What social media sentiment tells us about the ebb and flow of a city's moods

Nancy Etcoff, Ph.D.*
Harvard Medical School
netcoff@mgh.harvard.edu

Hugo Liu, Ph.D.*
Hedonometrics
hugo@media.mit.edu

Introduction

Social media sentiment affords new ways to glimpse the affective dynamics of cities and communities because content is timestamped and often geolocated.

As part of the City of Santa Monica's Wellbeing Project, a multi-year effort to understand the subjective well-being of city inhabitants, we designed a social media sentiment index called the Positivity Pulse. The index analyzes and compiles the sentiment from real-time tweets emanating from within the geographic boundary of Santa Monica, and also from the Greater LA basin, as a control. Along with related urban experiments such as the Public Face I art monument (2008), and the MIT Mood Meter (2012), the Positivity Pulse is designed to reflect the ebb and flow of social mood of the community -- be it 71 (upbeat) or 92 (euphoric) -- back onto city inhabitants.

Method

The index collects tweets from Twitter in real-time, and filters to keep tweets which fall within one of two target geographies. The Santa Monica (test) group consists of tweets 1) belonging to a user whose profile location is set to "Santa Monica, CA," or 2) having a geotag originating from within Santa Monica city limits. The Greater Los Angeles (control) group consists of tweets having a geotag originating within a 50km radius of downtown Los Angeles. The Greater LA corpus allows us to control for regional effects like weather and seasonality, and for regional linguistic conventions.

Tweet text is analyzed using the Linguistic Inquiry and Word Count (LIWC) program. In previous work, LIWC word choice has been shown to predict physical and mental health (Gottschalk & Gleser 1969; Pennebaker & Chung 2007). Using LIWC's Positive Emotion (e.g. reassure, entertaining) and Negative Emotion (e.g. ignorance, distraction) word classes, tweets are labeled using the net valence quantity, (P-N), where P is a count of positive words and N is a count of negative words. A tweet is positive if $(P-N)>0$, negative if $(P-N)<0$, and otherwise it is labeled neutral. For a given set of tweets, we calculate an aggregate measure called the "Pos/Neg Ratio," which is the ratio of tweets labeled positive and tweets labeled negative.

* equal authorship, alphabetical order

A second technique is to tally the use of first person singular (I, me, my) and first person plural pronouns (us, ours) in a given tweet, and to classify a tweet as “We” or “I” based on its pronoun content. Pronoun use has been shown to be a sensitive barometer of social engagement / disengagement. Predominant use of first person singular pronouns such as “I” has been found in texts and speech of those feeling depressed, self focused, and disengaged, while first person plural pronouns such as “we” mark the speech and text of those feeling engaged and socially connected (Pennebaker et al. 2003; Rude et al. 2010). For a given sets of tweets, we calculate an aggregate measure called the “We/I Ratio,” which is the ratio of tweets labeled “We” and tweets labeled “I.”

Results

We analyzed a corpus of 10K tweets from Santa Monica and 590K tweets from the Greater LA area, collected over the same six day period, grouped into hours.

We randomly sampled 1000 tweets from this corpus and automatically classified each tweet using LIWC as positive, negative, or neutral. The automated classifications were scored against a gold standard where each tweet was manually labeled by a single human judge as skewing positive, negative, or neutral. Overall, the LIWC classification of tweets achieved 76% accuracy, 91% precision, and 61% recall when compared to the gold standard.

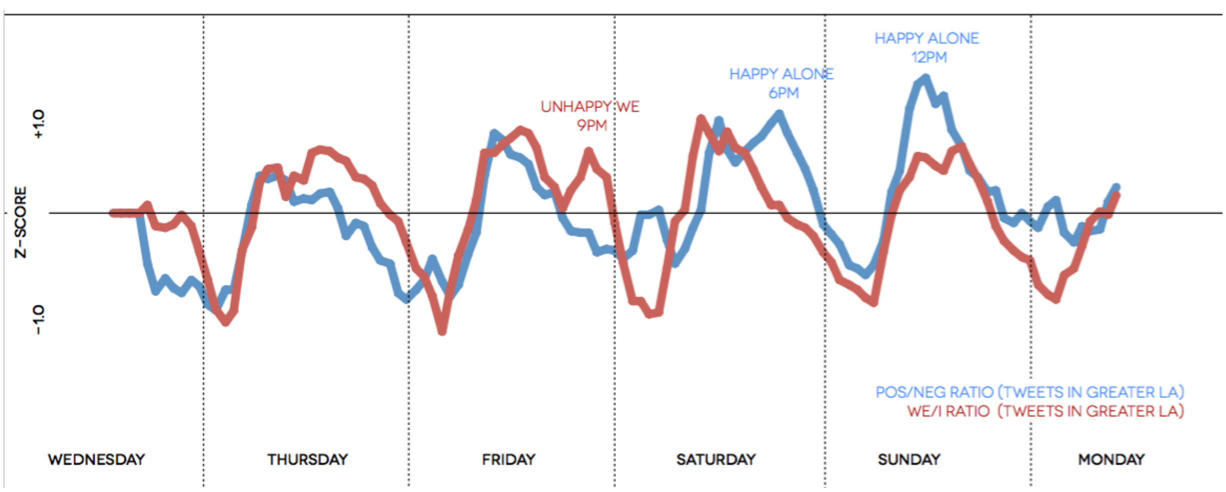

Figure 1

Connecting “we” to positivity. Using the 590K tweet Greater LA corpus, we examine the relationship between pronoun usage and LIWC analyzed mood, plotted in Figure 1. We found that We/I Ratio had a strong positive correlation to LIWC Pos/Neg Ratio. This agrees with the intuition that feeling socially engaged is linked to positive emotions, and together engagement and positive emotions creates positivity.

There are a few deviations between the plots that we can only speculate about, but can test empirically. For example, why did Friday nights yield an “unhappy we,” that is, high social engagement but lower moods, while Saturday evening and Sunday morning yielded a “happy alone” effect or less engagement but more positivity?

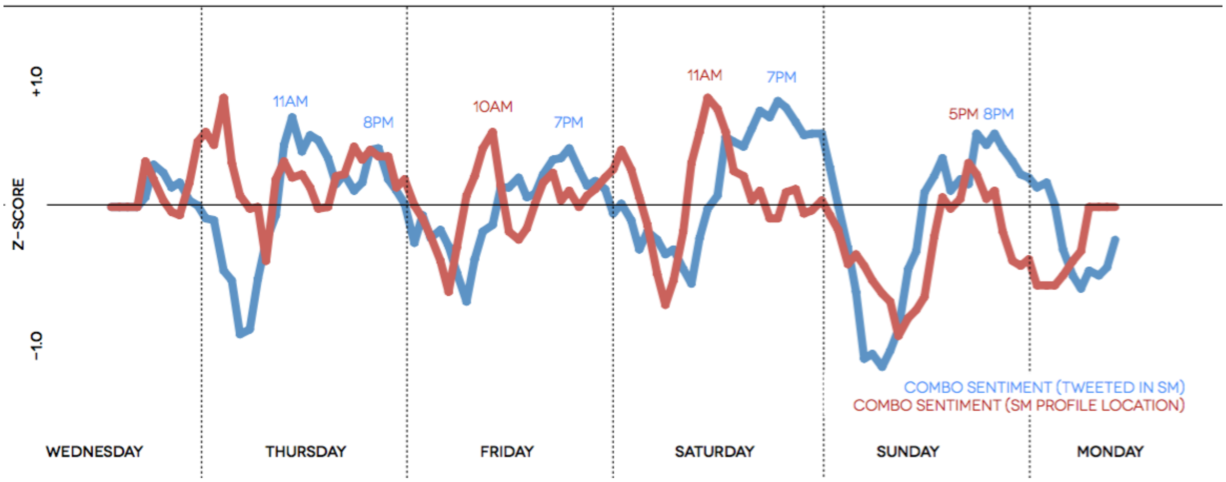

Figure 2

Residents vs. visitors. A challenging problem is knowing whether a Twitter user is a resident or visitor of the City of Santa Monica, short of a formal opt-in process. Santa Monica has many daytime visitors, such as tourists, shoppers, or those who work at offices in the city but live elsewhere, and many nighttime visitors. Of the corpora we collected, the corpus of tweets by users whose profile location was Santa Monica, CA arguably represents residents. The corpus of tweets sent from within city limits arguably represents a mixture of residents and visitors.

One potentially interesting data story emerges when we compare the positivity over time for these two groups (Figure 2). During the weekend, residents have peaks at Friday 10am and Saturday 11am, while the residents/visitors mixture have peaks at 7pm Friday and Saturday Due to the small sample this represents we can only offer speculation that this temporal pattern represents different drivers of engagement and positive affect for residents versus visitors.

Santa Monica vs. Greater LA. A possible and exciting outcome of the Positivity Pulse would be an improved understanding of the ebb and flow of the positivity of the Santa Monica community over various times of day and days of week.

As a unique community, we can expect Santa Monica to have slightly different positivity patterns than the surrounding LA communities. We visualize these differences by plotting the Positivity Pulse for Santa Monica against the positivity of Greater LA (Figure 3).

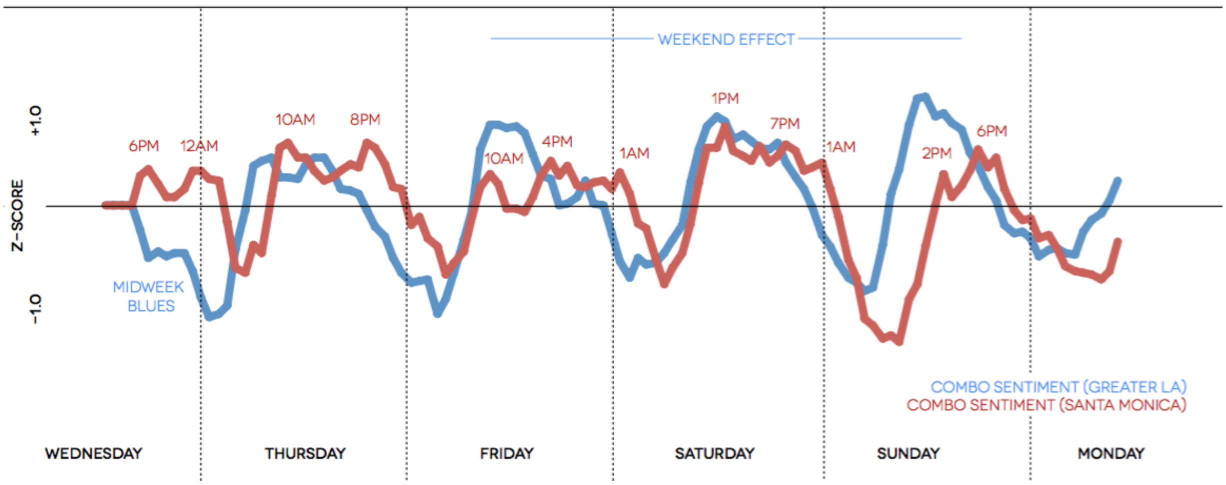

Figure 3

Two differences stand out. First, Santa Monica experiences multiple pronounced positivity peaks per day, indicating perhaps more stimuli for positivity, such as a view of the sunrise and sunset, ample nightlife, or greater leisure in the lives of some residents.

Second, Greater LA seems to exhibit a weekend effect, whereby Friday, Saturday and Sunday contains higher amplitude positivity peaks than the weekdays, and there are some Midweek Blues wednesday night. Santa Monica community positivity does not seem to exhibit a weekend effect, suggesting perhaps a different relationship to the weekend or more continuous enjoyment of life across weekdays and weekend.

Conclusion

The Positivity Pulse would have many purposes both theoretical and practical. Self-report surveys are an important time-tested methodology for gauging well-being but are limited by response biases and by limitations in hedonic introspection. The Positivity Pulse would not replace traditional surveys but provide converging and comparison data. While traditional surveys mine explicit responses to specific questions, the Positivity Pulse mines implicit data that is generated spontaneously. It requires no work or new effort from the participants.

Social media data mining through the Positivity Pulse has the potential to provide cities with a dynamic tool to guide policy and to study the effects of policy changes or neighborhood improvements in real time. The major barrier it faces are public fears about privacy and uncertainty about the uses to which their data would be put.

References

Gottschalk, L. A., & Gleser, G. C. (1969). *The measurement of psychological states through the content analysis of verbal behavior*. Univ of California Press.

Hernandez, J., Hoque, M. E., Drevo, W., & Picard, R. W. (2012). Mood meter: counting smiles in the wild. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing* (pp. 301-310). ACM.

Pennebaker, J. W., & Chung, C. K. (2007). Expressive writing, emotional upheavals, and health. *Handbook of health psychology*, 263-284.

Pennebaker, J. W., Mehl, M. R., & Niederhoffer, K. G. (2003). Psychological aspects of natural language use: Our words, our selves. *Annual review of psychology*, *54*(1), 547-577.

Rude, S. S., Durham-Fowler, J. A., Baum, E. S., Rooney, S. B., & Maestas, K. L. (2010). Self-report and cognitive processing measures of depressive thinking predict subsequent major depressive disorder. *Cognitive Therapy and Research*, *34*(2), 107-115.

Wilhelmer, R., Bismarck, J., & Maus, B. Feel-o-meter: Stimmungsgasometer (2008).