博士学位請求論文

Compositional Approach for Automatic Recognition of Fine-Grained Affect, Judgment, and Appreciation in Text

(テキスト中の感情・判断・評価認識のための構成的アプローチ)

By
Alena Neviarouskaya

A DISSERTATION SUBMITTED TO
THE GRADUATE SCHOOL OF INFORMATION SCIENCE
AND TECHNOLOGY

In Partial Fulfillment of the Requirements for the
DEGREE OF DOCTOR OF PHILOSOPHY

At the
University of Tokyo

Supervisor: Professor Mitsuru Ishizuka
(石塚　満　教授)

December 2010

Abstract

Sharing feelings, pleasant or painful impressions, showing sincere empathy or indifference, exchanging tastes and points of view, advancing moral values, expressing praise or reprehension are indispensable for full-value and effective social interplay between people. With rapidly growing online sources (news, blogs, discussion forums, product or service reviews, social networks etc.) aimed at encouraging and stimulating people's discussions concerning personal, public, or social issues, there is a great need in development of robust computational tools for the analysis of people's preferences and attitudes. Sentiment or subjectivity analysis is nowadays a rapidly developing field with a variety of emerging approaches targeting the recognition of sentiment reflected in written language. Automatic recognition of positive and negative opinions and classification of text using emotion labels have been gaining increased attention of researchers. However, the topic of recognition of fine-grained attitudes expressed in text has been ignored. Attitude types (namely, affect, judgment, and appreciation) define the specifics of appraisal being expressed: distinct types of personal emotional states; positive and negative appraisal of person's character, behavior, skills; and aesthetic evaluation of semiotic and natural phenomena (events, artifacts etc.), correspondingly.

In this thesis, first we describe the developed Affect Analysis Model (AAM) that is based on rule-based linguistic approach for classification of sentences using nine emotion labels (anger, disgust, fear, guilt, interest, joy, sadness, shame, and surprise) or neutral. We demonstrate the results of AAM evaluation on two data sets represented by sentences from diary-like blog posts. Averaged accuracy of our system is up to 81.5 percent in fine-grained emotion classification (nine emotion labels and neutral) and up to 89.0 percent in polarity-based classification.

As lexicon-based systems strongly depend on the availability of sentiment-conveying terms in their databases, in order to overcome the problem of lexicon coverage, we introduce original methods for building and expanding sentiment lexicon (SentiFul) represented by sentiment-conveying words that are annotated by sentiment polarity, polarity scores and weights. The main

features of the SentiFul are as follows: (1) it is built using not only methods exploring direct synonymy, antonymy, and hyponymy relations, but also innovative methods based on derivation and compounding with known lexical units (the originality and valuable contribution lie in the elaborate patterns/rules for the derivation and compounding processes that have not been considered before); (2) it is larger than the existing lists of sentiment words; (3) it includes polarity scores, in contrast to most existing sentiment dictionaries that lack assignments of degree or strength of sentiment. Our AttitudeFul database contains lexicon necessary for fine-grained attitude analysis; it includes attitude-conveying terms, extensive sets of modifiers, contextual valence shifters, and modal operators, which contribute to robust analysis of contextual attitude and its strength.

In this thesis, we introduce novel compositional linguistic approach for attitude recognition in text. There are several aspects that distinguish our Attitude Analysis Model (@AM) from other systems. First, our method classifies individual sentences using fine-grained attitude labels (nine for different affective states, two for positive and negative judgment, and two for positive and negative appreciation), as against other methods that mainly focus on two sentiment categories (positive and negative) or six basic emotions. Next, our Attitude Analysis Model is based on the analysis of syntactic and dependency relations between words in a sentence; the *compositionality principle* (the rules of *polarity reversal, aggregation*, *propagation*, *domination*, *neutralization*, and *intensification*, at various grammatical levels); a novel linguistic approach based on the rules elaborated for semantically distinct verb classes; and a method considering the hierarchy of concepts. As distinct from the state-of-the-art approaches, the proposed compositional linguistic approach for automatic recognition of fine-grained affect, judgment, and appreciation in text (1) is domain-independent; (2) extensively deals with the semantics of terms, which allows accurate and robust automatic analysis of attitude type, and broadens the coverage of sentences with complex contextual attitude; (3) processes sentences of different complexity, including simple, compound, complex (with complement and relative clauses), and complex-compound sentences; (4) handles not only correctly written text, but also informal messages written in an abbreviated or expressive manner; and (5) encodes the strength of the attitude and the level of confidence, with which the attitude is expressed, through numerical values in the interval [0.0, 1.0]. The performance of our Attitude Analysis Model was evaluated on data sets represented by sentences from different domains. @AM achieved high level of accuracy on sentences from personal stories about life experiences, fairy tales, and news headlines, outperforming other methods on several measures. In fine-grained attitude classification (14 labels) our system achieved averaged accuracy of 62.1 percent, and in coarse-grained classification (3 labels) – 87.9 percent.

Using Affect Analysis Model and Attitude Analysis Model, we have developed several applications: AffectIM (Instant Messaging application integrated with AAM), EmoHeart (application of AAM in 3D world Second Life), iFeel_IM! (innovative real-time communication system with rich emotional and haptic channels), and web-based @AM interface. We believe that the output of our systems can contribute to the robustness of the following society-beneficial and analytical applications: public opinion mining, deep understanding of a market and trends in consumers' subjective feedback, attitude-based recommendation system, economic and political forecasting, affect-sensitive and empathic dialogue agent, emotionally expressive storytelling, integration into online communication media and social networks.

Acknowledgments

I would like to express deep gratitude to my advisor Prof. Dr. Mitsuru Ishizuka for his kindness and care, encouragements and support during years of study, and great opportunity to be involved in the research of such a prominent Laboratory. I would also like to thank Associate Prof. Dr. Helmut Prendinger from National Institute of Informatics for the ideas that stimulated interesting discussions and inspired my work, for his kindness and great support. Their highly valuable creative and intelligent advices and suggestions allowed me to overcome the research challenges during my PhD course and greatly contributed to this thesis.

I also thank the members of my review committee: Prof. Dr. Keikichi Hirose, Prof. Dr. Kiyoharu Aizawa, Prof. Dr. Masaru Kitsuregawa, Associate Prof. Dr. Masashi Toyoda, and Associate Prof. Dr. Takeshi Naemura for their helpful comments and insights related to this thesis.

I am very grateful to all people of the Ishizuka Laboratory for their friendliness, sincere support and efforts towards me, and for their utter devotion (worthy of respect) to research. Particularly, I express my keen appreciation to Junichiro Mori, Jaewon Hur, Arturo Nakasone, Mostafa Al Masum, Eiko Kin, Danushka Bollegala, Manuel M. Martinez, Hugo Hernault, Werner Breitfuss, Li Haibo, and Tacoa Francisco Antonio for making my graduate student experience more pleasurable and successful. I am also indebted to the Assistant Prof. Hiroshi Dohi, Laboratory secretary Ms. Meiko Fujita, and Associate Prof. Dr. Helmut Prendinger's administrative assistant Ms. Hiroko Tokuda for their help in many official procedures. I am saying hearty thanks to Ms. Keiko Nakagaki (formerly employed at the Office of International Relations). She considerably assisted all international students from our Graduate School with formal procedures of living in Japan and studying at the University of Tokyo.

I would like to express my deep appreciation to Dr. Alessandro Valitutti and Dr. Diego Reforgiato for their kind help in Affect database creation. I wish also to express my gratitude to Dr. Dzmitry Tsetserukou, Dr. Shaikh Mostafa Al Masum, Manuel M. Martinez, Zoya Verzhbitskaya, Hutchatai Chanlekha, Nararat Ruangchaijatupon, Tatsiana Tsahelnik, Mariya Kalinova Sokolova, and Liudmila Vakkhova, who have contributed to the annotations of the database entries and sentences, for their efforts and time. Special thanks also go to Dr. Ulrich Apel and Dr. Boris Brandherm for their efforts and assistance in organizing the user study on AffectIM; to Cui Xiaoke, Tananun Orawiwattanakul, and Farzana Yasmeen for their work on EmoHeart promotion in Second Life; to all the participants of the demonstrations and experiments.

I would like to thank Ministry of Education, Culture, Sports, Science and Technology of Japan (MEXT) for providing me with Monbukagakusho Scholarship that made my study at the University of Tokyo possible. The cultural and scientific experiences I have got in the University of Tokyo allowed me to develop my spirit and to cognize the world with open and confident look.

On a more personal note, I would like to thank my parents, Vyachaslau Neviarouski and Valyantina Neviarouskaya, my brother Vyachaslau and his wife Tatsiana for their love, care and belief in my success. I express particular and sincere gratitude to my soul mate and wonderful husband Dzmitry for his devoted love, patience, encouragements, and support with many things contributed to this research.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Attitudes and Their Unique Role in Social Interactions

Social interaction among people is an essential aspect of every society, a strong foundation for the development and self-actualization of a person, as well as for the establishment of genuine interpersonal relationships and communities. Attitudes play the role of a sensitive catalyst, which fosters lively interactions between human beings and assists in the development and regulation of interpersonal relationships. Sharing feelings, pleasant or painful impressions, showing sincere empathy or indifference, exchanging tastes and points of view, advancing moral values, expressing praise or reprehension are indispensable for full-value and effective social interplay between people.

The centrality of emotions in social life is manifested by the rich history of theories and debates about emotions and their nature (Solomon 1973; Ekman 1977; Frijda 1986; Ortony, Clore and Collins 1988; Izard 1993; Lazarus 1994; Roseman and Smith 2001). The expression of emotions shapes social interactions by providing observers a rich channel of information about the conversation partner (Ekman 1993) or his social intentions (Fridlund 1992), by evoking positive or negative responses in others (Dimberg and Ohman 1996), and by stimulating other's social behaviour. By accentuating the functional role of emotions, Frijda (1986, 1994) argues that they preserve and enhance life, and Lutz (1988) emphasizes their communicative, moral, and cultural purposes.

Thurstone (1931) considered people's opinions and beliefs to be verbal expressions of attitude with differing degrees of favorableness or unfavorableness toward the attitude object. By the late 1950s, the multi-component view of attitude was adopted almost universally (Ajzen and Fishbein 1980). Attitudes were viewed as complex systems comprising the person's beliefs about the object, his feelings toward the object, and his action tendencies with respect to the object. Lazarus (1994) was one of the initiators of the appraisal theory, according to which emotion is an appraisal of the world. This theory is uncompromisingly devoted to the conceptual nature of emotions concerned with cognitive contents.

According to the Appraisal theory proposed by Martin and White (2005), attitude is represented by the following three types that define the specifics of appraisal being expressed:

(1) *Affect* – personal emotional state or reaction.

(2) *Judgement* – ethical appraisal of person's character, behaviour, skills etc. according to various normative principles.

(3) *Appreciation* – aesthetic evaluation of semiotic and natural phenomena, events, objects etc.

Expressions of attitude accompany us throughout the span of our lives and color the way we build and maintain the basis for interactions with people in a society.

1.2 Research Motivation and Objectives

> '*Attitudinal meanings tend to spread out and colour a phase of discourse as speakers and writers take up a stance oriented to affect, judgment or appreciation*.' Martin and White (2005: 43)

With rapidly growing online sources (news, blogs, discussion forums, product or service reviews, social networks etc.) aimed at encouraging and stimulating people's discussions concerning personal, public, or social issues, there is a great need in development of robust computational tools for the analysis of people's preferences and attitudes. The examples of attitude expressions are: '*This bill will in fact stifle small business, raise taxes, and further destroy the economy of this Nation*' (comment to a news); '*Great budget car with excellent gas mileage*' (product review); '*For me every* *minute on my horse is alike an hour in heaven!*' (personal experience); '*The room was immaculate and top line comfortable!*' (service review).

The following society-beneficial and analytical applications may be driven by the attitude-sensing system: public opinion mining, deep understanding of a market and trends in consumers' subjective feedback, attitude-based recommendation system, economic and political forecasting, affect-sensitive and empathic dialogue agent, emotionally expressive storytelling, integration into online communication media (IM, 3D virtual world etc.) and social networks (e.g., Facebook, Twitter).

The main objectives of our research are:

(1) Fine-grained classification of sentences using attitude types:

Affect: nine emotions defined by (Izard 1971): 'Anger', 'Disgust', 'Fear', 'Guilt', 'Interest', 'Joy', 'Sadness', 'Shame', and 'Surprise'.

For example: '*When I first saw that you could have a chance to swim with dolphins I was very excited*' ('Joy'); '*The fleas were entirely my fault, as I brought three cats to the house*' ('Guilt').

Judgment: positive and negative judgment: 'POS jud' and 'NEG jud'.

For example: '*My Mum is brilliant when it comes to baking and making cakes!!*' ('POS jud'); '*How can people be so mean to hurt an innocent little lovable animal that is just like any other animal*' ('NEG jud').

Appreciation: positive and negative appreciation: 'POS app' and 'NEG app'.

For example: '*I've always thought of life as a precious gift that we should spend as best as we can*' ('POS app'); '*While it is convenient, I think how unfriendly those little cups are for the environment*' ('NEG app').

(2) Novel way of deep attitude analysis based on the compositional approach and the semantics of terms.

(3) Analysis of the strength of the attitude in the interval [0.0, 1.0].

(4) Determination of the level of confidence, with which the attitude is expressed, in the interval [0.0, 1.0].

(5) Development of applications driven by attitude-sensing system.

1.3 Background and Related Work

Issues of recognition, interpretation, synthesis, and representation of affect have been extensively investigated by researchers in the field of affective computing (Picard 1997). A wide range of modalities has been considered, including affect in speech (Cahn 1990; Slot, Cichosz and Bronakowski 2008; Wu, Yeh and Chuang 2009), facial display (Di Fiore et al. 2008; Maglogiannis, Vouyioukas and Aggelopoulos 2009), body posture and gestures (see multimodal approach proposed by Castellano, Kessous and Caridakis (2008)), and physiological activity (Rigas et al. 2007). Recently, textual information has been gaining increased attention of researchers interested in studying different kinds of subjective phenomena, including sentiment, subjectivity, opinions, emotions, and attitudes. According to Reilly and Seibert (2003), sentiment-related information can be encoded *lexically* within the actual words of the sentence, *syntactically* by means of subordinate clauses, and *morphologically* through changes in attitudinal shades of word meaning using suffixes (especially, in languages with rich inflectional system, such as Russian or Italian). In order to analyse these phenomena communicated through written language, researchers in the areas of natural language processing and computational linguistics have proposed a variety of approaches, methodologies, and techniques.

Various approaches to subjectivity, sentiment or affect analysis on different textual composition levels have been proposed:

(1) Word level: Subasic and Huettner (2001); Kamps and Marx (2002); Riloff, Wiebe and Wilson (2003); Turney and Littman (2003); Baroni and Vegnaduzzo (2004); Andreevskaia and Bergler (2006); Strapparava, Valitutti and Stock (2007).

(2) Synset level: Esuli and Sebastiani (2006); Wiebe and Mihalcea (2006).

(3) Phrase level: Wilson, Wiebe and Hoffmann (2005).

(4) Clause or sentence level: Olveres, Billinghurst, Savage and Holden (1998); Boucouvalas (2003); Liu, Lieberman and Selker (2003); Yu and Hatzivassiloglou (2003); Mulder, Nijholt, den Uyl and Terpstra (2004); Read (2004); Kim and Hovy (2005); Neviarouskaya, Prendinger and Ishizuka (2007c, 2010b, 2011); Moilanen and Pulman (2007); Alm (2008);

Aman and Szpakowicz (2008); Choi and Cardie (2008); Ghazi, Inkpen and Szpakowicz (2010).

(5) Paragraph or document level: Subasic and Huettner (2001); Turney (2002); Pang, Lee and Vaithyanathan (2002); Mishne (2005); Kim and Hovy (2006); Leshed and Kaye (2006); Mihalcea and Liu (2006); Nadeau, Sabourin, De Koninck, Matwin, and Turney (2006).

1.3.1 Lexical Resources

To support applications relying on the recognition of textual subjectivity, semantic orientation, and affective language, researchers have created different lexical resources: subjective (Wilson et al. 2005), polarity (Hatzivassiloglou and McKeown 1997; Esuli and Sebastiani 2006; Neviarouskaya, Prendinger and Ishizuka 2009), affective (Strapparava and Valitutti 2004), and appraisal (Argamon, Bloom, Esuli and Sebastiani 2007) lexicons.

Methods for extracting and annotating subjective terms include: machine learning approaches examining the conjunction relations between adjectives (Hatzivassiloglou and McKeown 1997); clustering adjectives according to distributional similarity based on a small amount of annotated seed words (Wiebe 2000); pattern-bootstrapping algorithms to extract nouns (Riloff et al. 2003); consideration of web-based mutual information in ranking the subjective adjectives (Baroni and Vegnaduzzo 2004); bootstrapping algorithm employing a small set of seed subjective terms and an online dictionary, plus filtering the candidates based on a similarity measure (Banea, Mihalcea and Wiebe 2008); and morphosyllabic sentiment tagging (Moilanen and Pulman 2008).

A useful sentiment lexicon would contain assignments of polarity orientation (positive and negative) and also the strength of sentiment or, in some cases, the degree of centrality to the sentiment category. To determine the word-level strength of sentiment, Latent Semantic Analysis (LSA) (Turney and Littman 2003), the pointwise mutual information (PMI) technique (Turney and Littman 2003; Read 2004), and methods employing WordNet (Miller 1990) structure relations (Kamps and Marx 2002; Kim and Hovy 2004; Andreevskaia and Bergler 2006) have been proposed.

The subjectivity lexicon developed by Wilson et al. (2005) is comprised by over 8000 subjectivity clues annotated by type (strongly subjective / weakly subjective) and prior polarity (positive/negative/both/neutral). Hatzivassiloglou and McKeown (1997) created a list of 1336 adjectives manually labeled as either positive or negative. They assumed that, given the set of adjectives with predetermined orientation labels (positive or negative) and the pairs of adjectives conjoined using the following conjunctions: '*and*', '*or*', '*but*', '*either-or*', '*neither-nor*', it is possible to predict the orientation of two conjoined adjectives. A log-linear regression model that was automatically constructed based on the constraints on the orientations of conjoined adjectives from the corpus, in combination with the supplementary morphology rules, predicted whether two conjoined adjectives are of the same or opposite orientation with a high level of accuracy (82 percent). Wiebe (2000) proposed a method for identifying strongly subjective adjectives clustered according to distributional similarity. Two bootstrapping algorithms aimed at the generation of the lists of subjective nouns by exploiting the extraction patterns (e.g., '*expressed* <dobj>', '*voiced* <dobj>' etc.), which are discovered to be associated with 20 seed subjective nouns (e.g., '*delight*', '*embarrassment*' etc.), are described in (Riloff et al. 2003). The main assumption behind this bootstrapping approach is that words of the same semantic class appear in similar pattern contexts. Baroni and Vegnaduzzo (2004) proposed to rank a large list of adjectives according to a subjectivity score by employing a small set of manually selected subjective adjectives and computing the mutual information of pairs of adjectives (a seed adjective and an adjective to be ranked) using frequency and co-occurrence frequency counts on the web. To assign subjectivity labels to word senses, methods relying on distributional similarity (Wiebe and Mihalcea 2006) and on semi-supervised minimum cut algorithm (Su and Markert 2009) have been proposed.

Turney and Littman (2003) proposed an approach to measure the semantic orientation of a given word based on the strength of its association with a set of seven context-insensitive positive words (e.g., '*good*', '*excellent*' etc.), minus the strength of its association with a set of seven negative words (e.g., '*bad*', '*poor*' etc.). The researchers compared two different statistical measures of word association, PMI and LSA, and found that the method relying on PMI is less accurate and less stable than the LSA method. The limitations of these statistical methods include: size of the corpora required for good performance, long processing time, and the problem of word sense disambiguation. Kamps and Marx (2002) investigated the measures for affective and emotive aspects of meaning obtained from the structure of the WordNet lexical database. As the meaning of a concept in WordNet is determined by its position relative to other concepts, the researchers have decided to evaluate individual words (specifically, adjectives) by determining their relation (or distance) to the words '*good*' and '*bad*', and assigning the values in the interval [-1, 1]. To determine word-level sentiment, Kim and Hovy (2004) developed two models, which expand the list of a small amount of seed verbs and adjectives, manually annotated with sentiment labels (positive or negative), by means of exploration of the basic semantic relations in WordNet, such as synonymy and antonymy relations. For each word, both positive and negative strengths were computed. Research on mining WordNet for fuzzy sentiment was conducted by Andreevskaia and Bergler (2006), who proposed a method for extracting sentiment-bearing adjectives from WordNet using a set of positive and negative seed words. After expanding the list of seed words by their synonyms, antonyms, and hyponyms found in WordNet, the algorithm processed all WordNet glosses and extracted the terms, which contained in their definitions the sentiment-conveying words from the compiled list.

Motivated by the assumption that 'different senses of the same term may have different opinion-related properties', Esuli and Sebastiani (2006) developed a SentiWordNet lexicon based on WordNet (Miller 1990) synsets comprised from synonymous terms. Three numerical scores (*Obj*(*s*), *Pos*(*s*), and *Neg*(*s*), which range from 0.0 to 1.0 and in sum equal to 1.0), characterizing to what degree the terms included in a synset are objective, positive, and negative, were automatically determined based on the proportion of eight ternary classifiers that assigned the corresponding label to the synsets of adjectives, adverbs, nouns, and verbs by quantitatively analysing the glosses associated with them.

Aimed at introducing the hierarchy of 'affective domain labels', Strapparava and Valitutti (2004) created WordNet-Affect, a lexicon of affective concepts, based on the subset of WordNet synsets. Affective labels for the concepts related to emotional state, moods, traits, situations evoking emotions, or emotional responses were assigned to the WordNet-Affect entries (e.g., '*happy*' – EMOTION, '*aggressiveness*' – TRAIT etc.). The Appraisal lexicon developed by (Argamon et al. 2007) contains adjectives and adverbs annotated by attitude type (affect, judgment, appreciation) and orientation.

Most lexicon-based systems for sentiment analysis face the difficulty of assigning the sentiment scores to words that are not available in their databases. To deal with the limitation in lexicon coverage, we will therefore propose methods to automatically build and expand the sentiment lexicon represented by sentiment-conveying words, which are annotated by sentiment polarity, polarity scores and weights. Although many researchers have already attempted to extract and score new words through synonymy and antonymy relations, the derivation of new sentiment lexemes by manipulation with morphological structure of words, as well as compounding using sentiment-conveying terms as key elements, have not been well explored. To the best of our knowledge, the only works employing morphological analysis for sentiment tagging of words are (Moilanen and Pulman 2008) (English words are transformed and compared with known sentiment lemmas and affixes) and (Ku, Huang and Chen 2009) (the polarity of Chinese opinion-related compound words is predicted based on the analysis of their morphological structure).

In our work, we approach the problem from the opposite direction: based on the sentiment-scored lemmas and types of affixes, new words (adjectives, adverbs, nouns, and verbs) are automatically built and scored. We also developed the algorithm to extract sentiment-conveying compounds, and elaborated the rules for scoring them based on the patterns, according to which the compounds are created. In addition to the sentiment-related entries, we are proposing to collect (1) modifiers and functional words that influence the contextual sentiment (or its strength) of phrases or sentences, and (2) modal operators that play the role of indicators of the confidence degree, with which the opinion or attitude statement is expressed.

1.3.2 Sentiment Analysis in Text

To analyse sentiment (positivity and negativity) of a phrase or a sentence, researchers mainly focused on machine-learning and rule-based linguistic approaches.

Machine-learning approach

An unsupervised statistical method for the task of separating opinions from facts and classifying opinions as positive, negative, or neutral, using Naïve Bayes classifier, proposed by Yu and Hatzivassiloglou (2003), resulted in a high accuracy (up to 91 percent) at a sentence level. The best performance was observed when words, bigrams, trigrams, part-of-speech, and polarity were included in the feature set. The decision on the polarity of a sentence was based on the number and strength of semantically oriented words in the sentence. Kim and Hovy (2005) built a classifier that identified all sentences expressing polarity in a given text based on strong markers of opinion such as certain modal verbs, adjectives and adverbs. To automatically distinguish prior and contextual polarity of individual words and phrases in sentiment expressions, Wilson et al. (2005) employed a machine learning method with not only lexical (e.g., word, modification) and polarity (e.g., negation, polarity shifter) features, but also syntactic structure features.

In order to overcome the problem of strong dependency of machine learning techniques on domain, topic and time, Read (2005) constructed a corpus of text marked-up with emoticons and developed the emoticon-trained classifier aimed at sentiment classification. While this classifier performed well (up to 70 percent accuracy) on the articles extracted from the constructed corpus, it was not very effective in predicting the polarity of movie reviews and news. Read (2005) inferred that there exists the language-style dependency in sentiment classification.

The model of integration of machine learning approach with compositional semantics was proposed by Choi and Cardie (2008). A dependency tree-based method for sentiment classification of Japanese and English subjective sentences using conditional random fields (CRF) with hidden variables was recently introduced by Nakagawa, Inui and Kurohashi (2010). This approach relies on the lexicon (sentiment polarity expressions and polarity reversing words), dependency parser, and a probabilistic model to handle interactions between hidden variables.

Rule-based linguistic approach

To analyse contextual sentiment (polarity) of a phrase or a sentence, rule-based linguistic approaches (Nasukawa and Yi 2003; Mulder et al. 2004; Moilanen and Pulman 2007; Shaikh, Prendinger and Ishizuka 2007; Subrahmanian and Reforgiato 2008) have been proposed.

There is a strong tie between our approach to attitude analysis from text with the work of Moilanen and Pulman (2007) on sentiment composition. In their work, Moilanen and Pulman (2007) propose a theoretical composition model employing deep dependency parsing, sentiment propagation, polarity reversal, and polarity conflict resolution within various linguistic constituent types at various grammatical levels. The experiments with the developed lexical system revealed the crucial dependency on a wide-coverage lexicon, accurate parsing, and sentiment sense disambiguation in a compositional approach to sentiment analysis. The significant difference between our approaches lies in the levels of classification (polarity-based classes, or positive/negative/neutral, in (Moilanen and Pulman 2007) versus fine-grained classes, or nine emotions and four polarity-based labels for judgment and appreciation, and neutral, in our work).

1.3.3 Affect Analysis in Text

Lexical approach

An approach to analysing affect content in free text using fuzzy logic techniques was proposed by Subasic and Huettner (2001). Some researchers employed a keyword-spotting technique to recognize emotions in text (Olveres et al. 1998; Strapparava et al. 2007) or expressed in a multi-modal way (for example, speech signals along with textual content (Chuang and Wu 2004)). This method is fast; however, the use of a purely word-level analysis model cannot cope with cases where affect is expressed by phrases requiring complex phrase/sentence-level analyses, as words are interrelated and influence each other's affect-related interpretation (as in the sentence '*I use the ability to breathe without guilt or worry*'), or when a sentence carries affect indirectly through underlying meaning (for example, '*I punched my car radio, and my knuckle is now bleeding*').

Machine-learning approach

With the aim to classify blog sentences by six basic emotions (Ekman 1993), Aman and Szpakowicz (2008) employed a machine-learning model that utilized corpus-based features (unigrams) and the following emotion lexicons: Roget's Thesaurus (Jarmasz and Szpakowicz 2001) and WordNet-Affect (Strapparava and Valitutti 2004). The text-based emotion prediction problem in the domain of children's fairy tales was explored by Alm, Roth, and Sproat (2005) using a supervised machine-learning approach. As the researchers did not have sufficient training data to classify sentences according to fine-grained distinct emotions, in their preliminary study, Alm et al. (2005) focused only on three categories: neutral, positive emotion, and negative emotion. In her dissertation, Alm (2008) described the refined and improved feature set, and presented the results of experiments on fine-grained emotion classification of text using a hierarchical sequential model. A hierarchy-based machine learning method that considers the relations between neutrality, polarity and emotion of a sentence was implemented by Ghazi et al. (2010). To automatically recognize emotions in news headlines, Katz, Singleton and Wicentowski (2007) employed a supervised system based on unigram model, and Strapparava and Mihalcea (2008) proposed several methods using LSA and Naïve Bayes classifier trained on the corpus of blog posts annotated by emotions. Researchers also applied statistical language modelling techniques to analyse moods conveyed through online diary-like posts (Mishne 2005; Leshed and Kaye 2006; Mihalcea and Liu 2006; Keshtkar and Inkpen 2009).

The weak points of the machine learning methods for sentiment or affect analysis include: large corpora required for meaningful statistics and good performance; dependency on topic and domain; neglect of some prepositions, negation, modal, and condition constructions; disregard of syntactic relations and semantic dependencies in sentences; and long processing time.

Commonsense-based approach

An approach for understanding the underlying semantics of language using large-scale real-world commonsense knowledge was proposed by Liu et al. (2003), who incorporated the affect sensing engine into an affectively responsive email composer called EmpathyBuddy. The architecture of the affect sensing engine includes (1) Model Trainer that consists of three sequential modules: Linguistic Processing Suite that includes part-of-speech tagging, phrase chunking, constituent parsing, subject-verb-object-object identification, and semantic class generalization; Affective Commonsense Filter and Grounder module that filters affective commonsense from the whole corpus using emotion ground keywords and tags emotion keywords with 'grounds' in preparation for training the models; and Propagation Trainer that propagates the affect valence from the emotion grounds to concepts related through commonsense relations; and (2) Text Analyzer that is represented by five sequential modules: Text Segmenter; Linguistic Processing Suite; Story Interpreter; Smoother; and Expressor. Each parsed sentence is evaluated with the prepared models; and weighted scoring function generates a six-tuple score (the classification is based on the Ekman's (1993) set of six basic emotions).

Rule-based linguistic approach

Advanced rule-based linguistic approaches targetting textual affect recognition at the sentence level are described in (Boucouvalas 2003; Chaumartin 2007; Shaikh, Prendinger and Ishizuka 2009).

Boucouvalas (2003) developed the Text-to-Emotion Engine based on word tagging and analysis of sentences. The proposed system uses a small set of emotions, the six basic types defined by Ekman (1993). The emotion extraction engine can analyse input text from a chat environment, identify the emotion communicated, and deliver the parameters necessary to invoke an appropriate expressive image on the user's display. However, the proposed system employs a parser that generates emotional output only if an emotional word refers to the person himself/herself and the sentence is in present continuous or present perfect continuous tense. We believe that such limitations greatly narrow the potential of textual emotion recognition. As a result, sentences like '*Onion pie is disgusting*' and '*It was the most joyous feeling!*' are disregarded by the parser despite the fact that they evidently carry affect.

Chaumartin (2007) developed a rule-based system relying on the lexicon from WordNet-Affect (Strapparava and Valitutti 2004) and SentiWordNet (Esuli and Sebastiani 2006), and applied it to affect sensing in news headlines. The rule-based linguistic approach for textual affect sensing inspired by the rules of the OCC model of emotions (Ortony et al. 1988) was introduced in (Shaikh et al. 2009) and applied in the Emotion Sensitive News Agent system.

The weakness of most affect recognition systems integrated with chat (Olveres et al. 1998; Boucouvalas 2003) or e-mail (Liu et al. 2003) browsers, or analyzing diary-like blogs (Aman and Szpakowicz 2008), is that they do not take into account crucial aspects of informal online conversation such as its specific style and evolving language. In order to account for the peculiarity of online messaging and blogs, and to ensure satisfactory results on real examples, we investigated style, linguistic, and interactional features of online communication (see Section 1.4 for details), and took these into consideration in constructing our Affect Analysis Model and Attitude Analysis Model relying on rule-based linguistic approach.

1.3.4 Attitude Analysis in Text

Early attempt to focus on distinct attitude types (affect, judgment, and appreciation) in the task of textual attitude analysis was made by Taboada and Grieve (2004), who determined a potential value of adjectives for affect, judgement and appreciation by calculating the PMI with the pronoun-copular pairs '*I was (affect)*', '*He was (judgement)*', and '*It was (appreciation)*'. However, affect-conveying adjectives (e.g., '*joyful*', '*depressed*') may equally well occur not only with first person pronouns, but also with third person pronouns, thus describing emotional states experienced by oneself or by other person. Whitelaw, Garg and Argamon (2005) used a machine learning technique (SVM) with fine-grained semantic distinctions in features (attitude type, orientation) in combination with 'bag of words' to classify sentiment of movie reviews. However, the concentration only on adjectives expressing appraisal and their modifiers greatly narrows the potential of the Whitelaw et al. (2005) approach to sentiment analysis.

1.4 Features of Language in Online Communication Media

Many Internet users adopt online communication not only to conduct business but also to keep in touch with their family and friends, to seek emotional support, or to search for new interesting relationships. Nowadays, Instant Messaging (IM) and social networks have proven to be the most popular online applications.

In order to construct a practical and usable attitude-sensing system, we investigated the style of communication and the linguistic and interactional features of real-time conversations. Linguistic features of online communication media (chats, IMs, blogs, discussion forums, social networks), such as emoticons, unconventional spellings, representations of spoken language features, regional dialect features etc. have been extensively studied by the linguists and sociolinguists (Androutsopoulos 2006; Herring 2008). The main problem in messaging is that people cannot easily keep up with the evolving language. Although some of the abbreviations, such as ASAP ('*as soon as possible*'), FYI ('*for your information*'), or TIA ('*thanks in advance*'), are widely known, most of the acronyms are only used within the context of online environments. Examples include: BC ('*because*'), 2l8 ('*too late*'), CUL ('*see you later*'), etc. Participants often use different levels of abbreviations, and hence find it annoying when abbreviations are used without surrounding context to help the correct understanding of their meaning. During the study conducted by Grinter and Eldridge (2001), the teenager subjects reported using several different abbreviations for the same words (for example, '*2moro*', '*2morra*', '*tomor*', and '*2morrow*' for '*tomorrow*'), which makes text messages difficult to parse.

Successful computer-mediated communication, particularly within the IM environment, diary-like blogs, and social networks, depends on the use of various symbolic conventions, such as emoticons (to portray emotion states or communicative behaviour), capital letters or asterisks (to emphasize words), special symbols, etc. Trends show that IM users are increasingly turning to such expressive textual cues to supplement the lack of nonverbal (visual and aural) cues (Hu, Wood, Smith and Westbrook 2004; Derks 2007). Derks (2007) examined the use of emoticons (short symbols that resemble facial displays) in text-based computer-mediated communication, and observed that online messages are often replete with emoticons to fill the conversational gaps and to give additional social and emotional meaning. The study showed that: (1) the most common motives for emoticon use are 'expressing emotion', 'strengthening a message', and 'expressing humour'; (2) most emoticons are used towards friends, as compared to strangers; and (3) more emoticons are used in positive than in negative contexts (spontaneously as well as intentionally) (Derks 2007: 112).

1.5 Contributions of This Work

As was mentioned earlier, the recall of the lexicon-based systems for sentiment analysis strongly depends on the availability of sentiment-conveying words in their databases. In our research we propose methods to automatically build and expand the sentiment lexicon (SentiFul) represented by sentiment-conveying words, which are annotated by sentiment polarity, polarity scores and weights. The main features of the SentiFul are as follows: (1) it is built using methods exploring direct synonymy and antonymy relations, hyponymy relations, morphologic modifications and compounding with known lexical units (the originality and valuable contribution lie in the elaborate patterns/rules for the derivation and compounding processes that have not been considered before); (2) it is larger than the existing lists of sentiment words; (3) it includes polarity scores, in contrast to most existing sentiment dictionaries that lack assignments of degree or strength of sentiment. Our AttitudeFul database contains lexicon necessary for fine-grained attitude analysis; it includes attitude-conveying terms, extensive sets of modifiers, contextual valence shifters, and modal operators, which contribute to robust analysis of contextual attitude and its strength.

In this work, we introduce novel compositional linguistic approach to attitude recognition in text. There are several aspects that distinguish our work from other approaches to sentiment analysis. First, our method classifies individual sentences using fine-grained attitude labels (nine for different affective states, two for positive and negative judgment, and two for positive and negative appreciation), as against other methods that mainly focus on two sentiment categories (positive and negative) or six basic emotions. Next, our Attitude Analysis Model is based on the *compositionality principle*, a novel linguistic approach based on the rules elaborated for semantically distinct verb classes, and a method considering the hierarchy of concepts. Our compositional approach to automatic recognition of fine-grained affect, judgment, and appreciation in text extensively deals with the semantics of terms, which allows accurate and robust automatic analysis of attitude type, and broadens the coverage of sentences with complex contextual attitude. Our method is capable of (1) processing sentences of different complexity, including simple, compound, complex (with complement and relative clauses), and complex-compound sentences, and (2) handling not only correctly written text, but also informal messages written in an abbreviated or expressive manner. Moreover, our Attitude Analysis Model encodes the strength of the attitude through numerical value in the interval [0.0, 1.0], and determines the level of confidence, with which the attitude is expressed.

1.6 Outline of the Thesis

This thesis presents the core of our research on textual attitude analysis, demonstrates the results of this work, and describes the developed applications. The thesis consists of three main parts:

Part I. RECOGNITION OF FINE-GRAINED EMOTIONS IN TEXT

Chapter 2 describes the basis for affective text classification, namely, selection of nine emotion categories, and development of Affect database.

Chapter 3 details the algorithm behind the Affect Analysis Model and provides the examples of emotion sensing in sentences of different complexity.

Chapter 4 demonstrates the results of evaluation of our Affect Analysis Model on two data sets represented by sentences from diary-like blogs.

Part II. RECOGNITION OF AFFECT, JUDGMENT, AND APPRECIATION IN TEXT

Chapter 5 describes the proposed methods for generation of a sentiment lexicon (SentiFul), and the creation of a lexicon for attitude analysis (AttitudeFul).

Chapter 6 explains the core of our compositional linguistic approach to recognition of fine-grained affect, judgment, and appreciation (Attitude Analysis Model).

Chapter 7 describes the evaluation of performance of our Attitude Analysis Model on data sets represented by sentences from different domains: personal stories about life experiences, fairy tales, and news headlines.

Part III. APPLICATIONS

Chapter 8 contains the description of the developed applications (AffectIM, EmoHeart, iFeel_IM!, and web-based @AM interface).

Finally, Chapter 9 concludes this thesis with the discussion of the obtained results and future work.

Part I

Recognition of Fine-Grained Emotions in Text

Chapter 2

Basis for Affective Text Classification

In this Chapter we focus on the basis for affective text classification as an important task in the development of a system for automatic recognition of emotions conveyed in written language.

2.1 Emotion Categories

We had analysed emotion categorizations proposed by theorists and listed in (Ortony and Turner 1990) (see Table 2.1 for details). As the result of our investigation, we have decided to use the subset of emotional states defined by Izard (1971) for affect categorization: 'Anger', 'Disgust', 'Fear', 'Guilt', 'Interest', 'Joy', 'Sadness' ('distress'), 'Shame', and 'Surprise'.

Izard's theory (Izard 1971) postulates the existence of discrete fundamental emotions with their motivational, phenomenological properties, and personal meanings. To support his theory, prominent psychologist presented series of original cross-cultural, developmental, and socio-psychological investigations of facial patterning, emotion recognition, and emotion labelling. According to Izard (1971), there are two ways in which the fundamental emotions can be represented or operationally defined: (1) in facial behavior and (2) with concepts or verbal labels (via words). He proved that 'each of the fundamental emotions, in its pure form, can be represented in a unique pattern of facial activity or facial behavior' and is 'associated with a corresponding set of symbols or verbal labels'. Additionally, Izard (1971) assumed that 'the fundamental emotions are innate, universal phenomena', and evaluated this hypothesis in the light of cross-cultural research.

Table 2.1 A selection of the lists of 'basic' emotions (adapted from (Ortony and Turner 1990))

Reference	Fundamental emotion	Basis for inclusion
Arnold	Anger, aversion, courage, dejection, desire, despair, fear, hate, hope, love, sadness	Relation to action tendencies
Ekman, Friesen, and Ellsworth	Anger, disgust, fear, joy, sadness, surprise	Universal facial expressions
Frijda	Desire, happiness, interest, surprise, wonder, sorrow	Forms of action readiness
Gray	Rage and terror, anxiety, joy	Hardwired
Izard	Anger, contempt, disgust, distress, fear, guilt, interest, joy, shame, surprise	Hardwired
James	Fear, grief, love, rage	Bodily involvement
McDougall	Anger, disgust, elation, fear, subjection, tender-emotion, wonder	Relation to instincts
Mowrer	Pain, pleasure	Unlearned emotional states
Oatley and Johnson-Laird	Anger, disgust, anxiety, happiness, sadness	Do not require propositional content
Panksepp	Expectancy, fear, rage, panic	Hardwired
Plutchik	Acceptance, anger, anticipation, disgust, joy, fear, sadness, surprise	Relation to adaptive biological processes
Tomkins	Anger, interest, contempt, disgust, distress, fear, joy, shame, surprise	Density of neural firing
Watson	Fear, love, rage	Hardwired
Weiner and Graham	Happiness, sadness	Attribution independent

From the nine emotions mentioned, we distinguish three types of affective states: (1) positive ('Interest' and 'Joy'); (2) negative ('Anger', 'Disgust', 'Fear', 'Guilt', 'Sadness', and 'Shame'); and (3) ambiguous ('Surprise', depending on context).

2.2 Affect Database

In this Section we describe Affect database created using MySQL (*http://www.mysql.com*) to support the handling of abbreviated language and the interpretation of affective features of emoticons, abbreviations, and words by an automatic emotion recognition system.

2.2.1 Building the Lexical Resource

Our Affect database includes the following tables: Emoticons, Abbreviations, Interjections, Adjectives, Adverbs, Nouns, Verbs, and Modifiers.

For the accumulation of relevant and most often used emoticons and abbreviations (along with their transcriptions), we employed five online dictionaries dedicated to and describing such data. Only entries that occurred in at least three sources were selected. In this way, we collected 364 emoticons, both of American and Japanese style (for example, ‘:”>’ and ‘=^_^=’ for ‘*blushing*’), and the 337 most popular acronyms and abbreviations, both emotional and non-emotional (for example, ‘BL’ for ‘*belly laughing*’, ‘gj’ for ‘*good job*’, and ‘4U’ – ‘*for you*’). As interjections, such as ‘*alas*’, ‘*wow*’, ‘*yay*’, ‘*ouch*’, etc. are specific indicators of communicated emotion caused by unexpectedness, a long-awaited joyful event, or pain, they were collected as well.

The next category consists of words conveying affective content. People use emotion words in particular contexts to negotiate aspects of social reality and to create that reality. Lutz (1988) argues that ‘in particular cultures and contexts, emotion words may be used to theorize about events, to moralize about or to judge them, and to advance one’s interests by defining the situation in a particular way’. From WordNet-Affect (Strapparava and Valitutti 2004), we have taken 1627 words — adjectives (635), nouns (521), verbs (274), and adverbs (197) — that refer directly to emotions, moods, traits, cognitive states, behaviour, attitudes, and sensations. Moreover, we added to our database 434 words that carry the potential to elicit affective states in humans (for example, ‘*beautiful*’, ‘*disaster*’, ‘*break*’, ‘*deceive*’, ‘*violate*’ etc.). These words are considered as indirect emotion words describing the objects and situations that lead to some emotional reactions.

Further, we included 112 modifiers (e.g., ‘*very*’, ‘*extremely*’, ‘*slightly*’, ‘*hardly*’, ‘*less*’, ‘*not*’, etc.) into our database, as they influence the strength of related words and phrases in a sentence.

2.2.2 Annotations of Database Entries

Three independent annotators (non-native English speakers studying at the Graduate School of Information Science and Technology, the University of Tokyo) were asked to manually label the entries of the database using nine emotion categories (‘Anger’, ‘Disgust’, ‘Fear’, ‘Guilt’, ‘Interest’, 'Joy', 'Sadness', 'Shame', and 'Surprise') and intensities. Emotion intensity values range from 0.0 to 1.0, and describe the intensity degree of affective states from 'very weak' to 'very strong'. Annotators conformed to our guidelines with the description of emotional state gradation within intensity levels (see Appendix A for details). For example, '*cheerful*', '*glad*', '*happy*', '*joyful*', and '*elated*' all correspond to the 'Joy' emotional state, but to a different degree of intensity (0.2, 0.4, 0.6, 0.8, and 1.0, correspondingly).

2.2.2.1 Annotations of Emoticons and Abbreviations

Emoticons and abbreviations were transcribed and related to named affective states (with intensity), whereby each entry was assigned to only one category. The inter-rater agreement on the assigned category was calculated using Fleiss' Kappa statistics (Fleiss 1971). Fleiss' Kappa works for any number of rates giving categorical ratings to a fixed number of items. It can be interpreted as expressing the extent to which the observed amount of agreement among raters exceeds what would be expected if all raters made their ratings completely randomly. Shortly, Kappa gives a measure for how consistent the ratings are. The Kappa coefficient k is defined as:

$$k = \frac{\overline{P} - \overline{P}e}{1 - \overline{P}e}, \tag{1}$$

where $\overline{P}$ is the relative observed agreement among raters; and $\overline{P}e$ is the probability that agreement is due to chance. The Kappa scoring ranges between 0.0 and 1.0, poor and complete agreement, respectively.

The measured Kappa coefficients for emoticons and abbreviations are 0.94 and 0.93, respectively, showing strong annotation reliability (supposedly, due to unambiguous transcriptions provided along with these symbolic cues).

The percentage distributions of emoticons and abbreviations according to resulting affective labels (majority vote) are as follows (in descending order):

(1) Emoticons: 'Joy' – 45, 'Sadness' – 23, 'Fear' – 11, 'Anger' – 7, 'Surprise' – 7, 'Disgust' – 3, 'Shame' – 3, 'Interest' – 1, and 'Guilt' – 0 percent.

(2) Abbreviations: 'Joy' – 74, 'Guilt' – 7, 'Surprise' – 6, 'Disgust' – 4, 'Fear' – 3, 'Anger' – 2, 'Sadness' – 2, 'Shame' – 2, and 'Interest' – 0 percent.

In the resulting intensity estimation for each affect-related entry, variance of data from the annotator mean was taken into consideration. In statistics, the variance is considered as a measure of spread, that is how far the values deviate from the mean. The variance σ^2 of a set of values (in our case, intensity values given by three annotators) is defined as:

$$\sigma^2 = \frac{\sum (x - \bar{x})^2}{n}, \tag{2}$$

where $\bar{x}$ is the mean; n is the number of data values ($n = 3$); and x stands for each data value in turn.

If the variance was less than or equal to 0.027, the resulting intensity was measured as the average of intensities given by three annotators. Otherwise, the intensity value responsible for exceeding the variance threshold was removed, and only the remaining values were taken into account.

By way of example, some emoticons and abbreviations extracted from the developed database are listed in Table 2.2.

Table 2.2 Examples of emoticons and abbreviations extracted from Affect database

Type	Symbolic representation	Meaning	Category	Intensity
Emoticons (American style)	:-)	happy	Joy	0.6
	:-o	surprise	Surprise	0.8
	:-S	worried	Fear	0.4
Emoticons (Japanese style)	\(^O^)/	very excited	Joy	1.0
	(~_~)	grumpy	Anger	0.3
	m(._.)m	bowing, thanks	Thanks	-
Abbreviations	JK	just kidding	Joy	0.3
	4gv	forgive	Guilt	0.6
	PPL	people	-	-

2.2.2.2 Annotations of Affect-Related Words

Considering the fact that some affective words may express more than one emotion state, annotators could relate words to more than one category. For instance, in the annotation of the word '*frustrated*', both 'Anger' and 'Sadness' emotions are involved, with intensities 0.2 and 0.7, respectively (Table 2.3).

Table 2.3 Examples of words taken from Affect database

Affective word	Part of speech	Category	Intensity
enthusiasm	Noun	Interest	0.8
		Joy	0.5
astonished	Adjective	Surprise	1.0
frustrated	Adjective	Anger	0.2
		Sadness	0.7
discomfit	Verb	Anger	0.1
		Sadness	0.7
		Shame	0.3
remorsefully	Adverb	Guilt	0.8
		Sadness	0.5

Assignments of emotion labels to the same word might differ among annotators. We faced the difficulty of employing Fleiss' Kappa coefficient (Fleiss 1971) to measure inter-rater agreement here, as the important requirement of using it is that each entry needs to be assigned to only one of possible categories. For the resulting labelling, we only considered emotion categories that occurred in the assignments of at least two annotators. The most frequent emotion labels in resulting sets were 'Joy' and 'Sadness' (34.3 percent and 30.0 percent of overall number of affective words, respectively) whereas the least frequent was 'Guilt' (3.1 percent). The distribution of affective words with one, two, and three emotion labels is 67 percent, 29 percent, and 4 percent, respectively. Only one word (adjective '*aggravated*') was annotated by four resulting emotion labels ('Anger:0.5', 'Disgust:0.5', 'Sadness:0.3', 'Fear:0.1').

Regarding the emotion intensity annotations of affective words, we observed interesting statistics within each of the nine emotion categories. The percentages of cases with valid (not exceeding the threshold) variance of given intensities within each emotion category are as follows (in descending order): 'Shame' – 57.8, 'Guilt' – 51, 'Anger' – 49.8, 'Fear' – 42.7, 'Disgust' – 39.2, 'Surprise' – 27.8, 'Sadness' – 26.6, 'Joy' – 18.8, and 'Interest' – 8.6 percent. The annotators easily agreed in intensity assignments to 'Shame', 'Guilt', and 'Anger' categories, in contrast to frequent disagreement in cases of 'Interest', 'Joy', and 'Sadness'. We can only speculate that disagreement is related to the huge diversity of 'joyful' and 'sad' synonymous words with different emotional colorations, and due to the fuzziness of the 'interest' concept (some of psychologists do not consider 'interest' as an emotional state at all).

As to the indirect affective words that possibly induce emotional states through indication of emotional causes or responses (notions of direct and indirect affective words were used by Strapparava et al. (2006)), about 300 nouns from WordNet-Affect (Strapparava and Valitutti 2004) along with the categories, to which they correspond, were kindly provided by Dr. Alessandro Valitutti. Some examples are given in Table 2.4.

Table 2.4 Examples of indirect emotion nouns from WordNet-Affect and their annotations

Noun	**Emotion categories from WordNet-Affect with their annotations from our Affect database in brackets [ ]**	**Resulting emotion labels and intensities defined automatically**
brave	Pride – [Joy:0.4] Admiration – [Joy:0.6; Surprise:0.5]	[Joy:0.6; Surprise:0.5]
refusal	Sadness – [Sadness:0.9] Anger – [Anger:0.9] Resentment – [Anger:0.6] Disappointment – [Sadness:1.0]	[Anger:0.9; Sadness:1.0]
well-being	Satisfaction – [Joy:0.3] Joy – [Joy:0.9]	[Joy:0.9]

In order to label indirect emotion words using our fine-grained categories, the annotations of direct affective words that (1) represent emotion labels in WordNet-Affect and (2) are already included in Affect database (see annotations in brackets [ ] in the middle column of Table 2.4) were automatically analysed. The maximum intensity within the same emotion label from Affect database was taken as the resulting intensity for that emotion state in final annotations (see some results in the last column of Table 2.4).

2.2.2.3 Assignments of Coefficients to Modifiers

Adverbs of degree have an impact on neighbouring verbs, adjectives, or another adverb, and are used to mark that the extent or degree is either greater or less than usual (Biber, Johansson, Leech, Conrad, Finegan and Quirk 1999). In (Benamara, Cesarano, Picariello, Reforgiato and Subrahmanian 2007), the authors use adverbs of degree to modify the score of adjectives in sentiment analysis. In our work, such adverbs along with some of the prepositions constitute the set of modifiers. Two annotators gave coefficients for intensity degree strengthening or weakening (from 0.0 to 2.0) to them, and the result was averaged (Table 2.5).

Table 2.5 Examples of modifiers with coefficients of intensity degree strengthening or weakening

Category (coefficient range)	**Modifier**	**Coefficient**
Adverb of affirmation (from 1.0 to 2.0)	*certainly*	1.2
Adverb of doubt (from 0.0 to 1.0)	*arguably*	0.5
Strong intensifying adverb (from 1.0 to 2.0)	*immensely*	1.8
Weak intensifying adverb (from 0.0 to 1.0)	*slightly*	0.2
Negation (0.0)	*hardly*	0.0

Chapter 3

Affect Analysis Model (AAM)

This Chapter describes a novel rule-based linguistic approach to fine-grained affect recognition from text. Typically, researchers in the sentiment analysis field deal with grammatically and syntactically correct textual input. By contrast, our analysis of affect expressed through written language is inspired by the evolving language, style, and specifics of Instant Messaging conversations and diary-like blog posts. Our Affect Analysis Model copes with not only correctly written text, but also informal messages written in an abbreviated or expressive manner (Neviarouskaya et al. 2007a, 2007c, 2011). The proposed algorithm consists of five main stages: (1) symbolic cue analysis; (2) syntactic structure analysis; (3) word-level analysis; (4) phrase-level analysis; and (5) sentence-level analysis. The architecture of the Affect Analysis Model (AAM) is presented in Figure 3.1. Our method is capable of processing sentences of different complexity, including simple, compound, complex (with complement and relative clauses), and complex-compound sentences. Affect in text is classified into nine emotion categories (‘Anger’, ‘Disgust’, ‘Fear’, ‘Guilt’, ‘Interest’, ‘Joy’, ‘Sadness’, ‘Shame’, and ‘Surprise’),or neutral.

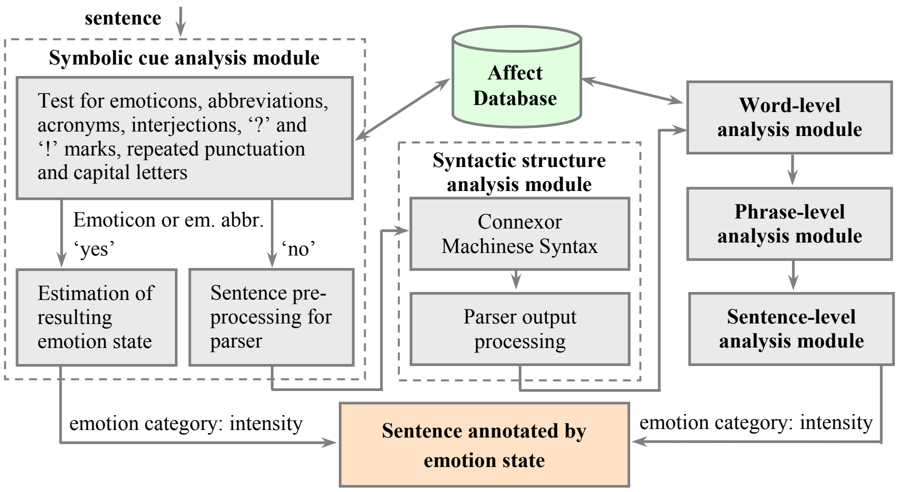

Figure 3.1 Architecture of the Affect Analysis Model

3.1 Symbolic Cue Analysis

In the *first stage*, the sentence is tested for occurrences of emoticons, abbreviations, acronyms, interjections, 'question mark' and 'exclamation mark', repeated punctuation, and capital letters. First, punctuation marks in a sentence are delimited from words in order to disambiguate sentence punctuation marks from those belonging to emoticons. The detected 'exclamation mark', repeated punctuation, and capital letters are considered as an emphasis of the communicated emotion.

If there is an emoticon or abbreviation related to an emotional state, no further analysis of affect in text is performed based on the simplifying assumption that the emoticon (or abbreviation) dominates the affective meaning of the entire (simple or compound) sentence. It is known that people type emoticons and emotional abbreviations to show actual feeling (e.g., '*I have taken the exams timetable already :S* [worry; Fear:0.4]'), or to avoid misleading the other participants, for instance, after irony, joke, or sarcasm (e.g., '*Thank you so much for your kind encouragement :-(* [Sadness:0.8]' or '*If you miss the meeting, I will hunt you down and murder you :)* [Joy:0.6]'). In a face-to-face communication sarcasm is conveyed by a positive tone or a smile and a negative message (Planalp and Knie 2002). Similarly, emoticons 'can create ambiguity and express sarcasm online by varying the polarity of the emoticon and the polarity of the message' (Derks 2007: 63). On the other hand, if there are multiple emoticons or emotion-relevant abbreviations in the sentence, we determine the prevailing (or dominant) emotion based on the following two (simplifying) rules:

(1) When emotion categories of the detected emoticons (or abbreviations) are the same (e.g., ‘*G* [grin; Joy:0.6] *it is nice song too ;-)* [winking; Joy:0.3]’), the higher intensity value is taken for this emotion.

(2) When they are different (e.g., ‘*I did not save that song :S* [worry; Fear:0.4]*, please send it once more ;”>* [blushing; Shame:0.5]’), the category (and intensity) of the emoticon occurring last is considered dominant.

As interjections are added to text to reflect an author’s feelings, as in the sentences ‘*Oh no, I forgot that the exam was today!*’ and ‘*But anyways, yay!*’, they are analysed as well. In case of an interrogative sentence, we process it further at subsequent stages in order to identify whether the question expresses a strong emotion or not. While some researchers (Boucouvalas 2003) ignore such sentences, we believe that questions, like ‘*Why do you irritate me so greatly?*’ may carry emotional content.

It is important to emphasize here that we distinguish two ways of assigning an emotional value to the sentence. In one case (as described above), the affective information is provided by emotion-related emoticons or abbreviations, and in the other one by the lexical meaning propagated through rules to the sentence level. If there are no emotion-relevant emoticons or abbreviations in a sentence, we prepare the sentence for parser processing by replacing non-emotional abbreviations and acronyms by their proper transcriptions found in the database (e.g., ‘*I m* [am] *stressed bc* [because] *i have frequent headaches*’). In such a way, the issue of correct processing of abbreviated text by syntactic parser is resolved.

3.2 Syntactic Structure Analysis

The *second stage* of Affect Analysis Model algorithm is devoted to syntactic structure analysis, and it is divided into two main subtasks:

(1) Sentence analysis by the syntactic parser, Connexor Machinese Syntax (*http://www.connexor.eu/technology/machinese/machinesesyntax/*), developed by the Connexor Oy company.

(2) Parser output processing.

Connexor Machinese Syntax provides a full analysis of texts by showing how words and concepts relate to each other in sentences, with competitive speed and accuracy. This tool assigns meaning-oriented syntactic structure to text, thus helping analytic applications understand text beyond the level of words, phrases, and entities. The parser returns exhaustive information for analysed sentences, including word base forms (lemmas), parts of speech, dependency functions representing relational information between words in sentences, syntactic function tags, and morphological tags. An example of Connexor Machinese Syntax output for the sentence '*Chewy colorful bears are like a tasty little rainbow.*' is shown in Table 3.1.

Table 3.1 An example of Connexor Machinese Syntax output

Token id	Text	Lemma	Syntactic relations and dependencies	Syntax and morphology
1	*Chewy*	*chewy*	attr:>2	@A> %>N A ABS
2	*colorful*	*colorful*	attr:>3	@A> %>N A ABS
3	*bears*	*bear*	subj:>4	@SUBJ %NH N NOM PL
4	*are*	*be*	main:>0	@+FMAINV %VA V PRES
5	*like*	*like*	man:>4	@ADVL %EH PREP
6	*a*	*a*	det:>9	@DN> %>N DET SG
7	*tasty*	*tasty*	attr:>8	@A> %>N A ABS
8	*little*	*little*	attr:>9	@A> %>N A ABS
9	*rainbow*	*rainbow*	pcomp:>5	@<P %NH N NOM SG
10	.	.		

When handling the parser output, we represent the sentence as a set of primitive clauses (either independent or dependent). Each clause might include Subject formation (SF), Verb formation (VF) and Object formation (OF), each of which may consist of a main element (subject, verb, or object) and its attributives and complements. The developed algorithm can detect not only subjects represented by noun phrases, but also subjects represented by gerund (non-finite verb form) as in the sentence '*Walking on the beach is a pleasure*', by an infinitive as in the sentence '*To offend the youngest child is an obscene action*', or by a full clause, introduced by '*that*', itself containing a subject and a predicate like in the sentence '*That tomorrow weather will be sunny is great*'. For the processing of complex or compound sentences, we build a so-called 'relation matrix', which contains information about dependences that the verbs belonging to different clauses have.

3.3 Word-Level Analysis

After handling the result from the previous analysis stage, the system transfers the data to the *third stage*, word-level analysis. For each word (found in Affect database) of a sentence, the affective features of a word are represented as a vector of emotion state intensities e = [Anger, Disgust, Fear, Guilt, Interest, Joy, Sadness, Shame, Surprise]. Here are three examples: e('*rude*') = [0.2,0.4,0,0,0,0,0,0,0]; e('*brotherly*') = [0,0,0,0,0,0.2,0,0,0]; and e('*love*') = [0,0,0,0,0.8,1.0,0,0,0]. In the case of a modifier, the system identifies its coefficient (e.g., coeff('*barely*') = 0.4).

Adjectives and adverbs have two forms that indicate degrees of comparison: comparative form and superlative form. The comparative form, which is made by adding '-*er*' or a preceding '*more*' to the positive form, either shows a greater degree than the positive form or makes a comparison between two persons or things. The superlative form, which is made by adding '-*est*' or a preceding '*most*' to the positive form, indicates the greatest degree of a quality or quantity among three or more persons or things. As our Affect database contains words only in their dictionary form, one important system function at this stage is to increase the intensity of the emotion vector of an adjective (e.g., e('*glad*') = [0,0,0,0,0,0.4,0,0,0]), or emotional adverb, if it is in comparative or superlative form, by multiplication by values 1.2 or 1.4, respectively (e.g., e('*gladder*') = [0,0,0,0,0,0.48,0,0,0] and e('*gladdest*') = [0,0,0,0,0,0.56,0,0,0]). Two persons were involved in the procedure of defining these multipliers. After annotators had manually assigned intensities to the set of words (e.g., '*good*', '*better*', '*best*'), multipliers were derived from the averaged assignments.

3.4 Phrase-Level Analysis

In the *fourth stage*, phrase-level analysis is performed. The purpose of this stage is to detect emotions involved in phrases, and then in Subject, Verb, and Object formations (for definitions, see Section 3.2). Words in a sentence are interrelated and, hence, each of them can influence the overall meaning and affective bias of a statement. We have defined rules for processing general types of phrases with regard to affective content:

(1) Adjective phrase: modify the vector of adjective (e.g., e('*extremely doleful*') =

coeff(‘*extremely*’) * e(‘*doleful*’) = 2.0 * [0,0,0,0,0,0,0.4,0,0] = [0,0,0,0,0,0,0.8,0,0]).

(2) Noun phrase: output vector with the maximum intensity within each corresponding emotional state in analysing vectors (e.g., e1=[0..0.7..] and e2=[0.3..0.5..] yield e3=[0.3..0.7..]). For instance, e(‘*brotherly love*’) = [0,0,0,0,0.8,1.0,0,0,0] where e(‘*brotherly*’) = [0,0,0,0,0,0.2,0,0,0] and e(‘*love*’) = [0,0,0,0,0.8,1.0,0,0,0]. In the rare case of words with opposite polarities, the resulting vector will contain mixed emotions (e.g., e(‘*annoying care*’) = [0.3,0,0,0,0.2,0.2,0,0,0] where e(‘*annoying*’) = [0.3,0,0,0,0,0,0,0,0] and e(‘*care*’) = [0,0,0,0,0.2,0.2,0,0,0]).

(3) Verb plus adverbial phrase: output vector with the maximum intensity within each corresponding emotional state in analysing vectors (e.g., e(‘*to shamefully deceive*’) = [0,0.4,0,0,0,0,0.5,0.7,0] where e(‘*shamefully*’) = [0,0,0,0,0,0,0,0.7,0] and e(‘*to deceive*’) = [0,0.4,0,0,0,0,0.5,0,0]).

(4) Verb plus noun phrase: if verb and noun phrase have opposite polarities (e.g., ‘*to deceive hopes*’, ‘*to enjoy bad weather*’), consider vector of verb as dominant (for instance, e(‘*to deceive hopes*’) = [0,0.4,0,0,0,0,0.5,0,0] where e(‘*to deceive*’) = [0,0.4,0,0,0,0,0.5,0,0] and e(‘*hope*’) = [0,0,0,0,0.1,0.3,0,0,0]); if polarities are the same (e.g., ‘*to celebrate victory*’, ‘*to hate crying*’), output vector with maximum intensity in corresponding emotional states.

(5) Verb plus adjective phrase (e.g., ‘*is very kind*’, ‘*feel bad*’): output vector of adjective phrase, as adjectives can come only after ‘stative’ verbs, which do not express actions, and they always refer to and qualify the subject of the sentence.

The rules for modifiers are as follows:

(1) Adverbs of degree increase or decrease emotional intensity values.

(2) Negation modifiers such as ‘*no*’, ‘*not*’, ‘*never*’, ‘*any*’, ‘*nothing*’ and connector ‘*neither…nor*’ cancel (set to zero) vectors of the related words, i.e., ‘neutralize the emotional content’ (e.g., positive vector of ‘*exciting*’ is neutralized due to ‘*nothing*’ in ‘*Yesterday I went to a party, but nothing exciting happened there*’). We use this rule as an initial heuristic, as it is problematic to find pairs of opposite emotions (except for ‘Joy’ and ‘Sadness’), in contrast to straightforward reversing of the polarity in polarity-based classification.

(3) Prepositions such as '*without*', '*except*', '*against*', '*despite*' cancel vectors of related words (for example, the phrase '*despite his endless demonstrations of rude power*' and the sentence '*I climbed the mountain without fear*' are neutralized due to prepositions).

Statements beginning with words like '*think*', '*believe*', '*sure*', '*know*', '*doubt*' or with modal verbs such as '*can*', '*may*', '*must*', '*need*', '*would*' etc. are not considered by our system, as they express a modal attitude towards the proposition. Conditional clause phrases beginning with '*after*', '*although*', '*as if*', '*as though*', '*before*', '*even if*', '*even though*', '*if*', '*if only*', '*unless*', '*whether*', '*when*', '*whenever*', etc. are disregarded as well (e.g., '*I eat when I'm angry, sad, bored…*', or '*If only my brain was like a thumbdrive, how splendid it would be*').

Each of the Subject, Verb, or Object formations may contain words conveying emotional meaning. During this stage, we apply the described rules to phrases detected within formation boundaries. Finally, each formation can be represented as a unified vector encoding its emotional content.

3.5 Sentence-Level Analysis

In the *fifth* and *final stage*, the overall emotion of a sentence and its resulting intensity degree are estimated. Our algorithm enables processing of different types of sentences, such as: simple, compound, complex (with complement or relative clauses), or complex-compound.

3.5.1 Emotion Vector of a Simple Sentence (or a Clause)

The emotion vector of a simple sentence (or a clause) is generated from Subject, Verb, and Object formation (SF, VF, and OF, respectively) vectors resulting from phrase-level analysis. The main idea here is to first derive the emotion vector of Verb-Object formation relation. It is estimated based on the 'verb plus noun phrase' rule described above. In order to apply this rule, we automatically determine polarities of Verb and Object formations using their unified emotion vectors (particularly, non-zero-intensity emotion categories). For instance, polarity of '*to calm disobedient child*' is positive based on polarity of a verb, which dominates negative polarity of object '*disobedient child*'.

The estimation of the emotion vector of a clause (Subject plus Verb-Object formations) is then performed in the following manner:

(1) If polarities of Subject formation and Verb formation are opposite (e.g., Subject formation = '*my darling*', Verb formation = '*smashed*', Object formation = '*his guitar*'; or Subject formation = '*troubled period*', Verb formation = '*luckily comes to an end*'), we consider the vector of the Verb-Object formation relation as dominant. For example, negative Subject formation '*mother's disapproval*' and positive Verb formation '*calmed*' in a sentence '*Mother's disapproval calmed disobedient child*' yield domination of positive emotion vector of '*calmed disobedient child*'.

(2) Otherwise, we output the vector with maximum intensities in corresponding emotional states of vectors of Subject and Verb-Object formations.

Let us consider the processing of Subject formations themselves containing Verb-Object formation (for example, '*To offend the neighbour*' in '*To offend the neighbour is an unfriendly behaviour*') or a full clause, Subject plus Verb-Object formations (for example, '*tomorrow weather will be sunny*' in '*That tomorrow weather will be sunny is great*'). In such cases, first we estimate the emotion vector of main Subject formation, formed by Verb-Object formation or an embedded clause, using rules described above. Then, we estimate the resulting emotion vector of a whole sentence.

It is important to note that our system enables the differentiation of the strength of the resulting emotion depending on the tense of a sentence and availability of first person pronouns. We introduce this idea based on our findings from the literature on psychology studies. Taking tense into account is very important, as 'emotions typically occur in response to an event, usually a social event, **real**, **remembered**, **anticipated**, or **imagined**' (Ekman 1993: 386) (emphasis added by authors). As Ekman states, 'sometimes when people give an account of an emotional experience they unexpectedly begin to re-experience the emotion' (Ekman 1993: 392). The genuine emotion expressions display that an emotion is now felt, whereas so-called referential expressions occur most often when people talk about past or future emotional experiences. Therefore, we assume that the strength of emotions conveyed by text depends on tense (e.g., strongest emotion for present tense, weakened emotion for past tense, and the weakest emotion for future tense).

As to first person pronouns, people tend to use them to 'more directly portray the speaker as the experiencer of the emotion' (Lutz 1990), and to underline the strength of an emotion. Many researchers neglect these phenomena. They ignore the difference between '*I am charmed by the cherry blossoms of Japan*' versus '*The* c*herry blossoms of Japan are charming*' (we think that emotion conveyed through the first sentence is stronger than in the case of the second one), and some of them completely disregard sentences in past or future tense and without first person pronouns (Boucouvalas 2003).

According to our proposal, the emotion vector of a simple sentence (or of a clause) is multiplied by the corresponding empirically determined coefficient of intensity correction (Table 3.2).

Table 3.2 Coefficients of intensity correction

Tense	**First person pronouns (FPP)**	
	yes	no
present	1 ('*My vase is broken*')	0.8 ('*She is annoying*')
past	0.8 ('*He made me angry*')	0.4 ('*It was the most joyous feeling*')
future	0.4 ('*I will enjoy the trip to Egypt*')	0 ('*The game will definitely bring them triumph*')

The dominant emotion of the sentence is determined according to the emotion state with the highest intensity within the emotion vector. However, if there are several emotion states with the same maximum intensity in the resulting vector, we use a function that selects the prevailing emotion randomly. Let us consider the example of processing the following simple sentence: '*My darling smashed his favourite guitar without regret*' (Figure 3.2), where emotion vector e = [Anger, Disgust, Fear, Guilt, Interest, Joy, Sadness, Shame, Surprise]; SF, VF, and OF mean Subject, Verb, and Object formations, respectively; the superscripts ^{0}, $^-$, and $^+$ indicate 'neutral', 'negative', and 'positive' polarities, respectively.

	word:	word-level:	phrase-level:	
SF:	*my*	e^0 = [0,0,0,0,0,0,0,0,0]	e^+ = [0,0,0,0,0,0.7,0,0,0]	
	darling	e^+ = [0,0,0,0,0,0.7,0,0,0]		
VF:	*smashed*	e^- = [0,0,0.6,0,0,0,0.8,0,0]	e^- = [0,0,0.6,0,0,0,0.8,0,0]	e^- = [0,0,0.6,0,0,0,0.8,0,0]
	without	modif. coeff=0.0	e^0 = [0,0,0,0,0,0,0,0,0]	
	regret	e^- = [0,0,0,0.2,0,0,0.1,0,0]		
OF:	*his*	e^0 = [0,0,0,0,0,0,0,0,0]	e^0 = [0,0,0,0,0,0,0,0,0]	e^+ = [0,0,0,0,0,0.6,0,0,0]
	favourite	e^+ = [0,0,0,0,0,0.6,0,0,0]	e^+ = [0,0,0,0,0,0.6,0,0,0]	
	guitar	e^0 = [0,0,0,0,0,0,0,0,0]	e^0 = [0,0,0,0,0,0,0,0,0]	

sentence-level:	
1.	(SF^+ and VF^-) yields domination of (VF and OF);
2.	(VF^- and OF^+) yields domination of VF;
3.	e (sentence) = e (VF^-) = [0,0,0.6,0,0,0,0.8,0,0];
4.	e (sentence) * coeff (tense:'past'; FPP:'yes') = [0,0,0.6,0,0,0,0.8,0,0] * 0.8 = [0,0,0.48,0,0,0,0.64,0,0]
5.	result ('*My darling smashed his favourite guitar without regret*'): Sadness:0.64.

Figure 3.2 Example of affect sensing in a simple sentence

3.5.2 Emotion Vector of a Compound Sentence

A compound sentence is composed of at least two independent clauses, but no dependent clauses. The clauses are joined by a comma and coordinate connector, or a semicolon with no conjunction. In order to estimate the emotion vector of a compound sentence, first, we evaluate the emotion vectors of its independent clauses. Then, we define the resulting vector of the compound sentence based on the following rules:

(1) With comma and coordinate connectors '*and*' and '*so*' (e.g., '*It is my fault, and I am worrying about consequences*', '*Exotic birds in the park were amazing, so we took nice pictures*'), or with a semicolon with no conjunction: output the vector with the maximum intensity within each corresponding emotional state in the resulting vectors of both clauses.

(2) With coordinate connector '*but*' (e.g., '*They attacked, but we luckily got away!*', '*It was hard to climb a mountain all night long, but a magnificent view rewarded the traveler in the morning*'): the resulting vector of a clause following after the connector is dominant.

3.5.3 Emotion Vector of a Complex Sentence

A complex sentence is a sentence with an independent clause and at least one dependent (embedded or subordinating) clause. The dependent clause is introduced by either a subordinate conjunction (e.g., '*as*', '*because*', '*if*', '*since*', '*that*', etc.) or a relative pronoun such as '*who*' or '*which*'. Some subordinating conjunctions, when used to introduce a phrase instead of a full clause become prepositions with identical meanings. In Section 3.4 we mentioned that in our Affect Analysis Model conditional clause phrases are neutralized due to specific prepositions or conjunctions. Therefore, the emotion vector of a dependent clause starting with one of these conjunctions represents a zero vector, and the vector of the independent clause forms the resulting emotion vector of such a complex sentence. If the subordinating clause in the complex sentence is connected to an independent clause through conjunctions such as '*as*', '*because*', '*since*', we take the maximum intensity within each corresponding emotional state in the resulting vectors of both clauses for the estimation of the resulting vector of the complex sentence.

We can distinguish two types of embedded clauses:

(1) Complement clauses.

(2) Relative clauses.

3.5.3.1 Sentences with Complement Clauses

Let us first look at the case of sentences with complement clauses. Special subordinating conjunctions, so-called complementizers (e.g., '*whether*', '*that*' etc.), introduce complement clauses (for example, '*I wonder whether we will go to the amusement park next weekend*' and '*We hope that you feel comfortable*'). There are basically three complementizers in English language: '*that*', '*for-to*' ('*for*' precedes the complement sentence and the '*to*' precedes the auxilliary constituent of the complement sentence), and what is known as 'POSS-ing' ('POSS' means the possessive suffix, which is affixed to the noun, and the '-ing' means the suffix attached to a verb stem) (Cairns H.S. and Cairns C.E. 1976: 58-62). Here are some examples below:

(1) With '*that*': '*Sam preferred that John take the blame*'.

(2) With '*for-to*': '*Sam preferred for John to take the blame*'.

(3) With 'POSS-ing': '*Sam preferred John's taking the blame. John resented Sam's telling the truth*'.

In order to process a sentence with a complement clause, first we derive the emotion vector of the complement clause (e.g., '*John take the blame*', '*John to take the blame*', '*John's taking the blame*' or '*Sam's telling the truth*'), then create Object formation for the main clause using this vector, and finally estimate the resulting emotion vector of the main clause with added Object formation. In brief, we represent such sentence as a simple one, using the following pattern: 'who-subject does-verb what-object', where object is represented as a complement clause.

3.5.3.2 Sentences with Relative Clauses

Our program is also able to process the complex sentences containing adjective (relative) clauses introduced by '*who*', '*whom*', '*whose*', '*that*', '*which*', or '*where*'. An adjective clause is a dependent clause that modifies a noun. Depending on the role (subject or object) that the relative pronoun plays in the embedded clause, sentences are called 'subject relatives' (see examples 1 and 2 below) or 'object relatives' (see example 3) (Cairns H.S. and Cairns C.E. 1976: 58-62). The following are examples of complex sentences with relative clauses:

(1) '*The wolf who ate the grandmother scared Little Red Riding Hood*'.

(2) '*The wolf who ate the grandmother who lived in the cottage scared Little Red Riding Hood*' (a case of multiple embedding).

(3) '*The wolf who the woodman killed scared Little Red Riding Hood*'.

In our algorithm, the sentences of such type are analysed in the following manner:

(1) First, the emotion vector of adjective clause is estimated.

(2) Then, this emotion vector is added to the Subject or Object formation of the main clause depending on the role of the word to which the adjective clause relates. For example, in a sentence '*The man who loved the woman robbed the bank*', the adjective clause '*who loved the woman*' relates to the subject '*man*'; and in the sentence '*The man robbed the bank where his beloved wife was working*', the adjective clause '*where his beloved wife was working*' relates to the object '*bank*'.

(3) Finally, the emotion vector of the whole sentence is estimated.

Figure 3.3 illustrates the steps by way of a complex sentence with multiple embedding of relative clauses: '*The policeman who loved his job ruined the life of the art student who had stolen famous painting, which was created by his favourite artist*'.

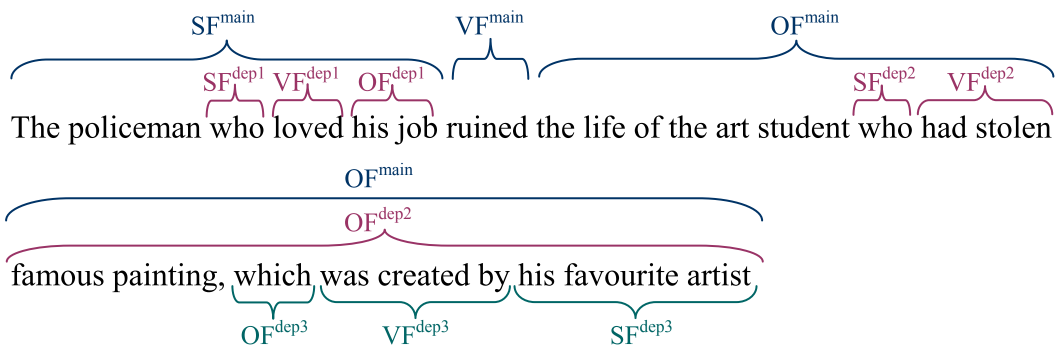

processing:

1. Emotion vectors of relative clauses, that do not have dependent clauses, are estimated and added to the corresponding Subject or Object Formations:

   1) e^{dep1} ('*who loved his job*') = coeff (tense:'past'; FPP:'no') * e^{dep1} (SF^{0dep1} & VF^{+dep1} & OF^{0dep1}) = 0.4 * [0,0,0,0,0.8,0.9,0,0,0] = [0,0,0,0,0.32,0.36,0,0,0] = e^{+dep1};

   SF^{main} = '*the policeman*' & e^{+dep1} = [0,0,0,0,0.32,0.36,0,0,0] = SF^{+main};

   2) e^{dep3} ('*which was created by his favourite artist*') = coeff (tense:'past'; FPP:'no') * e^{dep3} (SF^{+dep3} & VF^{0dep3} & OF^{0dep3}) = 0.4 * [0,0,0,0,0,0.6,0,0,0] = [0,0,0,0,0,0.24,0,0,0] = e^{+dep3};

   OF^{dep2} = '*famous painting*' & e^{+dep3} = [0,0,0,0,0,0.3,0,0,0.2] & [0,0,0,0,0,0.24,0,0,0] yield [0,0,0,0,0,0.3,0,0,0.2] = OF^{+dep2};

2. Then, described analysis procedure continues recursively till resulting emotion vector estimation:

   3) e^{dep2} ('*who had stolen famous painting, which was created by his favorite artist*') = coeff (tense:'past'; FPP:'no') * e^{dep2} (SF^{0dep2} & VF^{-dep2} & OF^{+dep2}) = 0.4 * ([0,0.2,0,0,0,0,0.5,0,0] & [0,0,0,0,0,0.3,0,0,0.2] yield [0,0.2,0,0,0,0,0.5,0,0]) = [0,0.08,0,0,0,0,0.2,0,0] = e^{-dep2};

   OF^{main} = '*the life of the art student*' & e^{-dep2} = [0,0.08,0,0,0,0,0.2,0,0] = OF^{-main};

   4) e^{main} ('*the policeman who loved his job ruined the life of the art student who had stolen famous painting, which was created by his favourite artist*') = coeff (tense:'past'; FPP:'no') * e^{main} (SF^{+main} & VF^{-main} & OF^{-main}) = 0.4 * ([0,0,0,0,0.32,0.36,0,0,0] & [0,0,0.7,0,0,0,0.9,0,0] & [0,0.08,0,0,0,0,0.2,0,0] yield [0,0.08,0.7,0,0,0,0.9,0,0])= [0,0.03,0.28,0,0,0,0.36,0,0] = e^{-main};

   5) result (sentence): Sadness:0.36.

Figure 3.3 Example of affect sensing in a complex sentence with multiple embedding of relative clauses

In Figure 3.3, the emotion vector is denoted by e = [Anger, Disgust, Fear, Guilt, Interest, Joy, Sadness, Shame, Surprise]; SF, VF, and OF represent Subject, Verb, and Object formations, respectively; the superscripts ^{0}, $^{-}$, and $^{+}$ indicate 'neutral', 'negative', and 'positive' polarities, respectively; ^{main} and ^{dep} mean belonging to 'main' and 'dependent' clauses, respectively.

3.5.4 Emotion Vector of a Complex-Compound Sentence

Sentences with at least two independent clauses and one or more dependent clauses are referred to as complex-compound sentences (for example, '*Max broke the china cup, with which Mary was awarded for the best song, so he regretted profoundly*'). While processing such type of sentences, first we generate emotion vectors of dependent clauses, then of complex sentences, and finally, we analyse the compound sentence formed by the independent clauses.

Chapter 4

Evaluation of the Affect Analysis Model Algorithm

In order to evaluate the performance of the Affect Analysis Model and to compare out method with related work, we conducted a set of experiments. In this Chapter we report and discuss the results of evaluation of the Affect Analysis Model algorithm.

4.1 Experiment with Our Collection of Sentences Extracted from Diary-Like Blogs

4.1.1 Data Set Description

As it is difficult to access logs of real IM sessions (due to privacy concerns), we initially investigated a collection of diary-like blog posts provided by BuzzMetrics, Inc. (*http://www.nielsenbuzzmetrics.com*). Here, we focused on online diary or personal blog entries, which are typically written in a free informal style and are rich in emotional colourations (Neviarouskaya, Prendinger and Ishizuka 2007b). Our observations suggest that every author practises a different style of writing. The most noticeable aspects of diary-like text are privacy, naturalism, and honesty in the expression of the author's thoughts and feelings. We concluded that the nature of such blog entries is reasonably close to online IM conversations (with the evident difference in the size of messages, however), and extracted 700 sentences from this Weblog Data

Collection in order to evaluate the emotion recognition algorithm (this annotated data set is freely available upon request).

Three independent annotators labelled the sentences with one of nine emotion categories (‘Anger’, ‘Disgust’, ‘Fear’, ‘Guilt’, ‘Interest’, ‘Joy’, ‘Sadness’, ‘Shame’, and ‘Surprise’), or neutral, and a corresponding intensity value. Additionally, we interpreted these fine-grained annotations using three polarity-based categories (positive emotion, negative emotion, and neutral) by merging ‘Interest’, ‘Joy’, and ‘Surprise’ in positive emotion category, and ‘Anger’, ‘Disgust’, ‘Fear’, ‘Guilt’, ‘Sadness’, and ‘Shame’ in negative emotion category. The reliability of the human raters’ annotations was measured using the Fleiss’ Kappa coefficient. The level of agreement on 700 sentences was moderate (0.47 in case of original annotations, and 0.59 in case of polarity-based annotations), and suggests that persons’ comprehension, interpretation, and evaluation of emotions are individualistic and might depend on personality type and emotional experience. As Davitz (1969: 85) stated, ‘In some respects, the experiences of each subject are undoubtedly unique. In fact, for any one person, even though experiences at different times are labelled by the same [emotional] term, these experiences are likely to differ somewhat from one another.’

To the best of our knowledge, there is no data set of sentences annotated with such an extensive set of labels (nine emotions and neutral); so there is no possibility to compare the agreements. It is obvious that the level of agreement on coarse-grained annotations (e.g., factual/subjective; positive/negative/neutral) is higher than on fine-grained annotations. As manual fine-grained annotations were interpreted using polarity-based categories by merging the emotions, this procedure could influence the agreement on polarity-based annotations. For example, annotators could assign ‘neutral’ label to non-emotional but having strong polarity sentences. If annotators were asked to provide polarity-based annotations for the sentence ‘*That place is one of the best places in Rochester for Mexican food, no lie*’, they would completely agree on ‘positive’ label, whereas fine-grained manual annotations were ‘neutral’/‘neutral’/‘Interest’.

For the evaluation of algorithm performance, we created the following gold standards:

(1) Sentences, on which at least two out of three human raters completely agreed:

- 656 sentences with fine-grained annotations (Fleiss’ Kappa coefficient is 0.51);
- 692 sentences with polarity-based annotations (Fleiss’ Kappa coefficient is 0.6).

(2) Sentences, on which all three human raters completely agreed:

- 249 sentences with fine-grained annotations (Fleiss' Kappa coefficient is 1.0);

- 447 sentences with polarity-based annotations (Fleiss' Kappa coefficient is 1.0).

The distributions of labels across gold standard sentences are summarized in Table 4.1.

Table 4.1 The distributions of labels across gold standard sentences

Annotations	Labels	At least two annotators agreed (number of sentences)	All three annotators agreed (number of sentences)
Fine-grained	Neutral	75	8
	Anger	59	17
	Disgust	30	9
	Fear	49	24
	Guilt	22	12
	Interest	43	8
	Joy	181	88
	Sadness	145	58
	Shame	9	3
	Surprise	43	22
	total	**656**	**249**
Polarity-based	Neutral	75	8
	Positive	270	177
	Negative	347	262
	total	**692**	**447**

4.1.2 Results based on gold standard sentences, on which at least two annotators agreed

To analyse the importance of words of different parts of speech in affect recognition, first we evaluated the performance of the Affect Analysis Model (AAM) with adjectives only, then we cumulatively added adverbs, verbs, and nouns to the algorithm. Averaged accuracy, precision, recall, and F-score at each step of this experiment are shown in Table 4.2 for each category. As was expected, the obtained results indicate that consideration of all content parts of speech plays a crucial role in emotion recognition from text. Two-tailed t-tests with significance level of 0.05 showed that the differences in accuracy between the preceding and the following algorithms are statistically significant ($p < 0.001$) in fine-grained as well as coarse-grained classifications, with the exceptional case of insignificant difference after adding the adverbs to the algorithm relying purely on adjectives.

The baseline for comparison (last raw in Table 4.2) is represented by the results of a simple method that selects the emotion with maximum intensity from the annotations of sentence tokens found in Affect database. Our AAM outperformed the baseline method (the difference is statistically significant in fine-grained and coarse-grained classifications: $p < 0.001$), thus demonstrating the contribution of the sentence parsing and our hand-crafted rules to the reliable recognition of emotions from text.

Table 4.2 Accuracy across sentences in the experiment with words of different parts of speech

Algorithm*	**Measure**	**Fine-grained categories**										**Merged labels**		
		Neut	**Ang**	**Disg**	**Fear**	**Guilt**	**Inter**	**Joy**	**Sad**	**Sh**	**Sur**	**Pos**	**Neg**	**Neut**
AAM +ADJ	Averaged accuracy	0.389										0.439		
	Precision	0.15	0.77	0.69	0.61	0.73	0.62	0.79	0.78	0.50	0.82	0.85	0.92	0.14
	Recall	0.79	0.17	0.30	0.29	0.36	0.37	0.47	0.25	0.44	0.33	0.49	0.33	0.79
	F-score	0.25	0.28	0.42	0.39	0.48	0.46	0.59	0.38	0.47	0.47	0.62	0.48	0.24
AAM +ADJ +ADV	Averaged accuracy	0.416										0.470		
	Precision	0.16	0.67	0.69	0.60	0.75	0.59	0.79	0.80	0.50	0.87	0.84	0.93	0.15
	Recall	0.77	0.17	0.30	0.31	0.41	0.37	0.50	0.28	0.44	0.47	0.53	0.36	0.77
	F-score	0.26	0.27	0.42	0.41	0.53	0.46	0.61	0.42	0.47	0.61	0.65	0.52	0.25
AAM +ADJ +ADV +VERB	Averaged accuracy	0.640										0.720		
	Precision	0.28	0.91	0.59	0.72	0.73	0.63	0.86	0.78	0.55	0.86	0.86	0.94	0.25
	Recall	0.65	0.36	0.63	0.69	0.50	0.81	0.73	0.57	0.67	0.72	0.80	0.67	0.65
	F-score	0.39	0.51	0.61	0.71	0.59	0.71	0.79	0.66	0.60	0.78	0.83	0.79	0.37
AAM +ADJ +ADV +VERB +NOUN	Averaged accuracy	**0.726**										**0.816**		
	Precision	0.46	0.83	0.63	0.76	0.75	0.56	0.87	0.78	0.57	0.85	0.85	0.92	0.41
	Recall	0.55	0.41	0.73	0.84	0.68	0.88	0.83	0.72	0.89	0.77	0.90	0.81	0.55
	F-score	0.50	0.55	0.68	0.80	0.71	0.68	0.85	0.75	0.70	0.80	0.87	0.86	0.47
Baseline +ADJ +ADV +VERB +NOUN	Averaged accuracy	0.546										0.692		
	Precision	0.09	0.62	0.55	0.59	0.57	0.36	0.63	0.64	0.50	0.67	0.69	0.81	0.08
	Recall	0.07	0.27	0.77	0.82	0.59	0.56	0.67	0.54	0.67	0.72	0.80	0.74	0.07
	F-score	0.08	0.38	0.64	0.68	0.58	0.44	0.65	0.59	0.57	0.70	0.74	0.77	0.07

* AAM stands for Affect Analysis Model; ADJ and ADV refer to adjectives and adverbs, respectively.

Next, we conducted a *functional ablation* experiment that aimed at evaluating our AAM with selectively removed functionality components: negation, neutralization due to modality, neutralization due to conditionality, modification by adverb-intensifiers, and intensity correction. We compared AAM with all functionalities, AAM without additional functionalities, and five approaches in which one specific functionality component was ablated from AAM. We believe that the remaining five AAM configurations would show the enhancement that each functionality adds to complete AAM rather than what is missing when each is removed. Table 4.3 includes the results of this experiment, showing that AAM mostly benefits from rules on negation and conditionality. Although no statistically significant differences in accuracy were found between the AAM with all functionalities and AAM algorithms with single additional functionality component removed, the statistical testing with significance level of 0.05 showed that the accuracy of AAM is significantly higher ($p < 0.01$ in fine-grained classification and $p < 0.05$ in coarse-grained classification) than the accuracy of AAM without all additional functionalities.

Table 4.3 Averaged accuracy across sentences from blogs in *functional ablation* experiment

Algorithm	Fine-grained categories	Merged labels
AAM with all functionalities	**0.726**	**0.816**
AAM w/o all additional functionalities	0.659	0.772
AAM w/o negation	0.688	0.790
AAM w/o modality	0.720	0.814
AAM w/o conditionality	0.707	0.808
AAM w/o modification by adverb-intensifiers	0.723	0.814
AAM w/o intensity correction	0.723	0.816

The analysis of errors in the assignment of polarity-based categories for the AAM (see Table 4.4) revealed that system requires common sense or additional context to deal with 28.5 percent of all errors. As human annotators labelled sentences only using fine-grained emotion categories and could assign 'neutral' to non-emotional but having strong polarity cases, we can consider the next type of error in the table (21.0 percent) as a non-strict one in the experiment with merged labels, where gold standard was based on fine-grained emotion annotations. In 9 percent of cases, where the system result did not agree with the gold standard due to the rule of neutralization of negated phrases, the solution would be to reverse the polarity of a statement; however, finding the pairs of opposite emotions might be problematic. The errors resulting from neutralization due to 'cognition-related' words comprise about 7 percent of errors. The failures also include some exceptional cases with connector '*but*', errors caused by the lack of relevant terms in Affect database, and incorrect results from the syntactic parser.

Table 4.4 Distribution of errors of AAM in experiment on sentences from blogs with merged labels

Error type	Error #	Error %	Sample sentence (gold standard; AAM result)*
Common sense or additional context	38	28.5	*It's true, my other friends' scanners work better*. (Sadness-NEG; Joy-POS) *What I hope is that he can understand how much I treasure this friendship*. (Sadness-NEG; Joy-POS)
Non-emotion (neutral) category, but with polarity	28	21.0	*Being rude is always out of style*. (neutral; Disgust-NEG) *That place is one of the best places in Rochester for Mexican food, no lie*. (neutral; Joy-POS)
Negation neutralization instead of negation reversal	12	9.0	*I don't care whether they like me at the cocktail parties, or not*. (Anger-NEG; neutral) *My job hunt isn't going so well, mainly because I don't have a job yet*. (Sadness-NEG; neutral)
Neutralization due to '*assume*', '*know*', '*think*'	9	6.8	*I always thought she liked my beard best*. (Joy-POS; neutral) *I tried explaining to him my outlooks on life last night, and I think that I upset him*. (Sadness-NEG; neutral)
Connector '*but*'	8	6.0	*It's still ugly, but at least it's moderately clean*. (Disgust-NEG; neutral)
Lexicon	8	6.0	*He's just lying*. (Anger-NEG; neutral)
Parser	6	4.5	*My son's team got 27 out of 30 questions right!* (Joy-POS; neutral)
Conflict (correct emotion is in the final vector of AAM, but is not dominant)	5	3.8	*I am always amazed, and angered, when I see people putting their infants in the front seat of their cars*. (Anger-NEG; Surprise-POS)
Neutralization due to '*can*', '*could*', '*may*', '*would*'	5	3.8	*A few weeks ago, I decided that I would pursue adopting a child through the foster care system*. (Joy-POS; neutral)
Neutralization due to negation	5	3.8	*I can't imagine how awful it will be to exist in this world two years from now*. (Fear-NEG; neutral)
Sense ambiguity	4	3.0	*The scene where the boys turned into donkeys was freaky*. (Surprise-POS; Anger-NEG)
Neutralization due to condition	4	3.0	*If I hated them they wouldn't be my friends would they?* (Anger-NEG; neutral)
Other	1	0.8	
Total, including double errors	133	100	

* Gold standard annotations and AAM results are given in the form 'Emotion - Polarity-based label for merged categories' in a last column.

We also evaluated the system performance with regard to intensity estimation. The percentage of emotional sentences (not considering neutral ones), on which the result of our system conformed to the fine-grained gold standard, according to the measured distance between intensities given by human raters (averaged values) and those obtained by the AAM is shown in Table 4.5. As seen in the table, our system achieved satisfactory results for emotion intensity estimation.

Table 4.5 Percentage of high agreement sentences according to the range of intensity difference between human annotations and output of algorithm

Range of intensity difference	[0.0 – 0.2]	(0.2 – 0.4]	(0.4 – 0.6]	(0.6 – 0.8]	(0.8 – 1.0]
Percentage of sentences, percent	48.5	32.2	15.9	3.4	0.0

4.1.3 Results based on gold standard sentences, on which all three annotators agreed

More accurate fine-grained and polarity-based annotations were obtained on the sentences, on which all three human annotators completely agreed. As seen in Table 4.6, the averaged accuracy of AAM on these sentences is higher (on 9 percent in case of fine-grained annotations and on 7.4 percent in case of polarity-based annotations) than on gold standard sentences, on which at least two annotators agreed (Table 4.2).

Table 4.6 Accuracy of AAM across gold standard sentences, on which all three annotators agreed

Algorithm	**Measure**	**Fine-grained categories**										**Merged labels**		
		Neut	**Ang**	**Disg**	**Fear**	**Guilt**	**Inter**	**Joy**	**Sad**	**Sh**	**Sur**	**Pos**	**Neg**	**Neut**
AAM	Averaged accuracy	**0.815**										**0.890**		
	Precision	0.26	0.92	0.83	0.91	0.83	0.44	0.95	0.88	0.67	0.86	0.93	0.99	0.15
	Recall	0.75	0.65	0.56	0.88	0.83	0.88	0.88	0.79	0.67	0.82	0.94	0.86	0.75
	F-score	0.39	0.76	0.67	0.89	0.83	0.58	0.91	0.84	0.67	0.84	0.94	0.92	0.25

4.2 Experiment with the Emotion Blog Data Set

4.2.1 Data Set Description

This data set was developed and kindly provided by Aman and Szpakowicz (2007). It includes sentences collected from blogs, which are characterized by rich emotional content and good examples of real-world instances of emotions conveyed through text. To directly compare the Affect Analysis Model with the machine learning methods proposed by Aman and Szpakowicz (2008), as the gold standard we considered their benchmark, which includes sentences annotated by one of six emotions ('Happiness' — in the description of this experiment we further use label 'Joy' instead, 'Sadness', 'Anger', 'Disgust', 'Surprise', and 'Fear') or neutral, on which two annotators completely agreed. The distribution of labels across sentences from the benchmark used in the experiment is shown in Table 4.7.

Table 4.7 Distribution of labels across sentences from benchmark used in the experiment

Labels	**Number of sentences**
Joy	536
Sadness	173
Anger	179
Disgust	172
Surprise	115
Fear	115
Neutral	600
total	**1890**

4.2.2 Results

As AAM is capable of recognition of nine emotions, and methods described in (Aman and Szpakowicz 2008) classify text to six emotions, in order to compare the results of our approaches we decided to reduce the number of our labels by mapping 'Interest' to 'Joy', and 'Guilt' and 'Shame' to 'Sadness'. The results of experiments are shown in Table 4.8, where AAM is compared to two classifiers trained using Support Vector Machines (the results of these classifiers are taken from Aman and Szpakowicz (2008)):

(1) 'ML with unigrams', which employs corpus-based features, namely, all unigrams that occur more than three times in the corpus, excluding stopwords.

(2) 'ML with unigrams, RT features, and WNA features', which combines corpus-based features with features based on the following emotion lexicons: Roget's Thesaurus (Jarmasz and Szpakowicz 2001) and WordNet-Affect (Strapparava and Valitutti 2004).

Table 4.8 Results of AAM compared to machine learning algorithms proposed by Aman and Szpakowicz (2008)

Algorithm	Measure	Label						
		Joy	Sadness	Anger	Disgust	Surprise	Fear	Neutral
AAM	Averaged accuracy				0.770			
	Precision	**0.846**	**0.673**	**0.910**	**0.946**	0.758	0.785	**0.698**
	Recall	**0.858**	**0.763**	**0.564**	**0.506**	**0.652**	**0.730**	**0.862**
	F-score	**0.852**	**0.715**	**0.697**	**0.659**	**0.701**	**0.757**	**0.771**
ML with unigrams	Precision	0.840	0.619	0.634	0.772	**0.813**	**0.889**	0.581
	Recall	0.675	0.301	0.358	0.453	0.339	0.487	0.342
	F-score	0.740	0.405	0.457	0.571	0.479	0.629	0.431
ML with unigrams, RT features, and WNA features	Precision	0.813	0.605	0.650	0.672	0.723	0.868	0.587
	Recall	0.698	0.416	0.436	0.488	0.409	0.513	0.625
	F-score	0.751	0.493	0.522	0.566	0.522	0.645	0.605

The obtained results (precision, recall, and F-score) revealed that our rule-based system outperformed both machine learning methods in automatic recognition of 'Joy', 'Sadness', 'Anger', 'Disgust', and 'neutral'. In case of 'Surprise' and 'Fear' emotions, 'ML with unigrams' resulted in higher precision, but lower recall and F-score than our AAM.

4.3 Summary

The salient features of the Affect Analysis Model are the following:

(1) Analysis of nine emotions on the level of individual sentences: this is an extensive set of labels if compared to six emotions mainly used in related work.

(2) The ability to handle the evolving language of online communications: to the best of our

knowledge, our approach is the first attempt to deal with informal and abbreviated style of writing, often accompanied by the use of emoticons.

(3) Foundation in database of affective words (each term in our Affect database was assigned at least one emotion label along with emotion intensity, in contrast to annotations of one emotion label or polarity orientation in other approaches), interjections, emoticons, abbreviations and acronyms, modifiers (which influence the degrees of emotion states).

(4) Vector representation of affective features of words, phrases, clauses, and sentences.

(5) Consideration of syntactic relations and semantic dependencies between words in a sentence: our rule-based method accurately classifies context-dependent affect expressed in sentences containing emotion-conveying terms, which may play different syntactic and semantic roles.

(6) Analysis of negation, modality, and conditionality: most researchers ignore modal expressions and condition prepositions, therefore, their systems show poor performance in classifying neutral sentences, which is, indeed, not an easy task.

(7) Consideration of relations between clauses in compound, complex, or complex-compound sentences: to our knowledge, AAM is the first system comprehensively processing affect reflected in sentences of different complexity.

(8) Emotion intensity estimation: in our work, the strength of emotion is encoded through numerical value in the interval [0.0, 1.0], in contrast to low/middle/high levels detected by some of other methods.

Our system showed promising results in affect recognition on real examples of diary-like blog posts: (1) on data set created by us, where at least two annotators agreed, averaged accuracy was 72.6 percent for fine-grained (nine categories, and neutral) emotion classification and 81.6 percent for polarity-based merged categories (positive, negative, and neutral); (2) on data set created by us, where all three annotators agreed, averaged accuracy was 81.5 percent for fine-grained emotion classification, and 89.0 percent for polarity-based merged categories; (3) on data set provided by Aman and Szpakowicz (2008), averaged accuracy was 77.0 percent for fine-grained (six categories, and neutral) emotion classification, and our system outperformed the method reported in related work in terms of precision, recall, and F-scores.

Part II

Recognition of Affect, Judgment, and Appreciation in Text

Chapter 5

Lexical Resources

In the first part of the Chapter we describe methods to automatically generate and score a new sentiment lexicon, called SentiFul, and expand it through direct synonymy and antonymy relations, hyponymy relations, morphologic modifications and compounding with known lexical units. We propose to distinguish four types of affixes (used to derive new words) depending on the role they play with regard to sentiment features: propagating, reversing, intensifying, and weakening. Besides derivation, we considered important process of finding new words such as compounding, which is a highly productive process, especially in the case of nouns and adjectives. We elaborated the algorithm for automatic extraction of new sentiment-related compounds from WordNet (Miller 1990) using words from SentiFul as seeds for sentiment-carrying base components and applying the patterns of compound formations.

The second part of the Chapter is devoted to an AttitudeFul database containing lexicon necessary for fine-grained attitude analysis. The importance of considering modifiers, contextual valence shifters, and modal operators, which are integral parts of the lexicon for robust attitude analysis, is also discussed.

5.1 SentiFul: Generating a Reliable Lexicon for Sentiment Analysis

5.1.1 Generating the Core of Sentiment Lexicon

The first step in building the lexicon of sentiment-conveying terms involves the collection of relevant content words (adjectives, adverbs, nouns, and verbs), and the assignment of prior polarity scores (positivity score and negativity score) to each lexical unit. By ‘sentiment polarity score’ we mean the strength or degree of intensity of sentiment (for example, ‘*cheerful*’, ‘*happy*’, and ‘*elated*’ have different strengths of positivity). In our work, for both opposite valences, the bounds of the polarity score are 0.0 (indicating the absence of given orientation of sentiment) and 1.0 (the utmost value of intensity).

For the generation of the core of our sentiment lexicon, we employ the extended version of Affect database (see Chapter 2 for details), which contains in total 2438 direct and indirect emotion-related entries: 918 adjectives (e.g., ‘*euphoric*’, ‘*hostile*’), 243 adverbs (e.g., ‘*luckily*’, ‘*miserably*’), 900 nouns (e.g., ‘*fright*’, ‘*mercy*’), and 377 verbs (e.g., ‘*reward*’, ‘*blame*’). The affective features of each distinct word in this database are encoded using nine emotions (‘Anger’, ‘Disgust’, ‘Fear’, ‘Guilt’, ‘Interest’, ‘Joy’, ‘Sadness’, ‘Shame’, and ‘Surprise’), and are represented as a vector of emotional state intensities that range from 0.0 to 1.0. Using emotional vectors, we interpreted the sentiment of Affect database entries by means of polarity scores and polarity weights. Polarity weight means rate of the number of positive (negative) emotions with intensity greater than 0.0 to the total number of emotions with intensity greater than 0.0 in the emotional vector (positive and negative weights add up to 1.0). We considered three emotions (‘Interest’, ‘Joy’, and ‘Surprise’) as having mainly positive orientation, and six emotions (‘Anger’, ‘Disgust’, ‘Fear’, ‘Guilt’, ‘Sadness’, and ‘Shame’) as negatively-valenced.

Positivity and negativity scores were calculated using (3) and (4). Based on (5) and (6), we derived the polarity weights.

$$Pos\_score = \left[\frac{\sum_{i=1}^{pos} Intensity(i)}{pos} \right], \tag{3}$$

$$Neg\_score = \left[\frac{\sum_{i=1}^{neg} Intensity(i)}{neg} \right], \tag{4}$$

$$Pos\_weight = \left[\frac{pos}{pos + neg} \right], \tag{5}$$

$$Neg\_weight = \left[\frac{neg}{pos + neg} \right], \tag{6}$$

where *Intensity* is the intensity value of the corresponding emotion in the emotional vector; *pos* (*neg*) is the number of positive (negative) emotions having *Intensity* > 0.0 in the emotional vector, respectively.

We named our sentiment database as 'SentiFul'. Some examples of SentiFul entries are listed in Table 5.1.

Table 5.1 Examples of words with sentiment annotations from SentiFul

Affective word	Part of speech	Non-zero-intensity emotions from Affect database emotional vector	Polarity scores		Polarity weights	
			Pos_score	***Neg_score***	***Pos_weight***	***Neg_weight***
tremendous	Adjective	'Surprise:1.0', 'Joy:0.5', 'Fear:0.1'	0.75	0.1	0.67	0.33
pensively	Adverb	'Sadness:0.2', 'Interest:0.1'	0.1	0.2	0.5	0.5
success	Noun	'Joy:0.9', 'Interest:0.6', 'Surprise:0.5'	0.67	0.0	1.0	0.0
regret	Verb	'Guilt:0.2', 'Sadness:0.1'	0.0	0.15	0.0	1.0

The main drawback of a sentiment analysis approach, which is purely relying on a lexicon of sentiment-conveying terms, is the lack of scalability, since the recall of the lexical method depends on the coverage of the database used. Thus, to expand SentiFul, we first investigated the possibility to take advantage of sense-level scores from SentiWordNet (version 1.0) (Esuli and Sebastiani 2006).

5.1.2 Examining the SentiWordNet

SentiWordNet was developed based on WordNet (Miller 1990) synsets comprised from synonymous terms. Motivated by the assumption that 'different senses of the same term may have different opinion-related properties', Esuli and Sebastiani (2006) developed a method employing eight ternary classifiers and quantitatively analyzing the glosses associated with synsets. Three numerical scores (*Obj*(*s*), *Pos*(*s*), and *Neg*(*s*)), which characterize to what degree the terms included in a synset are objective, positive, and negative, were automatically determined based on the proportion of classifiers assigning the corresponding label to the synset. The scores range from 0.0 to 1.0 and sum up to 1.0.

The question '*How reliable is SentiWordNet?*' arouse at the very beginning of its exploration, just after analyzing the scores of synsets that include the adjective '*happy*' (Table 5.2). Three out of six synsets are characterized by negativity predominance (*Neg*(*s*) is greater than both *Pos*(*s*) and *Obj*(*s*)); in two synsets the scores of positivity prevail (*Pos*(*s*) is greater than both *Neg*(*s*) and *Obj*(*s*)); and one synset is completely objective (*Obj*(*s*) = 1.0) in SentiWordNet. A sentiment analysis system employing a sense disambiguation algorithm might yield counter-intuitive results on the sentence '*Those were happiest days, I never felt such elation!*', if scores for the {happy(5), euphoric(1)} synset would be considered.

Table 5.2 SentiWordNet scores for synsets containing adjective '*happy*'

Synset with corresponding sense	*Pos*(*s*)	*Neg*(*s*)	*Obj*(*s*)
{happy(1)}: enjoying or showing or marked by joy or pleasure or good fortune; '*a happy smile*'; '*spent many happy days on the beach*'; '*a happy marriage*'	**0.625**	0.25	0.125
{happy(2), pleased(3)}: experiencing pleasure or joy; '*happy you are here*'; '*pleased with the good news*'	0.0	**0.75**	0.25
{happy(3), felicitous(2)}: marked by good fortune; '*a felicitous life*'; '*a happy outcome*'	**0.875**	0.0	0.125
{happy(4)}: satisfied; enjoying well-being and contentment; '*felt content with her lot*'; '*quite happy to let things go on as they are*'	0.0	**0.75**	0.25
{happy(5), euphoric(1)}: exaggerated feeling of well-being or elation	0.125	**0.5**	0.375
{happy(6), well-chosen(1)}: well expressed and to the point; '*a happy turn of phrase*'; '*a few well-chosen words*'; '*a felicitous comment*'	0.0	0.0	**1.0**

Let us now turn to the analysis of possibilities to extend the SentiFul lexicon using SentiWordNet. As we restricted polarity scores and polarity weights in SentiFul to distinct lexemes (sentiment features of different senses of a term are unified), we considered two approaches to derive scores for each lexeme from SentiWordNet (other approaches are described in (Alm 2008) and (Fahrni and Klenner 2008)):

(1) Method '*FS*': take *Pos*(*s*), *Neg*(*s*), and *Obj*(*s*) scores of first synset for each lemma in SentiWordNet.

(2) Method '*UNI*': calculate unified positivity and negativity scores for each lemma in SentiWordNet using (7) and (8); and derive weights of positivity, negativity, and objectivity based on (9), (10), and (11). As there are synsets where *Pos*(*s*) = *Neg*(*s*) > 0.0, all weights need to be normalized.

$$Uni\_Pos\_score = \left[\frac{\sum_{i=1}^{pos} Pos(s)(i)}{pos} \right], \tag{7}$$

$$Uni\_Neg\_score = \left[\frac{\sum_{i=1}^{neg} Neg(s)(i)}{neg} \right], \tag{8}$$

$$Pos\_weight = \left[\frac{pos}{senses} \right], \tag{9}$$

$$Neg\_weight = \left[\frac{neg}{senses} \right], \tag{10}$$

$$Obj\_weight = \left[\frac{obj}{senses} \right], \tag{11}$$

where *pos* is the number of lemma senses having *Pos*(*s*)(*i*) >= *Neg*(*s*)(*i*) and *Pos*(*s*)(*i*) > 0.0; *neg* is number of lemma senses having *Neg*(*s*)(*i*) >= *Pos*(*s*)(*i*) and *Neg*(*s*)(*i*) > 0.0; *obj* is the number of lemma senses having *Obj*(*s*)(*i*) = 1.0; *senses* is a total number of lemma synsets.

Using '*FS*' and '*UNI*' methods, we obtained scores for all 152050 distinct lemmas in SentiWordNet. In particular, total numbers of distinct lemmas having either *Obj*(*s*) <= 0.5 (from

'*FS*') or *Obj_weight* <= 0.5 (from '*UNI*') are 14918 and 37414, respectively. In order to evaluate the appropriateness of scores derived from SentiWordNet, we created a gold standard based on SentiFul entries (originating from the manually annotated Affect database) and their scores. For the gold standard we considered only those SentiFul entries that also occur in SentiWordNet: 750 adjectives, 237 adverbs, 894 nouns, 372 verbs. The evaluation was based on the comparison of the valence of the dominant score derived from SentiWordNet with the valence of the dominant score from the SentiFul gold standard.

The rule for the determination of valence of the dominant score for a lemma in the gold standard is as follows:

IF (*Pos_score* >= *Neg_score* AND *Pos_weight* > *Neg_weight*)
THEN valence = positive
ELSE IF (*Pos_score* > *Neg_score* AND *Pos_weight* = *Neg_weight*)
  THEN valence = positive
  ELSE IF (*Neg_score* >= *Pos_score* AND *Neg_weight* > *Pos_weight*)
    THEN valence = negative
    ELSE IF (*Neg_score* > *Pos_score* AND *Neg_weight* = *Pos_weight*)
      THEN valence = negative
      ELSE IF (*Pos_score* = *Neg_score* AND *Pos_weight* = *Neg_weight*)
        THEN valence = random
        ELSE IF (*Pos_score* > *Neg_score*)
          THEN valence = positive
          ELSE valence = negative.

To obtain the valence of the dominant score within scores derived from SentiWordNet using '*FS*' and '*UNI*' methods, we propose four ways:

(1) '*FS_strength*' (disregarding *Obj*(*s*)):

IF (*Pos*(*s*) > *Neg*(*s*)) THEN valence = positive
ELSE IF (*Neg*(*s*) > *Pos*(*s*)) THEN valence = negative
  ELSE IF (*Pos*(*s*) = *Neg*(*s*) = 0.0) THEN valence = neutral
    ELSE valence = random.

(2) '*FS_obj*':

IF (*Obj*(*s*) > 0.5) THEN valence = neutral

ELSE IF (*Pos*(*s*) > *Neg*(*s*)) THEN valence = positive

ELSE IF (*Neg*(*s*) > *Pos*(*s*)) THEN valence = negative

ELSE valence = random.

(3) '*UNI_strength*' (disregarding *Obj_weight*):

IF (*Uni_Pos_score* > *Uni_Neg_score*) THEN valence = positive

ELSE IF (*Uni_Neg_score* > *Uni_Pos_score*) THEN valence = negative

ELSE IF (*Uni_Pos_score* = *Uni_Neg_score* = 0.0) THEN valence = neutral

ELSE valence = random.

(4) '*UNI_weight*':

IF (*Obj_weight* > 0.5) THEN valence = neutral

ELSE IF (*Uni_Pos_score* >= *Uni_Neg_score* AND *Pos_weight* > *Neg_weight*)

THEN valence = positive

ELSE IF (*Uni_Pos_score* > *Uni_Neg_score* AND *Pos_weight* = *Neg_weight*)

THEN valence = positive

ELSE IF (*Uni_Neg_score* >= *Uni_Pos_score* AND *Neg_weight* > *Pos_weight*)

THEN valence = negative

ELSE IF (*Uni_Neg_score* > *Uni_Pos_score* AND *Neg_weight* = *Pos_weight*)

THEN valence = negative

ELSE IF (*Uni_Pos_score* = *Uni_Neg_score* AND *Pos_weight* = *Neg_weight*)

THEN valence = random

ELSE IF (*Uni_Pos_score* > *Uni_Neg_score*)

THEN valence = positive

ELSE valence = negative.

Table 5.3 includes some examples of the obtained results.

Table 5.3 Examples of the comparison of results from different methods with gold standard

Lemma (POS)	**Method**	***Pos_score***	***Neg_score***	***Pos_weight***	***Neg_weight***	**Dominant**	**Result**
weakness (Noun)	SentiWordNet sense #1	0.0	0.125				
	SentiWordNet sense #2	0.125	0.625				
	SentiWordNet sense #3	0.0	0.375				
	SentiWordNet sense #4	0.5	0.125				
	SentiWordNet sense #5	0.0	0.875				
	SentiFul gold standard	**0.0**	**0.2**	**0.0**	**1.0**	**negative**	
	'*FS_strength*'	0.0	0.125	-	-	negative	**hit**
	'*FS_obj*'	0.0	0.125	-	-	neutral	**neutral no hit**
	'*UNI_strength*'	0.5	0.5	0.2	0.8	random	**random hit**
	'*UNI_weight*'	0.5	0.5	0.2	0.8	negative	**hit**
congratulate (Verb)	SentiWordNet sense #1	0.25	0.125				
	SentiWordNet sense #2	0.0	0.125				
	SentiWordNet sense #3	0.0	0.375				
	SentiWordNet sense #4	0.0	0.5				
	SentiFul gold standard	**0.4**	**0.0**	**1.0**	**0.0**	**positive**	
	'*FS_strength*'	0.25	0.125	-	-	positive	**hit**
	'*FS_obj*'	0.25	0.125	-	-	neutral	**neutral no hit**
	'*UNI_strength*'	0.25	0.333	0.25	0.75	negative	**no hit**
	'*UNI_weight*'	0.25	0.333	0.25	0.75	negative	**no hit**

The results of the evaluation of different methods for obtaining scores for adjectives, adverbs, nouns, and verbs based on SentiWordNet are displayed in Figure 5.1. As seen from the diagrams, more accurate scores were obtained for adjectives in comparison with other parts of speech, and the worst results were obtained for scoring the verbs. The '*UNI*' method performed better than the method based on the consideration of scores of the first synset in SentiWordNet ('*FS*' method). The results we obtained when examining SentiWordNet were not satisfying, and hence we decided to explore other ways to extend the SentiFul lexicon.

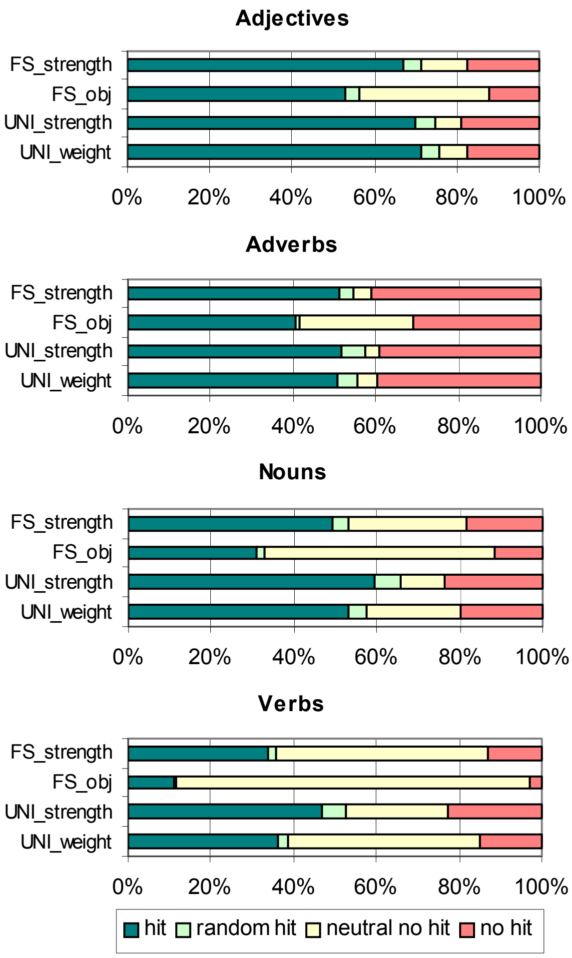

Figure 5.1 Accuracy of different methods for obtaining scores based on SentiWordNet

5.1.3 Methods for Expanding the SentiFul

5.1.3.1 Finding New Lexical Units through Direct Synonymy Relation

To find new sentiment-related words, the most direct way is to derive them through the synonymy relation with known lexemes. Undoubtedly, the deep meaning of any lexical unit is unique. However, we can take advantage of considering pairs of words that have similar senses and assign sentiment scores to them. The process of finding and scoring new words through a synonymy relation consists of three main steps, which are applied to adjectives, adverbs, nouns, and verbs independently. For the exploration of WordNet (Miller 1990) relations, we employed Java API for WordNet Searching (JAWS) publicly available at *http://lyle.smu.edu/~tspell/jaws*.

Step 1. Given a word from SentiFul, we derive all related synsets found in WordNet. For example, four synsets were found for verb ‘*congratulate*’: {‘*compliment*’, ‘*congratulate*’},

{‘*congratulate*’, ‘*felicitate*’}, {‘*pride*’, ‘*plume*’, ‘*congratulate*’}, and {‘*preen*’, ‘*congratulate*’}.

Step 2. In each multiple-word synset from the previous step, we retrieve words that are already included in SentiFul, then calculate averages of scores and weights within synsets that have new terms, and finally assign these values to remaining words within corresponding synset. For the above example, all synonyms of the verb ‘*congratulate*’, except ‘*compliment*’ in first synset and ‘*felicitate*’ in second synset, are already in SentiFul. Therefore, scores of ‘*congratulate*’ (*Pos_score* = 0.4, *Neg_score* = 0.0, *Pos_weight* = 1.0, and *Neg_weight* = 0.0) are propagated to ‘*complement*’ and ‘*felicitate*’. In case the verb ‘*pride*’ from third synset was new for SentiFul, we would take the averages of polarity scores and the averages of weights of both ‘*plume*’ and ‘*congratulate*’.

Step 3. After *Step 1* and *Step 2* are completed for all original SentiFul entries (we consider only their direct synonyms), we eliminate duplicates of new words, as they can obtain assignments from different synsets derived using different words from SentiFul, and calculate their new scores as averages of assignments of those redundantly produced words.

Relying on direct synonymy relations, we automatically extracted 4190 new words from WordNet (see examples in Table 5.4): 1122 adjectives, 107 adverbs, 1731 nouns, and 1230 verbs. We decided not to iterate the above procedure on these new words, because non-direct synonyms are not necessarily carrying similar sentiment features as original concepts (e.g., ‘*health*y’ – ‘*intelligent*’ – ‘*thinking*’).

Table 5.4 Examples of newly derived words based on direct synonymy relations

POS	**Lemma**	***Pos_score* / *Neg_score***	***Pos_weight* / *Neg_weight***
Adjective	*appealing*	0.333 / 0.033	0.833 / 0.167
	barbarous	0.0 / 0.625	0.0 / 1.0
	confounded	0.1 / 0.2	0.167 / 0.833
Adverb	*advantageously*	0.3 / 0.0	1.0 / 0.0
	frightfully	0.0 / 0.95	0.0 / 1.0
	poorly	0.0 / 0.334	0.0 / 1.0
Noun	*authority*	0.383 / 0.05	0.875 / 0.125
	defect	0.0 / 0.6	0.0 / 1.0
	impetuosity	0.65 / 0.65	0.5 / 0.5
Verb	*exhaust*	0.2 / 0.375	0.167 / 0.834
	glorify	0.3 / 0.0	1.0 / 0.0
	privilege	0.2 / 0.0	1.0 / 0.0

5.1.3.2 Examining Direct Antonymy Relation

Next step in enriching the SentiFul database is to analyze antonymy relations. We concentrated on the extraction of direct antonyms of words available in SentiFul from WordNet. Direct antonymous words are conceptual opposites that represent lexical pairs (indirect antonyms are not lexically paired).

Given a word from SentiFul and its class (adjective, adverb, noun, verb), we retrieve its direct antonyms from WordNet. Then, if these newly retrieved words are not available in SentiFul, we assign sentiment-related scores and weights to them based on the assumption that direct antonyms possess sentiment features that are opposite to those of the original word from SentiFul. Hence, the original *Pos_score* and *Neg_score* trade their places (same procedure for weights) in case of direct antonyms. If the same antonyms are retrieved using different original words from SentiFul, we calculate averages of scores and weights. For example, using the nouns '*falsehood*' and '*falsity*' from SentiFul, we retrieved duplicate entries of their direct antonym '*truth*' from WordNet, and calculated averages of the reversed scores and weights of the original words. The examination of direct antonymy relations of SentiFul entries allowed us to automatically extract 288 new words from WordNet (some examples are listed in Table 5.5): 123 adjectives, 13 adverbs, 73 nouns, and 79 verbs.

Table 5.5 Examples of newly derived words based on direct antonymy relations

POS	Word (direct antonym of words from SentiFul)	*Pos_score* / *Neg_score*	*Pos_weight* / *Neg_weight*
Adjective	*attractive* (*repulsive*, *unattractive*)	0.8 / 0.0	1.0 / 0.0
	maleficent (*beneficent*)	0.0 / 0.3	0.0 / 1.0
	wise (*foolish*)	0.9 / 0.0	1.0 / 0.0
Adverb	*carelessly* (*carefully*)	0.0 / 0.1	0.0 / 1.0
	honorably (*dishonorably*)	1.0 / 0.0	1.0 / 0.0
	painlessly (*painfully*)	0.4 / 0.0	1.0 / 0.0
Noun	*penalty* (*reward*)	0.0 / 0.2	0.0 / 1.0
	safety (*danger*)	0.7 / 0.0	1.0 / 0.0
	truth (*falsehood*, *falsity*)	0.9 / 0.0	1.0 / 0.0
Verb	*bless* (*curse*)	0.25 / 0.0	1.0 / 0.0
	defend (*attack*)	0.7 / 0.0	1.0 / 0.0
	deteriorate (*recuperate*)	0.0 / 0.2	0.0 / 1.0

5.1.3.3 Examining Hyponymy Relations

The most important semantic relation in organizing nouns in WordNet is a relation between lexicalized concepts, or a relation of subordination (Miller 1999). In WordNet the lexical hierarchy of nouns is represented using hypernym-hyponym relations between the appropriate synsets. At the top of the hierarchy, there are few generic terms, which can characterize many specific terms at the lower levels. The semantic relation known as hyponymy goes from the generic term to a more specific one, thus representing specialization (e.g., '*attainment*' => '*success*' => '*winning*'), whereas the hypernymy relation points in the opposite direction, i.e. from a specific term to a more generic one (e.g., '*winning*' => '*success*' => '*attainment*').

Miller (1999: 31) defines hyponymy as follows:

> '*When the features characterizing synset {A} are all included among the features characterizing synset {B}, but not vice versa, then {B} is a hyponym of {A}*.'

Hyponymy relation between nouns is of our particular interest, as we assume that sentiment features of a sentiment-conveying term (e.g., '*success*'), along with other features, are to some extent inherited by its hyponym (e.g., '*winning*'). On the other hand, hypernymy relation represents generalization, and it is not necessarily true that sentiment features of a sentiment-conveying term (e.g., '*success*') will characterize its hypernym (e.g., '*attainment*'), which is located at higher level of the lexical hierarchy.

Our algorithm for hyponymy retrieval from a lexical inheritance system of WordNet takes into account only one level of specialization. Given a noun from SentiFul, we automatically retrieve a list of corresponding hyponyms from WordNet, and propagate sentiment features (scores and weights) of the original term to its hyponyms. If the hyponymy relation of different nouns from SentiFul results in the same term, we eliminate duplicates and consider averages of their scores and weights as the resulting assignment. The examples of nouns retrieved from WordNet through examination of hyponymy relations are shown in Table 5.6. In total, 1085 new nouns were added to the SentiFul lexicon.

Table 5.6 Examples of nouns retrieved based on hyponymy relations

Retrieved word	*Pos_score* / *Neg_score*	*Pos_weight* / *Neg_weight*	Is a hyponym of [word from SentiFul]	*Pos_score* / *Neg_score*	*Pos_weight* / *Neg_weight*
amity	0.25 / 0.0	1.0 / 0.0	*friendliness*	0.3 / 0.0	1.0 / 0.0
			peace	0.2 / 0.0	1.0 / 0.0
aspersion	0.0 / 0.55	0.0 / 1.0	*attack*	0.0 / 0.7	0.0 / 1.0
			depreciation	0.0 / 0.4	0.0 / 1.0
betise	0.0 / 0.35	0.0 / 1.0	*error*	0.0 / 0.3	0.0 /1.0
			fault	0.0 / 0.6	0.0 /1.0
			mistake	0.0 / 0.15	0.0 /1.0
consonance	0.4 / 0.0	1.0 / 0.0	*harmony*	0.4 / 0.0	1.0 / 0.0
fiasco	0.0 / 0.5	0.0 / 1.0	*collapse*	0.0 / 0.5	0.0 / 1.0
reprehensibility	0.0 / 0.9	0.0 / 1.0	*evil*	0.0 / 0.9	0.0 / 1.0

5.1.3.4 Method Based on Morphological Modifications

We are proposing to expand our SentiFul lexicon through manipulations with morphological structure of known lemmas that result in the formation of new lexical units (Plag 2003). Adjectives, adverbs, nouns, and verbs form open classes, whereby membership is indefinite and unlimited (Biber et al. 1999). We can easily form new words playing with bases and affixes. Derivation is a process responsible for building new lexemes, by either adding derivational prefixes (attachments to the front of the base) or suffixes (attachments to the end of the base). Suffixes typically have less specific meanings than prefixes. The main contribution to meaning of many suffixes is that which follows from a change of the grammatical class.

We distinguish four types of affixes depending on the role they play with regard to sentiment features:

(1) *Propagating* affixes preserve sentiment features of the original lexeme and propagate them to newly derived lexical unit. For example:

‘*en-*’ + ‘*rich*’ => ‘*enrich*’;

‘*harmony*’ + ‘*-ous*’ => ‘*harmonious*’;

‘*scary*’ + ‘*-fy*’ => ‘*scarify*’.

(2) *Reversing* affixes change the orientation of sentiment features of the original lexeme. For example:

‘*dis-*’ + ‘*honest*’ => ‘*dishonest*’;

'*harm*' + '-*less*' => '*harmless*'.

(3) *Intensifying* affixes increase the strength of sentiment features of the original lexeme. For example:

'*super-*' + '*hero*' => '*superhero*';

'*over-*' + '*awe*' => '*overawe*'.

(4) *Weakening* affixes decrease the strength of sentiment features of the original lexeme. For example:

'*semi-*' + '*sweet*' => '*semisweet*'.

Table 5.7 summarizes our classification with respect to the type of an affix, class of a base lexeme (*a* stands for adjective, *adv* for adverb, *n* for noun, and *v* for verb), and class of a newly formed word.

Table 5.7 Our classification of affixes attached to a base lexeme to form new word

Type of affix	Prefix (+class of base lexeme); (class of base lexeme+) suffix	Examples
	Adjective formation	
Propagating	pro- (+*a*); (*a*+) -ish; (*v*+) {-able, -ant, -ent, -ible, -ing}; (*n*+) {-al, -en, -ful, -ic, -like, -type, -y}; (*v*/*n*+) {-ate, -ed, -ive, -ous}	*attacking*, *advanced*, *harmonious*, *careful*, *lovable*, *messy*
Reversing	{a-, ab-, an-, anti-, contra-, counter-, de-, dis-, dys-, il-, im-, in-, ir-, mal-, mis-, non-, pseudo-, un-, under-} (+*a*); (*n*+) -less	*intolerant*, *dishonest*, *misleading*, *guiltless*, *harmless*
Intensifying	{extra-, hyper-, mega-, super-, ultra-} (+*a*)	*superfine*
Weakening	semi- (+*a*)	*semisoft*
	Adverb formation	
Propagating	pro- (+*adv*); (*a*+) -ly; (*n*+) {-wise, -wards}	*charmingly*, *defectively*
Reversing	{a-, ab-, an-, anti-, contra-, counter-, de-, dis-, dys-, il-, im-, in-, ir-, mal-, mis-, non-, pseudo-, un-, under-} (+*adv*);	*imperfectly*, *ungratefully*
Intensifying	{extra-, hyper-, mega-, super-, ultra-} (+*adv*)	
Weakening	semi- (+*adv*)	
	Noun formation	
Propagating	{neo-, re-} (+*n*); (*v*+) {-age, -al, -ant, -ation, -ent, -ication, -ification, -ion, -ment, -sion, -tion, -ure}; (*a*+) {-ity, -ness}; (*n*+) {-ful, ist, -ship}; (*v*/*a*+) {-ance, -ence, -ee}; (*v*/*n*+) {-er, -ing, -or}; (*a*/*n*+) {-cy, -dom, -hood}; (*v*/*n*/*a*+) {-ery, -ry}	*awfulness*, *deceiver*, *offender*, *savagery*
Reversing	{anti-, counter-, dis-, dys-, in-, mal-, mis-, non-, pseudo-, under-} (+*n*)	*nonviolence*, *underachiever*
Intensifying	{arch-, hyper-, mega-, super-, ultra-} (+*n*)	*superego*
Weakening	{mini-, semi-} (+*n*); (n+) {-ette, -let}	*mini-recession*

Table 5.7 (Continued)

Type of affix	Prefix (+class of base lexeme); (class of base lexeme+) suffix	Examples
	Verb formation	
Propagating	{be-, co-, fore-, inter-, pre-, pro-, re-, trans-} (+*v*); {em-, en-} (+*n*/*a*); (*n*/*a*+) {-ate, -en, -fy, -ify, -ise, -ize}	*enrich*, *scarify*, *agonize*
Reversing	{de-, dis-, dys-, mis-, un-, under-} (+*v*)	*devalue*, *disagree*, *mistrust*
Intensifying	{out-, over-} (+*v*)	*outfight*, *overawe*

Our algorithm for building new words receives the following parameters: class of the base word, class (prefix or suffix) and type of the affix, affix, and the class of derived word. The processing is as follows (please see schematic illustration in Figure 5.2):

(1) given the class of the base word, the system successively extracts each corresponding lemma from SentiFul and its sentiment-related scores,

(2) depending on the affix class, affix is attached either to the front or to the end of the lemma to form new word;

(3) given the class of a derived word and the newly formed word itself, SentiFul is scanned on the presence of this lemma, and if the result is positive, this lemma is not considered for inclusion; else, WordNet is examined on the availability of this lemma, and if this word exists, it is considered for future inclusion to SentiFul along with sentiment-related scores.

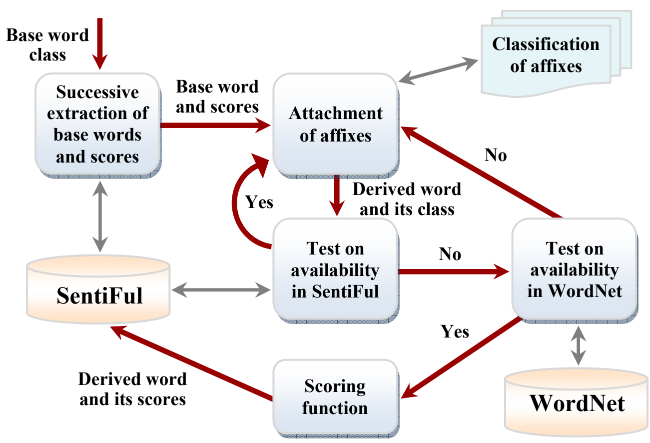

Figure 5.2 The algorithm of derivation and scoring of the new words

Based on the type of the affix and sentiment-related scores of the original word, the scoring function assigns polarity scores and weights to the derived word. In the case of *Propagating* affix, original scores and weights are transferred to the new word without variation. The original *Pos_score* and *Neg_score* trade their places (same procedure for weights) in case of a *Reversing* affix. If the affix belongs to *Intensifying* or *Weakening* type, the original scores are multiplied by 2.0 or 0.5, respectively.

In order to properly treat attachment of suffixes to base lexemes, we apply the following rules:

(1) Replace lexeme ending '*f*' (except the case of '*ff*') by '*v*' if suffix starts with '*a*/*e*/*i*/*o*/*u*/*y*'.

(2) Replace lexeme ending '*fe*' (except the case of '*ffe*') by '*v*' if suffix starts with '*a*/*e*/*i*/*o*/*u*/*y*'.

(3) Remove lexeme ending '*y*' if suffix starts with '*i*'.

(4) Replace lexeme ending '*y*', which follows the consonant, by '*i*'.

(5) Remove (noun or adjective) lexeme ending '*t*' or '*te*' before suffix '*cy*'.

(6) Remove lexeme ending '*e*' if suffix starts with '*a*/*e*/*i*/*o*/*u*/*y*'.

(7) Double lexeme ending '*b*/*d*/*f*/*g*/*l*/*m*/*n*/*p*/*r*/*s*/*t*/ *v*/*z*', which follows the vowel preceded by consonant, if suffix starts with '*a*/*e*/*i*/*o*/*u*/*y*'.

For example, while attaching the suffix '*-fy*' to the base lexeme '*beauty*' (noun), we replace the lexeme ending '*y*' by '*i*' to correctly derive verb '*beautify*' (rule 4); or in the case of base lexeme '*love*' (verb) and suffix '*-able*', we remove lexeme ending '*e*' to derive adjective '*lovable*' (rule 6).

Using this morphologically inspired method, we automatically derived and scored 4029 new words (see examples in Table 5.8): 1405 adjectives, 484 adverbs, 1800 nouns, and 340 verbs.

Table 5.8 Examples of morphologically modified words

POS	Lemma	*Pos_score* / *Neg_score*	*Pos_weight* / *Neg_weight*
Adjective	*lovable*	0.85 / 0.0	1.0 / 0.0
	reproachful	0.0 / 0.625	0.0 / 1.0
Adverb	*proficiently*	0.3 / 0.0	1.0 /0.0
Noun	*spoilage*	0.133 / 0.3	0.167 / 0.833
Verb	*beautify*	0.45 / 0.0	1.0 / 0.0

The *Propagating* type of affixes proved to be the most frequent and efficient in building words of all content parts of speech (Figure 5.3). The *Reversing* type of affixes played also significant role in the derivation process for adjectives and verbs, while *Intensifying* affixes brought noticeable effect only in building new verbs.

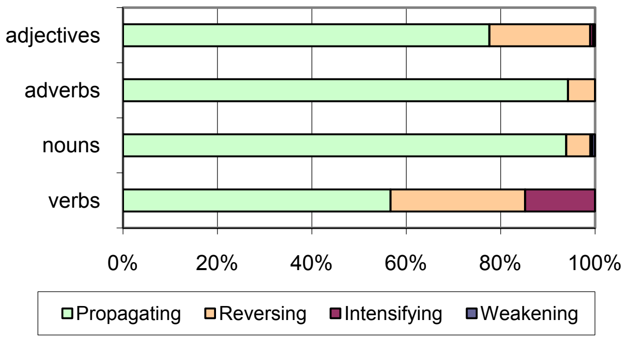

Figure 5.3 Percentage distribution of words derived by means of different affix types

The block diagram shown in Figure 5.4 indicates that adjectives, adverbs, and nouns were mainly derived by means of suffixes, whereas prefixes dominated in the case of verbs.

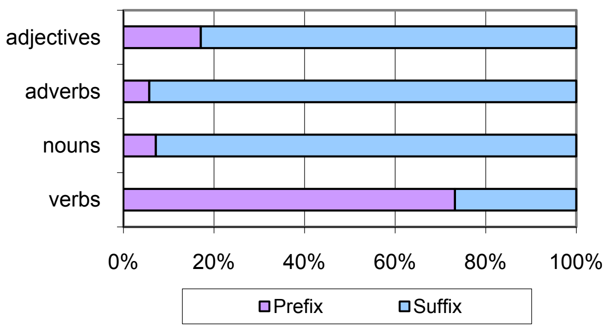

Figure 5.4 Percentage distribution of words derived by means of prefixes and suffixes

The most productive affixes to form new words are listed in Table 5.9.

Table 5.9 Top 10 most productive affixes to form adjectives, adverbs, nouns, and verbs

POS	Affixes and counts									
Adjective	-ed	-ing	un-	-able	-less	-ive	-y	-ful	-al	in-
	492	226	148	97	80	64	64	50	31	29
Adverb	-ly	un-	a-	in-	im-	dis-	-wise	-wards	-	-
	458	14	7	3	2	2	2	1	-	-
Noun	-er	-ing	-ness	-or	-ion	-ation	-ment	-ist	-ery	-ity
	607	367	340	79	75	53	45	37	34	32
Verb	re-	over-	-en	dis-	un-	de-	out-	mis-	-ize	-ise
	56	34	30	26	22	21	18	18	16	16

5.1.3.5 Compounding Using Known Sentiment-Carrying Base Components

Besides derivation, we considered important process of finding new words, such as compounding, which is a highly productive process, especially in the case of nouns and adjectives. Compounds are words that contain at least two roots. In other words, independently existing bases are combined to form new lexemes. Compounding functions as a linguistic economy-mechanism that allows expressing in a concise way something which would otherwise have to be rendered by means of a phrase (Meys 1975).

A number of different compounding patterns are attested in English. We analysed major patters of formation of noun compounds and adjectival compounds described in (Meys 1975; Biber et al. 1999; Plag 2003). The patterns, which are of our main interest, are summarized in Table 5.10 along with the examples illustrating sentiment-conveying compounds.

Although some compounds have idiosyncratic meanings which are different from the sum of the meanings of their parts, the meanings of many compounds may be systematically related to the meanings of their components via a number of different rules (Kaplan 1989). We assume that if a compound contains at least one base component that conveys sentiment features, we can predict the valence of this compound.

Table 5.10 Patterns of formation of noun compounds and adjectival compounds

Patterns	Structure in terms of paraphrasing	Examples of compound words	Valence-based interpretation	Rule
		Formation of noun compounds		
noun + noun	'modifier-head'	*love-affair* *death-feud*	pos-neutral => pos neg-neg => neg	Rule 1 Rule 4a
noun + noun/verb-er	'verb-object'	*life-saver* *peace-lover* *pain-killer*	neutral-pos => pos pos-pos => pos neg-neg => pos	Rule 1 Rule 5a Rule 5b
noun + verb-ing	'verb-object'	*law-breaking* *peace-keeping*	neutral-neg => neg pos-neutral => pos	Rule 1 Rule 1
adjective + noun	'modifier-head'	*poor-quality* *good-neighborliness* *no-nonsense*	neg-neutral =>neg pos-neutral => pos 'negation'-neg=> pos	Rule 1 Rule 1 Rule 2
verb + noun	'modifier-head'	*cry-baby*	neg-neutral => neg	Rule 1
verb-ing + noun	'modifier-head'	*loving-kindness*	pos-pos => pos	Rule 4a
pronoun + noun	'modifier-head'	*self-pity*	neutral-neg => neg	Rule 1
noun + preposition + noun	'modifier-head'	*wall-of-death*	neutral-neg => neg	Rule 1

Table 5.10 (Continued)

Patterns	Structure in terms of paraphrasing	Examples of compound words	Valence-based interpretation	Rule
		Formation of adjectival compounds		
noun + verb-ing	'verb-object'	*eye-gladdening* *fight-eliciting* *award-winning* *health-destroying* *quarrel-loving*	neutral-pos => pos neg-neutral => neg pos-pos => pos pos-neg => neg neg-pos => neg	Rule 1 Rule 1 Rule 5a Rule 5c Rule 5d
pronoun +verb-ing	'verb-object'	*self-destructing*	neutral-neg => neg	Rule 1
adjective + verb-ing	'modifier-head'	*pleasant-testing* *evil-smelling*	pos-neutral => pos neg-neutral => neg	Rule 1 Rule 1
adverb + verb-ing	'modifier-head'	*equally-damaging* *ever-loving* *badly-fitting*	neutral-neg => neg neutral-pos => pos neg-neutral => neg	Rule 1 Rule 1 Rule 1
noun + verb-en	'verb-PP'	*poverty-stricken* *fortune-favored* *snob-despised* *war-torn* *love-agonized*	neg-neutral => neg pos-pos => pos neg-neg => neg neg-neg => neg pos-neg => neg	Rule 1 Rule 6a Rule 6a Rule 6a Rule 6b
pronoun + verb-en	'verb-PP'	*self-convicted*	neutral-neg => neg	Rule 1
adjective + verb-en	'modifier-head'	*kind-hearted*	pos-neutral => pos	Rule 1
adverb + verb-en	'modifier-head'	*poorly-adapted* *well-merited* *ill-famed*	neg-neutral => neg pos-pos => pos neg-pos =>neg	Rule 1 Rule 4a Rule 4b
verb-en + preposition	'verb-preposion'	*broken-down*	neg-neutral => neg	Rule 1
adjective + verb	'modifier-head'	*easy-follow* *difficult-to-master*	pos-neutral => pos neg-pos => neg	Rule 1 Rule 4b
noun + adjective	'modifier-head'	*user-friendly* *money-mad* *crash-proof* *error-free*	neutral-pos => pos neutral-neg => neg neg-'valence shifter' => pos neg-'valence shifter' => pos	Rule 1 Rule 1 Rule 3 Rule 3
pronoun + adjective	'modifier-head'	*self-conscious*	neutral-pos => pos	Rule 1
adjective + preposition + pronoun	'adjective-PP'	*spurious-to-me* *good-for-nothing*	neg-neutral => neg pos-'negation' => neg	Rule 1 Rule 2
adjective + noun	'modifier-head'	*no-win*	'negation'-pos=> neg	Rule 2
adjective + adjective	'modifier-head'	*manic-depressive*	neg-neg => neg	Rule 4a
adverb + adjective	'modifier-head'	*highly-respectable* *critically-ill* *not-too-pleasant*	neutral-pos => pos neg-neg => neg 'negation'-pos => neg	Rule 1 Rule 4a Rule 2
verb + noun	'verb-object'	*cut-throat* *ban-the-bomb*	(indirect)neg-neutral => neg neg-neg => pos	Rule 1 Rule 5b
verb + adjective	'verb-adjective'	*get-rich-quick*	neutral-pos => pos	Rule 1
verb + adverb	'modifier-head'	*die-hard*	neg-(indirect)pos => pos	Rule 4b

Based on this assumption, we elaborated the algorithm for automatic extraction of new sentiment-related terms (particularly, compounds) from WordNet using words from SentiFul as seeds for sentiment-carrying base components, and patterns for formation of compounds (Table 5.10). It is important to note here that we restricted the algorithm to form compounds written with a hyphen. The rules for estimation of sentiment features (*Pos_score*, *Neg_score*, *Pos_weight*, and *Neg_weight*) of newly retrieved words are described below. The examples of compound words, the valence-based interpretations of their constituent parts, and the corresponding rules are also given in Table 5.10.

Rule 1. If one of the constituent elements of a compound conveys sentiment features, and another element, which is not 'negation' or 'valence shifter' word, is neutral, then sentiment-features are propagated to the whole compound. For example:

'*good*' & '*neighborliness*'

[*Pos_score* = 0.3, *Neg_score* = 0.0] & [neutral]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0]

=> '*good-neighborliness*'

[*Pos_score* = 0.3, *Neg_score* = 0.0]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0];

'*pound*' & '*foolish*'

[neutral] & [*Pos_score* = 0.0, *Neg_score* = 0.7]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0]

=> '*pound-foolish*'

[*Pos_score* = 0.0, *Neg_score* = 0.7]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0].

Rule 2. If one of the constituent elements of a compound conveys sentiment features, and another element is a 'negation' word, then sentiment features of the sentiment-conveying component are reversed and assigned to the whole compound. For example:

'*no*' & '*nonsense*'

[negation] & [*Pos_score* = 0.0, *Neg_score* = 0.5]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0]

=> ‘*no-nonsense*’

[*Pos_score* = 0.5, *Neg_score* = 0.0]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0];

‘*good*’ & ‘*nothing*’

[*Pos_score* = 0.3, *Neg_score* = 0.0] & [negation]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0]

=> ‘*good-for-nothing*’

[*Pos_score* = 0.0, *Neg_score* = 0.3]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0].

Rule 3. If the left-hand member of a compound conveys sentiment features, and the right-hand member is a ‘valence shifter’ (e.g., ‘*safe*’, ‘*free*’, ‘*proof*’’, etc.) or its derivative, then sentiment features of the sentiment-conveying component are reversed and assigned to the whole compound. For example:

‘*fail*’ & ‘*safe*’

[*Pos_score* = 0.0, *Neg_score* = 0.9] & [valence shifter]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0]

=> ‘*fail-safe*’

[*Pos_score* = 0.9, *Neg_score* = 0.0]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0];

‘*risk*’ & ‘*free*’

[*Pos_score* = 0.0, *Neg_score* = 0.567] & [valence shifter]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0]

=> ‘*risk-free*’

[*Pos_score* = 0.567, *Neg_score* = 0.0]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0].

Rule 4. If a compound is interpreted in such a way that one member modifies another member (so called ‘modifier-head’ structure), and both the ‘modifier’ and the ‘head’ are sentiment-conveying terms, then:

Rule 4a. If both components are predominantly positive (or negative), then their sentiment features (scores and weights) are averaged, and the result is assigned to the whole word. For example:

‘*loving*’ & ‘*kindness*’

[*Pos_score* = 0.9, *Neg_score* = 0.0] & [*Pos_score* = 0.6, *Neg_score* = 0.0]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0] [*Pos_weight* = 1.0, *Neg_weight* = 0.0]

=> ‘*loving-kindness*’

[*Pos_score* = 0.75, *Neg_score* = 0.0]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0];

‘*death*’ & ‘*feud*’

[*Pos_score* = 0.0, *Neg_score* = 0.65] & [*Pos_score* = 0.0, *Neg_score* = 0.4]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0] [*Pos_weight* = 0.0, *Neg_weight* = 1.0]

=> ‘*death-feud*’

[*Pos_score* = 0.0, *Neg_score* = 0.525]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0].

Rule 4b. If both components have contrasting sentiment features, then sentiment features of the ‘modifying’ member are considered as dominant and are propagated to the whole word. For example:

‘*ill*’ & ‘*famed*’

[*Pos_score* = 0.0, *Neg_score* = 0.467] & [*Pos_score* = 0.475, *Neg_score* = 0.0]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0] [*Pos_weight* = 1.0, *Neg_weight* = 0.0]

=> ‘*ill-famed*’

[*Pos_score* = 0.0, *Neg_score* = 0.467]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0].

Rule 5. If a compound corresponds to one of the patterns, which can be paraphrased as ‘verb + direct object’ (so called ‘verb-object’ structure), and both components are sentiment-conveying terms, then:

Rule 5a. If both 'noun' and 'verb/verbal' members are predominantly positive, then their sentiment features (scores and weights) are averaged and the result is assigned to the whole word. For example:

'*award*' & '*winning*'

[*Pos_score* = 0.55, *Neg_score* = 0.0] & [*Pos_score* = 0.8, *Neg_score* = 0.0]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0] [*Pos_weight* = 1.0, *Neg_weight* = 0.0]

=> '*award-winning*'

[*Pos_score* = 0.675, *Neg_score* = 0.0]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0].

Rule 5b. If both 'noun' and 'verb/verbal' members are predominantly negative, then their sentiment features (scores and weights) are averaged, and the inverted result is assigned to the whole word. For example:

'*pain*' & '*killer*'

[*Pos_score* = 0.0, *Neg_score* = 0.8] & [*Pos_score* = 0.0, *Neg_score* = 0.35]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0] [*Pos_weight* = 0.0, *Neg_weight* = 1.0]

=> '*pain-killer*'

[*Pos_score* = 0.575, *Neg_score* = 0.0]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0].

Rule 5c. If the 'noun' member is predominantly positive and the 'verb/verbal' member is predominantly negative, then sentiment features of the 'verb/verbal' member are considered as dominant and are propagated to the whole word. For example:

'*health*' & '*destroying*'

[*Pos_score* = 0.25, *Neg_score* = 0.0] & [*Pos_score* = 0.0, *Neg_score* = 0.65]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0] [*Pos_weight* = 0.0, *Neg_weight* = 1.0]

=> '*health-destroying*'

[*Pos_score* = 0.0, *Neg_score* = 0.65]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0].

Rule 5d. If the 'noun' member is predominantly negative and the 'verb/verbal' member is predominantly positive, then sentiment features of the 'noun' member are considered as dominant and are propagated to the whole word. For example:

'*quarrel*' & '*loving*'

[*Pos_score* = 0.0, *Neg_score* = 0.35] & [*Pos_score* = 0.9, *Neg_score* = 0.0]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0] [*Pos_weight* = 1.0, *Neg_weight* = 0.0]

=> '*quarrel-loving*'

[*Pos_score* = 0.0, *Neg_score* = 0.35]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0].

Rule 6. If a compound corresponds to the pattern, which can be paraphrased as 'verb-en by/with/in/from noun' (so called 'verb-PP' structure), were 'noun' member represents an agent, instrument, location etc., and both components are sentiment-conveying terms, then:

Rule 6a. If both components are predominantly positive (or negative), then their sentiment features (scores and weights) are averaged, and the result is assigned to the whole word. For example:

'*fortune*' & '*favored*'

[*Pos_score* = 0.7, *Neg_score* = 0.0] & [*Pos_score* = 0.6, *Neg_score* = 0.0]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0] [*Pos_weight* = 1.0, *Neg_weight* = 0.0]

=> '*fortune-favored*',

which is paraphrased as '*favored by fortune*',

[*Pos_score* = 0.65, *Neg_score* = 0.0]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0];

'*war*' & '*torn*'

[*Pos_score* = 0.0, *Neg_score* = 0.85] & [*Pos_score* = 0.0, *Neg_score* = 0.1]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0] [*Pos_weight* = 0.0, *Neg_weight* = 1.0]

=> '*war-torn*'

[*Pos_score* = 0.0, *Neg_score* = 0.475]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0].

Rule 6b. If both components have contrasting sentiment features, then sentiment features of the 'verbal' member (verb-en) are considered as dominant and are propagated to the whole word. For example:

'*love*' & '*agonized*'

[*Pos_score* = 0.9, *Neg_score* = 0.0] & [*Pos_score* = 0.0, *Neg_score* = 0.85]

[*Pos_weight* = 1.0, *Neg_weight* = 0.0] [*Pos_weight* = 0.0, *Neg_weight* = 1.0]

=> '*love-agonized*',

which is paraphrased as '*agonized by love*',

[*Pos_score* = 0.0, *Neg_score* = 0.85]

[*Pos_weight* = 0.0, *Neg_weight* = 1.0].

Based on the sentiment-conveying words from SentiFul and the above rules, we could automatically extract and annotate not only new noun compounds and adjectival compounds from WordNet, but also some adverbs (e.g., '*light-heartedly*') and verbs (e.g., '*goof-proof*', '*atom-bomb*'). During evaluation of newly derived compounds, we found out that few words were assigned incorrect sentiment features (e.g., negative adjective '*half-truth*' was given dominant *Pos_score* due to positive noun '*truth*'; positive verb '*trouble-shoot*' was given dominant *Neg_score* due to negative noun '*trouble*', etc.).

We also decided to add some neoclassical compounds automatically retrieved from WordNet to our SentiFul database. Neoclassical compounds are defined as forms in which lexemes of Latin or Greek origin are combined to form new combinations that are not attested in the original languages (Plag 2003). Key ending elements that have strongly affective content, such as '*-cide*' (meaning: '*murder*'), '*-itis*' (meaning: '*disease*'), and '*-phobe*' (meaning: '*fear*'), were considered. Compounds having these endings were automatically retrieved from WordNet. Sentiment features (*Pos_score*, *Neg_score*, *Pos_weight*, and *Neg_weight*) of the meaning word representing particular key ending element were assigned to the compounds derived by means of this element. For example:

'*genocide*', '*suicide*', etc. were given sentiment features of word '*murder*' (*Pos_score*=0.0, *Neg_score*=0.8, *Pos_weight*=0.0, *Neg_weight*=1.0); exceptions are '*viricide*' and '*virucide*';

'*appendicitis*', '*radiculitis*', etc. are characterized by sentiment features of word '*disease*' (*Pos_score*=0.0, *Neg_score*=0.3, *Pos_weight*=0.0, *Neg_weight*=1.0);

'*claurtrophobe*', '*technophobe*' were assigned sentiment features of word '*fear*' (*Pos_score*=0.0, *Neg_score*=0.9, *Pos_weight*=0.0, *Neg_weight*=1.0).

Compounding using known sentiment-carrying key elements allowed us to expand our SentiFul lexicon by 853 new words: 377 adjectives, 15 adverbs, 445 nouns (including 184 common compounds and 261 neoclassical compounds), and 16 verbs.

5.1.4 Evaluation of the SentiFul

5.1.4.1 Evaluation Based on Human Annotations

In order to evaluate the accuracy of the methods described in the previous Section, we randomly extracted 1000 terms from SentiFul, particularly, 200 terms from each of the five lists created by different methods, including techniques based on direct synonymy, antonymy, and hyponymy relations, derivation process, and compounding. We asked two human annotators to assign the dominant polarity label (positive, negative, or neutral) and the polarity score to each of the randomly retrieved word. As the gold standard we considered only those words where both annotators agreed on the polarity labels, excluding words with neutral label, as our methods were not designed to distinguish between neutral and sentiment-conveying terms. The statistical data on the manual annotations of 1000 SentiFul terms are given in Table 5.11.

Table 5.11 Statistical data on the manual annotations of 1000 SentiFul terms

Method	Interannotator Cohen's Kappa on 200 words	Words with complete agreement	Percentage distribution of labels, %			Number of words in the gold standard
			pos	neg	neutral	
Synonymy	0.78	179	27.9	69.8	2.2	175
Antonymy	0.66	156	44.2	26.3	29.5	110
Hyponymy	0.87	187	31.6	67.4	1.1	185
Derivation	0.91	191	35.6	60.7	3.7	184
Compounding	0.93	193	45.6	53.9	0.5	192

For the comparison with the gold standard annotations, the dominant polarity of each word was extracted from the SentiFul. The results of the evaluation of different methods with regard to polarity assignments are shown in Table 5.12.

Table 5.12 Results of evaluation of polarity assignments

Method	Accuracy, %	Precision, %		Recall, %		F-score, %	
		pos	neg	pos	neg	pos	neg
Synonymy	95.4	86.2	100	100	93.6	92.6	96.7
Antonymy	94.5	97.0	90.7	94.2	95.1	95.6	92.9
Hyponymy	98.9	96.7	100	100	98.4	98.3	99.2
Derivation	97.8	95.7	99.1	98.5	97.4	97.1	98.3
Compounding	99.5	98.9	100	100	99.0	99.4	99.5

As seen from the data, the method relying on antonymy relations yielded noisy results (29.5 percent of words, on which both annotators agreed, are neutral). The method based on compounding performed with the highest accuracy (99.5 percent) in assigning dominant positive or negative labels, followed by the methods considering hyponymy relations (98.9 percent), derivation process (97.8 percent), synonymy relations (95.4 percent), and antonymy relations (94.5 percent). With regard to positive and negative labels, the F-score of assigning negative label is greater than the F-score of assigning positive label in the case of four out of five methods, except the method based on antonymy relations. A possible explanation might be in the proportions of positive and negative labels in the 'gold standard' (Table 5.11).

The accuracy of the methods concerning different content words (adjectives, adverbs, nouns, and verbs) is given in Table 5.13. As seen from the table, more accurate labels were obtained for nouns and verbs in comparison with other parts of speech, and the worst results were obtained from labelling adverbs using antonymy relations and adjectives using derivation process and antonymy relations.

Table 5.13 Accuracy with regard to different parts of speech

Method	Accuracy, %			
	adjectives	adverbs	nouns	verbs
Synonymy	95.7	90.5	97.8	97.6
Antonymy	91.7	75.0	100	96.2
Hyponymy	-	-	98.9	-
Derivation	93.8	97.9	100	100
Compounding	100	100	98.8	100

The evaluations of the polarity scores were based on the Pearson measure of correlation between the polarity scores automatically assigned by each of the methods and the gold standard scores manually assigned by human annotators. We considered only those words, on the dominant polarity label of which our methods agreed with both annotators. Table 5.14 contains the Pearson measures of correlation between scores provided by each method and scores given by each annotator individually, showing mainly strong positive relationships ($r > 0.5$).

Table 5.14 Results of the evaluation of the polarity scores

Method	**Pearson's correlation coefficient (r)**	
	Scores of Annotator 1	**Scores of Annotator 2**
Synonymy	0.576	0.626
Antonymy	0.112	0.599
Hyponymy	0.498	0.618
Derivation	0.520	0.603
Compounding	0.617	0.757

The obtained results indicate that the methods based on compounding and synonymy relations achieved high accuracy in assigning appropriate polarity scores to sentiment-conveying terms; and the method relying on antonymy relations was the least accurate.

We analysed the erroneous outcomes of the derivation process. We found that, for example, the derivation algorithm assigned positive scores to the verb '*reprise*' ('*re*-'+'*prise*') and the nouns '*lovage*' ('*love*'+'-*age*') and '*truster*' ('*trust*'+'-*er*'), which were labeled as neutral by both human raters. The examples of mislabeled words include the adjectives '*chanceful*', '*fanciful*', and '*oddish*' (positive in SentiFul, while negative in the gold standard), and the adverb '*modestly*' (negative in SentiFul, while positive in the gold standard).

5.1.4.2 Evaluation Based on General Inquirer

Next we evaluated our SentiFul entries based on the polarity lexicon from General Inquirer (GI) (*http://www.wjh.harvard.edu/~inquirer/*). Particularly, we collected 4002 GI terms (distinct adjectives, adverbs, nouns, and verbs) labeled as "Positiv" (1813) and "Negativ" (2189). The gold standard for evaluation is based on the intersection of polarity-based GI and SentiFul; it includes in total 2223 words (957 positive and 1266 negative). The statistical data on the polarity lexicon from General Inquirer and the GI gold standard are given in Table 5.15.

Table 5.15 Number of entries in the polarity lexicon from General Inquirer and the gold standard

Part of speech	General Inquirer (GI)			GI gold standard		
	pos	neg	total	pos	neg	total
adjectives	706	766	1472	287	365	652
adverbs	50	19	69	25	10	35
verbs	388	654	1042	163	315	478
nouns	669	750	1419	482	576	1058
total	**1813**	**2189**	**4002**	**957**	**1266**	**2223**

In order to evaluate the SentiFul lexicon, we calculated the agreement (Cohen's Kappa coefficient) between the SentiFul annotations and GI gold standard. The agreement was substantial (k = 0.72) for all content words. The highest agreement was obtained on adverbs (0.81), followed by adjectives (0.79), nouns (0.7), and verbs (0.67).

The SentiFul annotations were found accurate (86.3 percent accuracy). The measured values of precision, recall, and F-score with regard to polarity labels (positive and negative) are given in Table 5.16. The results of evaluation based on different parts of speech (adjectives, adverbs, nouns, and verbs) are shown in Table 5.17.

Table 5.16 Accuracy with regard to polarity assignments

Measure	Polarity	
	pos	neg
Accuracy, %	**86.3**	
Precision, %	81.8	90.1
Recall, %	87.6	85.3
F-score, %	84.6	87.6

Table 5.17 The results of evaluation based on different parts of speech

Measure	Part of speech			
	adjectives	adverbs	nouns	verbs
Accuracy, %	89.4	91.4	85.0	84.5

5.2 AttitudeFul: a Lexicon for Fine-Grained Attitude Analysis

We built a lexicon for fine-grained attitude analysis (AttitudeFul) that includes:

(1) Attitude-conveying terms.

(2) Modifiers.

(3) 'Functional' words.

(4) Modal operators.

5.2.1 The Core of an Attitude-Conveying Lexicon

As a core of lexicon for attitude analysis (namely, for the analysis of affect, judgment, and appreciation in text), we employ the Affect database (Section 2.2) and the SentiFul database (Section 5.1). The affective features of each emotion-related word are encoded using nine emotion labels and corresponding emotion intensities that range from 0.0 to 1.0. The original version of SentiFul database, which contains sentiment-conveying adjectives, adverbs, nouns, and verbs annotated by sentiment polarity, polarity scores and weights, was manually extended using fine-grained attitude labels:

(1) *Affect categories*:

'Anger', 'Disgust', 'Fear', 'Guilt', 'Interest', 'Joy', 'Sadness', 'Shame', and 'Surprise', which are interpreted using 'POS aff' (positive affect) and 'NEG aff' (negative affect) polarity labels.

(2) *Judgment polarity labels*:

'POS jud' (positive judgment) and 'NEG jud' (negative judgment).

(3) *Appreciation polarity labels*:

'POS app' (positive appreciation) and 'NEG app' (negative appreciation).

Some examples of annotated attitude-conveying words are listed in Table 5.18. It is important to note here that some words may express different attitude types (affect, judgment, appreciation) depending on context; such lexical entries were annotated by all possible categories (e.g., adjective '*unfriendly*' in Table 5.18).

Table 5.18 Examples of attitude-conveying words and their annotations

Part of speech	Word	Attitude category	Intensity
Adjective	*honorable*	POS jud	0.3
	unfriendly	NEG aff (Sadness)	0.5
		NEG jud	0.5
		NEG app	0.5
Adverb	*gleefully*	POS aff (Joy)	0.9
Noun	*abnormality*	NEG app	0.25
Verb	*frighten*	NEG aff (Fear)	0.8
	desire	POS aff (Interest)	1.0
		POS aff (Joy)	0.5

5.2.2 Modifiers and 'Functional' Words

A robust attitude analysis method should rely not only on attitude-conveying terms, but also on modifiers and contextual valence shifters (this term was introduced by Polanyi and Zaenen (2004)), which are integral parts of the AttitudeFul lexicon.

We collected modifiers that have an impact on contextual attitude features of neighbouring words, related phrases, or clauses. The modifiers include:

(1) Adverbs of degree (e.g., '*significantly*', '*slightly*' etc.) and adverbs of affirmation (e.g., '*absolutely*', '*seemingly*') that have an influence on the strength of the attitude of related words. Two annotators gave coefficients for intensity degree strengthening or weakening (from 0.0 to 2.0) to each adverb, and the result was averaged (e.g., coeff('*perfectly*') = 1.9, coeff('*slightly*') = 0.2). The Pearson's correlation coefficient calculated between human annotations (r = 0.98) showed very strong positive relationship.

(2) Negation words (e.g., '*never*', '*nothing*', '*no*' etc.) that reverse the polarity of related statement.

(3) Adverbs of doubt (e.g., '*scarcely*', '*hardly*' etc.) and adverbs of falseness (e.g., '*wrongly*' etc.) that reverse the polarity of related statement.

(4) Prepositions (e.g., '*without*', '*despite*' etc.) that neutralize the attitude of related words.

(5) Condition operators (e.g., '*as if*', '*if*', '*even though*' etc.) that neutralize the attitude of related words.

In total, AttitudeFul currently contains 138 modifiers.

We distinguish two types of ‘functional’ words that can influence the contextual attitude and its strength:

(1) *Intensifying* type:

adjectives (e.g., ‘*resurgent*’, ‘*rapidly-growing*’, ‘*rising*’ etc.),

nouns (e.g., ‘*increase*’, ‘*increment*’, ‘*up-tick*’ etc.), and

verbs (e.g., ‘*to grow*’, ‘*to rocket*’ etc.),

which increase the strength of attitude of the related words.

(2) *Reversing* type:

adjectives (e.g., ‘*reduced*’ etc.),

nouns (e.g., ‘*termination*’, ‘*reduction*’ etc.), and

verbs (e.g., ‘*to decrease*’, ‘*to limit*’, ‘*to diminish*’ etc.),

which reverse the prior polarity of the related words.

Using the list of seed functional words, we analysed their basic semantic relations (synonymy and antonymy relations) in WordNet, and thus collected 240 relevant terms.

5.2.3 Modal Operators

> ‘*Modality… is… concerned with subjective characteristics of an utterance, and it could even be further argued that subjectivity is an essential criterion for modality. Modality could, that is to say, be defined as the grammaticalization of speaker's (subjective) attitudes and opinions.*’ Palmer (1986: 16)

Modality is related to assertions of probability, possibility, permission, intention, obligation and the like (Hoye 1997). Consideration of the modal operators in the tasks of opinion mining, sentiment and attitude analysis is very important, as they indicate the degree of a person’s belief in the truth of the proposition, which is subjective in nature. Modal expressions point to likelihood and clearly involve the speaker’s judgment. Modals are distinguished by the *confidence level*. Hoye [1997: 80] argues that ‘inference and confidence go “hand-in-hand” and are directly tied to the status of the speaker’s “knowledge”; the stronger the evidence, the more forceful can be the expression of the speaker’s resolve’.

We collected modal operators of two categories:

(1) Modal verbs (in total, 13 verbs).

(2) Modal adverbs (in total, 61 adverbs).

Table 5.19 includes the classification of the collected modal operators and their examples. Since modals are considered as indicators of the *confidence level* of expressed attitude, we asked three human annotators to assign the *confidence level*, which ranges from 0.0 to 1.0, to each modal verb and adverb, based on the corresponding predefined range of *confidence level* displayed in the last column of Table 5.19. The Pearson's correlation coefficients calculated between human annotations pairwise (r_1 = 0.98, r_2 = 0.99, r_3 = 0.99) showed very strong positive correlations. Three ratings of each modal operator were averaged.

Table 5.19 Classification of modal operators

Type	**Example (*confidence level*)**	**Range of *conf. level***
	Modal verbs	
Central modal auxiliaries of possibility	*may* (0.27), *can* (0.5)	[0.2-0.5]
Central modal auxiliaries of probability	*should* (0.6), *would* (0.8), *will* (0.9)	(0.5-1.0)
Central modal auxiliaries of certainty	*must* (1.0)	1.0
Modal 'marginals'	*dare* (0.5), *ought* (0.7), *need* (1.0)	[0.2-1.0]
	Modal adverbs	
Adverbs of doubt	*doubtfully* (0.1), *fishily* (0.1), *vaguely* (0.17)	[0.0-0.3)
Adverbs of possibility	*conceivably* (0.37), *supposedly* (0.5)	[0.3-0.5]
Adverbs of probability	*arguably* (0.63), *likely* (0.7)	(0.5-0.9)
Adverbs of certainty	*ultimately* (0.97), *definitely* (1.0), *indeed* (1.0)	[0.9-1.0]
Adverbs of trueness	*frankly* (0.9), *truthfully* (0.97), *veritably* (1.0)	[0.9-1.0]
Adverbs of falseness	*falsely* (1.0), *erroneously* (1.0)	[0.9-1.0]

The percentage distributions of modal verbs and modal adverbs according to the ranges of *confidence level* are given in Table 5.20. The set of modal operators as well as their *confidence levels* were added to the AttitudeFul lexicon to assist in the analysis of attitude expressed through written language.

Table 5.20 Percentage distributions of modal operators

Category	**Ranges of *confidence level***					
	[0.0-0.2)	**[0.2-0.4)**	**[0.4-0.6)**	**[0.6-0.8)**	**[0.8-1.0)**	**1.0**
Modal verbs	0.0	23.1	15.4	23.1	15.4	23.1
Modal adverbs	9.8	8.2	4.9	4.9	26.2	45.9

Chapter 6

Attitude Analysis Model (@AM)

The automatic analysis and classification of text using fine-grained attitude labels is the main task we address in our research. In this Chapter we introduce our system @AM (**AT**titude **A**nalysis **M**odel), which (1) classifies sentences according to the fine-grained attitude labels (nine affect categories: ‘Anger’, ‘Disgust’, ‘Fear’, ‘Guilt’, ‘Interest’, ‘Joy’, ‘Sadness’, ‘Shame’, ‘Surprise’; four polarity labels for judgment and appreciation: ‘POS jud’, ‘NEG jud’, ‘POS app’, ‘NEG app’; and ‘Neutral’); (2) assigns the strength of the attitude; and (3) determines the level of confidence, with which the attitude is expressed. @AM relies on the *compositionality principle*, a novel linguistic approach based on the rules elaborated for semantically distinct verb classes, and a method considering the hierarchy of concepts. Our compositional approach to automatic recognition of affect, judgment, and appreciation in text extensively deals with the semantics of terms, which allows accurate and robust automatic analysis of attitude type, and broadens the coverage of sentences with complex contextual attitude.

6.1 Architecture of @AM and Primary Stages of the Analysis

The architecture of the developed system for fine-grained attitude recognition in text is presented in Figure 6.1 (in total, @AM algorithm consists of 8460 lines of code in Java; see concise pseudo-code of the @AM algorithm in Appendix B). Given a text, the @AM manager transfers it to the ‘Sentence Splitter’ module, which splits the document into paragraphs, and paragraphs into sentences.

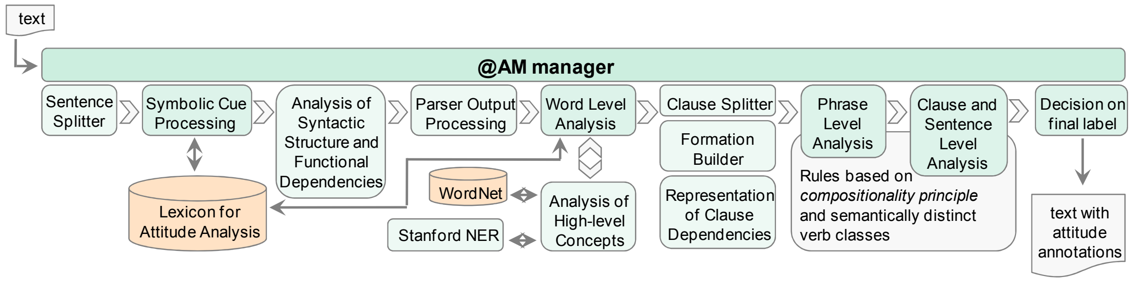

Figure 6.1 Architecture of @AM

Then, each individual sentence is processed by the system for the purpose of detecting the attitude, its strength, and confidence level. During the 'Symbolic Cue Processing' stage, the system analyses the occurrences of emoticons (e.g., 'B-)' ['*cool*'; Joy:04]), abbreviations and acronyms (e.g., '*4gv*' ['*forgive*'; Guilt:0.6]), interjections (e.g., '*alas*' [Sadness:0.3]), 'question mark' and 'exclamation mark', repeated punctuation, and capital letters. In order to disambiguate sentence punctuation marks from those belonging to emoticons, the algorithm delimits the punctuation marks from words. The procedure of detecting emotions based on the symbolic cues is detailed in Section 3.1. The issue of correct processing of abbreviated text by syntactic and dependency parser is resolved by replacing non-emotional abbreviations and acronyms by their proper transcriptions found in the database.

The analysis of syntactic structure and functional dependencies of a sentence is performed by the Connexor Machinese Syntax (*http://www.connexor.eu/technology/machinese/machinesesyntax/*), developed by the Connexor Oy company. This parser assigns meaning-oriented syntactic structure to the analysed sentence and provides the detailed description of sentence tokens (e.g., lemma, part of speech, syntactic function tags, and morphological tags) and dependency relations between them. An example of Connexor Machinese Syntax output for the sentence '*Chewy colorful bears are like a tasty little rainbow.*' is shown in Table 3.1 (see Section 3.2 for details).

6.2 Word Level Annotations and Analysis of High-level Concepts

On the 'Word Level Analysis' stage, the system checks the availability of the sentence tokens in the AttitudeFul database (see Section 5.2) and gets their annotations depending on the category (attitude-conveying word, modifier, 'functional word', modal operator).

In case of an attitude-conveying word, its attitude features are represented as a vector of attitude strengths (intensities): a = [POS jud, NEG jud, POS app, NEG app, Anger, Disgust, Fear, Guilt, Interest, Joy, Sadness, Shame, Surprise]. For example:

a('*high-spirited*') = [0.7 (POS jud),0,0,0,0,0,0,0,0,0.7 (Joy),0,0,0];

a('*nuisance*') = [0,0.8 (NEG jud),0,0.8 (NEG app),0,0,0,0,0,0,0.8 (Sadness),0,0];

a('*graceful*') = [0.3 (POS jud),0,0.3 (POS app),0,0,0,0,0,0,0,0,0,0].

If an attitude-conveying adjective or adverb is in a comparative or superlative form, then its attitude features are multiplied by values 1.2 or 1.4, respectively.

There are several categories of modifiers registered in the AttitudeFul database: adverbs of degree, adverbs of affirmation, negation words, adverbs of doubt, adverbs of falseness, prepositions, and condition operators. The coefficients for intensity degree strengthening or weakening are taken from the database and assigned to the adverbs of degree and adverbs of affirmation detected in a sentence. For example:

coeff('*slightly*') = 0.2 [adverb of degree];

coeff('*surely*') = 1.3 [adverb of affirmation].

Words belonging to other categories of modifiers are marked by the corresponding category retrieved from the database. For example:

'*nothing*' [negation];

'*doubtfully*' [adverb of doubt];

'*erroneously*' [adverb of falseness];

'*without*' [preposition];

'*although*' [condition].

After the word level annotations are taken from the database, the system turns to the analysis of high-level concepts, which will play the key role in the decision on final attitude label of a sentence. A high-level concept of each noun in the sentence is determined based on:

(1) Analysis of the sequence of hypernymic semantic relations of a particular noun in WordNet (Miller 1999).

(2) Annotations from the Stanford Named Entity Recognizer (Stanford NER) (Finkel, Grenager, and Manning 2005).

The hypernym-hyponym relation between lexicalized concepts is a semantic relation that organizes nouns into a lexical hierarchy in WordNet. Hypernymy relation is represented in WordNet by a pointer between the appropriate synsets; this relation means generalization and can be formulated by '*is a*' or '*is a kind of*' (e.g., '*happiness* [hyponym] *is an emotional state* [hypernym]').

Our algorithm for retrieval of high-level concepts from noun hierarchy of WordNet follows the trail of hypernymically related synsets and analyses a sequence of levels from a specific term at the lower level to a generic term at the top of a hierarchy. The following high-level concepts are of our main interest:

ABSTRACTION, *ACTIVITY*, *ANIMAL*, *ARTIFACT*, *ATTRIBUTE*, *BODY*, *COGNITION*, *COMMUNICATION*, *ENTITY*, *EVENT*, *FEELING*, *FOOD*, *GROUP*, *HUMAN*, *LOCATION*, *MAN*, *MOTIVATION*, *NATURAL OBJECT*, *NATURAL PHENOMENON*, *OBJECT*, *ORGANISM*, *PERSON*, *PLANT*, *POSSESSION*, *PROCESS*, *PSYCHOLOGICAL FEATURE*, *QUANTITY*, *RELATION*, *SHAPE*, *STATE*, *SUBSTANCE*, and *TIME*.

Given a noun from the sentence, our system retrieves its high-level concept (one of the listed above). For example:

'*student*' => *PERSON*;

'*miracle*' => *EVENT*;

'*decoration*' => *ARTIFACT*.

In order to determine high-level concepts of named entities contained in the sentence, our @AM system employs Stanford Named Entity Recognizer, which analyses the sentence and annotates the available named entities using the following labels: *PERSON*, *ORGANIZATION*, and *LOCATION*. For example, given a sentence '*It is Max's dream to become an engineer and to work for NASA in US*', Stanford NER recognizes '*Max*' as a *PERSON*, '*NASA*' as an *ORGANIZATION*, and '*US*' as a *LOCATION*.

As was mentioned earlier, the analysis of high-level concepts will contribute to the decision on final attitude label of a sentence.

6.3 Representation of Sentence Structure

The algorithms developed for the 'Clause Splitter', 'Formation Builder', and 'Representation of Clause Dependencies' modules are based on the analysis of parser output. Let us consider the following example (parser output is given in Table 6.1): '*The museum experience is better, when you have a guide, who really loves what he is doing*'.

Table 6.1 An example of Connexor Machinese Syntax output for the complex sentence

Token id	Text	Lemma	Syntactic relations and dependencies	Syntax and morphology
1	*The*	*the*	det:>3	@DN> %>N DET
2	*museum*	*museum*	attr:>3	@A> %>N N NOM SG
3	*experience*	*experience*	subj:>4	@SUBJ %NH N NOM SG
4	*is*	*be*	main:>0	@+FMAINV %VA V PRES SG3
5	*better*	*good*	comp:>4	@PCOMPL-S %NH A CMP
6	*,*	*,*		
7	*when*	*when*	tmp:>9	@ADVL %EH ADV WH
8	*you*	*you*	subj:>9	@SUBJ %NH PRON PERS NOM
9	*have*	*have*	tmp:>4	@+FMAINV %VA V PRES
10	*a*	*a*	det:>11	@DN> %>N DET SG
11	*guide*	*guide*	obj:>9	@OBJ %NH N NOM SG
12	*,*	*,*		
13	*who*	*who*	subj:>15	@SUBJ %NH <Rel> PRON WH NOM
14	*really*	*really*	meta:>15	@ADVL %EH ADV
15	*loves*	*love*	mod:>11	@+FMAINV %VA V PRES SG3
16	*what*	*what*	obj:>19	@OBJ %NH PRON WH
17	*he*	*he*	subj:>18	@SUBJ %NH PRON PERS NOM SG3
18	*is*	*be*	v-ch:>19	@+FAUXV %AUX V PRES SG3
19	*doing*	*do*	obj:>15	@-FMAINV %VA ING

The 'Clause Splitter' module detects the boundaries of the clauses (either independent or dependent) in the following manner:

(1) First, it finds the verbs contained in the sentence and involved in clauses as key elements based on the syntactic functional tags for a finite main predicator ('@+FMAINV'), a nonfinite main predicator ('@-FMAINV'), and a nonfinite clause as preposition

complement (‘@<P-FMAINV’). In the above example, the algorithm detects four main verbs: ‘*is*’ (id: 4), ‘*have*’ (id: 9), ‘*loves*’ (id: 15), and ‘*doing*’ (id: 19).

(2) Then, the system analyses the dependency functions, syntactic and morphological tags of the words related to these key verbs, and marks the beginning and ending of each clause.

As the result, ‘Clause Splitter’ module detects four clauses in the example sentence:

Clause 1 (main independent clause):

‘{*the museum experience is better*}’.

Clause 2 (subordinating clause, introduced by conditional conjunction and containing two embedded ‘object relative’ clauses):

‘{*when you have a guide,* {*who really loves* {*what he is doing*}}}’.

Clause 3 (‘object relative’ clause itself containing an embedded ‘object relative’ clause):

‘{*who really loves what he is doing*}’.

Clause 4 (‘object relative’ clause):

‘{*what he is doing*}’.

Using the data from the ‘Clause Splitter’, the ‘Formation Builder’ module represents each clause as a set of formations: Subject formation (SF), Verb formation (VF) and Object formation (OF), each of which may consist of a main element (subject, verb, or object) and its attributives and complements. The developed algorithm can detect not only subjects represented by noun phrases, but also subjects represented by gerund (non-finite verb form), by an infinitive, or by a full clause, introduced by ‘*that*’, itself containing a subject and a predicate. Figure 6.2 shows the schematic representation of the clauses and formations in the example sentence ‘*The museum experience is better, when you have a guide, who really loves what he is doing*’.

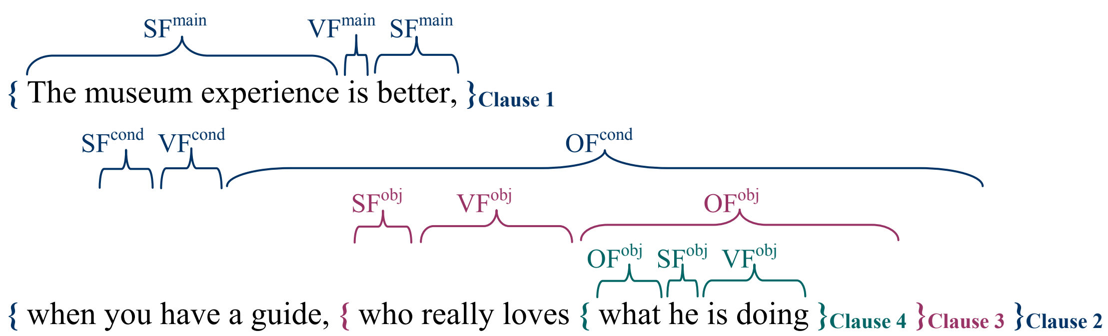

Figure 6.2 The schematic representation of the clauses and formations in the complex sentence

The 'Representation of Clause Dependencies' module is responsible for building a so-called 'relation matrix', which contains information about the dependencies between different clauses in a compound, complex, or complex-compound sentences (e.g., coordination, subordination, condition, contingency, etc.). Table 6.2 shows the 'relation matrix' for the above example.

Table 6.2 The example of a 'relation matrix' built for the complex sentence

Clause	**Verb id**	**Syntactic relation**
Clause 1	4	main:>0
Clause 2	9	tmp:>4
Clause 3	15	mod:>11 => obj:>9
Clause 4	19	obj:>15

6.4 Phrase, Clause, and Sentence Level Analyses

Different types of sentences, namely, simple, compound, complex (with complement or relative clauses), and complex-compound, are handled by the @AM modules aimed at processing the attitude information on various grammatical levels (phrase, clause, and sentence). Our @AM relies on the *compositionality principle*, a novel linguistic approach based on the rules elaborated for semantically distinct verb classes, and a method considering the hierarchy of concepts.

6.4.1 Compositionality Principle

> '*The full story of how lexical items reflect attitudes is more complex than simply counting the valences of terms.*' Polanyi and Zaenen (2004)

Words in a sentence are interrelated and, hence, each of them can influence the overall meaning and attitudinal bias of a statement. Our algorithm for attitude classification is designed based on the *compositionality principle*, according to which we determine the attitudinal meaning of a sentence by composing the pieces that correspond to lexical units or other linguistic constituent types governed by the rules of *polarity reversal, aggregation (fusion)*, *propagation*, *domination*, *neutralization*, and *intensification*, at various grammatical levels.

Polarity reversal means that a phrase or statement containing an attitude-conveying term/phrase with prior positive polarity becomes negative, and vice versa. The rule of *polarity reversal* is applied in three cases:

(1) Negation word-modifier in relation with an attitude-conveying statement:
‘*never*’ & POS(‘*succeed*’) => NEG(‘*never succeed*’).

(2) Adverb of doubt or adverb of falseness in relation with an attitude-conveying statement:
‘*scarcely*’ & POS(‘*relax*’) => NEG(‘*scarcely relax*’).

(3) ‘Functional’ word of *reversing* type in relation with an attitude-conveying statement:
adjective ‘*reduced*’ & POS(‘*enthusiasm*’) => NEG(‘*reduced enthusiasm*’).

In the case of judgment and appreciation, the use of the *polarity reversal* rule is straightforward (‘POS jud’ <=> ‘NEG jud’, ‘POS app’ <=> ’NEG app’). However, it is not trivial to find pairs of opposite emotions in the case of a fine-grained classification, except for ‘Joy’ and ‘Sadness’. Therefore, we assume that:

(1) The opposite emotion for three positive emotions, namely, ‘Interest’, ‘Joy’, and ‘Surprise’, is ‘Sadness’ (i.e. ‘POS aff’ => ‘Sadness’).

(2) The opposite emotion for six negative emotions, namely, ‘Anger’, ‘Disgust’, ‘Fear’, ‘Guilt’, ‘Sadness’, and ‘Shame’, is ‘Joy’ (i.e. ‘NEG aff’ => ‘Joy’).

The rules of *aggregation (fusion)* are as follows:

(1) If polarities of attitude-conveying terms in adjective-noun, noun-noun, adverb-adjective, adverb-verb phrases have opposite directions, mixed polarity with dominant polarity of a premodifier is assigned to the phrase:
POS(‘*beautiful*’) & NEG(‘*fight*’) => POS-neg(‘*beautiful fight*’);
NEG(‘*shamelessly*’) & POS(‘*celebrate*’) => NEG-pos(‘*shamelessly celebrate*’).

(2) Otherwise, the resulting polarity is based on the equal polarities of terms, and the strength of attitude is measured as a maximum between polarity scores (intensities) of terms (*max*(Score1,Score2)).

The rule of *propagation* is useful, as proposed in (Nasukawa and Yi 2003), for the task of the detection of local sentiments for given subjects. The ‘propagation’ verbs propagate the sentiment towards the arguments (e.g., positive polarity of the verb ‘*respect*’ is propagated to the object in ‘*to* *respect OBJ*'); the 'transfer' verbs transmit sentiments among the arguments (e.g., positive or negative polarity of the object is transmitted to the subject in '*SUBJ serves OBJ*'). The rule of *propagation* is applied when a verb of 'propagation' or 'transfer' type is used in a phrase/clause and sentiment of an argument that has prior neutral polarity needs to be investigated. For example: PROP-POS('*to admire*') & '*his behaviour*' => POS('*his behaviour*'); '*Mr. X*' & TRANS('*supports*') & NEG('*crime business*') => NEG('*Mr. X*').

The rules of *domination* are as follows:

(1) If polarities of a verb (this rule is applied only for certain classes of the verbs) and an object in a clause have opposite directions, then the polarity of verb is prevailing:
NEG('*to deceive*') & POS('*hopes*') => NEG('*to deceive hopes*').

(2) If a compound sentence joints clauses using the coordinate connector '*but*', the attitude features of a clause following after the connector are considered as dominant:
'NEG(*It was hard to climb a mountain all night long*)*, but* POS(*a magnificent view rewarded the traveler at the morning*).' => POS(whole sentence).

The rule of *neutralization* is applied when:

(1) The attitude-conveying statement is introduced by some of the prepositions:
'*despite*' & NEG('*worries*') => NEUTRAL('*despite worries*').

(2) The statement is conditional:
'*even if*' & NEG('*they were not perfect*') => NEUTRAL('*even if they were not perfect*').

The rule of *intensification* means strengthening (or weakening) the attitude score (intensity), and is applied when:

(1) Adverb of degree or adverb of affirmation relates to the attitude-conveying term:
Pos_score('*almost happy*') < Pos_score('*happy*') < Pos_score('*extremely happy*').

(2) Attitude-conveying adjective or adverb is used in a comparative or superlative form:
Neg_score('*ungrateful*') < Neg_score('*more ungrateful*') < Neg_score ('*most ungrateful*').

(3) 'Functional' word of *intensifying* type relates to the attitude-conveying statement:
Pos_score('*success*') < Pos_score('*rapidly-growing success*').

While applying the compositionality principle, we consecutively assign attitude features to words, phrases, formations, clauses, and finally, to the whole sentence.

6.4.2 Consideration of the Semantics of Verbs

All sentences must include a verb, because the verb tells us what action the subject is performing and object is receiving. In order to elaborate rules for the attitude analysis based on the semantics of verbs, we investigated VerbNet (Kipper et al. 2007), the largest on-line verb lexicon that is organized into verb classes characterized by syntactic and semantic coherence among members of a class. Based on the thorough analysis of 270 first-level classes of VerbNet and their members, 73 verb classes (1) were found useful for the task of attitude analysis, and (2) were further classified into 22 classes (1129 verbs) differentiated by the role that members play in attitude analysis and by rules applied to them. Our classification is given in Table 6.3.

Table 6.3 Verb classes defined for attitude analysis

Verb class	Verb samples	Examples of VerbNet classes
1 Psychological state or emotional reaction		
1.1 Object-centered (oriented) emotional state	*appreciate*, *distrust*	Admire-31.2, Care-88.1 etc.
1.2 Subject-driven change in emotional state (transitive)	*charm*, *inspire*, *bother*	Amuse-31.1 etc.
1.3 Subject-driven change in emotional state (intransitive)	*appeal to*, *grate on*	Appeal-31.4
2 Judgment		
2.1 Positive judgment	*bless*, *honor*	Judgment-33-pos.
2.2 Negative judgment	*blame*, *punish*	Judgment-33-neg. etc.
3 Favorable attitude	*accept*, *allow*, *tolerate*	Allow-64, Appoint-29.1 etc.
4 Adverse (unfavorable) attitude	*discourage*, *elude*, *forbid*	Forbid-67, Refrain-69 etc.
5 Favorable or adverse calibratable changes of state	*grow*, *decline*	Calibratable_cos-45.6
6 Verbs of removing		
6.1 Verbs of removing with neutral charge	*delete*, *remove*	Remove-10.1
6.2 Verbs of removing with negative charge	*deport*, *expel*	Banish-10.2-neg., Fire-10.10 etc.
6.3 Verbs of removing with positive charge	*evacuate*, *cure*	Banish-10.2-pos., Free-80 etc.
7 Negatively charged change of state	*break*, *crush*, *smash*	Break-45.1 etc.
8 Bodily state and damage to the body	*sicken*, *injure*	Change_bodily_state-40.8.4 etc.

Table 6.3 (Continued)

Verb class	Verb samples	Examples of VerbNet classes
9 Aspectual verbs		
9.1 Initiation, continuation of activity, and sustaining	*begin*, *continue*, *maintain*	Begin-55.1, Sustain-55.6 etc.
9.2 Termination of activity	*quit*, *finish*	Complete-55.2, Stop-55.4
10 Preservation	*defend*, *insure*	Defend-85
11 Verbs of destruction and killing	*damage*, *poison*	Destroy-44, Murder-42.1 etc.
12 Disappearance	*disappear*, *die*	Disappearance-48-2
13 Limitation and subjugation	*confine*, *restrict*	Limit-76, Subjugate-42.3
14 Assistance	*succor*, *help*	Help-72
15 Obtaining	*win*, *earn*	Get-13.5.1
16 Communication indicator/reinforcement of attitude	*guess*, *complain*, *deny*	Advise-37.9, Conjecture-29.5 etc.
17 Verbs of leaving	*abandon*, *desert*	Leave-51.2, Resign-10.11
18 Changes in social status or condition	*canonize*, *widow*	Orphan-29.7
19 Success and failure		
19.1 Success	*succeed*, *manage*	Succeed-74-pos.
19.2 Failure	*fail*, *flub*	Succeed-74-neg.
20 Emotional nonverbal expression	*smile*, *weep*	Nonverbal_expression-40.2
21 Social interaction	*marry*, *divorce*	Correspond-36.1, Marry-36.2 etc.
22 Transmitting verbs	*supply*, *provide*	Fulfilling-13.4.1 etc.

For each of our verb classes, we developed set of the rules that are applied to attitude analysis on the phrase/clause-level. Some verb classes (e.g., "*Psychological state or emotional reaction*", "*Judgment*", "*Bodily state and damage to the body*", "*Preservation*" etc.) include verbs annotated by attitude type, prior polarity orientation, and the attitude strength (score). The attitude features of phrases that contain positively or negatively charged verbs from such classes are context-sensitive and are defined by means of the rules designed for each of the class.

As an example, we provide short description and rules elaborated for the subclass "*Object-centered (oriented) emotional state*". The features are as follows: a subject experiences the emotions towards some stimulus; verb prior polarity is positive or negative; attitude is context-sensitive.

Verb-Object rules (subject is ignored):

(1) "Interior perspective" (subject's inner emotion state or attitude):

S & V-pos('*admires*') & O-pos('*his brave heart*') => (fusion, *max*(V_score,O_score)) => 'POS aff'.

S & V-pos('*admires*') & O-neg('*the mafia leader*') => (verb valence dominance, V_score) => 'POS aff'.

S & V-neg('*disdains*') & O-pos('*his honesty*') => (verb valence dominance, V_score) => 'NEG aff'.

S & V-neg('*disdains*') & O-neg('*the pathological liars*') => (fusion, *max*(V_score,O_score)) => 'NEG aff'.

(2) "Exterior perspective" (social/ethical judgment):

S & V-pos('*admires*') & O-pos('*his brave heart*') => (fusion, *max*(V_score,O_score)) => 'POS jud'.

S & V-pos('*admires*') & O-neg('*the mafia leader*') => (verb valence reversal, *max*(V_score,O_score)) => 'NEG jud'.

S & V-neg('*disdains*') & O-pos('*his honesty*') => (verb valence dominance, *max*(V_score,O_score)) => 'NEG jud'.

S & V-neg('*disdains*') & O-neg('*the pathological liars*') => (verb valence reversal, *max*(V_score,O_score)) => 'POS jud'.

(3) In case of neutral object => attitude type and prior polarity of verb, verb score (V_score).

Verb-PP (prepositional phrase) rules:

(1) In case of the negatively charged verb and PP starting with '*from*' => verb dominance:

S & V-neg('*suffers*') & PP-neg('*from illness*') =>

"interior perspective": (verb valence dominance, *max*(V_score,PP_score)) => 'NEG aff'.

"exterior perspective": (verb valence dominance, *max*(V_score,PP_score)) => 'NEG jud'.

S & V-neg('*suffers*') & PP-pos ('*from love*') =>

"interior perspective": (verb valence dominance, *max*(V_score,PP_score)) => 'NEG aff'.

"exterior perspective": (verb valence dominance, *max*(V_score,PP_score)) => 'NEG jud'.

(2) In case of the positively charged verb and PP starting with ‘*in*’/‘*for*’ => handle PP the same way as object (see above):

S & V-pos(‘*believes*’) & PP-neg(‘*in evil*’) =>

“interior perspective”: (verb valence dominance, V_score) => ‘POS aff’.

“exterior perspective”: (verb valence reversal, *max*(V_score,PP_score)) => ‘NEG jud’.

S & V-pos(‘*believes*’) & PP-pos(‘*in kindness*’) =>

“interior perspective”: (fusion, *max*(V_score,PP_score)) => ‘POS aff’.

“exterior perspective”: (fusion, *max*(V_score,PP_score)) => ‘POS jud’.

In the majority of rules the strength of attitude is measured as a maximum between attitude scores (for example, the attitude conveyed by ‘*to suffer from grave illness*’ is stronger than that of ‘*to suffer from slight illness*’).

In contrast to the rules of “*Object-centered (oriented) emotional state*” subclass, which ignore the attitude features of a subject in a sentence, the rules elaborated for the “*Subject-driven change in emotional state (transitive)*” disregard the attitude features of an object, as in sentences containing the members of this subclass the object experiences the emotion, and the subject causes the emotional state. Some examples are given below:

S(‘*Classical music*’) & V-pos(‘*calmed*’) & O-neg(‘*disobedient child*’) =>

“interior perspective”: (verb valence dominance, V_score) => ‘POS aff’.

“exterior perspective”: (verb valence dominance, V_score) => ‘POS app’.

S-neg(‘*Fatal consequences of GM food intake*’) & V-neg(‘*frighten*’) & O(‘*me*’) =>

“interior perspective”: (fusion, *max*(S_score,V_score)) => ‘NEG aff’.

“exterior perspective”: (fusion, *max*(S_score,V_score)) => ‘NEG app’.

The rules for the sentence containing the verb from the “*Subject-driven change in emotional state (intransitive)*” are as follows:

S-neg(‘*Max’s goal-seeking impudence*’) & V-pos(‘*appeals*’) & PP(‘*to his friends*’) =>

“interior perspective”: (verb valence dominance, V_score) => ‘POS aff’.

“exterior perspective”: (verb valence reversal, *max*(S_score,V_score)) => ‘NEG jud’.

The Verb-Object rules for the “*Judgment*” subclasses, namely “*Positive judgment*” and “*Negative judgment*” (judgment or opinion that someone may have in reaction to something), are very close to those defined for the subclass "*Object-centered (oriented) emotional state*". However, Verb-PP rules have some specifics: for both positive and negative judgment verbs, the system handles the PP starting with '*for*'/'*of*'/'*as*' the same way as object in the Verb-Object rules. For example:

(1) "*Positive judgment*" subclass:

S('*They*') & V-pos('*praised*') & O('*him*') & PP-pos ('*for his successful speech*') =>

"interior perspective": (fusion, *max*(V_score,PP_score)) => 'POS jud'.

"exterior perspective": (fusion, *max*(V_score,PP_score)) => 'POS jud'.

(2) "*Negative judgment*" subclass:

S('*He*') & V-neg('*blamed*') & O-pos('*innocent person*') =>

"interior perspective": (verb valence dominance, V_score) => 'NEG jud'.

"exterior perspective": (verb valence dominance, *max*(V_score,O_score)) => 'NEG jud'.

S('*They*') & V-neg('*punished*') & O('*him*') & PP-neg('*for his misdeed*') =>

"interior perspective": (fusion, *max*(V_score,PP_score)) => 'NEG jud'.

"exterior perspective": (verb valence reversal, *max*(V_score,PP_score)) => 'POS jud'.

Verbs from the "*Favorable attitude*" and "*Adverse (unfavorable) attitude*" classes have prior neutral polarity and positive or negative reinforcement, correspondingly, that means that they have an impact only on the attitude features of a non-neutral phrase (object in a sentence written in active voice, or subject in a sentence written in passive voice, or PP in case of some verbs). The rules are as follows:

(1) If verb belongs to the "*Favorable attitude*" class and the polarity of a phrase is not neutral, then the attitude score of the phrase is intensified (symbol '^' means intensification):

S('*They*') & [V pos. reinforcement]('*elected*') & O-pos('*fair judge*') =>

'POS app'; O_score^.

S('*They*') & [V pos. reinforcement]('*elected*') & O-neg('*corrupt candidate*') =>

'NEG app'; O_score^.

(2) If verb belongs to the "*Adverse (unfavorable) attitude*" class and the polarity of a phrase is not neutral, then the polarity of phrase is reversed and score is intensified:

S('*They*') & [V neg. reinforcement]('*prevented*') & O-neg('*the spread of disease*') =>

'POS app'; O_score^.

S-pos('*His achievements*') & [V neg. reinforcement]('*were overstated*') =>

'NEG app'; S_score^.

The members of the "*Favorable or adverse calibratable changes of state*" class describe positive or negative changes along a scale. They involve entities that themselves have a measurable attribute. The verbs that describe favorable changes have prior positive reinforcement and influence (intensify) the attitude strength, while the verbs characterizing adverse changes have negative reinforcement and, thus, reverse and intensify the attitude features of the subject. For example:

S-neg('*The level of crime*') & [V favorable change]('*is growing*') =>

'NEG app'; S_score^.

S-neg('*The level of crime*') & [V adverse change]('*is decreasing*') =>

'POS app'; S_score^.

Below are some examples of processing the sentences with verbs from "*Verbs of removing*" class:

(1) "*Verbs of removing with neutral charge*" subclass:

S('*The tape-recorder*') & [V neutral removal]('*automatically ejects*') & O('*the tape*') =>

'Neutral'.

S('*The safety invention*') & [V neutral removal]('*ejected*') & O('*the pilot*') & PP-neg('*from burning plane*') =>

(PP valence reversal, PP_score^) => 'POS app'.

(2) "*Verbs of removing with negative charge*" subclass:

S('*Manager*') & [V negative removal]('*fired*') & O-neg('*careless employee*') & PP('*from the company*') =>

(object valence reversal, *max*(V_score,O_score)) => 'POS app'.

(3) "*Verbs of removing with positive charge*" subclass:

S('*They*') & [V positive removal]('*evacuated*') & O('*children*') & PP-neg('*from dangerous place*') =>

(verb valence dominance, *max*(V_score,PP_score)) => 'POS app'.

There are two subclasses in the "*Aspectual verbs*" class: "*Initiation, continuation of activity, and sustaining*" and "*Termination of activity*". The members of the former subclass impact the attitude strength of the subject (in case of intransitive use) or object (in case of transitive use):

S-neg(‘*Next financial crisis*’) & [V initiation](‘*begins*’) =>

‘NEG app’; S_score^.

S(‘*They*’) & [V initiation](‘*resumed*’) & O-pos(‘*their courage*’) =>

‘POS app’; O_score^.

The “*Termination of activity*” verbs have negative reinforcement:

S(‘*They*’) & [V termination](‘*discontinued*’) & O-pos(‘*helping children*’) =>

‘NEG app’; O_score^.

Along with the modal verbs and modal adverbs, members of the “*Communication indicator/reinforcement of attitude*” class also indicate the confidence level or degree of certainty concerning given opinion. The features are as follows: a subject (communicator) expresses some statement with/without attitude; the statement is a PP starting with ‘*of*’, ‘*on*’, ‘*against*’, ‘*about*’, ‘*concerning*’, ‘*regarding*’, ‘*that*’, ‘*how*’ etc.; the ground is positive or negative; and the reinforcement is positive or negative. Using the sample sentences, we explain the rules:

(1) If the polarity of expressed statement is neutral, then the attitude is neutral:

S(‘*Professor*’) & [V pos. ground, pos. reinforcement, confidence:0.83](‘*dwelled*’) & PP-neutral(‘*on a question*’) =>

‘Neutral’.

(2) If the polarity of expressed statement is not neutral and the reinforcement of a verb is positive, then the score of the statement (PP) is intensified:

S(‘*Jane*’) & [V neg. ground, pos. reinforcement, confidence:0.8](‘*is complaining*’) & PP-neg(‘*of a headache again*’) =>

‘NEG app’; PP_score^; confidence:0.8.

(3) If the polarity of expressed statement is not neutral and the reinforcement of a verb is negative, then the polarity of the statement (PP) is reversed and its score is intensified:

S(‘*Max*’) & [V neg. ground, neg. reinforcement, confidence:0.2](‘*doubt*’) & PP-neg{‘*that*’ S-pos(‘*his good fortune*’) & [V termination](‘*will ever end*’)} =>

‘POS app’; PP_score^; confidence:0.2.

In the last example, to determine the attitude of PP, we apply the rule for the verb '*end*' from the "*Termination of activity*" subclass of the "*Aspectual verbs*" class, which reverses the non-neutral polarity of a subject (in case of intransitive use of verb) or object (in case of transitive use of verb).

6.4.3 Decision on Attitude Label

There are three main attitude-related concepts: affect (AFF), judgment (JUD), and appreciation (APP). Affect is a personal emotional state; judgment is a social or ethical appraisal of other's behaviour, character, skills, etc.; and appreciation is an evaluation of phenomena, events, objects. Based on these definitions, we measure the potential of a clause or sentence to convey affect, judgment, and appreciation.

The decision on the most appropriate final label for the clause, in case @AM annotates it using different attitude types according to the words with multiple annotations (e.g., a('*unfriendly*') = [0,0.5 (NEG jud),0,0.5 (NEG app),0,0,0,0,0,0,0.5 (Sadness),0,0]) or based on the availability of the words conveying different attitude types (e.g., '*I am delighted with the new comfortable house*', where a('*delight*') = [0,0,0,0,0,0,0,0,0,0.8 (Joy),0,0,0] and a('*comfortable*') = [0,0,0.1 (POS app),0,0,0,0,0,0,0,0,0,0]), is made based on the analysis of:

(1) Morphological tags of nominal heads and their premodifiers in the clause: first person pronoun (PRON, PERS, NOM/ACC/GEN, SG1/PL1), third person pronoun (PRON, PERS, NOM/ACC/GEN, SG3/PL3), demonstrative pronoun (PRON, DEM), reciprocal pronoun (PRON, RECIPR), nominative or genitive noun (N, NOM or GEN), etc.

(2) High-level concepts of nouns based on WordNet (WN) (see Section 6.2 for details): *ARTIFACT*, *FOOD*, *HUMAN*, *NATURAL PHENOMENON*, and other.

(3) High-level concepts of named entities based on the annotations from the Stanford Named Entity Recognizer (see Section 6.2 for details): *PERSON*, *ORGANIZATION*, and *LOCATION* entities.

The possible values of a potential of a clause or simple sentence to convey affect, judgment, and appreciation are 0 (indicating no potential) or 1 (indicating that the potential exists); and they are based on the analysis of features listed above.

We developed the algorithm using the following assumptions:

Potential(AFF) = 1, if a clause contains at least one of the following elements:

(1) Nominal first person pronouns (NomFPP): '*I*', '*we*', etc.

(2) Accusative first person pronouns (AccFPP): '*me*', '*us*', etc.

(3) Reflexive first person pronouns (RefFPP): '*myself*', etc.

Otherwise, *Potential*(AFF) = 0.

Potential(JUD) = 1, if a clause contains at least one of the following elements:

(1) Nouns characterized by the following high-level concepts: *PERSON*, *MAN*, and *HUMAN*.

(2) Named entity labelled by the Stanford NER as a *PERSON* or *ORGANIZATION*.

(3) Nominal third person pronouns (NomTPP): '*he*', '*she*', etc.

(4) Accusative third person pronouns (AccTPP): '*him*', '*them*', etc.

(5) Reflexive third person pronouns (RefTPP): '*herself*', '*himself*', etc.

(6) Relative wh-pronouns (RelWhP) '*who*' and '*whom*'.

(7) Genitive wh-pronoun (GenWhP) '*whose*' that is related to the noun from (1): '*whose student*'.

(8) Genitive first person pronoun (GenFPP) that is related to the noun from (1): '*my consultant*', '*our doctor*'.

(9) Genitive third person pronoun (GenTPP) that is related to the noun from (1): '*his bodyguard*', '*their assistants*'.

(10) Genitive first person pronoun (GenFPP) that is related to the named entity from (2): '*our Diana*'.

(11) Genitive third person pronoun (GenTPP) that is related to the named entity from (2): '*their NASA*'.

(12) Genitive named entity (GenNE) that is related to the noun from (1): '*John's friend*', '*Japan's Prime Minister*'.

(13) Genitive named entity (GenNE) that is related to the named entity from (2).

Otherwise, *Potential*(JUD) = 0.

Potential(APP) = 1, if a clause contains at least one of the following elements:

(1) Nouns characterized by the high-level concepts other than *PERSON*, *MAN*, and *HUMAN*: *ARTIFACT*, *EVENT*, *NATURAL OBJECT*, *PLANT*, *SHAPE*, and other.

(2) Named entity labelled by the Stanford NER as a *LOCATION*.

(3) Genitive wh-pronoun (GenWhP) '*whose*' that is related to the noun from (1): '*whose wedding*'.

(4) Genitive first person pronoun (GenFPP) that is related to the noun from (1): '*my appartment*'.

(5) Genitive third person pronoun (GenTPP) that is related to the noun from (1): '*her concert*'.

(6) Genitive first person pronoun (GenFPP) that is related to the named entity from (2): '*my Belarus*'.

(7) Genitive third person pronoun (GenTPP) that is related to the named entity from (2): '*their Hawaiian Islands*'.

(8) Genitive named entity (GenNE) that is related to the noun from (1): '*Mary's flowers*'.

(9) Genitive named entity (GenNE) that is related to the named entity from (2).

Otherwise, *Potential*(APP) = 0.

In case of the conflict between different attitude concepts (AFF, JUD, APP) in the final attitude vector of a clause, the algorithm takes into account only the attitude concepts with non-zero potentials (*Potential* = 1), and the following rules are applied:

(1) If attitude concepts with non-zero potentials are (AFF, JUD, APP), or (AFF, JUD), or (AFF, APP), then affect is considered as prevailing, and the affective label with the maximum attitude score in a clause vector is taken as final (in case of affective labels with equal scores, the affective label is selected randomly).

(2) If there is a conflict between JUD and APP concepts (JUD, APP) with non-zero potentials, then the judgment or appreciation label with the maximum attitude score in a clause vector is taken as final (in case of the labels with equal scores, the label is selected randomly).

For example, @AM outputs different attitude labels for the following sentences containing only one attitude-conveying word '*unfriendly*' (a('*unfriendly*') = [0,0.5 (NEG jud),0,0.5 (NEG app),0,0,0,0,0,0,0.5 (Sadness),0,0]): '*I feel highly unfriendly attitude towards me*', '*The salesperson was really unfriendly*', and '*Plastic bags are environment unfriendly*':

(1) *I* [NomFPP] *feel highly* [modifier: adverb of degree: 1.7] *unfriendly* [NEG aff (Sadness): 0.5; NEG jud: 0.5; NEG app: 0.5] *attitude* [WN: *COGNITION*] *towards me* [AccFPP] =>
=> 'NEG aff' ('Sadness'): 0.85.

(2) *The salesperson* [WN: *PERSON*] *was really* [modifier: adverb of degree: 1.55] *unfriendly* [NEG aff (Sadness): 0.5; NEG jud: 0.5; NEG app: 0.5] =>
=> 'NEG jud': 0.78.

(3) *Plastic bags* [WN: *ARTIFACT*] *are environment* [WN: *STATE*] *unfriendly* [NEG aff (Sadness): 0.5; NEG jud: 0.5; NEG app: 0.5] =>
=> 'NEG app': 0.5.

6.5 Walking through the Examples

In this Section we demonstrate the process of attitude analysis on the example sentences. The first example is '*My enthusiasm and excitement disappeared, when this unfair decision was made*' (Figure 6.3). This sentence is composed of two clauses: the main clause (Clause 1) and conditional clause (Clause 2). The attitude processing is given in Table 6.4. The first clause conveys negative affect ('Sadness' emotion). Although the second clause may convey judgment or appreciation according to the attitude vector of the adjective '*unfair*', the attitude of this clause is neutralized due to conditionality (conjunction '*when*'). As seen from Table 6.4, the overall attitude of this sentence is 'Sadness' with attitude score 1.0.

$$\{\ \overbrace{\text{My enthusiasm and excitement}}^{SF^{main}}\ \overbrace{\text{disappeared,}}^{VF^{main}}\ \}_{\textbf{Clause 1}}$$

$$\{\ \overbrace{\text{when this unfair decision}}^{SF^{cond}}\ \overbrace{\text{was made}}^{VF^{cond}}\ \}_{\textbf{Clause 2}}$$

Figure 6.3 Clauses and formations in the sentence '*My enthusiasm and excitement disappeared, when this unfair decision was made*'

Table 6.4 Attitude processing in the sentence '*My enthusiasm and excitement disappeared, when this unfair decision was made*'

Analysis of Clause 1 ('*my enthusiasm and excitement disappeared*'):

SF^{main} = {'*my enthusiasm and excitement*'}:

a('*enthusiasm*') = [0,0,0,0,0,0,0,0,0.8 (Interest),0.5 (Joy),0,0,0].

a('*excitement*') = [0,0,0,0,0,0,0,0,0.9 (Interest),0.6 (Joy),0,0,0].

Aggregation (fusion) rule: a(SF^{main}) = [0,0,0,0,0,0,0,0,0.9 (Interest),0.6 (Joy),0,0,0].

VF^{main} = {'*disappeared*'}:

Verb class: "*Disappearance*" [neutral polarity; negative reinforcement].

Rule of SF attitude reversal and intensification: a(Clause 1) = a(SF^{main} & VF^{main}) = reinforcement_coeff * reversal(a(SF^{main})) = 1.2 * [0,0,0,0,0,0,0,0,0,0,0.9 (Sadness),0,0] = [0,0,0,0,0,0,0,0,0,0,1.08 (Sadness),0,0].

Attitude label and score: 'Sadness': 1.0.

Analysis of Clause 2 ('*when this unfair decision was made*'):

SF^{cond} = {'*this unfair decision*'}:

a(SF^{cond}) = a('*unfair*') = [0,0.2 (NEG jud),0,0.2 (NEG app),0,0,0,0,0,0,0,0,0].

VF^{cond} = {'*when* [condition]... *was made*'}:

a(VF^{cond}) = neutral vector.

Rule of clause neutralization due to condition: a(Clause 2) = a(SF^{cond} & VF^{cond}) = neutral vector.

Attitude label and score: 'Neutral': 0.0.

Sentence-level result:

Aggregation (fusion) rule: a(Clause 1 & Clause 2) = [0,0,0,0,0,0,0,0,0,0,1.08 (Sadness),0,0].

Attitude label and score: 'Sadness': 1.0.

Next example is presented by the complex-compound sentence '*Some believe that GM foods could help to reduce global hunger, but a number of studies over the past decade have revealed that genetically modified foods can pose serious threats to the consumers' health and Earth's ecosystem, because novel and possibly hazardous genetic material disrupts the natural harmony and causes the uncontrolled biological pollution that may harmfully influence the human health and environment*'. Figure 6.4 shows the schematic representation of the output of the 'Clause Splitter' and the 'Formation Builder' modules: boundaries of the clauses and formations in the given sentence. The system detects eight clauses and analyses the relations between them. The steps of the attitude processing are summarized in Table 6.5. The result of attitude analysis revealed that the overall attitude of this sentence is negative appreciation with attitude score 1.0.

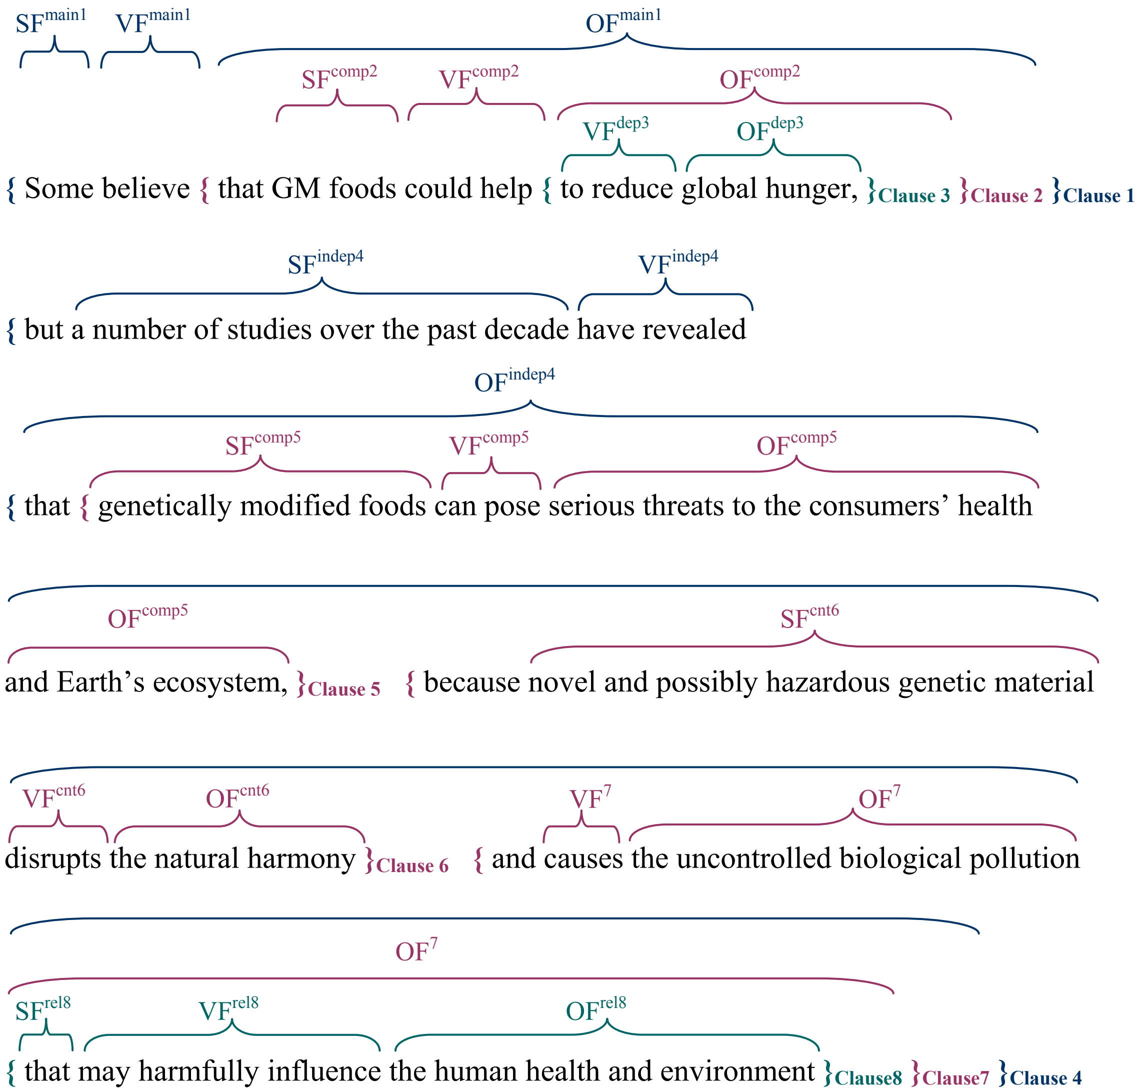

Figure 6.4 Clauses and formations in the complex-compound sentence containing eight clauses

Table 6.5 Attitude processing in the complex-compound sentence containing eight clauses

Steps in @AM algorithm
1. First, @AM analyses the clauses that do not have dependent clauses, namely: Clause 3 and Clause 8.
Analysis of Clause 3 ('*to reduce global hunger*'):
VF^{dep3} = {'*to reduce*'}: Verb class: "*Limitation and subjugation*" [neutral polarity; negative reinforcement]. OF^{dep3} = {'*global hunger*'}: a(OF^{dep3}) = a('*hunger*') = [0,0,0,0,0,0,0,0,0.5 (Interest),0,0.7 (Sadness),0,0]. Rule of OF attitude reversal: a(Clause 3) = a(VF^{dep3} & OF^{dep3}) = reversal(a(OF^{dep3})) = [0,0,0,0,0,0,0,0,0,0.7 (Joy),0,0,0]. Attitude label and score: 'Joy': 0.7.

Table 6.5 (Continued)

Steps in @AM algorithm
Clause 3 represents the Object Formation of Clause 2: OF^{comp2} = Clause 3.
Analysis of object-relative Clause 8 (‘*that may harmfully influence the human health and environment*’):
SF^{rel8} = {‘*that*’}: a(SF^{rel8}) = neutral vector. VF^{rel8} = {‘*may harmfully influence*’}: confidence(‘*may*’) = 0.27 [central modal auxiliary]. a(‘*harmfully*’) = [0,0,0,0,0,0,0,0,0,0,0.9 (Sadness),0,0]. a(‘*influence*’) = neutral vector [no classification]. a(VF^{rel8}) = a(‘*harmfully*’) = [0,0,0,0,0,0,0,0,0,0,0.9 (Sadness),0,0]. OF^{rel8} = {‘*the human health and environment*’}: a(OF^{rel8}) = neutral vector. a(Clause 8) = a(SF^{rel8} & VF^{rel8} & OF^{rel8}) = a(VF^{rel8}) = [0,0,0,0,0,0,0,0,0,0,0.9 (Sadness),0,0]. Attitude label and score: ‘Sadness’: 0.9. Confidence level: 0.27. Clause 8 is related to the Object Formation of Clause 7.
2. After adding the Clause 3 to the Object Formation of Clause 2, and the Clause 8 to the Object Formation of Clause 7, the @AM analyses Clause 2, Clause 5, Clause 6, and Clause 7.
Analysis of complement Clause 2 (‘*GM foods could help to reduce global hunger*’):
SF^{comp2} = {‘*GM foods*’}: a(SF^{comp2}) = neutral vector. VF^{comp2} = {‘*could help*’}: confidence(‘*could*’) = 0.37 [central modal auxiliary]. Verb class (‘*help*’): “*Assistance*” [attitude-conveying verb]. a(VF^{comp2}) = a(‘*help*’) = [0,0,0.5 (POS app),0,0,0,0,0,0,0,0,0,0]. OF^{comp2} = {‘*to reduce global hunger*’}: a(OF^{comp2}) = a(Clause 3) = [0,0,0,0,0,0,0,0,0,0.7 (Joy),0,0,0]. As both VF^{comp2} and OF^{comp2} have positive polarity, the aggregation (fusion) rule is applied: a(Clause 2) = a(SF^{comp2} & VF^{comp2} & OF^{comp2}) = fusion(a(VF^{comp2}),a(OF^{comp2})) = [0,0,0.5 (POS app),0,0,0,0,0,0,0.7 (Joy),0,0,0]. Decision on attitude label: *Potential*(AFF) = 0. *Potential*(APP) = 1 due to high-level concepts of ‘*food*’ (WN: *FOOD*) and ‘*hunger*’ (WN: *STATE*). Attitude label and score: ‘POS app’: 0.5. Confidence level: 0.37.

Table 6.5 (Continued)

Steps in @AM algorithm
Complement Clause 2 represents the Object Formation of Clause 1: OF^{main1} = Clause 2.
Analysis of complement Clause 5 ('*genetically modified foods can pose serious threats to the consumers' health and Earth's ecosystem*'):
SF^{comp5} = {'*genetically modified foods*'}: a(SF^{comp5}) = neutral vector. VF^{comp5} = {'*can pose*'}: confidence('*can*') = 0.5 [central modal auxiliary]. a('*pose*') = neutral vector [no classification]. a(VF^{comp5}) = neutral vector. OF^{comp5} = {'*serious threats to the consumers' health and Earth's ecosystem*'}: a(OF^{comp5}) = a('*threat*') = [0,0,0,1.0 (NEG app),0,0,1.0 (Fear),0,0,0,0,0,0]. a(Clause 5) = a(SF^{comp5} & VF^{comp5} & OF^{comp5}) = a(OF^{comp5}) = [0,0,0,1.0 (NEG app),0,0,1.0 (Fear),0,0,0,0,0,0]. Decision on attitude label: *Potential*(AFF) = 0. *Potential*(APP) = 1 due to high-level concepts of '*food*' (WN: *FOOD*), '*threat*' (WN: *ENTITY*), '*health*' (WN: *STATE*), '*earth*' (WN: *NATURAL OBJECT*), and '*ecosystem*' (WN: *GROUP*). Attitude label and score: 'NEG app': 1.0. Confidence level: 0.5. Complement Clause 5 is related to the Object Formation of Clause 4.
Analysis of Clause 6 ('*novel and possibly hazardous genetic material disrupts the natural harmony*'):
SF^{cnt6} = {'*novel and possibly hazardous genetic material*'}: coeff('*possibly*') = 0.45 [adverb of affirmation]. a('*hazardous*') = [0,0,0,0,0,0,0.4 (Fear),0,0,0,0,0,0]. a(SF^{cnt6}) = a('*possibly hazardous*') = coeff('*possibly*') * a('*hazardous*') = 0.45 * [0,0,0,0,0,0,0.4 (Fear),0,0,0,0,0,0] = [0,0,0,0,0,0,0.18 (Fear),0,0,0,0,0,0]. VF^{cnt6} = {'*disrupts*'}: Verb class ('*disrupt*'): "*Negatively charged change of state*" [attitude-conveying verb]. a(VF^{cnt6}) = a('*disrupt*') = [0,0,0,0,0,0,0,0,0,0,0.2 (Sadness),0,0]. OF^{cnt6} = {'*the natural harmony*'}: a(OF^{cnt6}) = a('*harmony*') = [0.4 (POS jud),0,0.4 (POS app),0,0,0,0,0,0,0.4 (Joy),0,0,0]. a(Clause 6) = a(SF^{cnt6} & VF^{cnt6} & OF^{cnt6}) = intensification(reversal(a(OF^{cnt6}))) = 1.2 * reversal(a(OF^{cnt6})) = 1.2 * [0,0.4 (NEG jud),0,0.4 (NEG app),0,0,0,0,0,0,0.4 (Sadness),0,0] = [0,0.48 (NEG jud),0,0.48 (NEG app),0,0,0,0,0,0,0.48 (Sadness),0,0]. Decision on attitude label: *Potential*(AFF) = 0.

Table 6.5 (Continued)

Steps in @AM algorithm

Potential(JUD) = 0.

Potential(APP) = 1 due to high-level concepts of '*material*' (WN: *SUBSTANCE*) and '*harmony*' (WN: *STATE*).

Attitude label and score: 'NEG app': 0.48.

Clause 6 is related to the Object Formation of Clause 4.

Analysis of Clause 7 ('*causes the uncontrolled biological pollution that may harmfully influence the human health and environment*'):

VF^7 = {'*causes*'}:

a('*cause*') = neutral vector [no classification].

a(VF^7) = neutral vector.

OF^7 = {'*the uncontrolled biological pollution that may harmfully influence the human health and environment*'}:

a('*pollution*') = [0,0,0,0.475 (NEG app),0,0,0,0,0,0,0,0,0].

a(Clause 8) = a('*that may harmfully influence the human health and environment*') = [0,0,0,0,0,0,0,0,0,0,0.9 (Sadness),0,0].

confidence(Clause 8) = 0.27.

a(OF^7) = a('*the uncontrolled biological pollution*' & Clause 8) = fusion(a('*pollution*'),a(Clause 8)) = [0,0,0,0.475 (NEG app),0,0,0,0,0,0,0.9 (Sadness),0,0].

a(Clause 7) = a(VF^7 & OF^7) = a(OF^7) = [0,0,0,0.475 (NEG app),0,0,0,0,0,0,0.9 (Sadness),0,0].

Decision on attitude label:

Potential(AFF) = 0.

Potential(APP) = 1 due to high-level concepts of '*pollution*' (WN: *STATE*), '*health*' (WN: *STATE*), and '*environment*' (WN: *LOCATION*).

Attitude label and score: 'NEG app': 0.475.

Confidence level: 0.27.

Clause 7 is related to the Object Formation of Clause 4.

3. After adding the Clause 2 to the Object Formation of Clause 1, and the Clauses 5-7 to the Object Formation of Clause 4, the @AM analyses Clause 1 and Clause 4.

Analysis of Clause 1 ('*some believe that GM foods could help to reduce global hunger*'):

SF^{main1} = {'*some*'}:

a(SF^{main1}) = neutral vector.

VF^{main1} = {'*believe*'}:

Verb class: "*Communication indicator/reinforcement of attitude*" [neutral polarity; positive ground; positive reinforcement].

confidence('*believe*') = 0.7.

Table 6.5 (Continued)

Steps in @AM algorithm
OF^{main1} = {‘*that GM foods could help to reduce global hunger*’}: OF^{main1} = Clause 2. Attitude label and score of Clause 2: ‘POS app’: 0.5. $a(OF^{main1})$ = [0,0,0.5 (POS app),0,0,0,0,0,0,0,0,0,0]. confidence(OF^{main1}) = 0.37. As reinforcement of verb ‘*believe*’ is positive, the intensification rule is applied: a(Clause 1) = $a(SF^{main1}$ & VF^{main1} & $OF^{main1})$ = intensification($a(OF^{main1})$) = reinforcement_coeff * $a(OF^{main1})$ = 1.2 * [0,0,0.5 (POS app),0,0,0,0,0,0,0,0,0,0] = [0,0,0.6 (POS app),0,0,0,0,0,0,0,0,0,0]. Attitude label and score: ‘POS app’: 0.6. Modal confidence level: 0.37. Verb confidence level: 0.7.
Analysis of Clause 4 (‘*a number of studies over the past decade have revealed that genetically modified foods can pose serious threats to the consumers’ health and Earth’s ecosystem, because novel and possibly hazardous genetic material disrupts the natural harmony and causes the uncontrolled biological pollution that may harmfully influence the human health and environment*’):
SF^{indep4} = {‘*a number of studies over the past decade*’}: $a(SF^{indep4})$ = neutral vector. VF^{indep4} = {‘*have revealed*’}: Verb class (‘*reveal*’): “*Communication indicator/reinforcement of attitude*” [neutral polarity; positive ground; positive reinforcement]. confidence(‘*reveal*’) = 1.0. OF^{indep4} = {‘*that genetically modified foods can pose serious threats to the consumers’ health and Earth’s ecosystem, because novel and possibly hazardous genetic material disrupts the natural harmony and causes the uncontrolled biological pollution that may harmfully influence the human health and environment*’}: OF^{indep4} = Clause 5 & Clause 6 & Clause 7. Attitude label and score of Clause 5: ‘NEG app’: 1.0. Attitude label and score of Clause 6: ‘NEG app’: 0.48. Attitude label and score of Clause 7: ‘NEG app’: 0.475. Modal confidence level of Clause 5: 0.5. Modal confidence level of Clause 7: 0.27. The aggregation (fusion) rule is applied: $a(OF^{indep4})$ = fusion(a(Clause 5),a(Clause 6), a(Clause 7)) = [0,0,0,1.0 (NEG app),0,0,0,0,0,0,0,0,0]. As reinforcement of verb ‘*reveal*’ is positive, the intensification rule is applied: a(Clause 4) = $a(SF^{indep4}$ & VF^{indep4} & $OF^{indep4})$ = intensification($a(OF^{indep4})$) = reinforcement_coeff * $a(OF^{indep4})$ = 1.2 * [0,0,0,1.0 (NEG app),0,0,0,0,0,0,0,0,0] = [0,0,0,1.2 (NEG app),0,0,0,0,0,0,0,0,0].

Table 6.5 (Continued)

Steps in @AM algorithm
Attitude label and score: 'NEG app': 1.0 [maximum value of attitude score]. Modal confidence level: 0.5 [maximum value between modal confidence levels of Clause 5 and Clause 7]. Verb confidence level: 1.0.
4. Finally, after processing the Clause 1 and Clause 4, which are independent clauses of a compound sentence with the coordinate connector '*but*', @AM analyses the overall attitude of the sentence.
Sentence = {'Clause 1, *but* Clause 4'}: As the clauses are joined by coordinate connector '*but*', the domination rule is applied: a(Sentence) = a(Clause 4). Attitude label and score: 'NEG app': 1.0. Modal confidence level: 0.5. Verb confidence level: 1.0.

Chapter 7

Evaluation of the @AM Algorithm

In this Chapter we describe the evaluation of the performance of our algorithm for fine-grained recognition of affect, judgment, and appreciation conveyed in text. The evaluation is based on a set of experiments on the data sets created in different domains: personal stories about life experiences, fairy tales, and news headlines.

7.1 Experiment with Our Collection of Sentences from the Experience Project

7.1.1 Data Set Description

The first experiment was conducted on the set of sentences extracted from personal stories about life experiences that were anonymously published on the social networking website *Experience Project* (*www.experienceproject.com*). This website represents an interactive platform that allows people to share personal experiences, thoughts, opinions, feelings, passions, and confessions through the network of personal stories. With over 6.5 million experiences accumulated (as of November 2010), *Experience Project* is a perfect source for researchers interested in studying different types of attitude expressed through text.

For our experiment we extracted 1000 sentences from various stories grouped by topics within 13 different categories, such as 'Arts and entertainment', 'Current events', 'Education', 'Family and friends', 'Health and wellness', 'Relationships and romance' and others, on the *Experience Project*.

The sentences were collected from 358 distinct topic groups, such as ‘I still remember September 11’, ‘I am intelligent but airheaded’, ‘I think bullfighting is cruel’, ‘I quit smoking’, ‘I am a fashion victim’, ‘I was adopted’ and others. The distribution of sentences across the categories, the samples of topic groups and the example sentences are given in Table 7.1.

Table 7.1 The distribution of sentences across the Experience Project categories

Category	# of sent.	Sample of topic groups	Example sentence
Arts and entertainment	60	I love films that make me think about life. I am prone to inappropriate laughter.	*The slow pace of the film only contributes to creating this special atmosphere.* *I laugh to relieve myself from extremely stressful situations.*
Culture and religion	60	I am African American. I am Indonesian.	*Too many children die cause of these gangs.* *Indonesia is a very beautiful country, it’s just like a box of candy with so many flavor inside.*
Current events	60	I am trying to be more green. I still remember September 11.	*If we aren’t careful, we’ll end up starting wars for water instead of wars for oil.* *The country lost hundreds of heroes that day, some of them firefighters, some of them just plain folks.*
Education	60	I am intelligent but airheaded. I want to improve my mind.	*As we get busier in life, it gets harder to keep track of everything.* *I feel that it’s a wasted day when I don’t learn something new.*
Family and friends	60	I was adopted. I love my kids.	*Searching for my birth father is a bit of a hopeless case.* *It’s not always easy raising children, but the happiness u* [you] *get from them is priceless.*
Food and drink	60	I hate when people chew loudly. I love coffee.	*Loud chewing, especially bubble gum, is very annoying.* *While it is convenient, I think how unfriendly those little cups are for the environment.*
Health and wellness	60	I quit smoking.	*I have been smoke free for just over a year and it’s nice to be able to breathe again.*
Jobs and personal finance	60	I am unemployed.	*My experience of working in the UK is that companies very rarely put any money into staff training.*
Lifestyle and style	60	I am a fashion victim.	*It’s torture trying to find jeans and slacks that fit right.*
Pets and animals	60	I think bullfighting is cruel.	*I think it’s senseless to put an animal through that kind of pain for no good reason.*
Recreation and sports	60	I love camping.	*There were wild horses walking around everywhere and deer that were so used to humans that they would eat out of your hands.*
Relationships and romance	60	I have a crush.	*My best friend at the time started fooling around with her which started my depression.*
Other	280	I have trouble dealing with criticism.	*My whole enthusiasm and excitement disappear like a bubble touching a hot needle.*

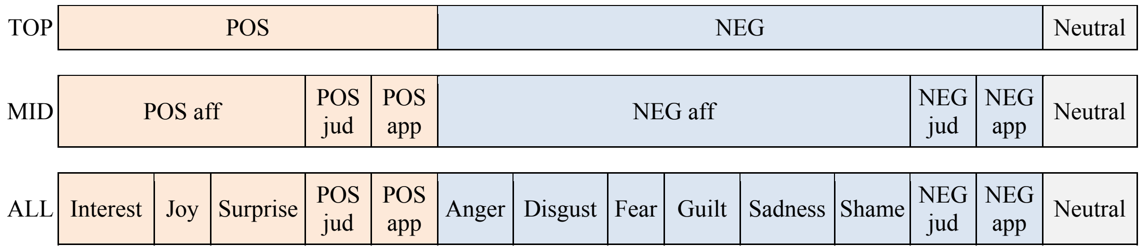

Figure 7.1 Hierarchy of attitude labels

We considered three hierarchical levels of attitude labels in our experiment (see Figure 7.1). Three independent annotators labelled the sentences with one of 14 categories from ALL level and a corresponding score (the strength or intensity value). These annotations were further interpreted using labels from MID and TOP levels. Fleiss' Kappa coefficient was used as a measure of reliability of human raters' annotations (Table 7.2). The agreement coefficient on 1000 sentences was 0.53 on ALL level, 0.57 on MID level, and 0.73 on TOP level.

Table 7.2 Inter-rater agreement on the data set from the Experience Project

Level	Complete data set		Gold standards (at least two annotators agreed)	
	Number of sentences	Fleiss Kappa	Number of sentences	Fleiss Kappa
ALL	1000	0.53	868	0.62
MID	1000	0.57	925	0.63
TOP	1000	0.73	997	0.74

Only those sentences, on which at least two out of three human raters completely agreed, were included in the gold standard for our experiment. Three gold standards were created according to the hierarchy of attitude labels (Table 7.2). Fleiss' Kappa coefficients are 0.62, 0.63, and 0.74 on ALL, MID, and TOP levels, correspondingly. Table 7.3 shows the distributions of labels in the gold standards.

Table 7.3 The distributions of labels across gold standard sentences

ALL level		**MID level**	
Label	**Number**	**Label**	**Number**
Anger	45	POS aff	233
Disgust	21	NEG aff	332
Fear	54	POS jud	66
Guilt	22	NEG jud	78
Interest	84	POS app	100
Joy	95	NEG app	29
Sadness	133	Neutral	87
Shame	18	**total**	**925**
Surprise	36		
POS jud	66	**TOP level**	
NEG jud	78	**Label**	**Number**
POS app	100	POS	437
NEG app	29	NEG	473
Neutral	87	Neutral	87
total	**868**	**total**	**997**

7.1.2 Results

The results of a simple method selecting the attitude label with the maximum intensity from the annotations of sentence tokens found in the AttitudeFul database were considered as the baseline. After processing each sentence from the data set by the baseline method and our @AM system, we measured averaged accuracy, precision, recall, and F-score for each label in ALL, MID, and TOP levels. The results are shown in Table 7.4.

As seen from the obtained results, our algorithm performed with high accuracy significantly surpassing the baselines in all levels of attitude hierarchy, thus demonstrating the contribution of the sentence parsing and our hand-crafted rules to the reliable recognition of affect, judgment, and appreciation from text. Two-tailed t-tests with significance level of 0.05 showed that the differences in accuracy between the baseline method and our @AM system are statistically significant ($p < 0.001$) in fine-grained as well as coarse-grained classifications.

Table 7.4 Results of the evaluation of performance of the baseline method and our @AM

Level	Label	Baseline method				@AM			
		Accuracy	Precision	Recall	F-score	Accuracy	Precision	Recall	F-score
ALL	Anger	0.437	0.742	0.511	0.605	**0.621***	0.818	0.600	0.692
	Disgust		0.600	0.857	0.706		0.818	0.857	0.837
	Fear		0.727	0.741	0.734		0.768	0.796	0.782
	Guilt		0.667	0.364	0.471		0.833	0.455	0.588
	Interest		0.380	0.357	0.368		0.772	0.524	0.624
	Joy		0.266	0.579	0.364		0.439	0.905	0.591
	Sadness		0.454	0.632	0.528		0.528	0.917	0.670
	Shame		0.818	0.500	0.621		0.923	0.667	0.774
	Surprise		0.625	0.694	0.658		0.750	0.833	0.789
	POS jud		0.429	0.227	0.297		0.824	0.424	0.560
	NEG jud		0.524	0.141	0.222		0.889	0.410	0.561
	POS app		0.349	0.150	0.210		0.755	0.400	0.523
	NEG app		0.250	0.138	0.178		0.529	0.310	0.391
	Neutral		0.408	0.483	0.442		0.559	0.437	0.490
MID	POS aff	0.524	0.464	0.695	0.557	**0.709***	0.668	0.888	0.762
	NEG aff		0.692	0.711	0.701		0.765	0.910	0.831
	POS jud		0.405	0.227	0.291		0.800	0.424	0.554
	NEG jud		0.458	0.141	0.216		0.842	0.410	0.552
	POS app		0.333	0.150	0.207		0.741	0.400	0.519
	NEG app		0.222	0.138	0.170		0.474	0.310	0.375
	Neutral		0.378	0.483	0.424		0.514	0.437	0.472
TOP	POS	0.732	0.745	0.796	0.770	**0.879***	0.918	0.920	0.919
	NEG		0.831	0.719	0.771		0.912	0.922	0.917
	Neutral		0.347	0.483	0.404		0.469	0.437	0.452

* Significant difference comparing with the baseline method, $p < 0.001$.

In the case of fine-grained attitude recognition (ALL level), the highest precision was obtained for ‘Shame’ (0.923) and ‘NEG jud’ (0.889), while the highest recall was received for ‘Sadness’ (0.917) and ‘Joy’ (0.905) emotions at the cost of low precision (0.528 and 0.439, correspondingly). The algorithm performed with the worst results in recognition of ‘NEG app’ and ‘Neutral’.

The analysis of the confusion matrix for the ALL level (Table 7.5) revealed the following top confusions of our system:

(1) ‘Anger’, ‘Fear’, ‘Guilt’, ‘Shame’, ‘NEG jud’, ‘NEG app’ and ‘Neutral’ were predominantly incorrectly predicted as ‘Sadness’ (for ex., @AM resulted in ‘Sadness’ emotion for the

sentence '*I know we have several months left before the election, but I am already sick and tired of seeing the ads on TV*', while human annotations were 'Anger' / 'Anger' / 'Disgust').

(2) 'Interest', 'POS jud' and 'POS app' were mostly confused with 'Joy' by our algorithm (e.g., @AM classified the sentence '*It's one of those life changing artifacts that we must have in order to have happier, healthier lives*' as 'Joy'(-ful), while human annotations were 'POS app' / 'POS app' / 'Interest').

Table 7.5 Data from a confusion matrix for ALL level

Actual label	**Incorrectly predicted labels (%), in descending order**
Anger	Sadness (28.9), Joy (4.4), Neutral (4.4), NEG app (2.2)
Disgust	Anger (4.8), Sadness (4.8), NEG jud (4.8)
Fear	Sadness (13.0), Joy (5.6), POS app (1.9)
Guilt	Sadness (50.0), Anger (4.5)
Interest	Joy (33.3), Neutral (7.1), Sadness (3.6), POS app (2.4), Fear (1.2)
Joy	Interest (3.2), POS app (3.2), Sadness (1.1), Surprise (1.1), Neutral (1.1)
Sadness	Neutral (3.8), Joy (1.5), Anger (0.8), Fear (0.8), Guilt (0.8), NEG app (0.8)
Shame	Sadness (16.7), Fear (5.6), Guilt (5.6), NEG jud (5.6)
Surprise	Fear (5.6), Neutral (5.6), Joy (2.8), POS jud (2.8)
POS jud	Joy (37.9), POS app (9.1), Interest (4.5), Sadness (1.5), Surprise (1.5), NEG jud (1.5), Neutral (1.5)
NEG jud	Sadness (37.2), Anger (3.8), Disgust (3.8), Neutral (3.8)
POS app	Joy (37.0), Neutral (9.0), Surprise (7.0), Interest (3.0), POS jud (3.0), Sadness (1.0)
NEG app	Sadness (44.8), Fear (13.8), Disgust (3.4), Surprise (3.4), Neutral (3.4)
Neutral	Sadness (29.9), Joy (13.8), Interest (3.4), Fear (2.3), POS jud (2.3), NEG app (2.3), NEG jud (1.1), POS app (1.1)

Our system achieved high precision for all categories on the MID level (Table 7.4), with the exception of 'NEG app' and 'Neutral', although high recall was obtained only in the case of categories related to affect ('POS aff', 'NEG aff'). These results indicate that affect sensing is easier than recognition of judgment or appreciation from text. TOP level results (Table 7.4) show that our algorithm classifies sentences that convey positive or negative sentiment with high accuracy (92 percent and 91 percent, correspondingly). On the other hand, neutral sentences still pose a challenge.

The importance of content words of different parts of speech in textual attitude analysis was evaluated in the next experiment, when first the sentences were annotated by the @AM algorithm taking into account attitude-conveying adjectives only, then by the @AM algorithms with cumulatively added attitude-conveying adverbs, nouns, and verbs. The averaged accuracies obtained in this experiment for each level of attitude hierarchy are given in Table 7.6.

Table 7.6 Results of the experiment with words of different parts of speech cumulatively added to the algorithm

Algorithm~	**Accuracy**		
	ALL	**MID**	**TOP**
@AM (ADJ)	0.325	0.357	0.491
@AM (ADJ, ADV)	0.347	0.376	0.516
@AM (ADJ, ADV, NOUN)	0.397*	0.452**	0.626**
@AM (ADJ, ADV, NOUN, VERB)	**0.621****	**0.709****	**0.879****

~ @AM stands for Attitude Analysis Model; ADJ and ADV refer to adjectives and adverbs, respectively.
* Significant difference comparing with the preceding method, $p < 0.05$.
** Significant difference comparing with the preceding method, $p < 0.001$.

As seen from Table 7.6, our algorithm for the attitude analysis benefits from the consideration of attitude-conveying words of all content parts of speech (adjectives, adverbs, nouns, and verbs). The statistically significant improvements in accuracy are observed on all classification levels after adding the nouns and verbs to the algorithm (two-tailed t-test with significance level of 0.05 was used to assess the significance of difference in accuracy between the preceding and the following algorithms). Although no statistically significant improvement in accuracy was obtained after adding the adverbs to the algorithm relying purely on adjectives, we think that the role of the attitude-conveying adverbs in attitude analysis should not be underestimated.

With the aim to evaluate our @AM algorithm with selectively removed functionality components, we conducted a *functional ablation* experiment. We focused on such @AM functionalities as:

(1) Polarity reversal by negations, modifiers, and functional words.

(2) Neutralization due to condition, preposition, and connector '*but*'.

(3) Adjustment of attitude labels based on analysis of pronouns, WordNet high-level concepts, and Stanford NER labels.

We compared @AM with all functionalities, @AM without additional functionalities, and three algorithms in which one specific functionality was ablated from the @AM algorithm. We believe that the specified three @AM configurations would show the enhancement that each functionality adds to complete @AM algorithm rather than what is missing when each is removed.

The results of the functional ablation experiment in terms of the accuracy are shown in Table 7.7. The most noticeable degradation effect was observed when all additional functionalities were ablated from the complete @AM algorithm. T-test with significance level of 0.05 showed that the accuracy of the algorithm with this configuration was significantly lower than the accuracy of @AM with all additional functionalities on the MID ($p < 0.05$) and TOP ($p < 0.01$) levels. The obtained results also revealed that the functionality related to polarity reversal due to negations, modifiers, and 'functional' words significantly contributes to the performance of complete @AM algorithm ($p < 0.05$) on the TOP level (polarity-based classification). On both ALL and MID levels our @AM mostly benefits from such functionality as adjustment of attitude labels based on analysis of pronouns, WordNet high-level concepts, and Stanford NER labels.

Table 7.7 Averaged accuracy obtained in the *functional ablation* experiment

Algorithm	**Accuracy**		
	ALL	**MID**	**TOP**
@AM with all functionalities	**0.621**	**0.709**	**0.879**
@AM w/o all additional functionalities	0.581	0.665*	0.830**
@AM w/o polarity reversal by negations, modifiers, and 'functional' words	0.609	0.692	0.843*
@AM w/o neutralization due to condition, preposition, and connector '*but*'	0.614	0.708	0.875
@AM w/o adjustment of attitude labels based on analysis of pronouns, WordNet high-level concepts, and Stanford NER labels	0.588	0.685	0.878

* Significant difference comparing with @AM with all functionalities, $p < 0.05$.
** Significant difference comparing with @AM with all functionalities, $p < 0.01$.

The analysis of errors in assigning the fine-grained attitude labels on ALL level (Table 7.8) revealed that in 32.3 percent of errors the @AM system confused the similar attitude states (e.g., 'Joy' and 'Interest', 'NEG jud' and 'Sadness', etc.), and in 30.1 percent of errors the system resulted in agreement with only one human annotator. We can consider these types of errors as non-strict ones, as even humans do not always agree with each other in assigning the attitude labels due to subjective interpretation of the attitude expressions.

Table 7.8 Distribution of errors of @AM in assigning the attitude labels (fine-grained or ALL classification level)

Error type	**Error**		**Sample sentence (gold standard — @AM label)**
	#	**%**	
Confused similar attitude states	106	32.3	*When I first saw that you could have a chance to swim with dolphins I was very excited*. (Joy — Interest) *Without love, we become detached and selfish*. (NEG jud — sadness)
Agreement with one annotator	99	30.1	*Basically, my water bottle is like an ever-faithful dog*. (POS app / Neutral / POS app — Neutral)
Common sense	63	19.1	*For me every minute on my horse is alike an hour in heaven!* (Joy — Neutral) *All through my life I've felt like I'm second fiddle*. (Sadness — Neutral
Correct label in the final vector, but not dominant	15	4.6	*My former boss was not good at communication and used manipulation and fear to motivate*. (NEG jud — Fear)
Sense ambiguity	12	3.7	*The planet has so many incredible things to offer*. (POS app — Surprise)
Lexicon	7	2.1	*Holding hands is such a simple and yet meaningful act between two people*. (POS app — Neutral)
Negation	6	1.8	*I couldn't let myself reach the depression level that I had reached five weeks ago*. (Sadness — Joy) *I can't even begin to describe how important they are to me, how a good day with them can erase any bad days I've had*. (POS jud — NEG jud) *Everyone who knows me knows that I just don't sleep well*. (Neutral — Sadness)
Connector '*but*'	5	1.5	*Sometimes I still struggle with depression but I've learned how to be successful*. (Sadness — Joy)
Condition	3	0.9	*I know that even though I panic at the thought of going to school, once I'm there it's not so bad*. (Fear — POS app)
Incorrect opposite emotion due to reversal	3	0.9	*And now, although I don't do bodily harm, I'm definitely not fun to be around if I'm woken up!* (Anger — Sadness)
Parser	2	0.6	*It should have been the greatest trip of my entire life, but in fact it was a total nightmare*. (NEG app — POS app)
Verb rule	2	0.6	*Zebra, Oreo, halfbreed, these names and more seemed to be my first name instead of my given – Mike – and over time, they ceased to bother me*. (Anger — Joy)
No analysis of modality	2	0.6	*Many people take digital pictures and need to enhance the photos using simple tools like color balance correction and red-eye reduction*. (Neutral — POS app)
Preposition '*without*'	1	0.3	*She is not without her faults*. (NEG jud — Neutral)
No neutralization of '*instead of*'	1	0.3	*Instead of doing a few things spectacular, I am doing many things mediocre*. (Guilt — Interest)
Other	2	0.6	
Total	329	100	

The attitude sensing system requires common sense or additional context to deal with sentences like '*All through my life I've felt like I'm second fiddle.*' (gold standard: 'Sadness'; @AM: 'Neutral') or '*For me every minute on my horse is alike an hour in heaven!*' (gold standard: 'Joy'; @AM: 'Neutral'), which caused 19.1 percent of the errors.

In 4.6 percent of misclassification cases, the correct attitude label was in the final vector, but it was not dominant, taking into account its intensity. The ambiguity of word senses was responsible for about 4 percent of the errors, highlighting the importance and necessity of word sense disambiguation. About 2 percent of the errors were due to the lack of attitude-conveying entries in the database.

For the correct interpretation of the expressed attitude, diversified approaches should be developed for the analysis of sentences containing negations (1.8 percent of errors), as simple negation (e.g., '*It is not fun*'), negation with modal operators (e.g., '*couldn't*'), or negation used in an idiomatic expression (e.g., '*can't even begin*' in '*I can't even begin to describe how important they are to me, how a good day with them can erase any bad days I've had.*') have a different influence on the contextual attitude (e.g., reversal of polarity, neutralization, or even no impact).

The rule on connector '*but*' has produced 1.5 percent of misclassifications, including the failure on the following sentence: '*Sometimes I still struggle with depression but I've learned how to be successful.*' (gold standard: 'Sadness'; @AM: 'Joy'). The errors resulting from the neutralization of conditional statements comprise about 1 percent of errors. As finding the pairs of opposite emotions is problematic, our @AM resulted in about 1 percent of errors due to application of the reversal rule (e.g., *And now, although I don't do bodily harm, I'm definitely not fun to be around if I'm woken up!* (gold standard: 'Anger'; @AM: 'Sadness')).

The failures carrying grammatical character and caused by the parser include 0.6 percent of the errors; the same number of errors was due to (1) the verb rule and (2) no analysis of the modal expressions. We also found out that the preposition '*without*' does not always neutralize the attitude of related words, and in some cases the rule of polarity reversal is applicable, like in the sentence '*She is not without her faults.*' (gold standard: 'NEG jud'; @AM: 'Neutral'). The analysis of errors also revealed that it is necessary to neutralize the attitude of a statement starting from the preposition '*instead of*'.

The evaluation of the attitude scores (intensities) was based on the Pearson measure of correlation between the intensities automatically assigned by the @AM and the intensities manually assigned by three human annotators. We considered only those sentences, where our @AM algorithm agreed on the attitude label, except 'Neutral' label, with each annotator individually.

Table 7.9 contains the Pearson measures of correlation (1) between intensities given by annotators pairwise; and (2) between intensities provided by our method and intensities given by each annotator individually. The Pearson's correlation coefficients calculated between attitude intensities given by human annotators pairwise indicate mainly strong positive relationships. The character of the measured correlations between intensities provided by each annotator and our @AM system varied depending on the annotator (moderate positive relationship in case of @AM versus Annotator 1 (An1); from moderate to weak positive relationship in case of @AM versus Annotator 2 (An2); and negligible relationship in case of @AM versus Annotator 3 (An3)). These results show that our @AM system achieved satisfactory results in assigning the attitude strength to the attitude-conveying sentences.

Table 7.9 The results of evaluation of the @AM in assigning the attitude intensities

Level	Pearson's correlation coefficient (r)					
	Between annotators			Between @AM and each annotator		
	An1-An2	An1-An3	An2-An3	@AM-An1	@AM-An2	@AM-An3
ALL	0.58***	0.46***	0.49***	0.34**	0.32**	0.11^
MID	0.56***	0.45***	0.47***	0.34**	0.28*	0.10^
TOP	0.51***	0.39**	0.41***	0.31**	0.23*	0.08^

^ no or negligible relationship, r = 0.01 to 0.19
* weak positive relationship, r = 0.20 to 0.29
** moderate positive relationship, r = 0.30 to 0.39
*** strong positive relationship, r = 0.40 to 0.69

7.2 Experiment with Sentences from Fairy Tales

7.2.1 Data Set Description

In our next experiment, we wanted to compare the performance of the Attitude Analysis Model with Alm's (2008) system that reportedly outperformed Liu's (2003) system on affect sensing in sentences from fairy tales. Following the same evaluation scenario as Alm (2008), we considered three hierarchical levels of affect labels in our experiment (see Figure 7.2).

Level							
TOP (2 labels)		Emotional (EM)					Neutral
MID (3 labels)		Positive emotions (POS)		Negative emotions (NEG)			Neutral
ALL (6 labels)	Alm's	Happy	Surprised	Anger-Disgusted	Fearful	Sad	Neutral
	Our	Joy	Surprise	Anger-Disgust	Fear	Sadness-Guilt-Shame	Neutral*

* Neutral, including 'Interest', 'POS jud', 'NEG jud', 'POS app', and 'NEG app'.

Figure 7.2 Affect hierarchy and set of labels

As @AM is capable of recognizing nine emotions, as well as positive and negative judgments and appreciations, in order to compare the results of our method with Alm's (2008) method, we reduced the number of labels by:

(1) Merging the following labels: 'Anger'-'Disgust', 'Sadness'-'Guilt'-'Shame'.

(2) Considering the sentences annotated by @AM using 'Interest', 'POS jud', 'NEG jud', 'POS app', and 'NEG app' labels as 'Neutral'.

We ran the experiment on the subset of 1207 sentences marked by high agreement (indicating that affect labels assigned by four human annotators for the sentence were identical), and a subset of sentences with neutral label (affect data from fairy tales were downloaded from *http://lrc.cornell.edu/swedish/dataset/affectdata/index.html*). As we did not have the subsets of neutral sentences used by Alm in her experiments, we randomly extracted them from the whole corpus of sentences that were labelled by human annotators as neutral (differences in data sets, however, might add some incomparability to the results). The size of a sample of neutral sentences varied at each hierarchical level and was determined based on the number of affective labels at each level by (12) (taken from (Alm 2008)):

$$\left\lceil \frac{|HA|}{|Ai|-1} \right\rceil, \tag{12}$$

where HA is the set of high agreement affect sentences in the whole corpus; Ai is the set of affect labels at a specific level i in the affect hierarchy.

7.2.2 Results

We compared the results of the @AM with Alm's (2008) LOOHAsnowtag method (supervised machine learning approach) and two baselines (see Table 7.10 partially taken from (Alm 2008)). The baselines are represented as follows:

(1) The ratio of neutrally labelled sentences (N-BL).

(2) The ratio of the most frequent affect label (Freq-BL).

Table 7.10 Accuracy across sentences from fairy tales in high agreement experiment (span of mean accuracy given for LOOHAsnowtag method)

Level	Data size	Baselines		Individual classification methods	
	Total number of sentences (number of neutral sentences)	N-BL	Freq-BL	LOOHAsnowtag	@AM
ALL	1448 (241)	17	31 ('Joy')	69-70	63.9
MID	1810 (603)	33	40 (NEG)	69-73	72.8
TOP	2414 (1207)	50	50 (any)	79	80.8

As seen from the obtained results (Table 7.10), @AM significantly outperformed both baselines. The observed increase in accuracy is inversely proportional to the number of possible labels at the hierarchical levels (the highest accuracy is at the TOP level with only two possible categories). On the fine-grained level our method was less accurate than the LOOHAsnowtag method. @AM resulted in a similar accuracy as the LOOHAsnowtag method on the MID level. Our method outperformed the Alm's method on the TOP level (with about 2 percent gain in accuracy).

However, as was mentioned above, there are two factors that may have an impact on the comparability of our methods:

(1) The differences in the classification labels.

(2) The use of different data sets of neutral sentences.

7.3 Experiment with News Headlines

7.3.1 Data Set Description

This data set was created for the SemEval-2007 task on 'Affective Text' (Strapparava and Mihalcea 2007). The test data set consists of 1000 news headlines independently labelled by six annotators using scores of six emotions ('Anger', 'Disgust', 'Fear', 'Joy', 'Sadness', and 'Surprise') by means of web-based interface with six slide bars (this data set is available at the SemEval-2007 web site: *http://nlp.cs.swarthmore.edu/semeval*). The emotion score interval is [0, 100], where 0 means the emotion is missing from the given headline, and 100 represents the maximum emotional load. The annotators were instructed to select the appropriate emotional scores for each headline based on the presence of words or phrases with emotional content, as well as the overall feeling invoked by the text, ensuring thus annotations of cases where multiple emotions are involved.

In this gold standard, the distribution of headlines according to the number of non-zero scores of particular emotions assigned is as follows: one emotion – 1.1 percent of headlines; two emotions – 19.1 percent of headlines; three emotions – 17.7 percent of headlines; four emotions – 21.6 percent of headlines; five emotions – 27.2 percent of headlines; and six emotions – 13.3 percent of headlines. Therefore, 62.1 percent of all headlines were annotated by at least four emotions.

7.3.2 Results

In this experiment, we evaluated the performance of our system based on the fine-grained and coarse-grained evaluation metrics proposed in (Strapparava and Mihalcea 2008). Fine-grained evaluations were based on Pearson measure of correlation between the system scores and the gold standard scores, averaged over all the headlines in the data set.

In order to produce more or less comparable results, we:

(1) Considered the final attitude vector resulting from the @AM as the overall annotation for each headline.

(2) Reduced the number of our labels to six by mapping 'Interest', 'POS jud', and 'POS app' to 'Joy'; and 'Guilt', 'Shame', 'NEG jud', and 'NEG app' to 'Sadness'.

(3) Scaled intensities in interval [0.0, 1.0] to scores in interval [0, 100]. However, it is important to note here that the concepts and functions of our 'emotion intensities' and 'emotion scores' used in the gold standard differ significantly, as intensity shows the strength of emotion involved while score in the gold standard indicate how much particular emotion is involved in the headline.

For the coarse-grained evaluations, each emotion was mapped to a 0/1 classification (0 = [0, 50), 1 = [50, 100]), and precision, recall, and F-score were calculated for each emotion.

We compared the performance of our method with the systems participating at the SemEval-2007 task on 'Affective Text'. The results of our @AM and other systems (reported in (Strapparava and Mihalcea 2008)) are shown in Table 7.11, where:

(1) @AM is our Attitude Analysis Model.

(2) WN-A is 'WordNet-Affect presence' method, which computes the scores based on the frequencies of the direct affective words found in the headlines.

(3) LSA SW is 'LSA single word' method, which measures the similarity between the given text and each emotion, where an emotion is represented as the vector of the specific word denoting the emotion (e.g., 'Joy').

(4) LSA ES is 'LSA emotion synset' method, which uses the synonyms from the WordNet synsets in addition to the word denoting an emotion.

(5) LSA AEW is 'LSA all emotion words' method, which extends the previous set by adding the words from all the synsets labeled with a particular emotion in WordNet-Affect.

(6) NB BLOG is a Naïve Bayes classifier trained on the corpus of blog posts annotated by emotions (methods (2)-(6) were developed by Strapparava and Mihalcea (2008)).

(7) SWAT (Katz, Singleton and Wicentowski 2007) is a supervised system, which is based on unigram model.

(8) UA (Kozareva, Navarro, Vazquez and Montoyo 2007) is a system, which calculates emotion scores using Pointwise Mutual Information.

(9) UPAR7 (Chaumartin 2007) is a rule-based system, which is based on linguistic approach using SentiWordNet (Esuli and Sebastiani 2006) and WordNet-Affect (Strapparava and Valitutti 2004).

Table 7.11 The results of @AM compared to knowledge-based and corpus-based systems participating in the task 'Affective Text' at SemEval-2007

System	Emotion labels						Average result
	Anger	Disgust	Fear	Joy	Sadness	Surprise	
Fine-grained evaluation (Pearson's correlation coefficient)							
@AM	30.10*	11.98	41.19*	25.93*	**44.05**	6.16	26.57*
WN-A	12.08	-1.59	24.86	10.32	8.56	3.06	9.54
LSA SW	8.32	13.54	29.56	4.92	8.13	9.71	12.36
LSA ES	17.80	7.41	18.11	6.34	13.27	12.07	12.50
LSA AEW	5.77	8.25	10.28	7.00	10.71	12.35	9.06
NB BLOG	19.78	4.77	7.41	13.81	16.01	3.08	10.81
SWAT	24.51	**18.55**	32.52	**26.11**	38.98	11.82	25.41
UA	23.20	16.21*	23.15	2.35	12.28	7.75	14.15
UPAR7	**32.33**	12.85	**44.92**	22.49	40.98*	**16.71**	**28.38**
Coarse-grained evaluation (Precision / Recall / F-score)							
@AM	23.81	0.00	39.54	34.67	31.34	**27.27**	26.11
	23.81	0.00	36.96	23.01	40.39	7.14	21.88
	23.81	-	**38.20**	27.66	**35.29**	11.32	**27.26**
WN-A	**33.33**	0.00	**100.00**	50.00	33.33	13.04	**38.28**
	3.33	0.00	1.69	0.56	3.67	4.68	1.54
	6.06	-	3.33	1.10	6.61	6.90	4.00
LSA SW	6.28	**2.41**	12.93	17.81	13.13	6.73	9.88
	63.33	70.59	**96.61**	47.22	55.05	67.19	66.72
	11.43	**4.68**	22.80	25.88	21.20	12.23	16.37
LSA ES	7.29	1.53	12.44	19.37	14.35	7.23	9.20
	86.67	64.71	94.92	72.22	58.71	89.06	77.71
	13.45	3.00	22.00	30.55	23.06	13.38	13.38
LSA AEW	6.20	1.98	12.55	18.60	11.69	7.62	9.77
	88.33	**94.12**	86.44	**90.00**	**87.16**	**95.31**	**90.22**
	11.58	3.87	21.91	30.83	20.61	14.10	17.57
NB BLOG	13.68	0.00	16.67	22.71	20.87	8.33	12.04
	21.67	0.00	3.39	59.44	22.02	1.56	18.01
	16.77	-	5.63	**32.87**	21.43	2.63	13.22
SWAT	12.00	0.00	25.00	35.41	32.50	11.86	19.46
	5.00	0.00	14.40	9.44	11.92	10.93	8.61
	7.06	-	18.27	14.91	17.44	11.78	11.57
UA	12.74	0.00	16.23	40.00	25.00	13.70	17.94
	21.60	0.00	26.27	2.22	0.91	16.56	11.26
	16.03	-	20.06	4.21	1.76	**15.00**	9.51
UPAR7	16.67	0.00	33.33	**54.54**	**48.97**	12.12	27.60
	1.66	0.00	2.54	6.66	22.02	1.25	5.68
	3.02	-	4.72	11.87	30.38	2.27	8.71

Best results are given in bold.
* Result close to the best one.

In the fine-grained evaluation, our system achieved the best result in recognition of 'Sadness' emotion, while SWAT was more successful in case of 'Disgust' and 'Joy', and UPAR7 in case of 'Anger', 'Fear' and 'Surprise'. Our @AM performed with the results very close to the best ones in recognizing 'Anger', 'Fear', and 'Joy' emotions. In terms of averages of Pearson's correlation coefficients for all emotions, UPAR7 showed the best performance (28.38), followed by our @AM system (26.57) and SWAT (25.41).

These results indicate that our @AM system showed good results in detecting emotions in news headlines, in spite of the facts that it was not initially developed for this particular task and its results are not completely comparable to the gold standard. The important point is that the annotators of headlines assigned emotion scores based on words in the sentence (e.g., '*Tsunami fears ease after quake'* was annotated in the gold standard by Anger:0, Disgust:0, Fear:79, Joy:13, Sadness:13, Surprise:0, with predominant 'Fear' emotion); in contrast, our system analyzing the sentence in consecutive stages outputs unified attitude vector, which for this example sentence does not contain 'Fear' emotion involved in the object '*tsunami fears*', as positive vector of the verb '*ease*' dominates, therefore the final vector has Fear:0 and Joy:8 ('Joy' with intensity 0.08). Probably, to get more comparable results from the @AM, one have to sum attitude vectors from the word level annotation stage (see Section 6.2 for details) and ignore the phrase and sentence level analysis stages, so that the final vector includes all possible emotions. Such approach would perhaps result in a higher correlation coefficient between the @AM scores and gold standard; however, there would not be much intelligence.

In the coarse-grained evaluation, @AM ensured the best F-scores for 'Anger', 'Fear', and 'Sadness', as well as in case of the average results, while the highest F-scores for 'Disgust', 'Joy', and 'Surprise' were achieved by LSA SW, NB BLOG, and UA systems, correspondingly.

7.4 Summary

Our Attitude Analysis Model, which is based on novel compositional linguistic approach, is the only system that classifies individual sentences using fine-grained attitude labels (nine for different affective states, two for positive and negative judgment, and two for positive and negative appreciation), as against other methods that mainly focus on two polarity-based categories (positive and negative) or six basic emotions. The AttitudeFul lexicon that contains attitude-conveying terms, modifiers, contextual valence shifters, and modals, the analysis of syntactic and dependency relations between words in a sentence, the proposed *compositionality principle* (the rules of *polarity reversal, aggregation (fusion)*, *propagation*, *domination*, *neutralization*, and *intensification*, at various grammatical levels), the rules elaborated for semantically distinct verb classes, and a method considering the hierarchy of concepts based on WordNet and StanfordNER — all contribute to the robustness of algorithm for analysis of contextual attitude conveyed in written language. Our Attitude Analysis Model is capable of encoding the strength of the attitude and the level of confidence, with which the attitude is expressed, through numerical values in the interval [0.0, 1.0].

As there was no annotated data set (with fine-grained attitude annotations) available for evaluation of our @AM algorithm, we created our annotated data set of sentences extracted from personal stories about life experiences. In our experiment, we considered three hierarchical levels of attitude labels (14 categories for ALL level, 7 for MID level, and 3 for TOP level). In fine-grained as well as coarse-grained classifications, our @AM algorithm performed with high accuracy (62.1 percent on ALL level, 70.9 percent on MID level, and 87.9 percent on TOP level) significantly surpassing the baselines in all levels of attitude hierarchy. Our experiment aimed at evaluating the importance of different parts of speech in textual attitude analysis showed that the consideration of attitude-conveying words of all content parts of speech (adjectives, adverbs, nouns, and verbs) plays a crucial role in attitude analysis. The *functional ablation* experiment revealed that (1) the most noticeable degradation effect is observed when all additional functionalities are ablated from the complete @AM algorithm; (2) on both ALL and MID levels our @AM mostly benefits from such functionality as adjustment of attitude labels based on analysis of pronouns, WordNet high-level concepts, and Stanford NER labels; and (3) the functionality related to polarity reversal due to negations, modifiers, and 'functional' words significantly contributes to the performance of complete @AM algorithm on the TOP level. Our @AM system achieved satisfactory results in assigning the attitude strength to the attitude-conveying sentences.

In order to compare the performance of our @AM system with state-of-the-art methods on available annotated data sets of sentences from fairy tales (Alm 2008) and news headlines

(Strapparava and Mihalcea 2008), we had to reduce the number of attitude labels, as @AM is capable of recognizing nine emotions, as well as positive and negative judgments and appreciations.

On the sentences extracted from fairy tales (Alm 2008), @AM (1) significantly surpassed both baselines on all levels; (2) outperformed the Alm's method on the TOP level (with about 2 percent gain in accuracy: 80.8 versus 79 percent); (3) resulted in a similar accuracy as the Alm's method on the MID level (72.8 versus 69-73 percent); and (4) was less accurate than the Alm's method on the fine-grained level (63.9 versus 69-70 percent). However, as was already mentioned in Section 7.2, the differences in the number of classification labels and in data sets of neutral sentences could influence the comparability of our methods.

Comparing the performance of our @AM system with eight systems from related work on the task of recognition of emotions in news headlines (Strapparava and Mihalcea 2008), we found that, even though there were some inconsistencies in comparing the results with gold standard, our system resulted in high level of accuracy, outperforming other methods on several measures.

In fine-grained classification, @AM (1) outperformed all other methods (except UPAR7 rule-based linguistic method) in terms of average result; (2) completely outperformed (6 out of 6 labels) the method based on presence of affective keywords and the method employing Naïve Bayes classifier; (3) mainly surpassed (4 or 5 out of 6 labels) the 'LSA single word', 'LSA emotion synset' and 'LSA all emotion words' methods, and the system calculating emotion scores using Pointwise Mutual Information; (4) performed equally well (3 out of 6 labels) as the supervised system based on unigram model; and (5) was less accurate (in case of 4 out of 6 labels) than UPAR7 rule-based linguistic method in terms of results for distinct emotion labels.

In coarse-grained classification, our method (1) significantly outperformed all other methods in terms of average result of F-score; (2) completely or mainly surpassed the method based on presence of affective keywords, UPAR7 rule-based linguistic method, the 'LSA single word' method, the method employing Naïve Bayes classifier, the supervised system based on unigram model, and the system calculating emotion scores using Pointwise Mutual Information in terms of F-score results for distinct emotion labels; and (3) was more accurate (in case of 3 out of 6 labels) than 'LSA emotion synset' and 'LSA all emotion words' methods in terms of F-score results.

Part III

Applications

Chapter 8

Applications of Affect Analysis Model and Attitude Analysis Model

Using the Affect Analysis Model and the Attitude Analysis Model, several applications have been developed. This Chapter contains the summary of the following applications: AffectIM, EmoHeart, iFeel_IM!, and web-based @AM interface.

8.1 AffectIM: Instant Messaging Application Integrated with the Affect Analysis Model

8.1.1 Motivation

Social interaction among people is an essential part of every society, and strong foundation for the development and self-actualization of a person, as well as for the establishment of genuine interpersonal relationships and communities. Nowadays, media for remote online communications, which provide new opportunities for social contact, grow rapidly, engage people, and gain great popularity among them. The online world of computer-mediated communication is an environment where people can virtually remain in touch with their relatives and friends to exchange experiences, share opinions and feelings, and satisfy their social need of interpersonal communication. The main motivations for 'residents' of chat rooms or virtual environments to connect to these media are seeking conversation, experimenting with a new communication media, and initiating relationships with other people. A study conducted by Peris et al. (2002) revealed that 'relationships developed online are healthy' and considered by people 'as real as face-to-face relationships'. Authors point to the fact that online communications may stimulate rather than inhibit social relations. Findings described in (Hu et al. 2004) indicate that there is a positive relationship between the amount of online media use and verbal, affective, and social intimacy, and that frequent online conversation actually encourages the desire to meet face-to-face, thus reinforcing personal interaction. To emphasize the realism and significance of social exchanges in such environments, Chayko (2002) proposed to use the term 'sociomental' rather than 'virtual'.

To establish a social and friendly atmosphere, people should be able to express emotions. The richness of emotional communication greatly benefits from the expressiveness of verbal (spoken words, prosody) and nonverbal cues (gaze, face, gestures, body pose) that enable auditory and visual channels of communication (Planalp 1999). All types of expressive means potentially carry communicative power and promote better understanding (Allwood 2002). The emotional significance of an utterance is accompanied, complemented and modified by vocal and visual cues.

In everyday life we communicate with each other through multiple informative channels. People in online environments tend to interact in a social way too. However, computer-mediated communication often lacks signals of face-to-face communication such as spoken language, intonation, gaze, facial expressions, gestures, and body language. Trends show that people often try to enrich their interaction online by introducing affective symbolic conventions or emphases into text (emoticons, capital letters etc.) (Reid 1991; Walther and D'Addario 2001; Hu et al. 2004; Yigit 2005; Derks 2007), by colouring emotional messages, or by manually controlling the expressiveness of avatars (graphical representations of users) in order to supplement the lack of paralinguistic cues. Despite the playful nature of these conventions, the expressions of emotion conveyed are, according to Reid (1991), 'not in any way thought to be shallow or ephemeral'. The results of the study described in (Derks 2007) imply that emoticons can serve as nonverbal surrogates for visual cues in face-to-face communication and certainly have an impact on online message interpretation.

One of the first attempts to study effects of conveying emotional expressions through communication in computer-mediated environment was done by Rivera, Cooke, and Bauhs (1996). The results of their experiment indicated that subjects allowed to use emoticons were more satisfied with the system than those subjects having conversations without these symbolic emotional expressions. The user study of ExMS (Persson 2003), messaging system that allows its users to concatenate and annotate avatar animations, showed that the interplay between pure text and animation significantly improved the expressiveness of messages, and that the users felt pride of being identified with their embodied representation. Previous studies showed that (1) the interplay between pure text and animation significantly improves the expressiveness of messages (Persson 2003), and (2) the visualization of emotions by avatars in online communication media increases user's enjoyment (Olveres et al. 1998), involvement, and feeling of a social presence (Fabri, Elzouki, and Moore 2007). However, in synchronous communication media users can be exposed to the 'conversational stress' caused by the necessity to handle both avatar control and immediate response simultaneously (Persson 2003).

The motivation behind our approach is to enrich social interactivity and emotional expressiveness of real-time messaging, where a machine is used as a communication channel connecting people and transmitting human emotions. Here, a key issue is to provide the automation of multiple expressive channels, so that the user does not have to care about visual self-presentation, as it is the case in standard IM systems where an avatar's emotion is selected manually, but can focus on the textual content of the conversation. We believe that an IM system, which provides automatic emotion recognition from text and automatic visualization of emotions by avatars, is more efficient (in other words, easy-to-control) than the system with manual selection of an emotion, which can shatter the flow of conversation by the necessity to handle both avatar control and immediate response simultaneously.

To make the user's experience in online communication enjoyable, exciting and fun, we have designed a web-based Instant Messaging system, AffectIM, and endowed it with emotional intelligence by integrating the Affect Analysis Model (see Chapter 3 for details), which can detect nine emotions (and their intensity) from text messages automatically.

8.1.2 System Architecture and User Interface

The AffectIM system was developed as a web-based application running in the Internet browser, so that the user does not have to download and install the system on a local computer. The system architecture is schematically represented in Figure 8.1.

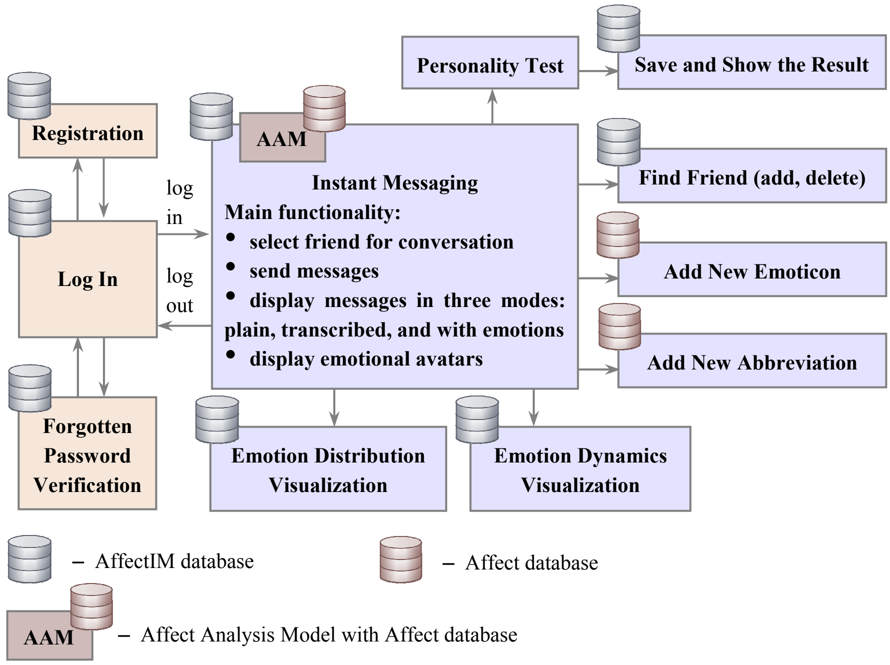

Figure 8.1 Architecture of the AffectIM system

In AffectIM, we propose to equip the user with an avatar, i.e. a graphical representation of the user in Instant Messaging. The avatar is endowed with the ability to express emotions and to exhibit social nonverbal behavior, whereby its behavior is based on textual affect sensing, and the interpretation of communicative functions conveyed by online conversations. We designed two avatars, one male and one female; so the graphical representative is automatically selected by the system according to the user's sex. The main window of the AffectIM system, showing an online conversation, is depicted in Figure 8.2. From the list of friends displayed in the left frame, the user selects the person (available online), whom he or she wishes to communicate with. The central frame allows the user to type and to send the messages. It displays the conversation flow in three modes: plain, transcribed, and with emotions. Further, it displays emotional avatars (own – to the left of conversation field, and friend's – to the right). Two buttons located under the avatar animation refer to the visualization of emotion distribution (either in a color bar or pie graph) and emotion dynamics (line graph). As the language of online communication is constantly evolving, AffectIM also provides the functionality to add new abbreviations, acronyms, and emoticons to the Affect database (see two buttons located to the left from the input text field in Figure 8.2).

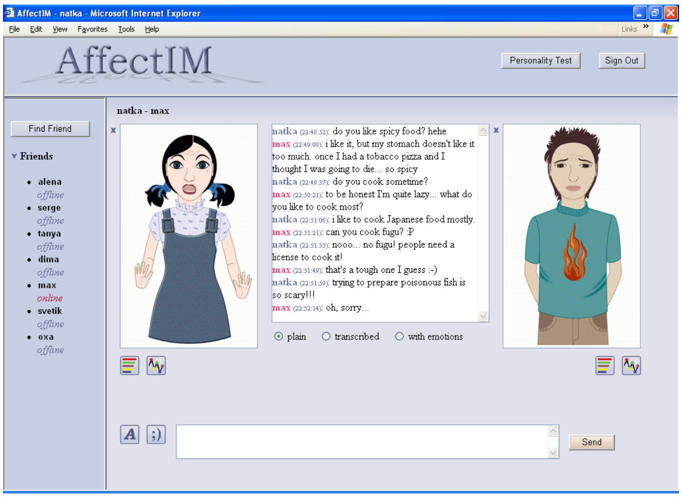

Figure 8.2 ‘Instant Messaging’ page of AffectIM

8.1.3 User Study of the AffectIM

We conducted the twenty-person user study aimed to evaluate ‘richness of experience’ and ‘affective intelligence’ of our AffectIM system (details are given in (Neviarouskaya et al. 2010c)). As the gold standard in communicating person’s emotions in the user study we considered an Instant Messaging interface allowing users to manually select emotions intended for visualization by avatar.

Our main hypotheses can be stated informally as follows:

The user experience with an Instant Messaging interface where avatar emotions have to be selected manually (Manual interface) is better than with an interface featuring expressive avatars, which convey user emotions based on the dominant emotion that is automatically recognized from text by our Affect Analysis Model (Automatic interface).

The user experiences with Manual and Automatic interfaces are better than with an interface with avatar emotions (quasi-) randomly displayed (Random interface).

The main hypotheses are tested by considering the following dimensions regarding users’ experience: (1) interactivity; (2) involvement (engagement); (3) sense of co-presence (sense of being together with another person in a shared computer-generated environment); (4) enjoyment; (5) affective intelligence; and (6) overall satisfaction. In addition to these main criteria, we asked participants to give us feedback on general questions (e.g. report on degree of necessity to look at an avatar of conversation partner; comments on helpfulness and frequency of usage of emotion selection function etc.).

Our main hypotheses do not address the question whether automatic emotion recognition is preferable over the technically simpler solution of having users select their intended emotion manually. While this issue is not at the core of our investigation, it is important to justify the need for automatic emotion recognition for Instant Messaging. Hence we formulate the following assertion. (We opt for 'assertion' rather than 'hypothesis', because we do not provide statistical evidence to support or refute the claim.)

> *Automatic emotion recognition in Instant Messaging is more efficient than manual selection of an emotion.*

Efficiency in Instant Messaging can refer to a variety of dimensions, including the success at conveying a user's emotion, the success of the overall conversation, the convenience of computer-mediated communication, and others, which were used for 'user experience'. Here, we assume a restrictive interpretation of 'efficiency', which only relates to the easiness of controlling the interface. In other words, a communication is considered the more efficient, the better the user can focus on writing IM text and the less other interface elements have to be operated. With this characterization of efficiency, it seems trivial that an automatic system outperforms a system that requires the user to manually select emotions. However, the very nature of IM has it that speed of communication is key to the popularity of this communication medium. In particular, an automatic interface releases the user from 'conversational stress' caused by the necessity to handle both avatar control and immediate response simultaneously (this issue was reported by Persson (2003)).

8.1.3.1 Subjects and Experimental Design

Twenty university students and staff (10 males, 10 females) took part in our study. All of them were computer literate and 19 persons had prior experience with computer based chat or Instant Messaging system. As IM application is language-based, the main requirement to the participants was to have an appropriate ability to converse in English. The distribution of subjects according to nationality aspect was as follows: 7 subjects from Japan; 3 from Germany; 2 – France; 2 – Iran; 2 –

China; 1 – Vietnam; 1 – Malaysia; 1 – Thailand; and 1 subject from Northern Africa. Subjects received 1,000 Yen for participation.

The experiment was designed as a within-subjects experiment. In particular, we compared the AffectIM interfaces based on three different settings (independent variable). For the user study, we prepared three versions of the system:

(1) Automatic (A-condition). In this interface, affect sensing from text is performed based on the developed Affect Analysis Model, and the recognized emotions are conveyed by the avatar expressions.

(2) Manual (M-condition). In this condition, no automatic emotion recognition from text is performed; however, users may select emotion (and its intensity) to be shown by avatars using 'select pop-up menus'. In the user study we consider the M-condition as a gold standard in communicating person's emotions (and evaluating 'affective intelligence' of the interfaces), as using this interface users explicitly indicate the emotion intended for visualization by avatar.

(3) Random (R-condition). Here, the avatars show a 'quasi-random' emotional reaction (explained in more detail below).

Regarding the R-condition, it is non-trivial to define an appropriate 'random' function. Specifically, given nine emotions and a 'neutral' state, we have to avoid a situation where some emotion is shown on each text entry, possibly adding too many emotional reactions to the IM conversation. On the other hand, we have to guarantee that all three conditions are in principle based on the same set of affective states (including neutral). In other words, we have to avoid a situation where users can experience some emotion in the A-condition or M-condition, which, by design, can never occur in the R-condition. In our study, we want to test the appropriateness of our automatically obtained emotions, rather than contrasting an 'emotional' with a 'non-emotional' condition. Hence we decided to inform the avatar reactions based on the following method, which we call 'quasi-random'. First, we process each sentence using the Affect Analysis Model, and then we apply two rules: (1) if the output is emotional, we run two functions that randomly select the emotion out of nine available emotions and its intensity, correspondingly; (2) for the case of 'neutral' output, we set the function that generates ‘neutral’ emotion with the probability of 60 percent or ‘random’ emotion with the probability of 40 percent.

In order to keep users’ attention on the conversation flow, we intentionally disabled additional functionality of the AffectIM interface (such as visualization of emotion distribution and dynamics, display of transcribed text or text annotated by emotions), and told subjects that we disabled these functions.

8.1.3.2 Materials, Apparatus and Procedure

In order to support active IM exchange during the experiment, we prepared three topics/scenarios (one for each condition):

1. **Traveling**. ‘*Please imagine that you and your partner are given 3000 Euros (about 500,000 Yen) in all, and asked to discuss, decide and agree about the country you will visit together. Who knows,… this might be your favorite country, or your partner will convince you of the benefit of choosing his or her country. When you agree on your travel destination, please imagine that you have already reached the destination. Now, you are allowed to do three things together: (1) to talk about, to agree on, and to visit only one place in this country; (2) to select and try only one national dish; (3) to buy only one souvenir in memory of this country.*’

2. **Exams**. ‘*We would like you to play a scenario, in which two students, that took the entrance examination to the University last week, just have learned their results and met each other. One of them had passed his/her examination successfully. However, the other student, whose dream was to enter and study in this University, failed his/her test. There are two roles: (1) student who passed the exam; (2) student who failed the exam. Both of them eager to talk and share their feelings about their situation. Please follow the role that experimenter will arbitrary indicate to you, and play this scenario using only your imagination (or real experience) and online communication.*’

3. **Food**. ‘*During the life, each person has been trying different kinds of cuisine. Everyone has his/her preferences. Please talk about the dishes, and try to discuss why the ingredients of particular dish harmonize and give savory taste, or, in contrast, mar the taste. What do you prefer: fast-food or home-made food? Maybe, interesting story of your own experience will come from your memory.*’

Each participant performed the IM session on a networked laptop (15 inch screen), which had the

AffectIM system installed, and a mouse pointer. As the current AffectIM system provides only one male and one female avatar, each pair of participants was composed by male and female subjects. In order to avoid awkward situations or silence during the experiment, all paired-up subjects knew each other, either as a fellow student or a friend.

Subjects were led to the experimental area individually, each guided by one experimenter. Before the IM session, all participants were given instructions and their AffectIM IDs and passwords. Each pair of participants was asked to have online conversations through three interfaces given in random order. We prepared the list of the orders of condition assignments (Automatic, Manual, and Random interfaces) for each pair (10 pairs of subjects) in advance. The order of three conditions for each pair was generated using a random function. The order of scenarios (‘Traveling’ => ‘Exams’ => ‘Food’) was fixed for all pairs, so that each scenario was associated with all three conditions throughout the user study. Each participant was asked to support the conversation flow continuously, and feel free to show emotions in his/her dialog. Participants were not informed about the type of interface they were using (A-condition or R-condition) during particular session, except in the case of the interface with manual selection of emotions (M-condition). After each interface condition, users filled in the corresponding page of the questionnaire and commented on their experience. After the participants completed the IM communication about the three topics and corresponding questionnaire, they were asked to answer some general questions about their experience with the IM system.

8.1.3.3 Results on Main Criteria and Responses to General Questions

The average duration of sessions on each interface was 10.1 minutes (minimum 8 and maximum 12.5 minutes), excluding the time needed to fill out the questionnaires. The 11 questions on the main criteria were answered based on a 7-item agreement Likert scale. It is worth looking at each measure in detail. As our study involved each subject being measured under each of three conditions (within-subjects design), in order to see if the apparent differences between interfaces are real or due to chance, we analyzed the data using two-factor ANOVA without replications (an extension of matched pair t-test) with significance level of 0.05. Additionally, we performed a post hoc power test (computation of achieved power, given significance level, sample size, and effect size) to measure the probability that a statistically significant difference would be found, when such a difference actually existed. The list of questions according to the analyzed criteria is given in Table 8.1. Mean scores, ANOVA results (p-values), effect sizes, and results of post hoc computation of achieved power are summarized in Table 8.2.

Table 8.1 The analyzed criteria and corresponding questionnaire items

Criteria	Questionnaire items
Interactivity	Q1: *The system was interactive*
Involvement (engagement)	Q2: *I felt it was important for my conversation partner that I responded after each his/her statement*
	Q3: *I was awaiting the replies of my conversation partner with true interest*
Sense of co-presence	Q4: *I felt if I were communicating with my conversation partner in the shared virtual space*
	Q5: *The system gave me the sense that the physical gap between us was narrowed*
Enjoyment	Q6: *I enjoyed the communication using this IM system*
Affective intelligence	Q7: *The system was successful at conveying my feelings*
	Q8: *The system was successful at conveying my partner's feelings*
	Q9: *The emotional behavior of the avatars was appropriate*
	Q10: *I understood the emotions of my communication partner*
Overall satisfaction	Q11: *I am satisfied with the experience of communicating via this system*

Table 8.2 Mean scores, ANOVA results, effect sizes, and results of post hoc computation of achieved power

Criteria and questionnaire items		**Mean scores**			**Results for pairs of conditions**								
					Random vs Manual			**Random vs Automatic**			**Manual vs Automatic**		
		R*	**M***	**A***	**p**	**Effect size**	**Power**	**p**	**Effect size**	**Power**	**p**	**Effect size**	**Power**
Interactivity	Q1	4.70	5.50	5.20	**0.068**	0.433	0.452	0.106	0.380	0.366	0.356	0.212	0.147
Involvement	Q2	5.50	5.35	5.15	0.691	0.090	0.067	0.376	0.203	0.138	0.551	0.136	0.089
	Q3	5.45	5.45	5.45	1.0	0.0	0.05	1.0	0.0	0.05	1.0	0.0	0.05
Sense of co-presence	Q4	4.25	4.60	4.95	0.273	0.606	0.189	**0.023**	0.555	0.653	0.232	0.276	0.217
	Q5	3.80	4.30	4.15	0.154	0.333	0.293	0.260	0.260	0.197	0.659	0.101	0.071
Enjoyment	Q6	4.50	4.85	4.80	0.246	0.267	0.206	0.301	0.238	0.173	0.881	0.034	0.052
Affective intelligence	Q7	3.40	4.50	4.45	**0.010**	0.640	0.775	**0.008**	0.667	0.807	0.883	0.033	0.052
	Q8	4.05	4.60	4.35	0.086	0.405	0.406	0.368	0.206	0.142	0.262	0.259	0.196
	Q9	3.30	4.50	4.60	**0.004**	0.729	0.871	**0.004**	0.742	0.882	0.772	0.066	0.059
	Q10	4.25	4.60	4.35	0.297	0.240	0.175	0.789	0.061	0.058	0.349	0.215	0.150
Overall satisfaction	Q11	4.25	4.60	4.60	0.232	0.276	0.216	0.309	0.234	0.169	1.0	0.0	0.05

* R, M, and A stand for Random, Manual, and Automatic conditions, correspondingly.
Significant and marginally significant results are given in bold.

The **interactivity** dimension was measured using the statement '*The system was interactive'*. The results on interactivity in the M-condition are a little bit higher than in A-condition, while the minimum level of interactivity was reported in R-condition. Judging from the data (Table 8.2), there is no significant difference in interactivity between the three interfaces. However, we can say that M-condition is marginally more interactive than R-condition (p (R-M) = 0.068). We are aware that the notion of 'interactivity' is ambiguous, and hence the result is hard to interpret. The question was intended as more exploratory in nature. Apparently, subjects tended to evaluate the condition that allowed them to manipulate the expressed emotion manually as most interactive. Based on the results of post hoc computation of achieved power and the analysis of corresponding graphs (power as a function of sample size), we estimated the approximate number of subjects needed to achieve power 0.8 for some of our insignificant results to become significant (it is generally accepted that power should be 0.8 or greater to have high chance of finding a statistically significant difference when there is one). Given moderate effect (0.380) and achieved power (0.366) in the case of measuring the difference in interactivity between R-condition and A-condition, we would need to conduct experiment with 55 subjects (instead of 20) to have high probability of getting the significant result and reporting that the A-condition is significantly more interactive than the R-condition. In the case of insignificant difference reported on interactivity between M-condition and A-condition by 20 subjects in our study, the effect size is small (0.212), and we would need to have approximately 175 subjects to achieve power 0.8.

The **involvement (engagement)** dimension was evaluated using two questionnaire items: '*I felt it was important for my conversation partner that I responded after each his/her statement*' and '*I was awaiting the replies of my conversation partner with true interest*'. The statistic analysis results (Table 8.2) showed that the reported involvement of all three systems does not differ significantly, showing that the level of engagement was almost the same.

The following two questionnaire items covering the aspects of space and togetherness are intended for evaluation of **sense of co-presence**, or social presence: '*I felt if I were communicating with my conversation partner in the shared virtual space*', '*The system gave me the sense that the physical gap between us was narrowed*'. The ANOVA results for the first questionnaire item (Q4 in Table 8.2) support the significance of the difference in the sense of co-presence felt in A-condition and R-condition (p (R-A) < 0.05). This result indicates that the A-condition gave a stronger feeling of communication in the shared virtual space than the R-condition. To achieve the result showing that users rate Q4 in the A-condition significantly higher than in M-condition, we would need to obtain scores from about 105 subjects. Having approximately 125 subjects assessing Q4, we would achieve high probability of finding significant difference between M-condition and R-condition. As seen from the ANOVA results for the second questionnaire item for sense of co-presence (Q5 in Table 8.2), no significant difference among three interfaces was reported by 20 subjects on this statement. However, we would most probably get the results indicating that Q5 was rated significantly higher in M-condition and A-condition in comparison with R-condition, if we conducted the user study with 70 and 115 subjects, correspondingly.

The level of **enjoyment** was evaluated using the statement '*I enjoyed the communication using this IM system*'. The higher levels of enjoyment (Q6 in Table 8.2) were reported during A-condition and M-condition in comparison with R-condition. However, ANOVA resulted in no significant differences among interfaces. Based on the graphs depicting the power as a function of sample size, we estimated the approximate number of subjects needed to statistically prove the fact that M-condition and A-condition are significantly more enjoyable than R-condition (110 and 140 subjects, correspondingly).

To evaluate **affective intelligence**, four statements (three – directly related to the system and one – indirectly related) were proposed to subjects in the questionnaire: '*The system was successful at conveying my feelings*' (Q7), '*The system was successful at conveying my partner's feelings*' (Q8), '*The emotional behavior of the avatars was appropriate*' (Q9), and '*I understood the emotions of my communication partner*' (Q10).

As seen from the results for Q7 (Table 8.2), the systems in A-condition and M-condition (with small prevalence of mean results in M-condition) were both more successful at conveying own feelings than the system in R-condition. Since M-condition is considered as a gold standard in communicating person's emotions, and ANOVA showed no significant difference between M-condition and A-condition, we might say that automatic emotion recognition system performed well enough to bring high affective intelligence to IM application. As was expected, significant differences were found between R-condition and M-condition (p (R-M) < 0.05), and between R-condition and A-condition (p (R-A) < 0.01).

While evaluating successfulness of the interfaces at conveying conversation partner's feelings (see Q8 in Table 8.2), the highest rate was given by subjects to M-condition, and the lowest – to R-condition. However, ANOVA for this criterion resulted in no significant difference among all interfaces. One user's comment regarding the emotional reactions of the partner's avatar was: '*I concentrated too much on the reactions of my avatar and not enough on that of my partner. Reading and thinking about the answer took away the concentration on the avatar*'. In the case of measuring the difference in successfulness of M-condition and R-condition at conveying partner's emotions, the effect (0.405) and power (0.406) were moderate. The sample size should be set to approximately 50 subjects in order to achieve high probability of finding the significant difference between these interfaces.

Interesting results were observed for the evaluation of appropriateness of emotional behavior of avatars. As seen from the mean scores and statistical data of ANOVA (Q9 in Table 8.2), results for A-condition and M-condition significantly prevailed those for R-condition (p (R-A) < 0.01; and p (R-M) < 0.01). Users' comments confirmed that during R-condition subjects sometimes couldn't understand why the avatars did not correspond to their words and reacted in 'wrong' ways. Although A-condition was rated a little bit higher than M-condition, no significant difference was detected between these interfaces.

The statement '*I understood the emotions of my communication partner*' measured the affective intelligence of the system indirectly, since people used to derive emotional content from text based on semantic information and their empathetic abilities. Emotional expressions of avatars may help to understand the partner's emotion clearer. As was expected, the highest rate was reported in M-condition, and the lowest – in R-condition, where participants might be confused, since sometimes emotions shown by the avatar contradict actual emotional content (see mean scores for Q10 in Table 8.2). However, no significant difference was found in partner's emotion comprehension among all three interfaces. A possible explanation for such results might be that a person typically relies on his/her own affective intelligence rather than on results of artificial affective intelligence. That is why the mean for R-condition appeared relatively high (4.25).

The **overall satisfaction** from using three AffectIM interfaces was evaluated using the statement '*I am satisfied with the experience of communicating via this system*'. As can be seen from the results on Q11 in Table 8.2, average scores for A-condition and M-condition were equal (4.60), whereas less satisfaction was reported for R-condition (4.25). The results of ANOVA showed no significant difference in overall satisfaction among interfaces. Using the graphs depicting the power as a function of sample size, we estimated the approximate number of subjects needed to statistically prove the fact that M-condition and A-condition overall cause more satisfaction than R-condition (105 and 140 subjects, correspondingly).

In addition to the main questionnaire items, after finishing communications through all three interfaces, participants were given general questions. To the question '*While online, do you use emoticons or abbreviations?*', 19 subjects answered positively. We inspected all automatically recorded dialogs, and found that the majority of participants used abbreviated language. Regarding emoticons, the total number of such symbolic cues was about 30 percent higher during A-condition (29 emoticons) than during M-condition and R-condition (19 and 18 emoticons, respectively).

Due to the specificity of online communication media, no one can guarantee that users would not make grammatical and syntactical mistakes in text because of fast typing. It is evident that this may decrease the performance of the emotion recognition system. In Figure 8.3, the results of answers to the question '*How often do you make spelling mistakes because of fast typing?*' are displayed. As seen from the graph, 70 percent of female subjects and 60 percent of male subjects reported high frequency (often or always) of making mistakes. While analyzing the recorded conversations, we detected the following misspelled words that influenced the result of AAM system and, therefore, on the displayed emotion: '*feiled*' instead of '*failed*'; '*dispointed*' instead of '*disappointed*'; '*dipressed*' instead of '*depressed*'; '*beter*' instead of '*better*'; '*promissing*' instead of '*promising*', etc.

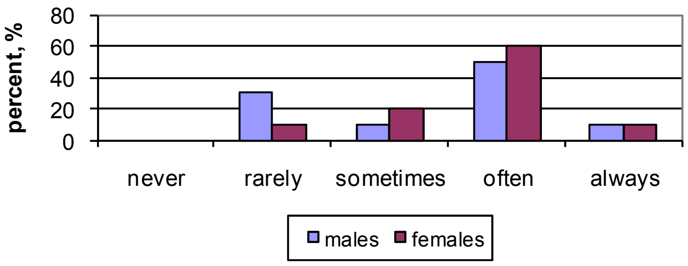

Figure 8.3 The reported frequency of making spelling mistakes because of fast typing

The participants' comments and answers to the question '*To what degree do you think is necessary to look at a graphical representation of the other communicating person?*' suggest that there are two types of IM users: (1) some are open to new features of IM, and consider animated graphical representation of a person helpful in understanding the partner's emotions and giving some sense of physical presence; (2) others tend to concentrate their attention on content, and prefer small emotional symbolic cues, like emoticons, to avatar expressions.

The participants were also asked to indicate whether manual selections of emotion state and intensity were helpful or not during M-condition. Only 30 percent of males and 60 percent of females answered positively. The result of answers to the question '*How often did you use this function, when you wanted?*' is represented as a bar graph in Figure 8.4. As seen from these data, female subjects used emotion selection function more consistently than male subjects.

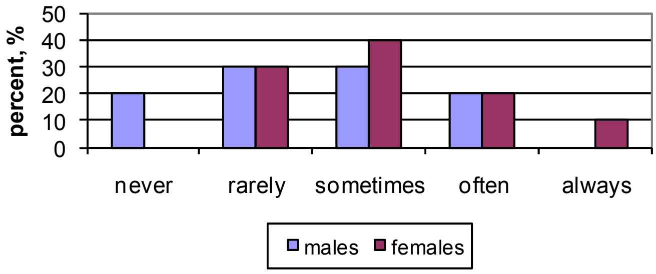

Figure 8.4 The reported frequency of usage of emotion 'select menu'

The users' opinions regarding the emotion 'select menu' aspect were very diverse. Some users criticized the type of pop-up menu, commenting that it was difficult to use, it took long time to select, and choosing emotion intensity on fine-grained scale was overwhelming. For more convenience, they proposed to replace pop-up menus by icons and spread them out. One of the subjects complained that emotion select menu disturbed the flow of the chat. Another reported problem is that since there is no preview of what the emotion expression looks like, it is unclear whether it matches the user's intention. Some subjects felt that basic emotions are too general and are not sufficient to convey emotion in many cases (like, for example, in phrase '*hey, just kidding*', where users typically use emoticon). Also, they suggested providing the possibility of showing more different or even mixed emotions (some state between sadness and joy). However, we think that displaying mixed emotional expressions would add more confusion and misinterpretation to the conversations.

Some subjects underlined positive aspects of manual selection of emotion states. They found this function interesting and helpful, because: (1) it offered the possibility to visually express feelings and better understand them, (2) it allowed preventing inappropriate emotional reaction of avatar, and (3) it guaranteed accuracy of communicated emotion. We can conclude that for sensitive conversation users would prefer manual control to avoid system mistakes that could sometimes harm the conversation.

Regarding the users' general impressions from using AffectIM system, they found automatic recognition and visual representation of emotions to be a good, promising idea and fun to watch. The representation of emotions through the avatars was interesting, clear and easy to understand, however, some participants reported the difficulty of distinguishing the displayed emotions and stated that the system would benefit from better customizable avatars or even abstract avatars like smileys. Several users commented that they would also like to express a wider range of emotions.

8.1.4 Summary and Discussion of Results

Based on the developed Affect Analysis Model we implemented an online IM application, AffectIM. To realize visual reflection of textual affective information, we have designed two animated avatars performing various expressive patterns (emotions, social behavior, and natural idle movements), contributing thus to greater interactivity. The developed AffectIM supports online communication, allows users to see the conversation flow in three modes (plain text, transcribed text, or text annotated with emotion), and visualizes the communicated emotions, emotion distribution and emotion dynamics.

We conducted a twenty-person user study (within subjects design)

(1) To find out whether the *user experience* with AffectIM interface featuring manual annotations (using emotion and intensity select menus) of avatar emotional expressions is better than the *user experience* with the AffectIM interface based on our Affect Analysis Model (system for automatic recognition of dominant emotion and its intensity on the level of distinct sentences), which drives the automatic display of the detected emotions by avatars.

(2) To support the hypothesis that the *user experiences* with the Manual and Automatic

AffectIM interfaces are better than the *user experience* with the AffectIM interface with randomly displayed avatar emotions.

Among the three interfaces, the Manual interface provides the most accurate way to annotate or accompany textual message with user emotion, because the user selects the to-be-conveyed emotion by him- or herself. For that reason, this condition was considered as the gold standard in our study. The first hypothesis has been rejected: no significant differences between Manual and Automatic interfaces were reported by the subjects in judging interactivity, involvement (engagement), sense of co-presence, enjoyment, affective intelligence, and overall satisfaction. The user experience with Manual interface is not statistically better than with Automatic interface. The results of the study indicate that our IM system with automatic emotion recognition function can achieve a level of affective intelligence (system is successful at conveying the users feelings, avatar expression is appropriate) that is comparable to the gold standard, where users select the label of the conveyed emotion manually.

However, this result does not provide sufficient argument for preferring an Automatic interface over a Manual interface. Some might argue that manual selection is the most intuitive way to annotate emotions. Here, we have to recall the specific nature of IM as a typically fast-paced activity, where users cannot easily divert their attention to other features of the interface, such as manual selection of emotion (Persson 2003) and its intensity. In the user study, we assert (rather than demonstrate statistically), that automatic emotion recognition and expression are more efficient, in the confined sense of less disruptive, than manual selection.

At the same time we keep in mind that occasionally, users may want to select the conveyed emotion manually. From informal subject remarks we learned that the IM application might benefit from an integration of automatic emotion sensing with manual control of emotional behavior of avatars in one interface, which will allow users to select between two modes depending on type and sensitivity of conversation, or to manually correct the automatic emotional expression of avatar.

The second hypothesis, which postulates that the quality of user experience with the Manual and Automatic interfaces is higher than with the (quasi) Random interface, has been statistically proven (significant differences were found) for such dimensions as: (1) interactivity (in the case of Manual vs Random interface), (2) sense of co-presence (the feeling of being together and communicating with another person in the shared virtual space) (in the case of Automatic vs Random interface), and affective intelligence (the successfulness of the system at conveying user feelings; the appropriateness of the emotional behavior of avatars) (in the case of Manual vs Random interface, and Automatic vs Random interface). While this result might not be surprising, the comparison is formally sensible, as in each condition, users should in principle be able to experience the same set of avatar expressions. The condition serves as a control condition to rule out the possibility that any expression of avatar emotion in the IM interface leads to a good user experience.

Alternatively, we could have compared the Automatic interface with an interface that shows no or a fixed emotion, such as low-intensity joy. However, the main focus of our work was to show that our automatic emotion recognition technique (the Affect Analysis Model) in IM performs on the level of the manual selection. We did not aim to show that the emotional avatars are (in some sense) better than non-emotional avatars, which has already been addressed widely in the literature (Olveres et al. 1998; Liu et al. 2003; Fabri et al. 2007).

The user study conducted on AffectIM showed that the IM system with automatic emotion recognition function was successful at conveying users' emotional states during communication online, thus enriching expressivity and social interactivity of online communications. Part of the participants considered the animated graphical representation of a person helpful in understanding the partner's emotions and giving some sense of physical presence. Users reported that their experience with AffectIM was fun and interesting.

Looking at all three interfaces, no significant differences were found for such dimensions as interactivity (except of significantly higher interactivity reported in Manual interface than in Random interface), involvement (engagement), enjoyment, and overall satisfaction. This has to be seen as a negative result of our study or as a sign that the mere presence of an avatar makes users feel involved, satisfied, and enjoy the interface. Effectively, we cannot provide a conclusive answer, since we did not compare avatar versus non-avatar versions of the system. The reason is similar to having no non-emotional condition, namely, those studies were already conducted by others (Olveres et al. 1998; Persson 2003).

There are also more practical reasons for some non-significant results. Based on the post hoc computation of achieved power and analysis of graphs depicting the power as a function of sample size, we estimated the approximate number of subjects needed to statistically prove the facts that (1) Automatic interface is significantly more interactive than Random interface (55 subjects); (2) Manual and Automatic interfaces are significantly more enjoyable than Random interface (110 and 140 subjects, correspondingly); (3) Manual and Automatic interfaces overall cause more satisfaction than Random interface (105 and 140 subjects, correspondingly). Therefore, to get better results, we would need to conduct a user study with more than 100 subjects. We concede that given a larger sample size (more than 175 subjects), Manual interface might be considered significantly more interactive than Automatic interface, which is explicable, as former interface allows subjects to manipulate the expressed emotion manually; Manual interface might outperform Automatic interface in the aspects of affective intelligence related to the emotional reactions of the partner's avatar ('*The system was successful at conveying my partner's feelings*') and to the comprehension of the partner's emotion ('*I understood the emotions of my communication partner*'). A study of such size, unfortunately, surpasses the possibilities of our Graduate School Research Laboratory.

8.2 EmoHeart: Conveying Emotions in Second Life

The 3D virtual world of Second Life (*http://secondlife.com*) imitates a form of real life by providing a space for rich interactions and social events. Second Life encourages people to establish or strengthen interpersonal relations, to share ideas, to gain new experiences, and to feel genuine emotions accompanying all adventures of virtual reality. Emotional expression is natural and very important for communication in real life, but currently rather cumbersome in Second Life, where expressions have to be selected and activated manually. Concretely, a user has to click on the animation gesture in the list, or type the predefined command following the symbol '/' in a textual chat entry. In order to breathe emotional life into graphical representations of users through the automation of emotional expressiveness, we applied the developed Affect Analysis Model to textual chat in Second Life (Neviarouskaya et al. 2010a).

8.2.1 Overview of the EmoHeart System

The architecture of the EmoHeart system is presented in Figure 8.5.

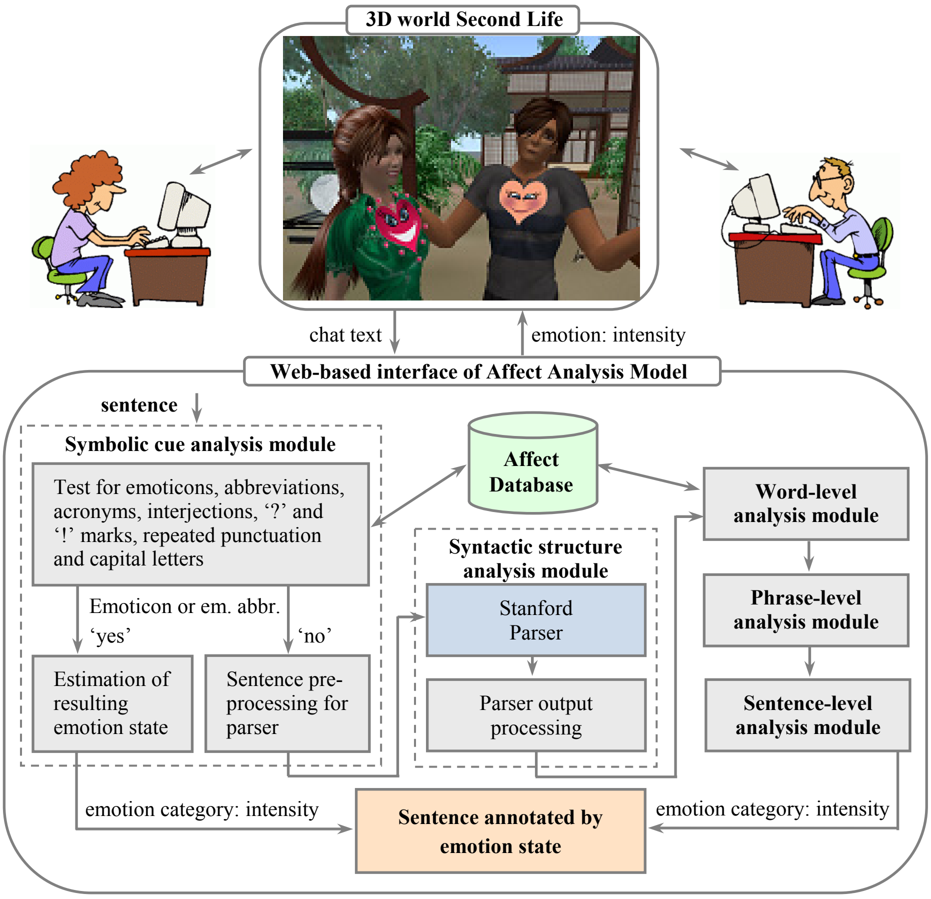

Figure 8.5 Architecture of the EmoHeart system

In order to make the EmoHeart system freely available for Second Life users, we decided to employ GNU GPL licensed Stanford Parser (De Marneffe, MacCartney, and Manning 2006) in place of the commercial parser, Connexor Machinese Syntax (see Section 3.2 for details), in the syntactic structure analysis stage of the Affect Analysis Model. The Stanford Parser is available at *http://nlp.stanford.edu/software/lex-parser.shtml*.

It is worth noting, however, that the accuracy of the Affect Analysis Model with the Connexor Machinese Syntax is higher in 6-8 percent than with the Stanford Parser on the data set described in Section 4.1 (see details of comparison in Table 8.3). This indicates that Stanford Parser employed for the syntactical structure analysis is less efficient. On the other hand, as we aim to freely distribute and apply our emotion recognition tool to textual messages in Second Life, we have to compromise on the performance of the system for the sake of free distribution.

Table 8.3 Comparison of accuracy of Affect Analysis Model employing different parsers (Connexor Machinese Syntax vs Stanford Parser)

Measure	Gold standard			
	At least two annotators agreed		All three annotators agreed	
	Fine-grained categories	Merged labels	Fine-grained categories	Merged labels
Accuracy of AAM with Connexor Machinese Syntax	0.726*	0.816*	0.815**	0.890**
Accuracy of AAM with Stanford Parser	0.649	0.747	0.751	0.814
Difference in percent, %	7.7	6.9	6.4	7.6

* Data taken from Table 4.3, Section 4.1.
** Data taken from Table 4.6, Section 4.1.

In Second Life, the Affect Analysis Model serves as the engine behind automatic visualization of emotions conveyed through textual messages. The control of the conversation in Second Life is implemented through the object called EmoHeart (invisible in case of neutral state) attached to the avatar's chest. The distributor of the EmoHeart object is located inside a (fictitious) Starbucks cafe of the Second Life replica of National Center of Sciences building in Tokyo (Second Life landmark: *http://slurl.com/secondlife/NIIsland/213/38/25/*). Once attached to the avatar, EmoHeart object (1) listens to each message of its owner, (2) sends it to the web-based interface of the Affect Analysis Model located on the server, (3) receives the result (dominant emotion and intensity), and visually reflects the sensed affective state through the animation of avatar's facial expression, EmoHeart texture (indicating the type of emotion), and size of the texture (indicating the strength of emotion, namely, 'low', 'middle', or 'high'). If no emotion is detected in the text, the EmoHeart remains invisible and the facial expression remains neutral.

Of the bodily organs, the heart plays a particularly important role in our emotional experience. People often characterize personal traits, emotional experiences, or mental states using expressions originating from word '*heart*' (for example, '*heartfelt*', '*warm-hearted*', '*heartlessness*', '*kind-heartedness*', '*broken-hearted*', '*heart-burning*', '*heart-to-heart*' etc.). The essence of emotional, moral, and spiritual aspects of a human being has long been depicted using heart-shaped symbol.

With the heart-shaped object of EmoHeart, we provide an additional channel for visualizing emotions in a vivid and expressive way. The examples of avatar facial expressions and EmoHeart textures are shown in Figure 8.6.

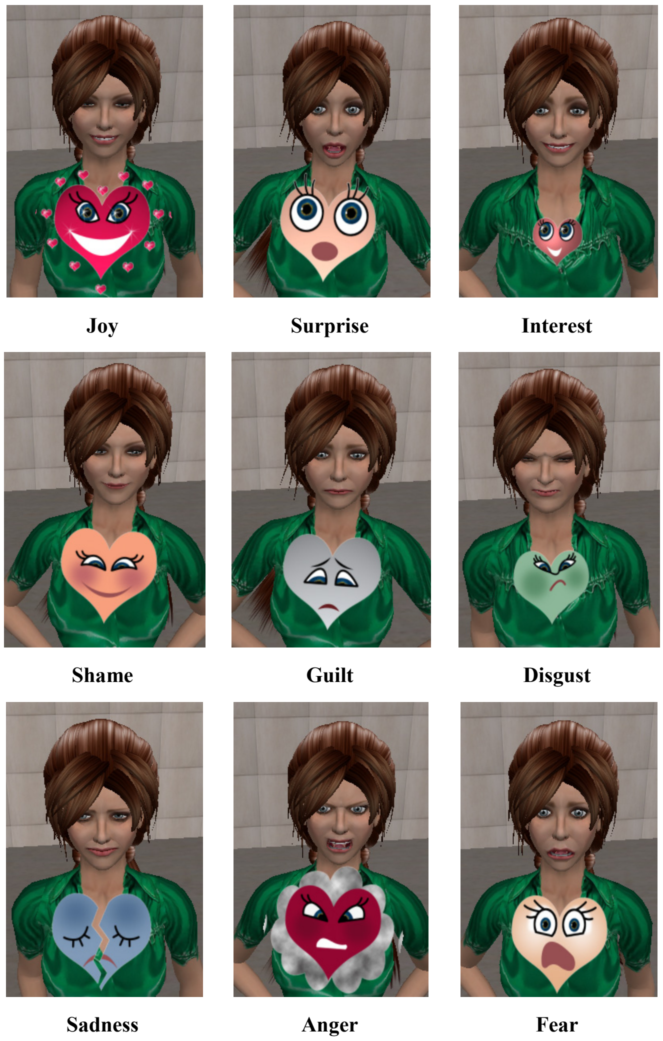

Figure 8.6 Examples of avatar facial expressions and EmoHeart textures

While designing EmoHeart textures, we followed the description of main characteristic features of the expressive means in relation to the communicated emotion (see Appendix C for details).

8.2.2 Analysis of the EmoHeart Log

We made EmoHeart available for Second Life users from December 2008 (see demonstration video at *http://www.youtube.com/watch?v=1TZ_74_LywE*). During a two month period (December 2008 – January 2009), we asked students to promote the EmoHeart object by visiting locations in Second

Life and engaging other Second Life residents in social communication. As a result, 89 Second Life users became owners of EmoHeart, and 74 of them actually communicated using it. Text messages along with the results from the Affect Analysis Model were stored in an EmoHeart log database. Some general statistics is given in Table 8.4. As seen from the table, the chat activity of users within two months (from 1 message to 2932 messages per user), as well as the length of a chat message in symbols (from 1 symbol to 634 symbols per message), varied significantly. In average, typical chat message included one sentence.

Table 8.4 Statistics on EmoHeart log of 74 users for period December 2008 – January 2009

Measure	Messages, number	Message length, symbols	Sentences, number
Total	19591 (for all users)	400420 (for all messages)	21396 (for all messages)
Minimal	1 (for user)	1 (for message)	1 (for message)
Maximal	2932 (for user)	634 (for message)	25 (for message)
Average	265 (per user)	20 (per message)	1.09 (per message)

From all sentences, 20 percent were categorized as emotional by the Affect Analysis Model and 80 percent as neutral (Figure 8.7). We observed that the percentage of sentences annotated by positive emotions ('Joy', 'Interest', and 'Surprise') essentially prevailed (84.6 percent) over sentences annotated by negative emotions ('Anger', 'Disgust', 'Fear', 'Guilt', 'Sadness', 'Shame'). We believe that this dominance of positivity expressed through text is due to the nature and purpose of online communication media, which allows people to exchange experiences, share opinions and feelings, and satisfy their social need of interpersonal communication. Harker and Keltner (2001) empirically verified that the tendency to express positive emotions creates more harmonious social relationships, which in turn fosters personal growth and well-being.

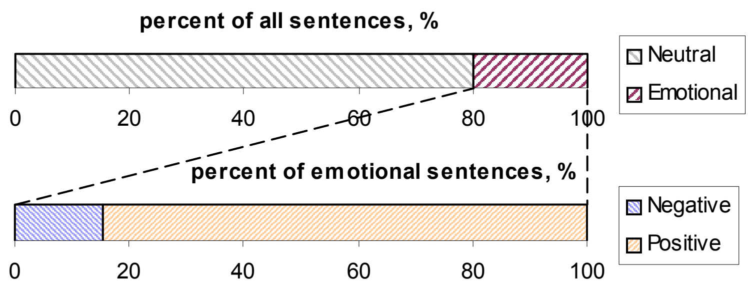

Figure 8.7 Percentage distribution of emotional (positive or negative) and neutral sentences

We analysed the distribution of emotional sentences from EmoHeart log data according to the fine-grained emotion labels from our Affect Analysis Model. We found that the most frequent emotion conveyed through text messages is 'Joy' (68.8 percent of all emotional sentences), followed by 'Surprise', 'Sadness' and 'Interest' (9.0 percent, 8.8 percent, and 6.9 percent, respectively). All remaining emotions individually do not exceed the level of 2.1 percent. The least frequent emotion detected from text messages is 'Shame' (0.6 percent of all emotional sentences).

8.3 iFeel_IM!: Innovative Real-Time Communication System with Rich Emotional and Haptic Channels

Driven by the motivation to enhance emotionally immersive experience of real-time messaging, we developed the system iFeel_IM! (**i**ntelligent system for **Feel**ing enhancement powered by affect sensitive **I**nstant **M**essenger), that employs haptic devices and visual stimulation to convey and augment the emotions experienced during online conversations. The philosophy behind the system is '***I feel*** [therefore] ***I am!***'. The architecture of the iFeel_IM! is presented in Figure 8.8.

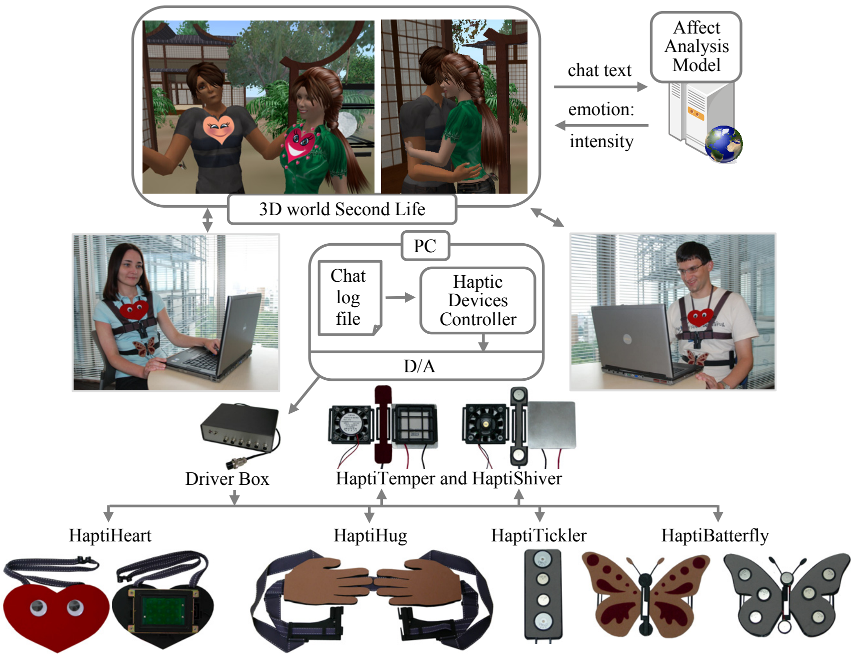

Figure 8.8 Architecture of the iFeel_IM! system

The system integrates:

(1) 3D world Second Life (with EmoHeart) as a platform for communication.

(2) Affect Analysis Model as an intelligent component for automatic emotion recognition from text messages.

(3) Innovative affective haptic interfaces providing additional nonverbal communication channels through simulation of emotional feedback and social touch (physical co-presence). iFeel_IM! users can not only exchange messages but also emotionally and physically feel the presence of the communication partner (e.g., family member, friend, or beloved person). In order to communicate through iFeel_IM! system, users have to wear the following affective haptic devices: HaptiHeart, HaptiHug, HaptiButterfly, HaptiTickler, HaptiTemper, and HaptiShiver.

In the iFeel_IM!, in addition to communication with the system for textual affect sensing (Affect Analysis Model), EmoHeart is responsible for sensing symbolic cues or keywords of 'hug' communicative function conveyed by text, and for visualization (triggering related animation) of 'hugging' in Second Life. The results from the Affect Analysis Model (dominant emotion and intensity) and EmoHeart ('hug' communicative function) are stored along with chat messages in a file on local computer of each user. The Haptic Devices Controller analyses these data in a real time and generates control signals for the Digital/Analog converter (D/A), which then feeds the Driver Box for haptic devices with the control cues. Based on the transmitted signal, the corresponding haptic device worn by user is activated.

We selected four distinct emotions that have strong physical features: ('Anger', 'Fear', 'Sadness', and 'Joy') for presentation through the haptic devices. The precision of AAM in recognition of these emotions is considerably higher than of other emotions (see Table 4.2 and Table 4.6 in Section 4.1). There are three types of affective haptic devices incorporated in the iFeel_IM!:

(1) HaptiHeart, HaptiButterfly, HaptiTemper, and HaptiShiver are intended for implicit elicitation of emotions.

(2) HaptiTickler directly evokes emotion.

(3) HaptiHug uses social touch to influence the mood and provide a sense of physical co-presence.

The developed heart imitator HaptiHeart consists of two modules (flat speaker and speaker holder) and is able to produce realistic heartbeat patterns according to emotion to be conveyed or elicited (sadness is associated with slightly intense heartbeat, anger with quick and violent heartbeat, and fear with intense heart rate), by means of the pre-recorded sound signal with low frequency that generates the pressure on the human chest through vibration of the speaker surface.

The HaptiButterfly was developed with the aim to evoke joy emotion. The idea behind this device is to reproduce effect of '*butterflies in the stomach*' (fluttery or tickling sensation) by means of the arrays of vibration motors attached to the abdomen area of a person.

To boost fear, HaptiShiver sends '*shivers up and down your spine*' through a row of vibration motors, and HaptiTemper sends '*chills up and down your spine*' through both cold airflow from a DC fan and the cold side of a Peltier element. HaptiTemper is also responsible for imitating 'warm' and 'hot' sensations of joy and anger emotions, respectively.

The HaptiTickler device for stimulation of joy emotion includes four vibration motors reproducing stimuli that are similar to human finger movements during rib tickling. The uniqueness of our approach is in (1) combination of the unpredictability and uncontrollability of the tickling sensation through random activation of stimuli, (2) high involvement of the social and emotional factors in the process of tickling (positively charged on-line conversation potentiates the tickle response).

The key feature of the developed HaptiHug is that it physically reproduces the hug pattern similar to that of human-human interaction. The hands for a HaptiHug are sketched from a real human and made from soft material so that hugging partners can realistically feel social presence of each other. The couple of oppositely rotating motors are incorporated into the holder placed on the user chest area. The Soft Hands, which are aligned horizontally, contact back of the user. Once 'hug' command is received, the couple of motors tense the belt, thus pressing the Soft Hands and chest part of the HaptiHug in the direction of a human body.

While developing the iFeel_IM! system, which was shortly summarized in this Section (details are given in (Neviarouskaya et al. 2009; Tsetserukou and Neviarouskaya 2010)), we attempted to bridge the gap between mediated and face-to-face communications by enabling and enriching the spectrum of senses such as vision and touch along with cognition and inner personal state.

8.4 Web-Based @AM Interface

We developed a web-based interface integrated with the Attitude Analysis Model (@AM) to provide users with a convenient tool enabling real-time online recognition of fine-grained attitudes (affect, judgment, and appreciation) conveyed in text. The screenshot of the web-based @AM interface is shown in Figure 8.9.

Figure 8.9 Screenshot of the web-based @AM interface

The @AM interface is divided into four functional areas:

(1) '@AM parameters' frame.

(2) '@AM functionality' frame.

(3) 'Textual attitude analysis' frame.

(4) 'Visualization of attitude statistics' frame.

The '@AM parameters' frame allows users to fix and regulate parameters (in acceptable bounds) for the @AM algorithm, namely, intensifying coefficients for all-capital words (e.g., '*HAPPY*'), adjectives and adverbs in a comparative or superlative degree (e.g., '*wiser*', '*wisest*'), adjectives (e.g., '*rapidly-growing*') and nouns (e.g., '*increase*') of the *intensifying* type, and a reinforcement coefficient for clause-level analysis. The system verifies the values of coefficients. For example, if a user assigns the value of a coefficient for superlative degree less than for comparative degree, then the system will warn the user with the following message: '*Coefficient for comparative degree cannot be higher than or equal to coefficient for superlative degree. Please modify these coefficients*.'

The '@AM functionality' frame empowers the user to modify the configuration of the Attitude Analysis Model, that is to enable or disable various functionality components of the @AM algorithm on different levels:

(1) Word level: intensification of all-capital words, adjectives and adverbs in a comparative or superlative degree.

(2) Parser: selection of a tool for syntactic and dependency parsing of a sentence.

(3) Phrase level: intensification by modifiers (adverbs of degree and adverbs of affirmation); intensification by adjectives and nouns of the *intensifying* type; reversal by modifiers (adverbs of doubt and adverbs of falseness); reversal by adjectives (e.g., '*reduced*') and nouns (e.g., '*termination*') of the *reversing* type; reversal by negative determiners (e.g., '*no*'); and neutralization by prepositions (e.g., '*without*').

(4) Clause/sentence level: reversal by negations; reversal by modifiers (adverbs of doubt and adverbs of falseness); neutralization by prepositions; neutralization due to condition; neutralization due to connector '*but*'; application of the rules for semantically distinct verb classes; and adjustment of the attitude label based on the analysis of personal pronouns, WordNet high-level concepts, and Stanford NER labels.

'Textual attitude analysis' frame is represented by two text boxes:

(1) The area for the user's textual input, where the user inserts the original text (sentences or paragraphs) to be analysed, or loads the sample sentences from the text files.

(2) The area for the output of the @AM system. The results of the attitude analysis (sentences annotated by the attitude type, attitude strength, confidence level, etc.) can be displayed in three formats: annotations on the sentence level, annotations on the clause level, and annotations on the word level. Figure 8.10 and Figure 8.11 show the examples of the @AM output for the sentences '*It's no wonder that this child has no respect for anyone*' and '*Audible chewing can be rather disgusting, especially if you are also trying to enjoy food*'.

Figure 8.10 @AM output for the sentence '*It's no wonder that this child has no respect for anyone*'

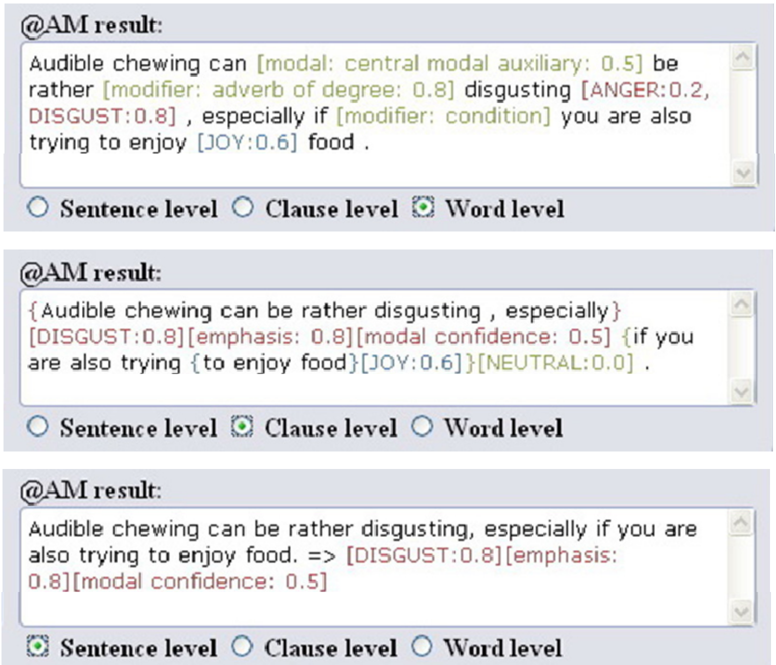

Figure 8.11 @AM output for the sentence '*Audible chewing can be rather disgusting, especially if you are also trying to enjoy food*'

'Visualization of attitude statistics' frame displays (1) the distribution of the attitude labels in the analysed text using either a pie chart or a bar chart; (2) the dynamics of the attitude from sentence to sentence over the whole text using a line chart. Both attitude dynamics and attitude distribution plots can be drawn using different number of labels depending on the level of attitude hierarchy (14 labels on ALL level, 7 labels on MID level, and 3 labels on TOP level).

Chapter 9

Discussion and Conclusions

Sentiment or subjectivity analysis is nowadays a rapidly developing field with a variety of emerging approaches targeting the recognition of sentiment reflected in written language. Recognition of positive and negative opinions and classification of text using emotion labels have been gaining increased attention of researchers. However, the topic of recognition of fine-grained attitudes expressed in text has been ignored. Attitude types (namely, affect, judgment, and appreciation) define the specifics of appraisal being expressed: distinct types of personal emotional states; positive and negative appraisal of person's character, behavior, skills etc.; and aesthetic evaluation of semiotic and natural phenomena, events, objects etc., correspondingly. In our research we developed robust computational tools for the analysis of fine-grained attitudes conveyed in text: Affect Analysis Model (AAM) for automatic recognition of nine emotional states and Attitude Analysis Model (@AM) for fine-rained attitude sensing (nine emotional states, positive and negative judgments, positive and negative appreciations).

To deal with limitation in sentiment lexicon coverage, we proposed original methods for expanding the sentiment lexicon. Our SentiFul database, which contains about 12900 sentiment-conveying words (it is larger than the existing lists of sentiment words), was automatically built using methods exploring direct synonymy and antonymy relations, hyponymy relations, and innovative methods based on morphologic modifications and compounding with known lexical units (the originality and valuable contribution lie in the elaborate patterns/rules for the derivation and compounding processes that have not been considered before). The evaluations of the proposed methods showed that they achieved high accuracy in assigning dominant polarity labels and polarity scores to the words. The method based on compounding performed with the highest accuracy in assigning dominant positive or negative labels, followed by the methods considering hyponymy relations, derivation process, synonymy relations, and antonymy relations (this method yielded noisy results). We believe that our innovative methods for derivation of new sentiment-related English terms (particularly, morphologic modifications and compounding) can be applicable to other languages, especially fusional languages that use bound morphemes and are characterized by a rich inflectional system. In order to support the analysis of contextual attitude and its strength, we created AttitudeFul database that contains attitude-conveying terms, extensive sets of modifiers, contextual valence shifters, and modal operators.

In this work, we introduced novel compositional linguistic approach to fine-grained attitude recognition in text. In contrast to other methods that mainly focus on two sentiment categories (positive and negative) or six basic emotions, our Attitude Analysis Model classifies individual sentences using fine-grained attitude labels (nine for different affective states, two for positive and negative judgment, and two for positive and negative appreciation). Currently, machine learning methods for sentiment or affect analysis suffer from the following weak points: large corpora required for meaningful statistics and good performance; neglect of some prepositions, negation, modal, and condition constructions; disregard of syntactic relations and semantic dependencies in sentences; and long processing time. Our @AM is domain-independent, and it greatly benefits from (1) the use of AttitudeFul lexicon; (2) the analysis of syntactic and dependency relations between words in a sentence; (3) the representation of sentence structure using Subject, Verb, and Object formations; (4) the proposed *compositionality principle* (the rules of *polarity reversal, aggregation (fusion)*, *propagation*, *domination*, *neutralization*, and *intensification*, at various grammatical levels); (5) the rules elaborated for semantically distinct verb classes; and (6) a method considering the hierarchy of concepts based on WordNet and StanfordNER. As distinct from the state-of-the-art approaches, the proposed compositional linguistic approach to automatic recognition of fine-grained affect, judgment, and appreciation in text (1) extensively deals with the semantics of terms, which allows accurate and robust automatic analysis of attitude type, and broadens the coverage of sentences with complex contextual attitude; (2) processes sentences of different complexity, including simple, compound, complex (with complement and relative clauses), and complex-compound sentences; (3) handles not only correctly written text, but also informal messages written in an abbreviated or expressive manner; and (4) encodes the strength of the attitude and the level of confidence, with which the attitude is expressed, through numerical values in the interval [0.0, 1.0].

We have conducted several experiments with AAM and @AM on the data sets of sentences from different domains: diary-like blog posts, personal stories about life experiences, fairy tales, and news headlines. Table 9.1 contains the summary of experimental results. Our AAM showed promising results in emotion recognition on real examples of diary-like blog posts; and @AM performed with high level of accuracy on sentences from personal stories about life experiences, fairy tales, and news headlines, outperforming other methods on several measures.

Table 9.1 The summary of experimental results

Data set	**Level of classification**	**Accuracy of our methods, %**	**Averaged accuracy of other best performed methods, %**		
Our collection of diary-like blog sentences (at least two annotators agreed)		AAM:	-		
	Fine-grained, 10 labels	72.6			
	Coarse-grained, 3 labels	81.6			
Our collection of diary-like blog sentences (all three annotators agreed)		AAM:	-		
	Fine-grained, 10 labels	81.5			
	Coarse-grained, 3 labels	89.0			
Emotion blog sentences		AAM:	ML with unigrams:	ML with unigrams, RT and WNA features:	
	Fine-grained, 7 labels	77.0	73.5*	70.2*	
Our data set from the Experience Project		@AM:	-		
	ALL level, 14 labels	62.1			
	MID level, 7 labels	70.9			
	TOP level, 3 labels	87.9			
Fairy tales		@AM:	LOOHAsnowtag:		
	ALL level, 6 labels	63.9	69-70		
	MID level, 3 labels	72.8	69-73		
	TOP level, 2 labels	80.8	79		
News headlines		@AM:	UPAR7:	SWAT:	UA:
	6 labels (fine-grained evaluation)	26.57**	28.38**	25.41**	14.15**
	6 labels (coarse-grained evaluation)	27.26***	8.71***	11.57***	9.51***

* Value calculated as average of precisions (see Table 4.8).
** Averaged Pearson's correlation coefficients.
*** Averaged F-scores.

Using the Affect Analysis Model and the Attitude Analysis Model, we have developed several applications: AffectIM (Instant Messaging application integrated with the AAM), EmoHeart (application of AAM in 3D world Second Life), iFeel_IM! (innovative real-time communication system with rich emotional and haptic channels), and web-based @AM interface.

We believe that the output of our Attitude Analysis Model can contribute to the robustness of the following society-beneficial and analytical applications: public opinion mining, deep understanding of a market and trends in consumers' subjective feedback, attitude-based recommendation system, economic and political forecasting, affect-sensitive and empathic dialogue agent, emotionally expressive storytelling, integration into online communication media (IM, 3D virtual world etc.) and social networks (e.g., Facebook, Twitter). The web-based system for attitude or opinion search may influence the decisions of product developers and potential customers, may enable analysts to examine the trends in public reactions to political decisions, social events, etc. The integration of attitude-sensing system into online communication media and social networks may have a great impact on the depth of emotional interpersonal connections in online community. Elderly and alone people may benefit from the interaction with empathic virtual agent that is able to sense person's emotional state from different modalities.

The primary objective for future research is to elaborate the algorithms for the extraction of an attitude holder, topic, causes/reasons, and consequences. The automatic detection of correlations between attitude and cause event, and the analysis of attitude and its consequences are new research topics and have a potentially strong impact on the robustness of a variety of real-world analytical applications.

Bibliography

Ajzen, I., and Fishbein, M. 1980. *Understanding Attitudes and Predicting Social Behavior*. Prentice-Hall, US. 1980.

Allwood, J. 2002. Bodily communication dimensions of expression and content. In B. Granstrom, D. House, and I. Karlsson (eds.), *Multimodality in Language and Speech Systems*, Kluwer Academic Publishers, Netherlands, pp. 7-26.

Alm, C. O. 2008. Affect in text and speech. PhD dissertation, University of Illinois at Urbana-Champaign.

Alm, C. O., Roth, D., and Sproat, R. 2005. Emotions from text: machine learning for text-based emotion prediction. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pp. 579-586.

Aman, S., and Szpakowicz, S. 2007. Identifying expressions of emotion in text. In *Proceedings of the Tenth International Conference on Text, Speech and Dialogue TSD 2007*, Plzen, Czech Republic, Lecture Notes in Computer Science, vol. 4629, Springer, pp. 196-205.

Aman, S., and Szpakowicz, S. 2008. Using Roget's Thesaurus for fine-grained emotion recognition. In *Proceedings of the Third International Joint Conference on Natural Language Processing IJCNLP 2008*, Hyderabad, India, pp. 296-302.

Andreevskaia, A., and Bergler, S. 2006. Mining WordNet for fuzzy sentiment: sentiment tag extraction from WordNet glosses. In *Proceedings of the Eleventh Conference of the European Chapter of the Association for Computational Linguistics*, pp. 209-216.

Androutsopoulos, J. 2006. Introduction: Sociolinguistics and computer-mediated communication. *Journal of Sociolinguistics* 10(4): 419-438.

Argamon, S., Bloom, K., Esuli, A., and Sebastiani, F. 2007. Automatically determining attitude type and force for sentiment analysis. In *Proceedings of the Third Language and Technology Conference*.

Banea, C., Mihalcea, R., and Wiebe, J. 2008. A bootstrapping method for building subjectivity lexicons for languages with scarce resources. In *Proceedings of the International Conference on Language Resources and Evaluations (LREC'08)*.

Baroni, M., and Vegnaduzzo, S. 2004. Identifying subjective adjectives through Web-based mutual information. In *Proceedings of the 7th German Conference on Natural Language Processing*.

Benamara, F., Cesarano, C., Picariello, A., Reforgiato, D., and Subrahmanian, V. S. 2007. Sentiment analysis: adjectives and adverbs are better than adjectives alone. In *Proceedings of the International Conference on Weblogs and Social Media*, OMNIPRESS, Boulder, CO, pp. 203-206.

Biber, D., Johansson, S., Leech, G., Conrad, S., Finegan, E., and Quirk, R. 1999. *Longman Grammar of Spoken and Written English*. Pearson Education Limited.

Boucouvalas, A. C. 2003. Real time text-to-emotion engine for expressive Internet communications. In *Being There: Concepts, Effects and Measurement of User Presence in Synthetic Environments*, Ios Press, pp. 306-318.

Cahn, J. E. 1990. The generation of affect in synthesized speech. *Journal of the American Voice I/O Society* 8: 1-19.

Cairns, H. S., and Cairns, C. E. 1976. *Psycholinguistics: A Cognitive View of Language*. New York: Holt, Rinehart, and Winston.

Castellano, G., Kessous, L., and Caridakis, G. 2008. Emotion recognition through multiple modalities: face, body gesture, speech. In C. Peter and R. Beale (eds.), *Affect and Emotion in Human-Computer Interaction*. LNCS 4868, Springer Heidelberg, pp. 92-103.

Chaumartin, F. 2007. Upar7: A knowledge-based system for headline sentiment tagging. In *Proceedings of SemEval-2007*, Prague, Czech Republic.

Chayko, M. 2002. *Connecting: How We Form Social Bonds and Communities in the Internet Age*. Albany, NY: State University of New York Press.

Choi, Y., and Cardie, C. 2008. Learning with compositional semantics as structural inference for subsentential sentiment analysis. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 793-801.

Chuang, Z.-J., and Wu, C.-H. 2004. Multi-modal emotion recognition from speech and text. *Computational Linguistics and Chinese Language Processing* 9(2): 45-62.

Davitz, J. R. 1969. *The Language of Emotions*. London: Academic Press.

De Marneffe, M.-C., MacCartney, B., and Manning, C. D. 2006. Generating typed dependency parses from phrase structure parses. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2006)*.

Derks, D. 2007. Exploring the missing wink: emoticons in cyberspace. Dissertation, Leiderdorp, The Netherlands: Open University of Netherlands.

Di Fiore, F., Quax, P., Vanaken, C., Lamotte, W., and Van Reeth, F. 2008. Conveying emotions through facially animated avatars in networked virtual environments. In *Proceedings of Motion In Games (MIG08)*, Springer LNCS 0302-9743, pp. 222-233.

Dimberg, U., and Ohman, A. 1996. Behold the wrath: psychophysiological responses to facial stimuli. *Motivation and Emotion* 20: 149-182.

Ekman, P. 1977. Biological and cultural contributions to body and facial movement. In J. Blacking (ed.), *Anthropology of the Body*, London: Academic Press, pp. 34-84.

Ekman, P. 1993. Facial expression and emotion. *American Psychologist* 48(4): 384-392.

Esuli, A., and Sebastiani, F. 2006. SentiWordNet: a publicly available lexical resource for opinion mining. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*, Genoa, Italy, pp. 417-422.

Fabri, M., Elzouki, S. Y. A., and Moore, D. 2007. Emotionally expressive avatars for chatting, learning and therapeutic Intervention. In J. Jacko (ed.), *Human-Computer Interaction*, Part III, HCII'07, pp. 275–285.

Fahrni, A., and Klenner, M. 2008. Old wine and warm beer: target-specific sentiment analysis of adjectives. In *Proceedings of the Symposium on Affective Language in Human and Machine*, AISB 2008 Convention, pp. 60-63.

Finkel, J. R., Grenager, T., and Manning, C. 2005. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43nd Annual Meeting of the ACL*, pp. 363-370.

Fleiss, J. L. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin* 76(5): 378-382.

Fridlund, A. J. 1992. The behavioral ecology and sociality of human faces. *Review of Personality and Social Psychology* 13: 90-121.

Frijda, N. 1986. *The Emotions*. *Studies in Emotion and Social Interaction*. Cambridge, New York: Cambridge University Press.

Frijda, N. 1994. Emotions are functional, most of the time. In P. Ekman and R. J. Davidson (eds.), *The Nature of Emotion: Fundamental Questions*. New York: Oxford University Press.

Ghazi, D., Inkpen, D., and Szpakowicz, S. 2010. Hierarchical versus flat classification of emotions in text. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pp. 140-146.

Grinter, R. E., and Eldridge, M. A. 2001. y do tngrs luv 2 txt msg? In *Proceedings of the Seventh European Conference on Computer-Supported Cooperative Work*, pp. 219-238.

Harker, L., and Keltner, D. 2001. Expressions of positive emotion in women's college yearbook pictures and their relationship to personality and life outcomes across adulthood. *Journal of Personality and Social Psychology* 80: 112-124.

Hatzivassiloglou, V. and McKeown, K. R. 1997. Predicting the semantic orientation of adjectives. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics and the 8th Conference of the European Chapter of the ACL*, pp. 174-181.

Herring, S. C. 2008. Language and the Internet. In W. Donsbach (ed.), *International Encyclopedia of Communication*, Blackwell Publishers, pp. 2640-2645.

Hoye, L. 1997. *Adverbs and Modality in English*. New York: Addison Wesley Longman Inc.

Hu, Y., Wood, J. F., Smith, V., and Westbrook, N. 2004. Friendships through IM: examining the relationship between Instant Messaging and intimacy. *Computer-Mediated Communication* 10 (1), Article 6.

Izard, C. E. 1971. *The Face of Emotion*. New York: Appleton-Century-Crofts.

Izard, C. E. 1993. Four systems for emotion activation: cognitive and noncognitive processes. *Psychological Review* 100(1): 68-90.

Jarmasz, M., and Szpakowicz, S. 2001. The design and implementation of an electronic lexical knowledge base. In *Proceedings of the 14th Biennial Conference of the Canadian Society for Computational Studies of Intelligence (AI 2001)*, Ottawa, Canada, pp. 325-333.

Kamps, J., and Marx, M. 2002. Words with attitude. In *Proceedings of Belgian-Dutch Conference on Artificial Intelligence*, pp. 449-450.

Kaplan, J. P. 1989. *English Grammar: Principles and Facts*. New Jersey: Prentice-Hall Inc.

Katz, P., Singleton, M., and Wicentowski, R. 2007. Swat-mp: the SemEval-2007 systems for task 5 and task 14. In *Proceedings of SemEval-20*07, Prague, Czech Republic.

Keshtkar, F., and Inkpen, D. 2009. Using sentiment orientation features for mood classification in blogs. In *Proceedings of the IEEE International Conference on Natural Language Processing and Knowledge Engineering (IEEE NLP-K'09)*, Dalian, China.

Kim, S.-M., and Hovy, E. 2004. Determining the sentiment of opinions. In *Proceedings of the International Conference on Computational Linguistics (COLING'04)*, pp. 1367-1373.

Kim, S.-M., and Hovy, E. 2005. Automatic detection of opinion bearing words and sentences. In R. Dale, K.-F. Wong, J. Su, and O. Y. Kwong (eds.), *Proceedings of the International Joint Conference on Natural Language Processing*, Lecture Notes in Computer Science, vol. 3651, Springer, Heidelberg, pp. 61-66.

Kim, S.-M., and Hovy, E. 2006. Identifying and analyzing judgment opinions. In *Proceedings of the Human Language Technology/North American Association of Computational Linguistics Conference*.

Kipper, K., Korhonen, A., Ryant, N., and Palmer, M. 2007. A large-scale classification of English verbs. *Language Resources and Evaluation* 42(1): 21-40.

Kozareva, Z., Navarro, B., Vazquez, S., and Montoyo, A. 2007. Ua-zbsa: a headline emotion classification through web information. In *Proceedings of SemEval-2007*, Prague, Czech Republic.

Ku, L.-W., Huang, T.-H., and Chen, H.-H. 2009. Using morphological and syntactic structures for Chinese opinion analysis. In *Proceedings of the International Conference on Empirical Methods in Natural Language Processing*, pp. 1260-1269.

Lazarus, R. 1994. From appraisal: the minimal cognitive prerequisites of emotion. In P. Ekman and R. J. Davidson (eds.), *The Nature of Emotion: Fundamental Questions*. New York: Oxford University Press.

Leshed, G., and Kaye, J. 2006. Understanding how bloggers feel: recognizing affect in blog posts. In *Extended Abstracts of the International Conference on Computer-Human Interaction*, pp. 1019-1024.

Liu, H., Lieberman, H., and Selker, T. 2003. A model of textual affect sensing using real-world knowledge. In *Proceedings of the International Conference on Intelligent User Interfaces*, pp. 125-132.

Lutz, C. 1988. *Unnatural Emotions*. Chicago: University of Chicago Press.

Lutz, C. 1990. Engendered emotion. In C. Lutz and L. Abu-Lughod (eds.), *Language and the Politics of Emotion*, Cambridge, UK: Cambridge University Press.

Maglogiannis, I., Vouyioukas, D., and Aggelopoulos, C. 2009. Face detection and recognition of natural human emotion using Markov random fields. *Personal and Ubiquitous Computing* 13(1): 95-101.

Martin, J. R. and White, P. R. R. 2005. *The Language of Evaluation: Appraisal in English*. London, UK: Palgrave.

Meys, W. J. 1975. *Compound Adjectives in English and the Ideal Speaker-Listener*. Amsterdam:

North-Holland Publishing Company.

Mihalcea, R., and Liu, H. 2006. A corpus-based approach to finding happiness. In *Proceedings of the AAAI Spring Symposium on Computational Approaches to Weblogs*.

Miller, G. A. 1990. WordNet: an on-line lexical database. *International Journal of Lexicography*, Special Issue 3(4): 235-312.

Miller, G. A. 1999. Nouns in WordNet. In C. Fellbaum (ed.), *WordNet: An Electronic Lexical Database*, Cambridge: The MIT Press, pp. 23-46.

Mishne, G. 2005. Experiments with mood classification in blog posts. In *Proceedings of the First Workshop on Stylistic Analysis of Text for Information Access*.

Moilanen, K., and Pulman, S. 2007. Sentiment composition. In *Proceedings of the Recent Advances in Natural Language Processing International Conference*, pp. 378-382.

Moilanen, K., and Pulman, S. 2008. The good, the bad, and the unknown: morphosyllabic sentiment tagging of unseen words. In *Proceedings of the ACL-08: HLT*, pp. 109-112.

Mulder, M., Nijholt, A., den Uyl, M., and Terpstra, P. 2004. A lexical grammatical implementation of affect. In *Proceedings of the Seventh International Conference on Text, Speech and Dialogue*, Berlin: Springer, pp. 171-178.

Nadeau, D., Sabourin, C., De Koninck, J., Matwin, S., and Turney, P. D. 2006. Automatic dream sentiment analysis. In *Proceedings of the Workshop on Computational Aesthetics at the Twenty-First National Conference on Artificial Intelligence*.

Nakagawa, T., Inui, K., and Kurohashi, S. 2010. Dependency tree-based sentiment classification using CRFs with hidden variables. In *Proceedings of the Human Language Technologies: The 11th Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, pp. 786-794.

Nasukawa, T., and Yi, J. 2003. Sentiment analysis: capturing favorability using natural language processing. In *Proceedings of the 2nd International Conference on Knowledge Capture*, pp. 70-77.

Neviarouskaya, A., Prendinger, H., and Ishizuka, M. 2007a. Analysis of affect expressed through the evolving language of online communication. In *Proceedings of the International Conference on Intelligent User Interfaces*, ACM Press, pp. 278-281.

Neviarouskaya, A., Prendinger, H., and Ishizuka, M. 2007b. Narrowing the social gap among people involved in global dialog: automatic emotion detection in blog posts. In *Proceedings of the International Conference on Weblogs and Social Media*, OMNIPRESS, Boulder, CO, pp. 293-294.

Neviarouskaya, A., Prendinger, H., and Ishizuka, M. 2007c. Textual affect sensing for sociable and expressive online communication. In A. Paiva, R. Prada, and R. W. Picard (eds.), *Proceedings of Second International Conference on Affective Computing and Intelligent Interaction*, Lecture Notes in Computer Science, vol. 4738, Springer, pp. 220-231.

Neviarouskaya, A., Prendinger, H., and Ishizuka, M. 2009. SentiFul: generating a reliable lexicon for sentiment analysis. In *Proceedings of International Conference on Affective Computing and Intelligent Interaction (ACII'09)*, Amsterdam, Netherlands (IEEE), pp. 363-368.

Neviarouskaya, A., Prendinger, H., and Ishizuka, M. 2010a. EmoHeart: conveying emotions in Second Life based on affect sensing from text. *Advances in Human-Computer Interaction*, *Special Issue on Emotion-Aware Natural Interaction*, Open Access Journal, Hindawi, Vol. 2010, Article ID 209801, doi:10.1155/2010/209801, January 2010, 13 pages.

Neviarouskaya, A., Prendinger, H., and Ishizuka, M. 2010b. Recognition of affect, judgment, and appreciation in text. In *Proceedings of the 23rd International Conference on Computational Linguistics (COLING 2010)*, Beijing, China, pp. 806-814.

Neviarouskaya, A., Prendinger, H., and Ishizuka, M. 2010c. User study on AffectIM, an avatar-based Instant Messaging system employing rule-based affect sensing from text. *International Journal of Human-Computer Studies*, Elsevier, 68(7): 432-450.

Neviarouskaya, A., Prendinger, H., and Ishizuka, M. 2011. Affect Analysis Model: novel rule-based approach to affect sensing from text. *International Journal of Natural Language Engineering* 17(1): 95-135, Cambridge University Press, doi:10.1017/S1351324910000239.

Neviarouskaya, A., Tsetserukou, D., Prendinger, H., Kawakami, N., Tachi, S., and Ishizuka, M. 2009. Emerging system for affectively charged interpersonal communication. In *Proceedings of ICROS-SICE International Joint Conference 2009 (ICCAS-SICE 2009)*, Fukuoka, Japan, pp. 3376-3381.

Olveres, J., Billinghurst, M., Savage, J., and Holden, A. 1998. Intelligent, expressive avatars. In *Proceedings of the First Workshop on Embodied Conversational Characters*, pp. 47-55.

Ortony, A., and Turner, T. J. 1990. What's basic about basic emotions? *Psychological Review* 97(3): 315-331.

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

Palmer, F. R. 1986. *Mood and Modality*. Cambridge: Cambridge University Press.

Pang, B., Lee, L., and Vaithyanathan, S. 2002. Thumbs up? Sentiment classification using machine learning techniques. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 79-86.

Peris, R., Gimeno, M. A., Pinazo, D., Ortet, G., Carrero, V., Sanchiz, M., and Ibanez, I. 2002. Online chat rooms: virtual spaces of interaction for socially oriented people. *CyberPsychology and Behavior* 5(1): 43-51.

Persson, P. 2003. ExMS: an animated and avatar-based messaging system for expressive peer communication. In *Proceedings of SIGGROUP Conference*, pp. 31-39.

Picard, R. 1997. *Affective Computing*. Cambridge, MA: The MIT Press.

Plag, I. 2003. *Word-Formation in English*. UK: Cambridge University press.

Planalp, S. 1999. *Communicating Emotion: Social, Moral, and Cultural Processes*. Cambridge, U.K.: Cambridge University Press.

Planalp, S., and Knie, K. 2002. Integrating verbal and nonverbal emotion(al) messages. In S. R. Fussell (ed.), *The Verbal Communication of Emotions: Interdisciplinary Perspectives*, Mahwah, NJ: Lawrence Erlbaum Associates, pp. 55-79.

Polanyi, L., and Zaenen, A. 2004. Contextual valence shifters. In *Working Notes of the AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications*.

Read, J. 2004. Recognising affect in text using Pointwise-Mutual Information. Master thesis, Brighton, UK: University of Sussex.

Read, J. 2005. Using emoticons to reduce dependency in machine learning techniques for sentiment classification. In *Proceedings of the ACL Student Research Workshop at ACL-05*, pp. 43-48.

Reid, E. M. 1991. Electropolis: communication and community on Internet Relay Chat. Honours Thesis, University of Melbourne, Australia.

Reilly, J., and Seibert, L. 2003. Language and emotion. In R. J. Davidson, K. R. Scherer, and H. H. Goldsmith (eds.), *Handbook of Affective Science*, pp. 535-559.

Rigas, G., Katsis, C. D., Ganiatsas, G., and Fotiadis, D. I. 2007. A user independent, biosignal based, emotion recognition method. In *User Modeling 2007*, Heidelberg: Springer Berlin, pp. 314-318.

Riloff, E., Wiebe, J., and Wilson, T. 2003. Learning subjective nouns using extraction pattern bootstrapping. In *Proceedings of the Conference on Natural Language Learning*, pp. 25-32.

Rivera, K., Cooke, N. J., and Bauhs, J. A. 1996. The effects of emotional icons on remote communication. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI 1996)*.

Roseman, I. J., and Smith, C. A. 2001. Appraisal theory. In K. Scherer, A. Schorr, and T. Johnstone (eds.), *Appraisal Processes in Emotion: Theory, Methods, Research*. Oxford, UK: Oxford University Press.

Shaikh, M. A. M., Prendinger, H., and Ishizuka, M. 2007. Assessing sentiment of text by semantic dependency and contextual valence analysis. In A. Paiva, R. Prada, and R. W. Picard (eds.), *Proceedings of Second International Conference on Affective Computing and Intelligent Interaction*, Lecture Notes in Computer Science, vol. 4738, Springer, pp. 191-202.

Shaikh, M. A. M., Prendinger, H., and Ishizuka, M. 2009. A linguistic interpretation of the OCC emotion model for affect sensing from text. In J. Tao and T. Tan (eds.), *Affective Information Processing*, Springer-Verlag London, pp. 45-73.

Slot, K., Cichosz, J., and Bronakowski, L. 2008. Emotion recognition with poincare mapping of voiced-speech segments of utterances. In *Proceedings of Ninth International Conference on Artificial Intelligence and Soft Computing (ICAISC'08)*, Springer, pp. 886-895.

Solomon, R. C. 1973. Emotions and choice. *Review of Metaphysics* 28(1): 20-41.

Strapparava, C., and Mihalcea, R. 2007. SemEval-2007 task 14: Affective Text. In *Proceedings of SemEval-2007*, Prague, Czech Republic, pp. 70-74.

Strapparava, C., and Mihalcea, R. 2008. Learning to identify emotions in text. In *Proceedings of the 2008 ACM Symposium on Applied Computing*, Fortaleza, Brazil, pp. 1556-1560.

Strapparava, C., and Valitutti, A. 2004. WordNet-Affect: an affective extension of WordNet. In *Proceedings of the International Conference on Language Resources and Evaluation*, pp. 1083-1086.

Strapparava, C., Valitutti, A., and Stock, O. 2006. The affective weight of lexicon. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*, Genoa, Italy, pp. 423-426.

Strapparava, C., Valitutti, A., and Stock, O. 2007. Dances with words. In *Proceedings of the International Joint Conference on Artificial Intelligence*, Hyderabad, India, pp. 1719-1724.

Su, F., and Markert, K. 2009. Subjectivity recognition on word senses via semi-supervised mincuts. In *Proceedings of the North American Association of Computational Linguistics (NAACL'09)*.

Subasic, P., and Huettner, A. 2001. Affect analysis of text using fuzzy semantic typing. *IEEE Transactions on Fuzzy Systems* 9(4): 483-496.

Subrahmanian, V. S., and Reforgiato, D. 2008. AVA: adjective-verb-adverb combinations for sentiment analysis. *IEEE Intelligent Systems* 23(4): 43-50.

Taboada, M., and Grieve, J. 2004. Analyzing appraisal automatically. In *Proceedings of AAAI Spring Symposium on Exploring Attitude and Affect in Text*, pp.158-161.

Thurstone, L. L. 1931. The measurement of attitudes. *Journal of Abnormal and Social Psychology* 26: 249-269.

Tsetserukou, D., and Neviarouskaya, A. 2010. iFeel_IM!: augmenting emotions during online communication. *IEEE Computer Graphics and Applications* 30(5): 72-80.

Turney, P. D. 2002. Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. In *Proceedings of the Association for Computational Linguistics 40th Anniversary Meeting*. Association for Computational Linguistics, pp. 417-424.

Turney, P. D., and Littman, M. L. 2003. Measuring praise and criticism: inference of semantic orientation from association. *ACM Transactions on Information Systems* 21(4): 315-346.

Walther, J., and D'Addario, K. 2001. The impacts of emoticons on message interpretation in computer-mediated communication. *Social Science Computer Review* 19(3): 324-347.

Whitelaw, C., Garg, N., and Argamon, S. 2005. Using appraisal groups for sentiment analysis. In *Proceedings of the 14th ACM International Conference on Information and Knowledge Management (CIKM)*, Bremen, Germany, pp. 625-631.

Wiebe, J. 2000. Learning subjective adjectives from corpora. In *Proceedings of 17th Conference of the AAAI*.

Wiebe, J., and Mihalcea, R. 2006. Word sense and subjectivity. In *Proceedings of the Joint Conference of the International Committee on Computational Linguistics and the Association for Computational Linguistics (COLING-ACL)*.

Wilson, T., Wiebe, J., and Hoffmann, P. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, Vancouver, Canada: ACL, pp. 347-354.

Wu, C.-H., Yeh, J.-F., and Chuang, Z.-J. 2009. Emotion perception and recognition from speech. In J. Tao and T. Tan (eds.), *Affective Information Processing*, Springer-Verlag London, pp. 93-110.

Yigit, O. T. 2005. Emoticon usage in task-oriented and socio-emotional contexts in online discussion board. Master Thesis, The Florida State University, US.

Yu, H., and Hatzivassiloglou, V. 2003. Towards answering opinion questions: separating facts from opinions and identifying the polarity of opinion sentences. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 129-136.

Appendix A

Emotional State Gradation within Intensity Levels

	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
ANGER	neutral		annoyed		irritated		indignant		angry		enraged
DISGUST	neutral			dislike				disgust			hate
FEAR	neutral			apprehensive				afraid			terrified
GUILT	neutral		regret		fault		guilt		compunction		self-reproach
INTEREST	neutral		interesting			attractive		involving			desirous
JOY	neutral		cheerful		glad		happy		joyful		elated
SADNESS	neutral		wistful		doleful		unhappy		sad		depressed
SHAME	neutral		confusion		embarrassment		shame		disgrace		dishonour
SURPRISE	neutral		unexpectedness		wonder		astonishment		surprise		amazement

Appendix B

Concise Pseudo-Code of the @AM Algorithm

function ATAManalysis (*Text*, @AM parameters (intensifying coefficients), @AM functionalitites)
Begin
    atam_output = “”
    P = SplitTextIntoParagraphs (*Text*)
    for each P_i in P
    start of loop
        paragraph_output = “”
        S = SplitParagraphIntoSentences (P_i)
        for each S_j in S
        start of loop
            sentence_output = “”
            result_from_symbolic_cue_analysis = ProcessSymbolicCues (S_j)
            if result_from_symbolic_cue_analysis is null
            then
                parser_output = ProcessParserOutput (GetParserResult (S_j))
                word_annotations = GetWordLevelAnnotations (parser_output, @AM parameters (intensifying coefficients), @AM functionalitites)
                C = SplitSentenceIntoClauses (parser_output)
                VFs = {}, SFs = {}, OFs = {}
                for each C_k in C
                start of loop
                    verb_formation = DefineVerbFormation (C_k, parser_output)
                    subject_formation = DefineSubjectFormation (C_k, parser_output)
                    object_formation = DefineObjectFormation (C_k, parser_output)
                    AddVFtoVFs (verb_formation, VFs)
                    AddSFtoSFs (subject_formation, SFs)
                    AddOFtoOFs (object_formation, OFs)
                end of loop
                relation_matrix = RepresentClauseDependencies (parser_output)

```
            VF_phrase_level_results = GetPhraseLevelResults (relation_matrix, VFs, @AM
            functionalitites)
            SF_phrase_level_results = GetPhraseLevelResults (relation_matrix, SFs, @AM
            functionalitites)
            OF_phrase_level_results = GetPhraseLevelResults (relation_matrix, OFs, @AM
            functionalitites)
            clause_level_results = GetClauseLevelResults (relation_matrix,
            VF_phrase_level_results, SF_phrase_level_results, OF_phrase_level_results,
            reinforcement coefficients, @AM functionalitites)
            sentence_level_result = GetSentenceLevelResult (relation_matrix,
            clause_level_results)
            sentence_output = sentence_level_result
        else sentence_output = result_from_symbolic_cue_analysis
        paragraph_output = paragraph_output + sentence_output
    end of loop
    atam_output = atam_output + paragraph_output
  end of loop
  return atam_output
End
```

Appendix C

Emotional States and Relevant Expressive Means

Emotion	Expressive means*
Anger	widely open eyes, fixated; pupils contracted; stare gaze; ajar mouth; teeth usually clenched tightly; rigidity of lips and jaw; lips may be tightly compressed, or may be drawn back to expose teeth
Disgust	narrowed eyes, may be partially closed as result of nose being drawn upward; upper lip drawn up; pressed lips; wrinkled nose; turn of the head to the side quasi avoiding something
Fear	widely open eyes; pupils dilated; raised eyebrows; open mouth with crooked lips; trembling chin
Guilt	downcast or glancing gaze; inner corners of eyebrows may be drawn down; lips drawn in, corners depressed; head lowered
Interest	eyes may be exaggeratedly opened and fixed; lower eyelids may be raised as though to sharpen visual focus; increased pupil size; sparkling gaze; mouth slightly smiling; head is slightly inclined to the side
Joy	'smiling' and bright eyes; genuinely smiling mouth
Sadness	eyelids contracted; partially closed eyes; downturning mouth
Shame	downcast gaze; blushing cheeks; head is lowered
Surprise	widely open eyes; slightly raised upper eyelids and eyebrows; the mouth is opened by the jaw drop; the lips are relaxed

* Data partially taken from (Izard 1971)

Publication List of the Author

Book chapter

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. Recognition of Fine-grained Emotions from Text: An Approach Based on the Compositionality Principle. *Modelling Machine Emotions for Realizing Intelligence: Foundations and Applications*, T. Nishida, L. Jain, and C. Faucher (eds.), Springer Smart Innovation, Systems and Technologies (SIST), Vol. 1, 2010, pp. 179-207.

Journal papers

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. Compositional Approach for Automatic Recognition of Fine-Grained Affect, Judgment, and Appreciation in Text. *Computational Linguistics*, MIT Press, 40 pages. (Submitted on January 7, 2011)

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. SentiFul: A Lexicon for Sentiment Analysis. *IEEE Transactions on Affective Computing*, 2011. (In press)

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. Affect Analysis Model: Novel Rule-based Approach to Affect Sensing from Text. *International Journal of Natural Language Engineering*, Cambridge University Press, Vol. 17, No. 1, 2011, pp. 95-135, doi:10.1017/S1351324910000239.

Dzmitry Tsetserukou and Alena Neviarouskaya. iFeel_IM!: Augmenting Emotions during Online Communication. *IEEE Computer Graphics and Applications*, Vol. 30, No. 5, September/October 2010, pp. 72-80. (Impact factor 1.9)

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. User Study on AffectIM, an Avatar-Based Instant Messaging System Employing Rule-Based Affect Sensing from Text. *International Journal of Human-Computer Studies*, Elsevier, Vol. 68, No. 7, July 2010, pp. 432-450. (Impact factor 2.38)

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. EmoHeart: Conveying Emotions in Second Life Based on Affect Sensing from Text. *Advances in Human-Computer Interaction*, *Special Issue on Emotion-Aware Natural Interaction*, Open Access Journal, Hindawi, Vol. 2010, Article ID 209801, doi:10.1155/2010/209801, January 2010, 13 pages.

Neviarouskaya A.V. Computer-Aided System for Providing Information-Advertising Services. *Scientific-Technical Journal "Informatics"*, No. 2, United Institute of Informatics Problems of the National Academy of Sciences of Belarus, Minsk, Belarus, 2004, pp. 133-136. (In Russian)

Papers presented at international conferences

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. Recognition of Affect, Judgment, and Appreciation in Text. In *Proceedings of the 23rd International Conference on Computational Linguistics (COLING 2010)*, Beijing, China, August 2010, pp. 806-814. (Acceptance rate 19 %) (**COLING 2010 Best Paper Award Finalist**; 4 finalists were selected from 815 submissions)

Dzmitry Tsetserukou and Alena Neviarouskaya. Innovative Real-Time Communication System with Rich Emotional and Haptic Channels. In *Proceedings of International Conference EuroHaptics 2010*, Amsterdam, the Netherlands, July 2010, pp. 306-313. (Acceptance rate 27 %)

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. @AM: Textual Attitude Analysis Model. In *Proceedings of NAACL-HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, ACL, Los Angeles, USA, June 2010, pp. 80-88.

Dzmitry Tsetserukou, Alena Neviarouskaya, Helmut Prendinger, Mitsuru Ishizuka, and Susumu Tachi. iFeel_IM: Innovative Real-Time Communication System with Rich Emotional and Haptic Channels. In *Proceedings of 28th ACM International Conference on Human Factors in Computing Systems (CHI 2010)*, Atlanta, GA, USA, April 2010, pp. 3031-3036. (Acceptance rate 33 %)

Dzmitry Tsetserukou and Alena Neviarouskaya. World's First Wearable Humanoid Robot that Augments Our Emotions. In *Proceedings of the First ACM International Conference on Augmented Human (AH 2010)*, Megeve, France, April 2010, pp. 47-53.

Dzmitry Tsetserukou, Katsunari Sato, Alena Neviarouskaya, Naoki Kawakami, and Susumu Tachi. FlexTorque: Innovative Haptic Interface for Realistic Physical Interaction in Virtual Reality. In *Proceedings of International Conference on Computer Graphics and Interactive Technologies (ACM SIGGRAPH Asia 2009)*, Article No. 10, Yokohama, Japan, December 2009. (Acceptance rate 29 %)

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. Semantically Distinct Verb Classes Involved in Sentiment Analysis. In *Proceedings of the IADIS International Conference on Applied Computing 2009*, Rome, Italy, November 19-21, 2009, pp. 27-34.

Dzmitry Tsetserukou, Alena Neviarouskaya, Helmut Prendinger, Naoki Kawakami, Mitsuru Ishizuka, and Susumu Tachi. iFeel_IM!: Affective Haptic Garment for Emotion Enhancement during Online Communication. In *Proceedings of Asiagraph 2009*, Tokyo, Japan, October 2009, pp. 120-123.

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. SentiFul: Generating a Reliable Lexicon for Sentiment Analysis. In *Proceedings of International Conference on Affective Computing and Intelligent Interaction (ACII'09)*, Amsterdam, Netherlands (IEEE), September 2009, pp. 363-368.

Jina Lee, Alena Neviarouskaya, Helmut Prendinger, and Stacy Marsella. Learning Models of Speaker Head Nods with Affective Information. In *Proceedings of International Conference on Affective Computing and Intelligent Interaction (ACII'09)*, Amsterdam, Netherlands (IEEE), September 2009, pp. 9-15.

Dzmitry Tsetserukou, Alena Neviarouskaya, Helmut Prendinger, Naoki Kawakami, Mitsuru Ishizuka, and Susumu Tachi. Affective Haptics in Emotional Communication. In *Proceedings of International Conference on Affective Computing and Intelligent Interaction (ACII'09)*, Amsterdam, Netherlands (IEEE), September 2009, pp. 181-186.

Alena Neviarouskaya, Dzmitry Tsetserukou, Helmut Prendinger, Naoki Kawakami, Susumu Tachi, and Mitsuru Ishizuka. Emerging System for Affectively Charged Interpersonal Communication. In *Proceedings of ICROS-SICE International Joint Conference 2009 (ICCAS-SICE 2009)*, Fukuoka, Japan, August 18-21, 2009, pp. 3376-3381.

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. EmoHeart: Automation of Expressive Communication of Emotions in Second Life. In *Proceedings of 3rd International Conference on Online Communities and Social Computing, held as part of HCI International 2009*, Springer LNCS 5621, San Diego, California, US, July 2009, pp. 584-592.

Dzmitry Tsetserukou, Alena Neviarouskaya, Helmut Prendinger, Naoki Kawakami, Mitsuru Ishizuka, and Susumu Tachi. iFeel_IM! Emotion Enhancing Garment for Communication in Affect Sensitive Instant Messenger. In *Proceedings of HCI International 2009*, *Human Interface*, *Part 1*, Springer LNCS, San Diego, California, US, July 2009, pp. 628-637.

Dzmitry Tsetserukou, Alena Neviarouskaya, Helmut Prendinger, Naoki Kawakami, Mitsuru Ishizuka, and Susumu Tachi. Enhancing Mediated Interpersonal Communication through Affective Haptics. In *Proceedings of 3rd International Conference on Intelligent Technologies for Interactive Entertainment (INTETAIN'09)*, Amsterdam, Netherlands, Springer LNICST, June 2009, pp. 246-251.

Marijn van Vliet, Alena Neviarouskaya, and Helmut Prendinger. Opinion Elicitation in Second Life. In *Proceedings of 3rd International Conference on Intelligent Technologies for Interactive Entertainment (INTETAIN'09)*, Amsterdam, Netherlands, Springer LNICST, June 2009, pp. 252-257.

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. Compositionality Principle in Recognition of Fine-Grained Emotions from Text. In *Proceedings of Third International Conference on Weblogs and Social Media (ICWSM'09)*, AAAI, San Jose, California, US, May 2009, pp. 278-281.

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. User Study of AffectIM, an Emotionally Intelligent Instant Messaging System. In *Proceedings of 8th International Conference on Intelligent Virtual Agents (IVA'08)*, Springer LNAI 5208, Tokyo, Japan, September 2008, pp. 29-36.

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. An Expressive Avatar for Instant Messaging Endowed with Emotional Intelligence. In *Proceedings of 7th International Conference on Intelligent Virtual Agents (IVA'07)*, Springer LNCS 4722, Paris, France, September 2007, pp. 395-396.

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. Textual Affect Sensing for Sociable and Expressive Online Communication. In *Proceedings of 2nd International Conference on Affective Computing and Intelligent Interaction (ACII'07)*, Springer LNCS 4738, Lisbon, Portugal, September 2007, pp. 220-231.

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. Recognition of Affect Conveyed by Text Messaging in Online Communication. In *Proceedings of 2nd International Conference on Online Communities and Social Computing (OCSC'07)*, held as part of HCI International 2007, Springer LNCS 4564, Beijing, China, July 2007, pp. 141-150.

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. Narrowing the Social Gap among People Involved in Global Dialog: Automatic Emotion Detection in Blog Posts. In *Proceedings of International Conference on Weblogs and Social Media (ICWSM'07)*, Omnipress, Boulder, Colorado, USA, March 2007, pp. 293-294.

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. Analysis of Affect Expressed through the Evolving Language of Online Communication. In *Proceedings of International Conference on Intelligent User Interfaces (IUI'07)*, ACM Press, Hawaii, USA, January 2007, pp. 278-281.

Neviarouskaya A.V. Software Development of Information-Advertising Catalogue Creation. In *Proceedings of International Conference on Modelling and Simulation (MS'2004)*, Minsk, Belarus, April 2004, pp. 373-376.

Papers presented at national conferences

Hisashi Sugawara, Alena Neviarouskaya, and Mitsuru Ishizuka. Affect Extraction from Text in Japanese. In *Proceedings of 23rd Annual Conference of the Japanese Society for Artificial Intelligence*, June 2009. (In Japanese)

Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. Towards Expressive Online Communication: Textual Affect Sensing and Visualization. In *Proceedings of 69th National Conference of Information Processing Society of Japan*, Vol. 4, Tokyo, Japan, March 2007, pp. 237-238.

Neviarouskaya A.V. Automation of the Enterprise Activity in the Field of Information Services. In *Proceedings of the Republican Conference of the Young Scientists of the National Academy of Sciences of Belarus*, Minsk, Belarus, October 2003, pp. 37-40. (In Russian)

Demonstration papers with full paper review process

Dzmitry Tsetserukou, Katsunari Sato, Alena Neviarouskaya, Naoki Kawakami, and Susumu Tachi. FlexTorque: Innovative Haptic Interface for Realistic Physical Interaction in Virtual Reality. In *Proceedings of International Conference on Computer Graphics and Interactive Technologies (ACM SIGGRAPH Asia 2009)*, Emerging Technologies, Yokohama, Japan, December 2009.

Dzmitry Tsetserukou, Alena Neviarouskaya, Helmut Prendinger, Naoki Kawakami, Mitsuru Ishizuka, and Susumu Tachi. iFeel_IM!: Affective Haptic Garment for Emotion Enhancement during Online Communication. In *Proceedings of Asiagraph 2009*, Tokyo, Japan, October 2009, p. 145.

Dzmitry Tsetserukou and Alena Neviarouskaya. Affective Haptic Garment Enhancing Communication in Second Life. In *Proceedings of International Conference on Affective Computing and Intelligent Interaction (ACII'09)*, Amsterdam, Netherlands (IEEE), September 2009, pp. 586-587.

Dzmitry Tsetserukou and Alena Neviarouskaya. Enchancing Mediated Interpersonal Communication through Affective Haptics. In *Proceedings of 3rd International Conference on Intelligent Technologies for Interactive Entertainment (INTETAIN'09)*, Amsterdam, The Netherlands, Springer LNICST, June 2009, p. 25.