The Language of Emotion in Short Blog Texts

Alastair J. Gill*, Robert M. French†, Darren Gergle*, Jon Oberlander‡

*Center for Technology and Social Behavior
Northwestern University
2240 Campus Drive,
Evanston, IL 60208, USA
{alastair|dgergle}@northwestern.edu

†LEAD-CNRS UMR 5022,
University of Burgundy
Pôle AAFE, BP 26513,
Dijon 21065, France
robert.french@u-bourgogne.fr

‡School of Informatics,
University of Edinburgh
2 Buccleuch Place,
Edinburgh, EH8 9LW, UK
J.Oberlander@ed.ac.uk

ABSTRACT
Emotion is central to human interactions, and automatic detection could enhance our experience with technologies. We investigate the linguistic expression of fine-grained emotion in 50 and 200 word samples of real blog texts previously coded by expert and naive raters. Content analysis (LIWC) reveals angry authors use more affective language and negative affect words, and that joyful authors use more positive affect words. Additionally, a co-occurrence semantic space approach (LSA) was able to identify fear (which naive human emotion raters could not do). We relate our findings to human emotion perception and note potential computational applications.

Author Keywords
Blogs, language, text analysis, emotion, affect.

ACM Classification Keywords
H.5.m [**Information Interfaces and Presentation (HCI)**]: Miscellaneous.

J4 [**Social and behavioral systems**]: Psychology.

INTRODUCTION
Successful social engagement often centers on understanding what others are experiencing and then acting appropriately. In technologically-mediated environments with few available cues, humans can still make fairly accurate judgments of others' emotional states [1,5,6]. In this paper, we focus on automatically detecting emotional information from the language of short blog texts, where explicit author mood information is not available (cf. Livejournal; [8,10]). Automatically extracted emotional information can be used in a variety of computational applications such as interfaces that can automatically detect and adapt to our emotional state, collaborative software that can identify interlocutor emotions, or applications that allow us to follow a friend's emotional state by processing their recent blog posts [2,11,12].

In this paper, we examine what emotional cues are available in a relatively impoverished computer-mediated environment. Theoretically, we test whether previous top-down content analysis findings are replicated across more specific emotion and linguistic categories [6 cf. 1]. In addition, we apply data-driven techniques previously used to classify opinion and mood [7,10,15,16] since these may be better generalized across different genres and applications.

Emotion
Following increasing interest relating to emotion in communication [e.g., 4,16], we note that Hancock et al. found that positive and negative emotion could accurately be perceived in a text-chat environment [6]. Using content analysis (Linguistic Inquiry and Word Count; LIWC), they found that authors portraying positive emotion used more exclamation marks and more words overall, whereas authors' portraying negative emotion used an increased number of affective words, words expressing negative emotion, and negations.

However, the study by Hancock et al. was limited to positive and negative emotions (happy vs. sad), the judges' ratings of emotion were based on a 30 minute interaction, and the emotions were acted out through a confederate. Extending this work, Gill et al. [5] used a corpus of personal blog texts, written by authors expressing genuine emotions: Naive raters could accurately identify four primary emotions (joy, disgust, anger and anticipation) across 50 and 200 word texts. Other work exploring automatic classification of mood from Livejournal blogs reached 66% accuracy, but was impeded by subjective author self mood assignment and short texts [10 cf. 8]. Additionally, affective states such as boredom and frustration have been classified from dialogue features in response to an automated tutoring system [2].

In this paper, we extend this work by examining automatic techniques for classifying emotion on a richer model of emotion. We adopt Plutchik's model of emotion [14; cf. 2] which describes emotion along activation (activity) and

CSCW'08, November 8–12, 2008, San Diego, California, USA.

evaluation (valence) dimensions and is considered well-suited to computational work [1]. It consists of eight primary emotions: Joy, Sadness, Acceptance, Disgust, Fear, Anger, Surprise, and Anticipation.

Text Analysis Techniques

The LIWC [13] content analysis tool was previously used to identify linguistic features of positive and negative emotion [12]. Since it counts occurrences of words according to pre-defined psychological and linguistic categories, it can be considered a top-down approach [cf. 9]. Using LIWC, we hypothesize that we will replicate previous findings for negative emotion (sadness), specifically more affective words, negative emotion words and negations [6]. We also expect joyful authors to use more positive emotion and positive feeling words, anticipating authors more optimism words, and an increase in fear, anger and sadness words with the corresponding emotional categories described in Plutchik's model.

We also explore data-driven techniques previously applied to mood [10], which assume that word semantics can be determined by their co-occurrence context over large corpora (e.g., Latent Semantic Analysis; LSA, [7]). Despite demonstrating human-like performance in synonym tasks, and classifying valence of movie reviews [7,10,15,16], critics have noted limitations in embodied or subcognitive ability [e.g., 3]. We apply Turney's technique for evaluating opinion to classifying emotions [15 cf. 10]. For LSA, we expect each emotion text to show greatest semantic similarity to the appropriate emotion semantic space, e.g., joyful texts to joyful emotion concept words, etc.

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*

* indicates concept words most similar to appropriate texts

Table 1: Emotion exemplar words

METHOD

Data Collection

We use texts extracted from real blogs (e.g., Livejournal, Xanga), previously used to study emotion perception by naive judges [5] (thus allowing direct comparison with human judge ability). 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 trained in personality psychology, with extensive exposure to the blogs. Author emotion information was not collected to avoid self-presentation biases, with the widely available author ratings of mood (e.g., Livejournal) often being subjective and not easily mapped to theoretical models of emotion [10]. From 135 texts, 20 were selected as expressing strong and clear emotional content (with the 4 'neutral' texts excluded from further evaluation).

Text Preparation

The 16 texts were analyzed in two versions: long (200 words) and short (the middle 50 words extracted from the long texts). Such short texts are typical of blog posts, but can limit the effectiveness of text classification techniques [10]. Additionally, emotion is unstable, and may not be consistently communicated in longer blog posts. Texts were submitted to LIWC [12], and their location in semantic space explored using Latent Semantic Analysis (LSA [7]). For co-occurrence analysis, 10 key words were extracted from each 50 or 200 word text using term frequency-inverse document frequency (TF-IDF; cf. extracting adjective-adverb phrases [15]).

Calculation of Semantic Space

Seven exemplar words represent each of Plutchik's eight basic emotions (Table 1; emotional concept words are in bold; [cf. 15]). Exemplar words for each emotion were derived from *Roget's II: The New Thesaurus* (3rd ed. Boston, MA: Houghton Mifflin), with ratings by 6 research assistants used to select the most similar synonyms to the emotion concept (ratings were summed; highest-scoring items used). For each of the 10 key terms extracted from the blog texts, we calculate a semantic distance to the exemplar words of each emotion. Here we treat each of the eight emotion concepts as independent dimensions [cf. 15]. LSA [7] was calculated via the University of Colorado at Boulder website (http://lsa.colorado.edu) using the default semantic space ('General Reading up to 1st year of college' TASA corpus; maximum number of factors, 300; comparison type 'term to term', i.e. word level).

Statistical Analysis

Linguistic variables (derived from LIWC or LSA) were entered into a regression model as dependent variables, with the expert emotion ratings for each of the 16 texts as the independent, categorical variable [cf. 6]. We treat each text as independent in these analyses, but 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). In addition we report the recall and precision for the LSA emotion classifications.

	ANOVA Fit Model			Mean scores for levels of categorical independent variable*							
	F	*DF*	*p*	Acceptance	Anger	Anticipation	Disgust	Fear	Joy	Sadness	Surprise
LIWC											
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}
Pos. Feel.	0.41	7	0.888	0.38 ^{a}	0.75 ^{a}	0.75 ^{a}	0.50 ^{a}	0.25 ^{a}	0.88 ^{a}	0.50 ^{a}	0.25 ^{a}
Optimism	1.09	7	0.401	0.00 ^{a}	0.25 ^{a}	0.25 ^{a}	0.88 ^{a}	0.63 ^{a}	0.25 ^{a}	0.00 ^{a}	0.25 ^{a}
Neg. Emotion	**6.36**	**7**	**0.000**	**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}
Anxiety	1.61	7	0.182	0.00 ^{a}	0.13 ^{a}	0.63 ^{a}	0.00 ^{a}	1.50 ^{a}	0.25 ^{a}	0.25 ^{a}	0.13 ^{a}
Anger	**5.82**	**7**	**0.001**	**0.63** ^{b}	**3.50** ^{a}	**0.13** ^{b}	**0.63** ^{b}	**1.75** ^{a,b}	**0.13** ^{b}	**0.38** ^{b}	**0.13** ^{b}
Sadness	**4.65**	**7**	**0.002**	**0.00** ^{b}	**0.88** ^{a,b}	**0.13** ^{b}	**0.13** ^{b}	**0.00** ^{b}	**0.13** ^{b}	**1.38** ^{a}	**0.00** ^{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}
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

Table 2: LIWC and LSA results by text emotion

Results and Discussion

The LIWC analysis (Table 2, top) includes the same variables as Hancock et al. (except word count), and also finer-grained categories for positive emotion (positive feeling, optimism) and negative emotion (anxiety, anger, sadness). However, we do not see a significant difference in use of negations according to emotion, but note a greater rate of affective language, especially positive emotion words. Counter to our hypotheses, Tukey HSD post-hoc tests across our more nuanced emotion categories show that anger – rather than sadness – texts use the highest proportion of affective terms (not significantly different to joy or sad texts) and more negative emotion words (not different to sad texts).

As predicted, more detailed LIWC word categories reveal that anger words are used more frequently by angry authors than in any other texts (except fear), with authors expressing sadness using a higher rate of sadness words (except authors expressing anger). However the use of anxiety words did not significantly vary. These results build upon and provide a more detailed linguistic and emotional view of previous findings for authors expressing negative emotion [6]. As predicted, positive emotion word rates were highest by joyful authors (significantly more than those expressing fear or surprise). Neither positive feeling nor optimism rates showed significant differences.

Turning now to the data-driven co-occurrence analysis (Table 2, bottom): These data show mean similarity distances of texts to exemplar emotion categories (LSA-Anger, etc; a higher mean indicates greater similarity). In terms of our hypotheses, here we note that only fear texts and joy texts are most semantically similar to their respective emotion exemplars, but in both cases Tukey HSD tests show that they are not distinct from other emotion texts. For example, LSA-Joy relates to acceptance, anger, anticipation, fear, sadness, surprise – as well as – joy texts. (In Table 1, exemplar words most similar to their respective texts are noted with '*' by emotion.)

Further, we examine LSA's performance in classifying emotions (where baseline performance is 12.5%): Across long and short texts, fear is the most accurately identified (Recall = .75), and shows fewest misclassifications (Precision = .38), with further exploration showing the majority of misclassifications as sadness. Surprise texts were classified correctly a quarter of the time (Recall = .25) with misclassifications (Precision = .17) mainly consisting of short disgust texts as surprise. Anticipation was also classified correctly a quarter of the time (Recall = .25), but with low precision (.08) where anger, surprise and sadness texts were misclassified as anticipation. It is relatively unsurprising that two negative emotions, fear and sadness, are misclassified by LSA, since presumably these concepts are located relatively closely in semantic space [cf. 3]. However, why might anticipation be confused with anger, surprise and sadness, and surprise with disgust? In both cases, these are positive emotions, so it may be that positive emotions are less explicitly vocalized (LIWC analysis showed that emotion language generally relates more with negative emotions). In addition, anticipation and disgust are passive emotions which may be less concrete and thus less easily conceptualized and classified [10].

In relating our linguistic findings to the human raters of emotion in these short blog texts [5], it is perhaps unsurprising that the top-down LIWC analysis found linguistic features relating to joy and anger, given that these were both relatively easily perceived by naive human judges, and LIWC dictionaries are based on human ratings. Counter to hypotheses based on prior work, anger appears more readily expressed through emotional language than sadness. Additionally, LIWC – like expert human raters – identified features of sad texts (sadness words), and LSA could classify fear texts, and to a lesser extent surprise and anticipation. We leave the exploration of these mechanisms in relation to naive/expert human rater ability to future work, but we expect classifiers combining top-down content analysis and data-driven semantic space approaches to be particularly fruitful [cf. e.g., 16].

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 may help technologies to develop a richer, more accurate understanding of the emotional content of existing written excerpts. In turn this may also be used to help imbue our technologies with a richer repertoire of techniques for inserting emotional content into their expressions.

CONCLUSION

Emotion is central to human interactions and has potential repercussions for the collaborative user experience. We explored linguistic characteristics of emotion in short (50 and 200 word) samples of blog texts coded by human judges. Automated content analysis (LIWC) revealed that angry authors used a larger portion of affective language and negative affect words, and that joyful authors used more positive affect words. A semantic space co-occurrence technique (LSA), like expert human raters, was also able to classify fear texts (to a lesser extent surprise and anticipation). Combining these techniques may enable more accurate detection of finer grained emotions and may be better suited to technology applications.

ACKNOWLEDGMENTS

We thank Scott Nowson for use of his blog corpus, and Art Graesser, Sidney D'Mello, Danielle McNamara, and our reviewers for their helpful comments. This work was supported in part by Région de Bourgogne FABER Post-Doctoral Fellowship (05512AA06S2469) (Gill), European Commission Sixth Framework grants NEST 516542 and NEST 029088 (French), and National Science Foundation grant IIS-0705901 (Gergle).

REFERENCES

1. Cowie, R., Douglas-Cowie, E., Tsapatsoulis, N., Votsis, G., Kollias, S., Fellenz, W., & Taylor, J.G. (2001) Emotion recognition in human-computer interaction. *IEEE Sig. Proc. Mag. Vol. 18(1)*, 32-80.
2. D'Mello, S., Craig, S, Witherspoon, A., McDaniel, B., & Graesser, A. (2008). Automatic detection of learner's affect from conversational cues. *User Modeling and User-Adapted Interaction*, *18*(1-2), 45-80.
3. French, R. M. & Labiouse, C. (2002). Why co-occurrence information alone is not sufficient to answer subcognitive questions. *Journal of Theoretical and Experimental Artificial Intelligence, 13*(4), 419-429.
4. Fussell, S. R. (Ed.) (2002). *The verbal communication of emotion: Interdisciplinary perspectives.* Mahwah, NJ: LEA.
5. Gill, A.J., Gergle, D., French, R.M., & Oberlander, J. (2008). Emotion rating from short blog texts. *Proc. ACM Conference on Human Factors in Computing Systems (CHI 2008),* 1121–1124.
6. 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.
7. 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.
8. Leshed, G. & Kaye, J. (2006). Understanding how bloggers feel: Recognizing affect in blog texts. *Work in Progress, Proc. ACM Conference on Human Factors in Computing Systems (CHI 2006), 1019-1024.*
9. Liu, H., Lieberman, H., & Selker, T. (2003). A model of textual affect sensing using real-world knowledge. *Proc. 8th International Conference on Intelligent User Interfaces,* 125-132.
10. Mishne, G. (2005). Experiments with mood classification in blog posts. *Style Workshop, 28th Annual International ACM Conference on Research and Development in Information Retrieval (SIGIR 2005).*
11. Nass, C., Jonsson, I.-M., Harris, H., Reaves, B., Endo, J., & Brave, S. (2005). Improving automotive safety by pairing driver emotion and car voice emotion. *Proc. ACM Conference on Human Factors in Computing Systems (CHI 2005),* 1973-1976.
12. Picard, R. (1997). *Affective Computing*. Cambridge, MA: MIT Press.
13. Pennebaker, J. & Francis, M. (1999). *Linguistic Inquiry and Word Count*. Mahwah, NJ: LEA.
14. Plutchik, R. (1994). *The psychology and biology of emotion*. New York, NY: HarperCollins.
15. Turney, P.D. (2002), Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. *Proc. ACL'02*, 417-424.
16. Wiebe, J., Wilson, T., & Cardie, C. (2005). Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation, 39,* 65-210.