OSN Mood Tracking: Exploring the Use of Online Social Network Activity as an Indicator of Mood Changes

James Alexander Lee
School of Engineering and
Digital Arts
University of Kent, UK
jal42@kent.ac.uk

Lu Bai
School of Engineering and
Digital Arts
University of Kent, UK
l.bai@kent.ac.uk

Christos Efstratiou
School of Engineering and
Digital Arts
University of Kent, UK
c.efstratiou@kent.ac.uk

Ubicomp/ISWC'16 Adjunct, September 12–16, 2016, Heidelberg, Germany

ACM 978-1-4503-4462-3/16/09...$15.00.
http://dx.doi.org/10.1145/2968219.2968304

Abstract
Online social networks (OSNs) have become an integral part of our everyday lives, where we share our thoughts and feelings. This study analyses the extent to which the changes of an individual's real-world psychological mood can be inferred by tracking their online activity on Facebook and Twitter. By capturing activities from the OSNs and ground truth data via experience sampling, it was found that mood changes can be detected within a window of 7 days for 61% of the participants by using specific, combined online activity signals. The participants fall into three distinct groups: those whose mood correlates positively with their online activity, those who correlate negatively and those who display a weak correlation. We trained two classifiers to identify these groups using features from their online activity, which achieved precision of 95.2% and 84.4% respectively. Our results suggest that real-world mood changes can be passively tracked through online activity on OSNs.

Author Keywords
Online Social Networks; Facebook; Twitter; Social Psychology; Emotion; Mood; Sentiment Analysis.

ACM Classification Keywords
H.3.5 [Information Storage and Retrieval]: Online Information Services

Introduction

Real-world communication between individuals is a multi-modal experience, involving much more than the exchange of verbal information. Every day social interactions are enriched by our natural ability to interpret visual or auditory clues that help us understand the emotional state of our counterpart [1]. In recent decades a significant part of our social lives have shifted into the digital world, through the use of online social networks (OSNs) such as Facebook and Twitter. However, compared to real-world social interactions, OSNs do not facilitate the discovery of emotional or mood changes of the users, unless they are explicitly declared by the users themselves. Considering the significant role that OSNs play in the daily lives of individuals, we believe that there is great value in transforming OSNs into an *affective communication* medium where mood or emotional changes can be passively communicated.

In this paper we explore the feasibility of inferring mood changes of OSN users by analysing their online activity on Facebook and Twitter. A number of previous studies have demonstrated that online activity can be a predictor for the detection of long-term psychological conditions [5] or a means for larger scale analysis of emotional trends in groups [12]. In this work we aim to identify signals from both Facebook and Twitter that can be exploited to detect individual user's mood changes within a relatively short time frame (7 days).

Affective communication may involve the understanding of the psychological state of users participating in an interaction. However, considering that OSN communication is typically asynchronous and sparse, we believe that real-time emotional detection in the context of OSNs is impractical and possibly unfeasible. Instead our aim is to capture psychological mood changes that can span over longer periods of time. The term "mood" is used to describe longer lasting changes of psychological state [7] and differs from the traditional concept of "emotion". Emotions such as anger, fear or happiness are short, intense episodes triggered by a specific person or event that can last a few minutes or hours. Mood is a state that lasts a lot longer and is less intense than emotion. A person's mood will generally remain stable for hours to days in length and does not usually have a trigger. In this work we demonstrate that OSN activity of users can contain signals that can reveal the mood changes of individual users. We do so by exploring the relationship between online activity and actual mood changes, captured through experience sampling over a period of one month.

More specifically, we conducted a study involving 36 university students lasting for one month. Using online activity traces captured from their Facebook and Twitter accounts as well as self reported daily logs of mood changes, we demonstrate that it is feasible to detect changes of mood within a window of 7 days, for 61% of the participants. We developed a machine learning classifier that can identify which users demonstrate such correlations, along with the type of correlation (positive / negative). With these results we demonstrate the feasibility of an automated technique that can discover OSN users whose activity is strongly correlated with mood change and predict the type of mood changes that can be observed, aggregated over a window of 7 days. To the best of our knowledge this is the first case of exploring links between OSN activities and real-world mood data captured through experience sampling. This approach allows us to find evidence about correlations between online activity and real-world mood for particular types of social networking users.

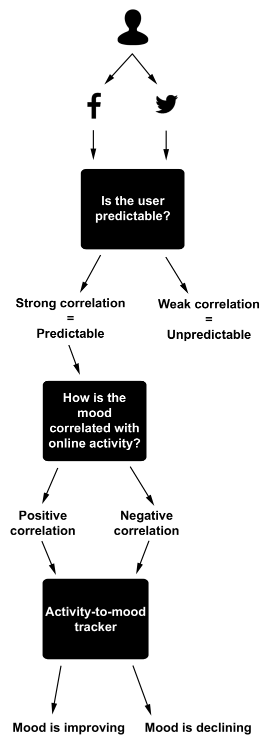

Figure 1: The mood tracking system. (i) User's Facebook and Twitter accounts are passively tracked. (ii) Identify whether this user's mood changes can be predicted through online activity. (iii) Identify whether their mood is positively or negatively correlated with their online activity. (iv) Online activity used to passively track mood changes.

Related Work

Research exploring psychological state online has shown varying results. Sentiment analysis on Facebook status updates has been shown to predict subjective well-being, such as satisfaction with life [9]. Sentiment analysis has also been used on Twitter to analyse and visualise mood across different locations in Brazil over a 24 hour period [10], to track changes in worldwide emotion in response to major social, political, cultural and economic events [3] and to detect emotions in the tweets of individuals [2]. It has been suggested that research which intends to analyse psychological state should be backed by ground truth data to be considered as an accurate representation of real life [5].

OSNs have also been used for purposes other than exploring psychological state. Using sentiment analysis on geotagged tweets, researchers have been able to track the spread of influenza [14]. Twitter has been used to detect and diagnose major depressive disorder by backdating a year of Twitter usage. This work used specific online features which are synonymous with real-world depression symptoms, such as depressive lexicon, interaction with other users and time of tweets representing insomnia [5]. Furthermore, features on Facebook such as sentiment, number of posts and overall usage patterns have been used to detect, characterise and predict postpartum depression in new mothers [4]. More recent work explores linguistic and interaction-based metrics on Reddit support communities to detect the shift from mental health discussion to suicidal thoughts [6].

In the majority of this work, the primary focus is on the analysis of aggregated data over large populations or relatively long time periods. Ground truth is typically coarse-grained, representing particular conditions or significant shifts in mental state. In contrast, in our work we are considering the case of short term (in the scale of days) individualised tracking of mood changes. Furthermore, we attempt to explore online activity in relationship to fine-grained ground truth, captured through experience sampling.

Motivation

This work aims to explore the feasibility of inferring changes in the mood of individuals through the analysis of their activity in OSNs. Our aim is to develop a system that can infer mood state signals in order to enhance online social interactions. Specifically, we envisage a system that can act as a "soft sensor" that can predict mood changes through online social data.

In order to achieve this we aim to develop a multilevel classifier that predicts the change of mood for OSN users. As seen in Figure 1, we consider a system that consists of:
(i) a classifier that can identify OSN users whose mood changes can be predicted through online activity, (ii) a classifier that can identify how their online activity is related to their current mood, and (iii) a tracker that maps online activity to mood changes. We attempt to construct such system experimentally. Specifically, using appropriately collected datasets, we explore whether mood changes can be correlated with particular features extracted from online activities on social networking sites. The results of this analysis is then used to train the classifiers that can identify users whose mood changes can be predicted using online data.

Data collection

Methodology

We conducted a study where we collected both online data and ground truth data about the mood changes of OSN users. Each participant was requested to share information with the research team about both their OSN activity and to report mood changes during the study.

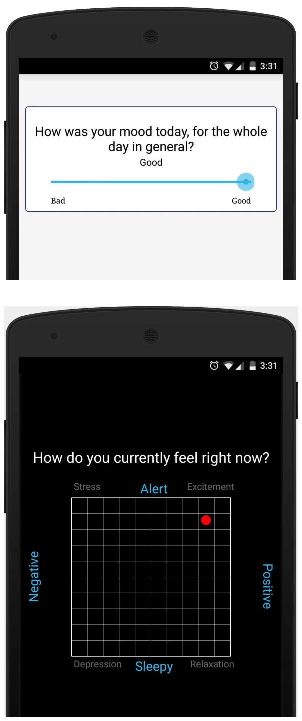

Figure 2: Surveying the participant's mood (top) and emotion (bottom) within the Easy M application.

Recruitment

The study was aimed at OSN users who maintain a relatively frequent interaction with Facebook and Twitter. Considering the general statistics about daily use of OSNs [11] we targeted younger adults between 18-25 years old who are considered more active online users. Specifically we targeted students at a British University. The study was advertised through a range of University forums and each interested student was asked to submit a short survey about how often they use OSNs. Of the 73 people who initially registered their interest, 36 were chosen to participate in the study.

We ran the study during the end of the academic year and over the summer break. Our intention was to ensure that students who participated would be part of the study during their exam period as well as the summer break. With this approach we anticipated that participants would demonstrate a wider variability of mood changes, possibly due to exam pressure and the relaxed summer break period that followed. Each participant was expected to be part of the study for approximately one month including part of their exam period and at least two weeks after the end of their exams. The average duration of participation was 28 days.

Ground Truth Data

The ground truth data collection required participants to install the smartphone application Easy M [1] for Android or PACO [2] for iOS. Both applications prompted the participants at 10pm every day to answer two questions: "How was your mood today, for the whole day in general?" and "How do you currently feel right now?". Participants could respond to Q1 on a 7-point Likert scale with clearly marked extremes: -3 (bad) - +3 (good) (Figure 2 (top)). The one dimensional input was selected to match with the commonly used *Positive and Negative Affect Schedule (PANAS)* to detect mood states [16]. Q2 was answered using an affect grid, initially proposed by Russell [13], in which the participant can easily record their emotion on a two dimensional grid: valance (x) and arousal (y) (Figure 2 (bottom)). The two questions were carefully worded to ensure the participants understood and could differentiate between the two answers that were required. By including questions about both their general mood throughout the day, and their emotional state at the time of question, our intention was to ensure that participants would not erroneously report their current emotional state as their daily mood.

In order to maximise the response rate for the ground truth collection we tried to keep the required interaction to the absolute minimum (one notification per day). 10pm was chosen for the notification time which is before most participants are likely to be asleep. At the end of the study, the overall response rate for the ground truth was 88%.

Self-reported mood changes are considered subjective. Particular people may tend to be more positive or negative in general about their mood. To accommodate for such bias we followed the same methodology as in [8]: We normalised the reported mood using the z-score as an objective metric of mood change: $z_{it} = \frac{x_{it}-\mu_i}{\sigma_i}$ where x_{it} is the mood reported by the participant i on date t, μ_i and σ_i are the average and standard deviation of mood for that participant over the duration of the study.

Online Data

Online data involved data about the participants' activities on Facebook and Twitter. Two crawlers were developed that collected online activity data from both OSNs. For Facebook, we used the Graph API to collect data about the participant's activity from their personal timeline and their home

[1] http://easy-m.io/
[2] http://pacoapp.com/

Facebook - posts	
Posts per window (avg.)	1.4
Days without posts per window (avg.)	5.69 days
Twitter - tweets	
Tweets per window (avg.)	11.2
Days without tweets per window (avg.)	6.39 days

Table 1: Sparsity of posts and tweets within a 7-day window. On average, almost six of the seven days are without posts and tweets from the participants.

feed. These included their own posts along with likes and comments on other people's posts. It is important to note that we were able to collect personal posts regardless of the privacy settings (i.e. public, friends only). The crawler also collected select profile information such as the participant's demographics. From Twitter, the crawler collected all of the participant's tweets, including replies and retweets, together with their friends and followers using their API. The crawlers were active for the duration of the study, collecting live data every hour which was later compiled into a single data feed for the whole duration of the study.

Data Cleaning
Following the end of the study, both datasets were cleaned. Specifically, days at the beginning and end of the study were trimmed where there were no ground truth reports submitted by the participant. Some participants unnecessarily reported their mood more than once per day, in which case the later time was used. Moreover, a number of participants were removed from the analysis using the following heuristics: (i) the same mood was reported for every day of the study, or (ii) after trimming, the dataset was less than 15 days long. After cleaning, the dataset consisted of a total of 16 participants including 406 individual days of data (avg. 25 days per participant) and 1,760 online actions (posts, likes, etc.) performed by the participants.

Mood detection

It is reasonable to expect that mood detection may not be possible for every OSN user. For some individuals their online activity can be more revealing of their psychological state than others. In our analysis, we tested different features, looking at the percentage of users for whom there was a statistically significant correlation ($p < 0.05$) with mood change (mood z-score).

Features
Using the OSN activity dataset, features were generated for each individual participant for the duration of the study. Values were calculated over a sliding window of n days, with $n-1$ days overlap. The optimal size of the window was estimated experimentally as shown in the "Results" section.

Sentiment Analysis
Sentiment analysis has been applied widely to discover the emotional context of messages exchanged online, including OSNs [8, 12, 15]. We employed sentiment analysis on the statuses and tweets that participants posted on Facebook and Twitter, using the Linguistic Inquiry Word Count (LIWC) dictionary. Through the LIWC toolkit we calculated the sentiment score S of each post using a method similar to [8]: $S = \frac{(n_{pos}-n_{neg})}{N}$ where n_{pos} is the number of positive words, n_{neg} is the number of negative words and N is the total number of words in a post ($S \in [-1, 1]$).

For each participant we estimated the average sentiment score per post within a window of 7 days and calculated the Pearson correlation with the mood z-score over the same window. The number of participants that demonstrated statistically significant correlations ($p < 0.05$) between their mood and the sentiment score was relatively low. Specifically for the Facebook posts, 38% of the participants show correlations between mood and sentiment score, and for Twitter only 6%. An interesting observation here is the clear difference in participant behaviour on the two OSNs. The same participants appear to be more revealing of their psychological state on Facebook than on Twitter. To some extent this is expected considering that Twitter is a public medium compared with Facebook which could be argued as being a more personal experience including friends and family. One possible explanation for the low performance of the sentiment score, is the relatively sparse dataset that

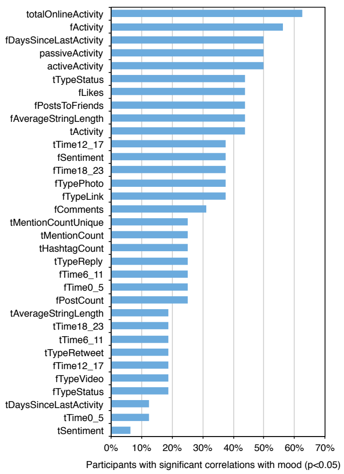

Figure 3: The percentage of participants who show significant correlations ($p < 0.05$) for each feature.

can be used for sentiment analysis. Within a window of 7 days, less than 1.5 days on average contain users' posts that can be analysed (Table 1). Twitter usage, for our participants in particular, was especially limited. Overall, sentiment score appeared to be a poor metric to identify real-world mood changes within a relatively short time frame.

Statistical Features

Using the datasets from Facebook and Twitter we calculated the most significant actions that a user can perform on the specific OSNs. These included counts of the following actions: Facebook status updates, likes, comments, posted links / photos / videos, Twitter posts, retweets, hashtag counts and mention counts. We consider these values as a reflection of the overall behaviour of each user on these OSNs. One important point is that the complete set of these features is significantly more dense than the plain textual posts / tweets that were considered for the sentiment analysis. This enabled us to work with denser datasets aggregated over a similar time window.

We expanded the feature set with a number of calculated / aggregate features that were derived from the plain features. Specifically, we calculated the total activity on each OSN ($fActivity$ and $tActivity$ for Facebook and Twitter), as a sum of all the counts of online actions that the user performs (e.g. posts, likes, comments, etc. on Facebook). We also calculated their overall online activity $totalOnlineActivity = fActivity + tActivity$. This is a reflection of how active each participant was online irrespective of the type of activity they performed. This was a very dense metric with values for almost every day of the study for every participant.

Motivated by the results from the work by De Choudhury et al [5] we enriched our set with features that captured the time of day that participants were active. Specifically, features $fTime0\_5$, $fTime6\_11$, etc. contain the average count of actions performed within the corresponding time windows: 00:00–05:00, 06:00–11:00 and so forth.

Feature	*Description*
f/tSentiment	Average sentiment score for posts and tweets
f/tAverageStringLength	Average length of post on Facebook / tweet
fActivity	# posts, likes, comments, posts to friends
tActivity	# of tweets
activeActivity	# posts, comments, tweets and replies
passiveActivity	# likes and retweets
fLikes	# likes on Facebook
fComments	# comments on Facebook
fPostsToFriends	# posts on friend's timelines
tHashtagCount tMentionCount	# hashtags and mentions in tweets
tMentionCountUnique	# unique mentions in tweets
fPostCount fTypeStatus/Link/ Photos/Video	# of posts and # of posts per type
tPostCount tTypeStatus/Reply/Retweet	# tweets and # of tweet per type
totalOnlineActivity	fActivity + tActivity
f/tTimeA_B	# of posts on Facebook / tweets within a time period (e.g. fTime6_11: # posts at 6:00–11:00)
f/tDaysSinceLastActivity	# days since the participant was last active on Facebook and Twitter

Table 2: Description of features

Finally, we wanted to capture the level of engagement of each participant with each OSN. Based on common experience with Facebook and Twitter, we classified the different types of online actions as "active" or "passive". Active actions contained actions such as submitting an original post or commenting on somebody else's post, while passive

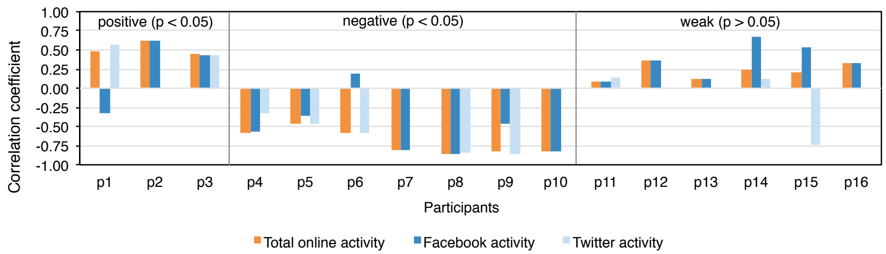

Figure 4: List of 16 participants with the correlation coefficients between mood and totalOnlineActivity, fActivity and tActivity. Participants on the left show a positive correlation, in the middle negative correlation, and on the right weak correlation.

actions were 'likes' or retweets. Generally, any action that required the user to type original text was deemed as "active" and all other actions were classified as "passive". The features $activeActivity$ and $passiveActivity$ contained the number of activities of these two classes. The aim behind these features was to explore whether more or less engaged interaction online is correlated with the mood of the user. The full featureset can be found in Table 2.

Results

We calculated Pearson's correlation for each participant between their mood changes and each feature in our set. We calculated the number of participants that demonstrated statistically significant results ($p < 0.05$) for each feature. Figure 3 shows the results for all these experiments. We can see that the $totalOnlineActivity$ is the feature where 61% of the participants demonstrate statistically significant correlation. These particular results were calculated over a sliding window of 7 days with 6 days overlap.

The optimal time window for the calculation of the features was estimated empirically. We ran multiple correlation tests using the percentage of users with statistically significant correlations between each feature and their mood as a metric. Increasing the size of the time window improves the performance of the $totalOnlineActivity$ feature with a peak at the 7-day window.

Based on our correlation results, we consider the $totalOnlineActivity$ aggregate feature as a good indicator of mood changes online. Participants in our study showed positive or negative correlation between their daily mood changes and their overall activity online (Figure 5), demonstrating that when some participants were experiencing a negative mood they were more active online, while for others a positive mood was also related to high online activity. Through informal interaction with some of our participants we received anecdotal confirmation that these results indeed match with their own perception of their online habits.

As seen in Figure 4, the correlation coefficients of different users can vary significantly. However, in order to develop a practical technique to detect mood changes, it is enough to identify the signage of the coefficient of each user, relying on online activity features. Indeed, knowing that a user has a negative correlation between $totalOnlineActivity$ and mood for example, can allow the tracking of mood change (whether it increases or decreases) using OSN activity data alone.

Discovering correlation types

In order to automate the tracking of mood changes through OSN activity, we need a mechanism to firstly, discover which users have strong or weak correlation between online activity and mood; and secondly, for those with strong correlation discover the signage of the coefficient between their online activity and mood changes. For this purpose we developed two machine learning classifiers to detect the different types of OSN users. We relied on the feature set that

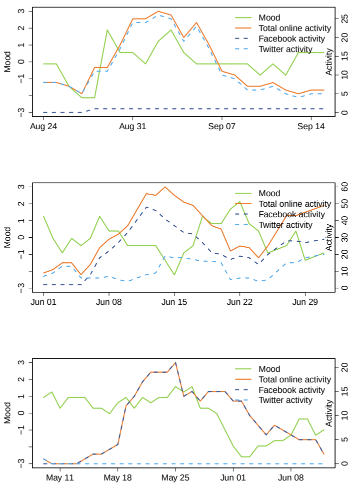

Figure 5: Positive (top) (c=0.45, p=0.03), negative (middle) (c=-0.46, p=0.01) and weak (bottom) (c=0.09, p=0.60) correlations between mood and totalOnlineActivity.

	Strong vs Weak	Pos. vs Neg.
Classifier	Random Forest	Voted Perceptron
Precision	95.2%	84.4%
Recall	94.7%	80.0%
F_1 score	0.947	0.763

Table 3: Classification results. (i) Discovering users with strong correlations between their mood changes and their online activity. (ii) classifying those with positive / negative correlations.

is shown in Table 2 and we used the correlation results to define the relevant classes: *strong* vs *weak* with respect to statistical significance, and *positive* vs *negative* with respect to coefficient signage.

Firstly, we aimed to develop a binary classifier that will identify users who may show statistically significant correlation against those with weak correlations between mood and online activity. We tried to select a minimum set of features that maximised the performance of the classifiers. We followed a *hill climbing* iterative approach, where we progressively eliminated features from the set and ran a 10-fold cross validation to measure the impact of each change. We eventually established a minimal feature set with the best performance, that consists of the following features:

- $lengthFAvg$: Average length of the Facebook posts
- $lengthTAvg$: Average length of the Twitter posts
- $activePassiveRatio$: Ratio of "active" actions (e.g. new posts) over "passive" actions (e.g. likes)
- $twitterFacebookRatio$: The ratio of Twitter actions over the Facebook actions of the user.

Conceptually this feature set captures the level of commitment of different users when interacting with OSNs. Users who actively interact with OSNs by posting are inherently more engaged than those who simply observe, like, or repost content. The length of their posts and the ratio of active vs passive actions were positively correlated with how predictable their mood is according to their online activity.

Using this minimal feature set we trained two classifiers. A "Strong vs Weak" classifier allows us to discover whose online activity reflects their mood changes, and a "Positive vs Negative" classifier which identifies the type of correlation. As shown in Table 3 the "Strong vs Weak" classifier demonstrates 95.2% precision and 94.7% recall, while the "Positive vs Negative" classifier demonstrates 84.4% precision and 80.0% recall. We acknowledge that the design of these classifiers is based on a relatively small dataset but the high accuracy demonstrates its feasibility to be applied on a larger scale.

The combination of these two classifiers along with the use of the *totalOnlineActivity* feature allows the design of the OSN mood tracking system (Figure 1). Using up to one month's online data, the two classifiers can identify with high accuracy the users who can be tracked and whether they demonstrate a positive/negative correlation between their online activity and mood. When the classification is achieved, the system uses a sliding window of one week's data to track the mood of the user.

Conclusions

In this paper we have demonstrated the feasibility of detecting real-world mood changes of people through their online activity on social networks such as Facebook and Twitter. To the best of our knowledge this is the first case of exploring correlations between activities over multiple OSNs and real-world mood data captured through experience sampling. The results of this work indicate that it is feasible to identify OSN users whose mood change can be tracked and then monitor their mood changes using their overall online activity. We believe that these results show how OSNs can act as a medium to facilitate affective communication online and demonstrates practical ways that this can be realised.

Acknowledgement

We would like to thank Dr. Neal Lathia for his support during the data collection using the EasyM app.

REFERENCES

1. M Argyle. 1972. Non-verbal communication in human social interaction. In *Non-verbal communication*. Cambridge U. Press, 243–270.
2. RC Balabantaray, M Mohammad, and N Sharma. 2012. Multi-class twitter emotion classification: A new approach. *International Journal of Applied Information Systems* 4, 1 (2012), 48–53.
3. J Bollen, H Mao, and A Pepe. 2011. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. In *Proceedings of ICWSM'11*. 450–453.
4. M De Choudhury, S Counts, EJ Horvitz, and A Hoff. 2014. Characterizing and predicting postpartum depression from shared facebook data. In *Proceedings of the 17th ACM conference on CSCW'14*. ACM, 626–638.
5. M De Choudhury, M Gamon, S Counts, and E Horvitz. 2013. Predicting Depression via Social Media. In *Proceedings of ICWSM'13*.
6. M De Choudhury, E Kiciman, M Dredze, G Coppersmith, and M Kumar. 2016. Discovering shifts to suicidal ideation from mental health content in social media. In *Proceedings of CHI'16*. ACM, 2098–2110.
7. P Ekkekakis. 2012. Affect, Mood, and Emotion. In *Measurement in sport and exercise psychology*, RC Eklund G Tenenbaum and A Kamata (Eds.). 321–332.
8. A DI Kramer. 2010. An unobtrusive behavioral model of gross national happiness. In *Proceedings of the CHI'10*. ACM, 287–290.
9. P Liu, W Tov, M Kosinski, DJ Stillwell, and L Qiu. 2015. Do Facebook Status Updates Reflect Subjective Well-Being? *Cyberpsychology, Behavior, and Social Networking* 18, 7 (2015), 373–379.
10. DN Prata, KP Soares, MA Silva, DQ Trevisan, and P Letouze. 2016. Social Data Analysis of Brazilian's Mood from Twitter. *International Journal of Social Science and Humanity* 6, 3 (2016), 179.
11. Pew Internet Project. 2014. Social Networking Fact Sheet. (2014). `http://www.pewinternet.org/fact-sheets/social-networking-fact-sheet/`
12. D Quercia, J Ellis, L Capra, and J Crowcroft. 2012. Tracking gross community happiness from tweets. In *Proceedings of CHI'12*. ACM, 965–968.
13. J A Russell, A Weiss, and GA Mendelsohn. 1989. Affect grid: a single-item scale of pleasure and arousal. *Journal of personality and social psychology* 57, 3 (1989), 493.
14. A Sadilek, HA Kautz, and V Silenzio. 2012. Modeling Spread of Disease from Social Interactions. In *Proceedings of ICWSM'12*.
15. N Wang, M Kosinski, DJ Stillwell, and J Rust. 2014. Can well-being be measured using Facebook status updates? Validation of Facebook's Gross National Happiness Index. *Social Indicators Research* 115, 1 (2014), 483–491.
16. D Watson, LA Clark, and A Tellegen. 1988. Development and validation of brief measures of positive and negative affect: the PANAS scales. *Journal of personality and social psychology* 54, 6 (1988), 1063.