Theory Driven Community Analytics and Influence on Community Success

Ryan Compton
University of California-Santa Cruz
Santa Cruz, CA 95060, USA
rcompton@ucsc.edu

Abstract
This paper describes two projects that analyze online community behaviors. The first discusses the modeling of emotional and factual language within argumentative forums. When applied to enterprise communities, this model is found to be a predictor of member satisfaction. The next project covers a unique perspective of member and owner content creation and curation behaviors across the lifecycle of online communities. Both members and owners are found to increase in content posting but the level of linking is role dependent.

CSCW '16 Companion, February 27 - March 02, 2016, San Francisco, CA, USA
ACM 978-1-4503-3950-6/16/02.
http://dx.doi.org/10.1145/2818052.2874355

Author Keywords
Online communities; emotions; facts; user satisfaction; enterprise; communities of practice; machine learning; lifecycle.

ACM Classification Keywords
H.5.3. Information interfaces and presentation; Group and Organization Interfaces; Computer-supported cooperative work; Evaluation/methodology; Web-based interaction.

Introduction
I am an online communities researcher using theory driven methods to explore online community behaviors. I aim to generate meaningful links that can be used to improve online experiences for people. I have explored both emotional language and lifecycle effects and how these affect the success of communities.

Emotional and Factual Language
I examined the role of *emotional* versus *factual* language in online enterprise communities. Enterprise communities have clear informational goals, e.g. for experts to answer novice's factual questions. However, prior work [3, 9, 10] suggests that emotional support is also critical for developing social relationships that are central to Communities of Practice. I assess the relative importance of emotional versus factual

Weights for Factual (top) and Emotional (bottom) Predictors.

Factual Predicting Features	Beta Weights	Standard Error
Pronoun	-4.9267	1.51E-86
"you"	-2.5759	1.13E-78
"I"	-2.1540	2.56E-53
All Percent	-0.8399	1.55E-07
Verb, 3rd person singular present	-0.0136	4.50E-13
Existential *there*	-0.0076	5.80E-08
Past	-0.0889	1.09E-15
Cardinal number	-0.0373	6.95E-33
Adjective, superlative	-0.0039	0.011241
Proper noun, plural	-0.0027	0.002036

Table 1: Main factual predictors are typically were more syntactical than lexical.

Emotional Predicting Features	Beta Weights	Standard Error
Personal pronouns	7.4971	1.04E-124
Exclamation Mark	0.1892	2.62E-25
Question Mark	0.1265	2.40E-27
Apostrophe	0.1126	5.00E-12
Religion	0.0841	2.02E-13
Filler	0.0552	1.50E-07
Symbols	0.0324	2.42E-08
Swear	0.0294	5.62E-13
Particle	0.0208	0.000261
Parenthesis	0.0177	2.61E-05

Table 2: Main emotional predictors included swear words and punctuation.

communication on community success by first adapting previously used algorithms[9] to detect emotional language within forum posts, and then evaluating the effects of emotional language on perceived community success.

To develop this emotion detection model, I needed a training corpus that captured the full range of emotional versus factual language. I collected 10,000 post-response pairs from a corpus of online forum debates about important societal issues such as abortion, religion, immigration, gay marriage, and so on[8]. To determine the emotionality of each debate response, human judges were recruited using Amazon Mechanical Turk. Turker judges were asked to make a judgment about the factual versus emotional basis of a given forum response. Each response was presented in the context of the initial forum post. Turkers were given a Likert probe and judgment scale of emotional to factual.

I used machine learning to model these Turker judgements using features developed from various linguistic lexicons such as LIWC[6], Emotion Lexicon[5], and the Subjectivity Lexicon[11]. Furthermore, I modeled *structural* features of language use, such as the role of parts of speech, question vs answers, and use of grammatical tense. We hypothesized that these might also signal emotional expression, and these were detected using a part of speech tagger[7]. A regression model that matched continuous output from the Turker annotations was best as a simple binary classification between factual or emotional text reduces the amount of information toward the intensity of the two dependent variables. Feature selection was needed in order to reduce the number of factors contributing noise to the model. Many regression models are susceptible to redundant explanatory factors, thus a VIF selection was implemented ensuring all factors had a VIF score below 5. The best model had an adjusted R^2 of 0.19, with tables 1 and 2 showing the weights of the explanatory factors. A higher positive output from the emotional model indicates a higher level of emotion, and a lower value is more factual.

This model was developed on an online debate corpus where there was extensive use of emotional language. To evaluate the model's validity in a communities context, I examined its ability to predict emotionality within an enterprise community. 1000 enterprise community posts were annotated for their level of emotional language using the same annotation procedure as the argumentation data set. The R^2 between the model predictions and the mean annotated value was 0.29, showing that the model had higher performance on the enterprise community dataset than on the argumentation data. This argues for the model's ability to generalize beyond the training dataset.

I used the model's output to assess how emotional versus factual communication affects community member satisfaction. It's observed that successful communities must meet members' needs; so I used a survey measure of member satisfaction[4] to access community success, rather than infer success from behaviors. To evaluate the contribution of emotionality, I constructed a Control Model that contained (non-emotional) structural variables that have been proposed as measures of community success.

	Model 1			Model 2			Model 3		
	Control Model			Control + Emotion			Interaction Model		
	Adj R^2		P	Adj R^2		P	Adj R^2		P
	0.09187		0.01381	0.1134		0.00329	0.1276		0.01726
	Std Coef.	SE	P	Std Coef.	SE	P	Std Coef.	SE	P
Intercept	4.01	3.09E-01	***	-139	5.80E+01	.	3.45	1.26E-01	***
Emotionality				-0.25084	1.87E+01	*			
Type							-110.74	2.58E+01	*
Owners	-0.20334	4.43E-03	*	-0.18253	4.36E-03	.	-0.1637	4.44E-03	
Contributors	-0.19375	3.49E-04	.						
Gini	-0.17282	4.62E-01	.						
Word Count							1921.1	9.41E-03	.
# Posts							-1446.4	6.56E-01	
# Comments	0.28944	1.02E-04	*	0.23604	9.05E-05	*	0.4638	2.73E-04	
# Views							-1906.3	7.09E-03	*
Emotionality * Type							-0.1029	8.33E+00	*
Emotionality * Word count							0.0099	3.04E-03	.
Emotionality * # Views							-0.1108	2.29E-03	*
Emotionality * # Posts							-41084	2.12E-01	

Table 3: Model 1 (Control) using the traditional measures for predicting member satisfaction. Model 2 (Control +Emotion) adds the Emotionality feature. Model 3 (Interaction) includes interaction variables between Emotionality and Control variables. In all cases models are derived using stepwise regression for feature selection ('*' indicates $p<0.001$, '*' indicates $p<0.05$, '.' $p<0.10$).**

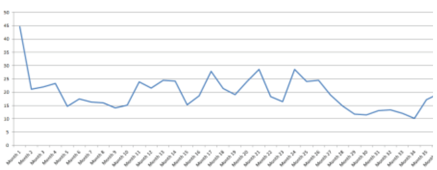

Figure 1: Overall content linking over time. X-axis is time by month iterations; Y-axis is the average amount of linking within a community.

I added emotionality to the Control Model to evaluate the emotionality's effect and created an additional model that includes the interaction of emotionality and the control variables. Table 3 shows these results. There were two main conclusions. Overall, factual rather than emotional posts predict member satisfaction, but this depends on community type. Contrary to expectations, Communities of Practice showed *less* satisfaction when members focused on emotional concerns. These results have important implications for both community tool design and the practices of community. For example, my emotional language detector might be used to warn moderators about overuse of emotional language and they could change their summarization and FAQ posts accordingly.

Understanding Lifecycle Effects on Content and Linking within Enterprise Online Communities

Various descriptive lifecycle models[1, 2] characterize community development, but there is little quantitative exploration of *how communities change over time*, like how they organize and structure extensive long-term content. I conducted a quantitative analysis of 2,010 successful communities over 36 months, analyzing 428,476 posts and 1,246,570 links. I explored *whether* content is organized, *who* organizes it, and which social media *tools* they use.

As expected, and consistent with most theories of community lifecycle, rates of content creation showed a gradual increase over time. Contrary to expectations, active linking (posts that included a URL) actually *decreases* over time. Figure 1 shows that linking rates *dropped* over time, despite having hundreds more posts to organize. Linking behaviors were significantly higher at the beginning of the time series than later as content builds up. The first half of the time period showed significantly more curation than the second (Kolmogorov-Smirnov D = 0.556, p = 0.004369).

Role analyses show that although community members create more content as a community matures (figure 2), members never take over full responsibility for linking (figure 3). These results challenge lifecycle models that argue member responsibility increases as communities mature. Technical implications include the need for more dedicated support for linking, to better exploit linking tools, and to encourage members to take more responsibility for organization.

Figure 2: Role average content creation over time; red line is owners, blue line is members

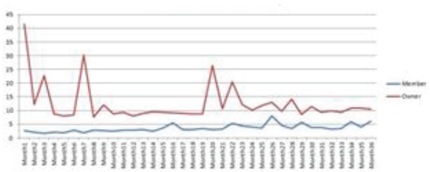

Figure 3: Role average linking over time; red line is owners, blue line is members

Future Work

The previous work on measuring emotionality relied on many various categorical results, but by developing a single emotional classifier, emotionality is able to be isolated to quantify the effects of emotion. Further applications, as the examination of emotionality over a lifecycle, have yet to be explored. There are still improvements to the model with using non-linear features and more advanced modeling techniques.

Examining role differences across content creation and linking is only the first step in the effort to examine how these roles differ and what these behaviors mean toward better community success. More analyses have been conducted on differences within tools available to users (Wikis, Forums, Blogs, etc.) that found differences across roles. This leads to certain tool design implications and further questions such as, how is the division of labor between the two roles being utilized to benefit communities?

Benefits of Attending Colloquium

Getting the opportunity to present my work to a panel of experts is one way I would benefit from attending this colloquium. A discussion with such diverse perspectives is a great way to expand my experience and knowledge when preparing to write my PhD thesis. The prespectives of experts within the field of CSCW is an upper tier of critque that can outline the merit of my previous work and the work that is yet to come. I welcome any opportunity to brainstorm further applications of my work.

References

[1] Iriberri, A. and Leroy, G. 2009. A life-cycle perspective on online community success. *ACM Comput. Surv.* 41, 2 (2009), 1–29.

[2] Kraut, R. et al. Building successful online communities: Evidence-based social design. *MIT*. (2012).

[3] Matthews, T. et al. 2012. Productive Interrelationships between Collaborative Groups Ease the Challenges of Dynamic and Multi-Teaming. *Proc. of CSCW*. 21, 4-5 (2012), 371–396.

[4] Matthews, T. et al. 2015. They Said What? Exploring the Relationship Between Language Use and Member Satisfaction in Communities. *Proc. of CSCW*. (2015).

[5] Mohammad, S.M. and Turney, P.D. 2010. Emotions evoked by common words and phrases: Using Mechanical Turk to cre-ate an emotion lexicon. *Proc. of NAACL*. (2010), 26–34.

[6] Tausczik, Y.. and Pennebaker, J.W. 2010. The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods. *Journal of Language and Social Psychology*. 29, 1 (2010), 24–54.

[7] Toutanova, K. et al. 2003. Feature-Rich Part-of-Speech Tagging with a Cyclic Dependency Network. *Proc. of NAACL*. (2003), 252–259.

[8] Walker, M.A. and et al. 2012. A Corpus for Research on Deliberation and Debate. *LREC*. (2012).

[9] Wang, Y. et al. 2012. To Stay or Leave? The Relationship of Emotional and Informational Support to Commitment in Online Health Support Groups. *Proc. of CSCW*. (2012).

[10] Wenger, E. et al. 2002. Cultivating communities of practice: A guide to managing knowledge. *Harvard Business*. (2002).

[11] Wilson, T. et al. 2005. Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. *Proc. of HLT/EMNLP*. (2005).