. V. Benjamins and D. Fensel. 1998. Knowledge engineering: Principles and methods, *Data & Knowledge Engineering*, 25(1-2):161-197.

- S. Y. Mei Lee, Y. Chen and C.-R. Huang. 2009. Cause Event Representations of Happiness and Surprise, proceedings of PACLIC 2009.
- T. Chklovski and P. Pantel. 2004. VerbOcean: Mining the Web for Fine-Grained Semantic Verb Relations”, proceedings of EMNLP-04.
- T. Danisman and A. Alpkocak. 2008. Feeler: Emotion Classification of Text Using Vector Space Model, proceedings of the AISB 2008 Convention, “Communication, Interaction and Social Intelligence”.
- W. Parrott. 2001. Emotions in Social Psychology, Psychology Press, Philadelphia.