

Developing Affective Lexical Resources

Alessandro Valitutti[♦], Carlo Strapparava, Oliviero Stock

ITC-irst, Trento, Italy

ABSTRACT

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. 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.

Keywords: *Affective Computing, NLP, Lexical Resources, WORDNET.*

Received 20 January 2004; received in revised form 31 March 2004; accepted 1 April 2004.

1. Introduction

Affective computing is advancing as a field that allows for a new way of human computer interaction. In the most sophisticated cases, the interaction modality is meant to be based on language. There is a wide perception that the future of HCI is in themes such as entertainment, emotions, aesthetic pleasure, motivation, attention, engagement, etc. Studying the relationship between natural language and affective information and dealing with their computational treatment is becoming crucial. Initial user studies show that computers' affective ability plays a vital role in improving interaction with users. This ability depends not only on the affective expressiveness, but also on the capacity to detect the affective state of the user. Researchers have tried detecting the user's affective state in many ways, such as through facial expressions, speech, physiology, and text. In particular, text is an important modality for sensing affect because the bulk of computer user interfaces today are textually based. Examples of such applications are synthetic agents for giving affective responses to the user input at the sentence level (e.g. an affective text analyzer architecture [Liu et

[♦] Corresponding Author:
Alessandro Valitutti,
ITC-irst, Via Sommarive 18,
38050, Povo, Trento, Italy.
Phone: +39 0461 314-568
Fax: +39 0461 302040
Email: alvalitu@itc.it

al. 2003], and affective text-to-speech systems). In other applications, the affective interaction is finalized to influence the emotional state of the user. We have worked in this direction, in particular addressing some aspects of computational humour. For instance, we have exploited NLP to build systems that were capable of inducing amusement and affecting the emotional state of users (e.g. HAHAcronym [Stock and Strapparava 2003]). For these applications it is necessary to have linguistic resources containing affective knowledge, which, unfortunately, cannot be found off-the-shelf. In order to develop a linguistic resource of this type, it is necessary first to consider the lexical components of the affective language.

In this paper we describe our efforts in this direction. We developed a preliminary version of a lexical knowledge base containing words in an affective lexicon connected with a set of affective concepts. This resource (named WORDNET-AFFECT) was developed starting from the lexical knowledge base WORDNET, through a selection and labeling of the affective concepts (represented by sets of synonyms). WORDNET-AFFECT was then extended taking into account OpenMind, a database of common sense sentences, in which there is a considerable amount of common sense knowledge [Singh et al. 2002]. Exploiting WORDNET-AFFECT and a word sense disambiguation algorithm [Magnini et al. 2002], we automatically chose an "affective-oriented" subset of OpenMind, named OpenMind-Affect. OpenMind-Affect is composed of sentences, patterns, and parsing trees containing affective concepts. In this way the affective lexicon is enriched by "contextual words", which do not directly refer to affective state (emotions and mood), but that are meaningful from the affect point of view.

2. Applicative relevance of affective lexicon

2.1 Affective verbal language and HCI

As claimed by Picard [1997], there are contexts in which human-computer interaction can be improved with communication that involves the use of emotional information. Affective communication can be performed with facial expressions, physiological and behavioral responses to the affective states, and natural language. We are particularly interested in the verbal component of the affective communication, for different reasons.

Regarding the applications in affect recognition, the nonverbal channel is insufficient for expressing the full range of human emotional experiences. According to Fussell, nonverbal cues can indicate what general class of emotions a person is feeling, but

they typically do not provide detailed information about that person's emotional state. "By seeing that someone is crying, for instance, we might assume that they are sad; by the extent of sobbing we might even be able to infer the intensity of the sadness. But the tears in and of themselves provide no information about the particular experience of sadness. Instead, verbal descriptions of emotional states can provide quite precise information about the specific form of an emotion, such as anger, depression, or happiness, that a person is experiencing" [Fussell 2003]. The richness of natural language makes it possible to express emotions in different modalities, and allows us to distinguish explicit communication (as in *introspective reports*) from unintentional communication (e.g. when the form or the lexical content of the expression reveals an appreciative or depreciative disposition of the speaker).

These characteristics of affective verbal language allow us to develop systems for the detection and modelling of user affective states. The affective user modelling can be used to realize an "emotional intelligence" consisting of the ability to recognize and appropriately respond to human emotions [Picard 2000]. For these applications, the behaviour of the system is not describable only in terms of affect recognition or expression, but consists of a complex adaptive interaction with the user.

In several applicative contexts the verbal output of the system depends not only on the actual affective state of the user, but especially on the desired one. For example, in computational humour and computational persuasion [Guerini, Stock and Zancanaro 2003], the system tries to put the user in the desired mental state (which, in the case of computational humour, is *amusement*). In this case, the affective verbal language is actively used by the program, instead of by the user.

In all HCI applications in which it is necessary to have an affective interaction, it is very useful to have an affective lexical resource. On this subject, let us consider some existing applicative research and systems.

2.2 Example applications

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 [1992], but it is originally based on the theoretical work of Andrew 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, availability of organized lexical resource allows us 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. They can be very useful for HCI systems performing text based affective user modeling. 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 emotion. The analysis is performed through a model built starting from *OpenMind Commonsense*, a large-scale collection of common sense knowledge. 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 [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 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 otherwise 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 a concept from which the system generates both the acronym and the expansion.

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.

3. The model for the organization of the affective lexicon

In order to organize the affective lexicon, it is necessary to have a model for the affective knowledge representing emotions, moods, attitudes, and traits. Using this model we should be able to identify a large number of affective concepts, organize them in a hierarchical structure and connect them via the lexicon. The past literature on affective lexicons guided us in the search of the more suitable model.

3.1 Limitations of the lexical semantics approach

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 semantics approach was 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 et Hamer. 1987].

The lexical semantics approach showed a number of important issues. 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. 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 some features such as “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 semantics 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”).

In our 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. Therefore, in order to build a structure for organizing the affective lexicon, we cannot use only information coming from the lexicon itself, but we need to get affective information from “extra-linguistic” sources, provided by recent scientific research on emotion. We preferred to look for an existing lexical resource and to enrich it with additional affective information by a semiautomatic tagging procedure.

3.2 The adoption of WORDNET as a model for the affective concepts

The sought-after model must exhibit an explicit representation of the concepts, wide lexical coverage and a simple correlation between concepts and words. We think that the lexical knowledge base WORDNET would be the best candidate for satisfying these requisites. In the following section we provide an overview on WORDNET and on its structural organization.

4. WORDNET and WORDNET DOMAINS

As far as the representation and organization of lexical information is concerned, a key concept is the idea of a Lexical Knowledge Base (LKB), proposed, among others, by [Briscoe 1991] and [Calzolari 1992] to provide information, mostly of a semantic nature, consistently structured and available electronically. A Lexical Knowledge Base is an evolution both from a Machine Readable Dictionary, in which one finds an electronic version of the paper dictionary, and from a Lexical Data Base, in which part of the information available in the text format dictionary has been extracted.

One of the most significant attempts to realize a large scale LKB is WORDNET¹, a thesaurus for the English language based on psycholinguistics principles and developed at the Princeton University by George Miller [Miller 1990; Fellbaum 1998]. WORDNET organizes lexical information in terms of word meanings, rather than word forms. It has been conceived as a computational resource, improving some of the drawbacks of traditional dictionaries, such as the circularity of the definitions and the ambiguity of sense references. English nouns, verbs, adjectives and adverbs (about 130,000 lemmas for all the parts of speech in version 1.6) are organized into synonym classes (*synsets*), each representing one underlying *lexical concept* (about 100,000 synsets). Lemmas (about 130,000 for version 1.6) are organized in synonyms classes (about 100,000 synsets). WORDNET can be described as a “lexical matrix” with two dimensions: a dimension for *lexical relations*, that is relations holding among words and thus language-specific, and a dimension for *conceptual relations*, which hold among senses (in WORDNET they are called synsets) and that, at least in part, we consider independent from a particular language. In Table 1 an example of a lexical matrix is reported. *Word form* refers to the physical utterance or inscription; *word meaning* refers to a lexicalized concept. F_1 and F_2 are synonymous, while F_2 is polysemous. Polysemy and synonymy are problems gaining access to information in the mental lexicon.

<i>Word meaning</i>	<i>Word forms</i>				
	F_1	F_2	F_3	...	F_n
M_1	E_{11}	E_{12}			
M_2		E_{22}			
M_3			E_{33}		
...				...	
M_m					E_{mn}

Table 1: WORDNET Lexical Matrix.

A synset contains all the words by means of which it is possible to express the synset meaning: for example the Italian synset {calcium, calcio, Ca} describes the sense of “calcio” as a chemical substance, while the synset {calcio, pedata} describes the sense of “calcio” as a leg movement. Here a list of some relations present in WORDNET follows.

¹ WORDNET is freely available, for research purposes, at <http://www.cogsci.princeton.edu/~wn/>.

4.1 WordNet Relations

Synonymy. The most important lexical relation for WORDNET is the similarity of meaning, since the ability to recognize synonymy among words is a prerequisite to build synsets and therefore meaning representation in the lexical matrix. Two expressions are synonymous if substitutivity is valid (in other words if the substitution of one with the other does not change the truth value of a phrase). It is important to note that defining synonyms in terms of substitutivity requires partitioning WORDNET into nouns, verbs, adjectives and adverbs. This is consistent with the psycholinguistic evidence that nouns, verbs, adjectives and adverbs are independently organized in the human semantic memory. Obviously if a word pertains to more than one synset, this gives an indication of its polysemy.

Antonymy. This is another familiar relation among words. It provides the organizing principle for adjectives. The antonym of a word w in general is $\text{not-}w$. However there can be exceptions to this interpretation: for instance, while “rich” and “poor” are antonyms, the statement that someone is not rich does not implies that he is poor.

Hyperonymy / Hyponymy. This corresponds to the well known ISA relation. In a different way from synonymy and antonymy, hyperonymy (and its inverse hyponymy) is a relation between meanings, so it holds among synsets. As an example the synset {apple tree} is a hyponymy of the synset {tree}, which in turn is an hyponymy of {plant}. This relation provides the organizing principle for the noun hierarchy. Given a Hyperonymy/Hyponymy hierarchy it is possible to calculate the “coordinate-terms” for a given synset. For example, among the “coordinate-terms” for {horse} there are the synsets {mule} and {zebra}, which are common hyponyms of the synset {equine, equid}.

Meronymy / Holonymy. This represents the relation between a whole and its parts. It is a relation among synsets. Three types of homonymic relations, along with their meronymic inverse, are used in WORDNET: member-of (e.g. {tree} is member-of {forest}); part-of of (e.g. {kitchen} is part-of {apartment}); substance-of of (e.g. {hydrogen} is substance-of {water H₂O}).

Entailment. This is a semantic relation used for defining the verb hierarchy. From a logic point of view a proposition P “entails” a proposition Q if there is no state of the

world in which P is true and Q is false. As an example the synset {snore} implies the synset {sleep}.

Troponymy. The entailment relation is at the base of the definition of the “troponymy” relation, which holds among verbs: in fact synset S1 is troponym of synset S2 if S1 implies S2 and if S1 is temporally co-extended with S2 (e.g. the synset {walk} is a troponym of the synset {move}).

ITC-irst has developed WORDNET DOMAINS, a multilingual extension of the well-known English WORDNET. This is a general tool: it is a multilingual lexical database where English and Italian senses are aligned. This multilingual extension [Artale et. al. 1997] [Magnini and Strapparava 1997] is based on the assumption that a large part of the conceptual relations defined for the English (about 72,000 ISA relations and 5,600 part-of relations) can be shared with Italian. From an architectural point of view, the Italian part of WORDNET DOMAINS implements an extension of the WORDNET lexical matrix to a “multilingual lexical matrix” through the addition of a third dimension relative to the language.

Some specific algorithms for use in computational humour have to be developed on top of WORDNET DOMAINS. A fundamental tool is an incongruity detector/generator: concretely we need to be able to detect or propose semantic mismatches between word meanings in the acronym context. This incongruity detector/generator exploits the domain augmentation described in the following section.




Fig.1: A portion of the domain hierarchy.

4.2 Augmenting WORDNET with Domain information

Domains have been used both in Linguistics (i.e. Semantic Fields) and in Lexicography (i.e. Subject Field Codes) to mark technical usages of words. Although this is useful information for sense discrimination, in dictionaries it is typically used for a small portion of the lexicon. WORDNET DOMAINS² is an attempt to extend the coverage of domain labels within an already existing lexical database, WORDNET (version 1.6). Synsets have been annotated with at least one domain label, selected from a set of about two hundred labels hierarchically organized. (Figure 1 shows a portion of the domain hierarchy).

² WORDNET DOMAINS and WORDNET-AFFECT are both freely available, for research purposes, at <http://wndomains.itc.it>.

<i>Sense</i>	<i>Synset and Gloss</i>	<i>Domains</i>
#1	depository financial institution, bank, banking concern, banking company (a financial institution...)	ECONOMY
#2	bank (sloping land...)	GEOGRAPHY, GEOLOGY
#3	bank (a supply or stock held in reserve...)	ECONOMY
#4	bank, bank building (a building. . .)	ARCHITECTURE, ECONOMY
#5	bank (an arrangement of similar objects...)	FACTOTUM
#6	savings bank, coin bank, money box, bank (a container...)	ECONOMY
#7	bank (a long ridge or pile...)	GEOGRAPHY, GEOLOGY
#8	bank (the funds held by a gambling house...)	ECONOMY, PLAY
#9	bank, cant, camber (a slope in the turn of a road...)	ARCHITECTURE
#10	bank (a flight maneuver. . .)	TRANSPORT

Table 2: WORDNET senses and domains for the word "bank".

We have organized about 250 domain labels in a hierarchy (exploiting Dewey Decimal Classification), where each level is made up of codes of the same degree of specificity: for example, the second level includes domain labels such as BOTANY, LINGUISTICS, HISTORY, SPORT and RELIGION, while at the third level we can find specialisation such as AMERICAN_HISTORY, GRAMMAR, PHONETICS and TENNIS.

Information brought by domains is complementary to what is already in WORDNET. First of all a domain may include synsets of different syntactic categories: for instance MEDICINE groups together senses from Nouns, such as `doctor#1` and `hospital#1`, and from Verbs such as `operate#7`. Second, a domain may include senses from different WORDNET sub-hierarchies (i.e. deriving from different "unique beginners" or from different "lexicographer files"). For example, SPORT contains senses such as `athlete#1`, deriving from `life_form#1`, `game_equipment#1`, from `physical_object#1`, `sport#1` from `act#2`, and `playing_field#1`, from `location#1`.

Finally, domains can have an important role in the design of Word Sense Disambiguation algorithms. In fact they may group senses of the same word into homogeneous clusters, with the side effect of reducing word polysemy in WORDNET.

Table 1 shows an example. The word "bank" has ten different senses in WORDNET 1.6: three of them (i.e. bank#1, bank#3 and bank#6) can be grouped under the ECONOMY domain, while bank#2 and bank#7 both belong to GEOGRAPHY and GEOLOGY, causing the reduction of the polysemy from 10 to 7 senses. The ITC-irst particular approach to Word Sense Disambiguation (called Word Domain Disambiguation) has shown good results at the Senseval-2 competition [Magnini *et al.*2001].

5. WORDNET-AFFECT

Our work on affective lexicon was focused on the realization of a resource that contains the set of affective concepts correlated to the 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 correspondent words. Moreover, WORDNET covers the entire English lexicon and provides an extraordinary large amount of conceptual distinctions. WORDNET 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 the "domain" label we attach to synsets in WORDNET DOMAINS, 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. The *concepts* that are *not emotional-affective* (e.g. those representing moods, situations eliciting emotions, or emotional responses) are characterized by other a-labels.

WORDNET-AFFECT was developed in two stages. The first consisted of the identification of the *core synsets*. The second step consisted of the extension of the core with the relations defined in WORDNET.

5.1 The development of the core of WORDNET-AFFECT

The first stage of the WORDNET-AFFECT development consists of collecting an initial set of affective words. To this aim, 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). In order to collect this material, we started from an initial set of psychological adjectives (in particular, affective adjectives). The collection was extended with the help of dictionaries. In a second step, 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, mental attitudes, physical or bodily states and feelings (such as *pleasure* or *pain*). Some examples terms and their category are given in Table 3.

<i>Category</i>	<i>Example Term</i>
Emotion	anger
Cognitive state	doubt
Personality	competitive
Behaviour	cry
Mental attitude	skepticism
Feeling	pleasure

Table 3: Categories and terms.

Discrete emotional information is characterized by an attribute whose value corresponds to one of the 24 emotional categories described by Elliot [1992]. Another attribute allows us to indicate one of the six basic emotions cited by Ekman [1992]. Dimensional emotional information needs two attributes denoting valence (that is, how

positive or negative a fixed emotional state is) and arousal (that is the level of emotional excitation).

Part of the information was collected from dictionaries and from scientific documentation on the psychology of emotions; the remaining information was inserted on an intuitive and arbitrary basis. The former kind of data was associated with references to the sources; the latter is the rough material for a subsequently critical review (for example, by psychologists or lexicographers). As an example, here is one of the frames from the database:

```
[name]: anger
[ita]: <rif src=c> rabbia, collera </rif>
      <rif src=wn sense=1> ira, collera, arrabbiatura, rabbia</rif>
      <rif src=wn sense=2> collera, ira, bile, furia, rabbia</rif>
[def]: <rif src=wn sense=1> (Psychology) a strong emotion; a feeling
      that is oriented toward some real or supposed grievance </rif>
      <rif src=wn sense=2> (Physiology) the state of being angry</rif>
[synonyms]: <rif src=wn sense=1>cholera, ire</rif>
            <rif src=wn sense=2>angriness</rif>
            <rif src=mw> fury, indignation, ire, mad, rage, wrath</rif>
[antonyms]: <rif src=mw>forbearance</rif>
[pos]: n
[posr]: <v>anger</v> <a>angry</a> <r>angrily</r>
[fundamental]: <rif src=d>a</rif>
[elliot]: anger
[valence]: -
[arousal]: 2
[ortony]: emotion
[notes]:
```

5.2 A-labels projections

By mapping the senses of terms in Affect to their respective synsets, the affective core was identified. Then we projected part of the affective information in the Affect database onto the corresponding senses of WORDNET-AFFECT, as value of an affective mark named *a-label* (in the appendix the complete list of a-labels is reported). The information projected was that of the Affect slot *ortony* (used to discriminate between different types of affective concepts). 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-label value to the whole set of affective synsets.

The opportunity to interface Affect with WORDNET allows us to outline different developments. On one hand it is possible to extend the collection through a search of synonyms and antonyms (performed on each of the terms of Affect that are contained in WORDNET). On the other hand it is useful to compare the affective information of the database with WORDNET hyperonym hierarchy restricted to the Psychology domain [Magnini and Cavaglia 2000], in order to propose enrichment in the structure of this semantic field. Our analysis of WORDNET synsets that contain Affect words suggested that synsets could be used to represent affective concepts. At the level of single affective concepts the characterization as synsets is quite accurate, though additional WORDNET relations, such as those resulting in the ISA hierarchy, are not always appropriate. We have given the name WORDNET-AFFECT to the subset of WORDNET that includes 1,314 synsets representing the senses of the entries in Affect.

	<i>#Nouns</i>	<i>#Adjectives</i>	<i>#Verbs</i>	<i>#Adverbs</i>	<i>#Total</i>
<i>#Synsets</i>	535	557	200	22	1314
<i>#Words</i>	1336	1472	592	40	3340

Table 4: Number of affective synsets and words, grouped by part of speech, in WORDNET-AFFECT.

5.3 Exploiting the WORDNET hierarchies

As explained in section 4, 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 yet represent affective concepts). If the resulting synsets are members of WORDNET-AFFECT, the answer is trivially affirmative. But in the case in which 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.

	#Nouns	#Adjectives	#Verbs	#Adverbs	#Total
similar-to	0	668	0	0	668
antonym	64	106	23	6	199
pertains-to (direct)	0	2	0	0	2
pertains-to (inverse)	16	0	0	0	16
derived-from (direct)	0	0	0	12	12
derived-from (inverse)	0	308	0	0	308
also-see	0	148	11	0	159
attribute	38	0	0	0	38
is-value-of	0	30	0	0	30

Table 5: Number of new affective synsets obtained applying WORDNET relations to the synsets of WORDNET-AFFECT.

6. OpenMind-Affect

Once a fixed set of affective concepts has been identified, represented by the synsets of WORDNET-AFFECT, it seems useful to distinguish those directly referring to affective states from those denoting, for example, their causes and consequences. For that it is very important to have a wide resource of common sense expressions. This resource can be exploited as stereotypical knowledge, which allows us to extract relevant information (e.g. events that typically causes specific emotions). With reference to the work of Liu et al. [2003], we have used OpenMind as a source of stereotypical knowledge. OpenMind is a wide common sense knowledge base containing sentences, linguistic patterns and parse trees. Unlike [Liu et al. 2003], sentences of OpenMind were annotated through a word sense disambiguation tool, developed at ITC-irst [Magnini et al. 2002], in order to associate each word with the corresponding WORDNET sense. In this way, it was possible to identify the sentences containing words with an affective meaning and to select an affectively significant subset of OpenMind, which we have called OpenMind-AFFECT. Using only words in Affect, we selected a set of 74,455 sentences in OpenMind. Using the words in WORDNET-AFFECT, we increased the size of this set to 171,657 sentences. This resource is employed as an environment for

experimentation about an affective lexicon. In particular, we aim at obtaining the following results:

1. increasing the collection of affective concepts. To this aim, we need to identify, in the sentences of OpenMind-Affect, new words related to well-known ones, in order to obtain new synsets to include in WORDNET-AFFECT.
2. collecting contextual information, such as events that typically cause specific emotions.
3. using contextual information in order to increase the affective knowledge of the lexical items, whenever possible.

In order to extract the contextual knowledge, it is necessary to exploit some linguistic patterns to connect words denoting affective states with contextual words. For example, the pattern X causes Y, where X denotes an event and Y refers to an emotional state, allows us to identify a typical cause of that emotion. The lexical semantics of emotional adjectives [Goy 2000] allows us to deduce some of these structures even if they are not explicitly present in the sentence. For example, the adjective “cheerful” in general may have different ways for denoting the emotional state (stative, manifestative, causative). Nevertheless, if it is included in the noun phrase “cheerful flower”, it assumes the causative reading and implicitly expresses the fact that the flower causes cheerfulness, which allows us to add a potentially new contextual concept to the set of the affective concepts.

7. Future work

At this stage of the work, we put our efforts towards the collection of an initial set of affective concepts. We performed manual tagging in order to have a *core*, then we employed some WORDNET relations in order to extend the core and to propagate the labelling.

Our next task consists of the definition of an appropriate procedure for the a-label annotation of the synsets. We are thinking of evaluating that annotation scheme with a test of reliability and performing an evaluation of the a-label values with an *intertagging agreement* procedure with multiple judges.

We aim to extend the number of affective synsets not only with WORDNET relations but also with machine learning techniques applied to large linguistic corpora. Finally, we want to test WORDNET-AFFECT through an evaluation of the applicative tools in which that resource will be used. For example, considering the affective text sensing

tools, WORDNET-AFFECT allows such systems to perform reasoning not only on the “word” level but on the “sense” level (through the synsets linked to each word in the resource). In this manner we think we can obtain an improvement in the performance of the affective understanding from texts. The a-label distinction between *emotion*, *mood*, and *trait* allows user-modeling systems to perform more complex distinctions and so improve the richness of the user model itself. We are thus working in the development of a number of prototypes with which we want to test the applicative potentiality of WORDNET-AFFECT.

8. Conclusions

WORDNET-AFFECT is potentially of wider use, as it allows for some reasoning capability and, as we pointed out, is also connected to relevant common sense knowledge, in textual format. An affective lexicon is *per se* an important resource for many applications, both based on language recognition and on language production. The potential applications in natural language processing are the basis for those in human-computer interaction. In particular, language recognition is employed in user modeling, and language production is necessary for the verbal communication of emotions.

In dealing with texts automatically, emotion-related contents can be retrieved or summarized or seriously classified only if we start from some level of lexical reasoning. If a system is able to perform some reasoning starting from emotion-inducing lexical entries, it may be substantially more sophisticated. In our case it may do so just because of the fact that the common sense thesaurus is linked to synonym sets. For instance we may have knowledge stating that people are afraid of earthquakes. At least some of the simplest statements in the common sense base can have been parsed and yielded semantic relations that involve affect. So whenever earthquakes are a topic, the system could also know they produce fear and therefore classify or summarize the text with this additional dimension. Along the same view all tasks that involve disambiguation can substantially exploit the resource.

The multilingual framework of WORDNET DOMAINS developed at IRST also accommodates WORDNET-AFFECT and this is an important resource for automatic translation, and can yield appropriate ways to overcome language gaps systematically.

On the basis of the lexical information, a system can understand explicitly stated information about the emotional state of the user or of someone else and so the system can use emotion expressions appropriately.

One class of potential applications concerns some (even limited) form of reasoning in a dialog system. For instance a system may be able to advise a user interested in choosing a movie taking into account the fact that he said he does not want to see a *scary* movie. In general the system may want to understand statements regarding the affective state of the user, a simple case is automatic analysis of a questionnaire. Whatever the indicators are (there are experiments with physiologically-based feedback), a system may adapt its expressions to the user.

If we consider system output, there are many aspects with potential for application of the affect lexical base. Multimodal interfaces are a privileged case: for instance a life-like character may appear as the agent delivering the message and the system may coordinate appropriate facial expression when emotional expressions are uttered; similarly the synthetic voice may reflect the emotional valence of a specific fragment of the message. Another case that can be interesting concerns Kinetic Types, where types move producing an emotional effect, normally in agreement with the contents.

Still other possibilities concern the use of persuasive expressions. A system may have to make reference to emotional concepts in order to persuade the audience to perform an action (see Guerini, Stock and Zancanaro 2003).

We see similar potential for games, in the future an essential motivating context for learning environments. Emotional concepts are at the basis of many games and if the interaction between player and system is going to become more natural and complex, they are likely to be expressed linguistically. In a system that helps group activity or in other cases of human-human interaction, it may be important that the system provides different messages to different participants, with appropriately different emotional expressions, even if conveying the same semantic contents.

Finally, we are particularly interested in automatic humour production.

Computational humour can have an important role in future interfaces. Humour plays on the beliefs and expectations of the hearer. By infringing on them, it causes surprise and then hilarity. Humour also encourages creativity. The change of perspective caused by humorous situations induces new ways of interpreting the same event. A typical use of humour can be found in advertisement. We aim at building semiautomatic (and later on fully automatic) systems for helping obtaining novel verbal humour expressions, for instance, in advertising applications. For that purpose we intend to use WORDNET-AFFECT with its connected common sense emotional expressions in a manner that is creative and disrupts conventions.

9. Acknowledgments

This work was realized as part of the TICCA project, funded by the Autonomus Province of Trento and the National Council for Research, and of the HUMAINE network of excellence, funded by the European Commission.

10. Appendix: A complete list of *a-labels* used in WORDNET-AFFECT

<i>A-Label</i>	<i>Examples</i>
EMOTION	noun "anger#1", verb "fear#1"
MOOD	noun "animosity#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 6: The *a-labels*.

EMOTION synsets are the only pure emotional synset: they refer directly to emotional states/processes.

From theories of emotions we have that moods have to be distinguished from emotions. Therefore, MOOD synsets are affective, but not emotional.

If personality traits influence attitude to have affective reactions to events, then the "TRAIT" synsets have an indirect reference to the affective states.

Affective states are often confused with cognitive ones because they have a mutual strict relation. COGNITIVE STATE synsets represent cognitive states that are often related to affect but form a distinct category.

Similar considerations have to be made for PHYSICAL STATE synsets.

Edonic signals (pleasure and pain) are a necessary component of the affective state but they are not emotions themselves. We introduced the *a-label* EDONIC SIGNAL in order to cope with these concepts.

Emotional states have causes and consequences. The causes are named ELICITING-EMOTION SITUATIONS and may be events, behaviors etc. (e.g. a "dangerous situation" may cause fear). The consequences of emotions are named "emotional responses", and may consist of behaviors, somatic changes, facial expressions etc. The EMOTIONAL RESPONSE synsets represent that set of responses.

BEHAVIOUR synsets represent behaviors that are either the cause or effect of affective states.

ATTITUDE synsets refer to "attitudes", which are complex mental states involving beliefs, feelings, values, and dispositions to act in certain ways.

A "sensation" is an unelaborated elementary awareness of stimulation. SENSATION synsets represent those kinds of concepts.

11. References

- [Artale et al. 1997] Artale, A, Magnini, B., Strapparava, C. (1997). WordNet for Italian and Its Use for Lexical Discrimination. *AI*IA97: Advances in Artificial Intelligence*. Springer Verlag.
- [Briscoe 1991] Briscoe, T. (1991). Lexical Issues in Natural Language Processing, in Klein E. and Veltman F. (eds.): *Esprit Symposium on Natural Language and Speech*, Berlin, Springer-Verlag.
- [Calzolari 1992] Calzolari N. (1992). Acquiring and Representing Semantic Information in a Lexical Knowledge Base, in Pustejovsky J. and Bergler S. (eds.) *Lexical Semantics and Knowledge Representation*, Springer-Verlag, Berkeley.
- [De Rosis and Grasso 1999]: De Rosis, F. and Grasso, F. (1999). Affective Text Generation. in Proc. of *International Workshop of Affect in Interactions (IWA199)*. Siena, Italy.
- [D'Urso and Trentin 1998]: D'Urso, V., Trentin, R. (1998). *Introduzione alla psicologia delle emozioni*. Ed. Laterza.
- [Ekman 1992]: Ekman, P. (1992). An argument for basic emotion. *Cognition and Emotion*, 6, 169-200.
- [Elliot 1992]: Elliot, C.D. (1992). *The Affective Reasoner: a process model of emotions in a multi-agent system*. Northwestern University doctoral thesis. Evanston, Illinois.
- [Fellbaum 1998]: Fellbaum, C., Ed. (1998). *Wordnet: an electronic lexical database*. MIT Press.

- [Fussell 2003]: Fussell, S. R., Ed. (2003). *The verbal communication of emotion: Interdisciplinary perspectives*. Mahwah, NJ:Lawrence Erlbaum Associates.
- [Goy 2000]: Goy, A. (2000). Lexical Semantics of Emotional Adjectives. In Feist, S., Fix, S., Hay, J., and Moore, J. (eds): *Linguistics in Cognitive Science: Proceedings of Student Conference in Linguistics 10*, MIT Working Papers in Linguistics 37, MIT Press.
- [Guerini, Stock and Zancanaro. 2003]: Guerini, M., Stock, O., and Zancanaro, M. (2003). Towards Intelligent Persuasive Interfaces. *Proc. of the IJCAI2003 Workshop on Natural Argumentation*. Acapulco, Mexico.
- [Liu et al. 2003]: Liu, H., Lieberman, H., Selker, T.(2003). A Model of Textual Affect Sensing using Real-World Knowledge. *Proceedings of the Seventh International Conference on Intelligent User Interfaces (IUI 2003)*,pp. 125-132. Miami, Florida.
- [Magnini and Cavaglia 2000]:Magnini, B. and Cavaglia, G. (2000). Integrating Subject Field Codes into WordNet. In: Gavrilidou, M., Crayannis, G., Markantonatu, S., Piperidis, S., Stainhaouer, G. (eds.): *Proceedings of LREC-2000, Second International Conference on Language Resources and Evaluation Theoretical Aspects of Computer Software*. Athens, Greece, 1413-1418.
- [Magnini and Strapparava 1997] Magnini, B., Strapparava, C. (1997). Costruzione di una base di conoscenza lessicale per litaliano basata su WORDNET. In *Proceedings of the XXVII Congresso Internazionale di Studi della Società di Linguistica Italiana Linguaggio e Cognizione*, Roma, Bulzoni.
- [Magnini et al. 2001] Magnini B, Strapparava C., Pezzulo G., and Gliozzo A. (2001). Using domain information for word sense disambiguation. In *Proc. of SENSEVAL-2*.
- [Magnini et al. 2002]: Magnini, B., Strapparava, C., Pezzulo, G., and Gliozzo, A. (2002). The Role of Domain Information in Word Sense Disambiguation. *Natural Language Engineering*, 8(4):359-373.
- [Miller 1990] Miller, G. (1990) An on-line lexical database. *International Journal of Lexicography*, 13(4):235-312.
- [Nowlis and Nowlis 1956]: Nowlis, V., Nowlis, H.H. (1956). The description and analysis of mood. *Annals of the New York Academy of Sciences*, 65, 345-355.
- [Ortony and Clore 1981]: Ortony, A. and Clore, G.L. (1981). Disentangling the Affective Lexicon. In *Proceedings of the Third Annual Conference of the Cognitive Science Society*, Berkeley.

- [Ortony et al. 1987]: Ortony, A., Clore G. L., and Foss, M. A. (1987). The psychological foundations of the affective lexicon. *Journal of Personality and Social Psychology*, 53, 751-766.
- [Ortony et al. 1988]: Ortony, A., Clore, G.L., Collins, A. (1988). The cognitive structure of emotions. Cambridge University Press, New York.
- [Picard 1997]: Picard, Rosalind (1997). *Affective Computing*. MIT Press.
- [Picard 2000]: Picard, R.W. (2000). *Toward computers that recognize and respond to user emotion*. *IBM Systems Journal*, 39(3&4):705-719.
- [Singh et al. 2002]: Singh, P., Lin, T., Mueller, E. T., Lim, G., Perkins, T., and Zhu, W. L. (2002). Commonsense: Knowledge acquisition from the general public. In *Proceedings of the First International Conference on Ontologies, Databases, and Applications of Semantics for Large Scale Information Systems*, Heidelberg. Springer-Verlag.
- [Stock and Strapparava 2003]: Stock, O. and Strapparava, C. (2003). Getting Serious about the Development of Computational Humor. In *Proceedings of International Joint Conference of Artificial Intelligence (IJCAI-03)*, Acapulco, Mexico.
- [Young and Hamer 1987]: Young, F.W. & Hamer, R.M. (1987) *Multidimensional Scaling: History, Theory and Applications* Erlbaum, New York.
- [Watson and Tellegen 1985]: Watson, D., Tellegen, A. (1985). Toward a consensual structure of mood. *Psychological Bulletin*, 98, 219-235.
- [Weigand 1998]: Weigand, Edda. (1998). *Contrastive Lexical Semantics*. John Benjamin Publishing.