U. Yasavur, J. Travieso, and C.L. Lisetti (in press). Sentiment Analysis Using Dependency Trees and Named-Entities. In Proceedings of the 27th International FLAIRS Conference, www.aaai.org, (Pensacola Beach, Florida, USA, May 2014).

Sentiment Analysis Using Dependency Trees and Named-Entities

Ugan Yasavur, Jorge Travieso, Christine Lisetti, Naphtali Rishe
School of Computing & Information Sciences
Florida International University
Miami, FL, 33199, USA

Abstract

There has recently been growing interest in valence and emotion sensing using a variety of signals. Text, as a communication channel, gathers a substantial amount of interest for recognizing its underlying sentiment (valence or polarity), affect or emotion (e.g. happy, sadness). We consider recognizing the valence of a sentence as a prior task to emotion sensing. In this article, we discuss our approach to classify sentences in terms of emotional valence. Our supervised system performs syntactic and semantic analysis for feature extraction. Our system processes the interactions between words in sentences using dependency parse trees, and it can identify the current polarity of named-entities based on on-the-fly topic modeling. We compared the performance of three rule-based approaches and two supervised approaches (i.e. Naive Bayes and Maximum Entropy). We trained and tested our system using the *SemEval-2007* affective text dataset, which contains news headlines extracted from news websites. Our results show that our systems outperform the systems demonstrated in *SemEval-2007*.

Introduction

In the last decade, emotion and sentiment analysis research has become a highly active field due to the increased necessity to recognize emotions, sentiments, opinions or affects conveyed through text. The possible applications which might benefit from recognizing affective information acquired from text, include but are not limited to sentiment analysis for customer reviews (Pang and Lee 2008), opinion mining (Li, Zhang, and Sindhwani 2009), reputation management systems (Yi and Niblack 2005), affective and natural language user interfaces, such as spoken dialog systems (Turunen et al. 2011; Yasavur, Lisetti, and Rishe 2013). Our system concentrates on sentence level emotion polarity recognition (i.e. positive or negative), which we consider as a prior task to emotion recognition.

We are using SemEval-2007 Task 14 affective text dataset, to evaluate our system (Strapparava and Mihalcea 2007). The dataset is composed of news headlines which were extracted from news websites (e.g. Google News, CNN). Headlines typically contain several words and are often written with the intention to provoke emotions to attract the attention from the readers. These characteristics of headlines make them suitable to use in emotion recognition and polarity classification tasks. The specific challenge is the small number of words available for the analysis. Although, there is a general intuition that all words can potentially convey affective meaning (Strapparava, Valitutti, and Stock 2006), the coverage of available lexical resources (Strapparava and Valitutti 2004; Esuli and Sebastiani 2006; Stone, Dunphy, and Smith 1966) falls short for annotating words in headlines. The coverage limitation is actually expected for headlines because of the small number of words.

Our intuition is that the importance or effect of each word for the overall polarity assessment is inversely proportional to the length of a headline. In other words, the contribution of each word to the emotional polarity of a headline increases while the number of words in a headline decreases. Based on this intuition, strictly using the contribution of each word for polarity assessment in a headline is crucial. In the lexical resources for affect or sentiment annotation, there are usually adjectives, verbs, adverbs and common nouns which are useful for emotion or polarity recognition. However, proper nouns (person, location and organization names) in headlines also evoke emotions and positive or negative sentiments for readers. Therefore it has an influence on the polarity of the whole sentence. Even proper nouns may dominate the other sentiment-bearing words which are retrieved from the lexical resources. For example ("*Asia seeks breakthrough on North Korea*"), there are two main sentiment bearing entities, *Asia* and *North Korea*. If we omit these two proper nouns, it is very likely that we will lose the opportunity to recognize the polarity of the sentence. In addition, some of the words or proper nouns may have more influence than others to overall sentiment polarity, in this case, *North Korea* dominates other sentiment bearing words in the headline and it becomes the main influence on the polarity.

The sentiment conveyed by proper nouns evolves by time and current events (e.g. natural or man-made disasters, economic developments, political developments). For example, Japan usually evokes positive sentiments on people but if there is a recent disaster, such as an earthquake or nuclear disaster, it evokes negative sentiments. The proper nouns in

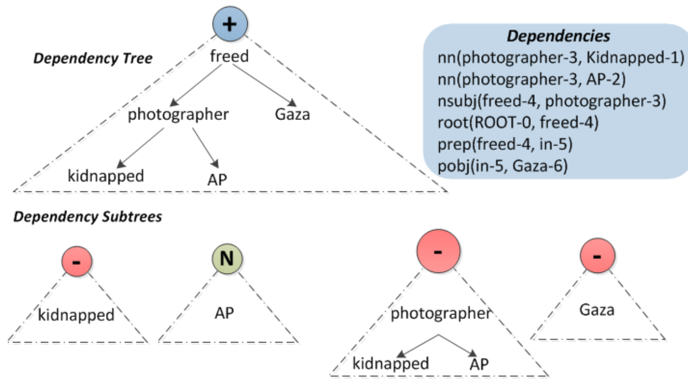

Figure 1: Dependencies

sentences can be identified by any standard named entity recognizers (Finkel, Grenager, and Manning 2005) which can recognize location, person and organization entities. The real difficulty is deciding the current sentiment polarity of the recognized named entities. We addressed this problem in our system which is discussed in the *Approach* section.

In sentiment polarity recognition, a sentence might contain positive or negative polarity words, but a sentence does not necessarily have the same polarity as a whole. To take into account the interactions between words instead of handling the words independently, our system implements dependency parsing and some linguistics rules, such as polarity reversal.

Related Work

Sentiment analysis has gathered the attention of many research groups from variety of areas, such as affective computing, linguistics and psychology. There is a substantial amount of work done using different approaches, that can be categorized in 2 main categories (i.e. rule-based and statistical approaches). Ruled-based systems usually try to apply linguistic compositionality rules (Neviarouskaya, Prendinger, and Ishizuka 2011) and create highly comprehensive lexicons (i.e. in conjunction with possible compositionality rules) (Neviarouskaya, Prendinger, and Ishizuka 2009) to perform sentiment analysis on text at sentence and phrase level. Machine learning approaches are also frequently applied to the sentiment analysis problem, such as linguistically inspired deep neural networks (Socher et al. 2013), graphical models (Nakagawa, Inui, and Kurohashi 2010) and classical supervised learning techniques (Alm, Roth, and Sproat 2005). In addition, micro-blogging websites also get the attention of researchers for sentiment analysis (Pak and Paroubek 2010; Kouloumpis, Wilson, and Moore 2011). In these studies, emoticons, as a different modality, are used for sentiment polarity recognition. In this paper, we discuss and compare our model to the systems presented at SemEval-2007, where the same corpus of news headlines was used. Our main emphasis is on affect of named-entities and prevalent compositionality in headlines.

Five teams participated to SemEval-2007 Task-14: Affective Text, with five systems for valence classification and three systems for emotion labeling. The CLAC system used a knowledge-based, domain-independent, unsupervised approach (Andreevskaia and Bergler 2007). It uses 3 knowledge sources, a list of sentiment-bearing words, a list of polarity reversal words, and a set of rules that define the results of combination of sentiment-bearing words with polarity reversal words. The CLAC-NB system uses a Naive Bayes classifier without feature extraction to assess the performance of this basic machine learning technique (Andreevskaia and Bergler 2007). The UPAR7 system (Chaumartin 2007) used a rule-based approach, where a list of words for high tech acronyms, celebrities was used in addition to SentiWordNet (Esuli and Sebastiani 2006) and WordNetAffect (Strapparava and Valitutti 2004) lexical resources.The SICS used a valence annotation approach based on a word-space model and set of seed words which is based on the idea of creating two points in a high-dimensional word space, one representing positive valence, the other representing negative valence, and projecting each headline into this space (Sahlgren, Karlgren, and Eriksson 2007). The SWAT system used a supervised methodology by implementing an uni-gram model trained to annotate. Moreover, they added an additional set of 1000 headlines for training (Katz, Singleton, and Wicentowski 2007). Therefore, the demonstrated systems in SemEval-2007 (Strapparava and Mihalcea 2007), did not consider the contribution of sentiment-bearing named-entities. Since the news headlines very commonly contain proper nouns, there is a good chance to improve the reported results, if the proper names are used in sentiment polarity determination.

Approach

We used five different approaches for classification after performing the feature extraction. We are using the Stanford dependency parser (De Marneffe et al. 2006), the part-of-speech tagger (Toutanova et al. 2003) and the named-entity recognizer (Finkel, Grenager, and Manning 2005) as software packages for feature extraction. We use MPQA Opinion Corpus (Wilson, Wiebe, and Hoffmann 2005) and General Inquirer (Stone, Dunphy, and Smith 1966) as lexical resources to identify sentiment orientation of words. To process the interaction between words, our system uses

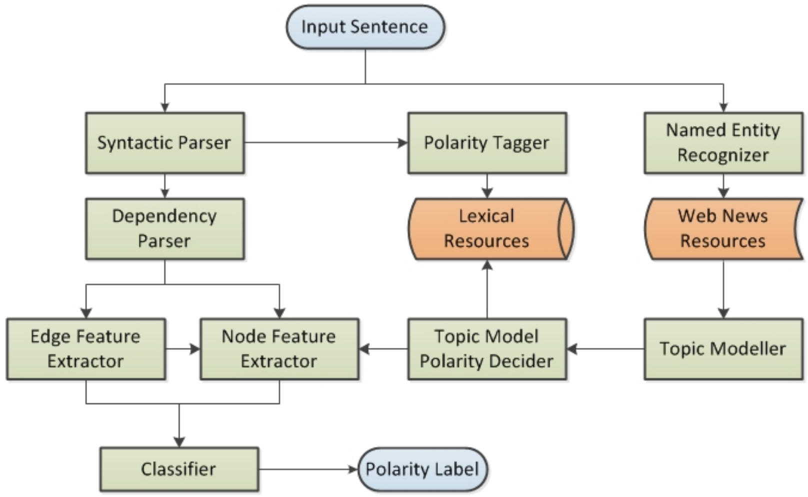

Figure 2: System Architecture

dependency parsing. The named-entity recognizer is used to recognize locations, organizations (e.g. company names, acronyms) and person names. Then, the syntactic parser is used for word-level analysis. The flow of information among the components of the system is depicted in Figure 2.

Named-Entities and Topic Modeling

We observed that most of the news headlines (82% of SemEval DataSet) contain named-entities which influence the overall sentiment polarity of each headline. Therefore, we needed to consider the influence of a named-entity to the overall polarity for an accurate sentiment classification. We are using the Stanford named-entity recognizer to find locations (e.g. Middle East, Gaza), organization names (e.g. Apple, Google, European Union) and person names (e.g. Obama, Madonna), all of which are named-entities. However, the real difficulty is, there is no lexical resource to annotate identified named-entities with the polarity information. To address this problem of finding the sentiment valence for named entities, we performed topic modeling on web news resources (i.e. Google News, CNN). As we showed in Figure 2, with each identified named-entity, the system searches on the web to find news articles and gathers top 30 most relevant news articles to the searched named-entity. Then the system performs topic modeling on the news articles retrieved from the web.

Topic modeling provides a simple way to analyze large volumes of unlabeled text. A ”topic” consists of a cluster of words that frequently occur together. Using contextual clues, topic models can connect words with similar meanings and distinguish between uses of words with multiple meanings. We used the Mallet[1] library for topic modeling.

A topic model with a cluster of 30 words derived from 30 news articles was retrieved from the web. In other words, we represented the recognized named-entity with 30 words to decide the polarity of the named-entity. The cluster of words was tagged using the sentiment polarity lexical resources. The topic model polarity decider component selects the aggregated overall polarity of a topic which actually represents the recognized named-entity. We are using a simple algorithm to decide the overall polarity. If the polarity of a word is positive, it adds 1. If it is negative it subtracts 1, if it is neutral, it does nothing. If the aggregated polarity is greater than 0, the polarity of named-entity is positive. If it is less than 0, the polarity of named-entity is negative.

For example, in Figure 1, *Gaza* (city name) is identified as a named-entity. Since it is a proper noun, it does not exist in our lexical resources. Our system would search the web to retrieve 30 news articles about *Gaza* and perform topic modeling on them to obtain the related word set. Our topic modeling component identified *”Israel, air strike, attack, bomb, war, Hamas, terror, refugee, crisis, rocket”* as top 10 words. As previously stated, our system searches for 30 words, all of which are then annotated using our lexical resources. Afterwards, the polarity score for the *Gaza* named-entity is automatically determined as negative using the previous algorithm.

Features and Interactions Between Words

The polarity of a whole sentence can not be calculated without considering interactions between words, as it can be seen in Figure 1 for the sentence ”*Kidnapped AP photographer freed in Gaza*”. There are two negative words (kidnapped, Gaza) and one positive polarity reversal word (freed). If the text is analyzed without considering interactions between words and dependencies, it is very likely that the results might end up with a negative classification. However, as it can be seen in Figure 1, even though sub-trees bears negative sentiment, the word ”freed” reverses the polarity of the whole sentence.

Our system uses two types of features, node features and edge features. Node features represent each node in dependency tree, and edge features represent the interactions between words. Node features include the prior polarity of a

[1] http://mallet.cs.umass.edu/

Node Features	
a	t_i
b	$t_i \& m_i$
c	$t_i \& m_i \& r_i$
d	$t_i \& s_i$
e	$t_i \& c_i$
f	$t_i \& f_i$

Table 1: Node Features

word which can be obtained from our lexical resources, polarity reversal, resulting polarity and pos-tag. Edge features represent interaction between sub-trees. For example, in Figure 1, kidnapped is negative, AP is neutral which is an organization name, photographer is neutral, the negative polarity of *kidnapped* propagates to the sub-tree where *photogtapher* is the root. Our system not only uses the tree structure for word interactions, but also uses dependencies. For this case, the first dependency, directly shows the relationship between the words *photographer* and *kidnapped*. For the entity *Gaza*, our system uses topic modeling, explained in the *Named-Entities and Topic Modeling* section that yielded a negative polarity. However, the word *freed* reverses the meaning of the whole sentence.

The polarity reversal feature of a node $r_i \epsilon \{0, 1\}$, represents whether or not it reverses the polarity of other words. A polarity reversal word list was prepared so that the property r_i in identified words it is set to 1 otherwise it is 0. The described list was constructed from General Inquirer in the same methodology used in (Choi and Cardie 2008). We collected words which belong to either NOTLW or DECREAS categories from General Inquirer (the dictionary contains 121 polarity reversing words). Choi and Cardie categorized polarity reversal words into two distinct categories: function-word negators, such as *not*, and content-word negators, such as *eliminate*. The polarity reversal of a phrase handles only content-word negators, and function-word negators are handled based on the result of Stanford dependency parser which gives the negation (*neg*) relation directly.

In Table 1 and Table 2, we show the features we used in this project. Features (a)-(f) in Table 1 represents node features for the *i*-th word, which is a node in the dependency tree. In Table 1, t_i denotes polarity of a node, m_i denotes prior polarity of a node. Prior polarity of a node $m_i \epsilon \{+1, 0, -1\}$ is the innate sentiment of a word obtained from the polarity lexical resources. As described before r_i represents polarity reversal word, s_i denotes surface form, c_i denotes coarse-grained part-of-speech (POS) tag, f_i denotes fine-grained POS tag. Features (A)-(D) in Table 2 represents dependency features for each dependency governor (head) and its dependent. In Table 2, t_g denotes polarity of governor, t_d denotes polarity of dependent, r_d polarity reversal word for dependent, m_d denotes prior polarity as in node features table, c_g denotes coarse-grained pos tag for governor, c_d denotes coarse-grained polarity tag for dependent and R denotes the dependency relationship type (e.g. negation).

After feature extraction including named-entities and interactions between words, we used rule-based and supervised systems to decide the polarity label of a sentence and compare the results.

Dependency Features	
A	$t_g \& t_d$
B	$t_g \& t_d \& r_d$
C	$t_g \& t_d \& r_d \& m_d$
D	$t_g \& t_d \& c_g \& c_d \& R$

Table 2: Dependency Features

Experiments

We conducted the experiments of sentiment classification in SemEval 2007 Affective text news headline corpora.

Data

The dataset consisted of news headlines from major news resources such as New York Times, CNN, and BBC News. The headlines are collected for two main reasons. First, the high load of emotional content, as they describe major national or worldwide events and are written in a style meant to attract the attention of the readers. Second, the structure of headlines was appropriate to make sentence-level annotations of emotions. Two different annotated datasets were made available: one is a development dataset consisting of 250 headlines, and the other with 1,000 headlines. For our experiments we only use valence (positive/negative annotation). The interval for the valence annotations was set to [-100, 100], where 0 represents a neutral headline, -100 represents a highly negative headline, and 100 corresponds to a highly positive headline. Even though, annotations are fine-grained, we used coarse-grained annotations with positive [50, 100], negative [-100, -50] and neutral (-50, 50). Moreover, we used 750 headlines for training and 500 headlines for testing.

It was reported that the test dataset was independently labeled by six annotators, who were instructed to select the appropriate emotions for each headline based on the presence of words or phrases with emotional content, as well as the sentiment polarity invoked by the headline. The agreement evaluations were conducted using the Pearson correlation measure, with an inter-annotator agreement of *78.01*.

Compared Methods

We compared five methods with a different set of features, 3 rule-based methods as used in (Nakagawa, Inui, and Kurohashi 2010) without using major polarity in training data and 2 supervised classification methods (Naive Bayes and Maximum Entropy) described below. A_i denotes the set of all the ancestor nodes of i-th word in the dependency tree, and val(x) is defined as:

$$val(x) = \begin{cases} +1, & (x > 1), \\ -1, & (x < 0), \\ 0 & (x = 0). \end{cases}$$

Method	Acc.	Prec.	Rec.	F1
Voting w/o PR	0.6613	0.4916	0.5378	0.5103
Voting w/ PR	0.6253	0.4553	0.5191	0.4851
Rule	0.644	0.4855	0.6195	0.5394
Naïve Bayes	0.6606	0.4921	0.6461	0.5564
Maximum Entropy	0.7673	0.5849	0.7411	0.6518

Table 3: Experiment Results. First column is method used, second column is accuracy, third column is precision, fourth column is recall and fifth column is F-1 score. PR stands for polarity reversal.

Voting without Polarity Reversal The polarity of the headline is decided by voting of each node's prior/innate polarity which also includes polarity of each named-entity which is obtained through topic modeling.

$$p = val(\sum_{i=1}^{n} m_i) \tag{1}$$

Voting with Polarity Reversal Similar to Voting without polarity reversal, except that the polarities of phrases which have odd numbers of reversal phrases in their ancestors are reversed before voting.

$$p = val(\sum_{i=1}^{n} m_i \prod_{j \epsilon A_i} (-1)^{r_j}) \tag{2}$$

Rule The sentiment polarity of a headline is deterministically decided based on rules, by considering the polarities of sub-trees. The polarity of the sub-tree whose ancestor is the i-th word is decided by voting the prior polarity of the i-th word and the polarities of the sub-trees whose ancestor nodes are the modifiers of the i-th word. The polarities of the modifiers are reversed if their governor phrase had a reversal word. The decision rule is applied from bottom to top, the polarity of root node is decided at last.

$$p = val(m_i + \sum_{j:h_j=i} t_i(-1)^{r_i}) \tag{3}$$

Supervised Classification We have used the Naive Bayes and Maximum Entropy methods for classification using the Mallet machine learning package. It is important to note that we used 750 news headlines as a training data and 500 as testing data.

Experiment Results

We have conducted the experiments using five different approaches. We presented our results in Table 3. The table columns from left to right shows, method used to decide polarity of a headline, accuracy, precision, recall and F1 measure. The performance of each method is measured with accuracy and F_β rate as shown in Equation 4 (Rijsbergen 1979).

$$F_\beta = \frac{(\beta^2 + 1) * precision * recall}{\beta^2 * precision + recall} \tag{4}$$

The performance of the rule-based approaches is close to each other, the Naive Bayes approach performed slightly better than the rule-based approaches. The Maximum Entropy classifier achieved better results compared to the other 4 methods.

When we compare our results with the accuracy of the systems presented in SemEval-2007 Affective Text task (Strapparava and Mihalcea 2007), our methodology outperformed the demonstrated systems. The highest reported accuracy and F1 score in SemEval are *0.5510* and *0.4243* respectively. Our best performing system's (i.e. maximum-entropy classifier) accuracy and F1 score are *0.7673* and *0.6518* (see Table 3) respectively. We believe that the main factor in improvement is taking into account the polarity information of the named-entities by performing topic modeling on news articles. It is easily observable that in news headlines there are many proper names and named entities, and the number of the words in each headline is little. If a system does not use sentiment polarity information for a named-entity, it skips important information for overall polarity of a headline. As we discuss in the *Named-Entities and Topic Modeling* Section, upon running the Stanford named-entity recognizer on the trial data provided in SemEval-2007 for example, we find that 82% percent of the headlines contain at least one named-entity.

Conclusion

In this paper, we presented our approach and experiments to perform sentiment polarity classification on the SemEval-2007 Affective Text data. Our system used some of the available sentiment polarity lexical resources and polarity reversal words for feature extraction. Also we performed topic modeling to decide sentiment polarity of each named-entity. Our approach shows that the consideration of named-entities has a positive effect on sentiment polarity classification. Our approaches outperformed the systems presented in SemEval-2007. Although we did not test it, we believe that our system can be used in other domains for sentence level sentiment polarity classification. As a future research, we are planning to test our system on different domains, such as movies reviews datasets.

Acknowledgments

This material is based in part upon work supported by the National Science Foundation under Grant Nos. SBIR IIP-1256641, SBIR IIP-1118610, I/UCRC IIP-1338922, AIR IIP-1237818, CREST HRD-0833093, I/UCRC IIP-0829576.

References

Alm, C. O.; Roth, D.; and Sproat, R. 2005. Emotions from text: machine learning for text-based emotion prediction. In *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, 579–586. Association for Computational Linguistics.

Andreevskaia, A., and Bergler, S. 2007. Clac and clac-nb: Knowledge-based and corpus-based approaches to sentiment tagging. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, 117–120. Prague, Czech Republic: Association for Computational Linguistics.

Chaumartin, F.-R. 2007. Upar7: A knowledge-based system for headline sentiment tagging. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, 422–425. Prague, Czech Republic: Association for Computational Linguistics.

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*, 793–801. Association for Computational Linguistics.

De Marneffe, M.-C.; MacCartney, B.; Manning, C. D.; et al. 2006. Generating typed dependency parses from phrase structure parses. In *Proceedings of LREC*, volume 6, 449–454.

Esuli, A., and Sebastiani, F. 2006. Sentiwordnet: A publicly available lexical resource for opinion mining. In *Proceedings of LREC*, volume 6, 417–422.

Finkel, J. R.; Grenager, T.; and Manning, C. 2005. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, 363–370. Association for Computational Linguistics.

Katz, P.; Singleton, M.; and Wicentowski, R. 2007. Swat-mp:the semeval-2007 systems for task 5 and task 14. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, 308–313. Prague, Czech Republic: Association for Computational Linguistics.

Kouloumpis, E.; Wilson, T.; and Moore, J. 2011. Twitter sentiment analysis: The good the bad and the omg! In *ICWSM*.

Li, T.; Zhang, Y.; and Sindhwani, V. 2009. A non-negative matrix tri-factorization approach to sentiment classification with lexical prior knowledge. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 1 - Volume 1*, ACL '09, 244–252. Stroudsburg, PA, USA: Association for Computational Linguistics.

Nakagawa, T.; Inui, K.; and Kurohashi, S. 2010. Dependency tree-based sentiment classification using crfs with hidden variables. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, 786–794. Association for Computational Linguistics.

Neviarouskaya, A.; Prendinger, H.; and Ishizuka, M. 2009. Sentiful: Generating a reliable lexicon for sentiment analysis. In *Affective Computing and Intelligent Interaction and Workshops, 2009. ACII 2009. 3rd International Conference on*, 1–6. IEEE.

Neviarouskaya, A.; Prendinger, H.; and Ishizuka, M. 2011. Affect analysis model: novel rule-based approach to affect sensing from text. *Natural Language Engineering* 17(1):95–135.

Pak, A., and Paroubek, P. 2010. Twitter as a corpus for sentiment analysis and opinion mining. In *LREC*.

Pang, B., and Lee, L. 2008. Opinion mining and sentiment analysis. *Foundations and trends in information retrieval* 2(1-2):1–135.

Rijsbergen, C. J. V. 1979. *Information Retrieval*. Newton, MA, USA: Butterworth-Heinemann, 2nd edition.

Sahlgren, M.; Karlgren, J.; and Eriksson, G. 2007. Sics: Valence annotation based on seeds in word space. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, 296–299. Prague, Czech Republic: Association for Computational Linguistics.

Socher, R.; Perelygin, A.; Wu, J. Y.; Chuang, J.; Manning, C. D.; Ng, A. Y.; and Potts, C. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1631–1642.

Stone, P. J.; Dunphy, D. C.; and Smith, M. S. 1966. The general inquirer: A computer approach to content analysis.

Strapparava, C., and Mihalcea, R. 2007. Semeval-2007 task 14: Affective text. In *Proceedings of the 4th International Workshop on Semantic Evaluations*, 70–74. Association for Computational Linguistics.

Strapparava, C., and Valitutti, A. 2004. Wordnet affect: an affective extension of wordnet. In *LREC*, volume 4, 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*, 423–426.

Toutanova, K.; Klein, D.; Manning, C. D.; and Singer, Y. 2003. Feature-rich part-of-speech tagging with a cyclic dependency network. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology-Volume 1*, 173–180. Association for Computational Linguistics.

Turunen, M.; Hakulinen, J.; Ståhl, O.; Gambäck, B.; Hansen, P.; Rodríguez Gancedo, M. C.; de la Cámara, R. S.; Smith, C.; Charlton, D.; and Cavazza, M. 2011. Multimodal and mobile conversational health and fitness companions. *Computer Speech & Language* 25(2):192–209.

Wilson, T.; Wiebe, J.; and Hoffmann, P. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, 347–354. Association for Computational Linguistics.

Yasavur, U.; Lisetti, C.; and Rishe, N. 2013. Modeling brief alcohol intervention dialogue with mdps for delivery by ecas. In *Intelligent Virtual Agents*, 92–105. Springer.

Yi, J., and Niblack, W. 2005. Sentiment mining in webfountain. In *Data Engineering, 2005. ICDE 2005. Proceedings. 21st International Conference on*, 1073–1083. IEEE.