Using Roget's Thesaurus for Fine-grained Emotion Recognition

Saima Aman School of Information Technology and Engineering University of Ottawa, Ottawa, Canada

saman071@site.uottawa.ca

Abstract

Recognizing the emotive meaning of text can add another dimension to the understanding of text. We study the task of automatically categorizing sentences in a text into Ekman's six basic emotion categories. We experiment with corpus-based features as well as features derived from two emotion lexicons. One lexicon is automatically built using the classification system of Roget's Thesaurus, while the other consists of words extracted from WordNet-Affect. Experiments on the data obtained from blogs show that a combination of corpus-based unigram features with emotion-related features provides superior classification performance. We achieve Fmeasure values that outperform the rulebased baseline method for all emotion classes.

1 Introduction

Recognizing emotions conveyed by a text can provide an insight into the author's intent and sentiment, and can lead to better understanding of the text's content. Emotion recognition in text has recently attracted increased attention of the NLP community (Alm et al., 2005; Liu et al, 2003; Mihalcea and Liu, 2006); it is also one of the tasks at Semeval-2007¹.

Automatic recognition of emotions can be applied in the development of affective interfaces for

Stan Szpakowicz School of Information Technology and Engineering University of Ottawa, Ottawa, Canada ICS, Polish Academy of Sciences Warszawa, Poland szpak@site.uottawa.ca

Computer-Mediated Communication and Human-Computer Interaction. Other areas that can potentially benefit from automatic emotion analysis are personality modeling and profiling (Liu and Maes, 2004), affective interfaces and communication systems (Liu et al, 2003; Neviarouskaya et al., 2007a) consumer feedback analysis, affective tutoring in e-learning systems (Zhang et al., 2006), and textto-speech synthesis (Alm et al., 2005).

In this study, we address the task of automatically assigning an emotion label to each sentence in the given dataset, indicating the predominant emotion type expressed in the sentence. The possible labels are happiness, sadness, anger, disgust, surprise, fear and no-emotion. Those are Ekman's (1992) six basic emotion categories, and an additional label to account for the absence of a clearly discernible emotion.

We experiment with two types of features for representing text in emotion classification based on machine learning (ML). Features of the first type are a corpus-based unigram representation of text. Features of the second type comprise words that appear in emotion lexicons. One such lexicon consists of words that we automatically extracted from Roget's Thesaurus (1852). We chose words for their semantic similarity to a basic set of terms that represent each emotion category. Another lexicon builds on lists of words for each emotion category, extracted from WordNet-Affect (Strapparava and Valitutti, 2004).

We compare the classification results for groups of features of these two types. We get good results when the features are combined in a series of ML experiments.

¹ Affective Text: Semeval Task at the 4th International Workshop on Semantic Evaluations, 2007, Prague (nlp.cs.swarthmore.edu/semeval/tasks/task14/summary.shtml).

2 Related Work

Research in emotion recognition has focused on discerning emotions along the dimensions of valence (positive / negative) and arousal (calm / excited), and on recognizing distinct emotion categories. We focus on the latter.

Liu et al. (2003) use a real-world commonsense knowledge base to classify sentences into Ekman's (1992) basic emotion categories. They use an ensemble of rule-based affect models to determine the emotional affinity of individual sentences. Neviarouskaya et al. (2007b) also use rules to determine the emotions in sentences in blog posts; their analysis relies on a manually prepared database of words, abbreviations and emoticons labeled with emotion categories.

Since these papers do not report conventional performance metrics such as precision and recall, the effectiveness of their methods cannot be judged empirically. They also disregard statistical learning methods as ineffective for emotion recognition at sentence level. They surmise that the small size of the text input (a sentence) gives insufficient data for statistical analysis, and that statistical methods cannot handle negation. In this paper, we show that ML-based approach with the appropriate combination of features can be applied to distinguishing emotions in text.

Previous work has used lexical resources such as WordNet to automatically acquire emotion-related words for emotion classification experiments. Starting from a set of primary emotion adjectives, Alm et al. (2005) retrieve similar words from WordNet utilizing all senses of all words in the synsets that contain the adjectives. They also exploit the synonym and hyponym relations in WordNet to manually find words similar to nominal emotion words. Kamps and Marx (2002) use WordNet's synset relations to determine the affective meaning of words. They assign multidimensional scores to individual words based on the minimum path length between them and a pair of polar words (such as "good" and "bad") in WordNet's structure.

There is also a corpus-driven method of determining the emotional affinity of words: learn probabilistic affective scores of words from large corpora. Mihalcea and Liu (2006) have used this method to assign a happiness factor to words depending on the frequency of their occurrences in happy-labeled blogposts compared to their total frequency in the corpus.

In this paper, we study a new approach to automatically acquiring a wide variety of words that express emotions or emotion-related concepts, using *Roget's Thesaurus* (1852).

3 Emotion-Labeled Data

We have based our study on data collected from blogs. We chose blogs as data source because they are potentially rich in emotion content, and contain good examples of real-world instances of emotions expressed in text. Additionally, text in blogs does not conform to the style of any particular genre per se, and thus offers a variety in writing styles, choice and combination of words, as well as topics. So, the methods learned for discerning emotion using blog data are quite general and therefore applicable to a variety of genres rather than to blogs only.

We retrieved blogs using seed words for all emotion categories. Four human judges manually annotated the blog posts with emotion-related information - every sentence received two judgments. The annotators were required to mark each sentence with one of the eight labels: happiness, sadness, anger, disgust, surprise, fear, mixed-emotion, and no-emotion. The mixedemotion label was included to handle those sentences that had more than one type of emotion or whose emotion content could not fit into any of the given emotion categories. Sample sentences from the annotated corpus are shown in Fig. 1.

We measured the inter-annotator agreement using Cohen's (1960) kappa. The average pair-wise agreement for different emotion categories ranged from 0.6 to 0.79. In the experiments reported in this paper, we use only those sentences for which there was agreement between both judgments (to form a benchmark for the evaluation of the results of automatic classification). The distribution of emotion categories in the corpus used in our experiments is shown in Table 1.

This was the best summer I have	
ever experienced.	(happiness)
I don't feel like I ever have that	
kind of privacy where I can talk	
to God and cry and figure things	
out.	(sadness)
Finally, I got fed up.	(disgust)
I can't believe she is finally here!	(surprise)

Fig 1. Sample sentences from the corpus

Emotion Class	Number of sentences
Happiness	536
Sadness	173
Anger	179
Disgust	172
Surprise	115
Fear	115
No-emotion	600

 Table 1. Distribution of emotion classes

4 A Baseline Approach

We are interested in investigating if emotion in text can be discerned on the basis of its lexical content. A naïve approach to determining the emotional orientation of text is to look for obvious emotion words, such as "happy", "afraid" or "astonished". The presence of one or more words of a particular emotion category in a sentence provides a good premise for interpreting the overall emotion of the sentence. This approach relies on a list of words with prior information about their emotion type, and uses it for sentence-level classification. The obvious advantage is that no training data are required.

For evaluation purposes, we took this approach to develop a baseline system that counts the number of emotion words of each category in a sentence, and then assigns this sentence the category with the largest number of words. Ties were resolved by choosing the emotion label according to an arbitrarily predefined ordering of emotion classes. A sentence containing no emotion word of any type was assigned the *no emotion* category. This system worked with word lists² extracted from *WordNet-Affect* (Strapparava and Valitutti, 2004) for six basic emotion categories.

Table 2 shows the precision, recall, and Fmeasure values for the baseline system. As we have seven classes in our experiments, the class imbalance makes accuracy values less relevant than precision, recall and F-measure. That is why we do not report accuracy values in our results.

The baseline system shows precision values above 50% for all but two classes. This shows the usefulness of this approach. This method, however, fails in the absence of obvious emotion words in the sentence, as indicated by low recall values. Thus, in order to improve recall, we need to increase the ambit of words that are considered emotion-related. An alternative approach is to use ML to learn automatically rules that classify emotion in text.

Class	Precision	Recall	F-Measure
Happiness	0.589	0.390	0.469
Sadness	0.527	0.283	0.368
Anger	0.681	0.262	0.379
Disgust	0.944	0.099	0.179
Surprise	0.318	0.296	0.306
Fear	0.824	0.365	0.506
No-emotion	0.434	0.867	0.579

 Table 2. Performance metrics of the baseline system

5 Approach Based on Machine Learning

We study two types of features: corpus-based features and features based on emotion lexicons.

5.1 Corpus-based features

The corpus-based features exploit the statistical characteristics of the data on the basis of the ngram distribution. In our experiments, we take unigrams (n=1) as features. Unigram models have been previously shown to give good results in sentiment classification tasks (Kennedy and Inkpen, 2006; Pang et al., 2002): unigram representations can capture a variety of lexical combinations and distributions, including those of emotion words. This is particularly important in the case of blogs, whose language is often characterized by frequent use of new words, acronyms (such as "lol"), onomatopoeic words ("haha", "grrr"), and slang, most of which can be captured in a unigram representa-

² Emotion words from WordNet-Affect (http://www.cse.unt.edu/~rada/affectivetext/data/WordNet AffectEmotionLists.tar.gz)

Similarity Score	Happiness	Sadness	Anger	Disgust	Surprise	Fear
16	family, home, friends, life, house, loving, partying, bed, pleasure, rest, close, event, lucks, times	crying, lost, wounds, bad, pills, falling, messed, spot, unhappy, pass, black, events, hurts, shocked	pride, fits, stormed, abandoned, bothered, mental, an- ger, feelings, distractions	shock, dis- gust, dislike, loathing	plans, catch, expected, early, slid, slipped, ear- lier, caught, act	nervous, cry, terror, panic, feelings, run, fog, fire, turn, police, faith, battle, war, sounds
14	love, like, feel, pretty, lovely, better, smiling, nice, beautiful, hope, cutest celebrations, warm, desires	ill, bored, feeling, ruin, blow, down, wrong, awful, evil, worry, crushing, bug, death, trouble, dark	hate, burn, upset, dislike, wrong, blood, ill, flaws, bar, defects, bit- ter, growled, black, slow	hate, pain, horrifying, ill, pills, sad, wear, blood, appalling, end, work, weighed, regrets, bad	left, swing, noticed, worry, times, amazing, stolen, break, interesting, attention	falling, life, stunned, pay, broken, hate, blast, times, hanging, hope, broken, blood, blue
12	gift, treats, adorable, fun, hug, kidding, bigger, great, lighting, won, stars, enjoy, favourite, social, divine	defeat, nasty, boring, ugly, loser, end, victim, sick, hard, serious, aggravating, bothering, burning	lose, throw, offended, hit, power, feel, flaring, pills, broken, life, forgot, rant- ing	feel, fun, lies, drawn, lose, missed, de- prived, lack, sighs, defeat, down, hurt, tears, insulted	realize, pick, wake, sense, jumped, new, late, magic, omen, forget, popped, feel, question, late, throw	fearful, spy, night, upset, feel, chased, hazardous, tomorrow, victim, grim, terrorists, apprehensive

Table 3. Emotion-related words automatically extracted from Roget's Thesaurus

tion. Another advantage of a unigram representation is that it does not require any prior knowledge about the data under investigation or the classes to be identified.

For our experiments, we selected all unigrams that occur more than three times in the corpus. This eliminates rare words, as well as foreign-language words and spelling mistakes, which are quite common in blogs. We also excluded words that occur in a list of stopwords - primarily function words that do not generally have emotional connotations. We used the SMART list of stopword³, with minor modifications. For instance, we removed from the stop list words such as "what" and "why", which may be used in the context of expressing *surprise*.

5.2 Features derived from *Roget's Thesaurus*

We utilized *Roget's Thesaurus* (Jarmasz and Szpakowicz, 2001) to automatically build a lexicon

of emotion-related words. The features based on an emotion lexiconrequire prior knowledge about emotion relatedness of words. We extracted this knowledge from the classification system in *Roget's*, which groups related concepts into various levels of a hierarchy. For a detailed account of this classification structure, see Jarmasz and Szpakowicz (2001).

Roget's structure allows the calculation of semantic relatedness between words, based on the path length between the nodes in the structure that represent those words. In case of multiple paths, the shortest path is considered. Jarmasz and Szpakowicz (2004) have introduced a similarity measure derived from path length, which assigns scores ranging from a maximum of 16 to most semantically related words to a minimum of 0 to least related words. They have shown that on semantic similarity tests this measure outperforms several other methods.

To build a lexicon of emotion-related words utilizing *Roget's* structure, we need first to make two decisions: select a primary set of emotion words starting with which we can extract other similar

³ SMART stopwords list. Used with the SMART information retrieval system at Cornell University

⁽ftp://ftp.cs.cornell.edu/pub/smart/english.stop)

Model	Class	Precision	Recall	F-Measure	Baseline F-Measure
	Happiness	0.840	0.675	0.740	0.469
Unigrams	Sadness	0.619	0.301	0.405	0.368
	Anger	0.634	0.358	0.457	0.379
	Disgust	0.772	0.453	0.571	0.179
	Surprise	0.813	0.339	0.479	0.306
	Fear	0.889	0.487	0.629	0.506
	No-emotion	0.581	0.342	0.431	0.579
Roget's Thesaurus	Happiness	0.772	0.562	0.650	0.469
	Sadness	0.574	0.225	0.324	0.368
	Anger	0.638	0.246	0.355	0.379
	Disgust	0.729	0.297	0.421	0.179
(ICI) I catales	Surprise	0.778	0.243	0.371	0.306
	Fear	0.857	0.470	0.607	0.506
	No-emotion	0.498	0.258	0.340	0.579
Unigrams + RT Features	Happiness	0.809	0.705	0.754	0.469
	Sadness	0.577	0.370	0.451	0.368
	Anger	0.636	0.419	0.505	0.379
	Disgust	0.686	0.471	0.559	0.179
	Surprise	0.717	0.374	0.491	0.306
	Fear	0.831	0.513	0.634	0.506
	No-emotion	0.586	0.512	0.546	0.579
Unigrams + RT Features + WNA Features	Happiness	0.813	0.698	0.751	0.469
	Sadness	0.605	0.416	0.493	0.368
	Anger	0.650	0.436	0.522	0.379
	Disgust	0.672	0.488	0.566	0.179
	Surprise	0.723	0.409	0.522	0.306
	Fear	0.868	0.513	0.645	0.506
	No-emotion	0.587	0.625	0.605	0.579

* Highest precision, recall, and F-measure values for each class are shown in bold

Table 4 ML Classification Results

words, and choose an appropriate similarity score to serve as cutoff for determining semantic relatedness between words.

The primary set of words that we selected consists of one word for each emotion category, representing the base form of the name of the category: {*happy, sad, anger, disgust, surprise, fear*}.

Experiments performed on Miller and Charles similarity data (1991), reported in Jarmasz and Szpakowicz (2004), have shown that pairs of words with a semantic similarity value of 16 have high similarity, while those with a score of 12 to 14 have intermediate similarity. Therefore, we select the score of 12 as cutoff, and include in the lexicon all words that have similarity scores of 12 or higher with respect to the words in the primary set. This selection of cutoff therefore serves as a form of feature selection. In Table 3, we present sample words from the lexicon with similarity scores of 16, 14, and 12 for each emotion category. These words represent three different levels of relatedness to each emotion category. We are able to identify a large variety of emotion-related words belonging to different parts of speech that go well beyond the stereotypical words associated with different emotions. We particularly note some generic neutral words, such as "feel", "life", and "times" associated with many emotion categories, indicating their conceptual relevance to emotions.

5.3 Features derived from *WordNet-Affect*

WordNet-Affect is an affective lexical resource that assigns a variety of affect-related labels to a subset

of *WordNet* synsets comprising affective concepts. We used lists of words extracted from it for each of the six emotion categories.

6 Experiments and Results

We train classifiers with unigram features for each emotion class using Support Vector Machine (SVM) for predicting the emotion category of the sentences in our corpus. SVM has been shown to be useful for text classification tasks (Joachims, 1998), and has previously given good performance in sentiment classification experiments (Kennedy and Inkpen, 2006; Mullen and Collier, 2004; Pang and Lee, 2004; Pang et al., 2002). In Table 4, we report results from ten-fold cross-validation experiments conducted using the SMO implementation of SVM in Weka (Witten and Frank, 2005). In each experiment, we represent a sentence by a vector indicating the number of times each feature occurs.

In the first experiment, we use only corpusbased unigram features. We obtain high precision values for all emotion classes (as shown in Table 4), and the recall and F-measure values surpass baseline values for all classes except no-emotion. This validates our premise that unigrams can help learn lexical distributions well to accurately predict emotion categories.

Next, we use as features all words in the emotion lexicon acquired from *Roget's Thesaurus* (RT). The F-measure scores beat the baseline for four out of seven classes. When we combine both corpus-based unigrams with RT features, we can increase recall values across all seven classes.

Finally, we add features from WordNet-Affect to the feature set containing corpus unigrams and RT features. This leads to further improvement in overall performance. Combining all features, we achieve highest recall values across all but one class. The resulting F-measure values (ranging from 0.493 to 0.751) surpass the baseline values across all seven classes. This increase was found to be statistically significant (paired t-test, p=0.05).

7 Discussion

We observe that corpus-based features and emotion-related features together contribute to improved performance, better than given by any one type of feature group alone.

Any automatic way of recognizing emotion should inevitably take into account a wide variety of words that are semantically connected to emotions. While some words are obviously affective, many more are only potentially affective. The latter derive their affective property from their associations with emotional concepts. For instance, words like "family", "friends", "home" are not inherently emotional, but because of their wellknown semantic association with emotion concepts, their presence in a sentence can be taken as an indicator of emotion expression in the sentence. We can interpret the results as indicators of how much correlation the classifiers can find between the features and the predicted class. Considering our best results using all features, we find that this correlation is highest for the "happy" class, indicated by a precision of 0.813 and recall of 0.698, the highest among all classes. We can therefore conclude that it is easier to discern happiness in text than Ekman's other basic emotions.

8 Conclusions

Working on a corpus of blog sentences annotated with emotion labels, we were able to demonstrate that a combination of corpus-based unigram features and features derived from emotion lexicons can help automatically distinguish basic emotion categories in written text. When used together in an SVM-based learning environment, these features increased recall in all cases and the resulting F-measure values significantly surpassed the baseline scores for all emotion categories.

In addition, we described a method of building an emotion lexicon derived from *Roget's Thesaurus* on the basis of semantic relatedness of words to a set of basic emotion words for each emotion category. The effectiveness of this emotion lexicon was demonstrated in the emotion classification tasks.

References

- Cecilia O. Alm, Dan Roth, and Richard Sproat, Emotions from text: machine learning for text-based emotion prediction. In *Proceedings of Joint Conference on HLT/EMNLP*, pages 579-586, Vancouver, Canada, Oct 2005.
- M.M. Bradley and P.J. Lang, *Affective norms for English words (ANEW): Instruction manual and affective*

ratings, Technical Report C-1, The Center for Research in Psychophysiology, University of Florida, 1999.

- J. Cohen, A coefficient of agreement for nominal scales, *Educational and Psychological Measurement*, 1960, 20 (1): 37–46.
- Paul Ekman, An Argument for Basic Emotions, Cognition and Emotion, 6, 1992, 169-200.
- Mario Jarmasz and Stan Szpakowicz, The Design and Implementation of an Electronic Lexical Knowledge Base. In *Proceedings of the 14th Biennial Conference of the Canadian Society for Computational Studies of Intelligence (AI 2001)*, Ottawa, Canada, June 2001, 325-333.
- Mario Jarmasz and Stan Szpakowicz, Roget's Thesaurus and Semantic Similarity. N. Nicolov, K. Bontcheva, G. Angelova, R. Mitkov (eds.) Recent Advances in Natural Language Processing III: Selected Papers from RANLP 2003, John Benjamins, Amsterdam/Philadelphia, Current Issues in Linguistic Theory, 260, 2004, 111-120.
- Thorsten Joachims, Text categorization with support vector machines: Learning with many relevant features. In *Proceedings of the European Conference on Machine Learning (ECML-98)*, pages 137–142.
- Jaap Kamps, Maarten Marx, Robert J. Mokken, and Marten de Rijke, Words with attitude, In Proceedings of the 1st International Conference on Global Word-Net, pages 332-341, Mysore, India, 2002.
- Alistair Kennedy and Diana Inkpen, Sentiment Classification of Movie Reviews Using Contextual Valence Shifters. *Computational Intelligence*, 2006, 22(2):110-125.
- Hugo Liu, Henry Lieberman, and Ted Selker, A model of textual affect sensing using real-world knowledge. In Proceedings of the ACM Conference on Intelligent User Interfaces, 2003, 125–132.
- Hugo Liu, and P. Maes, What Would They Think? A Computational Model of Attitudes. In *Proceedings of the ACM International Conference on Intelligent User Interfaces, IUI 2004*, 38-45, ACM Press.
- Rada Mihalcea and Hugo Liu, A corpus-based approach to finding happiness, In *Proceedings of the AAAI Spring Symposium on Computational Approaches for Analysis of Weblogs*, Stanford, CA, USA, March 2006.
- G. Miller and W. Charles. Contextual correlates of semantic similarity. *Language and Cognitive Processes*, 6(1):1-28, 1991.

- T Mullen and N Collier. Sentiment analysis using support vector machines with diverse information sources. In Dekang Lin and Dekai Wu, editors, *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing (EMNLP-2004)*, pages 412–418, Barcelona, Spain.
- Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. Analysis of affect expressed through the evolving language of online communication. In *Proceedings of the 12th International Conference on Intelligent User Interfaces (IUI-07)*, pages 278-281, Honolulu, Hawaii, USA, 2007a.
- Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka, Narrowing the Social Gap among People involved in Global Dialog: Automatic Emotion Detection in Blog Posts, In *Proceedings of the International Conference on Weblogs and Social Media (ICWSM 2007)*, pages 293-294, Boulder, CO, USA, March 2007b.
- A. Ortony, G.L. Clore, and A. Collins, The cognitive structure of emotions. New York: Cambridge University Press, 1988
- Bo Pang and Lillian Lee, 2004. A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL'04)*, Barcelona, Spain, pages 271-278.
- Bo Pang, Lillian Lee, and S. Vaithyanathan, Thumbs up? Sentiment classification using machine learning techniques, In *Proceedings of the 2002 Conference* on Empirical Methods in Natural Language Processing, Philadelphia, PA, 2002, 79–86.
- Peter Mark Roget, *Roget's Thesaurus of English Words and Phrases*. Harlow, Essex, England: Longman Group Limited, 1852.
- Carlo Strapparava and A Valitutti, WordNet-Affect: an affective extension of WordNet. In *Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC-2004)*, Lisbon, 2004, 1083-1086.
- Ian H. Witten and Eibe Frank. Data Mining: *Practical Machine Learning Tools and Techniques* (2nd ed.), Morgan Kaufmann, San Francisco, 2005. (www.cs.waikato.ac.nz/ml/weka/)
- L. Zhang, J. Barnden, R. Hendley, and A. Wallington, Exploitation in Affect Detection in Open-Ended Improvisational Text. In *Proceedings of the ACL Workshop on Sentiment and Subjectivity in Text*, 2006, pages 47-54, Sydney, Australia.