The Way I Talk to You: Sentiment Expression in an Organizational Context

Jiang Yang, Lada A. Adamic, Mark S. Ackerman
University of Michigan
Ann Arbor, MI 48109, USA
{yangjian, ladamic, ackerm}@umich.edu

Zhen Wen, Ching-Yung Lin
IBM T. J. Watson Research Center
Hawthorne, NY 10532, USA
{zhenwen, chingyung}@us.ibm.com

ABSTRACT

Sentiment is a rich and important dimension of social interaction. However, its presence in computer-mediated communication in corporate settings is not well understood. This paper provides a preliminary study of people's expression of sentiment in email conversations in an organizational context. The study reveals that sentiment levels evolve over time during the process of newcomers' socialization, that sentiment varies according to tie-strength with the recipient, and that sentiment patterns can be indicative of one's position in the corporate social network as well as job performance. These findings shed light on the complex and dynamic nature of sentiment patterns, and would inspire further explorations and applications of sentiment analysis in organizations.

Author Keywords

Sentiments; Computer-Mediated-Communication (CMC); organization science; social network; socializing; email.

ACM Classification Keywords

H.5.m [Information interfaces and presentation (e.g., HCI)]: Miscellaneous;

INTRODUCTION

Understanding social interaction in computer-supported collaborative systems is crucial in order to better design tools and systems to support communication and collaboration in groups. A wealth of literature has attempted to do so from a variety of angles: studying, e.g., the intensity of social relationships in social networking sites (SNS) [5], social exchange activities in online communities [10], and new members' socialization in SNS [2], online groups [1], and organizations [3].

Recently, sentiment analysis has added another dimension to the study of online interactions. It has been used to track aggregate trends in social media, from understanding aggregate mood of a nation to predicting product reviews and sales (e.g. [4, 6, 7]). However, the expression of sentiment as a rich and important dimension of social interaction, so far, has rarely been explored in an organizational context. This short paper aims to investigate several basic aspects of sentiment expression patterns to shed light on the complex and dynamic nature of sentiment in computer mediated interactions in an organizational context.

By applying sentiment analysis to communication data within a large global company, we sought to answer the following: whether the level of sentiment expressed evolves over time; whether it depends on the strength of tie with the recipient of the communication; and whether it varies by individuals' status within the organization. Each of these questions will be addressed in the Analysis section.

DATA AND APPROACH

Data Description

Our empirical study resides in a global company devoted to information technologies and consulting, with over 400,000 employees across more than 200 countries. In particular, we use a dataset of 15 million email conversations of 8,592 volunteer employees who agreed to share their outgoing messages when signing up for an Intranet service. They were recruited via snowball sampling, by allowing existing users to invite their contacts to join the service. As a result, users in our pool are scatted throughout the corporation's different divisions and countries. A prior study based on the same dataset compared the network characteristics and job roles of the sampled users to the rest of the firm and found minimal differences between the two distributions [9].

Each log entry of the communication records specifies the sender and receiver(s), a timestamp, the subject, and the content of the body of the message. To preserve privacy, the email addresses of users are hashed and the content is transformed into a vector of term frequencies. In order to focus on social interaction patterns, we eliminated spam and mass email announcements (refer to [4] for more information about the pre-processing of the dataset).

We also collected financial performance data for a subset of employees (more than 10,000), all of whom are dedicated business consultants and thus are comparable to each other in terms of job role and performance metrics. These consultants generate revenue by logging "billable hours" and [9] found the ability to generate revenue is an appropriate measure for productivity. Therefore for each

CHI'12, May 5–10, 2012, Austin, Texas, USA.

consultant, performance is measured by the revenue he/she generated during a fiscal year.

It is important to note that our data does not capture a myriad of factors affecting sentiment expression: the individuals' personality, conversation topics, or global time-varying factors such as the financial performance of the company. Therefore each of the measured factors: tenure with the company, strength of tie, and organizational status, can be expected to individually account for at most a very small fraction of the observed variance in sentiment expression.

Methods

We employ a commonly used method for evaluating sentiment, based on frequency of occurrence of words in text based on a subjectivity lexicon [8]. The lexicon includes a list of over 8,000 subjectivity clues that have been annotated as strongly positive (++, e.g., "successful"), weakly positive (+, e.g., "suggestion"), weakly negative (-, e.g., "wrong"), or strongly negative (--, e.g., "suffering"). At the aggregate level, we calculate the frequency of all subjectivity terms to quantify the degree of subjectivity of a set of conversations (see Eq. 1, where F is the frequency count). Thus we measure sentiment-level for each user along each of the four dimensions separately, from strongly positive to strongly negative.

$$sentiment\_level = \sum F_{\text{sentiment\_terms}} / \sum F_{\text{all\_terms}} \quad \text{(Eq 1)}$$

Term-frequency is a simplified method for sentiment analysis, and we recognize that it cannot precisely detect the sentiment of an individual conversation. However, by extracting aggregate patterns over many conversations, we can capture the degree to which an individual tends to express each kind of sentiment.

Since our subjectivity lexicon contains only English terms, we confine our analysis to emails written in English. We further limit the analysis to the 50% of employees who are in the US, as well as 2491 additional employees from Canada, the United Kingdom, and India, retaining nearly 80% of our sample. This restriction minimizes language issues, such as switching to another language to communicate different types of information. Finally, we conduct correlation and regression tests to explore patterns in expression of each polarity and intensity of sentiment separately[1].

ANALYSIS AND RESULTS

Sentiments in Newcomers' Socialization

Upon joining a new group and environment, an individual often goes through a period of socializing and adaptation [3]. Sentiment, as an important attribute of communication, might reflect this adjustment process to some degree. We therefore examine whether people change the polarity and amount of sentiment they express as they gain more experience in the company.

We assess the sentiment level expressed by an employee in their email messages for each month since they first appeared in our dataset. Employees are encouraged to keep all emails, which are automatically archived by the corporate email system. These records are included in our dataset upon the agreement of the volunteers. Even if some individual emails have been deleted, we expect this to have limited effect on the measurement of aggregate sentiment level. In addition, we confine our analysis to the most recent three years, to minimize the occurrence of missing data, and to draw from a sample of employees who had been hired relatively recently and exposed to similar developments within the company. The resulting dataset contains email records of 2586 users who joined since 2007 and had been with the company for at least 12 months.

Figure 1 presents how the sentiment levels for new employees evolve in each of the four dimensions. In general, expression of strong sentiment increases slightly within the first year (--: ρ=.029, sig.<10^{-7}; ++: ρ=.026 sig.<10^{-6}), while people use weak sentiments more frequently and consistently (sig. -: .08; +: .49). A possible explanation is that newcomers are more tentative in expression of sentiment, but gradually become somewhat more comfortable in voicing their opinion. Comfort level may also increase within individual relationships. We therefore examine whether sentiment level changes over time as one continues to communicate with the same contact. We selected user-contact pairs that had exchanged more than 50 messages and computed the sentiment level for each batch of 10 sequential messages. There are small correlations between time-order and

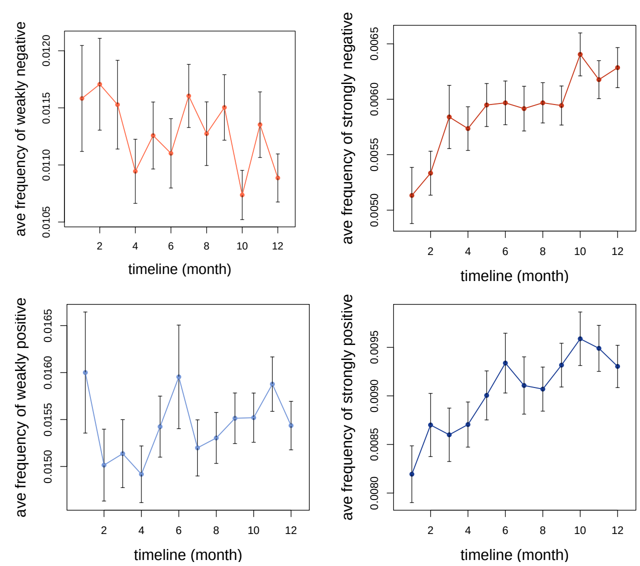

Figure 1. Change of sentiment levels since joining

---

[1] While normally the use of multiple tests would be a concern, the large sample size produces statistically significant results even if the p-values are inflated by an appropriate factor using the Bonferroni correction.

sentiment levels (-: ρ =.02, sig.<10^{-3}; --: ρ=.017, sig.<10^{-3}; +: ρ=.014, sig<10^{-2}; ++: ρ=.01, sig.<.05), suggesting that people tend to feel more comfortable expressing sentiment as they continue to communicate with the same individual.

Sentiments in Social Networks

Prior work (e.g., [5]) found that tie strength is reflected in the type and frequency of interaction in online settings. We thus further analyze whether the strength of a tie relates to the expression of sentiment in this organizational context.

We define tie-strength by the ranked position of a contact in terms of communication frequency. Thus tie-strength is defined relative to communication with one's remaining social contacts. In order to obtain enough observations in terms of both size of the social network and number of conversations, we only include the 3590 users who communicated with more than 100 contacts and sent more than 4000 messages. Please note this is also necessary since our relative tie-strength definition would only make sense when social networks are comparable in size. In particular, we group employee's contacts into three sets: 1~10^{th}, 11~40^{th}, >40^{th}, which correspond to weak, moderate, and strong ties in their social networks.

Figure 2 shows how frequently people express different sentiment levels to a tie of a specified strength. Sentiment levels are normalized for each individual's basic sentiment profile. Except for the weak-positive dimension, there are significant differences between different conversation pairs of different tie-strength (sig.<10^{-7}~10^{-16}). In fact, there are weak correlations between tie-strength rank and sentiment levels on negative (-: ρ = -.05, sig.<10^{-7}; --: ρ = -.06, sig.<10^{-10}) and strong-positive dimensions (ρ = .10, sig.<10^{-16}). In general, people tend to be more negative when conversing along strong ties and more positive when communicating with infrequent contacts.

Sentiments and Social Network Size

The size of one's social network, namely the number of contacts one has communicated with, is an indicator of an employee's level of activity and engagement within the organization. We study whether people with different sizes of social networks exhibit different patterns of sentiment expression. The number of contacts and sentiment levels were computed within a 6-month period, and we confirmed that the results are consistent when the start and duration of the observation period is varied. For the 4114 users included in our observation period (the one we selected to present here), the mean of the network size is 138 and median is 98.

With the exception of the strong-negative dimension, there are weak correlations between the number of contacts and the level of sentiment expressed. In particular, people with larger social networks tend to express more positive sentiments (+: ρ=.11, sig.<10^{-10}; ++: ρ = .07, sig.<10^{-5}), and slightly fewer negative (ρ= -.04, sig.= .005). Since network size and the length of time with the company are correlated (ρ=.10, sig.<10^{-8}), we ran regression of sentiment_level over both variables and found that the network size variable shows a consistent relationship with sentiment, even controlling for length of time in the company.

In order to better observe the pattern, we divided employees into three equal-sized groups, according to the size of their social networks: small, medium, and large. As shown in Figure 3, the trend is generally consistent with the correlation pattern: people who maintain larger social networks tend to be more positive and less negative than those with smaller social networks.

Sentiments as Indicator of Performance

Prior work has found that social network size is not significantly correlated with one's performance in an organization, but having strong connections to authorities is positively associated with revenue [9]. We further examine whether sentiment is correlated to job performance within the company. As explained, we confine this analysis to 549 individuals in our sample who reside within the business consultant group, who have the same job role, and who are evaluated using the same metric: the amount of yearly revenue they generate for the company. In particular, this measure naturally combines two dimensions of a business consultant's performance: value of a unit working hour (hourly pay) and productivity (total consulting work one has done).

For simplified illustration, we divided the sample into high,

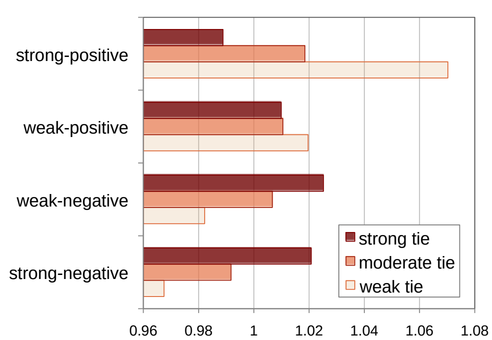

Figure 2. Sentiment levels (normalized) by tie-strength, ranked by communication frequency

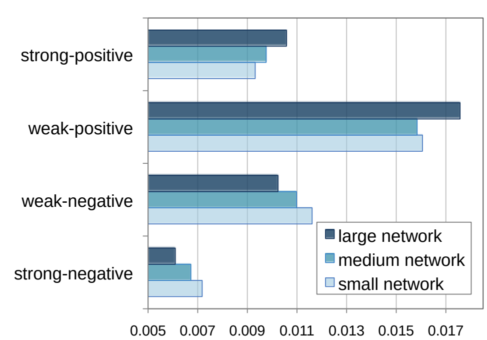

Figure 3. Sentiment levels and social network size

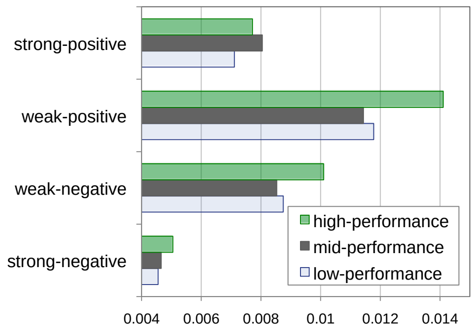

Figure 4. Sentiment levels by performance-groups

middle, and low performance groups. Figure 4 shows that the high-performance group more frequently expresses weak positive and negative sentiment than the other two groups (sig. +: =.005; -: =.01), but there is no significant difference between mid and low performance groups. In fact, there is a positive correlation between performance (log-revenue) and expression of weak-positive sentiment (ρ=.09,sig.<.05), which may suggest positive reinforcement between performance and sentiment.

DISCUSSION

In this paper we presented a preliminary exploration of several basic aspects of sentiment expressed in CMC conversations in an organizational context. The results suggest that sentiment expression depends on the individuals' tenure and performance in the company, as well as their relationships with others. These dynamic and multifaceted characteristics of sentiment could have important implications for designing collaborative systems.

First, during the initial socializing process, newcomers adjust their sentiment levels. Through their first year, they increase their use of strong sentiment words, both positive and negative. In addition, between pairs of individuals who continue to communicate, sentiment is also increasingly expressed. This implies that sentiment expression increases as employees become more comfortable with their environment and build stronger social bonds with other people. This dynamic aspect of sentiment might suggest using sentiment as a dimension to assess newcomers' adjustment progress and status. Moreover, some design features to support proper sentiment expression might be helpful to alleviate some newcomers' shyness in communication, for example, indicating others' sentiment levels might encourage newcomers to adapt to the norms within the participants' organizational context.

Within one's social network, people tend to use positive sentiment words when communicating with infrequent contacts and express negative sentiment more frequently to their strong-tie contacts. This might indicate a higher comfort level in expressing negative thoughts in close relationships, while putting up a positive front to those one does not know well. But when email messages are forwarded, the strength of tie context may be lost. Tools that flag strong sentiment in email may allow users to write and forward email more carefully in intercultural contexts.

Finally, sentiment expression differs according to one's social position in the organization. In particular, people with larger social networks tend to be more positive while less negative. Higher-performers tend to express more sentiment in their messages. Although it is difficult to identify the casual relationship between sentiment and these status variables, we tend to believe that they form some mutual reinforcement dynamics. This kind of association should be worth mentioning in employee training. In addition, giving users feedback on how their sentiment expression compares to others' of similar position and tenure with the company should help users be aware of and adjust their sentiment expression.

ACKNOWLEDGMENTS

The University of Michigan authors were partially supported by NSF IIS-0948639 and NSF IIS-0746646.

REFERENCES

1. Arguello, J., Butler, B., Elisabeth, J., Kraut, R., Ling, K.S., Rose, C., and Wang, X. *Talk to Me: Foundations for Successful Individual-Group Interactions in Online Communities*. in *CHI* (2006).
2. Burke, M., Marlow, C., and Lento, T.M. *Feed me: motivating newcomer contribution in social network sites*. in *CHI* (2009), 945-954.
3. Chan, D. and Schmitt, N., Interindividual differences in intraindividual changes in proactivity during organizational entry: A latent growth modeling approach to understanding newcomer adaptation. *Journal of Applied Psychology* **85**, 2 (2000), 190-210.
4. Diakopoulos, N.A. and Shamma, D.A. *Characterizing debate performance via aggregated twitter sentiment*. in *CHI* (2010), 1195-1198.
5. Gilbert, E. and Karahalios, K. *Predicting tie strength with social media*. in *CHI* (2009), 211-220.
6. Mishne, G. and Rijke, M.d., *Predicting movie sales from blogger sentiment*, in *AAAI 2006 Spring Symposium on Computational Approaches to Analysing Weblogs (AAAI-CAAW 2006)*. 2006.
7. Pang, B., Lee, L., and Vaithyanathan, S. *Thumbs up? Sentiment Classification using Machine Learning Techniques*. in *Proc.Empirical methods in natural language processing*, ACL (2002), 79-86.
8. Wilson, T., Wiebe, J., and Hoffmann, P. *Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis*. in *Proc. HLT-EMNLP*, (2005).
9. Wu, L., Lin, C.-Y., Aral, S., and Brynjolfsson, E., *Value of Social Network -- A Large-Scale Analysis on Network Structure Impact to Financial Revenue of Information Technology Consultants*. 2009.
10. Yang, J., Ackerman, M., and Adamic, L. *Virtual Gifts and Guanxi: Supporting Social Exchange in a Chinese Online Community*. in *Proc.CSCW*, ACM (2011), 45-54.