Someone to Talk To

Kiril Koroveshovski and Sonja Gievska

Faculty of Computer Science and Engineering
Ss. Cyril and Methodius University
Skopje, 1000, R. Macedonia

ABSTRACT

This paper reflects upon the challenges surrounding the efforts in recognizing and classifying user's affective state. A suitable set of rules for contextual valence shifting has a central role in the proposed lexical-based approach for automatic emotion detection, which utilizes a diverse set of publicly available lexical resources. To evaluate the strengths and weaknesses of the embedded algorithm for word valence assignment, an experimental study with a suitable dataset was conducted and the performance results are discussed. A prototype multimodal mobile application that steers the conversational dialogue aligned with user's affective states will also be presented.

Keywords: Automatic Emotion Detection, Affective Lexical Analysis, Contextual Valence Shifters, Emotional Intelligence

INTRODUCTION

The growing interest in the topics of individual and social behavior has generated an extensive research aimed at capturing such behavior. Future research opportunities that are expected to tackle the challenges of perceptive, cognitive and affective aspects of interaction design in novel ways. Guiding the interaction aligned with "human nature" holds great potential to be fruitfully applied to areas from e-learning and personalized recommendation engines to persuasive technologies and social collaboration.

Heralding emotions as a crucial component in a number of high-level cognitive activities, such as memory, decision making and social cognitive abilities, Damasio's work on the biological origins of consciousness (Damasio, 1999) has shown that emotions, uniquely defined in each person by a complex and rich plethora of inner neural and body states, delineate and pervade our behavior, thoughts and actions. Our knowledge of the phenomenon might have been deepened on some levels, though recognition of invisible and visible manifestation of human affective states is a challenging computational task; not an easy task for humans too. How difficult is to express your emotions or describe the emotional state of others? Face recognition, speech pattern analysis, physiological sensors technologies and semantic lexical analysis are expected to lead the way toward affective interaction design.

Computational linguistics approaches have been applied to a range of challenging problems with impact outside the language technology field, including affective interaction. Language use has long been a favorite testing ground for a number of theories in psychology, psychiatry and criminology relating the emotions to causes of different kinds of behavior. A variety of approaches have been proposed, which make exclusive use or a fusion of the following techniques: statistical lexical analysis, lexicons of affective words and phrases, contextual-based rules to resolve the word sense disambiguation, and machine learning techniques for training classifiers for emotion assessment. In the

https://openaccess.cms-conferences.org/#/publications/book/978-1-4951-2109-8

wake of the results from related studies, our research proposes a lexical-based method that utilizes publicly available lexical resources. The strength of the word emotional valence plays a pivotal role in our approach for detecting the dominating emotion. Our current attention has been given to the contextual relations between words in a sentence that can change the direction or neutralize the initially assigned emotional valence of a word.

Exhibiting emotional intelligence towards your conversational partner is what humans do; by being able to sense, interpret and adapt to users' emotions the systems can generate responses, offer comfort and advice or just be "someone to talk to". This research examines the ways in which intelligent technologies may facilitate detection of six basic emotions through lexical analyses of user's reflections on her daily experiences and affective states as a basis for emphatic response and potentially helpful suggestions by the system. This paper discusses the two essential parts of our research: (1) a lexical-based approach for real-time emotion detection; and (2) a multimodal Android application that is developed to provide the user with a comfortable way to express her emotions through dialog with the system. The results of our initial exploratory study of the performance and potential usability advantages of the proposed method will be presented. Future research directions for strengthening our research will be indicated. We consider the results of the performance evaluation of our method to be strong evidence that in conjunction with our previous research on personality prediction (Markovikj et al., 2013) points to the potential of lexical analysis of user-generated text in modeling users (e.g., personality traits, affective states).

RELATED RESEARCH

The problem of affective analysis of text has received a considerable research attention in the intelligence community in the last decade. In a literature review we can identify recurring approaches to the problem of interest, drawing both on lexical-based (Calvo and Sunghwan, 2013; Strapparava and Mihalcea, 2008; Chaumartin, 2007) and machine learning techniques (Inkpen et al. 2009; Kim and Valitutti, 2010). Carlo Strapparava and Valitutti have laid the background to a large number of research related to sentiment and emotional analysis from the perspective of language clues as a predictor of inner emotional state (Strapparava and Valitutti, 2004). A gold standard dataset of 1250 news headlines with emotion annotations, SemEval 2007 task, was created and subjected to analysis from numerous systems for automatic emotion evaluation. Six emotional categories namely, anger, surprise, happiness, sadness, disgust and fear were to be detected.

Few systems have been proposed in (Strapparava and Mihalcea, 2008), which have utilized a lexicon of direct and indirect affective words, mainly extracted from WordNetAffect, and a variation of Latent Semantic Analysis (LSA). The algorithm based on detection of the presence of WN-Affect lexicon has shown best results in terms of average precision 38.28%, although exhibited very low recall (avg. 1.54). The results of the three LSA-based approaches have shown significantly lower precisions below 10% and high recall spanning from 66.72 to 90.22 depending over emotions. A different dataset consisted of 8761 emotion-annotated blogs from LiveJournal.com was used for training a Naïve Bayes classifier, which have exhibited average precision and recall, 12.04% and 18.01, respectively.

The work closely related to our research is University of Paris 7 (UPAR 7) system that has participated in the same SemEval 2007 emotion annotation task and reported an average precision of 27.6 % and recall of 5.68 (Chaumartin, 2007). While reaching for similar publicly available resources and tools for emotion lexical analysis, the method underlying UPAR 7 system is quite different. Individual word rating that deeply relies on WordNet synsets, detection of contrasts (sentiment of "good" vs. "bad"), and global sentence rating were at the core of the system that obtains high accuracy results, on average 89.43%.

There are few mobile applications exploring objectives related to the one envisioned in our research. The feasibility of crowdsourcing is investigated in (Morris and Picard, 2012) as a way of assisting people when dealing with stressful situations. On demand human workforce as opposed to affective computation is employed to detect and respond to users' emotional states. Another Android application, Emotion Sense, that offers users a way to be aware of the reasons for their mood change is a result of a research exploring the relations between smartphone sensing technologies and behavior change interventions (Lathia et al., 2013). In the third related work, users' usage patterns on Twitter are proposed as a method for detecting emotions (Herdem, 2012). Few hypothetical scenarios that suggest contacts with friends when a user needs emotional support are discussed. The system is yet to be realized and tested.

RESOURCES AND TOOLS FOR AFFECTIVE LEXICAL ANALYSIS

A set of publicly available resources and appropriate tools have been used in our system for automatic emotion detection; some of them previously reported in related studies, although our algorithm depends on a unique combination of custom implementations and modified methods to suit the objective of our study. A short introduction of the lexical resources and tools and rationale for their inclusion follows:

WordNet-Affect[1], an extension of WordNet that contains a subset of synsets representing moods, situations eliciting emotions or emotional responses is frequently used in sentiment and emotion text analysis (Strapparava and Valitutti, 2004). A subset of the lexicon, containing the affective words directly referring to the six emotional states of interest – joy, sadness, anger, fear, disgust, and surprise, were used in our system.

AFINN-111[2], a list of 2477 English words and phrases annotated with their valence rating, an integer value between -5 and 5.

Stanford parser offers highly optimized Probabilistic Context-Free Grammars (PCFG), lexicalized dependency parsers and lexicalized PCFG parser for performing lexical analysis of a sentence (De Marneffe and Manning, 2012). Finding typed dependencies i.e. grammatical relations between the words in a sentence has provided assistance in the semantic (contextual) analysis of a text.

Java API for WordNet (JAWS) have been used by our system to retrieve data and find the base form of a given word from the WordNet database, version 2.1 and 3.0.

SharpNLP provides a collection of natural language processing tools. In particular, sentence splitter, tokenizer, part-of-speech (POS) tagger, and grammatical word tagging (e.g., adverbs, adjectives, verbs, pronouns) have been incorporated in our algorithm for affective text analysis.

OUR LEXICAL-BASED APPROACH FOR AUTOMATIC EMOTION DETECTION

The process of emotion detection goes through several stages as presented in Figure 1, each one requires a selection of suitable technique or a tool. The processing starts by comparing the user input against the *lexicon of phrases*, which is a list of all phrases in AFINN list with the inclusion of our own addition of 20 phrases deemed to be relevant to our research (e.g., pass away). By identifying the affective phrases first, a potential source for mistakes in affective valence assignment has been avoided; the interpretation of a phrase could be quite different if the words of a phrase are processed separately.

Once the recognition of affective phrases is performed, a set of preprocessing techniques from SharpNLP is employed to extract the affective words from the user input. The incorporated techniques for preprocessing include *sentence detection*, so that each sentence is processed independently starting with a *tokenization* i.e. extracting words from sentences. Then, *POS tagging* of words is performed to identify their corresponding grammatical types. In our initial investigation, only the verbs, nouns, adjectives and adverbs are included in the subsequent semantic analysis.

WordNet API is used to get the *base form of a word*. Our custom stemmer was employed to reduce a word to its

[1] http://wndomains.fbk.eu/wnaffect.html
[2] http://www2.imm.dtu.dk/pubdb/views/publication_details.php?id=6010

root, which is matched against our *lexicon of affective words*. The lexicon combines the words from two resources: all words related to the the six emotions from WordNet Affect and all words from the Afinn list. Six valences are assigned to each word, one for each emotion under investigation. The initial valence is calculated using the following formula:

$$\textit{Valence} = 5 + 3{*}\textit{Afinn\_val},$$

where ***Afinn_val*** is the the absolute value of the valence assigned by the Afinn list. The *valence assignment* by which the negative sign of the Affnn valences is eliminated, making the range of valences to span between 5 and 20, reflects our current objective, namely to detect the most prominent emotion in user input as opposed to its sentiment polarity.

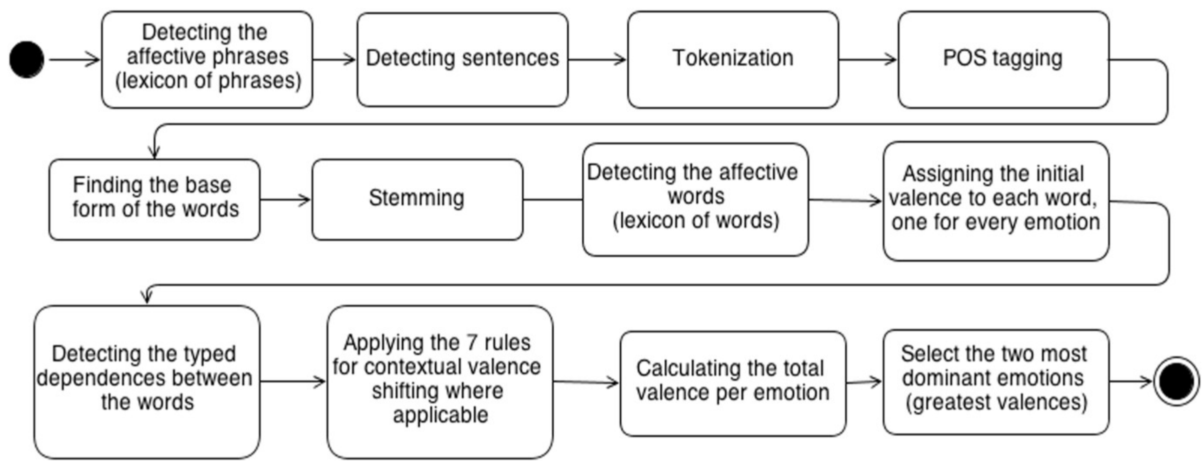

Figure 1. Steps in the lexical-based approach for automatic emotion detection

Rules for Contextual Valence Shifting

Stanford typed dependency parser was employed to identify the grammatical relations between the words in a sentence. There are 53 grammatical relations between a governor and a dependent (De Marneffe and Manning, 2012), only those related to our *seven rules for contextual valence shifting* were considered. Adding or subtracting a predetermined number of points following the rules underlying our method adjusts the valences as follows:

Negations. If a positively valenced word is in relation with a negation (e.g., not, never, nothing, none, neither), its valence is shifted into the valence of the opposite emotion, if such an emotion exists (e.g. opposite of joy is sadness). The valence is modified towards a neutral position, if the opposite emotion does not exist.

Amplifiers. The positive valence of a word is raised by 5, if the word is in a relation with a lexical item that intensifies its positive attitude (e.g., deeply, always, best, clearly, strongly).

Attenuators. If a word is in relation with a lexical item that lowers its emotion strength (e.g., rather, lack, least), its valence is lowered by 5.

Neutralizers. If a word is in relation with a connector word (e.g., however, although, but), its initial valence is neutralized, set to 0.

Root: If a word is identified as a root of the dependency graph of a sentence, 6 points are added to its previously assigned valence. Not to belittle or even diminish the importance of the emotions carried by other words while overemphasizing the sentence root, we have opted for increasing the valence by addition, which is quite different from the multiplication rule (3-5 times) adopted by UPAR7 system (Chaumartin, 2007). We put forward again the notion that our approach seeks to find the dominant emotion as opposed to detect the general attitude i.e., sentiment expressed in the user input.

Negative shifters. If a word is in relation with a verb (e.g., fail, omit, neglect) or a noun (e.g., failure, neglect) that modifies the initial valence of the word in the opposite direction, its valence is changed with the valence of the opposite emotion, if such emotion exists. In addition, the associated valence for anger of the word is increased by 5, to reflect the fact that something was expected by the user but had failed to realize, which can often attribute to feelings such as aggravation and frustration, in our coarse-grained classification closely related to anger.

Conditional tense. If a sentence is written or spoken in conditional tense and the relations between a word and the verb "would" is auxiliary, then the valence of the word is set to 0.

The final stages in the algorithm are dedicated to calculating the total valence for the entire text. The valences of the words are aggregated (summed) into six scores, one for each of the six emotions, each score is normalized by dividing it by the maximum valence value. The two topmost emotions (greatest valences) are selected to characterize the user input in an affective sense.

PERFORMANCE EVALUATION

The ISEAR[3] (International Survey on Emotion Antecedents and Reactions) project, a collaborative endeavor directed by K. R. Scherer and H. Wallbott, has produced a dataset that was used for evaluating the performance of our algorithm. The ISEAR dataset is a collection of a large reference corpus of personal reports on situations related to seven emotions – joy, fear, anger, sadness, disgust, shame and guilt, which were solicited from more than 3000 students from all over the world. For the purpose of exploring the performance standing of our method, we have used a subset of the ISEAR dataset that relates to joy, fear, anger, sadness, and disgust. In the absence of test data for the sixth emotional category, surprise, additional set of data was obtained from another dataset, "Classic literary tales annotated for affective contents" [4]. The final dataset consisted of 5287 annotated statements with the following distribution over the six emotions: 1265 - anger, 1023 - fear, 1081 - sadness, 756 - disgust, 1068 - joy, 94 - surprise.

We would like to note that in the current implementation of the algorithm, two situation may arise causing the system to fail in detecting the primary emotion: 1) if the two topmost emotions represent two opposing emotions of the spectrum, and 2) if the two emotions have very close valences (less than 10%). A confusion matrix in Table 1 shows the classification results for the testing dataset; the grey-shaded cells along the diagonal refer to the correctly classified cases, while the other numbers represent the misclassified cases. We could trace limitations in the performance of the typed dependency parser, which resulted in 495 (9.36%) unidentified cases. Close inspection of data has confirmed the existence of 501 (9.4% of the misclassified cases) to belong to the second conflicting situation in which the two topmost emotions had valence differences below 10% (margin set by the designers). The likeliest explanation of the rather high percentage of such cases could be the fact that people frequently express multiple emotions when reflecting on events and happenings in their life (e.g., when discussing death of a close one, the dominant emotion may be sadness, alongside anger or fear). We are currently exploring a hybrid approach that relies on employing machine learning techniques to complement our approach and eliminate the problems related to misclassified cases for future explorations.

Table 1: Confusion matrix for the six emotions

	Surprise	***Anger***	***Joy***	***Sadness***	***Fear***	***Disgust***	***None***
Surprise	90	0	0	1	1	0	2

3 http://www.affective-sciences.org/system/files/ webpage/ISEAR.zip.

4 http://people.rc.rit.edu/~coagla/affectdata/

Anger	3	1180	29	13	6	3	31
Joy	1	7	901	8	11	1	139
Sadness	1	15	92	893	2	1	77
Fear	12	15	41	17	835	3	100
Disgust	1	15	83	18	21	472	146

The results presented in Table 2 show high accuracy for five emotions on average 0.83. The best performance metrics correspond to the emotion categories of *surprise*, *anger* and *joy*; one could postulate the reason being the frequent use of words strongly indicative of the three emotions. The classifier was not as accurate for *disgust* 0.62 and the recall appeared to be on the low side as well (0.76). The result is inline with the performance metrics reported in the related research. The lack of differences in the metrics for *disgust* may in part reflect the fact that possible targets for this emotion are wide-ranging and very subjective (e.g., food, insects, horrible events); what is a source for *disgust* for one person could be an amusement to another.

The results suggest that the lexicon-based approach in combination with the rules for adjusting valence, based on context analysis inspired by the work of L. Polanyi and A. Zaenen (Polanyi and Zaenen, 2006), improves the performance of the affective assessment. A study presented in (Calvo and Sunghwan, 2013) reports on the performance evaluation of one dimensional and three categorical methods for emotion detection, employed with four datasets including the two used in our study. Our results show better precision for the four emotion categories in common. In terms of recall, our method yields better performance for *fear*, 0.89 as opposed to 0.263 obtained by their dimensional method, and *sadness*, 0.92 as opposed to the reported 0.491. We report very close recall values for *anger* and *joy* to the ones reported in their study (1 and 0.98, respectively).

Table 1: Accuracy, precision and recall for each emotion and their average

Emotion	**Accuracy**	**Precision**	**Recall**
surprise	0.96	0.52	0.97
anger	0.93	0.81	0.97
joy	0.84	0.55	0.86
sadness	0.83	0.77	0.92
fear	0.82	0.72	0.89
disgust	0.62	0.95	0.76
Average	**0.83**	**0.72**	**0.90**

Comparative performance analysis of the results of our exploratory study with other related research studies is challenged as usual by the differences in datasets. The performance results of ours, and the related studies point to the challenges in emotional assessment. Finding the valence of the words and phrases in any given text is just a starting point in the process of detecting the most prominent emotion (Morsy and Rafea, 2012; Ghazi et al., 2012). New challenges arise in the analysis; the foremost among these is the difficulty in semantic interpretation of a text that requires inclusion of additional indicators of word properties (Osherenko, 2008; Smith and Lee, 2012). Word sense ambiguity is an obstacle to any semantic lexical analysis, which merits the inclusion of additional contextual indicators of word property. Multi-entity topics, identifying the target of a sentiment and complex sentence structures that do not necessarily reveal user's emotion in actual situations, but rather speculations, other people's opinion, irony, etc. invite further research efforts.

The research on affective analysis points to the challenges in emotional assessment. Word sense ambiguity is an obstacle to any semantic lexical analysis, which merits the inclusion of additional contextual indicators of word property. We are continuing to refine the mechanisms for contextual analysis and explore the feasibility of considering other discourse features that may attribute to higher performance gains in emotion assessment.

MULTIMODAL MOBILE WELL-BEING APPLICATION

Identifying the emotions of a person based on her description of events, experiences, opinions or interests can be helpful for personalizing and guiding the interaction dialogue in a manner that is sensitive and adaptive to the user's affective states. To demonstrate the usefulness of our approach, an Android mobile application *TalkToMe* was developed as a first attempt in a realization of an envisioned application that cares about users' mental welfare and promotes healthy habits for alleviating everyday stress.

Our prototype application uses the Facebook SDK to extract personal information (e.g., name, gender, work- or study- related data, birthday, hobbies, events). A user can log into the application as anonymous or by using her Facebook account, which determines the extent to which the interaction will be personalized. The application employs speech recognition and text-to-speech technologies supported by Android to provide a multimodal interface for interacting with the user.

The dialogue starts with a personalized greeting (e.g., *Good evening, Kiril*) and a return greeting response from the user is expected for further interaction. The greeting phase is usually followed by a question about the day, aligned with the contextual Facebook data, such as the name of employer or school (e.g., *Today is Wednesday. How was your day at GS1*?). If Facebook connection is blocked, general contextual information such as time of day, day of the week, holiday is used (e.g., *How was your day? Did you have a nice weekend?*). A good bye greeting from the user at any time will close the application. The dialogue progresses respecting the boundaries on how far the user wants to elaborate. The system offers no advice or suggestion if the user does not welcomes it; a "No" or a 15 sec user's inactivity is considered a closure of the discussion on a certain topic.

All user utterances or typed text is analyzed by the system for automatic emotion detection, which is an implementation of our lexical-based approach as a web service. The communication with the service is realized with JSON objects. The processing time requires 5-10 sec depending on the text length. Being mindful of potential concerns related to privacy, trust and misuse, our current version does not store the history of user's conversations, although the potential for their inclusion in our further research efforts will be explored.

In case the system encounters the two previously mentioned situations related to too close or opposing emotions, the system does not favor either one and postpones the detection of the dominant emotion, waiting further responses from the user i.e., by encouraging the user to elaborate on her previous statement. If the system is successful in discriminating between emotions and detects the dominant one, it replies with one of the predefined set of statements, questions or suggestions that has been compiled in consultation with an expert in psychology, with a long expertise in working with post-trauma patients. System's responses are nonjudgmental, randomly selected from a set of statements appropriate to the detected emotion (fig. 2).

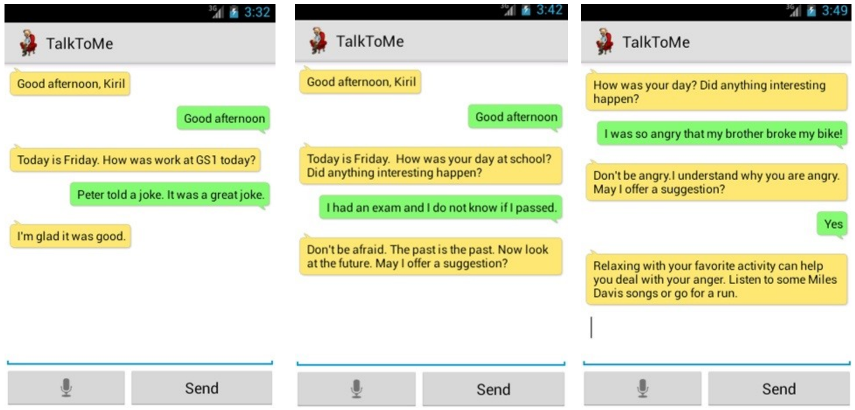

Figure 1. Excerpts of the communication dialogue for three different emotions: joy, fear and anger

A selection of popular motivational quotes from movies and books are included to accompany the system's response attuned with the detected emotion. System responses follow very general recommendations:

- Sharing user's happiness (e.g., "*Good, I'm glad that you're happy*.").
- Being compassionate and empathetic when they are sad (e.g., " *Sorry to hear that*.").
- Suggesting a positive perspective when dealing with difficult events (e.g., "*Look at the brighter side of things*").
- Suggesting offline social encounters with friends, or relaxing with a favorite activity (watching a movie, sports, reading a book) to alleviate stressful situations.
- Instead of negativity or focusing on failures when dealing with fear, the system reminds the user to recognize her strengths (e.g. "*Don't be afraid of your fears. They're not there to scare you, but to let you know that something is worth it*").
- When the dominant emotion is disgust, ignoring the cause or making a change is proposed (e.g. "*You have to identify the problem and try to solve it or ignore it.*").

An exploratory usability study was conducted with 24 participants with varying background experience. The age of the participants span between the ages of 12 to 66. Each participant has been asked to use the application for a number of tries over several days; each trial lasting for a short period of time, no more than 15min. Summarizing the results of the qualitative analysis, we would like to emphasize the high users' satisfaction with the way the interface initiates and carries the conversation. The system responses were rated high for their variety, wit and potential helpfulness. The multimodal interaction and the "word bubble" graphical interface were found to be engaging and a comfortable setting for sharing their reflections on the events and their inner feelings. Voice pattern recognition techniques offer a potential for complementing the current text-based method for emotion detection. The early results of the formative evaluation and overall qualitative gains are currently analyzed for reporting in a separate research paper. A large-scale long-term usability study under controlled experimental condition is in place

CONCLUSIONS

This paper explores the feasibility of real-time detection of six emotions using a lexical-based approach for affective analysis. A multimodal mobile application promoting mental health and well being by inviting people to share some reflections on their current emotional state have been developed to demonstrate the potential usefulness of the proposed approach. We have reported the encouraging performance results that merits further research efforts.

While we do not believe that persuasive technologies can substitute the need for a human-human conversation or a professional help when required, we conclude that future advances in affective interaction will require uniting experts from a number of different fields, from sociology and psychology to linguistic and computer science.

Acknowledgments. This research is part of the project MUSE-Modeling Users Societies Environments Year 2013/2014, partially funded by the Faculty of Computer Science and Engineering, University of “Ss. Cyril and Methodius”, Skopje, Macedonia.

REFERENCES

Calvo, R. A., Sunghwan, M. K. (*2013*), “*Emotions in Text*: *Dimensional and Categorical Models”*, COMPUTATIONAL INTELLIGENCE, Volume 29 No.3, pp. 527-543.

Chaumartin, F. (2007), *“UPAR7: A knowledge-based system for headline sentiment tagging”*, in: Proceedings of 4th International Workshop on Semantic Evaluations. pp. 422-425.

Damasio, A. (1999), “*The Feeling of What Happens: Body, Emotion and the Making of Consciousness*”, Harcourt Brace.

De Marneffe M.C., Manning D. C. (2012), “*Stanford typed dependencies manual”*, v. 2.0.4. Technical report, Palo Alto.

Ghazi D., Inkpen D., Szpakowicz S. (2012), *“Prior versus contextual emotion of a word in a sentence*”, in: Proceedings of the 3rd Workshop on Computational Approaches to Subjectivity and Sentiment Analysis. pp. 70-78.

Herdem, K. C. (2012), *“Reactions: Twitter based mobile application for awareness of friends' emotions”*, in: Proceedings of 2012 ACM Conference on Ubiquitous Computing. pp. 796-797.

Inkpen D., Keshtkar F., Ghazi D. (2009), “*Analysis and generation of emotions in texts”*, in: Proceedings of International Conference on Knowledge Engineering: Principle and Technique. pp. 3-13.

Kim, M. S., Valitutti A. (2010), “*Evaluation of unsupervised emotion models to textual affect recognition*”, in: Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text. pp. 62 – 70.

Lathia, N., Pejovic, V., Rachuri, K., Mascolo, C., Musolesi, M., Rentfrow, P., “*Smart phones for large-scale behavior change interventions*”, IEEE J. PERVASIVE COMPUTING Volume 12 No. 3, pp. 2-9.

Markovikj, D., Gievska, S., Kosinski, D., Stillwell, D.S. (2013), *“Mining Facebook Data for Predictive Personality Modelling”*, in: Proceedings of 7th International AAAI Conference on Weblogs and Social Media, AAAI Press.

Morris, R.R., Picard, R.W. (2012), “*Crowdsourcing Collective Emotional Intelligence*”, in: Proceedings of Collective Intelligence.

Morsy, S. A., Rafea A. (2012), “*Improving document-level sentiment classification using contextual valence shifters*”, in: Lecture notes in computer science. No. 7337. pp. 253-258.

Osherenko, A. (2008), “*Towards semantic affect sensing in sentences*”, in: Proceedings of the AISB 2008 Symposium on Affective Language in Human and Machine. pp. 41-44.

Polanyi, L., Zaenen, A. (2006). “*Contextual Valence Shifters”*, in: Computing Attitude and Affect in Text: Theory and Application, Shanahan, J.G. Qu, Y., Wiebe, J., editors, Volume 20, Springer, pp. 2-8.

Smith P., Lee M. (2012), “*A CCG-based approach to fine-grained sentiment analysis*”, in: Proceedings of the 2nd Workshop on Sentiment Analysis where AI meets Psychology. pp. 3-16.

Strapparava, C., Mihalcea, R. (2008), “*Learning to identify emotions in text*”, in: Proceedings of the 2008 ACM symposium on Applied computing. pp. 1556-1560.

Strapparava, C., Valitutti, A. (2004), “*WordNet-Affect: an affective extension of WordNet*”, in: Proceedings of 4th International Conference on Language Resources and Evaluation. pp. 1083-1086.