

Not All Moods are Created Equal!

Exploring Human Emotional States in Social Media

Munmun De Choudhury Scott Counts Michael Gamon

Microsoft Research, One Microsoft Way, Redmond, WA 98051, USA
{munmund, counts, mgamon}@microsoft.com

Abstract

Emotional states of individuals, also known as *moods*, are central to the expression of thoughts, ideas and opinions, and in turn impact attitudes and behavior. As social media tools are increasingly used by individuals to broadcast their day-to-day happenings, or to report on an external event of interest, understanding the rich ‘landscape’ of moods will help us better interpret and make sense of the behavior of millions of individuals. Motivated by literature in psychology, we study a popular representation of human mood landscape, known as the ‘circumplex model’ that characterizes affective experience through two dimensions: *valence* and *activation*. We identify more than 200 moods frequent on Twitter, through mechanical turk studies and psychology literature sources, and report on four aspects of mood expression: the relationship between (1) moods and usage levels, including linguistic diversity of shared content (2) moods and the social ties individuals form, (3) moods and amount of network activity of individuals, and (4) moods and participatory patterns of individuals such as link sharing and conversational engagement. Our results provide at-scale naturalistic assessments and extensions of existing conceptualizations of human mood in social media contexts.

Introduction

Social media tools including Twitter continue to evolve as major platforms of human expression, allowing individuals across the globe to share their thoughts, ideas, opinions and events of interest with others. While such content sharing can be objective in nature, it can also reflect emotional states from personal (e.g., loneliness, depression) to global scales (e.g., thoughts about a political candidate, musings about a newly released product or the global economy) (Bollen et al., 2011; Thelwall et al., 2011). We are interested in understanding these emotional states, or *moods*, of individuals at a large scale, manifested via their shared content on social media.

Human emotions and mood have been a well-studied research area in psychology (Mehrabian, 1980). Generally speaking, moods are complex patterns of cognitive processes, physiological arousal, and behavioral reactions (Kleinginna et al., 1981). Moods serve a variety of purposes.

They arouse us to action and direct and sustain that action. They help us organize our experience by directing attention, and by influencing our perceptions of self, others, and our interpretation and memory of events (Tellengen, 1985). By intensifying experiences, moods help identify self-relevant events (Tompkins, 1981). In summary, moods play a critical role in our everyday lives, fundamentally directing our attention and responses to environment, framing our attitudes and impacting our social relationships.

Given the important role of moods in human behavior, a growing body of literature has emerged in the social network/media research community (Mishne et al., 2006; Bollen et al., 2011; Golder et al., 2011) that aims to mine temporal and semantic trends of affect, or detect and classify sentiment in order to better understand polarity of human opinions on various topics and contexts. However, so far, researchers have predominantly looked at analyzing only Positive Affect and Negative Affect, which may miss important nuances in mood expression. For instance, *annoyed* and *frustrated* are both negative, but they express two very different emotional states. A primary research challenge, therefore, is finding a principled way to identify a set of words that truly represent emotional states of individuals.

Further, the ‘landscape’ of emotional states also encompasses an activation component in the context of an environmental stimulus, with some emotions being more *arousing* than others (Mehrabian, 1980; Russell, 1980). This is an important aspect that has not received attention in the social media community so far. For example, *depressed* is higher in arousal than *sad*; though both represent negative affect. Hence moods can be defined as a combination of values on both a valence and an activation dimension, and such characterizations have been advocated extensively in psychology (Russell, 1980; Tellengen, 1985). A popular representation is the ‘circumplex model’ (see Figure 1), that shows moods in 2-dimensional topology defined by *valence* (*x*-axis) and *activation* (*y*-axis): a spatial model in which affective concepts fall in a circle.

The main contributions of this paper are as follows. (1) First, using mechanical turk studies and forays into the psychology literature, we present a systematic method to identify moods in social media, that *capture the broad range of individuals’ emotional states*. Our list comprises

over 200 moods; and for each mood we include both the valence and activation dimension. (2) Second, we analyze these moods in the context of behavioral attributes that define an individual's actions in social media, including mood usage levels, linguistic diversity of shared content, network structure, activity rates and participatory patterns (link sharing / conversational engagement). (3) Finally, through these studies, we provide naturalistic validation of the mood landscape at a collective scale using the valence and activation dimensions, and provide comparisons to findings from the psychology literature, that extend existing conceptualizations of human moods.

Background Literature

Considerable research in psychology has defined and examined human emotion and mood (e.g., Ekman, 1973), with basic moods encompassing positive experiences like *joy* and *acceptance*, negative experiences like *anger* and *disgust*, and others like *anticipation* that are less clearly positive or negative. Of notable importance is the role of intensity or *activation* in emotion and mood, defined as the psychophysiological arousal of the mood in response to an associated stimulus. Together with *valence* (i.e., the degree of positivity/negativity of a mood), these two attributes characterize the structure of affective experience (ref. PAD emotional state model by Mehrabian, 1980; also the circumplex model of affect by Russell, 1980). In these works, authors utilized self-reports of affective concepts to scale and order emotion types on the pleasure-displeasure scale (valence) and the degree-of-arousal scale (activation) based on perceived similarity among the terms.

Turning to analyses of mood in social media, early work focused on sentiment in weblogs. Mihalcea et al. (2006) utilized happy/sad labeled blog posts on LiveJournal to determine temporal and semantic trends of happiness. Similarly, Mishne et al., (2006) utilized LiveJournal data to build models that predict levels of various moods and understand their seasonal cycles according to the language used by bloggers. More recently, Bollen et al. (2011) analyzed trends of public moods (using a psychometric instrument POMS to extract six mood states) in light of a variety of social, political and economic events. Nguyen (2010) proposed models to infer emotional patterns in blogs using normative emotional scores of English words.

Research involving affect exploration on Facebook and Twitter has looked at trends of use of positive and negative words (Kramer, 2010), sentiment extraction from posts based on linguistic features (Barbosa et al., 2010; Koulopis et al., 2011), sentiment classification, as well as sentiment flow in networks (Miller et al., 2011). Recently, Golder et al., 2011 studied how individual mood varies from hour-to-hour, day-to-day, and across seasons and cul-

tures by measuring positive and negative affect in Twitter posts, using the lexicon LIWC (<http://www.liwc.net/>).

Limitations. Despite widespread interest, it is clear that the notions of affect and sentiment have been rather simplified in current state-of-the-art, often confined to their valence measure (positive/negative), with the six moods in (Bollen et al., 2011) being an exception. However, as indicated, the psychology literature suggests that moods are likely to have a 'richer landscape' beyond just their valence. That is, the activation component is equally important, and the inter-relatedness of valence and activation in a systematic fashion is important in conceptualizing affect. Additionally, prior literature on social media has primarily focused on studying affect trends, or alternatively, classifying sentiment. Little attention has been directed towards understanding how the expression of moods is associated with the behavior of individuals, e.g., their linguistic usage of moods, social ties, activity levels, interaction patterns and so on. Through analyzing moods in social media using notions of valence and activation in a circumplex model, we propose that studying collective mood expression is a central contribution of this research.

Identifying Social Media Moods

We begin by discussing our methodology to identify representative mood words – signals that would indicate individuals' broad emotional states. We then characterize the moods by the two dimensions of *valence* and *activation*. In the next part of this section, we present our method of inferring values of these dimensions for all mood words.

What are Representative Moods?

With the goal of developing a mood lexicon in a principled manner that is relevant to social media, we started with the following sources from prior literature in which explicit mood expressions have been identified and investigated.

Sources. (1) Our primary source was a lexicon known as ANEW (Affective Norms for English Words) that provides a set of normative emotional ratings for ~2000 English words (Bradley and Lang, 1999), including valence and activation measurements. (2) Our second source was LIWC (Linguistic Inquiry & Word Count: <http://www.liwc.net/>), wherein we focused on sentiment-indicative categories like positive / negative emotions, anxiety, anger and sadness. (3) Third, we used a list of "basic emotions" provided by (Ortony and Turker, 1990) that included approximately 100 words, including words like *fear*, *contentment*, *disgust*, etc. (4) Our fourth source was the Emotion Annotation and Representation Language (EARL) dataset that classifies 48 emotions in various technological contexts (<http://emotion-research.net/projects/humaine/earl>). (5) Finally, to complement these sources with mood words in online contexts,

we used the list of moods from the blogging website LiveJournal (<http://www.livejournal.com/>). On LiveJournal, blog authors can tag their posts with appropriate mood words.

Mechanical Turk Study for Mood Identification. A number of words in this preliminary list pooled from the five sources bore the notion of positive/negative feelings, but were not appropriate as a mood. For example, “pretty” and “peace” both represent positive affect, but are not convincingly moods. Hence we performed a filtering task to identify mood-indicative words from this candidate list.

This filtering task progressed in two parallel phases. In one phase, two researchers (fluent English speakers) were asked to rate each of these words on a Likert scale of 1 – 7, where 1 indicated “not a mood at all” and 7 indicated “absolutely a mood”. This gave us a set of high quality baseline ratings. In the other parallel phase, we set up a similar mood-rating collection task using the framework provided by Amazon’s Mechanical Turk (AMT) interface (<http://aws.amazon.com/mturk/>). Like the researchers, the turkers were asked to rate the words on the 1 – 7 Likert scale. Each word was rated by 12 turkers, and we considered only those turkers who had a greater than 95% approval rating and were from the United States.

Using the ratings from the turkers and the researchers, we constructed a list of all those words where both the median turk rating and the median researcher rating was at least 4 (mid-point of the scale), and the standard deviation was less than or equal to 1. This gave us a final set of 203 mood words that were agreed upon by both parties to be mood-indicative (examples include: *excited, nervous, quiet, grumpy, depressed, patient, thankful, bored*).

Inferring Mood Attributes Values

Given the final list of representative moods, our next task was to determine the values of the valence and activation dimensions of each mood. For those words in the final list that were present in the ANEW lexicon, we used the source-provided measures of valence and activation, as these values in the ANEW corpus had already been computed after extensive and rigorous psychometric studies. For the remaining words, we conducted another turk study, to systematically collect these measurements.

Like before, we considered only those turkers who had at least 95% approval rating history and were from U.S. For a given mood word, each turker was asked to rate the valence and activation measures, on two different 1 – 10 Likert scales¹ (1 indicated low valence/low activation, while 10 indicated high valence/high activation). We thus collected 24 ratings per mood – 12 each for valence and activation.

¹ The choice of this Likert scale was made to align with the scales of valence and activation in the ANEW lexicon.




Figure 1. Moods represented on the valence-activation circumplex model. The space has four quadrants, e.g., moods in Q1 have higher valence and higher activation.

Finally, we combined the ratings per mood for valence and activation separately, and used the corresponding mean ratings as the final measures for the two attributes (Fleiss-Kappa measure of inter-rater agreement was 0.65).

Circumplex model

The outcome of the two phases of AMT studies was a set of 203 words, with each word characterized by both the valence and activation dimensions. Figure 1 illustrates the *circumplex model* (yellow circle) that results when each mood word (shown as a square) is plotted in 2-dimensions defined by its mean valence (x -axis) and activation (y -axis) ratings. Several features of this resulting space align with similar circumplex models found in prior psychology literature (Russell, 1980). First, the number of different mood words are fairly equally distributed in the four quadrants Q1 – Q4 (starting from top right, moving counterclockwise). Second, while the valence ratings cover almost the entire range between 1 and 10, there are fewer mood words with very high activation or very low activation, compared to valence. Also note that words of neutral valence (e.g., *quiet*) tend to be of lower activation. This is reasonable, as typically more extreme moods (e.g., *infuriated*) generate higher activation. Overall this circumplex model provides us a fine-grained psychometric instrument for study of mood expression in Twitter.

Collecting Labeled Mood Data

We focused on collecting data on moods from the popular social media Twitter. Because of its widespread use, Twit-

ter can be seen as a large and reliable repository for observing the rich ensemble of moods we consider in this paper. We utilized the Twitter Firehose that is made available to us via our company's contract with Twitter. We focused on a full year's worth of Twitter posts posted in English, from Nov 1, 2010 to Oct 31, 2011. Since there is likely to be a considerable volume of Twitter posts that does not reflect moods, and due to the scarcity of mood-labeled ground truth, our major challenge was to eliminate as many non-mood-indicative posts as possible, while simultaneously avoiding labor-intensive manual labeling of posts with moods. We hoped to yield a high precision / low false positive set of posts that truly captured moods on Twitter.

To tackle this challenge, we observed that Twitter users share posts with hashtagged moods, often with the hashtag at the end of the tweet, which might serve as labels for constructing our mood dataset. Consider the following post for instance: “#iphone4 is officially going to be on verizon!!! #excited”. In this light, we followed prior work where the authors used Twitter’s hashtags and smileys as labels to train sentiment classifiers (Davidov et al., 2010). We collected posts which have one of the moods in our mood lexicon in the form of a hashtag at the end of a post. By this process, our labeled mood dataset comprised about 10.6 million tweets from about 4.1 million users.

Verifying Quality of Mood Data

Collection of such labeled data on moods can be relatively easy since it avoids manual annotations or computationally intensive machine learning, but how reliable is it to consider mood hashtags at the end of posts as true indicators of an individual’s emotional state? To answer this question,

we first gathered responses from a set of Twitter users via a study on Amazon’s Mechanical Turk. The study intended to determine for how many cases a hashtagged mood word occurring at the end of a Twitter post truly captures an individual’s mood (without the presence of external signals). Specifically, we displayed a Yes/No question alongside a Twitter post, to which the turker indicated whether the (highlighted) mood hashtag at the end of the post indeed reflected the author’s sentiment. Like before, we again considered U.S. turkers with greater than 95% approval rating history, and then added the requirement of using Twitter at least five times a week (consuming content).

Separately, we also compared the quality of our mood data to a naïve method of spotting mood words anywhere in a Twitter post. Like before, we used Amazon’s Mechanical Turk to determine for how many cases a mood word present anywhere in a post indicated the author’s sentiment. For both the studies, this exercise was conducted over 100 posts per study, and each post was rated by 10 different turkers.

The studies indicated that in 83% of the cases, hashtagged moods at the end of posts indeed captured the users’ moods; while for posts with moods present anywhere, only 58% captured the emotional states of the corresponding users (Fleiss-Kappa measures of inter-rater agreement for both studies were 0.68 and 0.64 respectively), thus providing a systematic verification of the quality of our labeled mood dataset.

Usage Analytics of Moods

Our first study of mood exploration on Twitter data is based on analyzing the circumplex model of moods in terms of the moods’ usage frequencies. We illustrate these mood usage frequencies (count over all posts) on the circumplex model in Figure 2, where the size of squares (i.e., moods) is proportional to its frequency. We note that the usages of moods in each of the quadrants is considerably different (the differences between each pair of quadrants were found to be statistically significant based on independent sample *t*-tests: $p < 0.0001$). The overall trend shows that moods in Q3 (low valence, low activation) tend to be used extensively (*sad, bored, annoyed, lazy*), along with a small number of moods in Q1, of relatively higher valence and activation (*happy, optimistic*). Overall, usage frequencies of lower valence moods exceed those of higher valence moods.

We hypothesize the presence of a “broadcasting bias” behind these observations. Since individuals often use Twitter to broadcast their opinions and feelings on various topics, it is likely that the mood about some information needs to be of sufficiently low or high valence to be worth reporting to the audience. This appears to be particularly true with respect to positive valence terms, with mildly




Figure 2. Circumplex model showing usage frequencies of moods used as hashtags at the end of Twitter posts: larger squares represent higher frequency of usage.

positive moods expressed only rarely. The observation that lower valence moods are shared more often might be due to individuals seeking social support from their audiences in response to various happenings externally as well as in their own lives. The observation that lower activation words dominate usage of these lower valence moods could reflect the simple fact that people experience these moods more than their high activation counterparts (people are *bored* more frequently than they are *infuriated*) or that it is less acceptable to express extreme negative moods.

We also explore how the linguistic content associated with usage of various moods relates to their valence and activation. Like before, we show the moods (as squares) on the circumplex model (Figure 3). The color of the squares indicate a mood’s normalized entropy, defined as the entropy of the textual content (i.e., unigrams over all posts associated with the mood), divided by the total number of posts expressing the mood. In the figure, lighter shades indicate higher entropy. We observe that squares on the right side of the circumplex model (Q1, Q4) tend to have higher entropy than left (Q2, Q3) (statistically significant based on an independent sample *t*-test). This indicates that while positive moods tend to be shared across a wide array of linguistic context (topics, events etc.), negative moods tend to be shared in a limited context, confined to limited topics.

Sociality and Moods

In our next study, we intend to investigate the relationship between the nature of moods expressed and how “social” an individual is, referred to as “sociality”.

For the purposes of this paper, we define sociality to be the ratio of the number of followers (inlinks) to the number

of followees (outlinks) of an individual. We assume that individuals with a ratio close to 1 would be the most “social” since this implies roughly equal engagement on the part of those individuals in both outgoing and incoming social and information exchanges. When the ratio is significantly less than 1, it may indicate that the individual is not interesting enough to others and thereby possibly less “social” to the Twitter audience. On the other hand a high ratio much greater than 1 indicates that the individual is likely an elite user (e.g., a celebrity or news source) who typically is more of a broadcaster and thus also is not very “social” in the context of Twitter. Nevertheless, we note that this measure of sociality does not incorporate the absolute number of followers and followees (i.e., two people with sociality ratios of 20/20 and 2000/2000 respectively would be considered equally ‘social’). We argue, however, that sociality of an individual is likely to be a function of the #followers and #followees together, since they define the structure of his/her ego-network. Hence the ratio is more appropriate than the absolute values. Certainly, other network-centric measures to study the relationship between followers and followees could be used (notably, clustering coefficient), and would be worth exploring in the future.

We again use the circumplex model to represent relationship between moods and sociality. In Figure 4, we show the circumplex model where each square is a mood and is represented by a color on an RGB scale, indicating the mean #followers/#followees ratio of individuals who shared the mood. A red square denotes a low #followers/#followees ratio of the individuals who shared the mood, while blue denotes a high #followers/#followees ra-




Figure 3. Circumplex model showing entropies of moods in terms of the content of posts: higher valence moods (shown in lighter shades) have higher entropy.




Figure 4. Circumplex model of moods showing the relationship between mood expression and sociality (defined as #followers/#followees ratio).

the respective background probabilities (1st and 3rd lines in (a)). On the other hand, several highly positive moods (higher valence) co-occur more frequently than usual with a link (2nd line in (a)). This is likely to indicate that when individuals choose to express moods in a link-bearing post, it is more frequently a positive than a negative mood. In general, mood expression, positive or negative, is quite low in posts with @-replies (4th line in (a)).

2. Low activation moods tend to co-occur less frequently with a link, compared to the respective background probability (1st line in (b)), while high activation moods appear to be relatively more frequently co-occurring with links (2nd line in (b)). Finally, moods co-occurring with @-replies are overall quite infrequent, irrespective of their activation measures. In fact, a large number of moods co-occur with @-replies with a probability much lower than the respective background probability.

Based on these findings, it appears that there is a difference in the manner in which moods are shared in various participation contexts. For link sharing, when links and moods co-occur, those moods tend to be both positive and




Figure 6. Mood distribution over valence and activation, for two types of participatory behavior: external information sharing, and conversational engagement.

of moderate to high activation. Replies, on the other hand, show little in the way of mood expression.

Discussion

In this paper we extended literature on mood/affect analysis in social media by studying more than 200 fine-grained human moods at a large scale, and in light of not only their valence but also their activation. We utilized the notion of the circumplex model popular in the psychology literature (Russell, 1980) and studied various aspects of the relationship between mood expression and human behavior in social media: mood usage, social ties, activity level, and participatory patterns (link sharing, conversational engagement). To the best of our knowledge, *naturalistic validation of the circumplex model* and quantifying mood usage at scale was done for the first time in this paper.

Our findings indicated both similarities and differences with prior literature. Although the circumplex model had rather equal distribution of moods in each quadrant in terms of the valence and activation dimensions (drawn from and consistent with the literature), mood usage frequencies as well as diversity of linguistic content in mood sharing varied considerably across the circumplex model (unknown in prior literature). While negative and low activation moods were more frequent, a few positive and high activation moods showed high use. At the same time, the linguistic use of positive moods seemed to span a variety of contexts / topics (high entropy), while that of negative moods seemed concentrated in a limited context. The addition of usage frequencies and linguistic use to valence and activation in this manner thus helped to validate the topology of the circumplex model in the social media context and showed that *all moods are after all not created equal!*

Further, we observed a relationship between the social ties formed by individuals and the moods expressed: positive moods were correlated with highly ‘social’ people. This confirmed findings in the literature that show that in our social interactions, moods help regulate social relationships (Averill, 1976), and can stimulate pro-social behavior (Vittengl and Holt, 1998). Similarly, the finding that highly active individuals tend to share more positive moods is supported by evidence in prior literature indicating that daily activities and social events tend to be correlated with positive affect (Clark et al., 1988).

Findings from the participatory patterns and mood expression analysis further revealed insights that augment prior literature. Frequent co-occurrence of positive moods and link sharing indicated that the external information sharing behavior typically bears positivity. This might be related to the well-studied observation that positive moods are contagious and propagate in a network (Lykken, 2000). Finally, although positive moods have historically been found to be associated with different types of conversations

and communication processes (Buck, 1984; Mehl et al., 2010), our observations indicated otherwise: moods, whether positive or negative, do not co-occur frequently in an @-reply post. The reason for this departure from previous research could be due to a number of aspects of Twitter, e.g., usage norms wherein people are likely to communicate little about their moods through an explicit hashtag when engaged in conversation directed toward someone, compared to a context when they are reaching out to a broad audience on some issue of interest.

Our work provides ample opportunity for future research. A limitation of current work is that we have analyzed mood characteristics across *all* Twitter users, without making distinctions between *user types*. For instance, the mechanisms of mood expression might be different for celebrities compared to ordinary individuals. Finally, while we have comprehensively considered a variety of mood words that typically occur on Twitter, the lexicon could be augmented in the future with moods that have ‘evolved’ on these novel communication modalities and have not been investigated traditionally (e.g., ‘#yay’ would indicate a positive, moderate activation mood, frequent on Twitter). From an applications perspective, this research provides promising avenues to build mood classifiers useful in search and advertising, or to enable organizations track behavior of populations in health, socio-economics or urban development domains.

Conclusions

Moods impact human behavior, responses to environmental stimuli, and social ties. Naturally, it is imperative to understand characteristics of moods on social media in order to decipher the collective behavior of large populations. In this paper, we identified and studied a variety of moods that frequent Twitter posts. We used the dimensions of valence and activation to represent moods in the circumplex model and studied the topology of this space with respect to mood usage, network structure, activity, and participatory patterns. This work provides naturalistic validation of the circumplex model using Twitter data, extends this conceptualization of human mood to incorporate usage levels and linguistic diversity, and provides evidence connecting mood usage to social structure and interaction behaviors.

References

Averill, J.R. (1976). Emotion and anxiety: Sociocultural, biological, and psychological determinants. In M. Zuckerman & C.O. Spielberger (Eds.). *Emotion and anxiety: New concepts, methods, and applications* (pp. 87-130).

Barbosa, L. & Feng, J. (2010). Robust sentiment detection on Twitter from biased and noisy data. *Proc. COLING 2010*. 36-44.

Bollen, J., Mao, H., & Pepe, A. (2011). Modeling Public Mood and Emotion: Twitter Sentiment and Socio-Economic Phenomena. In *Proc. ICWSM 2011*.

Bradley, M.M., & Lang, P.J. (1999). Affective norms for English words (ANEW). Gainesville, FL. *The NIMH Center for the Study of Emotion and Attention*.

Buck, R. (1984). *The communication of emotion*. New York: Guilford.

Clark, Lee A. & Watson, David (1988). Mood and the mundane: Relations between daily life events and self-reported mood. *Journal of Personality and Social Psychology*, Vol 54(2), 296-308.

Davidov, D., Tsur, O., & Rappoport, A. (2010). Enhanced Sentiment Learning Using Twitter Hashtags and Smileys. *Proceedings of the 23rd International Conference on Computational Linguistics Posters*, (August), 241-249.

Ekman, P. (1973). Cross-cultural studies of facial expressions. In P. Ekman (Ed.), *Darwin and facial expression: A century of research in review* (pp. 169-229).

Golder, S. A., & Macy, M. W. (2011). Diurnal and Seasonal Mood Vary with Work, Sleep and Daylength Across Diverse Cultures. *Science*. 30 Sep 2011.

Kleinginna, P.R., & Kleinginna, A.M. (1981). A cataloged list of motivation definitions with a suggestion for consensual definition. *Motivation and Emotion*, 263-291.

Kouloumpis, E., Wilson, T., & Moore, J. (2011). Twitter Sentiment Analysis: The Good the Bad and the OMG! In *Proc. ICWSM 2011*.

Kramer, A. D. I. (2010). An unobtrusive behavioral model of “gross national happiness”. In *Proc. CHI 2010*. 287-290.

Lykken, D. (2000). *Happiness: The nature and nurture of joy and contentment*. New York: St. Martin's Press Griffin.

Mehl MR, Vazire S, Holleran SE, Clark CS. (2010). Eavesdropping on happiness: well-being is related to having less small talk and more substantive conversations. *Psychol Sci*. 2010 Apr;21(4):539-41.

Mehrabian, Albert (1980). Basic dimensions for a general psychological theory. pp. 39-53.

Mihalcea, R., & Liu, H. (2006). A corpus-based approach to finding happiness. In *Proceedings of computational approaches for analysis of weblogs, AAAI Spring Symposium*.

Miller, M., Sathi, C., Wiesenthal, D., Leskovec, J., & Potts, C. (2011). Sentiment Flow Through Hyperlink Networks. In *Proc. ICWSM 2011*.

Mishne, G. & de Rijke, M. (2006). Capturing Global Mood Levels using Blog Posts. In *AAAI 2006 Spring Symposium on Computational Approaches to Analyzing Weblogs*.

Nguyen, T. (2010). Mood patterns and affective lexicon access in weblogs. In *Proceedings of the ACL 2010 Student Research Workshop (ACLstudent '10)*. 43-48.

Ortony, A., & Turner, T. J. (1990). What's basic about basic emotions? *Psychological Review*, 97, 315-331.

Russell, James A. (1980). A circumplex model of affect. *Journal of Personality and Social Psychology*: 39, pp. 1161-1178.

Tellegen, A. (1985). Structures of mood and personality and their relevance to assessing anxiety, with an emphasis on self-report. In A. H. Tuma & J. D. Maser (Eds.), *Anxiety and the anxiety disorders* (pp. 681-706).

Thelwall, M., Buckley, K., & Paltoglou, G. (2011). Sentiment in Twitter events. In *J. Am. Soc. Inf. Sci. Technology*.

Tompkins, R.D. (1981). *Before it's too late...: The prevention manual on drug abuse for people who care*. Englewood Cliffs, NJ: Family Information Center.

Vittengl, J. R., & Holt, C. S. (1998). A time-series diary study of mood and social interaction. *Motivation and Emotion*, 22(3), 255-275. Springer.