Natural Language Engineering **17** (1): 95–135. © Cambridge University Press 2010
doi:10.1017/S1351324910000239

Affect Analysis Model: novel rule-based approach to affect sensing from text

ALENA NEVIAROUSKAYA[1],
HELMUT PRENDINGER[2] and MITSURU ISHIZUKA[1]

[1]*Department of Information and Communication Engineering, University of Tokyo, Engineering Building 2, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8656, Japan*
[2]*National Institute of Informatics, 2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430, Japan*
e-mails: lena@mi.ci.i.u-tokyo.ac.jp, ishizuka@i.u-tokyo.ac.jp, helmut@nii.ac.jp

(*Received 24 September 2009; revised 13 June 2010; accepted 22 July 2010; first published online 16 September 2010*)

Abstract

In this paper, we address the tasks of recognition and interpretation of affect communicated through text messaging in online communication environments. Specifically, we focus on Instant Messaging (IM) or blogs, where people use an informal or garbled style of writing. We introduced a novel rule-based linguistic approach for affect recognition from text. Our Affect Analysis Model (AAM) was designed to deal with not only grammatically and syntactically correct textual input, but also informal messages written in an abbreviated or expressive manner. The proposed rule-based approach processes each sentence in stages, including symbolic cue processing, detection and transformation of abbreviations, sentence parsing and word/phrase/sentence-level analyses. 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 (or neutral). The strength of the resulting emotional state depends on vectors of emotional words, relations among them, tense of the analysed sentence and availability of first person pronouns. The evaluation of the Affect Analysis Model algorithm showed promising results regarding its capability to accurately recognize fine-grained emotions reflected in sentences from diary-like blog posts (averaged accuracy is up to 77 per cent), fairy tales (averaged accuracy is up to 70.2 per cent) and news headlines (our algorithm outperformed eight other systems on several measures).

1 Introduction

People in virtual communities use words on screens to exchange pleasantries and argue, engage in intellectual discourse, conduct commerce, exchange knowledge, share emotional support, make plans, brainstorm, gossip, feud, fall in love, find friends and lose them, play games, flirt, create a little high art and a lot of idle talk.

Rheingold Howard (1993)

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

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

Smith 2001). Social interaction among people is an essential aspect of every society, a strong foundation for the development and self-actualization of a person, and for the establishment of genuine interpersonal relationships and communities. In computer-mediated communication world people tend to interact in a social way: they virtually remain in touch with their relatives and friends to exchange experiences, share opinions and feelings and satisfy their social need of interpersonal communication. Researchers (Peris *et al.* 2002) argue that online communication may stimulate rather than inhibit social relations, and chat users find it a medium for rich, intense and interesting experiences.

In our research we address the tasks of recognition and interpretation of affect communicated through written language in online communication environments, specifically, in Instant Messaging (IM) and diary-like blog posts, where people tend to use an informal and abbreviated style of writing, often accompanied by the use of emoticons (short symbols that resemble facial displays). The informal nature of language (e.g. use of acronyms and emoticons, underuse of apostrophes, etc.) pose a challenge for syntactic and dependency parsers employed as part of a pipeline in a variety of NLP tasks, including affect sensing in online communication media (see recent work of Foster (2010), who investigated the parser performance on discussion forum posts and attempted to handle problematic phenomena by transforming the input sentence or the material, on which the parser was trained). To address this challenge in the task of affect recognition from text, we propose a rule-based linguistic approach designed to deal with not only grammatically and syntactically correct textual input, but also informal messages written in an abbreviated or expressive manner (Neviarouskaya, Prendinger and Ishizuka 2007a).

Our Affect Analysis Model (AAM) performs fine-grained classification of sentences using ten categories: nine emotions ('Anger', 'Disgust', 'Fear', 'Guilt', 'Interest', 'Joy', 'Sadness' ('distress'), 'Shame' and 'Surprise') defined by Izard (1971) and neutral. The proposed rule-based approach processes each sentence in stages, including symbolic cue processing, detection and transformation of abbreviations, sentence parsing and word/phrase/sentence-level analyses. Each analysed sentence is automatically annotated with emotion (or neutral) label and numerical intensity (the strength of emotion).

We evaluate the performance of AAM on two data sets of sentences from diary-like blog posts (our data set, and data set provided by Aman and Szpakowicz 2007), on sentences from fairy tales (Alm's 2008 data set) and news headlines (corpus provided by Strapparava and Mihalcea 2007). The AAM shows promising results regarding its capability to accurately recognize fine-grained emotions reflected in sentences from diary-like blog posts (averaged accuracy is up to 77 per cent), fairy tales (averaged accuracy is up to 70.2 per cent) and news headlines (our algorithm outperformed eight other systems on several measures).

The remainder of the paper is structured as follows. In Section 2, we report on related work. The special features of online communication are described in Section 3. Section 4 discusses the development of Affect database. Our AAM, its evaluation and experimental results are detailed in Sections 5 and 6. Finally, in Section 7, we discuss the performance of the proposed method and conclude the paper.

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

2 Background and related work

Language is about something, does something, and is something in itself; the content and conduct of emotional communication are integrally related.

Donald Brenneis (1990: 114)

Issues of recognition, interpretation 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, facial display, posture and physiological activity. Recently, textual information has been gaining increased attention of researchers interested in studying different kinds of affective phenomena, including sentiment, subjectivity and emotions. In order to analyse affect communicated through written language, researchers in the area of natural language processing 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 *et al.* (1998); Boucouvalas (2003); Liu, Lieberman and Selker (2003); Yu and Hatzivassiloglou (2003); Mulder *et al.* (2004); Read (2004); Kim and Hovy (2005); Neviarouskaya, Prendinger and Ishizuka (2007c); Moilanen and Pulman (2007); Alm (2008); Aman and Szpakowicz (2008); Choi and Cardie (2008);
(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 *et al.* (2006).

2.1 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), affective (Strapparava and Valitutti 2004), appraisal (Argamon *et al.* 2007) and polarity (Hatzivassiloglou and McKeown 1997; Esuli and Sebastiani 2006; Neviarouskaya, Prendinger and Ishizuka 2009) lexicons.

Hatzivassiloglou and McKeown (1997) created a list of 1,336 adjectives manually labelled as either positive or negative. The subjectivity lexicon developed by Wilson *et al.* (2005) is comprised by over 8,000 subjectivity clues annotated by type (strongly subjective/weakly subjective) and prior polarity (positive/negative/both/neutral). 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

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

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 assigning the corresponding label to the synset. The SentiFul database, which contains an extensive list of sentiment-conveying adjectives, adverbs, nouns and verbs annotated by sentiment polarity, polarity scores/intensities and weights, was introduced in Neviarouskaya *et al.* (2009). The Appraisal lexicon developed by Argamon *et al.* (2007) contains adjectives and adverbs annotated by attitude type (affect, judgment, appreciation) and orientation. 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.).

2.2 *Sentiment analysis in text*

2.2.1 *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 per cent) 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 per cent 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 with hidden variables was recently introduced by Nakagawa, Inui and Kurohashi (2010). This approach relies on the lexicon

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

(sentiment polarity expressions and polarity reversing words), dependency parser and a probabilistic model to handle interactions between hidden variables.

2.2.2 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; Subrahmanian and Reforgiato 2008) were proposed.

There is a strong tie between our approach to affect 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 differences between our approaches lie in (1) the levels of classification (fine-grained classes, or nine emotions, and neutral, in our work versus polarity-based classes, or positive/negative/neutral in Moilanen and Pulman 2007); (2) the use of content words playing the role of reversers of a non-neutral polarity of related words (e.g. '*reduce*', '*prevention*') in their algorithm. However, it is not trivial to find pairs of opposite emotions in the case of a fine-grained classification, so such reversers cannot be applied here in a straightforward manner.

2.3 Affect analysis in text

2.3.1 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 emotion in text (Olveres *et al.* 1998; Strapparava *et al.* 2007) or expressed in a multi-modal way (e.g. speech signals along with textual content, Chuang and Wu 2004). 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 (e.g. '*I punched my car radio, and my knuckle is now bleeding*').

2.3.2 Machine-learning approach

With the aim to classify blog sentences by six basic emotions (Ekman 1993), Aman and Szpakowicz (2008) used 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

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

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 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. 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 Latent Semantic Analysis (LSA) technique 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).

The weak points of these methods include: large corpora required for meaningful statistics and good performance; neglect of some prepositions, negation, modal and condition constructions (although the work of Sokolova and Szpakowicz (2005) on analysis and classification of strategies of successful and unsuccessful electronic negotiations is not closely related to our work on fine-grained classification of emotion sentences, we acknowledge here that they used negations and modal verbs in their machine-learning algorithm); disregard of syntactic relations and semantic dependencies in sentences; and long processing time.

2.3.3 Rule-based linguistic approach

Advanced rule-based linguistic approaches targetting textual affect recognition at the sentence level are described in Boucouvalas (2003), Liu *et al.* (2003) and Chaumartin (2007). 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. 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. 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.

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

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 analysing 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 3 for details), and took these into consideration in constructing our AAM relying on rule-based linguistic approach.

3 Features of 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, IM has proven to be one of the most popular online applications. As stated by Shiu and Lenhart (2004), younger Internet users employ IM in greater numbers and more ardently than older generations.

In order to construct a practical and usable 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, which makes text messages difficult to parse.

Successful computer-mediated communication, particularly within the IM environment or diary-like blogs, 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 non-verbal (visual and aural) cues (Hu *et al.* 2004; Derks 2007). In her dissertation, 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

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

humour'; (2) most emoticons are used towards friends, as compared to strangers; (3) more emoticons are used in positive than in negative contexts (spontaneously as well as intentionally) (Derks 2007: 112).

4 Affect database

In this section, we describe Affect database created to support the handling of abbreviated language and the interpretation of affective features of emoticons, abbreviations and words by an automatic emotion recognition system. Our Affect database includes the following tables: Emoticons, Abbreviations, Adjectives, Adverbs, Nouns, Verbs, Interjections and Modifiers.

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. 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). In the resulting intensity estimation for each affect-related entry, variance of data from the annotator mean was taken into consideration. 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.

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 (please see the examples of such dictionaries of emoticons[1] and abbreviations[2]). Only entries that occurred in at least three sources were selected. In this way, we collected 364 emoticons, both of American and Japanese style (e.g. ':">' and '=^_^=' for '*blushing*'), and the 337 most popular acronyms and abbreviations, both emotional and non-emotional (e.g. '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.

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). The measured Kappa coefficients for emoticons and abbreviations are 0.94 and 0.93, respectively, showing strong annotation reliability

[1] Online dictionary of emoticons. http://www.netlingo.com/smileys.php.
[2] Online dictionary of messaging abbreviations. http://www.abbreviations.com/acronyms/CHAT.

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

Table 1. *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	–	–

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

By way of example, some emoticons and abbreviations extracted from the developed database are listed in Table 1.

The next category consists of words conveying affective content. From WordNet-Affect (Strapparava and Valitutti 2004), we have taken 1,627 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 (e.g. '*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.

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 (see Table 2).

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 per cent and 30.0 per cent of overall number of affective words, respectively) whereas the least frequent was 'Guilt' (3.1 per cent). The distribution of affective words with one, two and three emotion labels is 67 per cent, 29 per cent

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

Table 2. *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

and 4 per cent, 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 per cent. 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, Valitutti and Stock 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 3.

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 3) 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 3).

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. 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 *et al.* 1999). In Benamara *et al.* (2007), the

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

Table 3. *Examples of indirect affective 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
Well-being	Satisfaction – [Joy:0.3] Joy – [Joy:0.9]	[Joy:0.9]
Refusal	Sadness – [Sadness:0.9] Anger – [Anger:0.9] Resentment – [Anger:0.6] Disappointment – [Sadness:1.0]	[Anger:0.9; Sadness:1.0]

Table 4. *Examples of modifiers with coefficients of intensity degree strengthening or weakening*

Category (coefficient range)	Modifier	Coefficient
Adverb of affirmation (1.0–2.0)	Certainly	1.2
Adverb of doubt (0.0–1.0)	Arguably	0.5
Strong intensifying adverb (1.0–2.0)	Immensely	1.8
Weak intensifying adverb (0.0–1.0)	Slightly	0.2
Negation (0.0)	Hardly	0.0

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 (see Table 4).

5 The Affect Analysis Model

We have developed a novel algorithm for analysing affect expressed by text messages (Neviarouskaya *et al.* 2007a, 2007c). The proposed algorithm consists of five main stages: (1) symbolic cue analysis; (2) syntactic structure analysis; (3) word-level analysis; (4) phrase-level analysis; (5) sentence-level analysis. The architecture of the AAM is presented in Figure 1.

5.1 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.

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

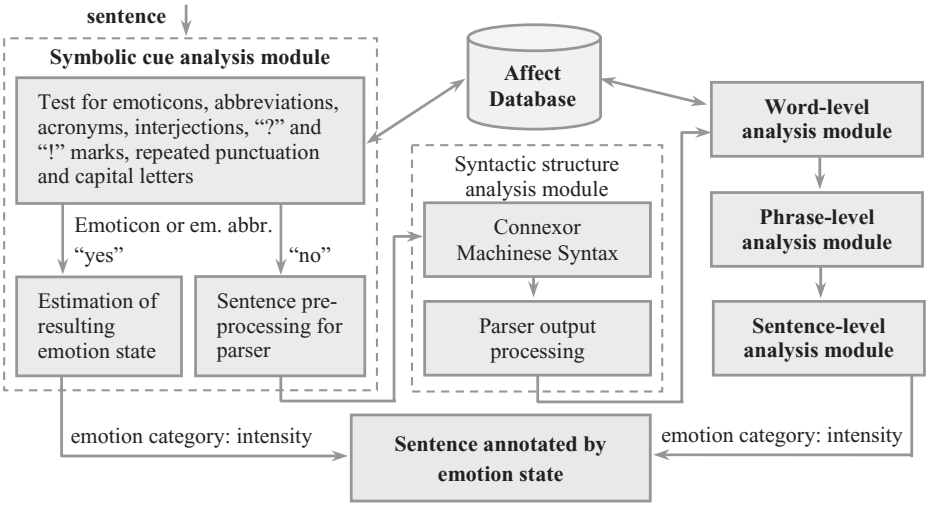


Fig. 1. Architecture of the Affect Analysis Model.

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 the 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.

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

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.

5.2 Syntactic structure analysis

The *second stage* of AAM 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,[3] 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.

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, Verb formation and Object formation, 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.

5.3 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

[3] Connexor Machinese Syntax. http://www.connexor.eu/technology/machinese/machinesesyntax/

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

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. 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.

5.4 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 5.2). 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

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

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 (e.g. 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.

5.5 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.

5.5.1 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

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

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 (e.g. '*To offend the neighbor*' in '*To offend the neighbor is an unfriendly behaviour*') or a full clause, Subject plus Verb–Object formations (e.g. '*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 cherry 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). The analysis of personal pronouns as contributors to the immediacy of disclosure in electronic negotiations (please see details in Sokolova and Szpakowicz 2006) revealed the importance of their use in language patterns for the tactical moves of influence strategies.

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

Table 5. *Coefficients of intensity correction*

Tense	First person pronouns (FPP)	
	Yes	No
Present	1	0.8
	(‘*My vase is broken*’)	(‘*She is annoying*’)
Past	0.8	0.4
	(‘*He made me angry*’)	(‘*It was the most joyous feeling*’)
Future	0.4	0
	(‘*I will enjoy the trip to Egypt*’)	(‘*The game will definitely bring them triumph*’)

	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.

Fig. 2. Example of affect sensing in a simple sentence.

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 (see Table 5).

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*’ (see Figure 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.

5.5.2 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

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

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.

5.5.3 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 5.4 we mentioned that in our AAM 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.

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 (e.g. '*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 and Cairns 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*'.

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

(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.

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 and Cairns 1976). The followings 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.

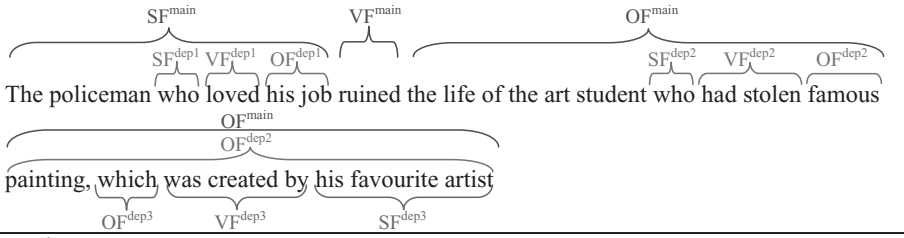
Figure 3 illustrates the steps by way of a complex sentence with multiple embedding of relative clauses: '*The policeman who loved his job ruined the life of the art student who had stolen famous painting, which was created by his favourite artist*'.

In Figure 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; $^{\text{main}}$ and $^{\text{dep}}$ mean belonging to 'main' and 'dependent' clauses, respectively.

5.5.4 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 (e.g. '*Max broke the china cup,*

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press



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 * [0,0,0,0,0,0.8,0.9,0,0,0] = [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.

Fig. 3. Example of affect sensing in a complex sentence with multiple embedding of relative clauses.

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.

6 Evaluation of the Affect Analysis Model algorithm

In order to evaluate the performance of the AAM and to compare out method with related work, we conducted a set of experiments on the data sets created in different domains: diary-like blogs, fairy tales and news headlines.

6.1 Experiments with our collection of sentences extracted from diary-like blogs

6.1.1 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.[4] 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,

[4] Weblog Data Collection. BuzzMetrics, Inc. http://www.nielsenbuzzmetrics.com

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

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[5] from this Weblog Data Collection in order to evaluate the emotion recognition algorithm.

Three independent annotators labelled the sentences with one of nine emotion categories listed in Section 1 (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 gold standards, on which at least two out of three human raters completely agreed: 656 sentences with fine-grained annotations (Fleiss' Kappa coefficient is 0.51) and 692 sentences with polarity-based annotations (Fleiss' Kappa coefficient is 0.6). The percentage distributions of labels across gold standard sentences are as follows (in descending order):

(1) 'Joy' – 28, 'Sadness' – 22, 'Neutral' – 11, 'Anger' – 9, 'Fear' – 7, 'Interest' – 7, 'Surprise' – 7, 'Disgust' – 5, 'Guilt' – 3 and 'Shame' – 1 per cent;
(2) 'Negative' – 50, 'Positive' – 39 and 'Neutral' – 11 per cent.

[5] This annotated data set is freely available upon request.

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

6.1.2 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 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 6 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 6) 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.

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 7 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.

The analysis of errors in the assignment of polarity-based categories for the AAM (see Table 8) revealed that system requires common sense or additional context to deal with 28.5 per cent 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 per cent) as a non-strict one in the experiment with merged labels, where gold standard was based on fine-grained emotion annotations. In nine per cent 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.

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

Table 6. *Accuracy across sentences from blogs in the experiment with words of different parts-of-speech*

Algorithm*	Measure	Fine-grained categories										Merged labels		
		Neutral	Anger	Disgust	Fear	Guilt	Interest	Joy	Sadness	Shame	Surprise	Positive	Negative	Neutral
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, V)	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, V, N)	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, V, N)	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, ADV, V and N refer to adjectives, adverbs, verbs and nouns, respectively.

Table 7. *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 errors resulting from neutralization due to 'cognition-related' words comprise about seven per cent 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.

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 9. As seen in the table, our system achieved satisfactory results in emotion intensity estimation.

6.2 Experiment with the emotion blog data set developed by Aman and Szpakowicz (2007)

6.2.1 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 AAM 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',[6] '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 10.

6.2.2 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'

[6] In the description of this experiment we further use label 'Joy' instead of 'Happiness'.

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

Table 8. *Distribution of errors of AAM in experiment on sentences from blogs with merged labels (gold standard annotations and AAM results are given in the form 'Emotion – Polarity-based label for merged categories' in a last column)*

Error type	# of errors	% of errors	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)
'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	

Table 9. *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 (%)	48.5	32.2	15.9	3.4	0.0

Table 10. *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

to 'Joy', and 'Guilt' and 'Shame' to 'Sadness'. The results of experiments are shown in Table 11, 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 11. *Results of AAM compared to machine learning algorithms proposed by Aman and Szpakowicz (2008)*

Algorithm	Measure	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

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

TOP		emotional EM					neutral N
MID		positive POS		negative NEG			neutral N
ALL	Alm's	Happy	Surprised	Angry-Disgusted	Fearful	Sad	neutral
	our	Joy-Interest J-I	Surprise S	Anger-Disgust A-D	Fear F	Sadness-Guilt-Shame S-G-S	neutral N

Fig. 4. Affect hierarchy and set of labels.

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.

6.3 Experiment with sentences from fairy tales

6.3.1 Data set description

In our next experiment, we wanted to compare the performance of the AAM with Alm's (2008) system that reportedly outperformed Liu *et al.* (2003) system on 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 4).

We ran the experiment on the subset of 1,207 sentences marked by high agreement[7] (indicating that affect labels assigned by four human annotators for the sentence were identical), and a subset of sentences with neutral label. 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 was 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 Equation 1 (taken from Alm 2008),

$$\left[\frac{|HA|}{|Ai|-1}\right], \tag{1}$$

such that 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.

6.3.2 Results

We compared the results of the AAM with Alm (2008) LOOHAsnowtag method (supervised machine learning approach) and two baselines (see Table 12 partially taken from Alm 2008). The ratio of neutrally labelled sentences was considered as first baseline (N-BL), while the second baseline (Freq-BL) was represented by the ratio of the most frequent affect label.

[7] Affect data from fairy tales. http://lrc.cornell.edu/swedish/dataset/affectdata/index.html

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

Table 12. *Accuracy across sentences from fairy tales in high agreement experiment (span of mean accuracy given for LOOHAsnowtag method)*

Level	Data size Total # of sentences (# of N sentences)	Baselines		Individual classification methods			
		N-BL	Freq-BL	LOOHAsnowtag	AAM	AAM_LEX	AAM_LEX_WI
ALL	1448 (241)	17	31 (J-I)	69–70	68.2	70.1	70.2
MID	1810 (603)	33	40 (NEG)	69–73	73.3	74.6	75.7
TOP	2414 (1207)	50	50 (any)	79	77.6	78.5	79.9

As seen from Table 12, AAM resulted in a similar accuracy level as the LOOHAsnowtag method, outperforming 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). The two classes of errors of AAM method (first-class errors in the case of high-agreement affect sentences at the ALL level; second-class errors in the case of neutral sentences at the TOP level) are analysed in Table 13.

The largest group of all AAM errors (52.7 per cent of first-class errors; 67.1 per cent of second-class errors) occurred in the cases where additional common sense knowledge or context are required for correct interpretation of conveyed emotion. The ambiguity of word senses was responsible for about 20 per cent of second-class errors, highlighting the importance and necessity of word sense disambiguation. About 14 per cent of first-class errors were so-called 'conflict' errors, meaning that the correct emotion was in the final vector of AAM, but it was not dominant, taking into account its intensity. The errors carrying grammatical character and caused by the parser include about 11 per cent of first-class errors and 4.5 per cent of second-class errors. The rules for negations, neutralizations and combinations of the SF, VF and OF together caused about 14.5 per cent of the first-class errors.

As our approach relies on the affect lexicon, and its recall depends on the coverage of the lexicon used (about seven per cent of errors in high agreement affect sentences at the ALL level were due to the lack of affect entries in database), we improved the performance of our system (see results of AAM_LEX in Table 12, where LEX means 'lexicon') by expanding the Affect database with some sample words from experimental data. The decision for further improvement (see results of AAM_LEX_WI in Table 12, where WI means 'without Interest' or 'neutral label instead of Interest') was based on the fact that about 8.5 per cent of errors in neutral sentences at the TOP level (se Table 13) were comprised of the cases where AAM system resulted in 'Interest' emotion that was not considered as a label in the Alm's (2008) set of affect labels.

6.4 Experiment with news headlines

In addition to the experiments on sentences from diary-like blog posts and fairy tales, we have also evaluated AAM algorithm on news headlines.

6.4.1 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 1,000 news headlines[8] 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. 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

[8] Data set is available at the SemEval-2007 web site. http://nlp.cs.swarthmore.edu/semeval

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

Table 13. *Distribution of errors of AAM in high agreement experiment on sentences from fairy tales*

Error type	# of errors	% of errors	Sample sentence (gold standard; AAM result)
First-class errors in high agreement affect sentences at the ALL level			
Common sense or additional context	218	52.7	When the seven kids saw that, they came running to the spot and cried aloud: 'The wolf is dead!' (Joy; Sadness) "Ah!" said the father, "what fears we have had for you!" (Joy; Fear)
Conflict (correct emotion is in the final vector of AAM, but is not dominant)	57	13.8	He started back, quite bewildered with the fright which the goloshes of Fortune had caused him. (Surprise; Fear)
Parser	45	10.9	So the father gave him his blessing, and with great sorrow took leave of him. (Sadness; Joy)
Lexicon	29	7.0	After waiting awhile, she went to Mother Holle and said, "I am so homesick, that I cannot stay with you any longer, for although I am so happy here, I must return to my own people." (Sadness; neutral)
Neutralization due to 'would', 'can', 'could', 'may', 'might'	17	4.1	If I did not take pity on you and save you, you would be lost. (Sadness; neutral) Can such happiness be imagined? (Joy; neutral)
Negation neutralization instead of negation reversal	11	2.7	But Tiny was not at all pleased; for she did not like the tiresome mole. (Sadness; neutral)
'But'	8	1.9	The bean thanked him most prettily, but as the tailor used black thread, all beans since then have a black seam. (Joy; neutral) At this, too, the youngest sister was terribly frightened, but the eldest always silenced her. (Fear; neutral)
Neutralization due to negation, 'could not', 'cannot'	8	1.9	Then there was no end to the rage and disappointment of Tom Thumb and Hunca Munca. (Anger; neutral) "Summer cannot show a more beautiful sight," she exclaimed, while her eyes sparkled. (Joy; neutral)
Neutralization due to 'think', 'believe', 'know'	7	1.7	I can trust you, for I believe that I do love you. (Joy; neutral)

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

Rules for combining the SF, VF, OF	5	1.2	It became a splendid flower-garden to the sick boy, and his little treasure upon earth. (Joy; sadness)
Neutralization due to 'when', 'as if'	4	1.0	The branches of the old willow-tree rustled in the wind, and large water-drops fell from his green leaves as if the old willow were weeping. (Sadness; neutral)
Other	5	1.2	
Total, including double errors	414	100	
Second-class errors in neutral sentences at the TOP level			
Common sense or additional context	222	67.1	My father grieved when I was gone. (neutral; Sadness)
Sense ambiguity	66	19.9	It was very cold; but the little creature did not really feel it, till the light in the garret went out, and the tones of music died away. (neutral; Sadness) They had all many things to relate, especially the shirt collar, who was a terrible boaster. (neutral; Fear)
'Interest' emotion (not considered in Alm's set of labels)	28	8.5	He has great possessions, but still he longs for more – everything must bow before the mighty Olaf Glob. (neutral; Interest)
Parser	15	4.5	It could easily be seen that she was a very lovely girl, and as yet she was not engaged. (neutral; Joy) Something of this kind Anthony felt when Molly talked to him of old times. (neutral; Joy)
Total, including double errors	331	100	

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 per cent of headlines; two emotions – 19.1 per cent of headlines; three emotions – 17.7 per cent of headlines; four emotions – 21.6 per cent of headlines; five emotions – 27.2 per cent of headlines; and six emotions – 13.3 per cent of headlines. Therefore, 62.1 per cent of all headlines were annotated by at least four emotions.

6.4.2 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 emotion vector resulting from AAM as the overall annotation for each headline; (2) reduced the number of our labels to six by mapping 'Interest' to 'Joy', and 'Guilt' and 'Shame' 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 AAM and other systems (reported in Strapparava and Mihalcea 2008) are shown in Table 14, where (1) AAM is our Affect 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 labelled 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 *et al.* 2007) is a supervised system, which is based on unigram model; (8) UA (Kozareva *et al.* 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).

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

Table 14. *Results of AAM 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)				
AAM	**32.88**	12.97	44.86	21.25	**43.57**	6.40	26.99
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)				
AAM	25.00/19.05	0.00/0.00	55.39/39.13	47.50/16.81	39.13/34.62	**50.00**/7.14	36.17/19.46
	21.62	–	**45.86**	24.84	**36.74**	12.50	**28.31**
WN-A	**33.33**/3.33	0.00/0.00	**100.00**/1.69	50.00/0.56	33.33/3.67	13.04/4.68	**38.28**/1.54
	6.06	–	3.33	1.10	6.61	6.90	4.00
LSA SW	6.28/63.33	**2.41**/70.59	12.93/**96.61**	17.81/47.22	13.13/55.05	6.73/67.19	9.88/66.72
	11.43	**4.68**	22.80	25.88	21.20	12.23	16.37
LSA ES	7.29/86.67	1.53/64.71	12.44/94.92	19.37/72.22	14.35/58.71	7.23/89.06	9.20/77.71
	13.45	3.00	22.00	30.55	23.06	13.38	13.38
LSA AEW	6.20/**88.33**	1.98/**94.12**	12.55/86.44	18.60/**90.00**	11.69/**87.16**	7.62/**95.31**	9.77/**90.22**
	11.58	3.87	21.91	30.83	20.61	14.10	17.57
NB BLOG	13.68/21.67	0.00/0.00	16.67/3.39	22.71/59.44	20.87/22.02	8.33/1.56	12.04/18.01
	16.77	–	5.63	**32.87**	21.43	2.63	13.22

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

Table 14. *Continued*

System	Emotion labels						Average result
	Anger	Disgust	Fear	Joy	Sadness	Surprise	
SWAT	12.00/5.00	0.00/0.00	25.00/14.40	35.41/9.44	32.50/11.92	11.86/10.93	19.46/8.61
	7.06	–	18.27	14.91	17.44	11.78	11.57
UA	12.74/21.60	0.00/0.00	16.23/26.27	40.00/2.22	25.00/0.91	13.70/16.56	17.94/11.26
	16.03	–	20.06	4.21	1.76	**15.00**	9.51
UPAR7	16.67/1.66	0.00/0.00	33.33/2.54	**54.54**/6.66	**48.97**/22.02	12.12/1.25	27.60/5.68
	3.02	–	4.72	11.87	30.38	2.27	8.71

In the fine-grained evaluation, our system achieved the highest results in recognition of 'Anger' and 'Sadness' emotions, while SWAT was more successful in case of 'Disgust' and 'Joy', and UPAR7 in case of 'Fear' and 'Surprise'. In terms of averages of Pearson's correlation coefficients for all emotions, UPAR7 showed the best performance (28.38), followed by our AAM system (26.99) and SWAT (25.41). These results indicate that our AAM 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 analysing the sentence in consecutive stages outputs unified emotion 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 AAM, one have to sum emotion vectors from 'Word-level analysis' stage (see Section 5.3 for details) and ignore 'Phrase-level analysis' 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 AAM scores and gold standard; however, there would not be much intelligence.

In the coarse-grained evaluation, AAM ensured the best F-scores for 'Anger', 'Fear' and 'Sadness', as well as in case of average results, while the highest F-scores for 'Disgust', 'Joy' and 'Surprise' were achieved by LSA SW, NB BLOG and UA systems, correspondingly.

In summary, the evaluation of the AAM algorithm showed promising results regarding its capability to recognize affective information in sentences from different domains: blogs, fairy tales and news titles. However, the current main limitations of the developed affect recognition module must be noted, namely: strong dependency on the resource of lexicon (Affect database) and the commercially available syntactic parser; no disambiguation of word senses; and disregard of contextual information.

7 Conclusions

This paper introduced a novel rule-based linguistic approach for affect recognition from text. Typically, researchers in the sentiment analysis field deal with grammatically and syntactically correct textual input. In contrast, our analysis of affect expressed through written language is inspired by the evolving language, style and specifics of IM conversations and diary-like blog posts. For textual input processing, our AAM copes with not only correctly written text, but also informal messages written in an abbreviated or expressive manner. In order to support the handling of abbreviated language and the interpretation of affective features of emoticons, abbreviations and words, a special Affect database was created.

The proposed rule-based algorithm for affect sensing from text processes each sentence in stages, including symbolic cue processing, detection and transformation

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

of abbreviations and acronyms, sentence parsing and word/phrase/sentence-level analyses. 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 (or neutral). The strength of the resulting emotional state depends on emotion vectors of words, relations among them, tense of the analysed sentence and availability of first person pronouns.

The salient features of the AAM 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 dependences 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 (gold standard where at least two annotators agreed), averaged accuracy was 72.6 per cent for fine-grained (nine categories and neutral) emotion classification and 81.6 per cent for polarity-based merged categories; (2) on data set provided by Aman and Szpakowicz (2008), averaged accuracy was 77.0 per cent 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. On the sentences extracted from fairy tales (Alm 2008), averaged accuracy of AAM was 70.2 per cent for fine-grained

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

emotion classification and 75.7 per cent for polarity-based merged categories (these results are slightly better than those reported in related work). Comparing the performance of the AAM 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.

Currently, the main limitations of the developed affect recognition module are: strong dependency on the lexicon resource (Affect database) and the commercially available syntactic parser; no disambiguation of word meanings; disregard of conversation history; and inability to recognize and process misspelled words in a sentence.

In our future study, we will investigate those issues and explore the possibilities to overcome the current limitations of the system. For example, the Affect database can be expanded automatically by considering (1) direct synonymy (e.g. '*congratulate*' – '*felicitate*') and antonymy (e.g. '*brave*' – '*faint-hearted*') relations and hypernym–hyponym (e.g. '*success*' – '*winning*') relations in WordNet; (2) compounding using known emotion-carrying base components (e.g. '*well-wishing*', '*ill-conditioned*', '*terror-haunted*'); (3) morphological modifications (e.g. '*harmony*'+'*-ous*' = >'*harmonious*', '*dis-*'+ '*honest*' = >'*dishonest*').

As our system relies on the lexicon and the language of online conversations is 'evolving', we are also planning to implement a procedure for the automatic updating of the Affect database. In addition, we will thoroughly examine different cases of the use of modals and condition prepositions in sentences, so as to be able to automatically distinguish emotional cases from neutral ones (in contrast to simple neutralization applied in current work). With respect to the rules for composition of emotion vectors of terms comprising phrases or clauses, we believe the approach aiming at learning rules from corpora would be useful.

To enrich the user's experience in online communication, to make it enjoyable, exciting and fun, we implemented a web-based IM application, AffectIM and endowed it with emotional intelligence by integrating the developed AAM. The findings of a twenty-person study conducted with our AffectIM system are described in Neviarouskaya, Prendinger and Ishizuka (2010).

Acknowledgments

We would like to acknowledge and thank Dr. Alessandro Valitutti and Dr. Diego Reforgiato for their kind help during the Affect database creation. We wish also to express our gratitude to Dr. Dzmitry Tsetserukou, Dr. Shaikh Mostafa Al Masum, Manuel M. Martinez, Zoya Verzhbitskaya, Hutchatai Chanlekha and Nararat Ruangchaijatupon, who have contributed to annotations of Affect database entries and sentences for their efforts and time. This research was partly supported by a JSPS Encouragement of Young Scientists Grant (FY2005-FY2007), an NII Joint Research Grant with the University of Tokyo (FY2007) and an NII Grand Challenge Grant (FY2008-FY2009).

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

References

Alm, C. O. 2008. *Affect in Text and Speech*. PhD Dissertation, Urbana, IL: 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, pp. 196–205. Lecture Notes in Computer Science, vol. 4629. Springer.

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 3rd Language and Technology Conference*.

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.

Brenneis, D. 1990. Shared and solitary sentiments. In C. Lutz and L. Abu-Lughod (eds.), *Language and the Politics of Emotion*. Cambridge, UK: Cambridge University Press.

Cairns, H. S., and Cairns, C. E. 1976. *Psycholinguistics: A Cognitive View of Language*. New York: Holt, Rinehart, and Winston.

Chaumartin, F. 2007. Upar7: A knowledge-based system for headline sentiment tagging. In *Proceedings of SemEval-2007*, Prague, Czech Republic.

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 Linguistic and Chinese Language Processing* **9**(2): 45–62.

Davitz, J. R. 1969. *The Language of Emotions*. London: Academic Press.

Derks, D. 2007. Exploring the missing wink: emoticons in cyberspace. Dissertation, Leiderdorp, The Netherlands: Open University of Netherlands.

Ekman, P. 1977. Biological and cultural contributions to body and facial movement. In: J. Blacking (ed.), *Anthropology of the Body*, pp. 34–84. London: Academic Press.

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*, pp. 417–422.

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

Fleiss, J. L. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin* **76**(5): 378–382.

Foster, J. 2010. "cba to check the spelling": Investigating parser performance on discussion forum posts. 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. 381–384.

Frijda, N. 1986. *The Emotions. Studies in Emotion and Social Interaction.* Cambridge, New York: Cambridge University Press.

Grinter, R.E., 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.

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*, pp. 2640–2645. Blackwell Publishers.

Howard, R. 1993. *The Virtual Community: Homesteading on the Electronic Frontier*. Menlo Park, CA: Addison-Wesley Publishing Co.

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).

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.

Katz, P., Singleton, M., and Wicentowski, R. 2007. Swat-mp: the SemEval-2007 systems for task 5 and task 14. In *Proceedings of SemEval-2007*, Prague, Czech Republic.

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*, pp. 61–66, Lecture Notes in Computer Science, vol. 3651. Springer, Heidelberg.

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.*

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.

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. 1990. Engendered emotion. In C. Lutz and L. Abu-Lughod (eds.), *Language and the Politics of Emotion.* Cambridge, UK: Cambridge University Press.

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

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. *Lexicography (Special Issue)* **3**(4): 235–312.

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.

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 Second 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*, pp. 220–231, Lecture Notes in Computer Science, vol. 4738. Springer.

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*, Amsterdam, The Netherlands, pp. 363–368. IEEE.

Neviarouskaya, A., Prendinger, H., and Ishizuka, M. 2010. User study on AffectIM, an avatar-based Instant Messaging system employing rule-based affect sensing from text. *Elsevier International Journal of Human-Computer Studies* **68**(7): 432–450.

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., Clore, G., and Collins, A. 1988. *The Cognitive Structure of Emotions.* Cambridge, UK: 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.

Picard, R. 1997. *Affective Computing.* Cambridge, MA: The MIT 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*, pp. 55–79. Mahwah, NJ: Lawrence Erlbaum Associates.

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press

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.

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.

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.

Shiu, E., and Lenhart, A. 2004. *How Americans use Instant Messaging*. Washington: Pew Internet and American Life Project.

Sokolova, M., and Szpakowicz, S. 2005. Analysis and classification of strategies in electronic negotiations. In B. Kegl and G. Lapalme (eds.), *Proceedings of the 18th Canadian Conference on Artificial Intelligence*, pp. 145–157, Lecture Notes in Artificial Intelligence, vol. 3501. Springer.

Sokolova, M., and Szpakowicz, S. 2006. Language patterns in the learning of strategies from negotiation texts. In L. Lamontagne and M. Marchand (eds.), *Proceedings of the 19th Canadian Conference on Artificial Intelligence*, pp. 288–299, Lecture Notes in Artificial Intelligence, vol. 4013. Springer.

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*, pp. 70–74, Prague, Czech Republic.

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*, pp. 423–426, Genoa, Italy.

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.

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.

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.

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.

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.

https://doi.org/10.1017/S1351324910000239 Published online by Cambridge University Press