

@AM: Textual Attitude Analysis Model

Alena Neviarouskaya University of Tokyo 7-3-1 Hongo, Bunkyo-ku Tokyo 113-8656, Japan lena@mi.ci.i.u-tokyo.ac.jp	Helmut Prendinger Nat. Institute of Informatics 2-1-2 Hitotsubashi Chiyoda Tokyo 101-8430, Japan helmut@nii.ac.jp	Mitsuru Ishizuka University of Tokyo 7-3-1 Hongo, Bunkyo-ku Tokyo 113-8656, Japan ishizuka@i.u-tokyo.ac.jp
--	--	---

Abstract

The automatic analysis and classification of text using fine-grained attitude labels is the main task we address in our research. The developed @AM system relies on compositionality principle and a novel approach based on the rules elaborated for semantically distinct verb classes. The evaluation of our method on 1000 sentences, that describe personal experiences, showed promising results: average accuracy on fine-grained level was 62%, on middle level – 71%, and on top level – 88%.

1 Introduction and Related Work

With rapidly growing online sources aimed at encouraging and stimulating people’s discussions concerning personal, public or social issues (news, blogs, discussion forums, etc.), there is a great need in development of a computational tool for the analysis of people’s attitudes. According to the Appraisal Theory (Martin and White, 2005), attitude types define the specifics of appraisal being expressed: affect (personal emotional state), judgement (social or ethical appraisal of other’s behaviour), and appreciation (evaluation of phenomena).

To analyse contextual sentiment (polarity) of a phrase or a sentence, rule-based approaches (Nasukawa and Yi, 2003; Mulder et al., 2004; Moilanen and Pulman, 2007; Subrahmanian and Reforgiato, 2008), a machine-learning method using not only lexical but also syntactic features

(Wilson et al., 2005), and a model of integration of machine learning approach with compositional semantics (Choi and Cardie, 2008) were proposed.

With the aim to recognize fine-grained emotions from text on the level of distinct sentences, researchers have employed a keyword spotting technique (Olveres et al., 1998; Chuang and Wu, 2004; Strapparava et al., 2007), a technique calculating emotion scores using Pointwise Mutual Information (PMI) (Kozareva et al., 2007), an approach inspired by common-sense knowledge (Liu et al., 2003), rule-based linguistic approaches (Boucouvalas, 2003; Chaumartin, 2007), machine-learning methods (Alm, 2008; Aman and Szpakowicz, 2008; Strapparava and Mihalcea, 2008), and an ensemble based multi-label classification technique (Bhowmick et al., 2009).

Early attempts to focus on distinct attitude types in the task of attitude analysis were made by Taboada and Grieve (2004), who determined a potential value of adjectives for affect, judgement and appreciation by calculating the PMI with the pronoun-copular pairs ‘*I was (affect)*’, ‘*He was (judgement)*’, and ‘*It was (appreciation)*’, and Whitelaw et al. (2005), who used machine learning technique (SVM) with fine-grained semantic distinctions in features (attitude type, orientation) in combination with “bag of words” to classify movie reviews. However, the concentration only on adjectives, that express appraisal, and their modifiers, greatly narrows the potential of the Whitelaw et al.’s (2005) approach.

In this paper we introduce our system @AM (ATtitude Analysis Model), which (1) classifies

sentences according to the fine-grained attitude labels (nine affect categories (Izard, 1971): ‘anger’, ‘disgust’, ‘fear’, ‘guilt’, ‘interest’, ‘joy’, ‘sadness’, ‘shame’, ‘surprise’; four polarity labels for judgment and appreciation: ‘POS jud’, ‘NEG jud’, ‘POS app’, ‘NEG app’; and ‘neutral’); (2) assigns the strength of the attitude; and (3) determines the level of confidence, with which the attitude is expressed. @AM relies on compositionality principle and a novel approach based on the rules elaborated for semantically distinct verb classes.

2 Lexicon for Attitude Analysis

We built the lexicon for attitude analysis that includes: (1) attitude-conveying terms; (2) modifiers; (3) “functional” words; and (4) modal operators.

2.1 The Core of Lexicon

As a core of lexicon for attitude analysis, we employ Affect database and extended version of SentiFul database developed by Neviarouskaya et al. (2009). The affective features of each emotion-related word are encoded using nine emotion labels (‘anger’, ‘disgust’, ‘fear’, ‘guilt’, ‘interest’, ‘joy’, ‘sadness’, ‘shame’, and ‘surprise’) and corresponding emotion intensities that range from 0.0 to 1.0. The original version of SentiFul database, which contains sentiment-conveying adjectives, adverbs, nouns, and verbs annotated by sentiment polarity, polarity scores and weights, was manually extended using attitude labels. Some examples of annotated attitude-conveying words are listed in Table 1. It is important to note here that some words could express different attitude types (affect, judgment, appreciation) depending on context; such lexical entries were annotated by all possible categories.

POS	Word	Category	Intensity
adjective	<i>honorable</i>	POS jud	0.3
		NEG aff (sadness)	0.5
	<i>unfriendly</i>	NEG jud	0.5
		NEG app	0.5
adverb	<i>gleefully</i>	POS aff (joy)	0.9
noun	<i>abnormality</i>	NEG app	0.25
verb	<i>frighten</i>	NEG aff (fear)	0.8
		POS aff (interest)	1.0
	<i>desire</i>	POS aff (joy)	0.5

Table 1. Examples of attitude-conveying words and their annotations.

2.2 Modifiers and Functional Words

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

We collected 138 modifiers that have an impact on contextual attitude features of related words, phrases, or clauses. They include:

1. Adverbs of degree (e.g., ‘*significantly*’, ‘*slightly*’ etc.) and adverbs of affirmation (e.g., ‘*absolutely*’, ‘*seemingly*’) that have an influence on the strength of attitude of the related words.

2. Negation words (e.g., ‘*never*’, ‘*nothing*’ etc.) that reverse the polarity of related statement.

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

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

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

Adverbs of degree and adverbs of affirmation affect on related verbs, adjectives, or another adverb. Two annotators gave coefficients for intensity degree strengthening or weakening (from 0.0 to 2.0) to each of 112 collected adverbs, and the result was averaged (e.g., $\text{coeff}(\textit{perfectly}) = 1.9$, $\text{coeff}(\textit{slightly}) = 0.2$).

We distinguish two types of “functional” words that influence contextual attitude and its strength:

1. *Intensifying* adjectives (e.g., ‘*rising*’, ‘*rapidly-growing*’), nouns (e.g., ‘*increase*’, ‘*up-tick*’), and verbs (e.g., ‘*to grow*’, ‘*to rocket*’), which increase the strength of attitude of related words.

2. *Reversing* adjectives (e.g., ‘*reduced*’), nouns (e.g., ‘*termination*’, ‘*reduction*’), and verbs (e.g., ‘*to decrease*’, ‘*to limit*’, ‘*to diminish*’), which reverse the prior polarity of related words.

2.3 Modal Operators

Consideration of the modal operators in the tasks of opinion mining and attitude analysis is very important, as they indicate a degree of person’s belief in the truth of the proposition, which is subjective in nature. Modal expressions point to likelihood and clearly involve the speaker’s judgment (Hoye, 1997). Modals are distinguished by the *confidence level*.

We collected modal operators of two categories:

1. Modal verbs (13 verbs).
2. Modal adverbs (61 adverbs).

Three human annotators assigned the *confidence level*, which ranges from 0.0 to 1.0, to each modal verb and adverb; these ratings were averaged (e.g., $\text{conf}(\text{'vaguely'}) = 0.17$, $\text{conf}(\text{'may'}) = 0.27$, $\text{conf}(\text{'arguably'}) = 0.63$, $\text{conf}(\text{'would'}) = 0.8$, $\text{conf}(\text{'veritably'}) = 1.0$).

3 Compositionality Principle

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

Polarity reversal means that phrase or statement containing attitude-conveying term/phrase with prior positive polarity becomes negative, and vice versa. The rule of *polarity reversal* is applied in three cases: (1) negation word-modifier in relation with attitude-conveying statement (e.g., 'never' & POS('succeed') => NEG('never succeed')); (2) adverb of doubt in relation with attitude-conveying statement (e.g., 'scarcely' & POS('relax') => NEG('scarcely relax')); (3) functional word of *reversing* type in relation with attitude-conveying statement (e.g., adjective 'reduced' & POS('enthusiasm') => NEG('reduced enthusiasm')). In the case of judgment and appreciation, the use of *polarity reversal* rule is straightforward ('POS jud' <=> 'NEG jud', 'POS app' <=> 'NEG app'). However, it is not trivial to find pairs of opposite emotions in the case of a fine-grained classification, except for 'joy' and 'sadness'. Therefore, we assume that (1) opposite emotion for three positive emotions, such as 'interest', 'joy', and 'surprise', is 'sadness' ('POS aff' => 'sadness'); and (2) opposite emotion for six negative emotions, such as 'anger', 'disgust', 'fear', 'guilt', 'sadness', and 'shame', is 'joy' ('NEG aff' => 'joy').

The rules of *aggregation (fusion)* are as follows: (1) if polarities of attitude-conveying terms in adjective-noun, noun-noun, adverb-adjective, adverb-

verb phrases have opposite directions, mixed polarity with dominant polarity of a descriptive term is assigned to the phrase (e.g., POS('beautiful') & NEG('fight') => POS-neg('beautiful fight'); NEG('shamelessly') & POS('celebrate') => NEG-pos('shamelessly celebrate')); otherwise (2) the resulting polarity is based on the equal polarities of terms, and the strength of attitude is measured as a maximum between polarity scores (intensities) of terms ($\text{max}(\text{score1}, \text{score2})$).

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

The rules of *domination* are as follows: (1) if polarities of verb (this rule is applied only for certain classes of verbs) and object in a clause have opposite directions, the polarity of verb is prevailing (e.g., NEG('to deceive') & POS('hopes') => NEG('to deceive hopes')); (2) if compound sentence joints clauses using coordinate connector 'but', the attitude features of a clause following after the connector are dominant (e.g., 'NEG(It was hard to climb a mountain all night long), but POS(a magnificent view rewarded the traveler at the morning).' => POS(whole sentence)).

The rule of *neutralization* is applied when preposition-modifier or condition operator relate to the attitude-conveying statement (e.g., 'despite' & NEG('worries') => NEUT('despite worries')).

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

1. adverb of degree or affirmation relates to attitude-conveying term (e.g., $\text{Pos_score}(\text{'extremely happy'}) > \text{Pos_score}(\text{'happy'})$);

2. adjective or adverb is used in a comparative or superlative form (e.g., $\text{Neg_score}(\text{'sad'}) < \text{Neg_score}(\text{'sadder'}) < \text{Neg_score}(\text{'saddest'})$).

Our method is capable of processing sentences of different complexity, including simple, compound, complex (with complement and relative clauses),

and complex-compound sentences. To understand how words and concepts relate to each other in a sentence, we employ Connexor Machine Syntax parser (<http://www.connexor.eu/>) that returns lemmas, parts of speech, dependency functions, syntactic function tags, and morphological tags. When handling the parser output, we represent the sentence as a set of primitive clauses. 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. For the processing of complex or compound sentences, we build a so-called “relation matrix”, which contains information about dependences (e.g., coordination, subordination, condition, contingency, etc.) between different clauses in a sentence.

The annotations of words are taken from our attitude-conveying lexicon. The decision on most appropriate label, in case of words with multiple annotations (e.g., word ‘unfriendly’ in Table 1), is made based on (1) the analysis of morphological tags of nominal heads and their premodifiers in the sentence (e.g., first person pronoun, third person pronoun, demonstrative pronoun, nominative or genitive noun, etc.); (2) the analysis of the sequence of hypernymic semantic relations of a particular noun in WordNet (Miller, 1990), which allows to determine its conceptual domain (e.g., “person, human being”, “artifact”, “event”, etc.). For ex., ‘I feel highly unfriendly attitude towards me’ conveys ‘NEG aff’ (‘sadness’), while ‘Shop assistant’s behavior was really unfriendly’ and ‘Plastic bags are environment unfriendly’ express ‘NEG jud’ and ‘NEG app’, correspondingly.

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

4 Consideration of the Semantics of Verbs

All sentences must include a verb, because the verb tells us what action the subject is performing and object is receiving. In order to elaborate rules for attitude analysis based on the semantics of verbs, we investigated VerbNet (Kipper et al., 2007), the largest on-line verb lexicon that is organized into verb classes characterized by syntactic and semantic coherence among members of a class. Based on the thorough analysis of 270 first-

level classes of VerbNet and their members, 73 verb classes (1) were found useful for the task of attitude analysis, and (2) were further classified into 22 classes differentiated by the role that members play in attitude analysis and by rules applied to them. Our classification is shown in Table 2.

Verb class (verb samples)
1 Psychological state or emotional reaction
1.1 Object-centered (oriented) emotional state (<i>adore, regret</i>)
1.2 Subject-driven change in emotional state (trans.) (<i>charm, inspire, bother</i>)
1.3 Subject-driven change in emotional state (intrans.) (<i>appeal to, grate on</i>)
2 Judgment
2.1 Positive judgment (<i>bless, honor</i>)
2.2 Negative judgment (<i>blame, punish</i>)
3 Favorable attitude (<i>accept, allow, tolerate</i>)
4 Adverse (unfavorable) attitude (<i>discourage, elude, forbid</i>)
5 Favorable or adverse calibratable changes of state (<i>grow, decline</i>)
6 Verbs of removing
6.1 Verbs of removing with neutral charge (<i>delete, remove</i>)
6.2 Verbs of removing with negative charge (<i>depart, expel</i>)
6.3 Verbs of removing with positive charge (<i>evacuate, cure</i>)
7 Negatively charged change of state (<i>break, crush, smash</i>)
8 Bodily state and damage to the body (<i>sicken, injure</i>)
9 Aspectual verbs
9.1 Initiation, continuation of activity, and sustaining (<i>begin, continue, maintain</i>)
9.2 Termination of activity (<i>quit, finish</i>)
10 Preservation (<i>defend, insure</i>)
11 Verbs of destruction and killing (<i>damage, poison</i>)
12 Disappearance (<i>disappear, die</i>)
13 Limitation and subjugation (<i>confine, restrict</i>)
14 Assistance (<i>succor, help</i>)
15 Obtaining (<i>win, earn</i>)
16 Communication indicator/reinforcement of attitude (<i>guess, complain, deny</i>)
17 Verbs of leaving (<i>abandon, desert</i>)
18 Changes in social status or condition (<i>canonize, widow</i>)
19 Success and failure
19.1 Success (<i>succeed, manage</i>)
19.2 Failure (<i>fail, flub</i>)
20 Emotional nonverbal expression (<i>smile, weep</i>)
21 Social interaction (<i>marry, divorce</i>)
22 Transmitting verbs (<i>supply, provide</i>)

Table 2. Verb classes defined for attitude analysis.

For each of our verb classes, we developed set of rules that are applied to attitude analysis on the phrase/clause-level. Some verb classes include verbs annotated by attitude type, prior polarity orientation, and the strength of attitude: “Psychological state or emotional reaction”, “Judgment”, “Verbs of removing with negative charge”, “Verbs

of removing with positive charge”, “Negatively charged change of state”, “Bodily state and damage to the body”, “Preservation”, and others. The attitude features of phrases, which involve positively or negatively charged verbs from such classes, are context-sensitive, and are defined by means of rules designed for each of the class.

As an example, below we provide short description and rules elaborated for the subclass “Object-centered (oriented) emotional state”.

Features: subject experiences emotions towards some stimulus; verb prior polarity: positive or negative; context-sensitive.

Verb-Object rules (subject is ignored):

1. “Interior perspective” (subject’s inner emotion state or attitude):

S & V+(*admires*) & O+(*his brave heart*) => (fusion, $\max(V_score, O_score)$) => ‘POS aff’.

S & V+(*admires*) & O-(*mafia leader*) => (verb valence dominance, V_score) => ‘POS aff’.

S & V-(*disdains*) & O+(*his honesty*) => (verb valence dominance, V_score) => ‘NEG aff’.

S & V-(*disdains*) & O-(*criminal activities*) => (fusion, $\max(V_score, O_score)$) => ‘NEG aff’.

2. “Exterior perspective” (social/ethical judgment):

S & V+(*admires*) & O+(*his brave heart*) => (fusion, $\max(V_score, O_score)$) => ‘POS jud’.

S & V+(*admires*) & O-(*mafia leader*) => (verb valence reversal, $\max(V_score, O_score)$) => ‘NEG jud’.

S & V-(*disdains*) & O+(*his honesty*) => (verb valence dominance, $\max(V_score, O_score)$) => ‘NEG jud’.

S & V-(*disdains*) & O-(*criminal activities*) => (verb valence reversal, $\max(V_score, O_score)$) => ‘POS jud’.

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

Verb-PP (prepositional phrase) rules:

1. In case of negatively charged verb and PP starting with *from* => verb valence dominance:

S & V-(*suffers*) & PP-(*from illness*) => interior: ‘NEG aff’; exterior: ‘NEG jud’.

S & V-(*suffers*) & PP+ (*from love*) => interior: ‘NEG aff’; exterior: ‘NEG jud’.

2. In case of positively charged verb and PP starting with *in*/*for*, treat PP same as object (see above):

S & V+(*believes*) & PP-(*in evil*) => interior: ‘POS aff’; exterior: ‘NEG jud’.

S & V+(*believes*) & PP+(*in kindness*) => interior: ‘POS aff’; exterior: ‘POS jud’.

In the majority of rules the strength of attitude is measured as a maximum between attitude scores of a verb and an object ($\max(V_score, O_score)$), because strength of overall attitude depends on both scores. For example, attitude conveyed by *to suffer from grave illness* is stronger than that of *to suffer from slight illness*.

In contrast to the rules of “Object-centered (oriented) emotional state” subclass, which ignore attitude features of a subject in a sentence, the rules elaborated for the “Subject-driven change in emotional state (trans.)” disregard the attitude features of object, as in sentences involving members of this subclass object experiences emotion, and subject causes the emotional state. For example (due to limitation of space, here and below we provide only some cases):

S(*Classical music*) & V+(*calmed*) & O-(*disobedient child*) => interior: ‘POS aff’; exterior: ‘POS app’.

S-(*Fatal consequences of GM food intake*) & V-(*frighten*) & O(*me*) => interior: ‘NEG aff’; exterior: ‘NEG app’.

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

S(*He*) & V-(*blamed*) & O+(*innocent person*) => interior: ‘NEG jud’; exterior: ‘NEG jud’.

S(*They*) & V-(*punished*) & O(*him*) & PP-(*for his misdeed*) => interior: ‘NEG jud’; exterior: ‘POS jud’.

Verbs from classes “Favorable attitude” and “Adverse (unfavorable) attitude” have prior neutral polarity and positive or negative reinforcement, correspondingly, that means that they only impact on the polarity and strength of non-neutral phrase (object in a sentence written in active voice, or subject in a sentence written in passive voice, or PP in case of some verbs).

Rules:

1. If verb belongs to the “Favorable attitude” class and the polarity of phrase is not neutral, then the

attitude score of the phrase is intensified (we use symbol '^' to indicate intensification):

S('They') & [V pos. reinforcement]('elected') & O+('fair judge') => 'POS app'; O_score^.

S('They') & [V pos. reinforcement]('elected') & O-('corrupt candidate') => 'NEG app'; O_score^.

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

S('They') & [V neg. reinforcement]('prevented') & O-('the spread of disease') => 'POS app'; O_score^.

S+('His achievements') & [V neg. reinforcement]('were overstated') => 'NEG app'; S_score^.

Below are examples of processing the sentences with verbs from "Verbs of removing" class.

"Verbs of removing with neutral charge":

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

S('The safety invention') & [V neutral rem.]('ejected') & O('the pilot') & PP-('from burning plane') => 'POS app'; PP_score^.

"Verbs of removing with negative charge":

S('Manager') & [V neg. rem.]('fired') & O-('careless employee') & PP('from the company') => 'POS app'; max(V_score, O_score).

"Verbs of removing with positive charge":

S('They') & [V pos. rem.]('evacuated') & O('children') & PP-('from dangerous place') => 'POS app'; max(V_score, PP_score).

Along with modal verbs and modal adverbs, members of the "Communication indicator/reinforcement of attitude" verb class also indicate the confidence level or degree of certainty concerning given opinion.

Features: subject (communicator) expresses statement with/without attitude; statement is PP starting with 'of', 'on', 'against', 'about', 'concerning', 'regarding', 'that', 'how' etc.; ground: positive or negative; reinforcement: positive or negative.

Rules:

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

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

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

S('Jane') & [V neg. ground, pos. reinforcement, confidence:0.8]('is complaining') & PP-('of a headache again') => 'NEG app'; PP_score^; confidence:0.8.

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

S('Max') & [V neg. ground, neg. reinforcement, confidence:0.2]('doubt') & PP-{'that' S+('his good fortune')} & [V termination]('will ever end') => 'POS app'; PP_score^; confidence:0.2.

In the last example, to measure the sentiment of PP, we apply rule for the verb 'end' from the "Termination of activity" class, which reverses the non-neutral polarity of subject (in intransitive use of verb) or object (in transitive use of verb). For example, the polarity of the following sentence with positive PP is negative: 'They discontinued helping children'.

5 Evaluation

In order to evaluate the performance of our algorithm, we conducted experiment on the set of sentences extracted from personal stories about life experiences that were anonymously published on the social networking website *Experience Project* (www.experienceproject.com). This website represents an interactive platform that allows people to share personal experiences, thoughts, opinions, feelings, passions, and confessions through the network of personal stories. With over 4 million experiences accumulated (as of February 2010), *Experience Project* is a perfect source for researchers interested in studying different types of attitude expressed through text.

5.1 Data Set Description

For our experiment we extracted 1000 sentences from various stories grouped by topics within 13 different categories, such as "Arts and entertainment", "Current events", "Education", "Family and friends", "Health and wellness", "Relationships and romance" and others, on the *Experience Project*. Sentences were collected from 358 distinct topic groups, such as "I still remember September 11", "I am intelligent but airheaded", "I think bullfighting is cruel", "I quit smoking", "I am a fashion victim", "I was adopted" and others.

TOP	POS					NEG						neutral		
MID	POS aff			POS jud	POS app	NEG aff						NEG jud	NEG app	neutral
ALL	interest	joy	surprise	POS jud	POS app	anger	disgust	fear	guilt	sadness	shame	NEG jud	NEG app	neutral

Figure 1. Hierarchy of attitude labels.

We considered three hierarchical levels of attitude labels in our experiment (see Figure 1). Three independent annotators labeled the sentences with one of 14 categories from ALL level and a corresponding score (the strength or intensity value). These annotations were further interpreted using labels from MID and TOP levels. Fleiss’ Kappa coefficient was used as a measure of reliability of human raters’ annotations. The agreement coefficient on 1000 sentences was 0.53 on ALL level, 0.57 on MID level, and 0.73 on TOP level.

Only those sentences, on which at least two out of three human raters completely agreed, were included in the “gold standard” for our experiment. Three “gold standards” were created according to the hierarchy of attitude labels. Fleiss’ Kappa coefficients are 0.62, 0.63, and 0.74 on ALL, MID, and TOP levels, correspondingly. Table 3 shows the distributions of labels in the “gold standards”.

ALL level		MID level	
Label	Number	Label	Number
anger	45	POS aff	233
disgust	21	NEG aff	332
fear	54	POS jud	66
guilt	22	NEG jud	78
interest	84	POS app	100
joy	95	NEG app	29
sadness	133	neutral	87
shame	18	total	925
surprise	36	TOP level	
POS jud	66		
NEG jud	78	Label	Number
POS app	100	POS	437
NEG app	29	NEG	473
neutral	87	neutral	87
total	868	total	997

Table 3. Label distributions in the “gold standards”.

5.2 Results

After processing each sentence from the data set by our system, we measured averaged accuracy, precision, recall, and F-score for each label within ALL, MID, and TOP levels. The results are shown in Table 4. The ratio of the most frequent attitude

label in the “gold standard” was considered as the baseline. As seen from the obtained results, our algorithm performed with high accuracy significantly surpassing the baselines on all levels of attitude hierarchy (except ‘neutral’ category on the TOP level, which is probably due to the unbalanced distribution of labels in the “gold standard”, where ‘neutral’ sentences constitute less than 9%).

ALL level				
Baseline	0.153			
Label	Accuracy	Precision	Recall	F-score
anger	0.621	0.818	0.600	0.692
disgust		0.818	0.857	0.837
fear		0.768	0.796	0.782
guilt		0.833	0.455	0.588
interest		0.772	0.524	0.624
joy		0.439	0.905	0.591
sadness		0.528	0.917	0.670
shame		0.923	0.667	0.774
surprise		0.750	0.833	0.789
POS jud		0.824	0.424	0.560
NEG jud		0.889	0.410	0.561
POS app		0.755	0.400	0.523
NEG app		0.529	0.310	0.391
neutral		0.559	0.437	0.490
MID level				
Baseline	0.359			
Label	Accuracy	Precision	Recall	F-score
POS aff	0.709	0.668	0.888	0.762
NEG aff		0.765	0.910	0.831
POS jud		0.800	0.424	0.554
NEG jud		0.842	0.410	0.552
POS app		0.741	0.400	0.519
NEG app		0.474	0.310	0.375
neutral		0.514	0.437	0.472
TOP level				
Baseline	0.474			
Label	Accuracy	Precision	Recall	F-score
POS	0.879	0.918	0.920	0.919
NEG		0.912	0.922	0.917
neutral		0.469	0.437	0.452

Table 4. Results of the system performance evaluation.

In the case of fine-grained attitude recognition (ALL level), the highest precision was obtained for ‘shame’ (0.923) and ‘NEG jud’ (0.889), while the highest recall was received for ‘sadness’ (0.917)

and ‘joy’ (0.905) emotions at the cost of low precision (0.528 and 0.439, correspondingly). The algorithm performed with the worst results in recognition of ‘NEG app’ and ‘neutral’.

The analysis of a confusion matrix for the ALL level revealed the following top confusions of our system (see Table 5): (1) ‘anger’, ‘fear’, ‘guilt’, ‘shame’, ‘NEG jud’, ‘NEG app’ and ‘neutral’ were predominantly incorrectly predicted as ‘sadness’ (for ex., @AM resulted in ‘sadness’ for the sentence ‘*I know we have several months left before the election, but I am already sick and tired of seeing the ads on TV*’, while human annotations were ‘anger’/‘anger’/‘disgust’); (2) ‘interest’, ‘POS jud’ and ‘POS app’ were mostly confused with ‘joy’ by our algorithm (e.g., @AM classified the sentence ‘*It’s one of those life changing artifacts that we must have in order to have happier, healthier lives*’ as ‘joy’(-ful), while human annotations were ‘POS app’/‘POS app’/‘interest’).

Actual label	Incorrectly predicted labels (%), in descending order
anger	sadness (28.9%), joy (4.4%), neutral (4.4%), NEG app (2.2%)
disgust	anger (4.8%), sadness (4.8%), NEG jud (4.8%)
fear	sadness (13%), joy (5.6%), POS app (1.9%)
guilt	sadness (50%), anger (4.5%)
interest	joy (33.3%), neutral (7.1%), sadness (3.6%), POS app (2.4%), fear (1.2%)
joy	interest (3.2%), POS app (3.2%), sadness (1.1%), surprise (1.1%), neutral (1.1%)
sadness	neutral (3.8%), joy (1.5%), anger (0.8%), fear (0.8%), guilt (0.8%), NEG app (0.8%)
shame	sadness (16.7%), fear (5.6%), guilt (5.6%), NEG jud (5.6%)
surprise	fear (5.6%), neutral (5.6%), joy (2.8%), POS jud (2.8%)
POS jud	joy (37.9%), POS app (9.1%), interest (4.5%), sadness (1.5%), surprise (1.5%), NEG jud (1.5%), neutral (1.5%)
NEG jud	sadness (37.2%), anger (3.8%), disgust (3.8%), neutral (3.8%)
POS app	joy (37%), neutral (9%), surprise (7%), interest (3%), POS jud (3%), sadness (1%)
NEG app	sadness (44.8%), fear (13.8%), disgust (3.4%), surprise (3.4%), neutral (3.4%)
neutral	sadness (29.9%), joy (13.8%), interest (3.4%), fear (2.3%), POS jud (2.3%), NEG app (2.3%), NEG jud (1.1%), POS app (1.1%)

Table 5. Data from a confusion matrix for ALL level.

Our system achieved high precision for all categories on the MID level (Table 4), with the exception of ‘NEG app’ and ‘neutral’, although high recall was obtained only in the case of categories

related to affect (‘POS aff’, ‘NEG aff’). These results indicate that affect sensing is easier than recognition of judgment or appreciation from text.

TOP level results (Table 4) show that our algorithm classifies sentences that convey positive or negative sentiment with high accuracy (92% and 91%, correspondingly). On the other hand, ‘neutral’ sentences still pose a challenge.

The analysis of errors revealed that system requires common sense or additional context to deal with sentences like ‘*All through my life I’ve felt like I’m second fiddle*’ (“gold standard”: ‘sadness’; @AM: ‘neutral’) or ‘*For me every minute on my horse is alike an hour in heaven!*’ (“gold standard”: ‘joy’; @AM: ‘neutral’).

We also evaluated the system performance with regard to attitude intensity estimation. The percentage of attitude-conveying sentences (not considering neutral ones), on which the result of our system conformed to the fine-grained “gold standard” (ALL level), according to the measured distance between intensities given by human raters (averaged values) and those obtained by our system is shown in Table 6. As seen from the table, our system achieved satisfactory results in estimation of the strength of attitude expressed through text.

Range of intensity difference	Percent of sentences, %
[0.0 – 0.2]	55.5
(0.2 – 0.4]	29.5
(0.4 – 0.6]	12.2
(0.6 – 0.8]	2.6
(0.8 – 1.0]	0.2

Table 6. Results on intensity.

6 Conclusions

In this paper we introduced @AM, which is so far, to the best of our knowledge, the only system classifying sentences using fine-grained attitude types, and extensively dealing with the semantics of verbs in attitude analysis. Our composition approach broadens the coverage of sentences with complex contextual attitude. The evaluation results indicate that @AM achieved reliable results in the task of textual attitude analysis. The limitations include dependency on lexicon and on accuracy of the parser. The primary objective for the future research is to use the results of named-entity recognition software in our algorithm.

References

- Cecilia O. Alm. 2008. *Affect in Text and Speech*. PhD Dissertation. University of Illinois at Urbana-Champaign.
- Saima Aman and Stan Szpakowicz. 2008. Using Roget's Thesaurus for Fine-Grained Emotion Recognition. *Proceedings of the Third International Joint Conference on Natural Language Processing IJCNLP 2008*, Hyderabad, India, pp. 296-302.
- Plaban Kumar Bhowmick, Anupam Basu, and Pabitra Mitra. 2009. Reader Perspective Emotion Analysis in Text through Ensemble based Multi-Label Classification Framework. *Computer and Information Science*, 2 (4): 64-74.
- Anthony C. Boucouvalas. 2003. Real Time Text-to-Emotion Engine for Expressive Internet Communications. *Being There: Concepts, Effects and Measurement of User Presence in Synthetic Environments*, Ios Press, pp. 306-318.
- Francois-Regis Chaumartin. 2007. UPAR7: A Knowledge-based System for Headline Sentiment Tagging. *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, Prague, Czech Republic, pp. 422-425.
- Yejin Choi and Claire Cardie. 2008. Learning with Compositional Semantics as Structural Inference for Subsentential Sentiment Analysis. *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 793-801.
- Ze-Jing Chuang and Chung-Hsien Wu. 2004. Multimodal Emotion Recognition from Speech and Text. *Computational Linguistic and Chinese Language Processing*, 9(2): 45-62.
- Leo Hoyer. 1997. *Adverbs and Modality in English*. New York: Addison Wesley Longman Inc.
- Carroll E. Izard. 1971. *The Face of Emotion*. New York: Appleton-Century-Crofts.
- Karin Kipper, Anna Korhonen, Neville Ryant, and Martha Palmer. 2007. A Large-scale Classification of English Verbs. *Language Resources and Evaluation*, 42 (1): 21-40.
- Zornitsa Kozareva, Borja Navarro, Sonia Vazquez, and Andres Montoyo, A. 2007. UA-ZBSA: A Headline Emotion Classification through Web Information. *Proceedings of the Fourth International Workshop on Semantic Evaluations*, pp. 334-337.
- Hugo Liu, Henry Lieberman, and Ted Selker. 2003. A Model of Textual Affect Sensing Using Real-World Knowledge. *Proceedings of the International Conference on Intelligent User Interfaces*, pp. 125-132.
- James R. Martin and Peter R.R. White. 2005. *The Language of Evaluation: Appraisal in English*. Palgrave, London, UK.
- George A. Miller. 1990. WordNet: An On-line Lexical Database. *International Journal of Lexicography*, Special Issue, 3 (4): 235-312.
- Karo Moilanen and Stephen Pulman. 2007. Sentiment Composition. *Proceedings of the Recent Advances in Natural Language Processing International Conference*, pp. 378-382.
- Matthijs Mulder, Anton Nijholt, Marten den Uyl, and Peter Terpstra. 2004. A Lexical Grammatical Implementation of Affect. *Proceedings of the Seventh International Conference on Text, Speech and Dialogue*, pp. 171-178.
- Tetsuya Nasukawa and Jeonghee Yi. 2003. Sentiment Analysis: Capturing Favorability using Natural Language Processing. *Proceedings of the 2nd International Conference on Knowledge Capture*, pp. 70-77.
- Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. 2009. SentiFul: Generating a Reliable Lexicon for Sentiment Analysis. *Proceedings of the International Conference on Affective Computing and Intelligent Interaction*, IEEE, Amsterdam, Netherlands, pp. 363-368.
- J. Olveres, M. Billinghamurst, J. Savage, and A. Holden. 1998. Intelligent, Expressive Avatars. *Proceedings of the First Workshop on Embodied Conversational Characters*, pp. 47-55.
- Livia Polanyi and Annie Zaenen. 2004. Contextual Valence Shifters. *Working Notes of the AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications*.
- Carlo Strapparava and Rada Mihalcea. 2008. Learning to Identify Emotions in Text. *Proceedings of the 2008 ACM Symposium on Applied Computing*, Fortaleza, Brazil, pp. 1556-1560.
- Carlo Strapparava, Alessandro Valitutti, and Oliviero Stock. 2007. Dances with Words. *Proceedings of the International Joint Conference on Artificial Intelligence*, Hyderabad, India, pp. 1719-1724.
- V.S. Subrahmanian and Diego Reforgiato. 2008. AVA: Adjective-Verb-Adverb Combinations for Sentiment Analysis. *Intelligent Systems*, IEEE, 23 (4): 43-50.
- Maite Taboada and Jack Grieve. 2004. Analyzing Appraisal Automatically. *Proceedings of American Association for Artificial Intelligence Spring Symposium on Exploring Attitude and Affect in Text*, pp.158-161.
- Casey Whitelaw, Navendu Garg, and Shlomo Argamon. 2005. Using Appraisal Groups for Sentiment Analysis. *Proceedings of the 14th ACM International Conference on Information and Knowledge Management*, CIKM, Bremen, Germany, pp. 625-631.
- Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing Contextual Polarity in Phrase-level Sentiment Analysis. *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, Vancouver: ACL, pp. 347-354.