

“Just the Facts”: Exploring the Relationship between Emotional Language and Member Satisfaction in Enterprise Online Communities

Ryan Compton

UC Santa Cruz
rcompton@ucsc.edu

Jilin Chen

Google*
chenjilin@gmail.com

Eben Haber

Couchbase, Inc.
eben@habers.us

Hernan Badenes

IBM Research
hbadenes@ar.ibm.com

Steve Whittaker

UC Santa Cruz
swhittak@ucsc.edu

Abstract

Online enterprise communities often have clear informational goals, e.g. for experts to answer novices’ factual questions. However, emotional support is also critical for developing online social relationships. But what is the optimal balance between informational and emotional communication for such communities? To address this, we develop and validate a model to assess the relative prevalence of emotional versus factual communication. We next test hypotheses about how this prevalence relates to community member satisfaction in enterprise communities. Overall, factual, not emotional, posts predict perceived member satisfaction. This relationship also depends on community subtype. Although prior work argues that Communities of Practice (CoPs) rely on frequent emotional communication, instead we found CoPs showed *less* satisfaction when members focused on emotional concerns. We discuss implications for both community tool design and the practices of community leaders.

Introduction

Recent work explores how the content of members’ conversations affects online community success. A key question concerns the impact of *emotional* versus *factual* language. In certain settings, such as health support communities, emotional language is vital (Wang, Kraut, and Levine 2012). In contrast, for goal-oriented communities that focus on providing answers to information requests, affective language may be less important. This study examines how the prevalence of emotional versus factual content relates to the perceived success of online enterprise communities. To assess the prevalence of emotional language, we first develop a generalized emotional language model, by extending previous methods used on open internet data. We then use this model to determine how emotional versus factual communication relates to community success. We evaluate success by surveying community members.

We examine the role of emotional language in the under-researched context of online enterprise communities. These are organizationally-sponsored communities that support collaboration, knowledge sharing, reuse of resources, ex-

pertise location, innovation and organizational change management. Their enterprise context means that such communities largely promote instrumental goals and informational interaction. Nevertheless, certain subtypes of enterprise online communities like Communities of Practice (CoP) are focused on social relationships and support. It may be that these CoPs require emotional interaction, whereas such interaction is a distraction for goal-oriented communities (Matthews et al. 2015).

Our modeling approach quantifies the amount of emotional versus factual language to test the following predictions: (1) Given the largely instrumental focus of enterprise communities, factual rather than emotional content will promote greater perceived success; (2) this relationship depends on community subtype, with CoPs having higher success when emotional interactions more prevalent.

Method

Adapting the Emotionality Model

We adapt previous work (Aman and Szapkowicz, 2007; Wang, Kraut, and Levine 2012) to develop a model that allows us to distinguish emotional vs factual posts. This involves the following steps: (a) develop a set of explanatory features; (b) construct a model using said features based on open internet data; (c) validate this model generalizes to enterprise communities.

Developing Explanatory Features

To accomplish the first two steps, we used the 10,000 post-response pairs from the Internet Argument Corpus (IAC) of online forum debates about important societal issues such as abortion, religion, immigration, gay marriage and so on (Walker et al. 2012). The corpus contains Factual vs. Emotional annotator judgments for each post based on an *emotionality* scale. We wanted to identify a set of explanatory linguistic features in the forum responses that predicted these emotionality judgments.

Following previous modeling work (Wang, Kraut, and Levine 2012), we used the following lexical resources.

LIWC v2007 (Tausczik and Pennebaker 2010) is a lexicon providing frequency counts of words indexing important psychological constructs, as well as relevant topics (Leisure, Work). The *Emotion Lexicon* (EmoLex) (Mohammad and Turney 2010) contains 14182 words classified into 10 emotional categories: Anger, Anticipation, Disgust, Fear, Joy, Negative, Positive, Sadness, Surprise, and Trust. The *Subjectivity Lexicon* is part of OpinionFinder (Wilson, Wiebe, and Hoffmann 2005). It consists of 8222 stemmed and un-stemmed words annotated by a group of trained annotators as either strongly or weakly subjective. However all of these lexical methods rely on annotated dictionaries that ignore syntax and conversational structure. We therefore used a part of speech (POS) tagger to count the relative frequencies of syntactic nouns, verbs, adjectives and adverbs, use of questions as well as tense and aspect information (Toutanova et al. 2003). A sample of features and their standardized coefficients (beta weights) for the emotionality regression model are shown in Table 1.

Lexical Features		
LIWC	Beta Weights	Standard Error
Pronoun		
“you”	-4.9267	1.51E-86
“I”	-2.5759	1.13E-78
“I”	-2.1540	2.56E-53
Tense		
Past	-0.0889	1.09E-15
Present	-0.0667	2.13E-13
Affect		
<i>Anxiety</i>	0.2151	4.00E-57
<i>Anger</i>	0.0253	3.04E-07
<i>Anger</i>	0.0166	2.18E-05
Topic/Informal Speech		
Cognitive Mechanism	-0.0238	0.0736
Swear	0.0294	5.62E-13
Filler	0.0555	1.50E-07
Punctuation		
Exclamation Point	-0.8399	1.55E-07
Exclamation Point	0.1892	2.62E-25
Question Mark	0.1265	2.40E-27
Emotion Lexicon		
Anger	0.0124	1.40E-16
Disgust	0.0169	1.19E-10
Joy	0.0311	1.12E-17
Sadness	0.0177	1.46E-07
Subjective Lexicon		
Weak Subjective	0.0985	0.016734
Strong Subjective	-0.0570	5.82E-67
Syntactic Features		
Noun		
<i>Singular</i>	-0.0462	4.25E-17
<i>Plural</i>	-0.0006	9.41E-38
Verb		
<i>Base form</i>	-0.0496	2.86E-17
<i>Past participle</i>	-0.0045	3.19E-25
Symbols	0.0324	2.42E-08

Table 1. Example features with weights for the emotionality regression model. Emotional predictors have positive weights and Factual predictors have negative weights.

Creating the Emotionality Detection Model

Our emotionality model uses linear regression, which was highly significant ($p < 2.2e-16$), with an Adjusted R^2 of 0.1968, and a Root Mean Squared Error (RMSE) of 1.38 for predicting the level of emotionality within the IAC. Comparing the human annotators’ standard deviation of 2.08 to the model’s RMSE, shows the model variance is comparable with human annotators.

Application Context

The online enterprise community dataset was collected from a single global enterprise. All communities used the same commercial software, which was available to all employees and enabled them to create and join community spaces. The software incorporated multiple social media tools, and communities often combined different tools, e.g. a wiki to orient new members and set community policy, a forum for Q/A and interpersonal interaction, along with blogs for commentary about pressing topics (Matthews et al., 2015). Usage was voluntary and widespread; there was a proliferation of communities in the company, with 166,000+ communities and 580,000+ distinct members over the five years when we collected data. Throughout the rest of the paper we will refer to the groups that use this software as *communities*.

To identify the subtype of each community, we surveyed members asking them to state the main goal of their community: 55% of communities self-identified as CoPs. CoPs are learning environments where employees with a shared interest communicate, build relationships, and share resources to assist with work. Like many open internet communities, CoP participants ask questions of experts but they also share common resources and experiences and mentor new members. 30% of communities self-identified as *Teams*, collaborating in a goal-directed way on a shared deliverable, similar to internet-based open source communities. These teams were typically developing new software or solutions to meet client needs. Other less common subtypes of communities included Technical Support (6%), Brainstorming (3%), and Recreation (2%). Our main focus in this paper is on the contrast between CoPs and other subtypes.

Survey Measures of Member Satisfaction

We wanted to explore relations between emotional language use and perceived community success. We assessed community success using a survey, only collecting posts for communities where multiple representative members first completed the survey. We surveyed community members as part of a larger research project (Matthews et al. 2015). Respondents represented a wide range of geographies, business divisions and roles, but were all within the same global enterprise. The current paper involves a subset of the survey and communities from that original study. Success was assessed using the most reliable survey question, the member satisfaction probe which asks community members “how well this community is meeting your

needs,” on a scale of 1=very poorly to 5=very well. This question was highly correlated with other related questions, e.g. ‘how successful is your community’ as well as being predictive of other behavioral success measures, including community participation and responsiveness. Details of the survey, demographics, sampling and reliability metrics are provided in (Matthews et al., 2015).

Linguistic Data

Recall that each community used multiple social media tools. For each community, we therefore collected all the content posted to their discussion forums, blogs and wikis over the community’s life. There were ~80,000 posts overall from 93 communities for which we had representative survey data. While all continents were represented, here we examine only English posts.

Behavioral Data

We also gathered behavioral data and metadata for each community. We collected the most common language independent structural behavioral measures of community success used in prior literature (Preece 2001). We include these as control variables in our models.

- **Membership:** # of leaders, # of members, # of contributors
- **Contribution:** # of posts across different community tools, # of wiki, blog and forum posts, # of blog and forum comments
- **Equality:** the gini measure of equality of contribution
- **Consumption:** # of wiki and blog views

Results

Cross Data Set Generalization

We next determined whether the emotionality model developed for online debate forums (see Table 1) generalized to the enterprise context. We therefore created a test set of annotated posts using the same procedure as for the IAC corpus, soliciting emotionality judgments for 1000 Communities posts selected at random from the enterprise data. We evaluated whether the model’s predictions for each post agreed with the judges’ emotionality ratings of that post. Model and judges’ ratings were highly correlated, $r = 0.54$ ($df = 991$, $p < 0.001$) and Kendall’s Tau = 0.37 ($p < 0.001$), showing that the emotionality model derived from open internet debates generalizes to the enterprise community data.

Predicting Member Satisfaction in Enterprise Communities

We next went on to test our hypotheses about the relations between emotionality and community success. Again we used regression methods, where multiple models’ performance was sequentially compared using Adjusted R^2 . To evaluate the role of emotionality, we first created a Control Model containing the following (language independent)

structural variables that have been proposed elsewhere as measures of community success (Preece 2001). Our first model used these structural factors to predict perceived member satisfaction. We next added emotionality to the Control Model to evaluate our prediction that greater emotional communication would relate to lower overall member satisfaction.

All data was centered and the resulting distributions were normal. We tested for multi-collinearity using variance inflation factor (VIF). Following standard procedures, variables with the highest VIF were removed until all variables were under a VIF threshold of 5. The *mean* and *standard deviations* for the control variables were: # members (1729, 2860), # contributors (116, 159), total # posts of all subtypes (605, 660), total wiki+blog+file views (37239, 62537), total comments (372, 553).

We then derived the Control model (Table 2, Model 1) using both-direction step-wise regression using AIC as a criterion. Using a both-direction step procedure is less biased than a one-way step. Stepwise selection led to the removal of Type, Members, Contributors, #Posts, Gini, Word Count, and Views variables for the Control model. Table 2 shows that the Control model (Model 1) has reasonable explanatory power (Adjusted $R^2=0.091$, AIC = 130.96) and is significant ($p=0.013$). # of Comments and Leaders are significant predictors of member satisfaction.

	Model 1			Model 2		
	Control Model			Control + Emotion		
	Adj R^2	P		Adj R^2	P	
	0.09187	0.01381		0.1134	0.00329	
	Std Coef.	SE	P	Std Coef.	SE	P
Intercept	4.01	3.09E-01	***	-139	5.80E+01	.
Emotionality				-0.2508	1.87E+01	*
Leaders	-0.2033	4.43E-03	*	-0.1825	4.36E-03	.
Contributors	-0.1937	3.49E-04	.			
Gini	-0.1728	4.62E-01	.			
# Comments	0.2894	1.02E-04	*	0.2360	9.05E-05	*

Table 2. Model 1(Control) using standard measures for predicting member satisfaction. Model 2 (Control +Emotion) adds the Emotionality feature. Emotionality improves predictive power, showing a negative relationship with satisfaction (** significance $p<0.05$, .’ $p<0.10$).

Facts Not Emotions Relate To Member Satisfaction

Adding the mean post emotionality (Model 2 in Table 2) of a community increases explanatory power (Adj $R^2=0.1134$), decreasing AIC in comparison to the Control Model ($\Delta AIC = -3.17$), and the model is a significant predictor of member satisfaction ($p=0.0032$). Confirming our prediction, the negative coefficient of the emotionality variable indicates that less emotional, i.e. more factual, content relates to satisfaction. It is important to note that this relationship depends on the *degree* of emotionality rather than the *valence* of emotions expressed; independent anal-

yses exploring positive versus negative emotions revealed no significant predictors. Nor did sentiment predict community success.

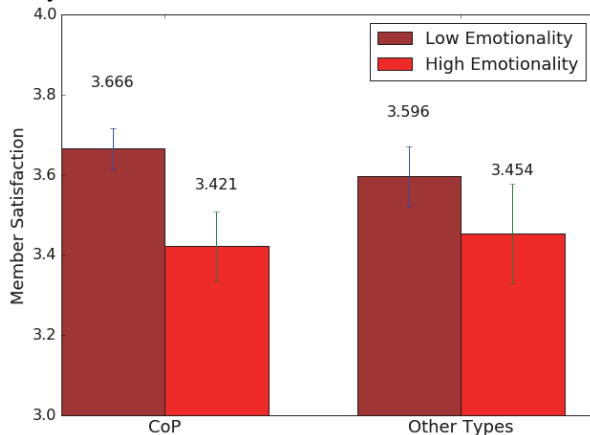


Figure 1. Satisfaction for Communities of Practice (CoP) vs other community subtypes contrasting communities with low and high emotional language use. Communities with greater emotionality show lower Satisfaction, with this difference being greater in CoPs than other community subtypes.

Emotional Language Has A Negative Relationship With Satisfaction In Communities Of Practice

We next examined whether the relationship between emotional language and member satisfaction depended on the *subtype* of community by using a median split of communities showing High versus Low Emotionality. Figure 1 shows how emotionality interacts with Community Subtype to influence satisfaction. It depicts satisfaction in CoPs with other subtypes of communities. The figure suggests that highly emotional language in CoPs has a *negative* relationship to satisfaction, a relationship that is less pronounced in other Subtypes of community. Using Hedge's G to calculate the effect size between the two groups shows that in CoPs ($g=0.714$, $CI=[0.617, 0.810]$) there is a strong difference, while in other community subtypes this difference is weaker ($g=0.314$, $CI=[0.176, 0.453]$).

Discussion

Theoretically our results are important, in showing that the relationship of emotional language is not always intuitive, instead depending on the precise context in which that language is used. Methodologically our work extends prior analyses of emotional community language use. By developing a single emotional classifier we were able to isolate and quantify the role of emotions, and explore how emotional language interacted with other variables such as community subtype.

Our results contrast with work on health support communities where greater use of emotional language is associated with member retention. But this discrepancy may result from the different goals of enterprise and support

communities. Successful enterprise CoPs may rely on factual language, with emotional language signaling a breakdown of communication.

There are limitations to our work however. Although our cross-sectional analysis suggests a relationship between emotionality and member satisfaction, it does not indicate the causal relationship between them. We also examined a single company and our results may not generalize outside this context. However this is only in respect to the relationship between emotional language and member satisfaction, as our model is derived from the open internet and shown to generalize across contexts.

Nevertheless, this increased understanding of the role of emotions could directly inform the design of community tools and practices. Community leaders might apply our results either by introducing policies concerning the use of emotional language or by moderating posts that are 'over-emotional'. And automatic tools that incorporate our emotionality model might help community leaders detect such posts.

Finally, by validating the emotionality model's ability to generalize outside its training domain, we open up new possibilities of exploring emotional language more generally across educational, therapeutic or political settings.

References

- Aman, S., and S. Szpakowicz,. 2007. Identifying Expressions of Emotion in Text. *TSD*.
- Matthews, T., and et al. 2015. They Said What? Exploring the Relationship Between Language Use and Member Satisfaction in Communities. *Proc. of CSCW*.
- Mohammad, S.M., and P.D. Turney. 2010. Emotions Evoked by Common Words and Phrases: Using Mechanical Turk to Create an Emotion Lexicon. *Proc. of NAACL*.
- Preece, Jenny. 2001. Sociability and Usability in Online Communities: Determining and Measuring Success. *Behaviour & Information Technology* 20 (5).
- Tausczik, Y.R, and J.W. Pennebaker. 2010. The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods. *Journal of Language and Social Psychology* 29 (1).
- Toutanova, K., and et al. 2003. Feature-Rich Part-of-Speech Tagging with a Cyclic Dependency Network. *Proc. of NAACL*.
- Walker, M. A., and et al. 2012. A Corpus for Research on Deliberation and Debate. *LREC*.
- Wang, Y., R.E. Kraut, and J. Levine. 2012. To Stay or Leave? The Relationship of Emotional and Informational Support to Commitment in Online Health Support Groups. *Proc. of CSCW*.
- Wilson, T., J. Wiebe, and P. Hoffmann. 2005. Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. *Proc. of HLT/EMNLP*.