WordNet-Affect: an Affective Extension of WordNet

Carlo Strapparava and Alessandro Valitutti

ITC-irst, Istituto per la Ricerca Scientifica e Tecnologica
I-38050 Povo Trento Italy
{strappa, alvalitu}@itc.it

Abstract

In this paper we present a linguistic resource for the lexical representation of affective knowledge. This resource (named WORDNET-AFFECT) was developed starting from WORDNET, through a selection and tagging of a subset of synsets representing the affective meanings.

1. Introduction

In this paper we present a linguistic resource for a lexical representation of affective knowledge. This resource (named WORDNET-AFFECT) was developed starting from WORDNET, through the selection and labeling of the synsets representing affective concepts.

Affective computing is advancing as a field that allows a new form of human computer interaction, in addition to the use of natural language. There is a wide perception that the future of human-computer interaction is in themes such as entertainment, emotions, aesthetic pleasure, motivation, attention, engagement, etc. Studying the relation between natural language and affective information and dealing with its computational treatment is becoming crucial.

For the development of WORDNET-AFFECT, we considered as a starting point WORDNET DOMAINS (Magnini and Cavaglià, 2000), a multilingual extension of WordNet, developed at ITC-irst. In WORDNET DOMAINS each synset has been annotated with at least one *domain label* (e.g. SPORT, POLITICS, MEDICINE), selected from a set of about two hundred labels hierarchically organized. A domain may include synsets of different syntactic categories: for instance the domain MEDICINE groups together senses from Nouns, such as `doctor#1` (i.e. the first sense of the word doctor) and `hospital#1`, and from Verbs such as `operate#7`.

For WORDNET-AFFECT, our goal was to have an additional hierarchy of "affective domain labels", independent from the domain hierarchy, with which the synsets representing affective concepts are annotated.

2. State of the Art in Affective Lexicons

The first attempts to build a lexical structure for affective terms concerned studying which terms are really representing emotions, and what classification criteria to consider. In particular, lexical semantic approaches are founded on the belief that "it is possible to infer emotion properties from the emotion words" (D'Urso and Trentin, 1998). This approach consists of three main steps. First, emotion words are collected from dictionaries (Weigand, 1998) or from literary and newspaper texts. Then, a fixed number of semantic contexts are fixed: e.g. pure emotion terms, personality trait terms, physical and cognitive state terms, etc. (Ortony et al., 1987). Finally, from each term a set of affective dimensions is extracted, using techniques such as factorial analysis (Nowlis and Nowlis, 1956) or multidimensional scaling (Young and Hamer, 1987).

The lexical semantic approach showed a number of important issues. Ortony and Clore (Ortony and Clore, 1981) reviewed the literature on emotion labels, and they suggested that the process used to select emotion words has not led to a domain of emotion words exclusively (e.g. the word *anger* refers to an emotion, *animosity* to a mood, and *confusion* to a cognitive state).

Another problem was outlined by (Watson and Tellegen, 1985): in the literature there is agreement only on two features: arousal (*excited*, *tense* versus *relaxed*, *sleepy*) and valence (*happy*, *glad* versus *sad*, *upset*). However, these two dimensions are not sufficient to individuate the whole spectrum of emotional concepts. Moreover, the techniques of the lexical semantic approach (e.g. factorial analysis and multidimensional scaling) don't allow us to distinguish different senses of the same word. For example, the word *surprise* may refer to a feeling ("the astonishment you feel when something totally unexpected happens to you"), to an event ("a sudden unexpected event"), or to an action ("the act of surprising someone"). Therefore, in order to build a structure for the affective lexicon organization, we cannot use only information coming from the lexicon itself, but we need to get affective information provided by recent scientific research on emotion. In the present approach to the affective lexicon, the center of interest is not to study the nature of emotions, but how the affective meanings are expressed in natural language.

3. WORDNET-AFFECT

Our work on affective lexicon was focused on the realization of a resource that contains a set of affective concepts correlated with affective words. The availability of the WORDNET database was an important starting point. The synset model is sufficiently simple to provide an intrinsic correlation between a concept and the corresponding words. Moreover, WORDNET covers the entire English lexicon and provides an extraordinarily large amount of conceptual distinctions. As well, it is particularly useful from a computational point of view because it was developed for easy access and navigation through its hierarchies. Starting from WORDNET we selected a subset of synsets (named WORDNET-AFFECT) suitable to represent affective concepts. We are actually aiming at exploiting the

expressivity of the WORDNET model without having to introduce modifications in the original structure. Therefore, we added additional information to the affective synsets without defining new ones. Similarly to our method for domain labels, we assign to a number of WORDNET synsets one or more affective labels (*a-labels*) that contribute to precise the affective meaning. For example, the affective concepts representing emotional state are individuated by synsets marked with the a-label EMOTION. There are also other a-labels for those concepts representing moods, situations eliciting emotions, or emotional responses (see Table 4 for a complete list of a-labels).

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.

3.1. The Developement of the *Core* of WORDNET-AFFECT

In order to have an initial set of affective words, a preliminary resource (named AFFECT) was manually realized.

AFFECT is a lexical database containing 1,903 terms directly or indirectly referring to mental (e.g. emotional) states. The main part of AFFECT consists of nouns (539) and adjectives (517). There is a smaller number of verbs (238) and a tiny set of adverbs (15). We started to collect a list of adjectives with the help of dictionaries. Later nouns were added through an intuitive correlation with the adjectives. In a similar way, verbs and adverbs were added. For each item, a frame was created in order to add lexical and affective information. Lexical information includes the correlation between English and Italian terms, parts of speech (pos), definitions, synonyms and antonyms. The attribute POSR relates terms having different pos but pointing to the same psychological category. For example, the adjective *cheerful* is semantically linked to the name *cheerfulness*, to the verb *cheer up* and to the adverb *cheerfully*. Affective information is a reference to one or more of the three main kinds of theories on emotion representation: discrete theories (based on the concept of cognitive evaluation), basic emotion theories and dimensional theories. According to the work of (Ortony et al., 1987), terms are classified in *emotional terms*, *non-emotional affective terms* (e.g. *mood*) and *non affective mental state terms*. Other terms are linked with *personality traits*, *behaviors*, *attitudes*, *physical* or *bodily states* and *feelings* (such as *pleasure* or *pain*). We named ORTONY the attribute used to indicate the affective category of the terms in the database. Some examples terms and their category are given in Table 1.

Category	Example Term
emotion	anger
cognitive state	doubt
trait	competitive
behaviour	cry
attitude	skepticism
feeling	pleasure

Table 1: Categories and terms

By mapping the senses of terms in AFFECT to their respective synsets, the "affective core" was identified. We selected a subset of WORDNET containing all synsets in which there are at least one word of the affective wordlist, and rejected those synsets that are not recognized as affective concepts.

An automatic check for coherence of the affective information inside the synsets was performed. In particular, we checked if there were synonyms with incompatible values of the affective information. The results have shown that the synsets are a good model for the representation of affective concepts.

Then, we projected part of the affective information from the AFFECT database onto the corresponding senses of WORDNET-AFFECT, as value of an affective mark (the a-label). The information projected was that of the AFFECT slot ORTONY (used to discriminate between different types of affective concepts, as explained above). This operation was not complete over all synsets of WORDNET-AFFECT, both because the value of the ORTONY slot was null for some of the AFFECT items, and because there are synsets manually added besides those individuated in AFFECT. For this reason, we proceeded to a further manual labeling, in order to assign the a-labels to the whole set of affective synsets.

3.2. The Extension of the *Core* with WORDNET relations

In WORDNET a fixed number of lexical (i.e. between words) and semantic (i.e. between synsets) relations are defined. Once we individuated the affective core, we studied if and at what extent, exploiting the WORDNET relations, the affective core of WORDNET-AFFECT could be extended.

For each relation, we examined if it preserves the affective meaning (i.e. if that relation, applied to the synset of WORDNET-AFFECT, generates synsets that still represent affective concepts). If the resulting synsets are members of WORDNET-AFFECT, the answer is trivially affirmative. But in the case where the relation generates synsets not included in the database, it should be necessary to proceed to manual checking. However, an exploratory examination allowed us to individuate a list of "reliable" relations (antonymy, similarity, derived-from, pertains-to, attribute, also-see), for which we assumed that the affective meaning is preserved for all items of WORDNET-AFFECT. Therefore, all synsets obtained by an application of those relations and not yet contained in WORDNET-AFFECT are, de facto, included in it.

For other relations (such as hyperonymy, entailment, causes, verb-group) we assumed that the affective meaning is only partially preserved. In that case it is necessary to manually filter the synsets in order to select those genuinely affective. In particular, it is useful to compare the affective information of the database with WORDNET hyperonym hierarchy restricted to the PSYCHOLOGY domain, in order to propose enrichment in the structure of this semantic field.

WORDNET-AFFECT at the moment contains 2,874 synsets and 4,787 words.

	#Nouns	#Adjectives	#Verbs	#Adverbs	#Total
#Synsets	763	1462	322	327	2874
#Words	1285	2293	657	552	4787

Table 2: Affective synsets and words, grouped by part of speech

	#Nouns	#Adjectives	#Verbs	#Adverbs	#Total
similar-to	-	573	-	-	573
antonym	57	83	23	6	169
pertains-to (*direct*)	2	-	-	-	2
pertains-to (*inverse*)	-	16	-	-	16
derived-from (*direct*)	-	12	-	-	12
derived-from (*inverse*)	-	-	-	308	308
also-see	-	138	11	-	149
attribute	-	38	-	-	38
is-value-of	30	-	-	-	30
Total	89	860	34	314	1297

Table 3: Affective synsets obtained applying WORDNET relations to the "core" of WORDNET-AFFECT

4. Applications

WORDNET-AFFECT is useful in all applications in which it is necessary to have an affective interaction. On this subject, let us consider some existing implemented research and systems.

Elliot's *affective reasoner*. This is a collection of Artificial Intelligence programs that reason about human emotion, and are embodied in multimedia computer agents. It was conceived and developed by Clark Elliot (Elliot, 1992), but it is originally based on the theoretical work of Andrew Ortony et al. (Ortony et al., 1988). The model on which the system was developed consists of a collection of 26 emotion categories related to eliciting conditions (events, objects and persons, actions) through a set of rules. The conditions determine the choice of the emotion and a corresponding emotional response, e.g. a convenient facial expression, for an embodied agent, or a verbal utterance, for a conversational agent. For the latter, the availability of an organized lexical resource would allow to enhance the verbal expressivity.

Information and tutoring tools. These systems use natural language generation to provide information on a particular subject, or to instruct how to perform some complex action. There are domains in which it is useful to produce messages that are empathetic to the hearer. In this case, the form of the messages is as important as the content. For example, when the message content produces an emotional effect on the subject, the form may offset the "unpleasant" information and stressing the "favourable" one, through mitigating or enhancing terms (such as detensifier and intensifier adverbs) (De Rosis and Grasso, 1999). For this purpose, an affective lexical resource can provide a wide spectrum of lexical variants of the same concept, with different affective weights.

Affective text sensing systems. These are programs for assessing the affective qualities of natural language. A new interesting approach, corpus-based, is that of (Liu et al., 2003). The affect of the text, at the sentence level, is classified into one of six basic categories of emotions. The analysis is performed through a model built starting from OpenMind Commonsense, a large-scale collection of common sense knowledge (`http://openmind.media.mit.edu`). Liu et al. chose a list of emotion words (named ground words) by which to bound a first set of affective sentences in OpenMind. These sentences contain other words on which the affective information of the ground words is propagated, with an attenuation factor. By these new words, a new set of affective sentences in OpenMind is individuated, and so on. This approach can be improved by increasing the number of ground words and by considering the senses of the words. Then, a lexical resource including the relation between affective words and concepts is required.

Computational humour. There are some situations where humour can play an important role in improving human-computer interaction (e.g. edutainment or frustration reduction). These are very difficult tasks, but there are some recent positive results in this direction. Stock and Strapparava (Stock and Strapparava, 2003) have worked at a concrete limited problem for the core of the European Project *HAHAcronym*. The main goal of HAHAcronym has been the realization of an acronym ironic re-analyzer and generator. The re-analyzer takes as input an acronym with its expansion, and gives as output some re-analyses of the same acronym with a humorous expansion. Making fun of existing acronyms amounts to basically using irony on them, desecrating them with some unexpected contrasting but nonetheless consistent sounding expansion. In this system, ironic reasoning is developed mainly at the level of acronym choice and in the incongruity resulting in the relation to the coherently combined words of the acronym expansion. The acronym generator is more complex than the re-analyzer. In this case, the input is some concepts (i.e. synsets) from which the system generates both the acronym and the expansion.

A-Labels	Examples
EMOTION	noun anger#1, verb fear#1
MOOD	noun animosisy#1, adjective amiable#1
TRAIT	noun aggressiveness#1, adjective competitive#1
COGNITIVE STATE	noun confusion#2, adjective dazed#2
PHYSICAL STATE	noun illness#1, adjective all_in#1
EDONIC SIGNAL	noun hurt#3, noun suffering#4
EMOTION-ELICITING SITUATION	noun awkwardness#3, adjective out_of_danger#1
EMOTIONAL RESPONSE	noun cold_sweat#1, verb tremble#2
BEHAVIOUR	noun offense#1, adjective inhibited#1
ATTITUDE	noun intolerance#1, noun defensive#1
SENSATION	noun coldness#1, verb feel#3

Table 4: A-Labels and corresponding example synsets

The availability of an affective lexical resource can improve this strategy by allowing the system to focalize the incongruity at the affective level. For re-analyzing, a positive or a negative valence value is attributed to the acronym, and then the expansion generation must include affective words (e.g. appreciative and depreciative words) with opposite valence. For acronym generation, the valence opposition should be applied to both the input concept and the acronym.

5. Future Work

The work described above is only a first step towards the development of this resource. One goal is that of extending the number of affective synsets, making use of the predefined WORDNET relations (in particular, cause and entailment relations) and finding correlations between affective labels and domain labels. It is also useful interfacing WORDNET-AFFECT with other linguistic resources containing common sense knowledge, in order to extract contextual lexicon (e.g. emotional responses and events that cause emotions). A good source of stereotypical knowledge is OpenMind Commonsense, a wide common sense knowledge base containing sentences, linguistic patterns and parse trees. WORDNET-AFFECT allows us to identify the sentences in OpenMind containing affective words corresponding to affective synsets.

Finally, for the organization of the a-labels we want to stress the importance of the affective lexicon in communication and persuasion. We pay attention mainly on *slanting lexicon* (e.g. *appreciative* and *depreciative* words; *intensifier* and *detensifier* adverbs). Slanted writing is that type of writing that springs from our conscious or subconscious choice of words and images. We may load our description of a specific situation with vivid, connotative words and figures of speech. These words have the capability to provide an affective connotation to the text and reveal the affective disposition of the speaker or induce an similar disposition on the recipient. They have an important role in persuasion and for this reason they are very used in advertisement.

We believe that enhancing the resource in this direction, we can make it more useful in affective computing and particularly in applications based on affective natural language processing.

WORDNET-AFFECT (like WORDNET DOMAINS) is freely available for research purposes.

6. References

De Rosis, F. and F. Grasso, 1999. Affective text generation. In *Proc. of International Workshop of Affect in Interactions (IWAI99)*. Siena, Italy.

D'Urso, V. and R. Trentin, 1998. *Introduzione alla psicologia delle emozioni*. Laterza.

Elliot, C. D., 1992. *The Affective Reasoner: a process model of emotions in a multi-agent system*. Ph.D. thesis, Northwestern University, Evanston, Illinois.

Liu, H., H. Lieberman, and T. Selker, 2003. A model of textual affect sensing using real-world knowledge. In *Proceedings of the Seventh International Conference on Intelligent User Interfaces (IUI 2003)*. Miami, Florida.

Magnini, B. and G. Cavaglià, 2000. Integrating subject field codes into wordnet. In *Proceedings of LREC-2000, the Second International Conference on Language Resources and Evaluation*. Athens, Greece.

Nowlis, V. and H. Nowlis, 1956. The description and analysis of mood. *Annals of the New York Academy of Sciences*, 65:345–355.

Ortony, A. and G. L. Clore, 1981. Disentangling the affective lexicon. In *Proceedings of the Third Annual Conference of the Cognitive Science Society*. Berkeley.

Ortony, A., G. L. Clore, and A. Collins, 1988. *The Cognitive Structure of Emotions*. New York: Cambridge University Press.

Ortony, A., G. L. Clore, and M. A. Foss, 1987. The psychological foundations of the affective lexicon. *Journal of Personality and Social Psychology*, 53:751–766.

Stock, O. and C. Strapparava, 2003. Getting serious about the development of computational humor. In *Proceedings of IJCAI 03, the Eighteenth International Joint Conference on Artificial Intelligence*. Acapulco, Mexico.

Watson, D. and A. Tellegen, 1985. Toward a consensual structure of mood. *Psychological Bullettin*, 98:219–235.

Weigand, E., 1998. *Contrastive Lexical Semantics*. John Benjamin Publishing.

Young, F. W. and R. M. Hamer, 1987. *Multidimensional Scaling: History, Theory and Applications*. New York: Erlbaum.