

# When technologies manipulate our emotions - Implications of the Facebook emotions study

*Rafael A. Calvo<sup>(1A)</sup>, Dorian Peters<sup>(1)</sup>, Sidney D'Mello<sup>(3B)</sup>*

(1) University of Sydney

(2) University of Notre Dame

On those rare occasions when an academic research study receives frenzied media attention, it usually indicates the topic has touched on societal fears in some way. On June 17, 2014 the Proceedings of the National Academy of Sciences (PNAS) published a paper by researchers at Facebook and Cornell University presenting evidence for widespread emotional contagion in online social networks [1]. The Facebook study manipulated the percentage of positive or negative posts appearing in Facebook user news feeds and reported that their manipulations had the anticipated effects of demonstrating emotional contagion at scale. The study contributes interesting results to an underserved area of research, but it triggered an understandable flurry of concern because of a failure to obtain consent from participants before attempting to influence their emotions. Public concern was also intensified by the fact that most people were unaware that Facebook already filters user news feeds (by necessity due to scale), that the company has such unprecedented reach, and that the manipulation involved something as personal as feelings and emotions.

The lack of awareness regarding information filtering supports an illusion of neutrality regarding technology design - the notion that computer programs are bereft of values or moral intent by default. This and other important issues have received less attention amid the flurry of criticism pertaining to research ethics (most of the 120 papers that Google Scholar identified as citing the study focused on ethics). Some open issues include the lack of transparency given restrictions to data access and the difficulty in systematic replication due to a lack of Facebook-level resources. This commentary takes a different approach by discussing the implications of the study on technology design and the criteria on which technologists should base their design decisions.

---

<sup>A</sup> RA Calvo is supported by Australian Research Council Future Fellowship. <sup>B</sup> S. K. D'Mello is supported by the National Science Foundation (NSF).

The Facebook-Cornell study emerged from an attempt to better understand widespread emotion contagion in social networks. The authors experimentally studied how two filtering algorithms influenced the emotional expressions of a large number of users (N=689,003) by manipulating the likelihood of positive or negative posts appearing in the users' News Feed. They then studied the emotional content of users' status updates which were ostensibly influenced by the emotional content in their News Feed. Emotion feelings were measured by the inclusion of emotional terms in the posts, computed with the Linguistic Inquiry and Word Count system (LIWC) [3] that provides psychologically-grounded lists of positive and negative emotional terms (amongst other categories). After a week, those in the positivity-reduced condition (for whom the number of positive posts was reduced) used fewer positive (0.1%) and more negative (0.04%) emotional terms compared to a control condition where a similar proportion of posts were reduced at random (i.e., without respect to emotional content). In contrast, when negative posts were reduced (negativity-reduced condition), users used more positive (0.06 %) and fewer negative (0.07%) emotional terms compared to the control condition. The authors' interpretation of these findings was that users *felt* more negative and positive emotions in the respective conditions – thus emotion contagion occurred.

If we accept the authors conclusions that emotion contagion did in fact occur<sup>a</sup>, this leads us to a larger question for technology: can design ever be emotionally neutral and if not, on what criteria should technologists base design decisions? Beyond the controversy surrounding the way

---

<sup>a</sup> It should be noted that there is some justification to question the central claim that small changes in emotional word use reflect corresponding changes in *felt* emotions. There might be a number of other reasons why people change their emotional expressions aside from emotional contagion, including empathic mimicry, social politics, and compassion. Because the evidence for a contagion effect produced by the study is already statistically small (an effect of .0001 sigma), it might be the case that even if other factors only contribute minimally, it may still be enough to put the claim of emotion contagion in question. If, as the authors state, “[mimicry] may well have been part of the effect” then how small would the resulting effect be were mimicry excluded? For example, there was less than a 0.1% difference in use of positive words in the negativity-reduced condition. Could empathic mimicry be driving 0.05% of this effect? In order to provide sufficient evidence that the data from the Facebook-Cornell study shows that emotions themselves (rather than just emotional expression) can be contagious on Facebook, more work will need to be done to clarify the prevalence of various factors underlying the observed effects. Investigating the various and nuanced motivations behind observed behavior is essential to truly understanding the impact of software designs on users' lives, yet it is, admittedly, a task not generally well suited to big data studies.

the Facebook study addressed informed consent, it is important not to miss the valuable contribution made by this study. This study contributes to a critical area of modern inquiry: how does digital experience and the design of it affect our emotions? This is significant because one issue neglected in the media discourse is that design, be it of a filter, interface, or algorithm, is arguably never neutral. For examples, newspapers use editorial guidelines to filter what information is published and search engines use algorithms to make these decisions. Every design decision must be based on some criteria. As researchers in Values-sensitive Design have made clear, the values and goals of designers and stakeholders will shape the design of any technology [4]. Thus, the obligation to understand the impacts of our design decisions, and to be transparent about what influences them, becomes imperative.

If design is known to affect emotions, above and beyond this particular study on emotion contagion (see [2] for a list of examples), how can we study these effects and how should we apply the knowledge gained? If software design is not neutral, how should a software designer make decisions about what is a “good” design --particularly when we are designing interfaces so closely linked to what we care most about: family, friends and relationships? For instance, if it is in fact *impossible* for Facebook *not* to filter information due to scale, how should the filter criteria be determined? Should filters only ever be randomized and not optimized for the user experience? Or can we look deeper and seek to support greater transparency and user autonomy; what if designers allowed users to make more of these decisions themselves? For example, what if users could set the parameters for their news feed filter or aspects of their search algorithms on their own? Transparency and autonomy seem underexplored opportunities for respecting individual differences and safeguarding against paternalism or misuse.

We also posit that, as we look for criteria upon which to base technology design decisions, we should be turning to the research on psychological wellbeing--that design decisions should seek to promote (rather than hinder) *thriving* (an area we call Positive Computing [2]). It is important to note that wellbeing is not defined simply as an increase in positive emotions. Other determinants include empathy, compassion, self-awareness, engagement, autonomy, and connectedness according to research in psychology. Negative emotions are also an important component of lasting wellbeing [5]. In this view, empathizing with a friend in need or receiving that empathy, may contribute more to one’s wellbeing than merely positive expression. Clearly,

emotional impact represents a rich, nuanced, and complex space of inquiry into which we are still only scratching the surface.

In summary, there is still much left to be understood about how our emotional lives play out in digital experience and how the design of systems, interfaces, and interactions shape our emotional experience. By publishing studies like the one mentioned herein, companies are helping contribute important knowledge, not just to the academic community, but also to those who care about the psychological impact of technology. We believe the controversy over the Facebook study was a useful reminder of how important it is to uphold ethical guidelines in research, and the important role technology plays in our emotional experience. However, we hope it will encourage rather than deter further research into understanding ourselves better and understanding how we, as computing professionals, can make design decisions that are of optimal benefit to society.

## References

- [1] Kramer, A, J.E. Guillory, and J.T. Hancock. "Experimental evidence of massive-scale emotional contagion through social networks. "Proceedings of the National Academy of Sciences (2014): 201320040.
- [2] Calvo RA and D. Peters. (2014) "Positive Computing: Technology for wellbeing and human potential" MIT Press.
- [3] Pennebaker, J. W., & Chung, C., M. Ireland, A. Gonzales, and R.J. Booth (2007). The development and psychometric properties of LIWC2007. *Austin, TX, LIWC*.
- [4] Friedman, B., Kahn, P. H., & Alan Borning. (2008). Value Sensitive Design and Information Systems. In K. E. Himma & H. T. Tavani (Eds.), *The Handbook of Information and Computer Ethics* (pp. 69–101). Wiley
- [5] Fredrickson, B. L., & Losada, M. F. (2005). Positive affect and the complex dynamics of human flourishing. *The American Psychologist*, 60(7), 678.