Increasing Feedback Receptivity by Moderating Negative Valence

Y. Wayne Wu
Department of Computer Science
University of Illinois at Urbana-Champaign
yuwu4@illinois.edu

C&C '17, June 27-30, 2017, Singapore, Singapore

ACM ISBN 978-1-4503-4403-6/17/06.
http://dx.doi.org/10.1145/3059454.3078694

Abstract
Content creators frequently encounter negative feedback that significantly harms their online feedback collection experience. My dissertation examines how negative information influences the perception of the feedback, and how we can increase content creators' resilience to negative feedback. For the next step, I plan to explore ways to solicit more positive content from feedback providers in online environment.

Introduction
Content creators collect feedback from various online crowds [17]. In comparison with friendsourcing feedback from peers, online crowds can quickly generate large quantity of feedback on demand without burning social capitals [10].

One key issue is that content creators frequently encounter feedback of negative valence in online platforms. People are more likely to leave negative feedback [13] and act in offensive ways [4] online. Prior work finds about 20% of all messages exchanged on a popular community site was negative feedback [18]. Negative feedback discourages community participation, harms content creators' affective states [17], and inhibits creativity [1].

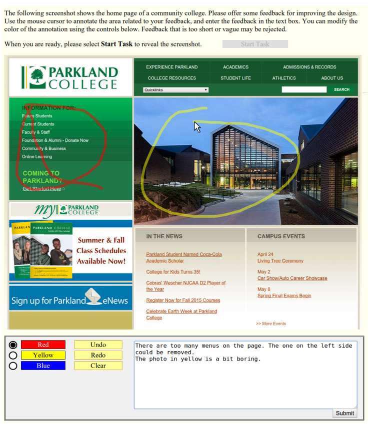

Figure 1. Task Interface for feedback collection. We logged all user actions to build the prediction model.

One common solution for platform designers is to simply remove the negative feedback. While this approach is very effective in limiting the influence of negative feedback, it may cause other issues, such as discarding constructive information in the feedback along the way [6] and discouraging feedback providers from future contribution [7]. To address this issue, my dissertation explores ways to increase content creators' resilience to negative feedback without removing or modifying the original feedback content.

How Information Cues Influence Feedback Perception

In the first component of my dissertation, we explore how information cues influence feedback perception. Unlike traditional face-to-face settings, where content creators can utilize information cues such as the provider's expertise and experience level to evaluate the feedback, online crowd services usually deliver the feedback with no supplementary information. This information opacity makes it difficult for content creators to evaluate the quality of the feedback, especially when conflicting views co-exist in the same feedback set.

Drawing on social transparency theory [12], we explored how information cues, including feedback provider's domain expertise and effort level, influence the perceived quality of the feedback [15]. We conducted a 3x3 full factorial experiment with two factors, namely *effort* and *expertise*, with 2,700 participants recruited from Amazon Mechanical Turk. We manipulated the information cues by telling the participants that the feedback providers had high / low / (no information given) levels of expertise / effort. After reviewing the provided cues, participants proceeded to rate the quality of the feedback. In total, five pieces of feedback of different intrinsic quality levels were rated by the participants. Our result shows, regardless of the intrinsic quality of the feedback, negative information cues lower the perceived quality rating of the feedback, and positive cues have no effect on the rating.

After revealing the influence of the information cues, we explored ways to provide the cues in online environment. Prior work has provided many ways to evaluate expertise, such as performance-based assessments [11], aptitude tests [5,8], or peer prediction [9]. In our work, we focus on providing effort cues. We took a machine learning approach and built an action logging framework that collected behavioral features while feedback providers were composing comments on designs (Figure 1). Later we built a predictive model from the collected data and achieved high prediction accuracy (92% for binary classification of high / low effort).

In this project, we observe negative cues have more powerful influences in the iterative design process than positive cues do. This finding inspired us to probe deeper into the influence of negative information in creative work setting. Specifically, we are interested in how negative feedback influences content creators' affective states, perceptions of the feedback and its providers, and revision extents. We proceeded to explore technological instrumentations that may mitigate the influence of negative feedback. This quest led us to two research projects.

Figure 2. Task interface for the story writing phase. Participants were told to compose a story for children based on the illustration.

Increasing Content Creators' Resilience via Valence Based Feedback Order

Content Creators frequently receive negative feedback in online communities. In this project [16], we show light weight instrumentations, such as valence based feedback ordering, can increase content creators' resilience to negative feedback while requiring minimum change to platforms.

In our study, participants wrote a short story for children based on a given illustration (Figure 2). Then we issued them three pieces of feedback, of which two had positive valence levels and one had negative. Then we tested three feedback orders, namely negative first, negative between, and negative last. We were interested in whether any of these orders could increase content creators' resilience to negative feedback. The experiment also examined the feedback orders in different feedback source conditions, including peers, experts, and anonymous providers. We explored whether the feedback orders would have the same influence in different source conditions.

Our results show presenting positive feedback first and negative last improves participants' affective states and their perception of the feedback. This pattern remains consistent across feedback source conditions. In this work, we show minor system changes, such as feedback orders, can exert a meaningful positive influence over content creators' experience in feedback collection.

Using Coping Activities to Increase Feedback Receptivity

Having explored lightweight mechanisms such as feedback reordering, we continued the line of work by studying standalone interventions that help participants cope with negative feedback. Specifically, we chose to examine three activities, namely self-affirmation, distraction, and expressive writing, all of which had been studied as effective emotional coping interventions in prior work. We were also interested in examining the activities with feedback sets of different valence levels, from all neutral to all negative.

In our experiment, we asked participants to write a short essay on a complex social issue. Later the participants revised their essays based on a set of feedback. Participants performed different coping activities and reviewing feedback set of different valence levels based on experimental conditions. During the experiment, we measured participants' affective states, their perceptions of the feedback set and its providers, and the revision extents.

Preliminary result analysis shows all three coping activities increase participants' resilience to negative feedback. Distraction improves participants' affective states across all valence balance conditions; expressive writing encourages essay revision in neutrally balanced conditions and improves affective states in negatively balanced conditions; self-affirmation improves affective states in negatively balanced conditions. Currently we are preparing the manuscript for submission to an ACM conference.

Encouraging Feedback Providers to Contribute More Positive Content

My previous projects explored the influence of negative feedback and the instrumentations that mitigate the influence on content creators. For the final component of my dissertation, I plan to explore ways to solicit positive valence comments from feedback providers.

Specifically, we are interested in mechanisms and background traces that may help generate more positive feedback. For mechanisms, currently we are considering two options. The first one is to ask participants to revisit their feedback after a one-day delay. Participants' emotional intensity decays over time [14], and they may adjust their language choice when they are less aroused. The second one is to build a scaffolding process that encourages task-involving rather than ego-involving feedback [2]. Prior work shows task-involving feedback increases participants' interest in the task and performance more than ego-involving feedback [3]. The scaffolding may prompt the participants to identify the design aspect they intend to criticize before composing the feedback, and discourage judgements directly threaten content creators' ego.

For traces, we intend to explore whether increasing social transparency may encourage feedback providers to compose more positive feedback. Prior work shows people act more negatively in online environment [13]. We want to explore background traces that may mitigate this influence. Our study may examine traces related to both the content creator's identity and the content creating process. For this part, we try to answer a series of research questions. Does sharing more background information about the content creator encourage more positive feedback? What kind of traces is more useful in this regard? Experience level, demographic info, or personalized details? Will sharing more information about the content creating process help?

Regarding methodology, we may conduct a full factorial experiment using mechanisms and background traces as factors. The feedback solicitation task may be intentionally designed to solicit negative valence feedback. One option may be to ask participants to provide feedback on low quality essays or graphic designs.

On the other hand, we may also consider conducting experiments in a naturalistic setting. We may implement a web browser plugin that encourage users to compose more positive feedback in online communities. The plugin can temporarily save the feedback and ask participants to revise it after a one-day delay. It can also provide scaffolding by rewriting and restructuring the HTML page when participants visit crowd feedback service sites. On the other hand, we can crawl content creators' public activity history and display a summary as a background trace. The plugin may record participants' behavioral data and text entries for later analysis.

Conclusion

My dissertation examines how negative information influences content creators in the iterative design process and how we can mitigate the influence of negative feedback. The last component of my dissertation will focus on encouraging feedback providers to write positive valence comments. By answering this series of questions related to negative feedback, I will contribute empirical knowledge and practical guidelines that improve the online environment for feedback collection.

References

1. Matthijs Baas, Carsten K W De Dreu, and Bernard A Nijstad. 2008. A meta-analysis of 25 years of mood-creativity research: Hedonic tone, activation, or regulatory focus? *Psychological Bulletin* 134, 6: 779–806. http://doi.org/10.1037/a0012815

2. Ruth Butler. 1987. Task-involving and ego-involving properties of evaluation: Effects of different feedback conditions on motivational perceptions, interest, and performance. *Journal of Educational Psychology* 79, 4: 474–482. http://doi.org/10.1037/0022-0663.79.4.474

3. Ruth Butler. 1988. Enhancing and undermining intrinsic motivation: The effects of task- involving and ego- involving evaluation on interest and performance. *British Journal of Educational Psychology* 58, 1: 1–14. http://doi.org/10.1111/j.2044-8279.1988.tb00874.x

4. Justin Cheng, Michael Bernstein, Cristian Danescu-Niculescu-Mizil, and Jure Leskovec. 2017. Anyone Can Become a Troll. *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing - CSCW '17*, ACM Press, 1217–1230. http://doi.org/10.1145/2998181.2998213

5. Steven Dow, Julie Fortuna, Daniel Schwartz, Beth Altringer, Daniel Schwartz, and Scott Klemmer. 2011. Prototyping dynamics. *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*, ACM Press, 2807. http://doi.org/10.1145/1978942.1979359

6. Dana R Ferris. 1997. The Influence of Teacher Commentary on Student Revision. *TESOL Quarterly* 31, 2: 315. http://doi.org/10.2307/3588049

7. Cliff Lampe, Rick Wash, Alcides Velasquez, and Elif Ozkaya. 2010. Motivations to participate in online communities. *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10*, ACM Press, 1927. http://doi.org/10.1145/1753326.1753616

8. Tanushree Mitra, C.J. J. Hutto, and Eric Gilbert. 2015. Comparing person- and process-centric strategies for obtaining quality data on Amazon Mechanical Turk. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*: 1345–1354. http://doi.org/10.1145/2702123.2702553

9. Drazen Prelec. 2004. A Bayesian truth serum for subjective data. *Science (New York, N.Y.)* 306, 5695: 462–6. http://doi.org/10.1126/science.1102081

10. Jeffrey M. Rzeszotarski and Meredith Ringel Morris. 2014. Estimating the social costs of friendsourcing. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, ACM Press, 2735–2744. http://doi.org/10.1145/2556288.2557181

11. James Shanteau, David J. Weiss, Rickey P. Thomas, and Julia C. Pounds. 2002. Performance-based assessment of expertise: How to decide if someone is an expert or not. *European Journal of Operational Research* 136, 2: 253–263.

http://doi.org/10.1016/S0377-2217(01)00113-8

12. H Colleen Stuart, Laura Dabbish, Sara Kiesler, Peter Kinnaird, and Ruogu Kang. 2012. Social Transparency in Networked Information Exchange : A Framework and Research Question. *Network Computing*: 451–460. http://doi.org/10.1145/2145204.2145275

13. John Suler. 2004. The Online Disinhibition Effect. *CyberPsychology & Behavior* 7, 3: 321–326. http://doi.org/10.1089/1094931041291295

14. Philippe Verduyn, Ellen Delvaux, Hermina Van Coillie, Francis Tuerlinckx, and Iven Van Mechelen. 2009. Predicting the duration of emotional experience: Two experience sampling studies. *Emotion* 9, 1: 83–91. http://doi.org/10.1037/a0014610

15. Y Wayne Wu and Brian P Bailey. 2016. Novices Who Focused or Experts Who Didn't? *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*, ACM Press, 4086–4097. http://doi.org/10.1145/2858036.2858330

16. Y Wayne Wu and Brian P Bailey. 2017. Bitter Sweet or Sweet Bitter? How Valence Order and Source Identity Influence Feedback Acceptance. *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition - C&C '17*. http://doi.org/10.1145/3059454.3059458

17. Yu-Chun Grace Yen, Steven P. Dow, Elizabeth Gerber, and Brian P. Bailey. 2016. Social Network , Web Forum , or Task Market ? Comparing Different Crowd Genres for Design Feedback Exchange. *ACM Designing Interactive Systems (DIS 2016)*: 773–784. http://doi.org/10.1145/2901790.2901820

18. Haiyi Zhu, Robert E Kraut, Yi-Chia Wang, and Aniket Kittur. 2011. Identifying shared leadership in Wikipedia. *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*, ACM Press, 3431. http://doi.org/10.1145/1978942.1979453