

Identifying Emotional Characteristics from Short Blog Texts

Alastair J. Gill (alastair@northwestern.edu)

Center for Technology and Social Behavior, Northwestern University
Evanston, IL 60208, USA

Robert M. French (robert.french@u-bourgogne.fr)

LEAD-CNRS UMR 5022, University of Burgundy
Dijon 21065, France

Darren Gergle (dgergle@northwestern.edu)

Center for Technology and Social Behavior, Northwestern University
Evanston, IL 60208, USA

Jon Oberlander (jon@inf.ed.ac.uk)

School of Informatics, University of Edinburgh
Edinburgh, EH8 9LW, UK

Abstract

Emotion is at the core of understanding ourselves and others, and the automatic expression and detection of emotion could enhance our experience with technologies. In this paper, we explore the use of computational linguistic tools to derive emotional features. Using 50 and 200 word samples of naturally-occurring blog texts, we find that some emotions are more discernible than others. In particular automated content analysis shows that authors expressing anger use the most affective language and also negative affect words; authors expressing joy use the most positive emotion words. In addition we explore the use of co-occurrence semantic space techniques to classify texts via their distance from emotional concept exemplar words: This demonstrated some success, particularly for identifying author expression of fear and joy emotions. This extends previous work by using finer-grained emotional categories and alternative linguistic analysis techniques. We relate our findings to human emotion perception and note potential applications.

Introduction

As humans, successful social engagement often centers on understanding what others are experiencing and then acting appropriately. One area where this is particularly salient is in the assessment of, and response to, another person's emotional state. This is such a basic underlying trait that even when interacting in technologically-mediated environments with few available cues, we are still able to make fairly accurate judgments of others' emotional states (to varying degrees of specificity) (Gill, et al, 2008; Hancock, et al., 2007; Cowie, et al. 2001). However, we are only recently beginning to develop an understanding of the degree to which emotional state can be detected in such environments, and developing accurate classification models for describing such cases. In this paper, we examine the use of computational methods to derive a richer set of emotional features that appear in naturally-occurring blog texts.

Recent investigations of affect in computer-mediated communication (CMC) environments such as text-chat, have used the LIWC text analysis tool (Pennebaker & Francis, 1999) to derive features characteristic of positive and negative emotion (Hancock, et al., 2007). While this work demonstrates the successful application of a generic dictionary-based text-analysis tool to the detection of positive and negative emotion in CMC, the ability to use such tools to detect differences in finer-grained emotion categories has yet to be demonstrated (cf. Liu et al. 2003 who categorize emotion according to 6 basic categories trained on human coded data). Another approach that has been successfully applied in other domains, such as detecting and classifying opinion and subjectivity, relies upon words with similar meanings co-occurring in similar contexts. This approach can give us access to higher-level semantic information that can be used to help classify a set of emotional concepts (Landauer & Dumais, 1997; Turney, 2002; Wiebe, et al, 2005; Pang, et al., 2002; Read, 2004). Further, such data-driven techniques are more likely to be generalisable across different areas and inform applications. In this paper we examine the applicability of this approach to detecting emotion in text.

We examine the language of emotion for two reasons: Firstly, we are interested in what emotional cues are available in the relatively impoverished CMC environment, and whether the previous findings of Hancock, et al. (2007) using the LIWC (Linguistic Inquiry and Word Count) text analysis tool can be replicated across more specific emotion categories; Second, we also examine the application of semantic space techniques to this area. In particular, we aim to identify emotion cues and develop computational descriptions, for example to bestow emotional abilities in a variety of applications such as embodied conversational agents (Ortony, 2002), or to capture fine-grained emotions of individuals and groups online (cf. Balog, et al. 2006; <http://ilps.science.uva.nl/MoodViews/>).

Emotion

Emotion is an individual response to stimuli. A number of approaches attempt to describe emotion, ranging from positive and negative affect and basic categories of emotion, to more detailed descriptions (cf. Ekman, 1982; Scherer, 2005). In this paper, we adopt a model of emotion derived from lists of words representative of emotional states, and which were then statistically grouped into eight primary emotions (Plutchik, 1994). These can be described as representing the extreme ends of four emotional continua: Joy-Sadness; Acceptance-Disgust; Fear-Anger; Surprise-Anticipation. These primary emotions can then be used to describe finer-grained secondary, and tertiary emotions, so that for example, a primary emotion like Acceptance is composed of secondary emotions Curiosity and Love. In the current work, we focus only on the eight primary emotions, represented as an activation-evaluation wheel shown in Figure 1 (Feldman Barrett, & Russell, 1998, derived from Plutchik, 1994). Increasing distance from the centre of the wheel indicates greater strength of the emotion, with evaluation (valence) increasing positive in emotions towards the right, and activity increasing in emotions towards the top of the wheel. This model is considered well-suited to computational work (Cowie, et al., 2001), has previously been used for rating emotion in speech (Makarova & Petrushin, 1999), and allows comparison with findings for valence (Hancock, et al, 2007). Note that in the activation-evaluation wheel the emotions have been aligned so that they are viewed as individual emotions, rather than belonging to a continuum, e.g., Joy-Sadness, (cf. Feldman Barrett, & Russell, 1998; Makarova & Petrushin, 1999).

In spite of – or perhaps because of – the complex individual nature of emotion, there has been a recent increase of interest in emotion in communication and language (Fussell, 2002; Weibe, et al. 2005; Makarova & Petrushin, 1999). In particular, we note the work by Hancock, et al. (2007), who asked participants in a text chat environment to express either positive (happy) or negative (unhappy) emotions to their naive conversational partner without explicitly describing their (projected) emotional state. They found that Naive judges (the text-chat partners) could accurately perceive their interlocutor's emotion, and were less likely to enjoy or want to meet the authors of negative messages relative to positive ones. Further, linguistic analysis (LIWC; Pennebaker & Francis, 1999) of the transcripts found that authors portraying positive emotion used a greater number of exclamation marks, and used more words overall, whereas authors' texts portraying negative emotion used an increased number of affective words, words expressing negative feeling, and negations. Punctuation features matched the self-reported strategies used by the portrayers of emotion to express emotion.

However, the study by Hancock et al. (2007), was limited to positive and negative emotions (happy vs. sad), the naive judges' ratings of emotion were based on a 30 minute interaction, and the emotions were acted out through a confederate. Gill et al. (2008) use a corpus of personal blog

texts, written by real authors expressing genuine emotions, to extend this previous work. They found that naive raters with little experience of using blogs are able to identify four of Plutchik's basic emotions (joy, disgust, anger and anticipation), showing relatively high agreement with expert judges. Additionally, rather than interacting for 30 minutes, the naive raters were able to achieve these judgments after reading 50 or 200 words of asynchronous blog text.




Figure 1: Activation-evaluation wheel.

Text analysis and word co-occurrence

The LIWC tool (Pennebaker & Francis, 1999) is a popular content analysis technique which counts occurrences of words according to pre-defined psychological and linguistic categories. The LIWC categories are grouped under four main dimensions: Linguistic Dimensions (e.g., word count, pronouns, negations, numbers) are values calculated directly from the text; Psychological Processes (e.g., positive or negative emotions) capture basic psychological processes; the Relativity dimension describes physical or temporal information (e.g., time and space); and Personal Concerns (e.g., occupation, leisure activities) address content topics of conversation. LIWC analysis has been successfully applied to a wide range of data, including determining the linguistic characteristics of emotion, personality, gender and genre (Hancock, et al. 2007; Nowson, et al. 2005). Given that the dictionaries which power the LIWC word count analysis have been derived by from human classifications, this can be considered a top-down approach (cf. Liu, et al. 2003).

In order to examine further whether a data-driven technique can be applied to the linguistic analysis and classification of emotion, we also adopt co-occurrence techniques previously applied to classifying opinion and subjectivity (Landauer & Dumais, 1997; Turney, 2002; Wiebe, et al. 2005; Pang, et al., 2002). This research originated as a way of improving the performance of document retrieval in an electronic database, by enabling the search to be performed on the basis of meaning or semantic-similarity rather than just by keywords. These techniques have been adopted to explore psychological phenomena, such as child language acquisition and reading difficulty and text cohesion (Landauer & Dumais, 1997).

In this paper, we examine two such techniques, Latent Semantic Analysis (LSA; Landauer & Dumais, 1997), and Hyperspace Analogue to Language (HAL; Lund, et al. 1995). The approach is based on the theory that they can determine the semantics (or, at least, some of the semantics) of a word by analyzing how closely other words physically co-occur with it over a large number of texts (e.g., the word “mother” will tend to cluster more closely with “child,” “father,” “birth,” and “baby” than say, with “clutch,” “carburetor”, and “gasoline”). Importantly, these programs have demonstrated human-like levels of performance in tasks such as English language learner synonym tasks (e.g., Landauer & Dumais, 1997), classifying the semantic orientation (good vs bad, etc.) of individual words and movie reviews (Turney, 2002). The limitations of these programs have also been discussed in the literature (e.g., Glenberg & Robertson, 2000; French & Labiouse, 2002). In particular, Bullinaria & Levy (2007) observe that “obviously, co-occurrence statistics *on their own* [original emphasis] will not be sufficient to build complete and reliable lexical representations”. In particular, we adopt the technique of Turney (2002) for evaluating affective orientation to emotions.

Method

Data collection

The blog texts were taken from a previously collected corpus (Nowson, et al. 2005). The texts were collected from real blogs extracted for a specified month. Permission for further use of each blog was granted by the authors before collection (Nowson, et al. 2005). The first 200 words of each post were classified as one of eight emotions (surprise, joy anticipation, acceptance, sadness, disgust, anger, fear) or neutral by six expert raters who had had extensive exposure to the texts. From 135 texts, 20 were selected as expressing strong and clear emotional content. This was based on all expert raters agreeing on the emotion assigned, and having the strongest emotion rating (2 for each emotion; and 4 for ‘neutral’, which we disregard in the current analysis). Figure 1 shows the emotion wheel used for text rating.

Text preparation

For each of these 20 texts we use two versions in the subsequent analysis: For the long version, we retain all 200 words; for the short version we extract the middle 50 words of the 200 word text, ignoring sentence boundaries. These are the same texts previously used for the naive rating of emotion (Gill et al. 2008). Analysis of the texts was performed by submitting them to the LIWC text analysis program (Pennebaker & King, 1999). To explore the location of these texts within semantic space, we use Latent-Semantic Analysis (LSA, Landauer & Dumais, 1997; <http://lsa.colorado.edu>) and an implementation of

Hyperspace Analogue to Language (HAL; Lund, et al. 1995; Huettig, et al. 2006; online version available at: http://www.cogsci.ed.ac.uk/~scottm/semantic_space_model.html). For this co-occurrence analysis, we do not use the whole 50 or 200 word sections from the blog texts, rather we extract 10 key words from each text using term frequency-inverse document frequency (TF-IDF; Belew, 2000; note that this is different to the approach of Turney, 2002, who extracted adjective-adverb phrases).

Table 1: Emotion exemplar words.

acceptance	fear	anger	joy
acceptance	fear	anger	joy
agreement	phobia	rage	delight
affirmation	terror	fury	bliss
admission	fright	outrage	rejoicing
adoption	scare	hatred	elation
approval	dread	tantrum	gaiety
assent	nightmare	animosity	glee
anticipation	sadness	disgust	surprise
anticipation	sadness	disgust	surprise
awaiting	depression	revulsion	unexpected
expectancy	sorrow	distaste	unforeseen
prospect	melancholy	aversion	astonishment
hope	woe	loathing	shock
promise	grief	dislike	amazement
apprehension	mourning	nausea	incredulity

Calculation of Semantic Space

Following Turney’s classification of sentiment, we use 7 exemplar words to represent each of Plutchik’s eight basic emotions (these can be found in Table 1, with the emotional concept word emboldened). The exemplar words for each emotion were derived from synonyms taken from *Roget’s II: The New Thesaurus* (1995), with ratings by 6 research assistants used to select the most similar words to the emotion concept. For each of the 10 key terms extracted from the blog texts, we calculate a semantic distance to the exemplar words for each emotion, using both LSA and HAL. Here we treat each of the eight emotion concepts individually (cf. Turney, 2002 who located sentiment of reviews between exemplars representing “good” and “bad” concepts). The following parameters were used for the calculation of semantic association:

- **HAL** was implemented using the British National Corpus (BNC), using a rectangular window of 7 words and distance between vectors calculated using cosine, as reported in Huettig et al. (2006).
- **LSA** (Landauer, & Dumais, 1997) uses the University of Colorado at Boulder website using the default semantic space derived from the ‘General Reading up to 1st year of college’ TASA corpus, and the maximum number of factors available (300). The comparison type used was ‘term to term’.

Table 2: LIWC results by text emotion.

	ANOVA Fit Model			Mean scores for levels of categorical independent variable*								
	<i>F</i>	<i>DF</i>	<i>p</i>	Acceptance	Anger	Anticipation	Disgust	Fear	Joy	Sadness	Surprise	
Affect	6.87	7	0.0002	3.38 ^{b,c}	7.25 ^a	2.73 ^{b,c}	3.73 ^{b,c}	3.13 ^{b,c}	5.65 ^{a,b}	4.78 ^{a,b}	1.50 ^c	
Pos. Emotion	2.55	7	0.041	1.50 ^{a,b}	1.63 ^{a,b}	1.75 ^{a,b}	2.23 ^{a,b}	1.25 ^b	4.28 ^a	1.90 ^{a,b}	1.38 ^b	
Neg. Emotion	6.36	7	0.0003	1.88 ^{b,c}	5.63 ^a	1.00 ^b	1.50 ^b	1.88 ^b	1.38 ^b	2.88 ^{a,b}	0.13 ^b	
Pronouns												
First Person	2.30	7	0.0606	6.65 ^{a,b}	10.75 ^{a,b}	12.23 ^a	3.38 ^b	7.13 ^{a,b}	6.15 ^{a,b}	9.40 ^{a,b}	7.43 ^{a,b}	
Third Person	1.40	7	0.2502	3.13 ^a	0.38 ^a	2.00 ^a	4.63 ^a	3.88 ^a	0.63 ^a	4.28 ^a	2.13 ^a	
Agreement												
Negation	1.88	7	0.1184	3.13 ^a	4.13 ^a	1.63 ^a	2.85 ^a	2.63 ^a	2.38 ^a	1.63 ^a	1.88 ^a	
Assent	1.51	7	0.2118	0.00	0.63 ^a	0.25 ^a	0.00 ^a	0.00 ^a	0.13 ^a	0.00 ^a	0.00 ^a	

*Tukey HSD comparison across all levels (differences between levels indicated by different superscript characters); Levels are emotions assigned by expert judges

Statistical Analysis

Linguistic variables (derived from LIWC, HAL or LSA) were entered into a regression model as the dependent variables. Expert emotion ratings for each of the 16 texts were entered as the independent, categorical variable (cf. Hancock et al. 2007; N.B. the 4 Neutral texts are omitted). In these analyses, we treat each text as independent, however we note that the short texts are in fact excerpts of the larger texts. Significant relationships within these statistical models are reported as ANOVAs, with Tukey HSD post-hoc tests used to identify significant differences between means (indicated by different superscript letters in the following tables).

Results and Discussion

Table 2 shows the result of analysis using LIWC: we include the same variables as Hancock et al. (except word count), and also include finer-grained categories for positive emotion (positive feeling, optimism) and negative emotion (anxiety, anger, sadness).

Unlike Hancock et al., we do not see a significant difference in use of negations according to emotion, however we do note a greater use of affective language overall, including an increased use of positive emotion words. From our more nuanced emotion categories, we note from the Tukey HSD post-hoc tests that texts demonstrating anger use the most affective terms (and more than surprise texts). In the case of positive emotion words, these are used most by joyful authors (more than those expressing fear or surprise). Negative emotion words are used more by angry authors than those expressing any other emotion, with the exception of sad authors. Examination of the more detailed LIWC emotion categories reveals that anger words are used more by angry authors than the authors of any other texts (except fearful authors), and that authors expressing sadness using more sadness words than other authors (except those expressing anger). Neither positive feeling nor optimism words showed significant difference across emotion texts.

Turning now to the co-occurrence analysis: Table 3 presents the results for each of the exemplar emotion categories for HAL and LSA (represented as HAL-Anger, LSA-Anticipation, etc.; in each case, the greater the mean score, the greater the semantic similarity). We note that for

Fear, both HAL and LSA analyses reveal that texts expressing this emotion have a greater semantic similarity to the Fear exemplar words. In the case of LSA, we also note that Joy texts are also rated as being most similar to the Joy concept exemplar words. In all three of these cases, although there are significant correlations present, Tukey HSD post-hoc tests reveal that Joy or Fear texts are not significantly more similar to the respective emotion exemplar words than a number of texts expressing other emotions (for HAL, texts expressing Fear were significantly more similar to the Fear exemplar words than Disgust texts, Acceptance and Surprise; for LSA, both texts expressing Fear and Joy were significantly more similar to Fear and Joy exemplars respectively than texts expressing Disgust). In other cases, LSA especially, appears to correctly identify the text emotion as being similar to the relevant emotion exemplar, however this is often alongside other texts expressing emotions (e.g., LSA-Anger, relates to Fear, Anger, Joy, and Surprise texts).

As can be seen from this example, Fear, Joy and Surprise texts – along with Anger texts – are regarded by LSA as semantically similar to the Anger exemplar concept. These results are counterintuitive, and somewhat surprising given the previous success of co-occurrence programs, like HAL and LSA, in areas such as synonym matching and assessing opinions from text (Landauer & Dumais, 1997; Lund et al. 1995; Turney, 2002; Bullinaria & Levy, 2007). However, it may be that many of the emotion terms which we have examined in this paper occur in similar contexts. Therefore, such terms may not be very well differentiated in co-occurrence semantic space, with such co-occurrence programs unable to identify meaning from experience, unlike humans (cf. Friedrich, 1993; Glenberg & Robertson, 2000; French & Labiouse, 2002; Bullinaria & Levy (2007). In relating our linguistic findings to the human raters of emotion from these short blog texts (Gill, et al. 2008), it is perhaps unsurprising that LIWC analysis found linguistic features relating to Joy and Anger, given that these were both relatively easily perceived by naive human judges. In the case of the HAL and LSA semantic space analysis, again it is unsurprising that Joy has linguistic correlates, however that both HAL and LSA both related the texts expressing Fear to the Fear exemplars is interesting. Given this result, we expect that texts expressing Fear contain words similar in meaning to “phobia”, “terror”, or “fright”. It may be that

Table 3: HAL and LSA semantic distances for each emotion concept by text emotion.

	ANOVA Fit Model			Mean scores for levels of categorical independent variable*							
	F	DF	p	Acceptance	Anger	Anticipation	Disgust	Fear	Joy	Sadness	Surprise
Sem. Similarity: HAL											
HAL-Acceptance	1.54	7	0.2001	0.21 ^a	0.24 ^a	0.23 ^{a,b}	0.22 ^a	0.18 ^a	0.19 ^a	0.24 ^a	0.21 ^a
HAL-Anger	3.43	7	0.0109	0.26^{a,b}	0.29^{a,b}	0.26^{a,b}	0.24^b	0.31^a	0.28^{a,b}	0.28^{a,b}	0.26^{a,b}
HAL-Anticipation	2.21	7	0.0696	0.28 ^a	0.30 ^a	0.31 ^a	0.27 ^a	0.30 ^a	0.29 ^a	0.29 ^a	0.26 ^a
HAL-Disgust	1.69	7	0.16	0.24 ^a	0.27 ^a	0.26 ^a	0.24 ^a	0.29 ^a	0.28 ^a	0.28 ^a	0.25 ^a
HAL-Fear	4.88	7	0.0015	0.29^b	0.31^{a,b}	0.30^{a,b}	0.26^b	0.36^a	0.31^{a,b}	0.32^{a,b}	0.28^b
HAL-Joy	2.33	7	0.0574	0.23 ^a	0.26 ^a	0.27 ^a	0.24 ^a	0.27 ^a	0.29 ^a	0.24 ^a	0.24 ^a
HAL-Sadness	2.88	7	0.0247	0.27^{a,b}	0.30^{a,b}	0.30^{a,b}	0.26^b	0.32^a	0.30^{a,b}	0.29^{a,b}	0.27^{a,b}
HAL-Surprise	2.06	7	0.089	0.27 ^a	0.29 ^a	0.28 ^a	0.26 ^a	0.31 ^a	0.30 ^a	0.29 ^a	0.26 ^a

*Tukey HSD comparison across all levels (differences between levels indicated by different superscript characters); Levels are emotions assigned by expert judges

	ANOVA Fit Model			Mean scores for levels of categorical independent variable*							
	F	DF	p	Acceptance	Anger	Anticipation	Disgust	Fear	Joy	Sadness	Surprise
Sem. Similarity: LSA											
LSA-Acceptance	2.74	7	0.0308	0.19^{a,b}	0.20^a	0.15^{a,b}	0.10^b	0.12^{a,b}	0.16^{a,b}	0.16^{a,b}	0.18^{a,b}
LSA-Anger	5.36	7	0.0009	0.21^a	0.24^a	0.21^a	0.13^b	0.26^a	0.24^a	0.20^{a,b}	0.22^a
LSA-Anticipation	5.34	7	0.0009	0.27^a	0.28^a	0.26^a	0.17^b	0.28^a	0.29^a	0.26^a	0.29^a
LSA-Disgust	1.85	7	0.1225	0.21 ^a	0.21 ^a	0.18 ^a	0.15 ^a	0.21 ^a	0.20 ^a	0.19 ^a	0.20 ^a
LSA-Fear	6.87	7	0.0002	0.24^a	0.27^a	0.24^{a,b}	0.16^b	0.32^a	0.28^a	0.25^a	0.27^a
LSA-Joy	4.85	7	0.0016	0.22^{a,b}	0.21^{a,b}	0.25^a	0.16^b	0.25^a	0.29^a	0.21^{a,b}	0.25^a
LSA-Sadness	6.30	7	0.0003	0.21^{a,b}	0.22^a	0.27^a	0.12^b	0.26^a	0.28^a	0.21^{a,b}	0.24^a
LSA-Surprise	6.02	7	0.0004	0.25^{a,b}	0.26^a	0.24^{a,b}	0.18^b	0.31^a	0.30^a	0.25^{a,b}	0.29^a

*Tukey HSD comparison across all levels (differences between levels indicated by different superscript characters); Levels are emotions assigned by expert judges

human judges for some reason do not expect such explicit references to fear, and therefore do not look for them. We leave the exploration of this to future work.

To summarize, first, by extending previous work examining the expression of emotion in CMC, we have shown emotion can be communicated linguistically in relatively short blog texts of 50 or 200 words, and that these emotions have replicated previous findings using automated content analysis. Secondly, we extend previous work to include data-driven co-occurrence techniques, and hope this can begin to inform computational approaches to emotion, for applications such as embodied conversational agents (e.g., Ortony, 2002). In particular we expect that future work combining machine learning approaches with the top-down content analysis and data-driven semantic space analysis will be particularly fruitful.

Together these results provide both theoretical and applied advances. At a theoretical level, this work further develops our understanding of the ways in which emotional characteristics can be articulated and comprehended in less rich environments such as blog texts. At an applied level the computational approaches examined in this work can be used in technologies to develop a richer understanding of the emotional content of existing written excerpts, and they may also be used to help imbue our technologies with a richer repertoire of techniques for inserting emotional content into their expressions. We also note potential future work and applications.

Conclusion

Emotion is at the core of understanding ourselves and those around us. In order to develop technologies that are capable of understanding or expressing emotion we need to further develop techniques and computational models that can automate the detection and expression of such emotions. In this paper, we have explored the use of computational linguistics techniques to derive and detect linguistic components that are correlated with human ratings of various emotional expressions. We used 50 and 200 word samples of naturally-occurring blog texts and found that some emotions are much more discernible than others. By using automated content analysis techniques we found that authors expressed anger using a larger portion of affective language and negative affect words. In addition to the content analysis approaches, we have demonstrated the use of co-occurrence semantic space techniques to classify texts via their distance from emotional concepts captured in exemplar words. This approach demonstrated some success, particularly for identifying author expression of fear and joy. In comparing these linguistic analyses to human emotion raters, we find that Anger and Joy from the LIWC analyses and Joy from the semantic space analyses are readily perceived by human judges. However, interestingly both HAL and LSA detect Fear, but human judges do not. Together this work extends prior studies by applying alternative linguistic analysis techniques to a finer-grained representation of emotion.

Acknowledgments

We thank Jonathan Ellis for preparing the rating materials, Scott Nowson for making his weblog corpus available, Jonathan Read for early comments about this work, Francisco Iacobelli for use of TF-IDF software and Scott McDonald for access to his co-occurrence software. We acknowledge support in part from Région de Bourgogne FABER Post-Doctoral Fellowship (05512AA06S2469) (first author), European Commission Sixth Framework grants NEST 516542 and NEST 029088 (second author), NSF IIS-0705901 (third author).

References

- Balog, K., Mishne, G., & Rijke, M., (2006) Why Are They Excited? Identifying and Explaining Spikes in Blog Mood Levels. In *Proceedings of the European Chapter of the Association of Computational Linguistic*.
- Belew, R. K. (2000). *Finding out about: A cognitive perspective on search engine technology and the WWW*. Cambridge, U.K.: Cambridge University Press.
- Bullinaria, J.A. & Levy, J.P. (2007). Extracting Semantic Representations from Word Co-occurrence Statistics: A Computational Study. *Behavior Research Methods*, 39, 510-526.
- Cowie, R., Douglas-Cowie, E., Tsapatsoulis, N., Votsis, G., Kollias, S., Fellenz, W., and Taylor, J.G. (2001) Emotion recognition in Human-Computer interaction. In *IEEE Sig. Proc. Mag. Vol. 18(1)*, 32-80.
- Ekman, P. (1982). *Emotion in the human face* (2nd ed.). New York: Cambridge University Press.
- Feldman Barrett, L., & Russell, J.A. (1998). Independence and bipolarity in the structure of affect. *J. Personality and Social Psychology*, 74, 967-984.
- French, R. M. & Labiouse, C. (2002). Four Problems with Extracting Human Semantics from Large Text Corpora. *Proceedings of the 24th Annual Conference of the Cognitive Science Society*.
- Friedrich, J. (1993), Primary Error Detection and Minimization (PEDMIN) Strategies in Social Cognition: A Reinterpretation of Confirmation Bias Phenomena, *Psychological Review*, 100 (2), 298-319.
- Fussell, S. R. (2002). The verbal communication of emotion: Interdisciplinary perspectives: Introduction and overview. In S. R. Fussell, (Ed.) *The verbal communication of emotion: Interdisciplinary perspectives*. Mahwah, NJ: Erlbaum .
- Gigerenzer, G. & Todd, P. M. (1999). Fast and frugal heuristics: the adaptive toolbox. In G. Gigerenzer, P. M. Todd, & the ABC Group (eds.) *Simple Heuristics That Make Us Smart*. NY: Oxford University Press, pp. 3-34.
- Gill, A.J., Gergle, D., French, R.M., and Oberlander, J. (2008). Emotion Rating from Short Blog Texts. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2008)*, pp 1121-1124.
- Hancock, J.T., Landrigan, C., & Silver, C. (2007). Expressing emotion in text-based communication. *Proc ACM Conference on Human Factors in Computing Systems (CHI 2007)*, 929-932.
- Huetting, F., Quinlan, P. T., McDonald, S. A. & Altmann, G. T. M. (2006). Models of high-dimensional semantic space predict language-mediated eye movements in the visual world. *Acta Psychologica*, 121, 65-80.
- Landauer, T.K., & Dumais, S.T. (1997). A solution to Plato's problem: The latent semantic analysis theory of the acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2), 211-240.
- Liu, H., Lieberman, H., and Selker, T. (2003). A model of textual affect sensing using real-world knowledge. In *Proceedings of the 8th international conference on Intelligent user interfaces*, pp. 125-132.
- Lund, K., Burgess, C., & Atchley, R.A. (1995). Semantic and associative priming in high-dimensional semantic space. In *Proceedings of the 17th Annual Conference of the Cognitive Science Society*, pages 660-665.
- Makarova, V., Petrushin V. A. (2002). RUSLANA: A Database of Russian Emotional Utterances. *Proc. Int. Conf. Spoken Language Processing (ICSLP 2002)*, pp. 2041-2044.
- Nowson, S., Oberlander, J., & Gill, A.J. (2005). Weblogs, genres and individual differences. In *Proceedings of the 27th Annual Conference of the Cognitive Science Society*, pp. 1666-1671.
- Ortony, A. (2002) On making believable emotional agents believable. In Trapp, R., Petta, P., Payr, S., (eds) *Emotions in humans and artefacts*, Cambridge, MA: MIT Press.
- Pang, B., Lee, L., & Vaithyanathan, S. (2002). 'Thumbs up? Sentiment Classification Using Machine Learning Techniques'. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2002)*, pp. 79-86.
- Pennebaker, J. and Francis, M. (1999). *Linguistic Inquiry and Word Count*. Mahwah, NJ: LEA.
- Plutchik, R. (1994). *The psychology and biology of emotion*. New York, NY: HarperCollins.
- Read, J. (2004). Recognising Affect in Text using Pointwise-Mutual Information. Unpublished M.Sc. Dissertation, University of Sussex, UK.
- (1995). *Roget's II: The New Thesaurus*, 3rd ed. Boston, MA: Houghton Mifflin.
- Scherer, K. R. (2005). What are emotions? And how can they be measured? *Social Science Information*, 44(4), 693-727.
- Turney, P.D. (2002), Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews, *Proc. 40th Annual Meeting of the Association for Computational Linguistics (ACL'02)*, pp. 417-424.
- Wiebe, J., Wilson, T., & Cardie, C. (2005). Annotating expressions of opinions and emotions in language. *Lang Resources and Evaluation*, 39, 65-210.